



**Πανεπιστήμιο Πειραιώς
Σχολή Οικονομικών, Επιχειρηματικών και Διεθνών Σπουδών
Τμήμα Οργάνωσης και Διοίκησης Επιχειρήσεων**

Διδακτορική Διατριβή

“The Complex Nature of Firms: A Meta-Theoretical Approach”

Η Σύνθετη Φύση των Επιχειρήσεων: Μία Μετα-Θεωρητική Προσέγγιση

του
Αλέξανδρου Δ.Κάμπη

Επιβλέπων: Καθ. Μιχαήλ Σφακιανάκης, Πρύτανης Πανεπιστημίου Πειραιώς

Πειραιάς, Απρίλιος 2026



ΠΑΝΕΠΙΣΤΗΜΙΟ ΠΕΙΡΑΙΩΣ
ΣΧΟΛΗ ΟΙΚΟΝΟΜΙΚΩΝ ΕΠΙΧΕΙΡΗΜΑΤΙΚΩΝ ΚΑΙ ΔΙΕΘΝΩΝ ΣΠΟΥΔΩΝ
ΤΜΗΜΑ ΟΡΓΑΝΩΣΗΣ ΚΑΙ ΔΙΟΙΚΗΣΗΣ ΕΠΙΧΕΙΡΗΣΕΩΝ

ΒΕΒΑΙΩΣΗ ΕΚΠΟΝΗΣΗΣ ΔΙΔΑΚΤΟΡΙΚΗΣ ΔΙΑΤΡΙΒΗΣ

(περιλαμβάνεται ως ξεχωριστή (δεύτερη) σελίδα στο σώμα της διδακτορικής διατριβής)

Δηλώνω υπεύθυνα ότι η διδακτορική διατριβή για τη λήψη του διδακτορικού τίτλου, του Τμήματος Οργάνωσης και Διοίκησης Επιχειρήσεων του Πανεπιστημίου Πειραιώς, με τίτλο The Complex Nature of Firms: A Meta-Theoretical Approach / Η Συνθετική Φύση των Επιχειρήσεων: Μια Μετα-Θεωρητική Προσέγγιση έχει συγγραφεί από εμένα αποκλειστικά και στο σύνολό της. Δεν έχει υποβληθεί ούτε έχει εγκριθεί στο πλαίσιο κάποιου άλλου διδακτορικού, μεταπτυχιακού προγράμματος ή προπτυχιακού τίτλου σπουδών, στην Ελλάδα ή στο εξωτερικό, ούτε είναι εργασία ή τμήμα εργασίας ακαδημαϊκού ή επαγγελματικού χαρακτήρα.

Δηλώνω επίσης υπεύθυνα ότι οι πηγές στις οποίες ανέτρεξα για την εκπόνηση της συγκεκριμένης εργασίας, αναφέρονται στο σύνολό τους, κάνοντας πλήρη αναφορά στους συγγραφείς, τον εκδοτικό οίκο ή το περιοδικό, συμπεριλαμβανομένων και των πηγών που ενδεχομένως χρησιμοποιήθηκαν από το διαδίκτυο. Παράβαση της ανωτέρω ακαδημαϊκής μου ευθύνης αποτελεί ουσιώδη λόγο για την ανάκληση του διπλώματός μου.

Υπογραφή Διδάκτορα

Ονοματεπώνυμο

Ημερομηνία



To my beloved family

Acknowledgements

This dissertation has benefited from the guidance, support, and encouragement of many people, to whom I would like to express my sincere gratitude.

First and foremost, I am deeply grateful to my supervisor, Professor Michael E. Sfakianakis, for his guidance, intellectual generosity, and steady support throughout this doctoral journey. His rigorous standards, thoughtful advice, and confidence in my work were invaluable at every stage of the research and writing process. It has been a privilege to work under his supervision.

I am also sincerely grateful to the members of my supervisory committee, Professor Dimitrios A. Georgakellos and Prof Nikolaos B. Georgopoulos, for their careful reading of the dissertation, their constructive comments, and their generous engagement with this project. Their observations helped strengthen the work in important ways.

I would also like to thank the University of Piraeus, and in particular the Department of Business Administration, for providing the institutional setting in which this research was conducted. I am grateful for the academic and administrative support that enabled the completion of this dissertation.

My deepest thanks go to my family, especially my mother, my father, and my sister, for their constant love, encouragement, and support throughout my studies. Their presence has been a source of strength at every stage of this journey.

I am also sincerely grateful to Christina for her patience, understanding, and steadfast support during the long and demanding period of this work. Her encouragement and presence meant a great deal throughout this process.

To all of them, I extend my sincere thanks.

-

ABSTRACT

This dissertation is a meta-theoretical inquiry into the nature of firms. Unlike conventional firm theory, which analyses firm behaviour within inherited frameworks, meta-theoretical inquiry examines the foundational assumptions determining what firms are taken to be, what can be known about them, and how they can legitimately be analysed. The motivation arises from a persistent pattern: despite decades of refinement, existing theories exhibit recurring explanatory limitations regarding adaptation, innovation, heterogeneity, and structural instability – consequences not of empirical anomalies but of assumptions inherited from a mechanistic paradigm. Because firms exhibit the defining properties of complex adaptive systems – emergence, path-dependence, relational constitution, fundamental uncertainty – reconceptualisation requires intervention at the level of foundational assumptions rather than incremental extension.

The dissertation adopts problematisation as its methodology: generating theoretical contributions through systematic assumption-challenging rather than gap-spotting. Where gap-spotting extends literatures while leaving foundational commitments intact, problematisation interrogates background assumptions, evaluates which warrant challenge, and develops coherent alternatives. This methodology is appropriate because the limitations of firm theory arise from an assumptive configuration requiring structural intervention rather than local repair.

From this stance emerges the primary research question: *How should firms be reconceptualised when the foundational assumptions of mainstream economic theorising are systematically problematised and reconstructed to be compatible with a complex adaptive systems perspective?* This question emerged through three recognitions: that explanatory tensions arise under assumptions suited to closed, stable, decomposable settings; that surface-level extensions cannot address tensions arising from deeper commitments; and that criticism alone is insufficient – alternatives require explicit replacement assumptions. Five subsidiary questions operationalise the inquiry: assumption identification and diagnostic evaluation (SQ1–2); and ontological, epistemological, and methodological reconstruction (SQ3–5). Through the FROM→TO transformation framework, the dissertation proceeds in three movements corresponding to the diagnostic, reconstructive, and evaluative stages of problematisation.

First, it identifies the assumptive architecture of mainstream firm theory. The architecture comprises fifteen interlocking assumptions across a Lakatosian hard core – foundational commitments treated as methodologically irrefutable – and protective belt. The hard core consists of five ontological assumptions (reductionism, determinism, substance ontology, closure, atomism) and four epistemological assumptions (certainty, universal laws, monism, perfect information). The protective belt consists of six derivatives: equilibrium, linearity, homogeneity, ergodicity, optimisation, decomposability. Three reinforcing feedback loops interlock at high-centrality nodes, producing a triple-locked architecture explaining why piecemeal reform fails: the preserved hard core regenerates challenged derivatives or absorbs modifications as special cases.

Second, it develops an alternative assumption ground through coordinated FROM→TO transformations. Ontologically: emergent wholeness replaces reductionism; process ontology replaces substance; relational primacy replaces atomism; thermodynamic openness replaces closure; the meso-level is restored as the privileged analytical scale. Epistemologically: fundamental uncertainty is recognised as ontological; pattern-based navigation replaces prediction-centred knowledge; scaffolded rationality replaces unbounded optimisation.

Methodologically: structured pluralism replaces monism. These alternatives form a mutually reinforcing architecture – emergence requiring relational constitution, process ontology necessitating openness, fundamental uncertainty following from non-ergodicity – forging a Lakatosian hard core as genuine paradigmatic alternative.

Third, it evaluates whether the alternative generates compelling theory. The reconceptualisation yields a transformed understanding: firms as historically evolving, open, multi-level process-patterns – relationally constituted configurations whose identity consists in self-reproducing activity rather than fixed essence, whose boundaries function as selective membranes, whose capabilities emerge from relational configurations, and whose evolution follows path-dependent trajectories. Properties that mechanistic assumptions render invisible – emergence, downward causation, path-dependence, adaptive tension – become analytically accessible.

This reconceptualisation transforms not only answers but the nature of questions themselves. In strategic management, sustainable advantage gives way to adaptive capacity; strategy becomes navigation rather than optimisation. In organisational design, control-based intervention gives way to enabling conditions for emergence. In economic methodology, theoretical monism dissolves into structured pluralism under explicit scope conditions.

The dissertation operates under explicit constraints that are design choices appropriate to paradigm-level work: pure theory without empirical testing, since foundational reconstruction must precede meaningful testing; scope confined to the firm-as-CAS at the meso-level; modest predictive ambitions under non-ergodicity. These constraints open space for subsequent programmes: empirical investigation of meso-level phenomena, operationalisation of structured pluralism, case-based research examining firms as process-patterns – programmes enabled by the reconstructive cascade.

Keywords: Theory of the Firm; Complex Adaptive Systems; Economic Methodology; Paradigm Reconstruction; Problematisation; Ontology and Epistemology; Meso-Level Analysis

JEL Classification Codes

- **B41** – Economic Methodology
- **D21** – Firm Behavior: Theory
- **L22** – Firm Organization and Market Structure
- **B52** – Institutional, Evolutionary, and Complexity Economics
- **D83** – Search, Learning, Information, and Knowledge

ΠΕΡΙΛΗΨΗ

Η παρούσα διδακτορική διατριβή συνιστά μία μετα-θεωρητική διερεύνηση (meta-theoretical inquiry) της φύσης των επιχειρήσεων. Σε αντίθεση με τη συμβατική θεωρία της επιχείρησης (conventional firm theory), η οποία αναλύει τη συμπεριφορά των επιχειρήσεων εντός ήδη κληρονομημένων πλαισίων (inherited frameworks), η μετα-θεωρητική διερεύνηση εξετάζει τις θεμελιώδεις παραδοχές (foundational assumptions) που καθορίζουν τι θεωρείται ότι είναι οι επιχειρήσεις, τι μπορεί να γνωσθεί γι' αυτές και με ποιους θεμιτούς τρόπους μπορούν να αναλυθούν. Το ερευνητικό κίνητρο προκύπτει από μία επαναλαμβανόμενη διαδικασία: παρά τις δεκαετίες θεωρητικής βελτίωσης (refinement), οι υφιστάμενες θεωρίες εμφανίζουν επαναλαμβανόμενους ερμηνευτικούς περιορισμούς ως προς την προσαρμογή (adaptation), την καινοτομία (innovation), την ετερογένεια (heterogeneity) και τη δομική αστάθεια (structural instability), περιορισμούς που δεν αποτελούν συνέπειες εμπειρικών ανωμαλιών (empirical anomalies), αλλά απορρέουν από παραδοχές κληρονομημένες από ένα μηχανιστικό εννοιολογικό-θεωρητικό πλαίσιο αρχών (mechanistic paradigm), δηλαδή από ένα πλαίσιο που θέτει τις αρχές ανάπτυξης θεωριών, υποδειγμάτων και εφαρμογών. Δεδομένου ότι οι επιχειρήσεις εμφανίζουν τις καθοριστικές ιδιότητες των σύνθετων προσαρμοστικών συστημάτων (complex adaptive systems), όπως η ανάδυση (emergence), η εξάρτηση από τη διαδρομή (path-dependence), η σχεσιακή συγκρότηση (relational constitution) και η θεμελιώδης αβεβαιότητα (fundamental uncertainty), η εννοιολογική τους αναδιαμόρφωση (reconceptualisation) απαιτεί παρέμβαση στο επίπεδο των θεμελιωδών παραδοχών και όχι βαθμιαία επέκταση (incremental extension).

Η διατριβή υιοθετεί ως μεθοδολογία την προβληματοποίηση (problematisation), δηλαδή τη δημιουργία θεωρητικών προσθηκών (generating theoretical contributions) μέσω συστηματικής αμφισβήτησης παραδοχών (systematic assumption-challenging) και όχι μέσω εντοπισμού κενών (gap-spotting). Ενώ ο εντοπισμός κενών επεκτείνει επιμέρους βιβλιογραφίες αφήνοντας ανέπαφες τις θεμελιώδεις δεσμεύσεις, η προβληματοποίηση διερευνά τις παραδοχές υποβάθρου (background assumptions), αξιολογεί ποιες χρήζουν αμφισβήτησης και αναπτύσσει συνεκτικές εναλλακτικές. Η μεθοδολογία αυτή είναι κατάλληλη, διότι οι περιορισμοί των υφιστάμενων θεωριών της επιχείρησης (firm theory) απορρέουν από μία διάταξη παραδοχών (assumptive configuration) που απαιτεί δομική παρέμβαση και όχι μεμονωμένη παρέμβαση.

Από αυτή τη σκοπιά αναδύεται το κύριο ερευνητικό ερώτημα (primary research question): Πώς θα πρέπει να αναδιαμορφωθεί εννοιολογικά η επιχείρηση όταν οι θεμελιώδεις παραδοχές της κυρίαρχης οικονομικής θεωρητικής ανάλυσης (mainstream economic theorising) προβληματοποιούνται συστηματικά και αναδομούνται ώστε να καταστούν συμβατές με την οπτική των σύνθετων προσαρμοστικών συστημάτων (complex adaptive systems perspective);

Το ερώτημα αυτό προέκυψε από τρεις διαπιστώσεις: πρώτον, ότι τα ερμηνευτικά προβλήματα αναδύονται υπό παραδοχές κατάλληλες για κλειστά, σταθερά και αποσυνθέσιμα περιβάλλοντα· δεύτερον, ότι οι αναλύσεις επιμέρους θεμάτων δεν μπορούν να επιλύσουν προβλήματα που απορρέουν από βαθύτερες θεμελιώδεις δεσμεύσεις και, τρίτον, ότι η κριτική από μόνη της δεν αρκεί, καθώς οι εναλλακτικές απαιτούν ρητές παραδοχές-υποθέσεις αντικατάστασης (replacement assumptions).

Πέντε δευτερεύοντα ερωτήματα (subsidiary questions) εξειδικεύουν τη διερεύνηση: ο εντοπισμός παραδοχών και η διαγνωστική αξιολόγηση (SQ1-SQ2), καθώς και η οντολογική, επιστημολογική και μεθοδολογική αναδόμηση (SQ3-SQ5). Μέσω του πλαισίου

μετασχηματισμού ΑΠΟ→ΠΡΟΣ (FROM→TO transformation framework), η διατριβή αναπτύσσεται σε τρεις κινήσεις, οι οποίες αντιστοιχούν στα διαγνωστικά, ανασκευαστικά (reconstructive) και αξιολογικά στάδια της προβληματοποίησης (problematisation).

Πρώτον, εντοπίζεται η αρχιτεκτονική των παραδοχών (assumptive architecture) της κυρίαρχης θεωρίας της επιχείρησης. Η αρχιτεκτονική αυτή συγκροτείται από δεκαπέντε αλληλοδιαπλεκόμενες παραδοχές, κατανεμμένες μεταξύ του «σκληρού πυρήνα» κατά Lakatos (Lakatosian hard core), δηλαδή θεμελιωδών δεσμεύσεων που αντιμετωπίζονται ως μεθοδολογικά αδιάψευστες (methodologically irrefutable), και της προστατευτικής ζώνης (protective belt). Ο σκληρός πυρήνας αποτελείται από πέντε οντολογικές παραδοχές: αναγωγισμός (reductionism), ντετερμινισμός (determinism), οντολογία της ουσίας (substance ontology), παραδοχή συστημικής πληρότητας και περιχαράκωσης (closure) και ατομισμός (atomism). Περιλαμβάνει επίσης τέσσερις επιστημολογικές παραδοχές: βεβαιότητα (certainty), καθολικοί νόμοι (universal laws), μονισμός (monism) και τέλεια πληροφόρηση (perfect information). Η προστατευτική ζώνη περιλαμβάνει έξι παράγωγες δεσμεύσεις: ισορροπία (equilibrium), γραμμικότητα (linearity), ομοιογένεια (homogeneity), εργοδικότητα (ergodicity), βελτιστοποίηση (optimisation) και αποσυνθεσιμότητα (decomposability). Τρεις βρόχοι ενισχυτικής ανατροφοδότησης (reinforcing feedback loops) διαπλέκονται σε κόμβους υψηλής κεντρικότητας (high-centrality nodes), παράγοντας μία τριπλά κλειδωμένη αρχιτεκτονική (triple-locked architecture) που εξηγεί γιατί η αποσπασματική μεταρρύθμιση (piecemeal reform) αποτυγχάνει: ο διατηρούμενος σκληρός πυρήνας αναπαράγει τις αμφισβητούμενες παράγωγες δεσμεύσεις ή απορροφά τις τροποποιήσεις ως ειδικές περιπτώσεις (special cases).

Δεύτερον, αναπτύσσεται ένα εναλλακτικό θεμελιώδες πλαίσιο παραδοχών (alternative assumption ground) μέσω συντονισμένων μετασχηματισμών ΑΠΟ→ΠΡΟΣ (FROM→TO transformations). Σε οντολογικό επίπεδο, η αναδύομενη ολότητα (emergent wholeness) αντικαθιστά τον αναγωγισμό (reductionism), η οντολογία της διαδικασίας (process ontology) αντικαθιστά την οντολογία της ουσίας (substance ontology), η σχεσιακή πρωτοκαθεδρία (relational primacy) αντικαθιστά τον ατομισμό (atomism), η θερμοδυναμική ανοικτότητα (thermodynamic openness) αντικαθιστά την παραδοχή συστημικής πληρότητας και περιχαράκωσης (closure), και το μεσο-επίπεδο (meso-level) αποκαθίσταται ως προνομιακή αναλυτική κλίμακα (privileged analytical scale). Σε επιστημολογικό επίπεδο, η θεμελιώδης αβεβαιότητα (fundamental uncertainty) αναγνωρίζεται ως οντολογική συνθήκη, η πλοήγηση βάσει προτύπων (pattern-based navigation) αντικαθιστά τη γνώση που εστιάζει στην πρόβλεψη (prediction-centred knowledge), και η πλαισιωμένη ορθολογικότητα (scaffolded rationality) αντικαθιστά την απεριόριστη βελτιστοποίηση (unbounded optimisation). Σε μεθοδολογικό επίπεδο, ο δομημένος πλουραλισμός (structured pluralism) αντικαθιστά τον μονισμό (monism). Οι εναλλακτικές αυτές συγκροτούν μία αμοιβαία ενισχυόμενη αρχιτεκτονική, καθώς η ανάδυση (emergence) προϋποθέτει σχεσιακή συγκρότηση (relational constitution), η οντολογία της διαδικασίας (process ontology) συνεπάγεται ανοικτότητα (openness), και η θεμελιώδης αβεβαιότητα (fundamental uncertainty) απορρέει από τη μη εργοδικότητα (non-ergodicity), συγκροτώντας έτσι έναν σκληρό πυρήνα κατά Lakatos (Lakatosian hard core) ως γνήσια εναλλακτική στο επίπεδο του εννοιολογικού-θεωρητικού πλαισίου.

Τρίτον, αξιολογείται κατά πόσον η εναλλακτική αυτή παράγει μία πειστική θεωρητική σύλληψη. Η εννοιολογική αυτή αναδιαμόρφωση (reconceptualisation) αποδίδει μία μετασχηματισμένη κατανόηση των επιχειρήσεων ως ιστορικά εξελισσόμενων, ανοικτών, πολυεπίπεδων διαδικασιακών προτύπων (process-patterns), δηλαδή ως σχεσιακά συγκροτημένων διαμορφώσεων (relationally constituted configurations), των οποίων η

ταυτότητα συνίσταται σε αυτοαναπαραγόμενη δραστηριότητα και όχι σε σταθερή ουσία (fixed essence), των οποίων τα όρια λειτουργούν ως εκλεκτικές μεμβράνες (selective membranes), των οποίων οι ικανότητες αναδύονται από σχεσιακές διαμορφώσεις (relational configurations), και των οποίων η εξέλιξη ακολουθεί τροχιές εξαρτώμενες από τη διαδρομή (path-dependent trajectories). Ιδιότητες που οι μηχανιστικές παραδοχές (mechanistic assumptions) καθιστούν αόρατες, όπως η ανάδυση (emergence), η αιτιακή επίδραση του όλου στα μέρη (downward causation), η εξάρτηση από τη διαδρομή (path-dependence) και η προσαρμοστική ένταση (adaptive tension), καθίστανται πλέον αναλυτικά προσβάσιμες.

Η εννοιολογική αυτή αναδιαμόρφωση μετασχηματίζει όχι μόνο τις απαντήσεις, αλλά και την ίδια τη φύση των ερωτημάτων. Στη στρατηγική διοίκηση (strategic management), το βιώσιμο πλεονέκτημα (sustainable advantage) παραχωρεί τη θέση του στην προσαρμοστική ικανότητα (adaptive capacity), ενώ η στρατηγική (strategy) καθίσταται πλοήγηση (navigation) αντί βελτιστοποίησης (optimisation). Στον οργανωσιακό σχεδιασμό (organisational design), η παρέμβαση που στηρίζεται στον έλεγχο (control-based intervention) παραχωρεί τη θέση της στη διαμόρφωση συνθηκών που επιτρέπουν την ανάδυση (enabling conditions for emergence). Στην οικονομική μεθοδολογία (economic methodology), ο θεωρητικός μονισμός (theoretical monism) υποχωρεί υπέρ ενός δομημένου πλουραλισμού (structured pluralism) υπό ρητώς καθορισμένους όρους ισχύος και εφαρμογής (explicit scope conditions).

Η διατριβή αναπτύσσεται υπό ρητούς περιορισμούς (explicit constraints), οι οποίοι συνιστούν επιλογές σχεδιασμού (design choices) κατάλληλες για εργασία στο επίπεδο του εννοιολογικού-θεωρητικού πλαισίου (Paradigm-level work): καθαρή θεωρία (pure theory) χωρίς εμπειρική δοκιμή (empirical testing), δεδομένου ότι η θεμελιώδης αναδόμηση πρέπει να προηγείται κάθε ουσιαστικού ελέγχου· περιορισμό του πεδίου στην επιχείρηση ως σύνθετο προσαρμοστικό σύστημα (firm-as-CAS) στο μεσο-επίπεδο (meso-level)· και περιορισμένες αξιώσεις πρόβλεψης (modest predictive ambitions) υπό συνθήκες μη εργοδικότητας (non-ergodicity). Οι περιορισμοί αυτοί δεν συνιστούν αδυναμίες, αλλά διαμορφώνουν τις προϋποθέσεις για περαιτέρω ερευνητικά προγράμματα, όπως η εμπειρική διερεύνηση φαινομένων του μεσο-επιπέδου, η περαιτέρω μεθοδολογική εξειδίκευση του δομημένου πλουραλισμού (structured pluralism) και η ανάπτυξη ερευνών μελετών περίπτωσης (case-based research) που προσεγγίζουν τις επιχειρήσεις ως εξελισσόμενα πρότυπα οργανωμένων διεργασιών (process-patterns). Τα προγράμματα αυτά καθίστανται εφικτά μέσω της αναδομητικής αλληλουχίας (reconstructive cascade) που εισάγει η διατριβή.

Λέξεις-κλειδιά: Θεωρία της Επιχείρησης (Theory of the Firm), Σύνθετα Προσαρμοστικά Συστήματα (Complex Adaptive Systems), Οικονομική Μεθοδολογία (Economic Methodology), Αναδόμηση Εννοιολογικού-Θεωρητικού Πλαισίου (Paradigm Reconstruction), Προβληματοποίηση (Problematization), Οντολογία και Επιστημολογία (Ontology and Epistemology), Ανάλυση Μεσο-Επιπέδου (Meso-Level Analysis).

Κωδικοί ταξινόμησης JEL

- B41 Οικονομική Μεθοδολογία (Economic Methodology)
- D21 Συμπεριφορά της Επιχείρησης: Θεωρία (Firm Behavior: Theory)
- L22 Οργάνωση της Επιχείρησης και Δομή της Αγοράς (Firm Organization and Market Structure)
- B52 Θεσμική, Εξελικτική και Οικονομική της Πολυπλοκότητας (Institutional, Evolutionary, and Complexity Economics)

TABLE OF CONTENTS

Glossary of Key Working Terms.....	xxii
Chapter 1 – Introduction: Reconstructing the Firm for a Complex World.....	1
1.1 The Problem: Why the Firm Remains Theoretically Elusive	1
1.2 Limits of Incremental Reform in Firm Theory	2
1.3 A Paradigmatic Diagnosis: When the Framework Becomes the Constraint.....	3
1.4 Why Complexity – and Why Now.....	4
1.5 Scope, Contribution, and Boundaries	7
1.5.1 Scope of the Inquiry	7
1.5.2 Positioning Within Complexity Economics	8
1.5.3 Forging a Hard Core	9
1.5.4 Research Questions	10
1.5.5 Nature of the Contribution.....	11
1.5.6 Boundaries and Non-Claims.....	12
1.6 Structure of the Dissertation	13

Chapter 2 – Methodological and Philosophical Foundations	17
Introduction	17
2.1 Positioning the Dissertation.....	19
2.1.1 The Categorical Distinction: Pure Theory vs Applied Research.....	19
2.1.2 Clarifying Conceptual Foundations: Research Programmes, Paradigms, and Meta-Theories	21
2.1.3 What Theory IS and IS NOT	24
2.1.4 Historical Precedents: Paradigm-Level Theoretical Work	27
2.1.5 The Complicated / Complex Distinction and Its Methodological Implications.....	30
2.2 Philosophical Foundations.....	35
2.2.1 Critical Realism as Philosophical Foundation.....	35
2.2.2 Knowledge, Pre-Understanding, and the Hermeneutic Circle	39
2.2.3 Theories as Perspectives: The Epistemic Status of Claims	41
2.2.4 Scientific and Methodological Pluralism	44
2.2.5 Axiological Positioning: Values and the Role of the Researcher.....	46
2.3 The Problematization Methodology	49
2.3.1 The Problematization Methodology Introduction	49
2.3.2 The Six Methodological Principles of Problematization.....	53
2.3.3 Precursor Methodologies: The Lineage of Problematization	56
2.3.4 Combining Problematization with First-Principles Methodology	58
2.3.5 The Typology of Assumptions: Levels of Assumption Engagement.....	61
2.3.6 The FROM→TO Transformation as Problematization Method	65
2.4 Phenomena Construction.....	72
2.4.1 What Is Phenomena Construction?	72
2.4.2 The Construction of “Firms as Complex Adaptive Systems”	75
2.5 The Relationship to Empirical Research.....	78
2.5.1 Why Foundational Propositions Are Not Directly Empirically Verifiable	78
2.5.2 Complex Phenomena Resist Conventional Empirical Testing.....	82
2.5.3 Assumptions Cannot Be Empirically Tested.....	86
2.5.4 Theory Precedes and Enables Empirical Research	90
2.5.5 Paradigm Shifts Are Gestalt Switches, Not Evidence Accumulation	93
2.6 Success Criteria for Pure Theory.....	99
2.6.1 Criteria Appropriate for Pure Theory	99
2.6.2 Criteria Inappropriate for This Work	105
2.6.3 Lakatosian Self-Appraisal	108
2.7 Research Design and Architecture	114
2.7.1 Overall Design: Systematic Paradigm Transformation.....	114
2.7.2 The Multi-Level Structure.....	119
2.7.3 The Assumptive Network as Unit of Analysis	123
2.7.4 Structural Conventions	127
2.8 Tools and Techniques.....	133
2.8.1 Transdisciplinary Orientation.....	133
2.8.2 Literature Review Methodology.....	141
2.8.3 Sources and Materials.....	144
2.8.4 Philosophical Argumentation Methods	146
2.8.5 Assumption Identification and Mapping	148
2.8.6 Managing Source Conflicts and Cross-Disciplinary Translation.....	150
2.8.7 Historical–Theoretical Analysis	151
2.8.8 Analytical Procedures.....	155
2.8.9 Visual and Diagrammatic Methods	158
2.8.10 Firm-Level Illustration as a Heuristic Device	160
2.9 Contributions, Limitations, and Future Directions	163
2.9.1 Contributions, Scope, and Methodological Origins	163
2.9.2 Limitations and Future Directions.....	167

Chapter 3 – The Diagnosis: Anatomy of the Mechanistic Paradigm.....	172
Introduction	172
3.1 The Assumption Network.....	173
3.1.1 Core Ontological Assumptions.....	174
3.1.2 Core Epistemological Assumptions	178
3.1.3 Derivative Assumptions	186
3.1.4 The Assumption Network.....	191
3.2 The Paradigmatic Lock-In: Why Piecemeal Reform Fails	196
3.2.1 Reinforcing Feedback Loops Between Assumptions	196
3.2.2 The Self-Stabilising Nature of Paradigms	200
3.2.3 Kuhn on Normal Science and Anomaly Accumulation	204
3.2.4 Lakatos on Degenerating Research Programmes	208
3.3 Methodological Symptoms.....	215
3.3.1 Physics Envy and Mathematical Formalism as Symptoms	215
3.3.2 The Optimization Fetish and Calculus-Based Methods	223
3.3.3 DSGE, Representative Agent, and Model Monism.....	231
3.3.4 “Elegance Over Truth” and the Aestheticization of Economics	238
3.3.5 The Pretense of Knowledge (Hayek)	242
3.4 Consequences	248
3.4.1 In Economic Theory: Predictive Failures and Brittleness	249
3.4.2 In Firm Management: The Illusion of Control	256
3.4.3 The 2008 Financial Crisis as Paradigmatic Failure.....	262
Conclusion.....	268

Chapter 4 – Ontological Critique: The Failure of Mechanistic Being.....	271
Introduction	271
4.1 Ontological Reductionism: The Mechanistic Premise That Wholes Equal Sums.....	272
4.1.1 The Reductionist Premise.....	273
4.1.2 Five Arguments Against Ontological Reductionism.....	275
4.2 Substance Ontology: The Metaphysical Foundations of Being Over Becoming	286
4.2.1 The Parmenidean Legacy	287
4.2.2 The Transformation Stakes.....	295
4.2.3 Counter-Arguments Addressed	296
4.3 The Atomistic Foundation: Elements Before Relations.....	300
4.3.1 The Atomistic Worldview in Economics	301
4.3.2 Four Arguments Against Atomism	301
4.3.3 The Transformation Stakes.....	309
4.3.4 Counter-Arguments Addressed	311
4.4 The Limits of Analytical Decomposition: When Analysis Destroys What It Studies	315
4.4.1 Analysis versus Synthesis: A Framework	316
4.4.2 Four Arguments Against Analytical Decomposition	317
4.4.3 Critical Nuances and Methodological Consequences	323
4.4.4 Counter-Arguments Addressed	325
Conclusion	328

Chapter 5 – Ontological Reconstruction: Process, Emergence, and Complex Adaptive Systems.....	330
Introduction	330
5.1 The Principle of Emergent Wholeness: When More Becomes Different.....	333
5.1.1 The Ontological Core and Defining Characteristics	334
5.1.2 Domains and Economic Applications	338
5.1.3 Emergence, the Meso-Level, and Synthesis.....	343
5.2 Process Ontology: Reality as “Becoming”	349
5.2.1 Philosophical Foundations: Substance versus Process.....	350
5.2.2 Economic Applications and the Meso-Level	357
5.2.3 Counter-Arguments and Responses	361
5.3 Substances as Stabilised Processes.....	365
5.3.1 Philosophical Foundations: From Atoms to Relations.....	366
5.3.2 Economic Applications: Embeddedness and Structuration.....	371
5.3.3 Counter-Arguments and Responses	375
5.4 Relational Ontology: The Constitutive Primacy of Relations	379
5.4.1 The Problem of Persistence.....	380
5.4.2 Mechanisms of Stabilisation.....	381
5.4.3 The Hegelian Resolution and Synthesis	385
5.5 Emergence as Meso-Level Phenomenon: The Missing Middle.....	391
5.5.1 The Missing Middle in Economic Theory.....	392
5.5.2 Meso-Level Phenomena: Routines, Networks, Industries, Institutions	394
5.5.3 The Three-Level Ontology and Synthesis	400
5.6 Thermodynamic Openness and Boundary Permeability.....	406
5.6.1 Thermodynamic Foundations: Open Systems and Dissipative Structures.....	408
5.6.2 The Firm as Metabolic Processor	411
5.6.3 Boundary Permeability and Dynamics.....	413
5.7 Network Architecture and Meso-Level Structure.....	422
5.7.1 Hierarchical Architecture and Near-Decomposability	423
5.7.2 Network Topology and Properties.....	431
5.7.3 The Firm as Network Cluster	436
5.7.4 Local Interactions, Global Emergence, and Synthesis	442
5.8 Interconnections – From Ontology to Epistemology, Methodology, and Structure	449
5.8.1 Ontology → Epistemology: Fundamental Uncertainty and Epistemic Complementarity	449
5.8.2 Ontology → Methodology: Expansionism over Reductionism	466
5.8.3 Ontology→Structure: The Necessity of Open Systems and Boundary Permeability	483
5.9 Dynamic Implications – Emergence Enables Novelty Generation	494
5.9.1 The Novelty Problem in Mechanistic Economics	495
5.9.2 Five Dynamic Implications	496
5.9.3 Counter-Arguments and the Generative Dynamics Framework	507
Conclusion	511

Chapter 6 – Conclusions, Contributions, Implications, Limitations, and Future Research Suggested.....	515
Introduction	515
6.1 Evaluation of the Dissertation’s Contribution.....	517
6.1.1 Nature and Level of Contribution.....	517
6.1.2 What the Dissertation Has Achieved.....	518
6.1.3 Criteria for Evaluation.....	521
6.2 What the Thesis Does Not Claim	523
6.2.1 It Does Not Propose a Unified "New Theory of the Firm"	524
6.2.2 It Does Not Claim Empirical Validation in This Dissertation	524
6.2.3 It Does Not Claim Ontological Emergence of Ideas	Error! Bookmark not defined.
6.2.4 It Does Not Reject Formal Modelling, Optimisation, or Prediction Per Se	525
6.2.5 It Does Not Offer a Managerial Manual or Policy Blueprint	525
6.2.6 It Does Not Present Itself as a Revolutionary Rejection of Prior Work.....	526
6.2.7 Its Scope Is Confined to the Firm at the Meso-Level.....	Error! Bookmark not defined.
6.3 Limitations and Conditions of Validity	527
6.3.1 The Cascade Condition.....	527
6.3.2 The Standards on Which Validity Rests.....	528
6.3.3 Acknowledged Vulnerabilities	529
6.3.4 Why These Limitations Are Productive	531
6.4 Implications Enabled by the Reconceptualisation	533
6.4.1 Implications for Economic Theory of the Firm.....	535
6.4.2 Implications for Methodology and Standards of Rigour	539
6.4.3 Implications for Management and Strategic Practice.....	541
6.4.4 Implications for Economic Policy and Institutional Analysis.....	543
6.4.5 Implications for the Philosophy and Methodology of Economics	545
6.4.6 Implications Amplified by Artificial Intelligence	548
6.5 Future Research Programmes	550
6.5.1 Theoretical Extensions	551
6.5.2 The Empirical Programme	552
6.5.3 Methodological Development	554
6.5.4 From Reconstruction to Empirical Purchase.....	555
6.5.5 Artificial Intelligence, Computation, and the Transformation of Firms	556
6.5.6 Comparative, Cross-Cultural, and Meta-Theoretical Frontiers	557
References List	560

LIST OF TABLES

Table 1: How Complex Adaptive System Properties Undermine Conventional Statistical Testing: Mechanisms and Methodological Implications.....	83
Table 2: Structural Features of Complex Economic Systems and Their Consequences for Causal Inference and Hypothesis Testing	84
Table 3: Chapter Structure and Transformative Logic of the Thesis: From Paradigm-Level Problematization to Theoretical Reconstruction	115
Table 4: Mapping Problematization Principles to Dissertation Implementation: Operationalising Methodological Critique Across Chapters	117
Table 5: Three-Level Paradigm Transformation: Ontological, Epistemological, and Methodological Shifts from Mainstream to Complex Adaptive Systems Foundations	119
Table 6: Ontological Assumptions of Mainstream Theory of the Firm and the Complex Adaptive System Phenomena They Structurally Eliminate	123
Table 7: Epistemological Assumptions of Mainstream Theory of the Firm and the Complex Adaptive System Phenomena They Structurally Eliminate	124
Table 8: Derivative Assumptions of Mainstream Theory of the Firm: Their Generative Foundations and the Complex Adaptive System Phenomena They Eliminate	124
Table 9: Four-Stage Assumption Transformation Protocol: Structure and Illustration Through the Case of Emergence.....	127
Table 10: Standardised Section Architecture: Structural Elements and Their Functions for Ensuring Argumentative Coherence.....	128
Table 11: Inter-Section Dependency Link Types: Functions and Illustrations of the Dissertation's Argumentative Connectivity Structure.....	128
Table 12: Key Terminological Conventions: Stipulated Definitions and Their Justifications for Maintaining Analytical Precision.....	129
Table 13: Literature Sufficiency Criteria: Operational Definitions for Evaluating Completeness of the Dissertation's Source Engagement	143
Table 14: Taxonomy of Assumption Types in Economic Theorising: Functions and Diagnostic Questions.....	149
Table 15: Contrasting Genealogical Orientations in Economic Theorising and Their Dominant Assumptive Emphases.....	152
Table 16: Typology of Visual-Analytical Diagrams Employed in the Dissertation: Purposes and Scope of Application	159
Table 17: Functions of Illustrative Examples in the Dissertation: Pedagogical, Heuristic, and Orientational Roles.....	161
Table 18: Logical Interdependencies Among Ontological Assumptions: The Mutually Reinforcing Network of Mainstream Theory of the Firm.....	177
Table 19: From Ontological Assumptions to Firm-Level Blind Spots: How Eliminated CAS Properties Render Key Organisational Phenomena Invisible	178
Table 20: Ontological-to-Epistemological Derivation: How Foundational Ontological Assumptions Generate Corresponding Epistemological Commitments	181
Table 21: From Epistemological Assumptions to Firm-Level Blind Spots: How Eliminated CAS Properties Obscure Key Dimensions of Firm Behaviour.....	182

Table 22: Epistemological Assumptions of Mainstream Theory of the Firm: Core Claims and the CAS Phenomena They Structurally Eliminate.....	185
Table 23: Generative Mapping of Derivative Assumptions: Tracing Each to Its Foundational Ontological and Epistemological Sources.....	188
Table 24: The Dual Role of Derivative Assumptions: Analytical Techniques Enabled and CAS-Relevant Phenomena Excluded.....	189
Table 25: Firm-Level Phenomena Eliminated by Derivative Assumptions: What Mainstream Theory Cannot Represent	190
Table 26: Summary of the Mechanistic Assumption Network: Distribution of Fifteen Assumptions Across Ontological, Epistemological, and Derivative Layers	190
Table 27: 3.1 Derivative assumptions and their generative foundations	191
Table 28: 3.2. Methodological symptoms and their emergence.....	192
Table 29: 3.3. Feedback loops within the assumption network	193
Table 30: 3.4. Assumptions and meso-level elimination mechanisms.....	195
Table 31: 3.5. Reinforcing feedback loops in the mechanistic paradigm	197
Table 32: 3.6. Semantic stabilisation through embedded assumptions.....	201
Table 33: 3.7. Methodological stabilisation through tool–assumption coupling	202
Table 34: 3.8. Institutional stabilisation through professional selection.....	202
Table 35: 3.9. Illustrative reform waves as stabilising accommodation	203
Table 36: 3.10. Predictive and policy anomalies.....	205
Table 37: 3.11. Conceptual anomalies	206
Table 38: 3.12. Institutional and historical anomalies.....	206
Table 39: 3.13. Technological and organisational anomalies	206
Table 40: 3.14. Structure of a Lakatosian research programme (Lakatos, 1970)	208
Table 41: Table 3.15. Progressive and degenerative programmes (Lakatos, 1970)	209
Table 42: 3.16. Hard-core commitment clusters in the inherited programme	209
Table 43: 3.17. Immunisation pattern: hard core preserved, protective belt adjusted	210
Table 44: 3.18. Post hoc accommodation sequence.....	210
Table 45: 3.19. Protective-belt proliferation in core model families	211
Table 46: 3.20. Calibration success and failure outside the calibration domain.....	212
Table 47: 3.21. Lakatosian criteria for a progressive alternative (Lakatos, 1970).....	213
Table 48: 3.22. Four physics-style appropriations and their diagnostic problem	216
Table 49: 3.23. From mechanistic assumptions to formalist method.....	217
Table 50: 3.24. Perceived benefits and systemic costs of physics-style formalism.....	218
Table 51: 3.25. Domain conditions: physics-style form versus economic reality	219
Table 52: 3.26. CAS features and the mechanistic commitments they contradict.....	220
Table 53: 3.27. Tractability-driven inversion of research priorities.....	221
Table 54: 3.28. The four appropriations and the source of failure.....	222
Table 55: 3.29. The cognitive shift.....	222

Table 56: 3.30. Optimisation as disciplinary template and its eliminations.....	224
Table 57: 3.30A. Why optimisation mis-specifies the firm as organisation	225
Table 58: 3.31. Cascade logic from assumptions to optimisation.....	225
Table 59: 3.32. Calculus requirements and their representational implications	226
Table 60: 3.32A. Process rationality versus solution rationality in innovation firms	227
Table 61: 3.33. Why path-dependent platform evolution defeats optimisation	229
Table 62: 3.33A. Optimisation requirements versus CAS properties	229
Table 63: 3.31. What “DSGE” operationally means and what it presupposes	232
Table 64: 3.32. DSGE’s equilibrium universe as an encoding of mechanistic assumptions	233
Table 65: 3.33. What the representative agent removes and what is lost thereby.....	234
Table 66: 3.34. Linearisation requirements and the dynamics removed.....	235
Table 67: 3.35. Operational symptoms of DSGE monism and their effects	236
Table 68: 3.36. Failure domains and the assumptions that generate them.....	237
Table 69: Mathematical Aesthetics as Substantive Theoretical Commitments: How Formal Tractability Requirements Exclude Empirically Relevant Phenomena	238
Table 70: Elegance Versus Empirical Adequacy: The Persistent Disconnect Across Canonical Mainstream Economic Models.....	239
Table 71: Mainstream Economics Evaluated Against Its Own Claimed Functions: Systematic Failure and the Circularity of Puzzle-Solving Success.....	239
Table 72: Features Excluded on Grounds of Mathematical Inelegance and Their Essential Roles in Complex Adaptive Systems	240
Table 73: The Aesthetic Displacement of Empirical Standards: Five Shifts in Evaluative Criteria Within Mainstream Economics	241
Table 74: Consequences of Aesthetic Dominance in Mainstream Economics: How Formal Elegance Displaces Empirical Reality	241
Table 75: Elegance as Paradigm Protection: Mechanisms by Which Mathematical Aesthetics Stabilise Mainstream Economic Theory	242
Table 76: Elements of Methodological Overreach: How Mainstream Economics Conflates Scientific Form with Scientific Validity.....	243
Table 77: Why Econometric Assumptions Fail in Complex Economies: Requirements and Their Points of Breakdown	244
Table 78: Presuppositions Underpinning Empirical Testability in Mainstream Economics: What Each Assumption Requires of Economic Systems.....	244
Table 79: Mainstream Economic Beliefs Versus Complex Adaptive System Realities: The Systematic Mismatch Undermining Policy and Prediction.....	245
Table 80: The Pretence of Knowledge Across Micro, Meso, and Macro Levels: Agent Claims and Their Collapse Under CAS Realities.....	245
Table 81: From Prediction to Navigation: Five Paradigmatic Shifts Toward CAS-Compatible Economic Reasoning.....	246
Table 82: Rhetorical Mechanisms of Paradigm Defence: How Mainstream Economics Deflects Fundamental Critique.....	246
Table 83: Mapping Hayekian Concepts to Their Complex Adaptive Systems Equivalents.....	247

Table 84: Assumption Families in Mainstream Macroeconomic Modelling: Analytical Tractability Gained and CAS Dynamics Suppressed	251
Table 85: From Mechanistic Suppression to Crisis Failure: How the Elimination of CAS Phenomena Produces Systematic Blind Spots in Crisis Prediction and Explanation.....	254
Table 86: 3.4.2-1. Imported assumptions and their managerial translation	256
Table 87: 3.4.2-2. Assumption-to-belief chain and its managerial consequence.....	257
Table 88: 3.4.2-3. CAS properties of organisations and the managerial implication	258
Table 89: 3.4.2-4. Mechanistic management principles and CAS alternatives.....	261
Table 90: Matrix 1. Ontological assumptions and crisis failure modes	262
Table 91: Matrix 2. Epistemological and methodological assumptions and predictive collapse....	263
Table 92: Risk-modelling presuppositions versus crisis dynamics.....	264
Table 93: Filtering mechanism: what was treated as “secondary” became decisive	265
Table 94: Pre-2008 warnings: mechanism-level anticipation (not date-level forecasting).....	265
Table 95: Objections and structural replies.....	266
Table 96: 3.4.3-A. The four appropriations (physics → economics) as failure mechanisms	266
Table 97: 3.4.3-B. The cognitive shift required by the 2008 stress test.....	267
Table 98: Table 4.1 – What Reductionist Ontology Systematically Eliminates from View.....	285
Table 99: From Substance to Process Ontology of the Firm: Five Foundational Shifts Toward a CAS-Compatible Conception.....	293
Table 100: Firm-Level Implications of the Ontological Shift: Reconceptualising Strategy, Capabilities, Innovation, Stability, and Failure from Substance to Process.....	294
Table 101: Being Versus Becoming Across Six Dimensions: How Substance and Process Ontologies Shape the Visibility of Firm Phenomena	295
Table 102: Six Core Conceptual Shifts from Being to Becoming: The Ontological Transformation from Substance to Process.....	299
Table 103: Diagnostic Summary of Atomistic Assumptions: Relational and Network Phenomena Forced Out of Mainstream Theory of the Firm.....	303
Table 104: Ontological Transformation from Atomistic Methodological Individualism to CAS Relational Ontology: Five Foundational Shifts.....	304
Table 105: Diagnostic Summary of Atomistic Presumptions About Firms and the Relational Phenomena They Mis-Specify	307
Table 106: What Atomism Makes Invisible: Seven Presumptions and the Network-Dependent Phenomena They Obscure or Misattribute.....	308
Table 107: Diagnostic Table: How Relational Causality Undermines the Implicit Requirements of Core Mainstream Analytical Tools	310
Table 108: The Three Pillars of the Mechanistic Ontology: Core Assumptions and the CAS Phenomena They Collectively Eliminate	310
Table 109: Atomistic vs Relational Commitments	314
Table 110: Reductionist vs Emergent Ontologies.....	314
Table 111: Matching Methodological Emphasis to System Structure: From Analysis and Synthesis to Their Explicit Integration in Mixed Systems	323

Table 112: Decomposition Methods in Mainstream Economics and Their Systematic Violation of Complex Adaptive System Properties.....	324
Table 113: From Decomposition to Synthesis: Six Methodological Shifts Toward a CAS-Compatible Analytical Framework.....	327
Table 114: Emergence Across Four Domains: What Arises and Why It Is Irreducible to Micro-Level Components.....	340
Table 115: Categorical Reclassification of Economic Reality: How Emergence Transforms Five Fundamental Economic Concepts from Orthodox to Systemic Categories	343
Table 116: Ontological Transformation from Reductionism to Emergence: Five Core Commitments Reconceived	348
Table 117: Comprehensive Transformation Map: Reconceiving Ten Fundamental Economic Concepts from Substance Ontology to Process Ontology.....	363
Table 118: Transformation Map from Atomism to Relational Ontology: Reconceiving Nine Foundational Dimensions of Economic Analysis	377
Table 119: From Substance to Stabilised Process: Reconceiving Nine Dimensions of Firm Ontology	389
Table 120: Framework Consequences of Recognising Routines as Causally Powerful Meso-Level Entities: From Individualist to Routine-Based Explanations	396
Table 121: Multi-Level Ontological Structure of the iOS Ecosystem: Illustrating Micro, Meso, and Macro Layers of a Platform as CAS	397
Table 122: Institutional Complementarities Across Liberal and Coordinated Market Economies: Context-Dependence Across Four Dimensions	398
Table 123: Route 128 Versus Silicon Valley: Institutional Divergence Across Five Dimensions as Evidence of Meso-Institutional Causation	399
Table 124: Three-Level CAS Analysis in Practice: Micro, Meso, and Macro Analytical Foci Illustrated Through the Pharmaceutical Industry	402
Table 125: From Two-Level Reduction to Three-Level Ontology: Reconceiving Eight Analytical Dimensions with the Meso-Level as Ontologically Autonomous.....	405
Table 126: Apple as Metabolic Network: Graduated Boundary Zones Illustrating the Permeability and Relational Nature of Firm Boundaries.....	414
Table 127: From Closed Walls to Open Gradients: Thirteen-Dimension Transformation of Firm Boundary Ontology	421
Table 128: From Flat Ontology to Nested CAS Architecture: Nine-Dimension Transformation Toward a Hierarchical, Multi-Level Ontology of the Firm.....	447
Table 129: From Atomistic to Topological: Eight-Dimension Transformation Reconceiving the Firm as a Network-Positioned Entity	447
Table 130: Ontological Significance of Meso-Level Clusters: Four Criteria Establishing Their Causal Autonomy and Irreducibility	448
Table 131: Two Complementary Modes of Explanation: Mechanistic and Organismic Approaches Contrasted by Direction and Explanatory Scope	460
Table 132: Coherent Versus Incoherent Ontology-Epistemology Pairings: Why Methodological Approach Must Match Ontological Commitment.....	461
Table 133: Table 5.8.1: The Epistemological Transformation	465

Table 134: Property-Based Method Selection: Mapping CAS-Appropriate Methods to Their Targeted System-Level Properties	475
Table 135: The Four Dimensions of Methodological Expansion	477
Table 136: Table 5.8.2: The Method–Ontology Bridge.....	482
Table 137: Table 5.8.2b: The Methodological Transformation.....	482
Table 138: Table 5.8.3a: From Production Functions to Ecosystems.....	489
Table 139: Four forms of economic irreversibility	490
Table 140: Table 5.8.3b: The Ontology → Structure Transformation.....	492
Table 141: Table 5.9.1: Recombination vs Ontological Novelty.....	496
Table 142: Table 5.9.2: Expansion of the Adjacent Possible.....	503
Table 143: Table 5.9.3: Negative vs Positive Feedback	503
Table 144: Table 5.9.4: Mechanisms of Novelty Amplification.....	505
Table 145: Failed Dynamic Hybrids	509
Table 146: Table 5.9.5: Dynamic Coherence.....	509
Table 147: Table 5.9.6: From Equilibrium Mechanics to Generative Dynamics	509
Table 148: Table 5.9.7: The Dynamic Transformation.....	509
Table 149: The Four Completed Paradigmatic Transformations: Summary of the Shift from Mechanistic to CAS Foundations Across Ontology, Epistemology, Methodology, and Dynamics	513

LIST OF FIGURES

Figure 1: Figure 4.1 – The Reductionist Inheritance Chain (From Physics to Economics to Management).	274
Figure 2: Figure 4.2 – Decomposable vs Non-Decomposable Systems: Why Analysis Works for Machines but Fails for CAS	276
Figure 3: Figure 4.3 – The Ontological Inversion: From Elements-Primary to Interactions-Primary	279
Figure 4: Figure 4.4 – The Meso-Level Erasure: Reductionism’s “Invisible Middle”	281
Figure 5: Figure 5.1. Substance and Process Ontology: Philosophical Foundations of Competing Economic Paradigms	351

Glossary of Key Working Terms

This glossary provides *working definitions as used in this dissertation*. Its purpose is to articulate meaning across a paradigmatic and meta-theoretical argument. Each term is also classified in brackets according to its primary disciplinary home within the thesis.

Abduction / Retroduction (Philosophy of Science & Methodology): Abduction is inference to the best explanation; reasoning from surprising observations to hypotheses that would render them expected. Peirce's "third form of reasoning" alongside deduction and induction. Logical form: "The surprising fact C is observed; but if A were true, C would be a matter of course; hence, there is reason to suspect that A is true." Retroduction, in critical realism, is the inferential movement from empirical phenomena to the underlying generative mechanisms that would explain them, asking: "What must the world be like for this phenomenon to be possible?" *References*: Peirce, C.S. (1903); Bhaskar, R. (1975/2008), pp. 24-27.

Absorptive capacity (Rationality, Uncertainty & Decision-Making): "The ability of a firm to recognize the value of new, external information, assimilate it, and apply it to commercial ends." Path-dependent: prior knowledge enables future learning; lack of early investment may cause "lockout." *Reference*: Cohen, W.M. & Levinthal, D.A. (1990), pp. 128-152.

Adaptation (Complexity Science & Systems Theory): The process through which entities modify their structures, behaviours, or strategies in response to environmental pressures or internal dynamics. In complex adaptive systems, adaptation occurs through decentralised mechanisms—agents responding to local feedback—rather than centralised design. Distinguished from mere reaction by involving structural or behavioural changes that persist and affect future responses. Key dimensions: (1) *Ontogenetic adaptation*: changes within an individual entity's lifetime; (2) *Phylogenetic adaptation*: population-level changes through variation and selection; (3) *Adaptive capacity*: the potential range of responses available given current structure. In CAS theory, adaptation is not optimisation toward a fixed target but ongoing adjustment in co-evolutionary dynamics where the "fitness landscape" itself shifts as agents adapt. *References*: Holland, J.H. (1995), Ch. 5.

Anomalies (*Original to thesis; Lakatosian usage*) (Philosophy of Science & Methodology): Anomalies are phenomena that cannot be explained within an existing framework without auxiliary adjustments. In Lakatosian terms, the accumulation of anomalies stresses a research programme and often triggers protective-belt expansions rather than revision of the hard core. *References*: (This thesis: Ch. 3 "diagnostic failures / anomalies"; Ch. 5 "prediction failure"). [Anomalies] [Research programme] (Kuhn, 1962; Lakatos, 1970).

Atomism / Holism (Ontology): *Atomism* is the ontological and methodological position that wholes are fundamentally composed of discrete, independent parts whose properties and interactions fully explain aggregate phenomena. In social science, methodological atomism holds that social facts must be explained by reference to individuals and their properties, treating persons as analytically prior to social relations. *Holism* is the counter-position that wholes possess properties irreducible to their constituent parts—that the whole is "greater than the sum of its parts." Holism maintains that systems must be understood as integrated totalities rather than decomposed into elements. *References*: Smuts, J.C. (1926); Nagel, E. (1961), pp. 380-398.

Attractor / Strange attractor (Complexity Science & Systems Theory): A set of states toward which a dynamical system evolves over time from nearby initial conditions. Represents the system's long-term behaviour or equilibrium regime. Types: *Fixed-point attractor*: system settles to single stable state; *Limit cycle*: system oscillates in periodic pattern; *Strange attractor*: bounded region in phase space exhibiting aperiodic, chaotic behaviour; sensitive dependence on initial conditions; fractal geometric structure. *References*: Lorenz, E.N. (1963; Ruelle, D. & Takens, F. (1971), pp. 167-192; Strogatz, S.H. (1994).

Autopoiesis (Complexity Science & Systems Theory): Self-production; the property of systems that continuously regenerate their own components and maintain their organisation through recursive processes. Autopoietic systems are operationally closed (self-referential in their organisation) while remaining structurally open (exchanging matter and energy with environment). Originally developed for biological cells; extended controversially to social systems by Luhmann. From Greek *auto* (self) + *poiesis* (creation, production). *References*: Maturana, H.R. & Varela, F.J. (1980; Luhmann, N. (1995).

Betweenness centrality (Network Science): A network measure quantifying how frequently a node lies on the shortest paths between other nodes. Nodes with high betweenness centrality serve as bridges or brokers, controlling information flow between otherwise disconnected regions. For a node v , betweenness centrality is the sum over all pairs of nodes of the fraction of shortest paths that pass through v . High-betweenness nodes connect different clusters; their removal can fragment networks. *Reference*: Freeman, L.C. (1977), 35-41.

Bifurcation / Bifurcation point (Complexity Science & Systems Theory): A qualitative change in system behaviour precipitated by small, smooth parameter variation, marking transition between dynamically distinct regimes. At bifurcation points, systems may shift from stable equilibria to oscillation, from periodicity to chaos, or between multiple attractors. Types include saddle-node, pitchfork, Hopf, and period-doubling bifurcations. Bifurcations are dynamical analogues of thermodynamic phase transitions. Provides mathematical foundation for "tipping points." *References*: Strogatz, S.H. (1994; Kuznetsov, Y.A. (2004).

Causal powers (Philosophy of Science & Methodology): The capacities of entities to act or be acted upon in virtue of their intrinsic structural properties. In critical realism, causal powers are possessed by objects regardless of their exercise—they represent potentialities that may or may not be actualised depending on conditions and countervailing powers. Powers are emergent properties of structured wholes, irreducible to their components. Powers exist transfactually (operating whether or not manifested); events are contingent actualisations of interacting powers in specific circumstances. *References*: Harré, R. & Madden, E.H. (1975); Bhaskar, R. (1975/2008).

Circular causation (Complexity Science & Systems Theory): A causal structure in which cause and effect are reciprocally related—A influences B, which influences A—forming closed causal loops. Central to Myrdal's "cumulative causation" (initial advantage leads to further advantage through reinforcing cycles) and structure-agency debates. Requires dynamic modelling approaches rather than linear causal inference. *References*: Myrdal, G. (1957); Archer, M.S. (1995).

Clustering coefficient (Network Science): Network measure quantifying the degree to which nodes cluster together—the extent to which neighbours of a node are themselves connected. For node v with k neighbours, local clustering coefficient = $2e/k(k-1)$, where e = edges between neighbours. High clustering (≈ 1) = cliquish structure; low clustering (≈ 0) = neighbours rarely connected. Small world networks combine high clustering with short path lengths. *Reference*: Watts, D.J. & Strogatz, S.H. (1998), 393, 440-442.

Co-evolution (Complexity Science & Systems Theory): The reciprocal evolutionary process whereby two or more entities or populations mutually shape each other's development trajectories through ongoing interaction. Each entity's adaptive landscape is modified by changes in the others, creating coupled dynamics where neither can be understood in isolation. Economic applications: (1) Firms and industries co-evolve with technologies, institutions, and markets; (2) Routines and capabilities co-evolve with competitive environments; (3) Supply and demand co-evolve through learning and innovation. Co-evolution implies that firm-environment boundaries are endogenous—firms partially constitute the environment to which they adapt. Whereas adaptation suggests entities responding to external pressures, co-evolution emphasises the *mutual constitution* of adapting entities. *References*: Kauffman, S.A. (1993); Nelson, R.R. (1994), 47-63; Lewin, A.Y. & Volberda, H.W. (1999), 519-534.

Complex adaptive systems (CAS) (Complexity Science & Systems Theory): Systems composed of heterogeneous agents that interact locally according to simple rules, generating emergent macro-level patterns irreducible to component behaviour. CAS exhibit adaptation through learning and selection, self-organisation without central control, and evolution through variation and retention. Distinguished from merely complicated systems by their capacity for novelty generation and their sensitivity to initial conditions and history. Key properties: emergence, self-organisation, non-linearity, feedback loops, path dependence, co-evolution, fitness landscapes. *References*: Holland, J.H. (1995); Miller, J.H. & Page, S.E. (2007).

Critical realism (Philosophy of Science & Methodology): A philosophy of science developed by Roy Bhaskar combining ontological realism (reality exists independently of our knowledge of it) with epistemological relativism (knowledge is socially produced and fallible) and judgemental rationality (some theories better approximate reality than others). Distinguishes three stratified domains: the *empirical* (experienced events), the *actual* (all events), and the *real* (generative mechanisms and structures). Science aims to identify causal mechanisms through retrodiction rather than merely cataloguing empirical regularities. *References*: Bhaskar, R. (1975/2008); Archer, M., Bhaskar, R., Collier, A., Lawson, T. & Norrie, A. (Eds.). (1998).

Cybernetics (Complexity Science & Systems Theory): The interdisciplinary study of control and communication in systems, whether biological, mechanical, or social. Coined by Norbert Wiener in 1948 (from Greek *kybernētēs*, steersman), cybernetics focuses on how systems use feedback to regulate themselves and achieve goals. Core concepts: feedback loops (circular causal processes); negative feedback (reduces deviation); positive feedback (amplifies deviation); information and control (regulation through information flow). *Reference*: Wiener, N. (1948).

Decomposability / Modularity (Network Science):: *Near-decomposability* is Simon's (1962) principle that complex systems exhibit hierarchical structure wherein interactions *within* subsystems are stronger and more frequent than interactions *between* subsystems. This separation of high-frequency (internal) from low-frequency (inter-modular) dynamics enables analytical tractability and evolutionary adaptability. *Modularity* is the degree to which system components cluster into discrete units with dense internal connections and sparse external ties. Modular architecture permits independent modification of subsystems without disrupting the whole. Simon's watchmakers parable: two watchmakers building thousand-piece watches; one uses modular subassemblies, the other doesn't. Interruptions devastate the non-modular approach while modular assembly recovers quickly. *References*: Simon, H.A. (1962); Baldwin, C.Y. & Clark, K.B. (2000).

Dissipative structures / Far from equilibrium (Complexity Science & Systems Theory): *Dissipative structures* are organised, non-equilibrium states that arise in open systems through continuous exchange of energy and matter with their environment. Named for their maintenance through entropy dissipation—they import low-entropy inputs and export high-entropy waste, sustaining internal order against thermodynamic decay. *Far from equilibrium* is the condition enabling dissipative structure formation; systems driven beyond equilibrium by energy flows may spontaneously organise into new patterns at critical thresholds (bifurcation points). *References*: Prigogine, I. (1977); Nicolis, G. & Prigogine, I. (1977); Prigogine, I. & Stengers, I. (1984).

Downward causation / Upward causation (Complexity Science & Systems Theory): Causal relationships across levels in hierarchically organised systems. *Upward causation*: micro-level components determine macro-level properties; consistent with reductionist explanation. *Downward causation*: higher-level patterns causally influence lower-level parts. Coined by Donald Campbell (1974), the concept addresses how emergent system-level properties constrain and shape the behaviour of constituent elements. Mechanisms include structural constraints, boundary conditions, selective pressures, and information control. *Reference*: Campbell, D.T. (1974). In F.J. Ayala & T. Dobzhansky (Eds.), (pp. 179-186).

Ecological rationality (Rationality, Uncertainty & Decision-Making): The proposition that rationality depends on fit between decision strategy and environmental structure, rather than adherence to context-independent logical rules. Simple heuristics can outperform complex strategies when matched to appropriate environments—"less is more" effects. *Reference*: Gigerenzer, G., Todd, P.M. & the ABC Research Group (1999).

Edge of chaos (Complexity Science & Systems Theory): A transition zone between ordered (frozen) and chaotic (random) dynamical regimes where complex adaptive systems exhibit maximal computational capacity, adaptability, and evolutionary potential. Coined by Norman Packard (1988) and formalised through Christopher Langton's (1990) cellular automata research. Stuart Kauffman extended the concept to biological evolution, arguing that living systems self-organise toward this critical boundary. *References*: Langton, C.G. (1990), 12-37; Kauffman, S.A. (1993).

Emergence / Emergent properties (Complexity Science & Systems Theory): The arising of novel properties, patterns, or behaviours at macro-level that are not present in, nor straightforwardly predictable from, micro-level components. Strong emergence implies irreducibility and potential downward causation; weak emergence allows in-principle derivation from micro-level specifications though such derivation may be computationally intractable. Emergence challenges reductionist explanation while avoiding mysticism through emphasis on relational organisation. *References*: Bedau, M.A. (1997), 375-399; Goldstein, J. (1999), 49-72.

Emergent Wholeness (*Original to thesis*) (Complexity Science & Systems Theory): Emergent wholeness refers to the property by which an integrated system exhibits coherent identity, behaviour, and causal efficacy that cannot be decomposed into, derived from, or predicted by the properties of its constituent parts taken in isolation. Within the thesis, emergent wholeness designates the firm-as-system as a unified, causally efficacious entity whose characteristics arise from organised interaction, coordination, and historical development rather than from aggregation of individual actions or transactions. The concept is distinguished from simple aggregation and from non-analytic holism. *References*: (This thesis: Ch. 4 “Ontology / Emergence”; Ch. 3 “Non-decomposability / Meso-level”). [Complex Adaptive Systems] [Emergence] (Holland, 1992/1995; Sawyer, 2005; Cilliers, 1998/2005).

Entropy (Complexity Science & Systems Theory): A measure with dual interpretations in thermodynamics and information theory. *Thermodynamic entropy* (Boltzmann): measure of microstates consistent with a macrostate; the Second Law establishes that entropy in isolated systems tends to increase ($S = k_B \ln W$). *Information entropy* (Shannon, 1948): measure of uncertainty in a random variable's outcomes ($H = -\sum p_i \log p_i$); high entropy = high uncertainty. "Gain in entropy always means loss of information" (G.N. Lewis). *Reference*: Shannon, C.E. (1948), 379-423.

Epistemology (Philosophy of Science & Methodology): The philosophical study of knowledge: its nature, sources, scope, and justification. Addresses what counts as valid knowledge, how knowledge claims can be warranted, and the relationship between knower and known. Distinct from ontology (what exists) and methodology (procedures for inquiry). *Reference*: Blaikie, N. (2007), pp. 18-22.

Ergodicity / Non-ergodicity (Complexity Science & Systems Theory): *Ergodic process*: a stochastic process where time averages equal ensemble averages—the long-run behaviour of a single trajectory matches the statistical behaviour across many independent realisations. Past states do not constrain future possibilities; the system "forgets" its history. *Non-ergodic process*: time averages diverge from ensemble averages; history matters, initial conditions persist, and population statistics do not predict individual trajectories. Path dependence, lock-in, and irreversibility characterise non-ergodic dynamics. Economic processes are fundamentally non-ergodic; mainstream equilibrium models inappropriately assume ergodicity. *References*: Peters, O. (2019); Davidson, P. (1991), 129-143.

Fallibilism (Philosophy of Science & Methodology): The epistemological principle that no belief or knowledge claim can be held with absolute certainty—that all human knowledge is provisional and potentially subject to revision. Fallibilism holds that the quest for certainty is misguided; what matters is maintaining openness to correction through inquiry and criticism. Distinct from scepticism: fallibilism does not deny that knowledge is possible, only that any belief is immune from potential revision. *References*: Peirce, C.S. (1897); Popper, K.R. (1963).

Feedback loops (positive/negative) (Complexity Science & Systems Theory): Circular causal pathways where system outputs return as inputs, amplifying or dampening subsequent behaviour. *Positive feedback* (self-reinforcing): output amplifies the process producing it, generating exponential growth, collapse, or lock-in (e.g., increasing returns, bandwagon effects). *Negative feedback* (self-correcting): output dampens the process, maintaining stability around equilibrium or target state (e.g., homeostasis, thermostats). *Reference*: Meadows, D.H. (2008), Ch. 1.

First-Principles Analysis (*Original to thesis*) (Complexity Science & Systems Theory): First-principles analysis is inquiry that begins by interrogating the most basic assumptions of a theory rather than accepting them as given. It evaluates what a theory can represent, explain, and exclude by examining its foundational ontological and epistemological commitments. *References*: (This thesis: Ch. 2 “Methodology / problematization”; Ch. 3 “first-principles architecture”; Ch. 4–5 “reconstruction”). [First-principles analysis] [Problematization] (Alvesson & Sandberg, 2011; Lakatos, 1970).

Fitness landscape (Complexity Science & Systems Theory): A conceptual and mathematical representation of the relationship between configurations and their fitness, visualised as a multidimensional topography where height represents fitness. Introduced by Wright (1932); formalised by Kauffman in the NK model of “tunably rugged” landscapes. Key concepts include adaptive walks (incremental improvements toward peaks), ruggedness (determined by epistatic interactions), and local optima (peaks from which no single change improves fitness). *Reference*: Kauffman, S.A. & Weinberger, E.D. (1989), 211-245.

Fundamental uncertainty (Rationality, Uncertainty & Decision-Making): Uncertainty that cannot be reduced to quantifiable risk—situations where outcomes are unknowable in probabilistic terms. Knight's distinction: “risk” is measurable uncertainty (randomness with known/estimable probabilities; insurable); “uncertainty” proper is unmeasurable ignorance (no basis for probability assignment). Keynes: “By ‘uncertain’ knowledge... I do not mean merely to distinguish what is known for certain from what is only probable. The game of roulette is not subject, in this sense, to uncertainty.” *References*: Knight, F.H. (1921); Keynes, J.M. (1921).

Generative mechanism (Philosophy of Science & Methodology): In critical realist ontology, the enduring structures, processes, and relations within entities that produce causal powers and tendencies. Generative mechanisms operate in the “domain of the real,” generating events in the “domain of the actual” that may or may not be observed empirically. Unlike Humean constant conjunctions, mechanisms explain *how* and *why* causation occurs through structural properties of objects. Mechanisms are (1) ontologically distinct from the events they generate, (2) may operate unexercised or unactualized, (3) interact with other mechanisms in open systems, producing contingent rather than invariant outcomes. *References*: Bhaskar, R. (1975/2008); Bhaskar, R. (1979).

Hermeneutic circle (Philosophy of Science & Methodology): The interpretive principle that understanding proceeds through iterative movement between parts and whole: comprehension of individual parts depends on prior grasp of the whole, while understanding of the whole emerges from interpretation of its parts. The circle is productive, not vicious; genuine understanding involves a "fusion of horizons" between interpreter and interpreted. *References*: Gadamer, H.-G. (1960/1989); Heidegger, M. (1927/1962), §32.

Hermeneutics (Philosophy of Science & Methodology): The theory and methodology of interpretation, originally developed for biblical and legal texts, subsequently generalised to all forms of human expression and understanding. *Methodological hermeneutics* (Schleiermacher, Dilthey) seeks to establish interpretation as a rigorous method for the human sciences. *Philosophical hermeneutics* (Heidegger, Gadamer) treats understanding as an ontological condition—human beings are fundamentally interpretive creatures. *References*: Gadamer, H.-G. (1960/1989); Ricoeur, P. (1981).

Homeostasis (Complexity Science & Systems Theory): Self-regulating processes by which biological systems maintain stability while adjusting to changing conditions. Coined by Walter Cannon (1926), building on Claude Bernard's *milieu intérieur* (1865). Key features: dynamic equilibrium (continuously adjusted within range); setpoints and feedback; negative feedback as primary mechanism. From Greek *homoios* (similar) + *stasis* (standing still)—staying *similar* rather than *the same*. *Reference*: Cannon, W.B. (1932).

Incommensurability (Philosophy of Science & Methodology): The condition whereby competing paradigms or theoretical frameworks lack a common measure or neutral standard for direct comparison. Kuhn (1962) and Feyerabend (1962) independently introduced the concept to challenge cumulative views of scientific progress, arguing that paradigm shifts involve changes in concepts, methods, and observational criteria that prevent straightforward translation between frameworks. Three dimensions (Kuhn): (1) *Methodological incommensurability*: different paradigms employ different standards of evaluation; (2) *Observational incommensurability*: perceptual experience is theory-laden, so practitioners "see" different things; (3) *Semantic incommensurability*: key terms change meaning across paradigms, obstructing translation. Incommensurability does not mean non-comparability—paradigms can be compared along many dimensions—but that no neutral algorithm determines rational theory choice. *References*: Kuhn, T.S. (1962/1970), pp. 147-150; Feyerabend, P. (1962).

Increasing returns (Rationality, Uncertainty & Decision-Making): Positive feedback whereby greater adoption makes further adoption more attractive. Mechanisms: learning by using, network externalities, scale economies, informational returns, technological interrelatedness. Implies multiple equilibria, path dependence, and potential lock-in. *Reference*: Arthur, W.B. (1994).

Irreversibility (Complexity Science & Systems Theory): The property of processes that cannot return to their initial state—time has a direction and certain transformations are one-way. In thermodynamics, the Second Law establishes that entropy increases in isolated systems, creating an "arrow of time." In CAS and evolutionary economics, irreversibility characterises path-dependent processes where history shapes future possibilities. Newtonian mechanics treats time as reversible; CAS ontology emphasises that real processes are fundamentally irreversible, generating novelty and historical specificity. Economic implications: sunk costs cannot be recovered; investments lock in commitments; learning cannot be "unlearned" costlessly; institutional paths constrain future options. *References*: Prigogine, I. & Stengers, I. (1984).

Knightian uncertainty (Rationality, Uncertainty & Decision-Making): Situations where probability distributions are unknown and unknowable, distinct from "risk" where probabilities can be estimated. Knight (1921): "risk" is measurable uncertainty; "uncertainty" proper is unmeasurable ignorance. True uncertainty is the source of entrepreneurial profit. *Reference*: Knight, F.H. (1921).

Lakatosian Protective Belt (Philosophy of Science & Methodology): In Lakatos's methodology of scientific research programmes (MSRP), the flexible ring of auxiliary hypotheses, initial conditions, and observational theories that surrounds and shields the programme's irrefutable *hard core* from direct empirical refutation. When empirical anomalies arise, modifications are made to the protective belt rather than to fundamental assumptions. A programme is *progressive* if belt modifications generate novel predictions subsequently corroborated; *degenerative* if modifications merely accommodate anomalies post hoc without predictive advance. *Reference*: Lakatos, I. (1970). In I. Lakatos & A. Musgrave (Eds.), pp. 91-196.

Lock-in (Rationality, Uncertainty & Decision-Making): A state where past choices constrain future options, making switching to alternatives costly. Emerges from increasing returns, network effects, switching costs, and complementary investments. Path dependence: initial conditions have persistent, sometimes irreversible effects. *Reference*: Arthur, W.B. (1989).

Mechanistic Scientific Paradigms (*Original to thesis*) (Philosophy of Science & Methodology): Mechanistic scientific paradigms are frameworks originating in classical mechanics that treat systems as decomposable, equilibrium-oriented, reversible, and governed by stable law-like relations—privileging reductionism, linear causality, and prediction as the dominant aims of explanation. *References*: (This thesis: Ch. 3 "mechanistic inheritance"; Ch. 4 "substance ontology"; Ch. 5 "predictive ideal"). [Mechanistic scientific paradigms] (Prigogine & Stengers, 1984; Arthur, 2014; Lawson, 2003).

Metaphysical conditions (Ontology / Philosophy of Science & Methodology): The fundamental assumptions concerning what exists, what kinds of entities are real, how causation operates, and what forms of structure and process are possible, which together delimit the space of admissible explanations within a theoretical framework. Metaphysical conditions operate *prior to* epistemology and methodology: they specify the nature of reality that a theory presupposes, thereby enabling certain explanatory strategies while rendering others incoherent or inadmissible. In the context of this dissertation, metaphysical conditions determine whether firms can be treated as static, decomposable entities or must be understood as historically constituted, relational, and emergent complex adaptive systems. *References:* (This thesis: Ch. 4 “Ontological hinge”; Ch. 3 “Mechanistic inheritance”; Ch. 5 “Epistemic incompatibilities”); Bhaskar, R. (1975/2008); Lawson, T. (2003); Kuhn, T.S. (1962/1970); Lakatos, I. (1970)

Meta-theoretical Work (*Original to thesis*) (Philosophy of Science & Methodology): Meta-theoretical work is inquiry that takes **theories themselves** as the object of analysis—examining the ontological assumptions, epistemological commitments, and methodological rules that structure how theories are constructed, stabilised, and evaluated—rather than developing or testing substantive models inside those theories. *References:* (This thesis: Ch. 1 “Nature of thesis”; Ch. 3 “Meta-theoretical diagnosis”; Ch. 6 “Contributions / positioning”). [Meta-theoretical work] (Kuhn, 1962; Lakatos, 1970; Dow, 2002; Lawson, 2003).

Methodological Individualism (Philosophy of Science & Methodology): The doctrine that social phenomena must be explained by showing how they result from individual actions, which in turn must be explained through reference to the intentional states motivating individual actors. Requires that explanations of 'macro' social phenomena be supplied with 'micro' foundations specifying action-theoretic mechanisms. Distinguished from *atomism* (fully pre-social characterisation of individuals) and from *methodological holism* (explanations invoking irreducible properties of collectives). *Reference:* Weber, M. (1922/1968).

Decomposability / Modularity (Network Science):: *Near-decomposability* is Simon's (1962) principle that complex systems exhibit hierarchical structure wherein interactions *within* subsystems are stronger and more frequent than interactions *between* subsystems. This separation of high-frequency (internal) from low-frequency (inter-modular) dynamics enables analytical tractability and evolutionary adaptability. *Modularity* is the degree to which system components cluster into discrete units with dense internal connections and sparse external ties. Modular architecture permits independent modification of subsystems without disrupting the whole. Simon's watchmakers parable: two watchmakers building thousand-piece watches; one uses modular subassemblies, the other doesn't. Interruptions devastate the non-modular approach while modular assembly recovers quickly. *References:* Simon, H.A. (1962); Baldwin, C.Y. & Clark, K.B. (2000).

Network effects (Rationality, Uncertainty & Decision-Making): The value of a product/service to users increases as more users adopt it. Types: direct (value from user numbers), indirect (value from complements), two-sided markets (platforms connecting distinct groups). Creates demand-side scale economies and winner-take-all dynamics. *Reference:* Katz, M.L. & Shapiro, C. (1985).

Nomological (Philosophy of Science & Methodology): Pertaining to laws (from Greek *nomos*, law). *Nomological explanation* accounts for phenomena by subsumption under general laws—the deductive-nomological (D-N) or "covering law" model proposed by Hempel and Oppenheim (1948). To explain why event E occurred is to show that E follows deductively from (a) statements of initial conditions and (b) universal laws. Characteristics: (1) *Lawlikeness*: explanans must include at least one general law; (2) *Deductive structure*: explanandum follows logically from explanans; (3) *Empirical content*: premises must be empirically testable; (4) *Symmetry thesis*: explanation and prediction share logical structure. Social phenomena rarely exhibit the universal regularities required for nomological explanation. *References*: Hempel, C.G. & Oppenheim, P. (1948).; Hempel, C.G. (1965).

Non-Linear Causality (Complexity Science & Systems Theory): Causal relationships in which effects are not proportional to causes, such that small inputs may produce large outputs (amplification), large inputs may produce minimal outputs (dampening), or identical inputs may produce different outputs depending on system state (context-dependence). Arises from feedback loops, interaction effects, threshold effects, and sensitive dependence on initial conditions. In complex adaptive systems, non-linear causality produces emergent phenomena, path dependence, and inherent limits to predictability. *Reference*: Strogatz, S.H. (2015).

Non-Linear Dynamics (Complexity Science & Systems Theory): The mathematical study of systems governed by equations where outputs are not linear functions of inputs. Unlike linear systems (where superposition applies), non-linear systems exhibit: *sensitive dependence on initial conditions* (the 'butterfly effect'); *bifurcations* (qualitative shifts from gradual parameter changes); *attractors* including strange attractors with fractal structure; *limit cycles* and *chaos*—deterministic yet unpredictable behaviour. Demonstrates that even simple deterministic systems can produce apparently random behaviour, establishing inherent limits to long-term prediction. *Reference*: Lorenz, E.N. (1963).

Non-linearity (Complexity Science & Systems Theory): The property of systems where outputs are not proportional to inputs; small causes can produce large effects and vice versa. Non-linear systems exhibit threshold effects, feedback amplification, and sensitivity to initial conditions. Most real-world complex systems are non-linear; linear models provide tractable approximations but miss essential dynamics. *Reference*: Strogatz, S.H. (1994).

Ontological Commitments (Ontology): The existential presuppositions of a theory—what kinds of entities the theory assumes to exist in order for its claims to be true. Following Quine: "A theory is committed to those and only those entities to which the bound variables of the theory must be capable of referring in order that the affirmations made in the theory be true." Ontological commitments are often implicit; meta-theoretical analysis surfaces these hidden assumptions. A theory's ontological commitments constrain what it can explain. *Reference*: Quine, W.V.O. (1948).

Ontological uncertainty (Rationality, Uncertainty & Decision-Making): Uncertainty arising from reality's fundamental nature rather than incomplete knowledge—genuine indeterminacy rather than epistemic limitation. Sources: open systems facing novel perturbations; emergence; human agency introducing novelty; reflexivity in social systems. Contrast with epistemic uncertainty: epistemic = observer's ignorance (system has definite state); ontological = genuine indeterminacy (future objectively open). *Reference*: Dequech, D. (2004).

Ontology: The branch of metaphysics concerned with the nature of being and existence—what kinds of entities exist and their fundamental categories. In philosophy of science, ontological commitments specify what a theory assumes to exist. Distinguished from epistemology (how we know) and methodology (how we investigate). *Reference:* Blaikie, N. (2007), pp. 13-18.

Open system / Closed system (Complexity Science & Systems Theory): *Open system:* a system that exchanges matter, energy, and information with its environment, enabling maintenance of organisation far from thermodynamic equilibrium. Open systems achieve dynamic steady states by importing negative entropy. *Closed system:* a system isolated from material exchange with its environment, though energy transfer may occur. Closed systems inevitably progress toward thermodynamic equilibrium and maximum entropy. *References:* von Bertalanffy, L. (1950); von Bertalanffy, L. (1968).

Paradigm (Philosophy of Science & Methodology): A constellation of shared beliefs, values, techniques, and exemplary achievements that define legitimate problems and methods for a scientific community. Paradigms are incommensurable—practitioners of different paradigms may talk past each other due to divergent assumptions about reality, causation, and evidence. Scientific revolutions occur when anomalies accumulate and a new paradigm displaces the old through gestalt shift rather than incremental proof. *Reference:* Kuhn, T.S. (1962/1970).

Paradigm-Level Doctoral Research (Philosophy of Science & Methodology) (*Original to thesis*): Paradigm-level doctoral research is research conducted at the level of foundational reconstruction, addressing ontological, epistemological, and methodological assumptions that structure an entire field of inquiry. Such research is evaluated primarily on internal coherence, analytical depth, and generative capacity rather than on incremental model development or bounded empirical testing. *References:* (This thesis: Ch. 1 “Nature of thesis”; Ch. 6 “contribution / limits”). [Paradigm-level doctoral research] (Kuhn, 1962; Lakatos, 1970).

Path dependence (Complexity Science & Systems Theory): The property whereby outcomes depend not only on current conditions but on the historical sequence of prior states and decisions. Early events, including accidents, can lock systems into trajectories that persist even when superior alternatives exist. Distinguished from mere history-dependence by irreversibility—past choices constrain future possibility space. *References:* Arthur, W.B. (1989); David, P.A. (1985).

Phase transition / Critical point (Complexity Science & Systems Theory): A qualitative change in a system's macroscopic state as a control parameter crosses a threshold value (the critical point). First-order transitions involve abrupt jumps; second-order (critical) transitions exhibit gradual change with diverging fluctuations and power-law scaling. Near critical points, systems show critical slowing down, increased variance, and long-range correlations—potential early warning signals. *Reference:* Scheffer, M. et al. (2009).

Phenomenology (Philosophy of Science & Methodology): A philosophical tradition and method, founded by Edmund Husserl, that investigates the structures of conscious experience as they present themselves. Phenomenology "brackets" (epoché) questions of external existence to focus on the essential structures of experience itself: intentionality, temporality, embodiment, and intersubjectivity. Methodological principle: "To the things themselves" (*zu den Sachen selbst*)—return to direct experience before theoretical construction. *References*: Husserl, E. (1913/1983); Merleau-Ponty, M. (1945/1962).

Positive heuristic / Negative heuristic (Philosophy of Science & Methodology): Methodological rules governing research within a Lakatosian research programme. The *negative heuristic* specifies what paths of research to avoid—directing the *modus tollens* away from the hard core. When empirical challenges arise, the negative heuristic forbids questioning the hard core, instead directing modification to the protective belt. The *positive heuristic* is a partially articulated set of suggestions on how to develop the research programme—how to modify, sophisticate, and extend the protective belt. *Reference*: Lakatos, I. (1970). In I. Lakatos & A. Musgrave (Eds.), pp. 91-196.

Positivism (Philosophy of Science & Methodology): A philosophical orientation holding that genuine knowledge derives exclusively from sensory experience and logical/mathematical reasoning, excluding metaphysical speculation. In its original Comtean form (1830s), positivism envisioned science progressing through theological and metaphysical stages to positive (scientific) maturity. Logical positivism (Vienna Circle, 1920s-30s) formalised this through the *verification principle*: statements are cognitively meaningful only if empirically verifiable or analytically true. Key tenets: (1) Unity of scientific method across natural and social sciences; (2) Fact-value distinction; (3) Covering-law model of explanation; (4) Theory-neutral observation language. Critiques by Popper, Quine, Kuhn, and Hanson undermined logical positivism's foundational claims. *References*: Comte, A. (1830-42); Carnap, R. (1928); Ayer, A.J. (1936).

Power law (distribution) (Network Science): A functional relationship where one quantity varies as a power of another: $P(x) \propto x^{-\gamma}$, appearing as a straight line on logarithmic axes. Unlike normal distributions, power laws are "scale-free"—exhibiting the same relative proportions at different magnitudes. Related forms include Pareto distribution, Zipf's law, and heavy-tailed distributions found in city sizes, firm sizes, income, citations, and network degrees. *Reference*: Newman, M.E.J. (2005).

Preferential attachment (Network Science): A network growth mechanism whereby new nodes connect to existing nodes with probability proportional to current degree—"the rich get richer." Originally described by Price (1976) as "cumulative advantage" in citation networks; generates power-law degree distributions and hub formation. *References*: Price, D.J. de Solla (1976); Barabási, A.-L. & Albert, R. (1999).

Problematisation (Philosophy of Science & Methodology): A research methodology for generating research questions by identifying and challenging the assumptions underlying existing literature, rather than merely "spotting gaps" in prior work. Aims to produce more interesting and influential theories by questioning what is typically taken for granted. Typology of assumptions includes: in-house, root metaphor, paradigm, ideology, and field assumptions. *Reference*: Alvesson, M. & Sandberg, J. (2011).

Procedural rationality (Rationality, Uncertainty & Decision-Making): Rationality assessed by decision-making *processes* rather than outcome optimality. Simon's distinction: substantive rationality = behaviour appropriate to achieving goals (evaluated by outcomes); procedural rationality = outcome of appropriate deliberation (evaluated by process quality). Key elements: satisficing; sequential attention; heuristics; aspiration adaptation. *Reference*: Simon, H.A. (1976).

Process ontology (Ontology): An ontological framework that takes processes, events, and relations as ontologically primary, treating apparently stable "things" as derivative patterns of activity. Process ontology holds that becoming is more fundamental than being—that reality consists of interconnected processes of change rather than enduring substances undergoing modification. *References*: Whitehead, A.N. (1929).

Recombination (Complexity Science & Systems Theory): The generation of novelty through new combinations of existing elements, components, or routines. In evolutionary economics, recombination (alongside mutation and selection) is a primary mechanism of economic change, explaining how innovation emerges from the reassembly of established knowledge, technologies, and organisational practices rather than *de novo* creation. Schumpeter (1934) defined innovation as "carrying out new combinations" of productive means—new products, methods, markets, supply sources, or organisational forms. Organisational dimensions: (1) *Knowledge recombination*: integrating diverse knowledge bases; (2) *Routine recombination*: adapting and combining established practices; (3) *Network recombination*: forming new inter-organisational linkages. *References*: Schumpeter, J.A. (1934); Nelson, R.R. & Winter, S.G. (1982); Kogut, B. & Zander, U. (1992).

Reconstruction Cascade (*Original to thesis*) (Philosophy of Science & Methodology): The reconstruction cascade is the methodological process by which foundational assumptions of mainstream economic theory are systematically identified, problematised, and replaced with alternatives consistent with complex adaptive systems. "Cascade" indicates that changes at one foundational level propagate through the theoretical structure. Stages: Identification → Problematisation → Replacement → Integration → Extension. *References*: (This thesis: Ch. 3 "Assumption system"; Ch. 4 "FROM/TO ontology"; Ch. 5 "epistemic recalibration"; Ch. 6 "contribution and extension"). [Reconstruction cascade] [Paradigm reconstruction] (Kuhn, 1962; Lakatos, 1970; Alvesson & Sandberg, 2011).

Reductionism (Philosophy of Science & Methodology): The methodological or ontological position that phenomena at higher levels of organisation can and should be explained entirely by lower-level components and processes. Methodological reductionism pursues decomposition as research strategy; ontological reductionism asserts higher-level properties are nothing but lower-level properties differently arranged. Opposed by emergentism and holism. *Reference*: Nagel, E. (1961), Ch. 11.

Relational ontology (Ontology): An ontological framework holding that relations are constitutive of entities rather than merely connecting pre-existing substances. In relational ontology, entities gain their identity and properties through their relations with other entities—what something *is* depends fundamentally on what it is *related to*. Emirbayer's "Manifesto for a Relational Sociology" argues that social analysis should begin with relations and processes rather than with pre-given individuals or social structures. *References*: Emirbayer, M. (1997); Donati, P. (2011).

Research programme (Philosophy of Science & Methodology): A series of interconnected theories sharing a common *hard core* of fundamental assumptions protected from falsification by a *protective belt* of auxiliary hypotheses. Progressive programmes generate novel predictions subsequently corroborated; degenerative programmes merely accommodate anomalies post hoc without predictive advance. Provides criterion for rational theory choice absent Kuhnian incommensurability. The *hard core* consists of irrefutable foundational propositions defining the programme's identity; the *protective belt* consists of adjustable auxiliary hypotheses modified to absorb empirical challenges. *Reference*: Lakatos, I. (1970), pp. 91-196.

Robustness (Rationality, Uncertainty & Decision-Making): The ability of a system to maintain function across a range of conditions without requiring adjustment. Robust systems continue working despite variation in parameters or operating conditions. "Robust yet fragile": robustness to one perturbation type may create fragility to others. *Reference*: Carlson, J.M. & Doyle, J. (2002).

Scale-free network (Network Science): A network whose degree distribution follows a power law $P(k) \propto k^{(-\gamma)}$. Unlike random networks, scale-free networks exhibit extreme heterogeneity with few highly connected "hubs." Generated through preferential attachment ("rich get richer"). Properties include robustness to random failure but vulnerability to targeted attack on hubs. *Reference*: Barabási, A.-L. & Albert, R. (1999).

Self-organisation (Complexity Science & Systems Theory): The spontaneous emergence of organised patterns, structures, or behaviours through local interactions among system components without external direction or centralised control. Arises in systems far from equilibrium through non-linear feedback dynamics. Distinguished from designed organisation by absence of blueprint or orchestrator; order emerges from the bottom up through iterative local interactions. *References*: Ashby, W.R. (1947); Kauffman, S.A. (1993); Prigogine, I. & Stengers, I. (1984). *Order Out of Chaos*. New York: Bantam.

Sensitivity to initial conditions (Complexity Science & Systems Theory): A defining property of chaotic dynamical systems whereby arbitrarily small differences in starting states lead to exponentially diverging trajectories over time. Discovered by Lorenz (1963); popularly known as the "butterfly effect." Implies that determinism does not guarantee predictability—chaotic systems are deterministic but practically unpredictable beyond short horizons. *References*: Lorenz, E.N. (1963).; Lorenz, E.N. (1993).

Small world networks (Network Science): Network structures exhibiting simultaneously high clustering (like regular lattices) and short average path lengths (like random graphs). Arise when a few random "shortcuts" are added to regular lattice structures, dramatically reducing average path lengths while preserving local clustering. Key characteristics: high clustering coefficient; short characteristic path length; "six degrees of separation" phenomenon. The Watts-Strogatz model: beginning with a regular ring lattice, edges are randomly "rewired" with probability p ; at low rewiring probabilities, path length drops precipitously while clustering remains high. *Reference*: Watts, D.J. & Strogatz, S.H. (1998), 440-442.

Stratified ontology (Philosophy of Science & Methodology): The critical realist view that reality consists of nested domains and emergent levels, each possessing distinctive causal powers irreducible to lower strata. Bhaskar distinguishes three domains: the *empirical* (observed experiences), the *actual* (events whether observed or not), and the *real* (generative mechanisms and structures). Reality is also stratified vertically—physical, chemical, biological, psychological, social—with each stratum exhibiting emergent properties. *References*: Bhaskar, R. (1975/2008); Archer et al. (1998).

Structural holes (Network Science): Gaps in information flow between non-redundant contacts in a network—the absence of ties between clusters. Actors bridging structural holes gain competitive advantage through: (1) early access to diverse, non-redundant information; (2) control over information flows between disconnected parties; (3) brokerage opportunities. Complements Granovetter's weak ties thesis with explicit network-structural mechanism. *References*: Burt, R.S. (1992); Burt, R.S. (2004).

Substance ontology (Ontology): An ontological framework, originating with Aristotle, in which reality is fundamentally composed of enduring substances—discrete entities possessing essential properties that persist through change. Substances are the primary category of being; properties, relations, and processes are ontologically secondary. Change is understood as the modification of substances' accidental properties while their essential nature remains stable. *References*: Aristotle. *Metaphysics*, Book VII; Whitehead, A.N. (1929).

Systems Thinking (Complexity Science & Systems Theory): A holistic intellectual orientation that emphasises understanding wholes through the interrelationships among their parts rather than through analysis of isolated components. Core principles: (1) *Interconnectedness*: system components are linked through mutual influence; (2) *Feedback awareness*: circular causation where effects become causes; (3) *Emergence*: system behaviour cannot be predicted from component properties alone; (4) *Dynamic complexity*: behaviour arises from interactions over time; (5) *Synthesis over analysis*: understanding the whole requires seeing patterns of relationship. *References*: Senge, P.M. (1990); Meadows, D.H. (2008). *Thinking in Systems: A Primer*. White River Junction, VT: Chelsea Green.

Tacit knowledge (Rationality, Uncertainty & Decision-Making): Knowledge difficult to articulate or codify—"we can know more than we can tell" (Polanyi). Embodied in skills and intuitions acquired through experience. Personal, practice-based, context-dependent, difficult to transfer. Source of competitive advantage because hard to imitate. *References*: Polanyi, M. (1966); Nonaka, I. & Takeuchi, H. (1995).

Teleological (Philosophy of Science & Methodology): Pertaining to explanation by reference to purposes, goals, or final causes (*telos* = end, purpose). Aristotle distinguished four types of cause: material, formal, efficient, and *final*—the "that for the sake of which" something exists or occurs. Teleological explanation answers "why?" by identifying the end or function served, contrasting with efficient-causal explanation that identifies antecedent causes. Teleology was central to Aristotelian natural philosophy but rejected by the mechanical philosophy of the Scientific Revolution as anthropomorphic projection. Neoclassical theory is implicitly teleological—agents pursue utility/profit maximisation, markets tend toward equilibrium. *References*: Aristotle. *Physics*, Book II; Mayr, E. (1974), 91-117.

Theoretical Pluralism (Philosophy of Science & Methodology): The position that multiple theoretical frameworks may be legitimately employed to understand complex phenomena, without requiring synthesis into a single unified theory. Justified by ontological complexity, epistemic humility, and heuristic value of alternative perspectives. Distinguished from eclecticism (unprincipled mixing) and relativism—pluralism requires theories to meet standards of internal coherence and empirical engagement. Caldwell's 'critical pluralism' requires frameworks to enable mutual criticism across paradigms. *References*: Caldwell, B.J. (1997); Dow, S.C. (2004).

Tipping point (Complexity Science & Systems Theory): A critical threshold at which a small perturbation triggers a large, often irreversible, shift in system state. Characterised by nonlinearity, potential hysteresis, and cascade effects. Near tipping points, systems exhibit critical slowing down, increased variance, and flickering between states. *Reference*: Scheffer, M. (2009).

Triple Lock (*Original to thesis*) (Philosophy of Science & Methodology): The Triple Lock denotes the mutually reinforcing configuration of ontological, epistemological, and methodological commitments that stabilise the dominant economic paradigm and render it resistant to fundamental critique. The lock operates because each dimension presupposes and validates the others: **Ontological lock**: atomistic agents, aggregation, equilibrium as the natural state; **Epistemological lock**: privileging deductive derivation and econometric testing while marginalising alternative forms of evidence; **Methodological lock**: mathematical formalism and methodological individualism as dominant standards, with prediction as the primary success criterion. Together, these commitments form a self-sealing architecture that absorbs anomalies without altering foundational assumptions. *References*: This thesis: Ch. 3 “Assumptive architecture”; Ch. 4 “Mechanistic ontology”; Ch. 5 “Predictive ideal / epistemic regime”). [Assumptive architecture] [Methodological monism] Kuhn (1962); Lakatos (1970); Lawson (2003); Arthur (1994).

Weak ties (Network Science): Social connections characterised by infrequent interaction, low emotional intensity, and limited mutual acquaintance—contrasted with strong ties of close friendship and family. Granovetter's (1973) "strength of weak ties" thesis demonstrates that weak ties provide superior access to novel information precisely because they bridge otherwise disconnected social clusters. *Reference*: Granovetter, M.S. (1973).

Chapter 1 – Introduction: Reconstructing the Firm for a Complex World

1.1 The Problem: Why the Firm Remains Theoretically Elusive

What are firms? For nearly a century after Coase posed the question, economic theory still lacks a coherent answer. Firms organise the vast majority of production in modern economies, coordinate knowledge that no individual possesses, generate the innovations upon which prosperity depends, and mediate continuously between markets, institutions, and technological change. Yet across dominant economic traditions, this central stakeholder remains and the theories that have been developed lack theoretical coherence, are theoretically unstable, appearing variously as a production function, a contractual nexus, or an aggregation device whose internal organisation is analytically secondary.

This instability is not a consequence of neglect. Over the past decades, economic theory has devoted substantial effort to improving its treatment of firms. Transaction cost economics, principal–agent models, property rights theory, behavioural extensions, evolutionary approaches, and institutional analyses have each addressed specific shortcomings of earlier formulations. These contributions have significantly enriched the literature. Nevertheless, persistent difficulties remain. Core features of real firms – organisational capability, adaptive coordination, innovation under uncertainty, resilience and failure, and the endogenous evolution of routines – continue to resist stable theoretical representation.

These difficulties are often interpreted as empirical anomalies or as indicators of incomplete modelling. Forecast errors, breakdowns of equilibrium reasoning, crisis blindness, and the fragility of predictive claims are commonly attributed to missing variables, imperfect information, or insufficient formalisation. In this view, further refinement – more data, better models, improved behavioural assumptions – appears as the natural remedy.

This dissertation advances a different diagnosis. It argues that the persistent elusiveness of the firm is not primarily an empirical or technical problem, but a structural one. The difficulty lies not in the absence of more elaborate theories or more sophisticated tools, but in the foundational assumptions that define what firms are taken to be, what can be known about them, and how they can legitimately be analysed. When these assumptions are misaligned with the properties of the object under study, theoretical and methodological refinements stabilise the problem, rather than resolving them.

The misalignment runs deeper than the theories themselves. Every question a theory poses, and every tool devised to answer it, encodes ontological and epistemological commitments – and determines the methodology they permit. In complex adaptive systems, such as firms and economies, these encoded assumptions, although mostly hidden, do much of the theoretical work before any formal analysis begins. When they are ill-suited to the system under investigation, the resulting theories may be internally consistent yet explanatorily empty. The problem, then, is not that we lack answers – it is that the questions themselves have been improperly specified.

Accordingly, the challenge addressed in this dissertation is not how to improve existing models of the firm, but how to reconsider the conceptual conditions under which firms are theorised in

the first place – to reconstruct the foundations that determine what kinds of questions about firms are meaningful, tractable, and scientifically legitimate, in a complex and uncertain world.

1.2 Limits of Incremental Reform in Firm Theory

A defining feature of the literature on the firm is its cumulative character. Successive research programmes have sought to correct the limitations of earlier approaches by introducing additional mechanisms, relaxing restrictive assumptions, or extending analytical scope. Information frictions, bounded rationality, incomplete contracts, institutional constraints, heterogeneous agents, and evolutionary dynamics have all been incorporated in response to recognised weaknesses of standard formulations.

While these efforts have yielded important insights, they have not produced a stable or unified theoretical understanding of the firm. Instead, they have generated an increasingly elaborate landscape of partially connected models, each addressing specific phenomena while leaving others unresolved. The firm remains theoretically fragmented: sometimes reduced to individual incentives, sometimes absorbed into market outcomes, and sometimes treated as a black box whose internal dynamics remain untheorised.

This pattern suggests a fundamental limit to incremental reform. Additions at the level of mechanisms or variables may improve local empirical fit, but they do not alter what counts as an explanation in the first place. Yet it is precisely at this deeper level that the boundaries of achievable theory are drawn. Theories are not self-contained constructions: they are deductions from more general theoretical frameworks, which themselves rest on still more fundamental commitments. At the base lies a worldview (*Weltanschauung*) – a set of core assumptions about the nature of reality, causation, and knowledge from which all subsequent theoretical claims ultimately derive. When these foundational commitments are misaligned with the properties of the system under study, no accumulation of refinements at lower levels can resolve the resulting mismatch.

Existing theories of the firm operate within a shared conceptual architecture that constrains what can appear as analytically real. This architecture is historically rooted in the mechanistic worldview that emerged during the Renaissance, which conceives of systems as decomposable, law-governed, and fully intelligible through the analysis of their constituent parts. Within this framework, firms are treated as equilibrium-seeking entities reducible to more fundamental units of analysis, such as contracts, incentives, or individual optimisation. Organisational coherence, emergent capability, and adaptive transformation are therefore not merely under-modelled; they are rendered ontologically invisible. Yet these are central characteristics of firms embedded in an increasingly complex, non-linear, and evolving economic world. What cannot be seen cannot be explained – and so these phenomena are relegated to auxiliary assumptions, treated as residual effects, or excluded from the explanatory domain altogether.

As a result, reform efforts frequently exhibit a compensatory character: when one assumption is relaxed, others are reinforced to preserve internal consistency. Behavioural departures from rationality are absorbed without challenging optimisation as the organising principle; institutional constraints are added without altering underlying notions of closure or equilibrium; evolutionary language is introduced without revising the ontology of the firm itself.

Such patterns are not accidental. They reflect what Kuhn identified as the stabilising logic of paradigms – self-reinforcing structures that define what counts as a legitimate problem, explanation, and solution. Within such frameworks, persistent anomalies are managed through auxiliary adjustments, rather than taken as grounds for reconsidering first principles. Meaningful progress, therefore, cannot be achieved through further piecemeal extension alone. It requires intervention at the level of first principles – where the commitments that shape theorising itself are brought into question.

1.3 A Paradigmatic Diagnosis: When the Framework Becomes the Constraint

To intervene at the level of first principles is to move from theory to paradigm. Paradigms are not individual models or hypotheses. They are coherent configurations of ontological assumptions, epistemological standards, and methodological norms that define what counts as a legitimate object of inquiry, what forms of knowledge are admissible, and what constitutes explanation or rigour within a field.

This dissertation adopts a paradigmatic diagnosis. It does not seek to adjudicate among competing theories of the firm by proposing a superior model within the existing framework. Instead, it interrogates the shared assumptions that structure those theories, asking whether they remain coherent once firms are recognised as open, evolving, and internally organised systems operating under fundamental uncertainty.

The dissertation aligns with traditions of paradigm analysis associated with Kuhn and Lakatos, drawing on their concepts in a precise and complementary manner. In Kuhn's account, a paradigm is the background framework of shared assumptions, exemplars, and standards that defines what counts as a legitimate problem, method, and explanation within a field; it structures normal science by guiding inquiry while also determining what can appear as an anomaly. Lakatos refines this picture by describing scientific development as competition between research programmes organised around a relatively stable hard core of foundational commitments, surrounded by a more flexible protective belt of auxiliary hypotheses that can be adjusted to accommodate anomalies and preserve the programme's coherence. Together, these ideas clarify the level at which the present inquiry is situated. The dissertation is not concerned with proposing yet another modification to the protective belt of firm theory, but with interrogating the hard core – the shared first-principles assumptions that stabilise the prevailing paradigmatic approach to the firm and delimit what can count as an explanation of organisational coherence, emergence, and adaptive transformation.

This paradigmatic stance does not entail rejection of existing economic theories. On the contrary, the dissertation treats them as analytically valuable but architecturally constrained. Their limitations are not failures of intelligence or technique, but consequences of assumptions inherited from earlier scientific paradigms whose applicability to complex organisational systems has never been systematically examined.

The task, therefore, is reconstructive rather than destructive. By making the underlying architecture explicit and subjecting it to systematic analysis, this dissertation develops the ontological, epistemological, and methodological foundations upon which firms can be coherently reconceptualised as complex adaptive systems.

1.4 Why Complexity – and Why Now

The Case for Complexity

The preceding section established that meaningful progress in firm theory requires intervention at the level of first principles – the foundational assumptions from which theories are deduced. This section argues that complexity science provides the appropriate basis for such intervention: not as theoretical fashion, nor as dissatisfaction with formal economics as such, but as a response to a structural mismatch between the inherited framework and the properties of the systems it purports to explain.

The Firm as Complex Adaptive System

The proposal to treat firms as complex adaptive systems arises from a convergence of empirical realities, analytical tensions, and conceptual developments that have rendered the mechanistic paradigm increasingly strained when applied to real organisational systems. A complex adaptive system is one composed of heterogeneous, interacting agents whose local interactions generate emergent properties at the system level – properties that cannot be derived from the components considered in isolation, no matter how complete the description. Such systems are open to their environments, exchanging energy, matter, and information across permeable boundaries that function as selective membranes rather than fixed walls. They evolve through path-dependent processes in which history is causally active: past choices sediment into routines, capabilities, and structures that constrain and enable present action. They operate under fundamental uncertainty – not merely imperfect information about a knowable future, but genuine indeterminacy arising from the continuous generation of novelty. They exhibit non-linear dynamics in which small causes can produce large effects and large interventions negligible ones, feedback loops that amplify or dampen perturbations, and downward causation through which system-level structures shape the behaviour of their components. They maintain themselves through adaptive tension at the edge of chaos – preserving sufficient structure for coordination yet sufficient flexibility for innovation – rather than settling into equilibrium. Firms, on this view, are not machines executing predetermined functions, but historically evolving, multi-level process-patterns: relationally constituted configurations of routines, capabilities, and coordination whose identity consists in self-reproducing patterns of activity rather than fixed essence, and whose properties emerge at the meso-level where neither micro-level reduction nor macro-level aggregation can reach.

The Mismatch Amplified

Complexity enters this dissertation as a response to a mismatch between the properties attributed to firms by dominant theories and the properties firms demonstrably exhibit. Modern firms operate in environments characterised by deep uncertainty, technological recombination, institutional instability, and rapid structural change. They learn, adapt, coordinate heterogeneous knowledge, and evolve through path-dependent processes that resist decomposition into isolated decisions or stable equilibria. These properties have always been constitutive of firm behaviour. What distinguishes the present moment is that accelerating complexity has amplified their visibility and significance to the point where they can no longer be dismissed as anomalies or relegated to the margins of theory.

These dynamics have become pervasive across sectors, scales, and institutional contexts. Innovation is increasingly systemic rather than discrete; competitive advantage is transient rather than stable; coordination relies on distributed architectures rather than centralised control; failure is often endogenous rather than externally induced. Under such conditions, treating the firm as a decomposable optimiser operating within a closed and well-defined environment becomes progressively less defensible as a general analytical stance.

The appeal to complexity is not new. Ideas associated with open systems, feedback, emergence, and adaptation have circulated across economics, organisation theory, and systems thinking for decades. Earlier engagements, however, were typically relegated to the margins of theory, treated as descriptive supplements, or absorbed into existing frameworks without altering their foundational commitments. Moreover, even where complexity was taken seriously, its practical implications could not be operationalised; the computational resources required to model, simulate, and implement complexity-informed approaches simply did not exist. Complexity was acknowledged rhetorically while remaining conceptually domesticated and practically inert.

Why Now: Five Convergent Developments

This dissertation argues that such domestication is no longer viable – and, for the first time, no longer necessary. Several developments have converged to make the limits of the inherited framework increasingly visible and its replacement increasingly feasible.

First, repeated economic crises and organisational failures have exposed the fragility of equilibrium-centred reasoning and predictive confidence. The 2008 financial crisis revealed not isolated modelling errors but the systematic consequences of applying mechanistic assumptions to adaptive, interconnected systems. Equilibrium, linearity, closure, ergodicity, and representative-agent assumptions failed simultaneously – not because of insufficient data or imperfect calibration, but because the framework's foundational commitments were misaligned with the system's actual properties. Subsequent disruptions – sovereign debt crises, supply chain collapses, pandemic-induced transformations – have reinforced the diagnosis. These are not anomalies to be absorbed through auxiliary adjustment; they are structurally produced outcomes of applying machine logic to complex adaptive systems.

Second, advances in organisational studies, evolutionary economics, and institutional analysis have documented the centrality of phenomena that cannot be reduced to individual optimisation. Routines, capabilities, organisational learning, and dynamic adjustment have been shown to operate at the meso-level – the intermediate scale between individual agents and aggregate outcomes – where the firm's distinctive properties actually reside. This meso-level has been systematically erased by the mechanistic framework's commitment to micro-to-macro aggregation. Complexity science restores it to analytical visibility, providing conceptual resources for theorising what the inherited architecture renders invisible.

Third, complexity science itself has matured. What began as suggestive metaphor has developed into rigorous conceptual architecture. The principles of emergence, self-organisation, feedback, path dependence, and adaptive dynamics are now sufficiently well-articulated to support theoretical reconstruction rather than impressionistic analogy. Formal tools – network analysis, agent-based modelling, non-linear dynamics – have been refined to

the point where they can be deployed with discipline. The charge that complexity is merely hand-waving no longer holds; the question is whether economists will engage with these developments or continue to insulate the discipline from them.

Fourth, the exponential growth in computational power has transformed complexity-informed approaches from abstract theoretical possibilities into implementable analytical practices. Simulating emergent dynamics, modelling heterogeneous agent interactions, and tracking path-dependent processes were computationally prohibitive even two decades ago. They are now routine. This shift has practical consequences: recommendations grounded in complex systems thinking are, for the first time, realisable in organisational and policy contexts.

The emergence of artificial intelligence amplifies both dimensions of this transformation. On one hand, AI provides unprecedented capacity to implement complexity-informed approaches. Pattern recognition under uncertainty, navigation through evolving possibility spaces, multi-scale analysis, scenario-based reasoning – these were previously constrained by computational limits that no longer bind. AI systems can, for the first time, make the insights of complex systems science actionable for real-world decision-making. On the other hand, the widespread deployment of AI is accelerating the very complexity that renders the inherited framework inadequate. Organisational boundaries are becoming more fluid, coordination architectures more distributed, adaptation cycles shorter, and emergent dynamics more consequential. The mechanistic framework – designed for stable, decomposable, predictable systems – provides no adequate conceptual resources for understanding these developments, let alone managing them at the speed now required. A CAS-informed reconceptualisation of the firm is therefore not merely theoretically preferable but practically necessary: it offers a coherent framework for understanding firms as entities that must navigate, incorporate, and be transformed by technologies that are themselves emergent, adaptive, and fundamentally unpredictable.

Fifth, recognition of the need for foundational reassessment is no longer confined to heterodox critics. Economics has reached a stage of theoretical maturity at which cumulative refinement has revealed its own limits. The proliferation of models, extensions, and qualifications has increased sophistication without delivering integration. Prominent voices from within the mainstream have increasingly acknowledged that something more fundamental may be required. When a paradigm's defenders begin to question its foundations, the conditions for productive reconstruction are present.

The convergence of these developments – empirical crises exposing mechanistic fragility, meso-level research documenting irreducible organisational phenomena, complexity science's maturation into rigorous architecture, computational capacity enabling implementation, AI simultaneously enabling and accelerating complexity, and growing mainstream recognition of foundational limits – creates conditions in which complexity can be taken seriously not merely as modelling heuristic but as foundational commitment. The question that remains is not whether complexity offers resources for reconceptualising the firm, but how this dissertation's particular contribution relates to the complexity economics tradition that has already achieved genuine reconstruction at the level of markets and economies. That question – concerning the scope, positioning, and distinctive contribution of the present work – is the subject of the following section.

1.5 Scope, Contribution, and Boundaries

1.5.1 Scope of the Inquiry

Firms are central to economic life, yet economic theory lacks a coherent account of what they are. Transaction cost economics asks why firms exist; agency theory asks how principals control agents; the resource-based view asks what makes firms persistently different. But a prior question remains unanswered: what kind of entity must the firm be assumed to be for such questions to make sense at all? This dissertation addresses that prior question.

The inquiry focuses on the firm itself, rather than on isolated markets, policies, or functional outcomes. Its aim is not to displace established analytical tools or models, but to examine the assumptions about firms that silently shape how such tools are used and interpreted. From this perspective, the theory of the firm functions as a connecting framework: it mediates between economic reasoning, organisational analysis, and practical decision-making.

The dissertation operates at the level of paradigm rather than model. Paradigms, in this usage, are not individual theories or schools of thought, but structured configurations of assumptions that define what kinds of entities firms are taken to be, what questions can meaningfully be asked about them, and what forms of explanation count as legitimate. Theories are developed within such frameworks; they do not define them.

A crucial distinction governs the analysis throughout. Complicated systems—such as aircraft or supply chains—are decomposable: their behaviour can be understood by analysing components in isolation and aggregating the results. Complex adaptive systems—such as firms, markets, and ecosystems—are not fully decomposable: their defining properties emerge from interaction and cannot be recovered through analytical separation alone. The mechanistic paradigm is appropriate for complicated systems; this dissertation addresses complex adaptive systems, for which that paradigm fails.

To conduct this investigation, the inquiry operates across three interdependent domains:

- **Ontology**, concerning what kind of entity the firm must be assumed to be.
- **Epistemology**, concerning what can be known about such an entity, and with what limits, given the ontological commitments.
- **Methodology**, concerning how inquiry must proceed given those ontological and epistemological commitments.

These domains form a dependency chain imposed by the object of inquiry itself. Ontological commitments about what the firm is determine what can coherently be claimed about it; both jointly determine what methods of inquiry are appropriate. This is not a conventional philosophical taxonomy but a constraint structure: get the ontology wrong, and epistemological and methodological errors follow necessarily.

1.5.2 Positioning Within Complexity Economics

Over the past three decades, Complexity Economics has emerged as a recognised field, associated particularly with the Santa Fe Institute and scholars such as Brian Arthur and Eric Beinhocker. This tradition has achieved genuine reconstruction: rejecting equilibrium as an imposed analytical requirement, embracing fundamental uncertainty, replacing representative agents with heterogeneous populations, and demonstrating path dependence and increasing returns through agent-based methods.

The existence of Complexity Economics with these achievements raises an obvious question: what does this dissertation contribute?

This dissertation is situated within that tradition—sharing its rejection of mechanistic assumptions and its embrace of emergence, adaptation, and fundamental uncertainty. The contribution lies not in challenging Complexity Economics, but in extending it to a domain it has not systematically reconstructed at the paradigmatic level: the firm itself.

Complexity Economics has developed primarily as a modelling tradition focused on markets and economies. Its characteristic contributions—the artificial stock market, technology lock-in models, network cascades—treat firms as adaptive but internally undifferentiated agents operating within larger complex systems. The question Complexity Economics characteristically asks is: how do macro-level patterns emerge from the interactions of heterogeneous agents? This dissertation pursues a different question: what must we assume about the nature of firms, the limits of knowledge about them, and the logic of inquiry into them if we take seriously that firms are themselves complex adaptive systems?

This difference entails three specific extensions.

First, the dissertation shifts the object of analysis from markets to firms. Where Complexity Economics theorises economies in which firms function as agents, this dissertation theorises firms as complex adaptive systems in their own right—asking what ontological commitments account for firm identity, boundaries, and persistence; what epistemological stance is appropriate for entities whose properties emerge from relational configurations; what methodological principles guide inquiry into systems that generate genuine novelty.

Second, the dissertation foregrounds the meso-level that Complexity Economics's macro-focus has not addressed. The intermediate organisational scale—where routines, capabilities, and coordination architectures constitute the firm's distinctive properties—operates between individual agents and aggregate market outcomes. Complexity Economics's models trace the micro-to-macro pathway; the meso-level is traversed rather than theorised. This dissertation attends to precisely that intermediate scale.

Third, the dissertation engages with theories that Complexity Economics has not systematically addressed: transaction cost economics, agency theory, the resource-based view, and dynamic capabilities. These firm-level literatures retain mechanistic foundations that Complexity Economics's macro-level reconstruction has left untransformed.

1.5.3 Forging a Hard Core

These extensions describe what the dissertation addresses. A deeper contribution concerns how it addresses them—specifically, the construction of philosophical architecture capable of constituting a genuine paradigmatic alternative.

Complexity Economics has two limitations when presenting such an alternative.

First, its commitments—although valid—have not been systematically consolidated through foundational philosophy. The mainstream tradition co-evolved with centuries of philosophical development: Cartesian mechanism, Newtonian determinism, Humean causation, and logical positivism all shaped the assumptions that became economic orthodoxy. This philosophical scaffolding gave the mainstream its resilience—its assumptions draw on systematic elaboration about reality, knowledge, and explanation. Complexity Economics has engaged with critical realism, process philosophy, and related traditions, but this engagement remains unevenly integrated. Its assumptions emerge primarily from modelling practice rather than from systematic philosophical reconstruction.

This dissertation addresses this limitation directly. Chapter 4 exposes why mechanistic ontological commitments fail when applied to complex adaptive systems; Chapter 5 constructs the alternative ontology and derives its epistemological and methodological consequences. Together, they consolidate the philosophical scaffolding that Complexity Economics has not yet systematically assembled.

Second, Complexity Economics has not forged its assumptions into a mutually reinforcing system capable of constituting a hard core in the Lakatosian sense. The assumptions it embraces—emergence, non-equilibrium, heterogeneity, fundamental uncertainty, path-dependence—represent genuine reconstructions, not mere relaxations. But they appear as a collection of departures from orthodoxy rather than a configuration whose elements entail and support one another.

This dissertation reconstructs these commitments as a mutually reinforcing system. The mechanistic paradigm's resilience derives not from individual assumptions but from their mutual reinforcement—a structure that absorbs challenges without fundamental change. A genuine alternative must exhibit the same property: internal coherence sufficient to organise inquiry on different foundations.

This represents a systematic attempt that, to the best of our knowledge, has not yet been carried out in this form: forging CAS-compatible assumptions into an interlocking philosophical architecture—grounded in foundational ontology, epistemology, and methodology—capable of constituting a genuine alternative to mechanistic firm theory.

1.5.4 Research Questions

Against this background, the dissertation is guided by the following primary research question:

How should firms be reconceptualised when the foundational assumptions of mainstream economic theory—identified through systematic problematisation—are transformed to be compatible with complex adaptive systems?

This question crystallised through three convergent recognitions: that mainstream theory exhibits recurrent inadequacy for observed firm behaviour in complex environments (problem recognition); that this inadequacy reflects paradigm-level assumptions rather than superficial gaps (diagnostic recognition); and that reconceptualisation requires comprehensive assumption transformation, not piecemeal modification (solution recognition).

Subsidiary Questions

The primary question decomposes into five sub-questions:

SQ1 (Identification): What are the foundational assumptions underlying mainstream economic understandings of firms?

SQ2 (Diagnosis): Why are these assumptions problematic for understanding and managing firms in complex environments?

SQ3 (Ontological Transformation): What alternative ontological assumptions are required to reconceptualise firms as complex adaptive systems?

SQ4 (Epistemological Transformation): What epistemological stance follows from this reconstructed ontology?

SQ5 (Methodological Transformation): What methodology is appropriate for CAS-consistent inquiry, given the ontological and epistemological commitments?

These questions constitute a dependency network rather than a linear sequence: SQ1 enables SQ2; SQ2 constrains admissible answers to SQ3–SQ5; the transformation questions are directional but also iterative, as epistemic limits illuminate ontological commitments and methodological requirements expose implicit epistemologies.

The questions correspond to chapters as follows: SQ1 and SQ2 are answered by Chapter 3's diagnostic architecture. SQ3 is answered across Chapters 4 and 5: Chapter 4 establishes why mechanistic ontology fails; Chapter 5 constructs the CAS-compatible replacement. SQ4 and SQ5 are answered by Chapter 5's derivation of epistemological and methodological implications.

The integration of these transformations into a coherent alternative is not a separate analytical operation but a structural property of the reconstruction itself: the chapters are organised so

that ontological commitments constrain epistemological possibilities, which constrain methodological options, creating the mutual entailment that paradigmatic coherence requires.

1.5.5 Nature of the Contribution

The dissertation's contribution is meta-theoretical. It does not advance a substantive theory of the firm—a model that explains or predicts specific firm behaviours. It reconstructs the foundational assumptions that determine what firms are taken to be, what can be known about them, and how they can legitimately be analysed. The output is an architecture: a CAS-consistent configuration of ontological, epistemological, and methodological commitments that defines what kind of entity the firm must be assumed to be, what constitutes an eligible answer to firm-theoretical questions, and what forms of inquiry remain coherent once those commitments are accepted. The CAS framework does not replace transaction cost economics, agency theory, or the resource-based view; it provides the foundational structure within which their insights can be preserved while their mechanistic commitments are transformed—and within which their scope conditions, relationships, and limits become explicit. The hard core forged in this dissertation provides the meta-level constraints within which theoretical diversity becomes productive rather than fragmentary.

Within this frame, the dissertation makes three interrelated contributions, each operating at the meta-theoretical level.

First, it identifies the assumptive architecture of mainstream firm theory and diagnoses why it generates persistent difficulties. Through systematic problematisation, the dissertation shows that recurring problems—instability, crisis blindness, weak treatment of innovation, marginalisation of organisational phenomena—originate not in insufficient data or incomplete formalisation but in an interlocking set of foundational assumptions inherited from a mechanistic paradigm. This diagnostic contribution establishes why piecemeal modification fails: three reinforcing feedback loops produce a triple-locked architecture in which challenging any single assumption while the others remain intact results in the preserved hard core regenerating the challenged derivative or absorbing the modification as a special case.

Second, it reconstructs the ontological assumptions required for firms to be coherently understood as complex adaptive systems. This reconstruction proceeds through coordinated FROM→TO transformations: emergent wholeness replaces reductionism; process ontology replaces substance; relational primacy replaces atomism; thermodynamic openness replaces closure; the meso-level is restored as the privileged analytical scale. The result specifies what firms must be assumed to be: historically evolving, open, multi-level process-patterns—relationally constituted configurations whose identity consists in self-reproducing patterns of activity rather than fixed essence, whose boundaries function as selective membranes, whose capabilities emerge from relational configurations, and whose evolution follows path-dependent trajectories.

Third, it recalibrates the epistemic standards that remain coherent once firms are conceptualised as complex adaptive systems. It demonstrates why prediction-centred, optimisation-based rigour becomes unreliable under conditions of emergence, non-ergodicity, and fundamental uncertainty, and articulates alternative standards—pattern recognition, navigation, historically

grounded explanation—that remain disciplined without presupposing predictability. This is not a retreat from rigour but its reconstitution for objects that generate genuine novelty.

These contributions are not independent outputs but components of a single reconstructive cascade. Just as the mechanistic assumptions mutually require and stabilise one another, the CAS-compatible assumptions form a mutually reinforcing architecture: emergence requires relational constitution; process ontology necessitates thermodynamic openness; fundamental uncertainty follows from emergence and non-ergodicity; structured pluralism follows from irreducible multi-scale dynamics. The contribution is thus to forge a Lakatosian hard core as genuine paradigmatic alternative—not a collection of relaxed constraints, but a self-reinforcing configuration of commitments whose mutual implications generate a research programme with its own heuristics and its own criteria of progress.

1.5.6 Boundaries and Non-Claims

The dissertation operates under explicit methodological constraints that are design choices appropriate to paradigm-level work.

First, it is pure theory without empirical testing. The dissertation establishes the conceptual conditions under which empirical research can proceed coherently—conditions that must be in place before meaningful empirical testing becomes possible. Empirical inquiry conducted within inherited frameworks would be circular, since the foundational assumptions the dissertation interrogates are precisely the commitments that determine what counts as evidence in the first place. Empirical work is enabled by the reconstruction but cannot precede it without reintroducing mechanistic assumptions by methodological default.

Second, it does not propose a new theory of the firm in the conventional sense—it does not advance a model that predicts specific firm behaviours or competes directly with transaction cost economics, agency theory, or the resource-based view on their own terrain. Rather, it operates at a prior level: it defines what kind of entity the firm must be assumed to be, specifies what constitutes an eligible answer to firm-theoretical questions, and reconstructs the ontological, epistemological, and methodological conditions under which such answers can be formulated, compared, and evaluated. The framework is architectural rather than totalising. It does not seek to unify existing theories into a single super-theory—such unification would contradict both the complexity of the object and the epistemological stance the dissertation develops—but to provide the foundational structure within which their insights can be situated, their scope conditions made explicit, and their relationships clarified.

Third, its scope is confined to the firm-as-CAS at the meso-level. The reconceptualisation applies to organisational forms that exhibit the defining properties of complex adaptive systems—non-decomposability, adaptive behaviour, path-dependence, relational constitution, and operation under fundamental uncertainty. Where interaction effects, feedback, and organisational learning are negligible, or where subsystems operate within stable, well-defined solution spaces, mechanistic or optimisation-based models may retain local and conditional validity. The claim is that such cases are special rather than representative once firms are understood as open, evolving systems.

Fourth, predictive ambitions remain necessarily modest under non-ergodicity and emergence. This is not a retreat from rigour but its reconstitution for objects that generate genuine novelty. Pattern recognition, navigation, and historically grounded explanation replace point prediction as the appropriate epistemic ideals.

Fifth, it does not present itself as revolutionary in the sense of rejecting prior work. Existing theories are treated as historically situated achievements whose limits become visible under new analytical conditions. The contribution preserves the insights of transaction cost economics, the resource-based view, and evolutionary approaches while making explicit how they relate, where they apply, and where their limits lie.

These constraints define the conditions under which the dissertation's claims are valid and open space for subsequent research programmes: empirical investigation of meso-level phenomena, operationalisation of structured pluralism, and case-based research examining firms as process-patterns—programmes enabled by the reconstructive cascade.

1.6 Structure of the Dissertation

Having established what the dissertation attempts and what it does not claim, this section describes how the argument is organised to achieve these aims.

The dissertation's structure mirrors the three stages of problematisation articulated in the abstract: diagnostic identification of the assumptive architecture generating persistent theoretical failures; reconstructive transformation of those assumptions across ontological, epistemological, and methodological domains; and evaluative assessment of whether the reconstructed framework generates more adequate theory. The dissertation unfolds in four connected movements corresponding to these stages:

1. **Orientation and research design** (Chapter 2)
2. **Diagnosis and critique** (Chapter 3)
3. **Reconstruction** (Chapters 4-5)
4. **Implications, and evaluation** (Chapter 6)

Chapter 2 — Methodological and Philosophical Foundations

Chapter 2 establishes the methodological architecture of the dissertation. It positions the inquiry within the tradition of pure economic theory—examining assumptions that structure inquiry rather than testing hypotheses within an assumed framework. The chapter articulates the philosophical foundations (critical realism, hermeneutic pre-understanding, perspectivalism) and develops problematisation as the governing methodology: targeting foundational assumptions directly rather than gap-spotting within accepted frameworks. This methodological commitment governs every subsequent chapter: each implements problematisation by rendering visible assumptions that normally operate as unquestioned background.

The chapter introduces the FROM→TO transformation framework as the operational device for paradigm-level reconstruction and derives the research questions as methodological products of this process. It articulates success criteria appropriate to paradigmatic work—internal coherence, problem-solving capacity, theoretical fruitfulness—and explains why criteria designed for empirical research (predictive accuracy, falsifiability) cannot be directly applied to foundational reconstruction, since such work establishes the conditions under which empirical testing becomes meaningful.

Chapter 3 — The Diagnosis: Anatomy of the Mechanistic Paradigm

Chapter 3 implements the diagnostic stage of problematisation. It reconstructs the assumptive architecture of mainstream firm theory as a network of fifteen interlocking assumptions organised across a Lakatosian hard core (five ontological, four epistemological) and protective belt (six derivative assumptions). The chapter explains why this architecture persists despite anomalies through three reinforcing feedback loops that produce a triple-locked system, traces how this lock-in manifests as methodological symptoms (physics envy, optimisation fetish, DSGE institutionalisation), and documents the consequences: predictive brittleness, managerial illusion of control, and crisis blindness—exemplified by the 2008 financial crisis, where equilibrium, linearity, and risk-model assumptions failed simultaneously. The chapter answers SQ1 and SQ2: what foundational assumptions underlie mainstream firm theory, and why are they problematic?

Chapter 3 remains diagnostic: it establishes why reconstruction is necessary without proposing the alternative.

Chapters 4–5 — The Reconstruction Cascade (Ontology → Epistemology → Methodology)

Chapters 4 and 5 implement the reconstructive stage of problematisation. Together, they execute the dissertation's central analytical operation: a cascade in which ontological reconstruction constrains epistemological possibilities, which in turn constrain methodological options. This ordering is not expository convenience but the thesis's governing logic—what firms are determines what can be known about them, and both jointly determine what methods can legitimately be used to study them. The two chapters answer SQ3, SQ4, and SQ5 respectively: what alternative ontological, epistemological, and methodological assumptions are required?

Chapter 4 begins at the only level where change can be decisive: first principles. It asks the prior question that existing theories typically leave implicit: what must the firm be assumed to be if it is to be understood as a complex adaptive system? The chapter problematises the trinity of mechanistic ontology—ontological reductionism, substance ontology, and atomism—showing that these form a mutually reinforcing system that renders the firm's defining properties analytically invisible. Under these assumptions, the meso-level at which routines, capabilities, and coordination architectures operate is eliminated by construction—the "Missing Middle" that Chapter 3 identified as economics' most consequential erasure. Reductionism decomposes wholes into additive parts, substance ontology freezes entities into stable essences, and atomism treats elements as independent and self-sufficient; together, this

triad does not merely simplify the firm but makes its defining properties—emergence, coordination, adaptation—analytically irrecoverable.

Chapter 4 establishes that reconstruction is logically necessary given the failures diagnosed. It answers SQ3 in its negative aspect: demonstrating what the firm cannot coherently be assumed to be. The positive reconstruction—specifying what the firm must be assumed to be instead—is the task of Chapter 5.

Chapter 5 performs the reconstructive work that Chapters 3 and 4 prepared, executing the full reconstruction cascade across all three domains: ontological, epistemological, and methodological. The cascade logic is explicit: the chapter first reconstructs what firms are (ontology), then derives what can be known about them (epistemology), then specifies how they must be studied (methodology).

The ontological reconstruction (answering SQ3) executes a coordinated set of transformations: emergent wholeness replaces reductionism; process ontology replaces substance; relational primacy replaces atomism; thermodynamic openness replaces closure; the meso-level is restored as the privileged analytical scale. The firm is re-specified as a historically evolving, open, multi-level process-pattern: a metabolic processor whose identity consists in self-reproducing patterns of coordination rather than fixed essence, and whose causal powers cannot be reduced to component-level descriptions.

From this reconstructed ontology, the epistemological reconstruction (answering SQ4) follows necessarily. The chapter problematises four pillars of mechanistic epistemology—Cartesian certainty, linear deterministic causality, ergodicity, and the conflation of risk with uncertainty—and demonstrates that the prevailing prediction-centred epistemology fails not because of weak forecasting performance, but because its enabling presuppositions do not hold under a CAS ontology. In their place, the chapter articulates CAS-consistent epistemic commitments: fundamental uncertainty recognised as ontological rather than temporary; pattern-based understanding replacing point prediction; distributed and scaffolded cognition replacing omniscient calculation; navigation replacing optimisation as the coherent logic of decision. The firm is reconceptualised from a calculator optimising over stable solution spaces to a navigation system steering through uncertain, co-evolving terrain.

From this epistemological recalibration, Chapter 5 derives the methodological consequences (answering SQ5). If firms are ontologically emergent and epistemically characterised by fundamental uncertainty, then methodological monism becomes incoherent. The chapter develops structured pluralism—not as a normative preference but as a necessity imposed by the object of inquiry. Ashby's Law of Requisite Variety formalises the principle: explanatory variety must match phenomenal variety. Because emergence generates irreducible properties at each level of analysis, no single model can span micro behaviour, meso interaction patterns, and macro outcomes simultaneously. The chapter establishes disciplining mechanisms—admissibility axioms, positioning criteria, transformation protocols, and context-specific validation standards—that provide rigour without reimposing monism.

The chapter concludes by tracing dynamic implications: endogenous novelty generation, path-dependence, non-ergodicity, downward causation, and adaptive tension—properties that

mechanistic assumptions render invisible but that become analytically accessible once the reconstruction cascade is accepted.

Chapter 6 — Conclusions, Contributions, Implications, Limitations, and Future Research

Chapter 6 implements the evaluative stage of problematisation. It assesses the dissertation against criteria appropriate to paradigmatic work: internal coherence, problem-solving capacity, theoretical fruitfulness, illumination of mechanisms, and comprehensiveness. It specifies what the thesis does not claim, then traces implications across five domains: economic theory of the firm, methodology and standards of rigour, management and organisational analysis, economic policy, and philosophy of economics.

The chapter articulates six conditions of validity delimiting where the framework applies, and identifies future research directions: empirical investigation of meso-level phenomena, operationalisation of structured pluralism, and case-based research examining firms as process-patterns.

Reading Guidance

The dissertation can be read according to evaluative focus:

- **For holistic understanding:** Chapters 1–3 provide orientation and diagnosis; Chapters 4–5 develop critique and reconstruction; Chapter 6 consolidates implications. This sequence displays the complete argumentative arc.
- **For methodological focus:** Begin with Chapter 2 to understand why problematisation rather than gap-spotting is required and how the FROM→TO framework operates, then proceed to Chapters 3–5.
- **For theory-of-the-firm focus:** Begin with Chapter 3's diagnosis of mechanistic limitations, then read Chapters 4–5 for critique and reconstruction, consulting Chapter 2 as needed.

Because the chapters function as components of an integrated argument rather than self-contained contributions, the dissertation's adequacy depends on whether the reconstruction is internally coherent, responsive to the problems it addresses, and generative of productive inquiry.

Chapter 2 – Methodological and Philosophical Foundations

Introduction

This chapter establishes the philosophical, epistemological, and methodological foundations that structure the inquiry undertaken in the dissertation. Its purpose is to make explicit the assumptions, analytical commitments, and evaluative criteria that govern how economic theories are examined, compared, and assessed throughout the study. Rather than advancing substantive claims about specific economic phenomena, this chapter clarifies the conceptual and methodological conditions under which such claims can be meaningfully formulated and evaluated.

Economic theories are embedded within broader frameworks of thought that shape what is considered observable, explainable, and methodologically legitimate. These frameworks – often inherited and left implicit – operate at a level prior to formal modelling or empirical testing. As a result, many persistent theoretical tensions and explanatory limitations in economics cannot be addressed through technical refinement alone, but require examination of the deeper assumptions that organise economic reasoning. Chapter 2 addresses this foundational layer directly.

The chapter therefore serves a *methodological* rather than a *substantive* role. It does not seek to adjudicate between competing economic theories, nor to propose an alternative theory of the firm or the economy. Instead, it develops a systematic approach for examining how different economic theories are constructed, what assumptions they rely upon, and how these assumptions shape their explanatory scope. This clarification is a necessary precondition for comparative theoretical evaluation.

Methodologically, the chapter positions the dissertation within the tradition of pure theoretical inquiry. The focus is on assumption analysis, first-principles reasoning, and problematisation rather than hypothesis testing or empirical validation. This positioning reflects the nature of the questions under examination. The assumptions addressed in this chapter – such as reductionism, equilibrium, determinism, closure, and aggregation – are not empirical propositions that can be directly tested. They are pre-empirical commitments that structure how theories are formulated and how evidence is interpreted. As such, they require philosophical and methodological analysis rather than empirical verification.

A central concern of the chapter is the treatment of analytical levels. Conventional economic methodology has tended to privilege either individual-level analysis or aggregate outcomes, often treating intermediate structures as analytically derivative or eliminable. This chapter introduces and justifies sustained attention to the intermediate (meso) level – comprising organisations, routines, networks, and institutional arrangements – as an analytically distinct domain. References to firm-level or organisational phenomena in this chapter are illustrative and classificatory only. They serve to clarify ontological and methodological distinctions, not to advance applied analysis or theoretical reconstruction.

The chapter also establishes a set of epistemic distinctions that govern the scope and claims of the dissertation. These include distinctions between explanation and evaluation, between intellectual and ontological emergence, and between methodological demonstration and empirical validation. Making these distinctions explicit is essential for maintaining analytical discipline and avoiding overextension of claims. The chapter thus defines not only what the methodology enables, but also what it does not attempt to establish.

Structurally, Chapter 2 performs three interrelated functions. First, it articulates the philosophical foundations relevant to economic inquiry, drawing on systems thinking, complexity science, and the philosophy of science to clarify alternative ontological commitments. Second, it develops a methodological framework – centred on problematisation and first-principles analysis – for examining economic theories as structured systems of assumptions rather than isolated models. Third, it specifies the evaluative criteria by which different economic theories can be systematically examined in subsequent chapters.

The framework developed here is applied in the following chapter to conduct a comparative evaluation of existing economic theories – both orthodox and heterodox – through the lens of Complex Adaptive Systems. That evaluation relies on the methodological apparatus established in this chapter and does not presuppose the superiority of any particular theoretical tradition. Chapter 2 therefore functions as the analytical spine of the dissertation, ensuring that subsequent evaluations are grounded in explicit methodological reasoning rather than implicit or unexamined presuppositions.

In this sense, Chapter 2 should be read as an enabling chapter. It does not seek to persuade through substantive conclusions, but to establish the conceptual and methodological conditions under which the remainder of the dissertation can be assessed on coherent and transparent terms. By making explicit what is typically left implicit, the chapter provides the foundation for systematic theoretical evaluation without pre-empting its outcomes.

2.1 Positioning the Dissertation

Chapter 2 establishes the philosophical, epistemological, and methodological foundations that structure the dissertation's inquiry. Before those foundations can be articulated, however, a prior question must be addressed: what kind of work is this dissertation, and by what standards should it be evaluated?

Section 2.1 answers this question by positioning the dissertation within the tradition of pure economic theory, clarifying the level at which it intervenes, and establishing the classificatory groundwork upon which subsequent methodological reasoning depends. Failure to clarify the genre of inquiry produces category errors in both critique and assessment. The problems addressed by this dissertation are paradigm-constitutive assumptions that cannot be revised incrementally or tested in isolation; appropriate evaluation criteria differ categorically from those applied to empirical research.

The argument proceeds through five subsections. Subsection 2.1.1 positions the dissertation within the pure theory tradition and establishes that this positioning is necessary, not merely permissible. Subsection 2.1.2 distinguishes research programmes, paradigms, and meta-theories, locating the dissertation's intervention at the paradigm level. Subsection 2.1.3 defines what constitutes theory and distinguishes genuine theoretical contribution from common substitutes. Subsection 2.1.4 demonstrates that paradigm-level theoretical work has distinguished precedents in Kuhn, Lakatos, Lawson, and Knight. Subsection 2.1.5 establishes the complicated/complex distinction, showing that firms exhibit properties rendering decomposition-based approaches categorically inappropriate.

Section 2.2 builds on this positioning by articulating the philosophical foundations appropriate to paradigm-level inquiry into complex phenomena.

2.1.1 The Categorical Distinction: Pure Theory vs Applied Research

This dissertation is positioned within the tradition of pure economic theory concerned with paradigmatic analysis, rather than applied or empirically oriented research. Before specifying particular methodological choices, it is therefore necessary to clarify the genre of inquiry undertaken and the evaluative criteria appropriate to it. Different forms of inquiry generate different kinds of knowledge and require different standards of justification; failure to distinguish among them produces category errors in both critique and assessment.

Economic methodology has long recognised a fundamental distinction between pure theory and applied or practical research. Pure theory establishes the conceptual apparatus through which empirical phenomena become intelligible, while applied research examines how this apparatus operates in particular contexts. As Bylund (2021) emphasises, pure theory constitutes the starting point of economic science, preceding and informing inductive and applied inquiry. This distinction is firmly rooted in the history of economic thought: Knight ([1921] 1985) described pure theory as “the first step toward a practical comprehension of the social system,” while Allin (1949) warned that conflating pure and applied theory produces systematic error – pure theory is judged by inappropriate empirical standards, while applied work proceeds without adequate conceptual foundations.

The epistemic status of pure theory follows directly from this role. Foundational theoretical assumptions occupy a different epistemic position from empirical hypotheses and therefore require different forms of evaluation. As Mill ([1843], in Machlup, 1955) observed, empirical verification pertains to the application of science, not to the establishment of its foundational concepts. Machlup (1955) further clarified that high-level generalisations are assessed not through independent empirical testing, but through their internal coherence, explanatory capacity, and ability to organise understanding.

Musgrave's (1981) distinction among negligibility, domain, and heuristic assumptions further illuminates this point. The foundational assumptions developed in pure theory function primarily as heuristic assumptions: simplifying constructs that enable systematic theorising rather than propositions awaiting direct empirical adjudication. Their validity lies in their capacity to render complex phenomena intelligible and to support coherent lines of inquiry.

From this perspective, theory does not compete with observation but enables it. As Mises (1949) argued, theoretical concepts function as mental tools that render reality intelligible; empirical experience alone cannot adjudicate between competing theoretical frameworks, since identical observations may be interpreted through fundamentally different conceptual lenses. Theoretical work must therefore precede empirical accumulation rather than emerge from it.

The necessity of pure theory is reinforced by disciplinary dynamics within economics. Akerlof's (2020) analysis of hardness bias demonstrates how methodological preferences for formal precision systematically disadvantage paradigm-level theorising, encouraging incremental refinement of existing frameworks while discouraging critical examination of foundational assumptions. Because evaluative standards are themselves embedded within paradigms, paradigm-level analysis cannot proceed through empirical accumulation alone.

This requirement becomes particularly salient when the phenomena under investigation exhibit essential complexity. As Hayek (1974) argued, complex social systems possess properties that emerge from interactions among multiple elements and therefore resist isolation, decomposition, and precise prediction. Theory may identify structural patterns and generative tendencies, but it cannot yield point predictions of specific outcomes; this limitation reflects the ontology of the subject matter rather than technical insufficiency.

Cilliers (2005) extended this insight by demonstrating that both system boundaries and system properties are theory-dependent. What constitutes "the firm" is not an objective datum awaiting discovery but a theoretical construction that determines what counts as relevant evidence. Empirical research necessarily presupposes such constructions and cannot establish them independently. Arrow (1994) and Kaufmann (1942) similarly emphasised that social and intentional phenomena cannot be reduced to categories appropriate to physical systems without systematic distortion.

Historically, the most consequential advances in economic thought have followed this path. Knight's distinction between risk and uncertainty, Mises's analysis of purposeful action, Hayek's theory of knowledge and spontaneous order, and later paradigm-level contributions by Kuhn, Lakatos, and Lawson all proceeded through conceptual reconceptualisation rather than empirical accumulation. None would satisfy criteria designed for applied research, yet each transformed the discipline's understanding of its subject matter.

Two objections might be anticipated. First, one might ask whether paradigm-level theoretical work is genuinely economics rather than philosophy. The answer lies in the subject matter: the assumptions examined concern firms, production, coordination, and economic organisation – topics central to economics regardless of the method by which they are addressed. Second, one might question whether pure theory without empirical testing can generate knowledge. The precedents examined in Section 2.1.4 demonstrate that it can: Knight, Hayek, Kuhn, and Lakatos all advanced understanding through conceptual work that preceded and enabled later empirical investigation. The issue is not whether empirical work matters – it does – but whether foundational reconstruction must await empirical confirmation. For the reasons outlined above, it cannot.

For these reasons, pure theory is not merely a permissible methodological choice for this dissertation but a necessary one. The problems addressed are not empirical anomalies but paradigm-constitutive assumptions embedded in the mechanistic framework governing the theory of the firm. Because these assumptions are mutually reinforcing and define the conditions of intelligibility for empirical work, they cannot be revised incrementally or tested in isolation.

Accordingly, this dissertation is explicitly positioned as pure theoretical work. Its contribution lies in the systematic examination of foundational assumptions and in the clarification of the ontological, epistemological, and methodological conditions under which coherent theories of the firm may be formulated without internal contradiction or category error. Consistent with Little's (1993) account of explanation in the social sciences, the aim is not law-like prediction but improved understanding of the causal mechanisms through which organisational phenomena emerge.

Finally, the dissertation intervenes at the meta-theoretical and methodological level, rather than advancing a single closed-form theory of the firm. It reconstructs the conceptual conditions under which alternative theoretical formulations can be assessed for coherence and adequacy, allowing theoretical conclusions to emerge as implications of analysis rather than as prescriptive solutions.

2.1.2 Clarifying Conceptual Foundations: Research Programmes, Paradigms, and Meta-Theories

Section 2.1.1 positioned this dissertation within the tradition of pure theory, where the central contribution lies in foundational analysis rather than empirical investigation. Pure theory, however, can operate at different levels of abstraction. Precision regarding the level at which this dissertation intervenes is therefore necessary both to render the contribution legible and to clarify the standards by which it should be evaluated.

Three concepts from the philosophy of science – research programmes, paradigms, and meta-theories – are frequently invoked in discussions of theoretical change, yet often conflated. This conflation is consequential, because the concepts refer to different analytical levels and imply different mechanisms of persistence, critique, and transformation. Clarifying their distinction is therefore a prerequisite for specifying the nature of the present intervention.

Research Programmes: Lakatos and Theoretical Persistence

Imre Lakatos developed the concept of a research programme to overcome limitations in both naïve falsificationism and Kuhnian accounts of scientific change. For Lakatos, science is not rationally reconstructed as isolated theory–test episodes; rather, the appropriate unit of appraisal is a sequence of theories organised around a shared core (Lakatos, 1970).

A research programme consists of a hard core – fundamental assumptions treated as methodologically irrefutable – and a protective belt of auxiliary hypotheses that can be modified in response to anomalies. Methodological rules specify which paths of research to avoid (negative heuristic) and which to pursue (positive heuristic). Anomalies therefore do not lead to immediate abandonment of a programme; they are typically absorbed through protective-belt adjustments, a structural feature rather than a defect of programme-based science.

Lakatos distinguished progressive from degenerating problemshifts. A programme is progressive when modifications generate novel theoretical content and new explanatory reach; it degenerates when adjustments become ad hoc, accommodating anomalies without independent gain (Lakatos, 1970, 1974). Importantly, falsification is not a simple relation between theory and evidence but depends on the availability of superior alternatives.

For the purposes of this dissertation, the research programme concept clarifies the resilience of mainstream economics. The neoclassical programme has repeatedly protected its hard core – optimisation, equilibrium tendencies, and methodological individualism – through successive protective-belt adaptations. This persistence reflects the internal logic of programme-based science rather than mere empirical obstinacy.

Paradigms: Kuhn and Disciplinary Matrices

Thomas Kuhn's concept of a paradigm operates at a different analytical level. Kuhn was concerned not with the logical structure of theory sequences but with the cognitive, normative, and sociological organisation of scientific communities. A paradigm is a disciplinary matrix: a constellation of shared symbolic generalisations, metaphysical commitments, values, techniques, and exemplars that define normal scientific practice (Kuhn, 1962/1970).

Paradigms shape not only what scientists believe but how they see. Kuhn emphasised that paradigm change has a gestalt-like character: competing paradigms are not merely alternative propositions but different ways of organising experience, often resistant to point-by-point translation. Changes in paradigm therefore involve shifts in evaluative standards as well as in substantive commitments.

Lakatos explicitly engaged Kuhn's analysis, proposing the research programme framework as a rational reconstruction of Kuhn's account (Lakatos, 1974). Taken together, the two perspectives clarify why foundational disputes cannot be resolved through local adjustment of auxiliary hypotheses alone. Within this combined framework, the limitations of mechanistic economics appear not as isolated anomalies but as paradigm-constitutive blind spots, particularly concerning emergence, uncertainty, and complexity.

Meta-Theories: Assumptions About Theorising Itself

Meta-theories operate at a still higher level of abstraction. They consist of assumptions about what counts as legitimate theorising: what kinds of entities exist, what constitutes explanation, how theory should relate to evidence, and what standards distinguish strong from weak theory (Ritzer, 1991). These assumptions typically remain implicit, functioning as background conditions that constrain the space of theoretical possibility.

Alvesson and Sandberg's (2011) typology clarifies how such commitments manifest in scholarly work, distinguishing in-house assumptions, root metaphors, paradigmatic ontological–epistemological–methodological assumptions, ideology assumptions, and field assumptions. Meta-theoretical commitments operate most powerfully at the paradigm and field levels, shaping what is taken to be “proper” inquiry before specific theories are even formulated.

Relationships Among the Concepts

Although distinct, research programmes, paradigms, and meta-theories are not independent. Paradigms encompass multiple research programmes, establishing shared commitments within which several programmes may coexist. Meta-theories, in turn, shape which paradigms appear scientifically respectable and which are marginalised.

All three involve assumptions, but at different depths. Research programme hard cores carry substantive theoretical commitments; paradigms carry broader commitments concerning methods, values, and exemplars; meta-theories concern the nature of theorising itself. Change at higher levels typically entails change at lower ones: paradigm reorientation affects programme hard cores, while meta-theoretical change reshapes paradigmatic standards.

What This Dissertation Targets

The purpose of these distinctions is to clarify the level at which the present dissertation intervenes. The dissertation operates primarily at the paradigm level. The limitations examined across subsequent chapters are not problems internal to a mechanistic framework but problems of that framework – constitutive assumptions that define what counts as legitimate analysis of firms and economic organisation.

The analysis is organised across three interrelated dimensions: (1) **Ontological** assumptions about what firms and economic systems are; (2) **Epistemological** assumptions about what can be known in open, complex systems and how knowledge claims should be warranted; and (3) **Methodological** assumptions concerning appropriate forms of analysis, including commitments to monism and exclusive formalism.

In Alvesson and Sandberg's (2011) terms, these are paradigmatic assumptions. Examining them entails what Kuhn described as a gestalt shift – not incremental adjustment within an existing framework, but reconsideration of how the object of inquiry is conceptualised and how knowledge of it is pursued. At the same time, paradigm-level analysis has implications at both

the meta-theoretical level (assumptions about economic theorising) and the research programme level (the status of hard-core commitments).

Why Paradigm-Level Analysis Is Necessary

Three considerations justify intervention at the paradigm level.

First, the limitations identified recur across multiple research programmes operating within the same mechanistic paradigm, even where those programmes diverge in detail. This recurrence indicates that inadequacy lies in shared foundational commitments rather than in programme-specific omissions.

Second, the challenged assumptions function as boundary-defining conditions of respectable economics. They delimit what is admitted as legitimate analysis independently of empirical plausibility, a hallmark of paradigmatic status.

Third, the issues at stake concern system-type classification rather than marginal extension. As Ackoff and Gharajedaghi (2003) argue, firms are purposeful social systems, and treating them as deterministic or merely complicated systems produces categorical mismatch – a mismatch that cannot be resolved through auxiliary adjustment alone.

Conclusion

This section distinguished three analytical levels central to theoretical change: Lakatosian research programmes, Kuhnian paradigms, and meta-theories concerning the nature of theorising itself. These concepts operate at different depths and imply different standards of evaluation.

The dissertation is positioned primarily at the paradigm level. The issues it examines are not anomalies within a mechanistic framework but questions concerning its constitutive assumptions. Accordingly, its contribution should be assessed not as a protective-belt modification, but as a systematic examination of paradigmatic commitments and their implications for understanding firms and economic organisation. Paradigm-level contributions are assessed by coherence, adequacy to the phenomena, and their capacity to illuminate underlying causal mechanisms rather than by criteria appropriate to hypothesis testing.

2.1.3 What Theory IS and IS NOT

Introduction

The preceding sections established that this dissertation operates within the pure theory tradition at the paradigm level. A further clarification is nevertheless required: what constitutes theory itself. This question is critical for two reasons. First, it defines the nature of the contribution advanced by the dissertation. Second, it provides criteria for distinguishing genuine theoretical work from scholarly outputs that merely resemble theory without fulfilling its explanatory function.

This section addresses the question in three parts. Part A examines influential definitions of theory, identifying common elements relevant to paradigm-level work. Part B draws on the critical literature to clarify what does not constitute theory. Part C establishes evaluative criteria for strong theoretical contribution.

Part A: What Theory IS

Definitions and Core Components

Despite disciplinary variation, influential definitions of theory converge on several core elements. Bacharach (1989, p. 496) defines theory as “a statement of relations among concepts within a set of boundary assumptions,” emphasising concepts, relations, and the conditions under which those relations hold. Wacker (1998) elaborates this structure into four components – definitions, domain, relationships, and predictions – while acknowledging that different forms of theoretical work may emphasise these elements unevenly. For pure theory, conceptual clarity and relational structure may precede predictive specification.

Fuhse (2022), drawing on Giere (2006), advances a perspectival conception: theories are “webs of concepts” that guide observation and interpretation rather than mirror reality directly. On this view, truth claims are not arbitrary but perspective-dependent, reflecting the theory-laden character of observation. Concepts do not merely describe phenomena; they participate in constructing what counts as the phenomenon.

This perspectival character extends to system boundaries. As Cilliers (2005) argues, the boundary of a complex system is not an objective datum but a theoretical decision that determines what is treated as internal or external. The decision to analyse “the firm” already presupposes a theoretical framework. Different theories therefore construct different objects of inquiry, rendering theoretical pluralism a necessary feature of research on complex phenomena rather than a defect to be eliminated.

Functions of Theory

Theory performs several interconnected functions.

Explanation is primary. Theory explains by identifying mechanisms, processes, or relational structures that generate observed patterns. As Margenau observed, explanation involves moving “beyond phenomena” into the constructional domain (quoted in Machlup, 1955). Explanatory assumptions therefore cannot be empirically verified in isolation; their value lies in their capacity to render phenomena intelligible.

Organisation is equally central. Theory organises disparate observations into coherent patterns. Knight ([1921] 1985) emphasised that pure theory isolates and defines essential characteristics of systems, enabling understanding of their tendencies. Without theoretical organisation, empirical findings remain disconnected and open to arbitrary interpretation (Mises, 1949).

Prediction plays a qualified role. For simple systems, point prediction may be possible. For complex phenomena, theory can predict only patterns rather than specific outcomes. As Hayek (1967) demonstrated, this limitation reflects the ontology of complex systems, not theoretical inadequacy. The pursuit of predictive simplicity in complex domains produces distortion rather than scientific advance.

Guidance constitutes a further function. Theory directs inquiry by identifying relevant variables, relationships, and mechanisms. As Machlup (1955) observed, foundational assumptions unify diverse observations under common explanatory principles, enabling systematic inquiry rather than ad hoc interpretation.

Theory as Mental Tools

Mises's (1949) conception of theory as a set of mental tools provides a unifying perspective. Theoretical concepts do not constitute factual knowledge themselves; they render facts intelligible by providing the conceptual apparatus through which reality can be grasped. Theory is necessarily abstract, and abstraction is a condition of generality. As Bylund (2021) summarises Knight's position, deductive theorising isolates essential characteristics of systems, enabling understanding of how they function in essence despite abstraction.

Part B: What Theory IS NOT

Clarifying what theory is not is equally instructive. Sutton and Staw's (1995) critique identified several common substitutes that create the appearance of theory without providing explanation.

References are not theory. Literature citations support arguments but do not constitute them. Extensive review, however sophisticated, remains review unless it advances explanatory relations.

Data are not theory. Empirical findings require interpretation through theoretical frameworks; correlations become meaningful only when embedded in accounts of generative mechanisms.

Diagrams are not theory. Visual representations may summarise theoretical relations, but boxes and arrows do not explain unless accompanied by conceptual accounts of why relationships obtain.

Hypotheses are not theory. Hypotheses specify expected relationships; theory explains why such relationships should exist.

Lists are not theory. Enumerating variables does not explain how or why they interact. Explanation requires organising factors into accounts of mechanisms and processes.

The persistence of these substitutes generates a pseudo-theory problem: scholarly work may appear theoretical while providing no explanatory insight. This distinction matters because theoretical contribution is frequently claimed without meeting theoretical standards. For the present dissertation, the critique functions as a diagnostic check: does the analysis merely catalogue or describe, or does it explain why observed patterns arise?

Part C: Criteria for Strong Theory

What distinguishes strong theory from weak? Weick (1995) proposed that strong theory “explains, predicts, and delights.” Explanation and prediction (understood as pattern prediction in complex systems) are familiar criteria. Delight refers to intellectual satisfaction: strong theory is parsimonious, revealing deep connections with economy of means; it is surprising, challenging taken-for-granted assumptions; and it is generative, opening new lines of inquiry.

This conception aligns with Little’s (1993) account of explanation in the social sciences. Rather than seeking law-like regularities, social science aims to identify causal mechanisms that generate observed patterns – mechanisms often unobservable and accessible only through theoretical abstraction.

Davis’s (1971) criterion of interestingness complements this view. A theory is interesting when it denies what its audience takes for granted: when “what seems to be X is in reality not-X.” Such theory advances understanding by challenging background assumptions rather than confirming them. This criterion connects directly to the problematisation methodology developed in Section 2.3, which proceeds through systematic assumption-challenging.

Conclusion

This section clarified what constitutes theory and how theoretical contribution should be evaluated. Theory explains by specifying relations among concepts within boundary assumptions (Bacharach, 1989). It organises understanding, enables pattern prediction in complex domains, and guides inquiry. Theory is not references, data, diagrams, hypotheses, or lists, though it may employ all of these in support of explanation. Strong theory explains, predicts in an appropriate sense, and delights; interesting theory challenges what is taken for granted.

The dissertation’s contribution is therefore not empirical testing or closed-form hypothesis specification. It operates at a paradigmatic and meta-theoretical register, examining boundary assumptions, assessing their coherence, and evaluating their adequacy to complex organisational phenomena. Criteria appropriate to hypothesis testing apply only indirectly, as implications for subsequent theoretical and empirical development rather than as evaluative standards for the present work

2.1.4 Historical Precedents: Paradigm-Level Theoretical Work

Introduction

The preceding sections established that this dissertation operates within the pure theory tradition (Section 2.1.1), intervenes at the paradigm level (Section 2.1.2), and satisfies recognised standards of theoretical contribution (Section 2.1.3). A natural question therefore arises: does such work possess recognised precedent within the scholarly record?

This section demonstrates that paradigm-level theoretical work without direct empirical testing has distinguished and discipline-shaping precedents. Four foundational contributions are

examined: Kuhn's *The Structure of Scientific Revolutions*, Lakatos's *The Methodology of Scientific Research Programmes*, Lawson's *Economics and Reality*, and Knight's *Risk, Uncertainty and Profit*. Each advanced paradigm-level insights through conceptual analysis and methodological critique rather than empirical testing, and each initially faced resistance precisely because it departed from prevailing empirical or predictive standards.

The purpose of this review is methodological calibration rather than historical exposition. Together, these works establish that assumption-challenging, non-empirical inquiry constitutes a recognised and legitimate mode of scientific contribution.

Four Historical Precedents

Kuhn's *The Structure of Scientific Revolutions* (1962/1970)

Kuhn fundamentally altered the philosophy of science by challenging the view of scientific progress as cumulative knowledge accumulation. Rather than proposing a new theory within existing frameworks, he reconceptualised how scientific development itself is understood.

His central claims – that science alternates between normal and revolutionary phases, that paradigm shifts involve gestalt-like reorientations, and that competing paradigms are often incommensurable – challenged prevailing meta-theoretical assumptions about rationality, progress, and evaluation (Kuhn, 1962/1970). Kuhn's analysis was entirely conceptual and historical, relying on reinterpretation of scientific episodes rather than experimentation or hypothesis testing.

Although initially criticised as irrationalist or sociologically reductive (Lakatos, 1970), Kuhn's work is now foundational across philosophy and history of science, demonstrating that paradigm-level theoretical work can exert enduring influence without conventional empirical validation.

Lakatos's *The Methodology of Scientific Research Programmes* (1970/1978)

Lakatos sought to reconcile Kuhn's insights with a notion of rational scientific appraisal. His contribution was methodological theory – a theory about how scientific theories should be evaluated – rather than a substantive scientific theory.

Lakatos distinguished between the hard core of research programmes, treated as methodologically irrefutable, and the protective belt of auxiliary hypotheses subject to modification. Scientific change, on this account, concerns sequences of theories rather than isolated theory–test episodes, and evaluation turns on whether programmes are progressive or degenerating (Lakatos, 1970; 1974).

The framework itself is non-empirical and non-testable; it provides criteria for appraisal rather than empirical claims about the world. Its sustained influence demonstrates that methodological theory can constitute a substantive scholarly contribution without direct empirical investigation.

Lawson's *Economics and Reality* (1997)

Lawson's work brought paradigm-level methodological critique directly into economics. Drawing on critical realism, he argued that mainstream economics systematically mischaracterises its subject matter by treating social reality as a closed system amenable to formal modelling and prediction.

Lawson contended that persistent predictive failure reflects categorical error rather than technical inadequacy: social systems are open, emergent, and historically structured, rendering mechanistic modelling inappropriate regardless of mathematical sophistication (Lawson, 1997). *Economics and Reality* introduced no new empirical data and conducted no empirical tests; its influence derives entirely from conceptual and methodological argument.

The work has reshaped debates in economic methodology and heterodox economics, illustrating how paradigm-level critique can reorient disciplinary self-understanding without conforming to conventional empirical standards.

Knight's *Risk, Uncertainty and Profit* ([1921] 1985)

Knight's work exemplifies pure economic theory in the tradition outlined in Section 2.1.1. His central contribution – the distinction between measurable risk and unmeasurable uncertainty – had far-reaching implications for theories of profit, entrepreneurship, and economic organisation.

Knight explicitly characterised his analysis as pure theory, aimed at clarifying “the nature and tendencies of the system” rather than providing empirical tests (Knight, [1921] 1985, p. 6). His method of decreasing abstraction relied on conceptual reasoning rather than data collection. Despite the absence of empirical validation, the risk–uncertainty distinction became foundational in economics and remains central to contemporary debates.

Common Features and Implications for Evaluation

These four works share several defining features of paradigm-level theoretical contribution.

First, they are conceptual rather than empirical: their primary contributions lie in reconceptualising phenomena rather than generating new data. Second, they are non-empirical but not non-evidential: each mobilised forms of evidence appropriate to its purpose, including historical analysis, logical argument, and theoretical reconstruction. Third, they are paradigm-challenging rather than gap-filling, questioning the adequacy of prevailing frameworks rather than extending them incrementally. Finally, each was initially contested yet later recognised as foundational.

These features clarify how paradigm-level theoretical work should be evaluated. Such work is not assessed by its capacity to generate testable predictions or empirical confirmations, but by its conceptual coherence, its adequacy to the phenomena it seeks to illuminate, and its capacity to reorient subsequent inquiry.

Conclusion

Paradigm-level theoretical work has distinguished precedents. Kuhn transformed the philosophy of science, Lakatos reshaped the evaluation of scientific theories, Lawson reoriented economic methodology, and Knight introduced concepts that remain central to economic analysis – all without conventional empirical investigation. Their success provides legitimation and guidance for evaluating theoretical contributions operating at this level.

Accordingly, the present dissertation should be assessed within this tradition: not by criteria appropriate to hypothesis testing, but by its coherence, its engagement with recognised conceptual problems, and its capacity to clarify the assumptions that structure existing approaches to the firm and economic organisation.

2.1.5 The Complicated / Complex Distinction and Its Methodological Implications

Introduction

The preceding sections positioned the dissertation within the pure theory tradition, clarified its paradigm-level intervention, and established the legitimacy of non-empirical, assumption-challenging inquiry. This section addresses a foundational classificatory question upon which subsequent methodological reasoning depends: *what kind of phenomena are firms?*

The distinction developed here is between complicated and complex systems. This distinction is ontological rather than terminological, carrying direct epistemological and methodological implications. Complicated systems, however intricate, remain amenable to analytical decomposition and precise prediction given sufficient information. Complex systems exhibit properties – emergence, non-linearity, path-dependence, and fundamental uncertainty – that render such approaches categorically inappropriate.

The purpose of this section is not to advance a substantive theory of the firm, but to establish the system-type classification within which methodological choices become intelligible. The section proceeds in three parts: Part A clarifies the complicated/complex distinction as a difference of kind rather than degree; Part B explains why mainstream economics persistently treats complex phenomena as if they were merely complicated; Part C outlines the methodological implications that follow from recognising complexity.

Part A: The Complicated/Complex Distinction

A Difference of Kind, Not Degree

The distinction between complicated and complex systems concerns the *nature of system relations*, not the number of variables or the difficulty of analysis. A complicated system – such as a mechanical watch or aircraft engine – can, in principle, be understood through analytical decomposition. By identifying components and specifying their relations, the

behaviour of the whole can be derived from the behaviour of its parts. Given sufficient information about initial conditions and boundary parameters, precise prediction is possible.

Complex systems – such as ecosystems, economies, and organisations – cannot be understood through decomposition alone. System-level properties arise from interactions among elements and are not derivable from component-level properties even in principle. As Hayek (1967, p. 26) observed, complexity is characterised by the minimum number of elements required for a pattern to exhibit all its characteristic attributes – a number that is necessarily large in social systems.

Dimensions of the Distinction

This difference of kind manifests across several analytically distinct but interrelated dimensions.

Decomposability. Complicated systems are decomposable: components can be analysed independently and recombined without loss. Complex systems are not decomposable in this way, because their properties arise from interactions that cannot be isolated. Simon's (1962) notion of “nearly decomposable” systems permits limited hierarchical analysis, but presupposes relatively stable subsystem boundaries – an assumption rarely satisfied in social systems, where boundaries are fluid, relational, and historically contingent.

Predictability. Complicated systems permit precise prediction given sufficient information. Complex systems permit only pattern prediction. As Hayek (1967, p. 27) argued, theory can predict the general character of structures that will form under specified conditions, but not their particular manifestations. This limitation reflects system ontology rather than informational deficiency.

Causality. Complicated systems exhibit predominantly linear and traceable causation. Complex systems exhibit non-linear, distributed, and circular causation, with feedback loops and mutual constitution. Causal relations cannot be reduced to simple chains without distortion (Stermann, 2000; Holland, 2014).

Boundaries. Complicated systems have clear and fixed boundaries. Complex systems have boundaries that are permeable, contested, and perspective-dependent. As Cilliers (2005b) emphasised, system boundaries are theoretical constructs that shape what is observed rather than pre-given empirical facts.

Knowledge Type. Complicated systems permit context-independent, formally specifiable knowledge (*episteme*). Complex systems require such knowledge to be complemented by contextual, integrative understanding (*synesis*). Drawing on Aristotle, Flyvbjerg (2001) argued that social complexity demands practical wisdom rather than purely technical rationality.

Woodill (2021) captures the epistemological reorientation these dimensions demand: complexity compels “a shift of focus from objects to relationships.” On this view, “there are no static or inert objects, only relationships... everything is dynamic, in motion, changing, and

interactive.” Traditional Western epistemology, oriented toward discrete objects with fixed properties, proves inadequate for phenomena constituted through dynamic relational processes.

Taken together, these dimensions indicate that complexity entails not analytical difficulty, but a fundamentally different mode of intelligibility.

Part B: Why Economics Treats Complex Phenomena as Complicated

Having established the categorical distinction, the question arises: why does economics persistently treat complex phenomena as if they were merely complicated?

Despite the conceptual clarity of the complicated/complex distinction, mainstream economics has persistently treated complex phenomena as if they were merely complicated. Several mutually reinforcing factors account for this tendency.

First, **physics envy** has played a decisive role. As Mirowski (1989) documented, neoclassical economics modelled itself on nineteenth-century physics, importing equilibrium concepts, mathematical formalism, and predictive aspirations independently of their suitability to social phenomena.

Second, **mathematical tractability** reinforces ontological simplification. Complicated-system assumptions enable optimisation, equilibrium analysis, and comparative statics; complex-system assumptions resist such techniques. Methodological investment in tractable mathematics therefore incentivises simplification rather than adequacy.

Third, the **prediction imperative** has exerted institutional pressure. As Hayek (1974) argued in his critique of the “pretence of knowledge,” economics has been driven to claim predictive precision that its subject matter does not permit. The demand for point prediction encourages the treatment of complexity as noise rather than as a defining property.

Finally, **methodological path-dependence** ensures reproduction. Once mechanistic assumptions become embedded in training, publication standards, and evaluation criteria, they persist independently of their explanatory adequacy. Treating complex phenomena as complicated becomes a background assumption rather than an explicit methodological choice.

Part C: Methodological Implications of Complexity Recognition

These disciplinary factors explain the persistence of misclassification. The question now becomes: what methodological implications follow from correctly recognising complexity?

Recognising that firms and economic organisations exhibit properties characteristic of complex systems entails specific methodological requirements. Analysis must be accompanied by synthesis, attending to how interactions generate emergent properties rather than decomposing wholes into independent parts. Prediction must be reframed as pattern prediction rather than point forecasting. Reasoning must incorporate history and path-dependence rather than assuming ahistorical equilibrium dynamics. Inquiry must remain sensitive to boundaries, recognising their permeability and co-evolution with environments. Models must

accommodate heterogeneity rather than relying on representative-agent abstractions. Explanation must attend to feedback and circular causation rather than assuming linear causal chains. And methodology itself must be reflexive, acknowledging that inquiry into complex systems is itself a complex adaptive process.

These implications function as methodological requirements – conditions of adequacy that any approach to complex phenomena must satisfy, without dictating specific models or techniques. They establish conditions under which methodological choices can be evaluated for coherence with the nature of the phenomena under investigation.

These requirements ground the methodological choices developed in subsequent sections. Problematisation (Section 2.3) follows from non-decomposability: isolated assumption critique proves ineffective when assumptions form networked structures. First-principles reasoning follows from path-dependence: surface modifications regenerate underlying commitments. Transdisciplinarity (Section 2.8) follows from system openness and cross-level interaction. The relationship to empirical testing (Section 2.5) follows from fundamental uncertainty and non-ergodicity. The analytical focus on meso-level phenomena follows from the locus of emergence. These are not arbitrary methodological preferences but requirements imposed by the ontology of the phenomena.

Conclusion

This section established the complicated/complex distinction as a foundational classificatory step with methodological consequences. The distinction is one of kind rather than degree: complicated and complex systems possess different properties and therefore require different modes of understanding. The persistent treatment of complex economic phenomena as merely complicated reflects disciplinary incentives and historical path-dependence rather than ontological adequacy.

Recognising complexity does not, by itself, yield a substantive theory of the firm. It clarifies the limits of certain approaches and the conditions under which alternative methodological strategies become intelligible. Subsequent sections build on this classificatory groundwork in developing the methodological orientation of the dissertation.

Forward References: Section 2.2 introduces the problematisation methodology appropriate to interconnected assumption structures. Section 2.3 develops a first-principles orientation consistent with complexity recognition. Section 2.5 examines the implications of complexity for empirical testing.

Conclusion of Section 2.1

Section 2.1 has positioned the dissertation within a recognised tradition of scholarly inquiry and established the classificatory groundwork upon which subsequent methodological reasoning depends. Five findings carry forward.

First, the dissertation operates within the tradition of pure economic theory, and this positioning is necessary rather than merely permissible. The problems addressed are paradigm-constitutive assumptions that are mutually reinforcing and define the conditions of intelligibility for

empirical work. Pure theory establishes the conceptual apparatus through which empirical phenomena become intelligible; it does not compete with empirical research but enables it. Foundational assumptions therefore require different forms of evaluation – internal coherence, explanatory capacity, and the ability to organise understanding – rather than direct empirical verification.

Second, the dissertation intervenes at the paradigm level, organised across three interrelated dimensions: ontological assumptions about what firms and economic systems are, epistemological assumptions about what can be known in open and complex systems, and methodological assumptions concerning appropriate forms of analysis. The limitations examined are not problems internal to a mechanistic framework but problems of that framework – constitutive blind spots concerning emergence, uncertainty, and complexity.

Third, theoretical contribution is defined by explanation – specifying relations among concepts within boundary assumptions – rather than by references, data, diagrams, hypotheses, or lists. Strong theory explains, enables pattern prediction in complex domains, and challenges what is taken for granted.

Fourth, paradigm-level theoretical work without direct empirical testing has distinguished precedents. Kuhn, Lakatos, Lawson, and Knight each advanced paradigm-level insights through conceptual analysis, each was initially contested, and each became foundational.

Fifth, the complicated/complex distinction establishes that firms exhibit properties – emergence, non-linearity, path-dependence, permeable boundaries, and fundamental uncertainty – that render decomposition-based and prediction-centred approaches categorically inappropriate. This persistent treatment of complex phenomena as merely complicated reflects physics envy, mathematical tractability imperatives, and methodological path-dependence rather than ontological adequacy.

These findings establish what kind of work the dissertation is and the evaluative criteria appropriate to it. Section 2.2 articulates the philosophical foundations – critical realism, hermeneutics, and perspectivalism – that ground the methodological choices to follow.

2.2 Philosophical Foundations

Section 2.1 positioned the dissertation within the pure theory tradition, clarified its paradigm-level intervention, and established the complicated/complex distinction as foundational classificatory groundwork. These positioning decisions, however, presuppose deeper philosophical commitments that must now be made explicit. Methodological choices do not float free of ontological and epistemological foundations; they derive their warrant from prior commitments concerning what exists, what can be known, and how knowledge claims can be justified.

Section 2.2 articulates those foundations. This articulation is necessary because paradigm-level critique – challenging the assumptions that constitute a framework rather than extending it – requires philosophical ground from which critique can proceed without collapsing into either naïve objectivism or self-undermining relativism. The dissertation must be able to assess whether inherited assumptions illuminate or distort their object, and this presupposes both that there is something to illuminate and that knowledge of it remains partial and revisable.

The argument proceeds through five subsections. Subsection 2.2.1 establishes critical realism as the philosophical foundation, explaining why it is preferred to positivism and strong constructivism, and articulating the ontological, epistemological, and inferential commitments that follow. Subsection 2.2.2 examines the hermeneutic dimension of inquiry – how pre-understanding shapes what can be seen and how reflexivity renders assumptions available to critique. Subsection 2.2.3 develops perspectivalism, establishing that the dissertation's claims concern comparative adequacy rather than exclusive truth. Subsection 2.2.4 articulates methodological pluralism, justifying the synthesis of multiple heterodox traditions. Subsection 2.2.5 addresses axiological positioning, making explicit the values that guide the research and the ethical stance toward paradigm critique.

Together, these subsections establish the philosophical conditions under which Section 2.3's problematisation methodology becomes intelligible and warranted.

2.2.1 Critical Realism as Philosophical Foundation

Introduction

The preceding sections positioned the dissertation within the pure theory tradition, clarified its paradigm-level intervention, and established the complicated/complex distinction as grounding subsequent methodological reasoning. Methodological choices, however, presuppose deeper philosophical commitments concerning what exists (ontology), what can be known (epistemology), and how knowledge claims can be warranted (methodology). This section makes those commitments explicit.

Critical realism, developed primarily by Roy Bhaskar and elaborated by scholars including Tony Lawson, Margaret Archer, and Andrew Sayer, provides the philosophical foundation for this dissertation. This choice is not arbitrary. Critical realism offers ontological and epistemological commitments compatible with the analysis of complex, open social phenomena in ways that dominant alternatives – positivism and strong constructivism – are not.

The section proceeds as follows. Part A explains why critical realism is preferred to its principal alternatives. Parts B and C articulate the ontological and epistemological commitments that follow. Part D develops retroduction as the inferential mode appropriate to theory operating

under conditions of openness, emergence, and fallibility. Part E establishes the philosophical warrant these commitments provide for the dissertation's transformation methodology.

Part A: Why Critical Realism?

The Inadequacy of Positivism

Positivism, influential in twentieth-century economics, holds that scientific knowledge consists primarily of verified empirical regularities and that unobservable entities should be excluded from scientific explanation. On this view, laws are understood as constant conjunctions – relations of the form “whenever A, then B” – and theories are valued insofar as they generate testable predictions.

For complex and open social phenomena, positivism is inadequate for three reasons.

First, complex social systems do not reliably exhibit constant conjunctions. Such regularities presuppose closed systems in which mechanisms operate in isolation and external influences can be excluded. Economic and organisational phenomena occur in open systems characterised by interacting mechanisms, contextual variation, and historical contingency. As Lawson (1997, Ch. 2) argues, the persistent predictive failures of mainstream economics reflect not insufficient data or technique but the categorical inappropriateness of expecting regularities of the positivist kind in open systems.

Second, positivism cannot accommodate emergence. Emergent properties – properties of wholes not possessed by their parts – cannot be captured by correlations among observables alone. Restricting scientific attention to regularities among observed events therefore marginalises phenomena central to organisational and economic life.

Third, the mechanisms that generate observed patterns are typically unobservable. Institutional structures, feedback effects, social relations, and causal powers cannot be directly inferred from surface regularities. Positivism's exclusion of unobservables therefore removes from inquiry precisely what explanation often requires.

The Inadequacy of Strong Constructivism

Strong constructivism avoids positivism's reliance on constant conjunctions by holding that reality is entirely socially constructed and that no observer-independent facts exist. While this position foregrounds interpretation and reflexivity, it generates a different problem. If all knowledge is constructed and no reality exists independently of our descriptions, the grounds for preferring one explanatory framework to another become unclear, resting on convention or power rather than adequacy to phenomena.

For paradigm-level critique, this is disabling. Assessing whether inherited assumptions illuminate or distort their object presupposes that there is something to illuminate or distort – some reality not exhausted by discourse. A purely constructivist stance therefore undermines the possibility of non-arbitrary theoretical evaluation.

Critical Realism as the Appropriate Foundation

Critical realism provides a middle path. It is realist in its ontology: a reality exists independently of our knowledge of it, and theoretical frameworks can be more or less adequate to that reality. It is critical in its epistemology: knowledge is always partial, fallible, and theory-laden, and direct access to reality is unavailable.

This combination – ontological realism coupled with epistemic fallibility – supports inquiry into complex social phenomena. It permits critique of existing frameworks without claiming final certainty, and it allows for theoretical improvement without assuming that any framework can deliver complete or definitive knowledge.

Part B: Ontological Commitments

Critical realism's distinctive contribution lies in its ontology. Three commitments are central for the methodological orientation adopted here.

Stratified Reality

Bhaskar (1975, p. 56–62; 1978, p. 13–15) distinguishes three domains of reality: the empirical (experiences and observations), the actual (events that occur whether or not they are observed), and the real (structures, mechanisms, and powers that generate events). Mechanisms may exist without being activated, and activation may or may not produce observable outcomes.

This stratification implies that scientific inquiry is not limited to identifying observed regularities. Its aim is to identify generative mechanisms operating in the real domain that give rise to events in the actual domain, some of which are experienced empirically. Reducing reality to experience – treating what is observed as all that exists – constitutes what Bhaskar termed the “epistemic fallacy.”

For organisational and economic phenomena, stratification is essential. Many of the most consequential causal powers – institutions, routines, structural constraints, and feedback processes – operate as generative mechanisms not directly accessible to observation.

Emergence

Critical realism provides robust grounding for emergence. Elder-Vass (2010, p. 16–23) defines emergent properties as causally efficacious properties of wholes that are not possessed by their parts and are not reducible to part-properties. This goes beyond the trivial observation that wholes differ from parts: emergent properties exert causal influence.

Archer (1995) adds a temporal dimension through the morphogenetic cycle. Social structures emerge from human action but subsequently condition action over time: “structure predates the actions which transform it,” while elaboration postdates action (Archer, 1995, p. 138, 168). This temporal stratification supports treating social entities as more than aggregations, endowed with emergent powers that shape future possibilities.

Open Systems

Lawson (1997; 2003) emphasises that social and economic systems are open. Mechanisms interact, contexts vary, and outcomes are contingent. In open systems, constant conjunctions are neither necessary nor expected. The same mechanism may produce different outcomes under different conditions, and multiple mechanisms may jointly generate observed events.

Openness explains both the limits of prediction and the inadequacy of modelling strategies that presuppose closure. It clarifies why explanatory success in social science cannot be equated with predictive precision.

Part C: Epistemological Commitments

Critical realism combines ontological realism with epistemological fallibility. Two commitments are central.

Fallibilism

All knowledge claims are provisional and revisable. Observation is theory-laden, and access to underlying mechanisms is indirect. Popper's (1963) principle of fallibilism – knowledge advances through conjecture and criticism rather than verification – is extended by critical realism to encompass the possibility that even apparently successful frameworks may misrepresent reality.

For the present dissertation, this implies that its claims are advanced as reasoned judgments open to revision, not as final or definitive truths.

Judgmental Rationality

Fallibilism does not imply relativism. Critical realism maintains that rational assessment among competing theories is possible even in the absence of conclusive proof. Bhaskar (1986, p. 72–75) termed this position “judgmental rationality.” Sayer (2000) elaborates criteria for such assessment, including internal coherence, explanatory power, empirical and practical adequacy (where relevant), and compatibility with established knowledge.

Judgmental rationality provides the basis for comparing frameworks in terms of their ability to illuminate mechanisms, handle emergence, and accommodate openness and uncertainty – without presupposing decisive empirical tests.

Part D: Retroduction as Inferential Mode

Critical realism's characteristic inferential mode is retroduction: reasoning from observed patterns and explanatory problems to the mechanisms that could generate them. Retroduction differs from deduction and induction. Deduction derives conclusions from premises; induction generalises from observations. Retroduction asks: *given that this phenomenon occurs, what must be the case for it to be possible?*

Bhaskar (1975; 1979) formalised retroduction as a transcendental form of argument. Danermark et al. (2002, p. 96–100) describe the process as involving identification of a phenomenon, theoretical redescription, proposal of candidate mechanisms, and comparative assessment of their explanatory adequacy. The outcome is not certainty but the provisional identification of generative structures.

This inferential mode aligns with the dissertation's methodological orientation. Where inherited assumptions generate persistent explanatory difficulties, retroduction provides a disciplined way to infer which underlying commitments or mechanisms could account for those difficulties and to assess alternative conceptualisations in terms of explanatory adequacy.

Part E: Philosophical Warrant for Transformation

These commitments provide warrant for the dissertation's transformation methodology. Stratified ontology legitimates inquiry into mechanisms not directly observable; the claim that mechanistic assumptions misdescribe generative structures becomes assessable rather than merely rhetorical. Emergence ontology legitimates treating firms as wholes with causal powers irreducible to individual actions. Open systems ontology legitimates the rejection of equilibrium-centred modelling and predictive aspirations inappropriate to open phenomena. Fallibilism frames the CAS framework as potentially more adequate representation rather than final truth. Judgmental rationality permits reasoned comparison without requiring conclusive proof.

The substantive demonstration of these transformations is the task of subsequent chapters. This section establishes only that such transformations are philosophically coherent, methodologically legitimate, and epistemically defensible.

Conclusion

Critical realism provides an appropriate philosophical foundation for the dissertation's methodological orientation. Its stratified ontology supports inquiry into unobservable mechanisms; its account of emergence grounds the causal efficacy of social wholes; its emphasis on open systems clarifies the limits of regularity and prediction; and its commitments to fallibilism and judgmental rationality enable comparative theoretical assessment without claims to final proof.

2.2.2 Knowledge, Pre-Understanding, and the Hermeneutic Circle

Introduction

The previous section established critical realism as the philosophical foundation for this dissertation, with commitments to stratified reality, emergence, and open systems. Yet inquiry is shaped not only by ontological and epistemological commitments but also by the fact that understanding always begins from prior understanding. This section examines the hermeneutic dimension of knowledge: how pre-understanding shapes what can be seen, which questions can be posed, and what answers can be recognised as meaningful.

The hermeneutic tradition – especially Martin Heidegger, Hans-Georg Gadamer, and Paul Ricoeur – shows that pre-understanding is not a contaminant to be eliminated but a condition of possibility for interpretation. Consistent with critical realism, the argument here is epistemological rather than ontological: hermeneutics clarifies conditions of knowing without collapsing being into knowing.

The section proceeds in four parts. Part A establishes pre-understanding as constitutive rather than eliminable. Part B explains the hermeneutic circle as the iterative relation between parts and wholes. Part C examines how mechanistic economics can function as a collective pre-understanding in disciplinary practice. Part D develops reflexivity as a strengthening discipline for paradigm-level critique.

Part A: Pre-Understanding as Constitutive

The central hermeneutic claim is simple: understanding always proceeds from prior understanding. We do not confront phenomena as blank slates; we bring assumptions, concepts, expectations, and interpretive frameworks that make recognition and sense-making possible.

Heidegger (1927/1962, p. 191) analysed this through the “fore-structures” of understanding: fore-having (background involvements), fore-sight (perspective), and fore-conception (conceptual resources). These are not optional add-ons; they are the preconditions of interpretation. Gadamer (1960/2004, p. 271–273) radicalised this insight through his rehabilitation of “prejudices” (*Vorurteile*) as productive pre-judgments. The Enlightenment's “prejudice against prejudice” (Gadamer, 1960/2004, p. 273) – the aspiration to view from nowhere – misunderstands the structure of understanding: interpretation always occurs within historically conditioned horizons.

This does not reduce knowledge to subjective bias. The relevant contrast is not “perspective vs no perspective” (the latter is impossible) but blind pre-understanding versus reflexive pre-understanding. Pre-understanding enables perception, organisation of evidence, and anomaly-

detection. The epistemic question becomes: which pre-understandings open phenomena, which close them off, and how can they be rendered available to critique? Gadamer's account of interpretive revision – through encounter and “fusion of horizons” (Horizontverschmelzung) (1960/2004, p. 305) – captures how pre-understandings can be corrected without presupposing neutrality.

Part B: The Hermeneutic Circle

The hermeneutic circle describes the reciprocal dependence of part–whole understanding: parts are intelligible only in relation to a whole, while wholes are grasped only through their parts. This is not a logical vice but the structure of interpretation itself. Gadamer (1960/2004, p. 291–299) illustrates the point in textual interpretation: word-meaning depends on sentence-meaning while sentence-meaning depends on word-meaning; the same iterative movement scales up to paragraphs, chapters, works, and traditions. Ricoeur (1981, p. 197) extends the circle beyond texts to action and institutions: practices can be “read” because they are meaning-laden, context-embedded, and consequence-structured.

For theoretical inquiry, the implication is methodological. Concepts derive meaning from a framework; frameworks derive content from their constitutive concepts. Anomalies are visible only relative to expectations supplied by a framework, while frameworks are revised through confrontation with recalcitrant phenomena. In paradigm-level critique, the circle implies that foundational assumptions cannot be evaluated in isolation, because their function and meaning depend on the wider configuration they jointly sustain. Critique therefore requires iterative movement between assumptions (parts) and their interlocking system (whole) until the dependency structure becomes explicit.

Part C: Mechanistic Economics as Collective Pre-Understanding

Kuhn's (1962/1970, p. 10) analysis of paradigms can be read hermeneutically: paradigms function as collective pre-understandings that shape what a community sees, what questions appear meaningful, what methods count as legitimate, and what constitutes an admissible answer. Paradigms are therefore not merely sets of propositions but interpretive infrastructures.

On this reading, mechanistic economics can function as collective pre-understanding within mainstream economics. Commitments such as equilibrium orientation, optimisation ideals, and the privileging of formal modelling often operate less as explicit theses and more as background standards of “rigour.” Through training, publication norms, and professional evaluation, such commitments can become invisible to those who hold them – precisely because they organise competence and legitimacy.

From a hermeneutic perspective, entrenched pre-understanding can become what Bachelard (1938/2002, p. 24) called an “**epistemological obstacle**”: not necessarily because it is false, but because its success in organising understanding may screen out phenomena that do not fit its categories. This is a diagnostic claim rather than a denunciation. Obstacles arise through taken-for-granted infrastructures, not necessarily through bad faith.

Akerlof's (2020, p. 405–409) “hardness bias” illustrates one mechanism by which such obstacles can be reinforced: methodological preferences for formal precision can channel attention toward questions that permit “hard” treatment, while discouraging engagement with phenomena that resist closure, measurement, or tractable formalism. In such conditions, methodological standards may function as gatekeeping devices, defining as inadmissible precisely those explanatory forms that challenge prevailing pre-understanding.

Part D: Making Pre-Understanding Visible and the Role of Reflexivity

If pre-understanding functions as background, the epistemic task becomes how to make it visible. The hermeneutic answer is: through contrast. Engagement with alternative conceptual vocabularies and explanatory strategies brings implicit assumptions to the surface, revealing them as contingent commitments rather than natural facts. This is why plural engagement can be epistemically productive: it discloses the limits and selectivities of one's own horizon.

Davis's (1971, p. 309) account of theoretical "interestingness" provides an allied point: theoretical significance often lies in challenging what is taken for granted. Contrast does not automatically validate the alternative; it forces the background to appear as background.

Reflexivity strengthens this process by applying the hermeneutic insight to inquiry itself. This dissertation necessarily proceeds from pre-understandings; reflexivity does not demand neutrality but explicitness about commitments and their enabling and constraining effects. This does not entail that all frameworks are equally adequate. Critical realism's judgmental rationality (Bhaskar, 1986, p. 72–75; Sayer, 2000) permits comparative assessment under fallibility. The hermeneutic point is that frameworks shape relevance, salience, and plausibility; rendering these shaping effects explicit strengthens theoretical rigour rather than weakening it.

Conclusion

This section examined the hermeneutic dimension of inquiry: how pre-understanding shapes what can be seen, which questions can be asked, and what answers can be recognised. Pre-understanding is constitutive rather than eliminable, and the hermeneutic circle structures interpretation in theoretical work no less than in textual understanding.

Mechanistic economics can function as collective disciplinary pre-understanding, shaping perception and standards of legitimacy in ways that may remain implicit. Under certain conditions, entrenched pre-understanding can become an epistemological obstacle (Bachelard, 1938/2002), reinforced by incentives such as hardness bias (Akerlof, 2020). Contrast with alternative vocabularies helps render assumptions visible, while reflexivity strengthens inquiry by clarifying commitments and limits without abandoning rational assessment.

2.2.3 Theories as Perspectives: The Epistemic Status of Claims

Introduction

The preceding sections established critical realism as the philosophical foundation (Section 2.2.1) and examined how pre-understanding shapes inquiry (Section 2.2.2). This section addresses a question implicit in both: what is the epistemic status of the theoretical framework this dissertation develops? Does the CAS perspective claim to be "true" while the mechanistic perspective is "false"? Or is something more nuanced at stake?

This section develops **perspectivalism** – the position that theories are not mirrors of reality but perspectival tools that reveal certain aspects of phenomena while necessarily obscuring others. This is not relativism: some perspectives are more adequate than others for particular purposes and domains. But it is a rejection of the view that one theory can exhaustively capture reality while all alternatives are simply wrong.

The section proceeds in four parts. Part A examines problems with the correspondence theory of truth for complex phenomena. Part B develops perspectivalism as an alternative. Part C discusses Rodrik's "a model, not THE model" framework. Part D articulates the implications for the epistemic claims advanced in this dissertation.

Part A: Problems with Correspondence Theory

The correspondence theory of truth holds that statements are true if they accurately describe facts, and that theories are true insofar as they correctly mirror reality. Applied to economics, this would suggest that there exists a single correct theory of the firm – one that accurately captures what firms really are. Competing theories would be competing truth claims; at most one could be correct.

For complex phenomena, the correspondence ideal encounters fundamental difficulties.

The problem of abstraction. All theories abstract from reality, selecting some features as relevant and ignoring others. No theory corresponds to reality in its full complexity. If correspondence requires completeness, no theory can be true. If it does not, the notion becomes indeterminate.

The problem of multiple true descriptions. The same reality may admit multiple valid descriptions depending on which aspects are foregrounded. Firms may be described as legal entities, collections of contracts, production functions, capability repositories, or complex adaptive systems. These descriptions are not strictly compatible, yet none is simply false.

The problem of underdetermination. Any finite body of evidence may be consistent with multiple theories. Evidence alone cannot uniquely determine which theory is true.

The problem of essential complexity. As Hayek (1967, p. 26–28) argued, complex phenomena may not admit of the kind of simple regularities that correspondence presupposes. A theory that captures pattern-level regularities without predicting particular outcomes represents genuine knowledge, not failed correspondence.

These difficulties do not imply that truth is irrelevant or that all theories are equally good. But they suggest that truth in complex domains cannot be exhausted by one-to-one correspondence alone.

Part B: Perspectivalism as Alternative

Ronald Giere (2006) develops **perspectival realism** as an alternative to both naïve correspondence and relativism. On this view, scientific theories provide perspectives on reality – ways of seeing that are constrained by what exists but shaped by the position from which inquiry proceeds.

Perspectivalism is realist: there is a reality independent of our theories, and theories can represent that reality more or less adequately. But representation is always from a perspective, never from nowhere. The relevant question is therefore not “does this theory correspond to reality?” but “what aspects of reality does this perspective reveal, and for what purposes is it adequate?”

This reframes theoretical competition. Mechanistic and CAS frameworks are not rivals in a zero-sum contest for exclusive truth. They are alternative perspectives that illuminate different aspects of economic reality. The methodological question becomes one of **comparative adequacy**, not absolute correspondence.

Nancy Cartwright (1999) sharpens this distinction by differentiating between **truth *in* models** and **truth *of* models**. A model may be internally true – its conclusions follow from its assumptions – without being literally true of the world. The relevant evaluative question is not

whether a framework is true in an absolute sense, but what it enables us to understand, explain, and anticipate – and in which domains.

Part C: Rodrik’s “A Model, Not THE Model”

Dani Rodrik (2015, p. 5–8) articulates perspectivalism in methodological terms for economics: the discipline is best understood as a collection of models rather than a single unified theory. The appropriate model depends on context. Expressions such as “it depends” or “context is all” are not evasions but acknowledgements that different analytical situations require different conceptual tools.

Rodrik’s formulation – “*a model, not THE model*” – captures the core point. The CAS framework is **a way** of understanding firms: one that foregrounds emergence, adaptation, and complexity. It is not *the* way, nor the sole framework that renders all alternatives incorrect. Likewise, the mechanistic framework is **a way** of understanding firms: one that foregrounds equilibrium, optimisation, and predictability. It too is not *the* way.

This does not imply that all frameworks are equally adequate. Rodrik is explicit that model choice must be disciplined by criteria of selection: assumptions should be appropriate to the context, mechanisms should plausibly operate in the domain under study, and conclusions should be relevant to the questions posed. Some models are therefore better than others for particular purposes. What follows is a distinction between “**better for this purpose**” and “**true while all alternatives are false.**”

Part D: Implications for This Dissertation

Perspectivalism provides the methodological warrant for the dissertation’s approach. The CAS framework is proposed as a **more adequate perspective** for understanding contemporary firms. This does not require proving the mechanistic framework false – demonstrating that equilibrium never obtains, that optimisation never occurs, or that prediction is always impossible.

Instead, the dissertation argues that the CAS perspective reveals features of firms – emergence, adaptation, self-organisation, path-dependence – that mechanistic perspectives systematically obscure or marginalise. These features are increasingly consequential for understanding contemporary economic phenomena. On this basis, the CAS framework is argued to be more adequate for present analytical purposes, even while mechanistic approaches may remain useful in restricted contexts.

The epistemic aspiration is therefore **better perspective**, not final or exclusive truth. This is more modest than claiming to have discovered the true theory of the firm while all predecessors were mistaken – but it is not arbitrary, subjective, or merely rhetorical. “Better perspective” is a substantive, assessable claim evaluated through comparison of what each framework reveals, explains, and enables.

Conclusion

This section established perspectivalism as a methodological commitment shaping how the dissertation conducts its inquiry. Theories are not mirrors of reality but perspectival tools; “a model, not THE model” applies to theoretical frameworks as well as formal models. The correspondence theory of truth encounters deep limitations in complex domains; perspectivalism offers a viable alternative that preserves realism and reasoned assessment without requiring impossible proofs.

For this dissertation, perspectivalism provides warrant for proposing the CAS framework without claiming to have falsified mechanistic economics. The epistemic aim is better perspective, not final theory.

2.2.4 Scientific and Methodological Pluralism

Introduction

The previous sections established critical realism as foundation (Section 2.2.1), hermeneutic pre-understanding as a condition of inquiry (Section 2.2.2), and perspectivalism as the epistemic stance toward claims (Section 2.2.3). This section develops **pluralism** – the position that complex phenomena may require multiple, irreducible perspectives and methods for their adequate understanding.

Pluralism is **not** relativism. It does not assert that all perspectives are equally valid or that no rational basis exists for preferring one framework to another. Rather, it recognises that the complexity of the phenomena themselves may exceed what any single theoretical lens or methodological approach can capture exhaustively.

The section proceeds in four parts. Part A explains why methodological monism fails for complex phenomena. Part B distinguishes forms of pluralism. Part C clarifies how pluralism operates without abandoning judgmental rationality. Part D demonstrates how pluralism shapes this dissertation’s research design.

Part A: Why Monism Fails for Complex Phenomena

Methodological monism – the view that there exists a single correct scientific method applicable to all domains – has long exercised powerful appeal. It promises unity, rigour, and progress. Physics has historically served as the monist exemplar; economics’ well-documented “physics envy” reflects this aspiration.

Friedrich Hayek (1967, p. 22–42) demonstrated that methodological monism fails precisely where phenomena are complex. Methods appropriate to simple phenomena – where the number of relevant variables is small, interactions are linear, and conditions can be controlled – become categorically inappropriate for phenomena of *organised complexity*, where numerous elements interact non-linearly.

For such phenomena, Hayek argued, scientific advance must take a different form: identifying patterns rather than point predictions, understanding types of structures rather than particular instances, and accepting that precise numerical values may be unavailable. The implication is methodological pluralism, not methodological retreat.

John Dupré (1993, p. 7–12) provides ontological support for pluralism. His account of the “disorder of things” rejects the assumption that reality forms a unified, hierarchically ordered system governed by a single set of fundamental laws. Instead, the world comprises a plurality of partially connected domains, each exhibiting its own kinds of regularities. If reality itself is ontologically plural, methodological monism is misguided in principle.

Nancy Cartwright (1999, p. 1–7) develops a complementary argument. Her “dappled world” is governed by local regularities, not universal laws. Stable laws operate only within carefully engineered “nomological machines” – highly controlled environments that rarely exist outside laboratories. For economics, this implies that the search for universal laws may be methodologically misplaced. Different contexts may legitimately require different theoretical frameworks.

Part B: Types of Pluralism

Three forms of pluralism are relevant to this dissertation.

Ontological pluralism holds that reality is not reducible to a single fundamental level. Multiple levels of organisation – physical, chemical, biological, psychological, social – possess their own properties, regularities, and causal powers. For understanding firms as complex adaptive systems, ontological pluralism is indispensable. Firms exist simultaneously at multiple levels: as legal entities, as collections of individuals, as configurations of routines, and as adaptive systems. These are not competing descriptions; they are different ontological levels of the same complex phenomenon.

Methodological pluralism holds that no single scientific method is appropriate for all domains. Different kinds of phenomena require different methods. For complex adaptive systems, this means pattern identification rather than point prediction, synthesis alongside analysis, and historical-institutional sensitivity rather than ahistorical modelling.

Integrative pluralism, as developed by Sandra Mitchell (2003), holds that integration across perspectives is possible and valuable, but integration occurs in application rather than through abstract theoretical unification. Theoretical pluralism is preserved, while integration happens pragmatically when addressing specific questions.

Part C: Pluralism and Judgmental Rationality

Pluralism does not entail relativism. Judgmental rationality (Section 2.2.1) remains intact: we can assess comparative adequacy even if no framework can be proven universally true.

Pluralism accepts multiple legitimate perspectives without accepting that all perspectives are equally good. Criteria such as explanatory power, empirical adequacy, coherence, and practical relevance remain operative. What pluralism rejects is the monist expectation that these criteria will uniquely select one framework for all purposes.

In Rodrik's (2015) terms, "it depends" is epistemic honesty, not evasion. If mechanisms operate differently across contexts, then conditional answers are not weakness but accuracy. Some models are better than others for particular purposes – the distinction is between "better for this purpose" and "true while all alternatives are false."

Part D: How Pluralism Shapes This Research

Pluralism shapes this dissertation's research design by justifying synthesis across multiple heterodox traditions rather than commitment to a single alternative framework. The dissertation integrates insights from:

- **Austrian economics:** dispersed knowledge, spontaneous order, entrepreneurial discovery, fundamental uncertainty
- **Evolutionary economics:** variation–selection–retention, routines, path-dependence, novelty
- **Institutional economics:** rules, norms, governance, bounded rationality
- **Complexity economics:** non-linearity, emergence, self-organisation, adaptation, networks

These traditions are not identical, but they share a rejection of mechanistic assumptions and a sensitivity to complexity. Pluralism permits synthesising their insights rather than choosing among them. The CAS framework developed in this dissertation constitutes an **integrative**

pluralist synthesis – not eclecticism, but structured integration grounded in shared rejection of mechanistic assumptions and unified through CAS concepts.

Distinction from Chapter 6. This section explains how pluralism shapes the **conduct of this research**. Chapter 6 argues that economics **should** embrace pluralism as a disciplinary norm. Here, pluralism is a research design commitment; there, it becomes a disciplinary recommendation. The distinction is deliberate and methodologically crucial.

Conclusion

Pluralism is established here as a methodological commitment guiding the dissertation's inquiry. Monism fails for complex phenomena because no single perspective or method can capture their full structure. Ontological, methodological, and integrative pluralism provide principled grounds for legitimate diversity. The dissertation's CAS framework exemplifies integrative pluralism in practice, synthesising Austrian, evolutionary, institutional, and complexity economics into a coherent alternative to mechanistic monism. Judgmental rationality remains intact: the claim is comparative adequacy, not universal truth.

2.2.5 Axiological Positioning: Values and the Role of the Researcher

Introduction

The preceding sections established the philosophical and methodological foundations of this dissertation: critical realism (Section 2.2.1), hermeneutic pre-understanding (Section 2.2.2), perspectivalism (Section 2.2.3), and pluralism (Section 2.2.4). This section addresses a further dimension that is often left implicit but is unavoidable in any serious inquiry: **axiology** – the values that guide research.

The positivist tradition held that science should be value-free. Weber (1904/1949) articulated this ideal as *Wertfreiheit*, while simultaneously recognising *Wertbeziehung* – the unavoidable role of values in selecting research problems. Subsequent scholarship has shown that values enter research not only in problem selection but also in methodological choices and interpretive framing.

This section adopts the position that values are unavoidable and that making them explicit enhances rather than diminishes rigour. As Myrdal (1958) argued, biases cannot be eliminated by silence; they can only be disciplined by exposure.

Part A: Values Guiding This Research

Five values guide this dissertation. Making them explicit enables readers to assess how they shape the analysis.

Intellectual honesty requires representing positions – including those under critique – as their proponents would recognise them. This dissertation adopts a *steelmanning* approach, engaging the strongest versions of mechanistic economics rather than caricatures.

Comprehensiveness means addressing the mechanistic paradigm systematically across ontological, epistemological, methodological, and relational dimensions, rather than selectively targeting weak assumptions while ignoring stronger ones.

Fairness requires recognising mechanistic economics as a serious intellectual tradition with genuine achievements. The claim advanced is not that mechanistic economics is useless, but

that its foundational assumptions become increasingly inadequate as economic complexity intensifies.

Transparency requires making commitments visible rather than presenting conclusions as if they followed inevitably from neutral observation. This dissertation makes explicit its philosophical, methodological, and theoretical commitments.

Humility requires recognising the provisional character of all theoretical claims, particularly in the context of complexity. As Cilliers (2005, p. 256, 262–263) emphasises, failure to acknowledge complexity is not merely a technical error but an ethical one. Claims to complete knowledge of complex systems are irresponsible.

Part B: Ethical Stance Toward Paradigm Critique

Paradigm critique carries ethical responsibilities because paradigms are embedded in scholarly communities and institutions.

Responsibility to represent fairly means that critique must be accurate, proportionate, and informed. Strawman arguments undermine both scholarly norms and the credibility of the critic.

Engagement with strongest versions means that effective critique must address the strongest defensible versions of a paradigm, not weak formulations.

Constructive transformation means that critique without alternative is irresponsible. Scholars cannot simply abandon a paradigm without replacement. This dissertation is therefore committed to constructive transformation: each FROM→TO shift replaces a mechanistic assumption with a CAS-compatible alternative addressing the same phenomenon.

Part C: Explicit Values Enhance Rigour

Values enter research at every stage: problem selection, method choice, interpretation, and presentation. Suppressing discussion of values does not remove them; it renders them invisible.

When values are explicit, readers can assess whether conclusions follow given the stated commitments, where alternative values might alter interpretation, and whether practice aligns with professed standards. This scrutiny is impossible when values remain hidden.

Following Harding (1991), this dissertation pursues **strong objectivity**: objectivity achieved through reflexive transparency rather than fictitious neutrality. Knowledge is situated; objectivity arises from examining how position shapes inquiry, not from denying that position exists.

Cilliers (2005) demonstrates that ethical considerations are intrinsic to complexity research. Because we cannot fully calculate or control complex systems, methodological and theoretical choices are unavoidably normative. The values articulated here – intellectual honesty, comprehensiveness, fairness, transparency, and humility – are therefore internal requirements of complexity-appropriate inquiry, not external moral add-ons.

Conclusion

This section has articulated the axiological positioning of the dissertation. Values are unavoidable; the choice is between hidden values and explicit ones. This dissertation adopts the latter, pursuing strong objectivity through transparency and reflexivity. By making its axiological commitments explicit, the dissertation enhances rather than compromises rigour.

Conclusion of Section 2.2

Section 2.2 has articulated the philosophical foundations that ground the dissertation's methodological orientation. Five findings carry forward.

First, **critical realism** provides an appropriate philosophical foundation for paradigm-level inquiry into complex social phenomena. Its stratified ontology – distinguishing the empirical, actual, and real domains – supports inquiry into generative mechanisms not directly accessible to observation. Its account of emergence grounds the causal efficacy of social wholes. Its emphasis on open systems clarifies why constant conjunctions and predictive precision are inappropriate standards. Epistemologically, fallibilism acknowledges that all knowledge claims are provisional, while judgmental rationality maintains that rational comparison remains possible without conclusive proof. Retroduction provides the inferential mode appropriate to theory under conditions of openness.

Second, **hermeneutic analysis** reveals that inquiry always proceeds from pre-understanding. This is not a defect but a condition of possibility for interpretation. The hermeneutic circle structures theoretical work: foundational assumptions cannot be evaluated in isolation because their meaning depends on the wider configuration they jointly sustain. Mechanistic economics can function as collective disciplinary pre-understanding that may become epistemological obstacle when it screens out phenomena that resist its categories. Reflexivity strengthens rigour.

Third, **perspectivalism** establishes the epistemic status of the dissertation's claims. Theories are not mirrors of reality but perspectival tools; “a model, not THE model” applies to theoretical frameworks. The correspondence theory encounters deep limitations in complex domains. The dissertation's epistemic aspiration is comparative adequacy – better perspective – not exclusive or final truth.

Fourth, **methodological pluralism** justifies the dissertation's synthesis of multiple heterodox traditions (Austrian, evolutionary, institutional, complexity). Monism fails for complex phenomena; different domains require different methods. The CAS framework constitutes an integrative pluralist synthesis, not eclecticism. Judgmental rationality remains operative.

Fifth, **axiological positioning** makes explicit the values guiding the research: intellectual honesty, comprehensiveness, fairness, transparency, and humility. Strong objectivity is achieved through reflexive transparency rather than fictitious neutrality. Values are internal requirements of complexity-appropriate inquiry.

These philosophical foundations do not adjudicate substantive disputes. They establish the conditions under which paradigm-level assumption-challenging can proceed with warrant. Section 2.3 operationalises these commitments through the problematisation methodology.

2.3 The Problematization Methodology

Section 2.2 articulated the philosophical foundations – critical realism, hermeneutic pre-understanding, perspectivalism, pluralism, and axiological transparency – that ground the dissertation’s inquiry. Those foundations establish what kinds of claims can be warranted and how knowledge proceeds under conditions of openness, emergence, and epistemic fallibility. They do not yet specify the logic of inquiry through which the dissertation’s theoretical contribution is developed. Section 2.3 supplies that specification by establishing problematization as the dissertation’s primary methodological stance.

This methodological articulation is necessary because the dissertation’s object is not an empirical gap to be filled but a paradigmatic structure of assumptions to be examined and reconstructed. Gap-spotting – the dominant mode of research question construction – produces novelty by extending literatures while leaving foundational commitments intact. Where the limitations concern the adequacy of assumptions themselves, gap-spotting is structurally insufficient. Problematization responds by treating theoretical advance as emerging from challenging the assumptions that organise a literature rather than extending it.

The argument proceeds through seven subsections. Subsection 2.3.1 defines problematization and contrasts it with gap-spotting. Subsection 2.3.2 elaborates the six methodological principles that discipline assumption-challenging. Subsection 2.3.3 situates problematization within a longer intellectual lineage spanning Socratic inquiry, dialectic, and Eastern traditions of category-disruption. Subsection 2.3.4 explains why problematization is combined with first-principles thinking to ensure depth-reach. Subsection 2.3.5 clarifies the levels at which assumptions are engaged and specifies the three-tier architecture of mechanistic assumptions that the dissertation targets. Subsection 2.3.6 introduces the FROM→TO transformation framework as the operational device for executing problematization at paradigmatic depth. Subsection 2.3.7 derives the dissertation’s research questions as methodological products of this process.

Together, these subsections establish not only the stance but the operational discipline through which Section 2.4’s phenomena construction and Section 2.5’s account of the relationship to empirical research become intelligible.

2.3.1 The Problematization Methodology Introduction

Section 2.2 established the philosophical commitments shaping this dissertation – critical realism, hermeneutic pre-understanding, perspectivalism, pluralism, and axiological transparency. The present section turns from foundations to methodology: the logic of inquiry through which the dissertation’s theoretical contribution is developed.

The dissertation adopts **problematization** as its primary methodological stance. As Hans-Georg Gadamer observed, understanding advances through questioning: “the path to all knowledge leads through the question” (Gadamer, 1960/1994, p. 363). Problematization, in Mats Alvesson and Jörgen Sandberg’s influential formulation, treats theoretical advance as emerging from **challenging the assumptions** that organise a literature, rather than extending that literature by filling empirical lacunae. This contrasts with the dominant “gap-spotting” approach, which constructs novelty by identifying missing variables, contexts, or populations while leaving underlying assumptions intact.

Problematization has deeper roots. Michel Foucault described it as an attempt “to think differently” rather than to legitimise what is already known (Foucault, 1985, p. 9), and as “the

development of a given into a question” (Foucault, 1984, p. 389). In this sense, problematization names a disciplined transformation: what appears as settled “fact” is re-located as a theoretical choice, and the taken-for-granted becomes accountable to inquiry.

This section proceeds in two parts. Part A defines problematization, clarifies what is meant by “assumptions,” and contrasts problematization with gap-spotting. Part B explains why problematization is the appropriate stance for paradigm-level work aimed at conceptual reconstruction rather than incremental refinement.

Part A: Definition and Recognition

Defining Problematization

Alvesson and Sandberg (2011, p. 247) define problematization as “generating research questions through a dialectical interrogation of one’s own familiar position, other stances, and the domain of literature targeted for assumption challenging.” Four elements are essential.

First, dialectical interrogation indicates engaged critique rather than external denunciation or routine extension: it requires serious encounter with existing positions in their strongest formulations while refusing to treat enabling assumptions as beyond scrutiny. Second, interrogation includes one’s own familiar position. Assumption-challenging cannot be purely outward-directed because implicit commitments often appear as “common sense” until contrasted with alternatives. As Slife and Williams (1995, p. 71) note, without contrast, implicit ideas do not appear as ideas at all. Third, problematization requires engagement with other stances. As Paul Feyerabend puts the point, critique requires alternative assumptions as an external reference standard (Feyerabend, 1978, p. 31). Finally, problematization is targeted: it is not generic skepticism but disciplined engagement with identifiable literatures whose assumptions can be specified and assessed.

The Nature of Assumptions

Problematization treats assumptions as boundary conditions of intelligibility: they shape what counts as an admissible question, what is accepted as explanation, and what is dismissed as noise. Samuel Bacharach’s definition makes the point explicit: theory is a statement of relations among concepts “within a set of boundary assumptions” (Bacharach, 1989, p. 496). If assumptions bound theory, then critique must be able to surface and evaluate those assumptions.

A key difficulty is that many influential assumptions are implicit. Gouldner’s distinction between explicit postulations and background assumptions remains instructive: theory often rests upon assumptions that are “un-postulated and un-labeled,” forming the background from which explicit claims emerge (Gouldner, 1970, p. 29). Problematization is directed precisely at this background level – rendering visible what typically operates as invisible infrastructure.

Assumptions also vary in depth. Some are local; others are deep commitments operating across domains. This matters because foundational assumptions not only guide answers; they delimit what a field can even recognise as a coherent question.

Contrast with Gap-Spotting

Alvesson and Sandberg (2011) identify **gap-spotting** as the dominant mode of research question construction in management and organisation studies, with clear analogues in economics. Gap-spotting typically generates questions of the form “we know X, but not yet

Y,” producing novelty by adding variables, extending contexts, or applying existing models to new settings – while preserving the accepted frame.

Empirical studies of publishing practices document the prevalence of this pattern. Locke and Golden-Biddle (1997) found that qualitative contributions in leading journals overwhelmingly justified novelty by portraying prior work as incomplete. Sandberg and Alvesson (2011) similarly show how research questions are often legitimated through “gap construction” rather than assumption interrogation.

Gap-spotting has virtues: it anchors work in ongoing conversations and yields sharply bounded contributions. But it has a structural limitation. By taking the underlying assumptions as given, it tends to reproduce and stabilise them. Alvesson and Sandberg (2011, p. 252) therefore argue that gap-spotting “reinforces rather than challenges the assumptions underlying existing literature.” Where the problem is internal incompleteness, this may be sufficient. Where the limitations concern the adequacy of the assumptions themselves, it is not.

Why Gap-Spotting Limits Theoretical Impact

Gap-spotting can inhibit assumption-challenging by defining novelty as “what is missing within the accepted frame,” thereby shifting attention away from whether the frame is adequate. Over time, this encourages incremental accumulation – complicating models without clarifying their explanatory commitments. Editorial reflections have repeatedly expressed concern about “footnote-on-footnote” accumulation and the scarcity of work that re-frames foundational questions (Daft & Lewin, 1990; 2008).

A deeper reason concerns what makes theory consequential. Murray S. Davis argues that theoretical significance is closely tied to challenging taken-for-granted assumptions: a theory is “interesting” when it denies what an audience assumes without collapsing into implausibility (Davis, 1971). This helps explain why theories that merely extend established commitments often appear obvious, while theories that deny everything appear absurd. Influential theories typically renegotiate the field’s enabling assumptions.

Subsequent studies support this logic. McKinley, Mone, and Moon (1999) show that attention-gaining theories tend to diverge from established work while remaining intelligible within the field. Bartunek, Rynes, and Ireland (2006) similarly report that assumption-challenging contributions are disproportionately perceived as important. Problematization takes these insights seriously by treating assumption re-examination as a systematic route to meaningful theoretical advance.

Scholarly Recognition

Problematization is a recognised methodology in theory-oriented scholarship. Its influence stems from translating a broad critical tradition into a systematic account of how assumptions can be identified and interrogated in ways that generate constructive theoretical development. Later work has extended the approach through discussions of problematizing research questions and reviews (Alvesson & Sandberg, 2013; 2020), consolidating problematization as a legitimate path to theory advancement where foundational assumptions matter.

Part B: Why Problematization Is Appropriate Here

The appropriateness of problematization depends on the kind of limitation a dissertation addresses. If a literature is broadly adequate but incomplete, gap-spotting may be sufficient. Problematization becomes appropriate when the core limitations concern the assumptive structures that define what counts as explanation.

In paradigm-level inquiry, the object of evaluation is not merely hypotheses but background commitments – ontological, epistemological, and methodological – that function as rules of admissibility. Such commitments are not straightforwardly testable like isolated empirical claims. They shape what counts as evidence, what counts as a mechanism, and what is recognised as anomaly. Accordingly, they are best assessed through criteria such as comparative adequacy, internal coherence, explanatory reach, and the kinds of phenomena they render visible or invisible – criteria that align naturally with assumption interrogation.

Problematization is therefore suited to conceptual reconstruction: making explicit what has become naturalised, questioning its adequacy for the phenomena at hand, and opening space for alternative assumptions that remain connected to recognised explanatory tasks. Critique is not an end in itself; it is a route to more coherent and capable theoretical articulation.

Finally, problematization fits the dissertation's intended stance: neither narrow novelty nor mere dissent, but a disciplined attempt to "think differently" in Foucault's sense – to transform givens into questions (Foucault, 1984; 1985) while remaining accountable to standards of fairness, clarity, and constructive orientation.

Conclusion

This section established problematization as the dissertation's methodological stance. Problematization generates theoretical contribution by challenging the assumptions that structure a literature, distinguishing it from gap-spotting approaches that produce novelty by filling lacunae within an accepted framework. Its logic is anchored in the hermeneutic primacy of questioning (Gadamer), in Foucault's account of turning givens into questions, and in scholarship showing that influential contributions typically renegotiate assumptions rather than merely extend established propositions (Davis; Alvesson & Sandberg).

Problematization is appropriate where the central limitations of existing work concern the adequacy of foundational assumptions rather than incomplete coverage of variables or contexts. Under those conditions, assumption-challenging becomes not an optional rhetorical strategy but a methodological requirement of serious paradigm-level theory work.

2.3.2 The Six Methodological Principles of Problematization

Introduction

The previous section established problematization as a recognised methodology for generating research questions through assumption-challenging, distinguishing it from gap-spotting approaches that extend literatures while leaving foundational commitments intact. This section specifies the methodological discipline of problematization by elaborating the six principles proposed by Mats Alvesson and Jörgen Sandberg (2011, p. 255–259).

Problematization is not algorithmic. Terence Deacon is explicit that assumption-challenging “cannot be reduced to a mechanical or even strictly analytical procedure” and necessarily involves a creative act (Deacon, 2000, p. 135). This does not imply arbitrariness. It implies that methodological discipline lies not in following a recipe but in sustaining a set of guiding commitments that structure interrogation without claiming procedural certainty.

The need for such disciplined interrogation is amplified by what George Akerlof terms disciplinary “sins of omission,” driven by economics’ “hardness bias”: a tendency to ignore problems that resist tractable formalisation (Akerlof, 2020). Problematization responds by creating structured space for questioning assumptions that remain powerful in professional reproduction while weakening in explanatory reach (Alvesson & Sandberg, 2011, p. 258). The six principles below are presented sequentially for clarity; in practice they are interdependent and iteratively revisited.

The Six Principles

Principle 1: Identifying a Domain of Literature

Problematization begins by specifying the domain to be targeted for assumption-challenging. This is not merely bibliographic delimitation but an interpretive act: deciding what counts as the relevant conversation and where its enabling assumptions are most clearly expressed (Alvesson & Sandberg, 2011, p. 256).

Rationale. Problematization typically requires selective depth rather than exhaustive coverage. Whereas gap-spotting often performs comprehensiveness as rigour, problematization depends on close engagement with canonical, programmatic, and path-defining texts – those that stabilise what is treated as “normal” within a research community.

Application. This dissertation targets mainstream economics’ theory of the firm, focusing on the mechanistic paradigm underlying neoclassical, transaction cost, and agency theories. Path-defining studies (Coase, 1937; Williamson, 1975, 1985; Jensen & Meckling, 1976) are examined alongside subsequent work that reproduces their assumptions.

Principle 2: Identifying and Articulating Assumptions

The core operation of problematization is making assumptions explicit – especially the background commitments that make a literature’s claims appear natural.

Rationale. Alvin Gouldner’s distinction remains instructive: explicit postulations often rest upon “unpostulated and unlabeled” background assumptions (Gouldner, 1970, p. 29). Problematization targets this level by transforming what is treated as fact into assumptions open to scrutiny (Alvesson & Sandberg, 2011, p. 256).

Assumption types. Alvesson and Sandberg distinguish in-house assumptions (local to a subfield), root metaphor assumptions (broad images shaping how the subject is imagined), paradigm assumptions (ontological, epistemological, methodological commitments), ideology assumptions (normative/moral presuppositions), and field assumptions (widely shared, often invisible because “everyone” shares them).

Assumption functions. Alan Musgrave distinguishes negligibility, domain, and heuristic assumptions (Musgrave, 1981). His warning about “assumption slippage” is methodologically crucial: assumptions may migrate from empirically meaningful negligibility claims to domain restrictions or mere heuristics, reducing testability while retaining rhetorical authority.

Application. Chapter 3 identifies fifteen mechanistic assumptions across all assumption types, analysing their functional roles using Musgrave’s typology. The CAS framework provides the necessary external standpoint for rendering these assumptions visible.

Principle 3: Evaluating Which Assumptions Are Worth Challenging

Problematization is not indiscriminate scepticism. It requires selection: which assumptions have sufficient theoretical leverage to warrant challenge?

Rationale. Alvesson and Sandberg (2011, p. 258) stress the recurring question: what is the theoretical potential of challenging this assumption?

Evaluation criteria. Scope and leverage (broad assumptions have higher potential impact); explanatory adequacy (assumptions merit challenge when they persist while contributing weakly to understanding); fruitfulness (the goal is conceptual opening); temporal dynamics (assumptions can become “conventional wisdom,” losing generative capacity); logical status (Musgrave’s distinctions matter against instrumentalist tolerance for unrealistic assumptions).

Application. Mechanistic assumptions are problematised because they are paradigmatic, widely shared, increasingly inadequate for complex adaptive systems, and have transformed from substantive claims into unfalsifiable domain restrictions.

Principle 4: Developing an Alternative Assumption Ground

Problematization is constructive. Challenging assumptions without developing alternatives produces critique without theoretical advance.

Rationale. Imre Lakatos’s point is directly relevant: in sophisticated practice there is “no falsification before the emergence of a better theory” (Lakatos, 1974, p. 119). Frameworks are displaced when an alternative offers explanatory surplus – accounting for what the old framework handled while illuminating what it obscured.

Application. The dissertation develops CAS-compatible alternatives to each mechanistic assumption through systematic FROM→TO transformations. These integrate Austrian, evolutionary, institutional, and complexity economics into a unified framework.

Principle 5: Considering Assumptions in Relation to Audience

Problematization operates within research communities, not outside them. Assumption-challenging must therefore be calibrated to audience structure, disciplinary norms, and institutional gatekeeping.

Rationale. Different sub-communities have different tacit commitments and different thresholds for what counts as acceptable critique. Hardness bias helps explain defensive

responses: where rigour is equated with “hard” styles, arguments requiring different kinds of evidence face additional barriers (Akerlof, 2020).

Application. The dissertation addresses multiple audiences by accurately representing mechanistic economics, acknowledging its achievements, framing CAS as a more adequate perspective rather than exclusive truth, and maintaining scholarly politeness while sustaining substantive critique.

Principle 6: Evaluating the Alternative Assumption Ground

The final principle asks whether the alternative assumption ground is likely to generate theory that is both compelling and productive.

Rationale. Murray S. Davis’s account provides a strong heuristic: theoretical moves are ignored if obvious and rejected if absurd; the “interesting” zone lies between (Davis, 1971). Additional criteria include fruitfulness (Van de Ven, 2007), explanatory surplus (Lakatos, 1974), and argument quality (Astley, 1985).

Application. The CAS framework is positioned as interesting rather than radical or trivial: it challenges equilibrium, optimisation, and decomposability while remaining connected to established heterodox traditions and generating novel explanatory capacity.

Conclusion

This section elaborated the six methodological principles through which problematization is conducted: delimiting a literature domain, articulating assumptions, selecting assumptions with high theoretical leverage, developing alternative assumption grounds, considering audience structure and disciplinary constraints, and evaluating whether the alternative is both compelling and fruitful.

Two meta-points follow. First, problematization is disciplined without being algorithmic: it cannot be reduced to a recipe, yet it is guided by identifiable methodological commitments. Second, in disciplines marked by hardness bias and institutionalised omission, systematic assumption-challenging is not an indulgence but a corrective: it reopens questions that methodological closure tends to foreclose.

2.3.3 Precursor Methodologies: The Lineage of Problematization

Introduction

Problematization is often presented as a contemporary methodology for generating research questions by challenging the assumptions underpinning established literatures. Yet the core gesture it formalises – treating the taken-for-granted as questionable and rendering the unquestioned visible – has a longer intellectual lineage. This section situates problematization within that lineage, showing that systematic assumption-challenging is not an idiosyncratic modern invention but a recognisable family of epistemic practices spanning philosophical traditions.

The purpose of this genealogy is methodological rather than antiquarian. It serves three functions: **(1) legitimization**, by showing that disciplined assumption-challenging has long been treated as a serious cognitive practice; **(2) enrichment**, by supplying additional vocabularies for describing how presuppositions are exposed without collapsing into mere provocation; and **(3) clarification of method**, by highlighting that in such work “method” is often best understood as an oriented path rather than a preset procedure – especially when the targets are deep assumptions embedded in professional common sense.

Two families of precursor practices are examined: a Western lineage associated with Socratic inquiry and dialectical thinking, and an Eastern lineage associated with disciplined disruption of habitual categories and heightened attention to relational “in-betweenness.” The aim is not to claim direct historical continuity, but to illuminate methodological kinship: distinct traditions converging on the epistemic task of making the background visible.

Part A: Ancient Greek Precursors

The Socratic Method (Elenchus)

A clear early template for assumption-challenging is the Socratic elenchus: systematic questioning directed at beliefs treated as obvious and self-authenticating. Through interrogation, implications are drawn out until tensions or contradictions surface, producing a cognitive state in which what seemed self-evident becomes doubtful – and therefore thinkable.

Three features are methodologically decisive. First, elenchus targets **secure convictions**, anticipating the central intuition of problematization that the most consequential assumptions are often those least noticed. Second, it is **structured** rather than arbitrary: critique is anchored in commitments already granted by the interlocutor, revealing instability from within rather than imposing standards from without. Third, contradiction is **productive**: refutation functions as a clearing that opens conceptual space for revised understanding.

Dialectic: Preservation Through Transcendence

Dialectical reasoning provides a second precursor: movement through positions that does not merely negate a target but preserves what is valid while exposing its limits. This is methodologically important for paradigm-level work because it avoids a common failure mode of critique – treating an established framework as mere error and proposing an alternative as total replacement. A dialectical stance supports critique-with-continuity: paradigms are treated as historically productive achievements that may nevertheless reach the limits of adequacy under altered conditions.

Maieutics: Inquiry as Elicitation

The Socratic “midwifery” motif adds an enabling dimension. Understanding is not simply transmitted as doctrine; it is drawn out through questioning. This matters because deep assumptions are rarely abandoned by direct refutation alone; they loosen when agents recognise, through guided interrogation, tensions between lived experience and inherited categories. Maieutic inquiry is therefore not only critical but generative: it facilitates the emergence of alternative intelligibility.

Defamiliarization

A further precursor is defamiliarization: showing that what appears natural may be historically contingent. The transition from mythic explanation to rational inquiry illustrates how explanatory frames can become invisible as frames. The methodological analogue is direct: where a framework has become indistinguishable from reality, critique begins by rendering it visible as a set of commitments rather than a mirror of the world.

Part B: Japanese Philosophical Precursors

Koan Practice as Disciplined Category-Disruption

In Zen training, the koan functions less as a puzzle than as a device for disrupting habitual conceptualisation. Its methodological significance lies in how it targets the limits of prevailing categories. Where ordinary reasoning proceeds within established frames, koan practice forces sustained confrontation with what cannot be resolved inside those frames, inducing disciplined doubt that can culminate in perspective transformation.

The analogy to problematization is precise: when assumptions function not merely as premises but as conditions of intelligibility – determining what questions can be asked and what answers can be recognised – further analysis within the same frame may deepen error rather than correct it. In such situations, the task is not incremental refinement but **category revision**: altering the grammar of explanation.

Ma (間): The Significance of the “Between”

Japanese aesthetics offers a second precursor in the concept of Ma – the interval or “between” that gives form and meaning to what it separates. Ma is not emptiness but structuring absence: relational spacing through which elements become intelligible.

This bears directly on paradigm-level assumption-challenging. Paradigms are not mere lists of discrete assumptions; they are structured ensembles in which assumptions mutually stabilise one another. A single assumption may appear defensible in isolation, yet derive its force from its position within a network. Ma provides a vocabulary for shifting attention from isolated propositions to the relational architecture that binds and protects them. The object of problematization becomes not only assumptions-as-nodes, but the configuration that makes them resilient.

Part C: Toward an Enriched Conception of Problematization

These precursor traditions do not replace contemporary problematization frameworks; they clarify what such frameworks presuppose and deepen the methodological vocabulary available for describing assumption-challenging. Four enrichment effects follow.

1. **From critique to transformation:** Socratic interrogation and koan-like disruption converge on the point that assumption-challenging is productive only when it changes what can be seen as a legitimate question and what can count as an adequate answer.
2. **From refutation to reconstruction:** dialectic formalises a disciplined transformation – preserving what earlier frameworks captured while revising what they obscure – supporting non-caricatural critique.
3. **From doctrine to elicitation:** maieutics reframes critique as enabling recognition; assumption-challenging brings latent tensions into articulate form rather than merely opposing a position.
4. **From isolated assumptions to relational architecture:** Ma directs attention to interdependence; paradigms persist not because each assumption is independently irresistible, but because the ensemble forms a stabilising structure. Problematization becomes, in this enriched sense, an architectural practice targeting configurations rather than propositions alone.

Conclusion

Problematization can be understood as a contemporary formalisation of a much older methodological impulse: to expose the unexamined, loosen the grip of inherited categories, and create conceptual space in which alternative understanding becomes possible. Ancient Greek precursors illuminate disciplined interrogation, dialectical preservation-through-transcendence, and inquiry as elicitation. Japanese precursors illuminate the methodological role of category disruption and relational “in-betweenness” – the architecture through which assumptions stabilise each other as a paradigm.

Taken together, these lineages clarify what is at stake in problematization: not merely generating new research questions, but enabling shifts in the conditions under which questions and answers become intelligible. In this sense, problematization is best understood not as a technical procedure appended to research, but as an epistemic practice: a way of thinking differently by making the background visible.

2.3.4 Combining Problematization with First-Principles Methodology

Introduction: Why Problematization Alone Is Insufficient

Problematization is indispensable because it renders assumptions visible and open to challenge. Yet, on its own, it does not guarantee that inquiry will reach the depth at which paradigms are generated and reproduced. Assumption-challenging can remain confined to surface propositions, local modelling conventions, or in-house theoretical commitments, leaving deeper generative structures intact.

Because this dissertation targets paradigmatic reconstruction rather than local refinement, methodological depth is a condition of success. Problematization is therefore combined with first-principles thinking – not as philosophical embellishment, but as an operational strategy for identifying and interrogating foundational commitments.

Part A: The Problem of Theoretical Depth

Theoretical knowledge is layered. Specific models and propositions are produced within broader frameworks, which themselves rest on deeper ontological and epistemological commitments (Simon, 1962; Ackoff, 1971). What appears as a technical modelling choice at the surface often depends on presuppositions about what exists, what counts as causation, and what constitutes valid knowledge.

Critiques that remain at surface levels are easily absorbed. In Lakatosian terms, they modify the protective belt while leaving the hard core intact (Lakatos, 1970). In Kuhnian terms, anomalies are managed within normal science unless foundational commitments are displaced (Kuhn, 1962/1970). If the aim is paradigmatic reconstruction, the methodology must be capable of identifying the generative commitments that stabilise and reproduce the paradigm.

Part B: First-Principles Thinking as a Depth-Reaching Method

First-principles thinking refers here to the disciplined identification of assumptions treated as starting points rather than conclusions – commitments that function as conditions of possibility for subsequent theorising.

Two features are decisive. First, first principles are typically tacit. They are presupposed rather than argued for, which is why paradigms appear natural to their adherents. Second, first principles are generative. They actively shape what can be modelled, what counts as explanation, and what is relegated to noise or anomaly.

Applied to economics, first-principles analysis allows the dissertation to move beyond disputes over particular models toward the foundational commitments organising conceptions of firms, markets, knowledge, and causation.

Part C: Why the Combination Is Necessary

Problematization and first-principles thinking perform complementary functions.

- **Problematization** supplies the stance and discipline of assumption-challenging, converting what is treated as given into an object of inquiry (Alvesson & Sandberg, 2011).
- **First-principles thinking** supplies depth-reach, enabling identification of structurally prior commitments and tracing how surface debates depend on them.

Neither is sufficient alone. First-principles analysis can identify foundations without compelling their revision; problematization can challenge assumptions without ensuring it has reached generative depth. Their combination enables directed assumption-challenging at the level required for paradigmatic reconstruction.

Part D: Implications for Research Design

The combined methodology structures the dissertation through four design moves:

1. **Depth identification:** locating foundational commitments shaping modelling, explanation, and interpretation.

2. **Relational mapping:** treating assumptions as interdependent rather than isolated.
3. **Coherent replacement:** constructing alternatives that form a consistent stance rather than ad hoc relaxations.
4. **Systematic reconstruction:** articulating transformations across ontology, epistemology, methodology, and firm conceptualisation.

Evaluation proceeds not through hypothesis testing, but through standards consistent with judgmental rationality: comparative adequacy, coherence, explanatory reach, and the capacity to render intelligible phenomena previously obscured (Section 2.2.1).

Conclusion

Paradigmatic critique and reconstruction require both assumption-challenging and depth-reaching capacity. By combining problematization with first-principles thinking, this dissertation targets not merely surface propositions but the structured commitments through which a research tradition becomes coherent, resilient, and – under changed conditions – potentially inadequate.

2.3.5 The Typology of Assumptions: Levels of Assumption Engagement

Introduction

Having established the combined methodology of problematization and first-principles thinking, this section clarifies both the levels at which assumptions are engaged and the specific architecture of assumptions that the dissertation targets. Not all assumption challenges are equivalent. Challenging a local modelling choice differs categorically from challenging the background commitments that define what counts as a legitimate model.

The purpose here is methodological transparency: to locate the dissertation's intervention accurately and to specify the structure of mechanistic assumptions with sufficient precision that readers can assess both the diagnosis and the proposed transformation.

Part A: The Alvesson and Sandberg Typology

Alvesson and Sandberg (2011) distinguish multiple assumption levels, ranging from narrow commitments internal to specific literatures to broader commitments structuring entire research communities:

- **In-house assumptions:** local to particular subfields.
- **Root metaphor assumptions:** dominant images shaping how phenomena are imagined.
- **Paradigm assumptions:** ontological, epistemological, and methodological commitments.
- **Field assumptions:** background commitments shared across multiple traditions.

The deeper and broader the assumption, the greater its potential theoretical impact – and the greater the interpretive care required to avoid polarisation or dismissal (Davis, 1971; Alvesson & Sandberg, 2011).

Part B: Where This Dissertation Operates

This dissertation engages assumptions primarily at **paradigm** and **field** levels, with root metaphors functioning as a bridge between them. The aim is not to refine a particular firm model within an existing framework, but to assess whether the dominant background commitments through which firms are conceptualised remain adequate under conditions of interdependence, non-linearity, and adaptation.

Accordingly, the dissertation's stance involves:

- **Root metaphor engagement:** examining metaphor as a structuring device (firm-as-machine versus firm-as-adaptive system).
- **Paradigm engagement:** interrogating ontological, epistemological, and methodological commitments.
- **Field engagement:** engaging assumptions shared across mainstream approaches despite local disagreements.

Part C: The Three-Tier Architecture of Mechanistic Assumptions

Tier 1: Foundational Assumptions (Hard Core)

Following Lakatos (1970, p. 91–196), the paradigm's hard core comprises nine foundational commitments, organised into ontological and epistemological categories.

Ontological Foundations (5):

1. **Reductionism** – Wholes are explained entirely by their parts; emergence is denied or treated as merely apparent.
2. **Determinism** – Future states follow necessarily from present states given sufficient knowledge; genuine novelty is excluded.
3. **Substance ontology** – Entities are stable things with fixed essential properties; change is secondary to persistence.
4. **Closure** – Systems are treated as isolated from their environment; external influences are controlled or ignored.
5. **Atomism** – Entities are self-contained units; relations are external and non-constitutive.

Epistemological Foundations (4):

6. **Certainty** – Knowledge is or can be complete; uncertainty reflects temporary ignorance.
7. **Universal laws** – Context-independent regularities govern phenomena; local knowledge is derivative.
8. **Monism** – One correct theory exists for each domain; theoretical competition yields convergence.
9. **Perfect information** – All relevant knowledge is or can be available to decision-makers.

These nine assumptions are not independent. They mutually require and stabilise one another: reductionism presupposes atomism; determinism depends on closure; substance ontology underwrites the stable regularities that universal laws describe; certainty follows from determinism plus perfect information.

Tier 2: Derivative Assumptions (Protective Belt)

From the Tier 1 hard core follow six derivative assumptions that pervade mainstream economic analysis:

10. **Equilibrium** – Systems tend toward stable states where forces balance; disequilibrium is transitional.
11. **Linearity** – Effects are proportional to causes; small causes produce small effects.
12. **Homogeneity** – Agents are interchangeable; structural differences are parametric variations.
13. **Ergodicity** – Time averages equal ensemble averages; the past is a reliable guide to the future.
14. **Optimisation** – Agents maximise objective functions subject to constraints.
15. **Decomposability** – Wholes can be analysed into independent parts without loss.

These derivatives constitute what may be called the “complexity-as-complicatedness” category error: if causation is proportional (linearity), analysis separable (decomposability), agents identical (homogeneity), and systems stabilising (equilibrium), then “complex” collapses into “complicated” – a distinction the CAS framework treats as fundamental.

Tier 3: Methodological Regime (Observable Practice)

Tier 1 and Tier 2 commitments manifest as observable methodological style, evaluation norms, and modelling defaults:

- Physics envy (pursuit of universal laws and predictive precision)
- Representative agent modelling (homogeneity + monism)
- Optimisation formalism as default analytical framework
- Elegance-over-truth aesthetic (parsimony valued over adequacy)
- Hardness bias (Akerlof, 2020): preference for tractable over important

Part D: The Triple Lock – Three Feedback Loops

The three-tier architecture is self-reinforcing through three interlocking feedback loops:

Loop 1 (Ontological): reductionism ↔ equilibrium ↔ closure

Decomposition into parts enables equilibrium analysis; equilibrium presupposes that external perturbations are controlled (closure); closure licenses treating wholes as reducible to parts.

Loop 2 (Epistemic): determinism ↔ ergodicity ↔ optimisation

Determinism licenses treating future states as knowable from present states; ergodicity ensures past patterns predict future patterns; optimisation presupposes that agents can identify best responses given deterministic payoffs.

Loop 3 (Structural): perfect rationality ↔ representative agent ↔ equilibrium

Perfect rationality justifies modelling all agents identically (representative agent); identical rational agents coordinate to equilibrium; equilibrium presupposes that deviations are corrected through optimal responses.

Interlock: Equilibrium appears in Loops 1 and 3; optimisation links Loops 2 and 3. The result is redundant stabilisation. Breaking one node triggers defence from multiple loops. The paradigm is **triple-locked against piecemeal reform**.

Part E: Why Piecemeal Transformations Fail

The three-tier network explains the historical pattern of “successful critiques, failed transformations”:

- Behavioural economics challenges rationality yet is absorbed as a “neoclassical repair shop” – anomalies are patched while equilibrium and optimisation logic remain intact.
- Evolutionary economics offers sustained attack on neoclassical assumptions yet remains marginal; without coordinated ontological transformation, equilibrium logic reasserts.
- Institutional economics adds institutions but often treats them as constraints on otherwise standard agents, leaving atomistic ontology untouched.

Network logic: Challenge one derivative while preserving hard-core commitments → the system either absorbs the challenge (as a special case) or regenerates the challenged derivative from intact foundations. The core regenerates the belt; the loops absorb perturbations; the regime re-stabilises.

Part F: The Missing Middle – Meso-Level Elimination

The mechanistic network produces a systematic blind spot: meso-level dynamics are eliminated. Micro (optimising individuals) connects directly to Macro (aggregates via representative agent), while meso (firms, networks, industries, institutions) is structurally suppressed.

Yet CAS dynamics are typically meso-phenomena. As Arthur (2013, p. 14) notes, key economic phenomena operate at multiple scales but “usually... in between the micro and macro.” Mechanistic elimination of the meso follows directly from reductionism (intermediate levels redundant), atomism (no constitutive relations), homogeneity (no structural differentiation), and universal laws (context and meso irrelevant).

This is not an accidental omission; it is a structural consequence of the assumption network. The dissertation’s focus on firms as complex adaptive systems is therefore a direct response to this meso-level blindness.

Conclusion

This section has located the dissertation’s intervention at paradigm and field levels – where assumption challenge has maximum impact but also maximum resistance – and has specified the mechanistic paradigm as a three-tier system: a Lakatosian hard core (Tier 1), a protective belt (Tier 2), and an observable methodological regime (Tier 3). The paradigm is not a list but a network held in place by a Triple Lock of feedback loops.

This architecture explains why piecemeal reforms fail: preserved foundations regenerate challenged derivatives or absorb them as special cases. Section 2.3.6 operationalises the solution: a FROM→TO transformation that targets the hard core coherently so the protective belt and methodological regime change as consequences rather than as isolated repairs.

2.3.6 The FROM→TO Transformation as Problematization Method

Introduction: Operationalizing Paradigm-Level Transformation

Problematization and first-principles thinking establish *what* must be challenged and *how deeply* inquiry must reach. A further requirement follows: an operational structure capable of carrying assumption-challenging through to reconstruction. Paradigm-level critique cannot proceed through scattered objections or isolated assumption “relaxations.” It requires a disciplined format that specifies what is challenged, on what grounds, and with what alternative.

The dissertation adopts the **FROM→TO transformation framework** as its operational device for executing problematization at paradigmatic depth. The framework functions as a methodological control mechanism that enforces precision, constructiveness, and coherence across the reconstruction.

Part A: The FROM→TO Template

Each transformation follows a strict bi-polar structure:

- **FROM:** an assumption as it operates within the targeted paradigm, articulated in its strongest plausible form and situated in its explanatory role.
- **TO:** a substantively different replacement assumption, grounded in a complexity-informed framework and designed to preserve what is explanatorily valuable while addressing systematic blind spots.

The TO is not a negation but a replacement. The transformation is progressive rather than denunciatory: the FROM is treated as historically productive under certain conditions; the TO is introduced as more adequate under conditions of openness, interdependence, and adaptation.

Part B: Why the FROM→TO Structure Is Methodologically Necessary

Precision in critique.

The FROM requirement forces specificity. It prevents critique from remaining impressionistic by requiring explicit articulation of what assumption is challenged and how it functions within the paradigm.

Obligatory reconstruction.

The TO requirement prevents critique from terminating in negation. Each challenge carries a constructive burden: an alternative assumption must be proposed, justified, and shown to improve adequacy for the dissertation’s target phenomena.

Comparative evaluation.

The paired structure enables evaluation by comparison rather than assertion. Readers can assess whether the TO resolves diagnosed limitations, preserves explanatory strengths, expands reach, and remains coherent with the dissertation’s overall stance.

This aligns with the dissertation's epistemic register: paradigms are evaluated through comparative adequacy, coherence, and explanatory reach rather than through falsification of isolated assumptions.

Part C: Implementation Across the Dissertation

To maintain discipline, the dissertation applies a repeated transformation grammar:

1. **FROM articulation:** the assumption is stated in its strongest defensible form.
2. **Diagnostic tension:** its limitations under relevant conditions are identified.
3. **TO construction:** an alternative assumption is developed and grounded.
4. **Implications:** consequences for explanation, interpretation, and inquiry are traced.

This standardisation prevents rhetorical drift and conceptual fragmentation.

Part D: Coherence and Ordering

Paradigmatic assumptions form coherent ensembles rather than independent propositions. The FROM→TO framework therefore imposes a consistency constraint: alternatives must be mutually compatible and must not reintroduce, indirectly, the commitments they replace.

While not a mechanical sequence, reconstruction respects an ordering logic: ontological commitments constrain epistemological stances, which in turn constrain methodological choices. Methods are thus derived as fitting the kind of object the firm is taken to be.

Conclusion

The FROM→TO framework operationalises problematization as constructive paradigm-level work. It ensures that critique is precise, reconstruction is obligatory, and alternatives form a coherent stance rather than a patchwork. The dissertation's contribution lies not only in challenging assumptions, but in deploying a disciplined transformation format capable of carrying assumption-challenging through to systematic reconstruction.

2.3.7 Research Questions: Derived Through Problematization Introduction: Questions as Methodological Products

In conventional research design, questions precede inquiry. This dissertation adopts the reverse orientation. Consistent with problematization, research questions are treated as *products of method*: they crystallised through systematic assumption identification, evaluation, and reconstruction.

When the object of inquiry is a paradigmatic structure of assumptions, research questions cannot be specified a priori without presupposing what must first be examined. The questions therefore emerge endogenously through iterative engagement with the literature.

Part A: Primary Research Question

How should firms be reconceptualized when the foundational assumptions of mainstream economic theorising are systematically problematized and reconstructed to be compatible with a complex adaptive systems perspective?

Why this question could not be posited a priori

This question did not exist as a fixed starting point. It formed through the convergence of three recognitions generated by the preceding methodological work:

1. **Recognition of inadequacy:** recurring explanatory tensions arise when firms are analysed under assumptions suited to closed, stable, and decomposable settings.
2. **Recognition of depth:** these tensions are not adequately addressed by “filling gaps” or extending models at the surface level; they arise from deeper commitments about what firms are and how economic knowledge is possible.
3. **Recognition of reconstruction requirement:** criticism alone is insufficient; a viable alternative requires explicit replacement assumptions capable of carrying explanatory burdens under conditions of openness, interdependence, and adaptation.

The primary question expresses these recognitions in a single formulation. It is thus best understood as the **methodological output** of the problematization process rather than its initial input.

Part B: Sub-Questions as Operational Requirements

The primary question is necessarily broad. To make it operational without reducing it to a sequence of independent tasks, it is decomposed into a set of sub-questions. Importantly, these sub-questions are not “add-ons.” They correspond to distinct requirements of problematization as a methodology of theoretical reconstruction.

SQ1. Assumption identification

What are the most consequential assumptions structuring mainstream economic theorising about firms?

This question arises from the first operational demand of problematization: one cannot challenge assumptions that remain unarticulated.

SQ2. Diagnostic evaluation

In what ways do these assumptions become inadequate when firms are treated as open, adaptive, and relational phenomena?

Problematization requires more than asserting deficiency; it requires diagnosing what the assumptions enable, what they exclude, and what tensions they generate.

SQ3. Ontological reconstruction

What alternative ontological commitments are required to conceptualise firms as complex adaptive systems?

Once diagnosis reaches ontological depth, reconstruction must respond at the same level; otherwise, revision remains superficial.

SQ4. Epistemological reconstruction

What alternative epistemological stance follows from a complexity-informed ontology of the firm?

If firms are reconceptualized as adaptive processes, the standards of knowledge appropriate to them must also be reconsidered.

SQ5. Methodological reconstruction

What methodological approach becomes appropriate when firm behaviour is analysed under a complexity-consistent ontology and epistemology?

This ensures that method is not imposed as an external preference but derived as a fit between object, knowledge, and inquiry.

These sub-questions should be read as **co-requirements** rather than as detachable steps. The dissertation's architecture reflects this: it is organised so that identification enables diagnosis, diagnosis motivates reconstruction, and reconstruction culminates in synthesis.

Part C: Implementation in the Dissertation (Rewritten – Submission-Safe)

To ensure that problematization operates as a disciplined method rather than as episodic critique, the dissertation applies the FROM→TO transformation through a **standardised transformation unit** across its substantive chapters. While the specific content of each transformation varies by topic – ontology, epistemology, methodology, and the theory of the firm – the underlying methodological grammar remains constant.

Each transformation proceeds through four analytically distinct but internally connected moments:

1. **FROM articulation.**

The assumption under examination is stated explicitly and in its strongest plausible form. This includes not only the assumption's propositional content, but also its functional role within the paradigm: what it enables, what it excludes, and how it structures explanation, modelling, and interpretation. This step enforces fidelity and guards against caricature.

2. **Diagnostic tension.**

The analysis identifies why the assumption becomes problematic under conditions central to the dissertation's object of inquiry – specifically, conditions of openness, interdependence, non-linearity, and adaptation. The aim is not to declare error, but to show how an assumption

that was productive under certain boundary conditions generates systematic tensions, blind spots, or misclassifications under others.

3. TO construction.

An alternative assumption is developed as a constructive replacement. The TO is not introduced as an arbitrary preference or as an isolated correction, but as a commitment grounded in a complexity-informed understanding of economic and organisational phenomena. It is developed with explicit attention to what it preserves from the FROM and what it revises.

4. Implications.

The consequences of the transformation are traced for explanation, interpretation, and research practice. This step demonstrates that the TO is not merely conceptually appealing but capable of carrying explanatory and analytical burdens.

The repeated use of this transformation logic performs an important methodological function. It reduces the risk of **rhetorical drift**, where critique becomes performative rather than analytical, and **conceptual drift**, where alternatives accumulate as loosely connected preferences rather than as components of a coherent stance.

Part D: Coherence Across Transformations (Rewritten – Submission-Safe)

A paradigm is not a collection of independent propositions. Its assumptions form a mutually reinforcing configuration: commitments about what exists, what can be known, and what counts as explanation tend to stabilise one another. A central methodological challenge for paradigm-level reconstruction is therefore avoiding **patchwork revision** – the introduction of local alternatives that resolve one problem while generating contradictions elsewhere.

The FROM→TO framework addresses this challenge by treating transformations as **mutually constraining** rather than independent. Each TO is developed under a coherence requirement: it must remain compatible with the dissertation's broader ontological, epistemological, and methodological commitments, and it must not reintroduce – explicitly or implicitly – the very assumptions the transformation is intended to revise.

This coherence discipline is particularly important given the dissertation's engagement with multiple theoretical traditions. Pluralist engagement does not imply "anything goes." It implies disciplined integration of complementary perspectives under explicit standards of adequacy and consistency. The FROM→TO framework functions as a safeguard against superficial eclecticism by requiring that alternative assumptions form a consistent stance rather than a set of ad hoc substitutions.

A further methodological implication concerns **ordering**. While the dissertation does not treat transformation as a mechanical sequence, it recognises a non-arbitrary direction of dependence: assumptions about what the world is like constrain assumptions about what can be known, which in turn constrain assumptions about how inquiry should proceed. This ordering reflects the critical realist intuition that the character of the object places limits on the kinds of knowledge and methods that are appropriate to it.

Accordingly, the dissertation's reconstruction proceeds with an internal logic that ensures later methodological proposals are anchored in prior ontological and epistemological re-specifications. Methods are not introduced as external techniques, but as responses fitted to the kind of object – the firm – that the reconstructed framework claims to be analysing. This is not a rigid pipeline, but a coherence constraint: methodological choices must remain consistent with the ontological and epistemological commitments that precede them.

Conclusion

The dissertation's research questions are methodological achievements rather than premises. They crystallised through disciplined problematization, were refined through first-principles analysis, and stabilised through systematic reconstruction. The questions instantiate the dissertation's methodological stance: inquiry begins not from unexamined givens, but from the work of making assumptions explicit, interrogating their consequences, and reconstructing them into a more adequate conceptualisation of the firm.

Conclusion of Section 2.3

Section 2.3 has established problematization as the dissertation's primary methodological stance and specified the operational discipline through which assumption-challenging proceeds. Six findings carry forward.

First, problematization generates theoretical contribution by challenging the assumptions that structure a literature rather than filling gaps within an accepted framework. It treats assumptions as boundary conditions of intelligibility – shaping what counts as admissible question, legitimate mechanism, and recognisable anomaly – and renders them available to critique through contrast with alternatives.

Second, problematization is disciplined without being algorithmic. The six principles – identifying a domain, articulating assumptions, evaluating theoretical leverage, developing alternatives, considering audience, and assessing fruitfulness – structure interrogation without claiming procedural certainty. In disciplines marked by hardness bias, systematic assumption-challenging reopens questions that methodological closure tends to foreclose.

Third, problematization has a longer lineage than its contemporary formulation suggests. Socratic elenchus, dialectical reasoning, maieutic inquiry, and Eastern traditions of category-disruption converge on a shared epistemic task: making the background visible and enabling shifts in what can count as question and answer. The concept of *Ma* directs attention from isolated assumptions to the relational architecture through which paradigms achieve resilience.

Fourth, problematization is combined with first-principles thinking to ensure depth-reach. Assumption-challenging that remains at surface levels modifies the protective belt while leaving the hard core intact. First-principles analysis identifies the generative commitments that stabilise and reproduce a paradigm, enabling critique at the level required for reconstruction.

Fifth, the dissertation engages assumptions at paradigm and field levels rather than local modelling choices. Assumptions at these depths are assessed through comparative adequacy, coherence, and explanatory reach rather than through isolated hypothesis testing. Engagement is relational: assumptions reinforce one another, and reconstruction proceeds through coordinated re-specification rather than piecemeal relaxation.

Sixth, the FROM→TO transformation framework operationalizes problematization as constructive paradigm-level work. Each transformation specifies what is challenged, on what grounds, and with what alternative – enforcing precision, obligatory reconstruction, and coherence across the dissertation's substantive chapters.

These methodological commitments produce the dissertation's research questions as outputs rather than inputs. The primary question – how the firm should be reconceptualized when foundational assumptions are systematically problematized – crystallised through the process Section 2.3 describes. Section 2.4 turns to the construction of the phenomenon through which this reconceptualization is pursued.

2.4 Phenomena Construction

Section 2.3 established problematization as the dissertation's methodological stance and specified the FROM→TO transformation framework through which assumption-challenging proceeds. Before foundational assumptions can be systematically examined, however, a prior clarification is required: what counts as "the firm" for theoretical purposes is not a neutral given but a constructed research phenomenon.

This clarification is necessary because the dissertation's contribution is not to apply complexity vocabulary to an already fixed object. It is to re-describe and reframe what counts as the firm – constructing "firms as complex adaptive systems" as a research phenomenon through which organisational dynamics become intelligible in ways that mechanistic framings systematically obscure.

The argument proceeds through two subsections. Subsection 2.4.1 establishes the foundational distinction: research phenomena are constituted through theoretical choices, not simply discovered. Observation is theory-laden; what becomes visible depends on the conceptual resources brought to inquiry. Subsection 2.4.2 specifies how the dissertation constructs its central phenomenon, foregrounding emergent properties, adaptive reconfiguration, relational dynamics, historical contingency, and the limits of prediction – features that mechanistic framings background or treat as residual.

Section 2.5 addresses the relationship between this constructed phenomenon and empirical research.

2.4.1 What Is Phenomena Construction?

Introduction

This section establishes a foundational methodological distinction: research phenomena are not simply discovered but are constituted through theoretical and methodological choices. This distinction is central to the present dissertation because its contribution is not to “apply” complexity vocabulary to an already fixed object, but to re-describe and reframe what counts as the firm for theoretical purposes. Treating “firms as complex adaptive systems” as a constructed research phenomenon is therefore not rhetorical; it is an explicit acknowledgement of the conditions under which social-scientific objects become intelligible and investigable.

The Nature of Phenomena Construction

Phenomena construction refers to the process through which researchers generate and stabilise the objects of investigation – what is taken to be the phenomenon requiring explanation. As Alvesson and Sandberg argue, research “not only includes empirical work aimed at more or less ‘given as fact’ phenomena” but also involves “the process of generating and establishing phenomena to investigate and theorize” (2024, p.1738). Importantly, this process is not a preliminary step but “integral to, both the empirical and theorizing phases in research” (Alvesson & Sandberg, 2024, p. 1738).

At a minimal level, a research phenomenon is what is being studied – “any problem, issue, or topic that is chosen as the subject of an investigation” (Van de Ven, 2016, p. 265). Methodologically, however, this “choice” is inseparable from the conceptual resources through which the world becomes visible. This aligns with the philosophical meaning of phenomenon as appearance: how something presents itself under a particular description and horizon of

meaning (Heidegger, 1927/1962, p. 51–63). What is “noteworthy,” “discernible,” or “regular” (Hacking, 1983, p. 221) is therefore not neutral but framework-dependent.

Accordingly, many philosophers and methodologists emphasise that research phenomena “are typically not just lying around waiting to be discovered,” but must be “created, further developed, and perhaps radically revised” (Bourdieu et al., 1991; Hacking, 1983, 1999). Van de Ven captures the point succinctly: “phenomena do not exist objectively ‘out there’; they are uniquely perceived and framed by different people” (2016, p. 265). The claim is not that firms are fictitious, but that what counts as “the firm” for explanatory purposes depends on how features are selected, organised, and related.

The Theory-Ladenness of Observation

The constructive character of phenomena follows from a familiar epistemological insight: observation is theory-laden. As Knight noted, “one does not hear or see much more than one understands... What we immediately, consciously, apprehend is the ‘meaning’” (1940, p. 17). What is taken to be the phenomenon is thus shaped by the conceptual resources brought to the world.

From a complexity perspective, this precludes any “view from nowhere.” As Cilliers observes, “we cannot have complete knowledge of complex systems; we can only have knowledge in terms of a certain framework... We choose our frameworks” (2005, p. 259). This choice is not a methodological failure but the condition of intelligibility in inquiry – especially where the object is open, relational, and adaptive.

Distinguishing Phenomena, Data, and Theory

A further clarification is essential. Phenomena are not identical with data, and neither is identical with theory. As Bogen and Woodward argue, “a phenomenon is something which is to be explained by theory, while data... are the evidence for the existence of phenomena” (1988, p. 306). Data are traces or indicators; phenomena are the explanatory targets inferred from and organised through those traces.

Kaplan’s distinction remains useful: data describe observed patterns, whereas theory explains why such patterns occur or are expected (1964, p. 349). Mintzberg’s reminder is methodological rather than polemical: “The data do not generate theory – only researchers do that” (1979, p. 584). Data, phenomena, and theory are therefore mutually shaping without collapsing into one another (Apel, 2011; Fuhse, 2022).

Theories as Perspectives

This account aligns with perspectival realism. Theories function as interpretive vantage points that organise what becomes visible as an explanatory object (Giere, 2006a, 2006b). As Fuhse puts it, theories are “webs of concepts” guiding observation and sense-making; access to “real features” is always mediated by perspective (2022, p. 99, 118). This does not imply equivalence of perspectives, but requires evaluation in terms of what a perspective enables us to see, explain, and connect empirically.

Implications

Four implications follow.

1. **Responsibility for constitution:** researchers are responsible not only for analysing phenomena but for how they constitute them; framing is a substantive epistemic commitment.
2. **Apparent contradiction as perspectival difference:** theoretical disagreements may reflect different constructions of the phenomenon rather than competing explanations of the same object.
3. **Legitimacy of reconstruction:** when dominant constructions systematically obscure relevant features, reconstructing the phenomenon is a defensible theoretical contribution.
4. **Perspectival without relativism:** construction does not imply “anything goes.” As Bourdieu et al. emphasise, originality presupposes a break with immediate “readings-off” of the real (1991, p. 14), while remaining accountable to the realities addressed.

Why Phenomena Construction Matters for This Dissertation

The phenomena construction perspective is essential here because definitions of research phenomena partly reflect researchers’ ontological and paradigmatic stance (Alvesson & Sandberg, 2024, p. 1742). The mechanistic paradigm constructs “the firm” as a production function, nexus of contracts, or governance structure – each a theoretically loaded framing that renders some properties visible while suppressing others. Recognising construction enables explicit scrutiny of the assumptions embedded in dominant firm constructions – and legitimates alternative constructions when dominant ones systematically fail.

Relation to Problematization

Phenomena construction and problematization are complementary operations. Problematization exposes the limits of established constructions by challenging their assumptions; phenomena construction builds a new object of inquiry on alternative assumptions. In this dissertation, problematization reveals that mechanistic firm constructions presuppose away precisely those properties that characterise complex adaptive systems; phenomena construction then develops “firms as CAS” as an alternative phenomenon designed to make those properties theoretically tractable.

This complementarity structures the dissertation’s methodology: Section 2.3’s problematization identifies what is foreclosed by mechanistic assumptions; Section 2.4’s phenomena construction specifies what becomes visible when those assumptions are transformed.

Avoiding Misunderstandings

Two misunderstandings warrant explicit exclusion.

First, phenomena construction does not imply “anything goes.” Framework choice carries responsibility; constructions remain constrained by their capacity to enable productive engagement with organisational reality (Cilliers, 2005, p. 264). Frameworks that fail to illuminate, that generate persistent anomalies, or that foreclose significant dynamics are legitimately open to critique and reconstruction.

Second, the claim is not anti-realist. Firms exist and constrain what can be said about them; what is denied is that the categories through which firms are understood are given by reality rather than theoretically constituted. The firm as production function and the firm as complex

adaptive system are different constructions of the same independently existing entities – constructions that render different properties salient and invite different explanatory strategies.

2.4.2 The Construction of “Firms as Complex Adaptive Systems”

Introduction

Having established that research phenomena are constituted through theoretical perspectives, this section specifies how the dissertation constructs its central phenomenon: the firm as a complex adaptive system (CAS). The aim is not to deny alternative constructions – such as the firm as production function, governance structure, or contractual nexus – but to explain why a complexity-informed construction is warranted for the organisational dynamics motivating the dissertation’s critique and reconstruction.

Defining the Constructed Phenomenon

Constructing “firms as CAS” means treating firms as systems composed of many interacting elements whose interactions generate adaptive behaviour and system-level properties irreducible to component-level descriptions. Holland’s canonical definition provides orientation: complex adaptive systems are “systems that involve many components that adapt or learn as they interact” (2006, p. 1). The claim is not that firms are merely complicated, but that they exhibit features characteristic of adaptive, interaction-driven systems.

CAS accounts converge on a set of family resemblances. Cilliers emphasises openness, far-from-equilibrium operation, and non-linear interaction (2005, p. 257). Holland highlights parallelism, conditional action, modularity, and evolution through recombination (2006, p. 1). The dissertation draws on these accounts as conceptual resources shaping what becomes salient when firm behaviour is approached through a CAS lens.

What the Construction Makes Salient

The CAS construction is not terminological novelty but selective visibility. It foregrounds organisational phenomena that mechanistic framings tend to background or treat as residual:

- **Emergent organisational properties:** system-level capabilities arise from interaction structures rather than individual intentions (Cilliers, 2005, p. 257).
- **Adaptive reconfiguration:** change involves reorganisation under uncertainty rather than movement between stable equilibria.
- **Relational and meso-level dynamics:** teams, routines, internal networks, and inter-firm relations become central explanatory sites.
- **Historical contingency and path dependence:** trajectories matter; not all relevant features are time-invariant.
- **Limits of prediction as an evaluative ideal:** explanatory success may rest on mechanisms and patterns rather than point forecasts.

- **Co-evolutionary dynamics:** firm and environment mutually shape through feedback rather than one-way “constraints.”

These features are not asserted as universal truths about all firms, but as properties that become theoretically legible once the firm is constructed as a complex adaptive phenomenon.

Internal Coherence of the Construction

A frequent misunderstanding is to treat CAS as a checklist of attributes. Phenomena construction requires more: the phenomenon is an integrated configuration in which properties hang together as an explanatory object. Emergence presupposes interaction structure; adaptation presupposes openness; and the epistemic stance appropriate to inquiry depends on assumptions about novelty and historical contingency.

This is why superficial invocations of “complexity” are insufficient. The dissertation’s construction specifies (i) what becomes visible under the CAS frame, (ii) what becomes less central or less plausible, and (iii) how this reframing reshapes the questions that firm theory can coherently ask.

The Role of Alternatives and “Otherness”

The CAS construction also performs a methodological function emphasised in hermeneutic and critical traditions: alternatives provoke awareness of tacit pre-understandings (Gadamer, 1960/1994, p. 299). As Feyerabend argued, critique requires “a set of alternative assumptions” as an external standard (1978, p. 31). The CAS construction is deployed in precisely this sense – not as a final truth claim, but as an alternative standpoint that renders mechanistic presuppositions visible.

Knowledge Limits as Constitutive

For firms construed as CAS, incompleteness of knowledge is not a defect awaiting better data but a constitutive feature of systems producing novelty through interaction. As Cilliers notes, there is no “framework for frameworks”; we choose frameworks and our knowledge remains provisional (2005, pp. 259, 264). This epistemic limit is built into the constructed phenomenon rather than treated as a temporary gap. Inquiry proceeds under acknowledged constraint – a stance aligned with the epistemic humility articulated in Section 2.2.

Use Within the Dissertation

The “firms as CAS” phenomenon is developed cumulatively across the dissertation’s substantive chapters. It is demonstrated through sustained theoretical work rather than asserted as definitional fiat. Empirical material is engaged not as a decisive test of foundational assumptions, but as a means of illustrating, stress-testing, and refining what the constructed phenomenon enables us to notice and explain.

The evaluative standard is therefore comparative: whether the CAS construction increases the intelligibility of organisational dynamics that remain persistently difficult to articulate under dominant framings.

Conclusion

This section has shown how the dissertation constructs “firms as complex adaptive systems” as a legitimate research phenomenon grounded in complexity-informed commitments. The construction does not deny the independent existence of firms; it denies that any single conceptualisation is given by reality itself. The contribution is thus not re-labelling, but disciplined reframing aimed at expanding what firm theory can coherently see, describe, and explain.

Conclusion of Section 2.4

Section 2.4 has established that research phenomena are constituted through theoretical perspectives and specified how the dissertation constructs its central phenomenon. Four findings carry forward.

First, phenomena construction is integral to theoretical work rather than a preliminary step. What counts as the phenomenon – what is taken to be noteworthy and requiring explanation – depends on the conceptual resources through which the world becomes visible. Observation is theory-laden: one does not see much more than one understands. This does not collapse into relativism. Perspectival realism maintains that theories function as interpretive vantage points accountable to reality, even as access is always mediated by perspective. Researchers bear responsibility for how they constitute their objects; framing is a substantive epistemic commitment.

Second, constructing “firms as complex adaptive systems” is not terminological novelty but selective visibility. The CAS construction foregrounds phenomena that mechanistic framings background: emergent properties arising from interaction structures, adaptive reconfiguration under uncertainty, relational and meso-level dynamics, historical contingency and path dependence, co-evolutionary dynamics, and the limits of prediction as an evaluative ideal. The construction is an integrated configuration in which properties hang together – emergence presupposes interaction structure, adaptation presupposes openness – rather than a checklist of attributes.

Third, the CAS construction performs a methodological function: it renders mechanistic presuppositions visible by providing an alternative standpoint from which tacit pre-understandings can be recognised as commitments rather than natural facts. Phenomena construction and problematization are complementary: problematization exposes what mechanistic constructions foreclose; phenomena construction develops what becomes visible under transformed assumptions. The evaluative standard is comparative: whether the CAS construction increases the intelligibility of organisational dynamics that remain difficult to articulate under dominant framings.

Fourth, the CAS construction has methodological consequences. Because complex adaptive systems generate novelty through interaction, equilibrium-based tools and prediction-oriented validation are of limited applicability. The appropriate methods are pattern-oriented understanding, mechanism tracing, and comparative process analysis – developed in subsequent chapters. For firms construed as CAS, knowledge limits are constitutive rather than temporary; inquiry proceeds under acknowledged constraint.

Section 2.5 clarifies why foundational propositions are not directly empirically verifiable and how theory precedes and enables empirical inquiry.

2.5 The Relationship to Empirical Research

Section 2.4 established that "firms as complex adaptive systems" is a constructed research phenomenon – not a neutral datum but a theoretical framing through which organisational dynamics become intelligible. This raises an unavoidable question: what is the relationship between such foundational theoretical work and empirical research?

This clarification is necessary because a common misreading would demand direct empirical verification of the dissertation's ontological and epistemological commitments. Such a demand applies evaluation standards appropriate to derived hypotheses, not to paradigm-level propositions that function as conditions of intelligibility for empirical work itself.

The argument proceeds through five subsections. Subsection 2.5.1 establishes why foundational propositions cannot be directly empirically verified. Subsection 2.5.2 shows why complex adaptive phenomena resist conventional empirical testing. Subsection 2.5.3 clarifies why assumptions cannot themselves be tested – they define what would count as a test. Subsection 2.5.4 demonstrates that theory precedes and enables empirical research rather than following from it. Subsection 2.5.5 explains why paradigm shifts are gestalt switches, not evidence accumulation.

Section 2.6 specifies the success criteria appropriate to pure theoretical work of this kind.

2.5.1 Why Foundational Propositions Are Not Directly Empirically Verifiable

Introduction

This section draws a methodological boundary essential for reading – and evaluating – the present dissertation correctly: foundational propositions (ontological, epistemological, and paradigm-level commitments) are not directly empirically verifiable in isolation. This is not an evasion of empirical accountability. It reflects a well-established position in the methodology and philosophy of science: empirical tests operate on *derived claims* (model-implied predictions, operational hypotheses), whereas foundational propositions function as *conditions of intelligibility* – the background commitments that make observation, operationalisation, and testing possible in the first place.

Accordingly, demanding direct "empirical verification" of the foundational propositions developed in Chapters 3–4 applies an evaluation standard appropriate to lower-level hypotheses, not to first-order commitments concerning what the firm is, what can be known about it under conditions of complexity, and what forms of inference are appropriate. The argument below clarifies this boundary using classic statements in economic methodology and general philosophy of science, and specifies what empirical relevance properly means for foundational work.

WHAT: Foundational Propositions and the Limits of Direct Testing

1) The Logical Status of Foundational Propositions

Foundational propositions are not predictions about specific events. They are structuring assumptions concerning entities, causation, openness, time, knowledge, and explanation. As such, they shape:

- what counts as a relevant observation,

- what is treated as signal versus noise,
- what can be operationalised, and
- what forms of generalisation are considered legitimate.

For this reason, classic methodologists insist that foundational assumptions cannot be subjected to the same kind of test as subordinate hypotheses. Frank Knight's claim – often read as provocative – is methodological rather than rhetorical: deep economic propositions are not “read off” from observation, yet they are not arbitrary; they articulate analytic commitments about social and mental reality (Knight, 1940).

Mises reaches a parallel conclusion by a different route. In the sciences of human action, experience presents itself as complex wholes rather than as controlled experiments isolating single causal elements; empirical material is therefore compatible with multiple coherent theoretical interpretations (Mises, 1949).

Key implication: empirical confrontation rarely tests a foundational proposition alone. It always tests a *package*: foundational commitments plus auxiliary assumptions, operationalisations, measurement conventions, and interpretive rules.

2) The Hierarchy of Hypotheses and “Testing at the Bottom”

Machlup's enduring contribution was to make explicit the hierarchical structure of theoretical systems. At the top are foundational and methodological commitments; below them lie intermediate modelling structures; only at the lowest levels do propositions become sufficiently close to observation to generate testable implications (Machlup, 1955).

Braithwaite expresses the same point more generally: as one ascends the hierarchy, concepts become increasingly theoretical and are connected to observation only through complex logical bridges; empirical testing therefore targets *lowest-level hypotheses* rather than foundational principles (Braithwaite, 1953).

Machlup's crucial clarification – worth retaining because it blocks a common misunderstanding – is that there is no method for isolating and independently verifying fundamental assumptions, and that this is true even in physics (Machlup, 1955). As Einstein repeatedly emphasised (via Einstein & Infeld), basic concepts are not uniquely determined by the world; they are “free creations” constrained only indirectly by empirical adequacy through the performance of the theoretical system as a whole.

Key implication: empirical testability is not an intrinsic property of individual statements, but a relational property of statements embedded within a theoretical system.

3) Mill's Distinction: Developing Principles vs Applying Them

John Stuart Mill provides one of the clearest formulations of the boundary at issue. In social science, he argues, attempting to induce general laws directly from “facts in the concrete” is futile; science proceeds by reasoning from assumptions, while empirical confrontation concerns the *application* of the system rather than direct verification of its foundations ((Mill, 1843, bk. IV, ch. ix, §6)

Mill is not anti-empirical. He explicitly endorses comparison between predicted implications and observed outcomes. What he rejects is the idea that foundational assumptions themselves are verified in isolation by empirical procedure.

Key implication: foundational work is assessed by coherence, explanatory organisation, and fruitfulness; empirical work primarily adjudicates applications, specifications, and derived implications.

4) Complexity, Underdetermination, and Interpretive Pluralism

For complex social phenomena, underdetermination is structural rather than accidental. As Mises emphasised, rival theoretical frameworks can often interpret the same events coherently because the framework itself determines what counts as salient evidence, causal structure, and explanation (Mises, 1949).

This does not imply “anything goes.” Rather, empirical material pressures paradigms indirectly through *comparative performance*: which framework organises anomalies without ad hoc repair, which illuminates mechanisms, which supports diagnosis and intervention, and which generates more progressive research programmes.

5) Popper, Conventions, and the Non-Empirical Elements of Science

Popper’s falsificationism is often invoked to demand universal testability. Yet even in Popper’s own account, science presupposes methodological conventions: what counts as an acceptable basic statement, when an observation report is treated as reliable, and how auxiliary assumptions are handled. Machlup shows that this is not inconsistency but logical necessity: the act of treating an observation as a “basic statement” depends on rules not themselves settled by that observation (Machlup, 1955).

Key implication: demands for direct empirical verification of foundational propositions misunderstand the layered structure of scientific reasoning.

6) Hayek: Pattern Prediction vs Point Prediction

For complex phenomena, epistemic ambition often shifts from predicting particular events to identifying pattern regularities under classes of conditions. Hayek’s distinction remains decisive: predicting that “a pattern of a certain kind will appear” can be significant and falsifiable even when point prediction is infeasible (Hayek, 1967).

Key implication: empirical relevance is not abandoned but re-specified. Mechanism and pattern adequacy replace point prediction as the universal benchmark.

WHY: Why This Matters for Evaluating This Dissertation

1) The Nature of the Claims in Chapters 3–4

The propositions developed in Chapters 3–4 are not of the form “Firm X will do Y at time Z.” They concern:

- what the firm is when conceptualised as a CAS-relevant object (ontology),
- what kinds of knowledge are feasible about such an object (epistemology), and

- what styles of inquiry are appropriate (methodology).

These are paradigm-level commitments. They are empirically assessable only indirectly, through the performance of theoretical systems built upon them, not through isolated verification of the foundations themselves.

2) What Counts as an Appropriate Empirical Relationship

A false binary must be rejected: either “verified by prediction tests” or “unaccountable.” The appropriate relationship is instead:

- foundational commitments evaluated by coherence, ontological plausibility, and capacity to organise phenomena without systematic blindness;
- derived propositions subjected to direct empirical confrontation;
- empirical material used for stress-testing – revealing misdescription, ad hoc repair, or explanatory failure.

This aligns with the classic methodological maxim (often attributed to Conant): frameworks are displaced not by “facts alone” but by better frameworks – those that organise facts with less distortion and greater intelligibility.

HOW: How This Dissertation Positions Empirical Work

Accordingly, empirical material in Chapters 3–4 functions primarily as:

1. **Illustrative and diagnostic** – revealing what mechanistic framings systematically suppress; and
2. **Comparatively discriminating** – showing how a CAS-informed construction renders salient regularities and failure modes more intelligible.

The dissertation does not present itself as a predictive testing programme. The appropriate evaluation criterion is whether the CAS-informed foundations:

- reduce explanatory distortions generated by mechanistic assumptions,
- make previously residual phenomena (emergence, path dependence, adaptation, non-ergodicity) theoretically central, and
- enable more progressive downstream inquiry (mid-range mechanisms, comparative cases, meso-level tracing).

Conclusion

Foundational propositions are not verified in the manner of subordinate hypotheses. This is not special pleading; it reflects the hierarchical structure of scientific reasoning and the mediated relationship between theory and observation. Empirical confrontation remains essential – but it operates most directly at the level of derived implications and more indirectly through

comparative theoretical performance, especially for complex phenomena where point prediction is neither feasible nor epistemically appropriate.

Section 2.5.2 can now extend this argument by showing why complex adaptive phenomena intensify underdetermination, strengthen the case for pattern and mechanism standards, and further constrain conventional verification ideals.

2.5.2 Complex Phenomena Resist Conventional Empirical Testing

Introduction

Section 2.5.1 established that foundational propositions cannot be empirically verified in isolation, even in principle, because they function as framing commitments rather than testable event-claims. This section extends the argument in a domain-specific way: even when inquiry moves closer to empirical engagement, *complex adaptive phenomena resist the core preconditions of conventional empirical testing*.

The issue is not that complexity merely makes empirical work more difficult. It is a categorical mismatch between:

- (a) what conventional empirical testing presupposes – stable distributions, isolable variables, repeatability, stationarity, and observer independence – and
- (b) what complex adaptive systems exhibit – emergence, non-linearity, historical contingency, openness, reflexivity, and evolving possibility spaces.

As Hayek observed, theories of complex phenomena are necessarily more limited in what they can deliver than theories of simple phenomena (Hayek, 1967, p. 22). This limitation is inherent to the nature of the object, not a contingent shortcoming of method.

This section proceeds in four parts:

- (A) CAS properties and the failure of standard testing assumptions;
- (B) complex causality and the collapse of variable-isolation logic;
- (C) the repeatability problem; and
- (D) ergodicity violations and their implications for inference.

Part A: CAS Properties and Testing Failure

WHAT: How CAS Properties Defeat Conventional Testing Assumptions

Conventional empirical testing typically presupposes some combination of:

- (i) stable underlying regularities,
- (ii) separable causal contributions,
- (iii) repeatable conditions,
- (iv) independence between observer and observed, and

(v) generalisable relationships across contexts.

Core properties of complex adaptive systems systematically violate these presuppositions.

Table 1: How Complex Adaptive System Properties Undermine Conventional Statistical Testing: Mechanisms and Methodological Implications

CAS property	Why conventional testing fails	Methodological implication
Non-ergodicity	Past data need not represent future distributions	Time-series inference becomes structurally fragile
Emergence	System-level effects are not reducible to micro variables	Reductionist testing misses the explanandum
Fundamental uncertainty	Possibility spaces evolve; outcomes not specifiable ex ante	Hypothesis spaces are incomplete by construction
Context-dependence	Regularities are local and conditional	Universal generalisation is unreliable
Reflexivity	Observation and modelling can alter system behaviour	Tests may reconfigure their targets

Non-ergodicity. In non-ergodic environments, the data-generating process itself changes over time; time averages and ensemble averages diverge. As a result, the inferential move from past frequencies to future probabilities loses its grounding (Davidson, 2012).

Emergence. Where the explanandum is a system-level property generated by interaction structure, testing strategies that abstract away those interactions target the wrong object. As Cilliers emphasises, complex systems exhibit behaviour that arises from interactions rather than from component properties taken separately (Cilliers, 2005).

Fundamental uncertainty. Knight's distinction between measurable risk and unmeasurable uncertainty is decisive here (Knight, 1921). In adaptive settings characterised by novelty and structural change, not all relevant contingencies can be enumerated in advance. Conventional hypothesis testing presupposes a stable outcome space; under fundamental uncertainty, that presupposition fails at the outset.

Context-dependence. In open systems, causal relations are conditional on institutional, technological, cultural, and network configurations. Even when regularities appear, they are often non-law-like and non-transportable across contexts (Lawson, 1997).

Reflexivity. In reflexive domains, participants' interpretations influence the phenomena being studied. Soros (2013) emphasises this feedback between beliefs and outcomes, while MacKenzie demonstrates how models can be performative, reshaping markets once adopted. Observation is therefore not neutral recording but part of the causal field.

WHY: Why These Obstacles Are Categorical, Not Merely Practical

These obstacles do not arise from insufficient technique or data. They follow from the nature of the object. If a system is open, adaptive, historically constituted, and interaction-driven, then the empirical ideal of isolating stable regularities under repeatable conditions is not simply difficult – it is often mis-specified.

Cilliers argues that complete knowledge of a complex system would require grasping the system together with its full environment, which is not humanly achievable (Cilliers, 2005). Hayek adds that the decisive difficulty lies not merely in the number of variables, but in identifying which facts are relevant at all – often an insurmountable problem in complex social systems (Hayek, 1967).

HOW: Implications for Empirical Orientation

The appropriate inference is not that empirical research is irrelevant, but that *conventional verification standards* – developed for stable, repeatable domains – are not the correct evaluative baseline for complex adaptive phenomena. Empirical contribution instead takes the form of pattern identification, mechanism tracing, configurational comparison, and historically informed analysis, rather than invariant, context-free, point-predictive tests.

Part B: Complex Causality

WHAT: Why Causal Complexity Defeats Variable-Isolation Logic

Complex adaptive systems display causal structures that undermine standard causal inference designs premised on separability:

Table 2: Structural Features of Complex Economic Systems and Their Consequences for Causal Inference and Hypothesis Testing

Feature	Implication for testing	Consequence
Multiple interacting causes	“Holding all else constant” becomes unrealistic	Controlled isolation is often impossible
Latent variables	Confounding is endemic	Attribution becomes fragile
Feedback loops	Cause and effect co-determine	Unidirectional causal claims lose clarity
Time delays	Effects detach from proximate causes	Temporal correlation misleads
Non-linearity	Small changes yield large effects	Linear approximations distort

Even sophisticated causal frameworks rely on identifying confounders and specifying stable causal structure. In CAS domains, these conditions are frequently unmet: relevant variables are unobserved, evolving, or structurally indeterminate (Pearl, 2000; Woodward, 2003).

Feedback intensifies the problem. Where reinforcing and balancing loops operate, causation becomes systemic and recursive rather than linear. As Sterman notes, feedback implies mutual determination across time (Sterman, 2000). Testing designs that assume one-way causation risk mistaking loops for lines.

WHY: Why “Physics-Style” Control Is Misplaced

Attempts to isolate single causal effects under such conditions risk producing an illusion of precision at the cost of ontological misfit. Hayek’s warning against the “pretence of knowledge” is directly relevant: claiming control-knowledge where none is possible leads to systematic error (Hayek, 1974).

HOW: What Replaces Variable Isolation

The more appropriate empirical stance is mechanism illumination under configurational complexity: tracing how interacting conditions generate patterns, how feedback cascades operate, and how small perturbations propagate through relational structures.

Part C: The Repeatability Problem

WHAT: Why CAS Cannot Be Reproduced in the Relevant Sense

Conventional testing presupposes repeatability: similar conditions should yield similar outcomes. CAS resist this requirement in three structural ways:

1. Systems are historically constituted and path-dependent.
2. Environments co-evolve rather than remain fixed backdrops.
3. Initial conditions are unrecoverable, while sensitivity implies divergence.

Even defining “the same conditions” is problematic, since boundaries and relevant variables are partly observer-constructed (Cilliers, 2005). The outcome is not epistemic nihilism but a different epistemic object: historically situated, pattern-oriented, mechanism-seeking knowledge.

WHY: Why the Gold Standard Cannot Be Imported Wholesale

Repeatability is not a universal virtue; it is a property of certain domains. Where it fails categorically, the appropriate methodological response is not to insist upon it, but to adopt standards suited to singular, path-dependent phenomena.

HOW: What Replaces Repeatability

In complex domains, robustness arises from cross-case pattern recurrence, mechanism plausibility, and configurational comparison rather than laboratory replication.

Part D: Ergodicity Violations and the Limits of Statistical Inference

WHAT: Why CAS Violate Ergodic Assumptions

A system is ergodic when time averages converge to ensemble averages and statistical structure is stable. CAS violate these conditions through non-stationarity, path-dependence, and structural transformation of the generative process itself.

In non-ergodic environments, historical data do not license the inference that the future will be drawn from the same distribution as the past (Davidson, 2012). This is why structurally transforming economies and organisations systematically defeat forecast-from-history approaches.

WHY: Why More Data Does Not Solve the Problem

The failure is not sample size but distribution drift and regime change. Where the underlying process evolves, the stability required by many inferential routines is absent.

Hayek adds that statistical methods often treat elements as if they were independent, thereby eliminating precisely the connectedness that generates complex behaviour (Hayek, 1967). Statistics remains descriptively useful but unreliable as a universal verification device for causal claims about interaction-driven systems.

HOW: From Point Prediction to Pattern Expectations

The appropriate epistemic ambition therefore shifts toward pattern-level expectations: identifying typical dynamics, generic failure modes, and structural tendencies under broad conditions. Hayek's distinction between point prediction and pattern prediction captures this reorientation (Hayek, 1967).

Conclusion

Complex adaptive phenomena resist conventional empirical testing because they violate the very preconditions such testing presupposes: ergodicity, isolable causal variables, repeatability, stationarity, and observer independence. These are not practical inconveniences but categorical mismatches between method and object.

The implication is not that empirical research is irrelevant, but that empirical relevance must be specified differently – through pattern identification, mechanism illumination, configurational comparison, and historically informed analysis rather than through expectations of invariant, point-predictive verification.

Section 2.5.3 turns from the nature of complex phenomena to the status of assumptions themselves, clarifying why paradigm-defining commitments are not objects of direct empirical testing but conditions of empirical intelligibility and application.

2.5.3 Assumptions Cannot Be Empirically Tested

Introduction

Sections 2.5.1 and 2.5.2 established two limits: (i) foundational propositions cannot be empirically verified in isolation, and (ii) complex adaptive phenomena violate the preconditions of conventional empirical testing. This section presses the point to its logical core: **assumptions themselves are not the kinds of propositions that empirical tests can adjudicate**, because assumptions function as **pre-empirical commitments** that structure what counts as an object, a variable, evidence, and even a “test.”

Understanding this prevents a category error. To demand empirical tests of a framework's ontological, epistemological, or methodological assumptions is to ask empirical procedures to do work **logically prior** to the empirical procedures themselves.

WHAT: The Nature and Function of Assumptions

What assumptions are

Assumptions are framing commitments that define a theory's domain and its admissible moves. Bacharach states this directly: assumptions “state the domain of the theory... [and] form the boundaries of the theory” (1989, p. 498). They are not simply another class of hypotheses; they are the background architecture that makes hypothesis-generation possible.

Bacharach's decisive methodological distinction follows: **“Assumptions are not testable. If they could be tested, they would be hypotheses”** (1989, p. 498). This is not a semantic trick. It marks a difference in logical role:

- **Assumptions** define admissible entities, relations, and standards of relevance.
- **Hypotheses** are derived claims within that defined space, assessed against evidence.

Machlup's vocabulary captures this functional status: fundamental commitments operate as “heuristic principles,” “basic postulates,” “useful fictions,” and “procedural rules” (1955, p. 9) – terms that emphasise what assumptions *do* (guide inquiry) rather than what they *report* as freestanding, testable claims.

What assumptions do

Assumptions perform four interconnected functions:

1. Define the domain of inquiry

They specify what sorts of entities exist for the theory (e.g., firm-as-production-function versus firm-as-adaptive-system) and what relations are admissible.

2. Specify relevance and evidential salience

They shape what counts as “signal” versus “noise,” what is measurable, what is explanatory, and what qualifies as anomalous.

3. Enable hypothesis formation

A hypothesis is intelligible only relative to a background of assumptions (about agency, information, time, causality, boundaries, and the like).

4. Delimit generalizability

Conclusions hold within the “world” carved out by the assumptions; outside it, claims may not merely fail – they may cease to be meaningful.

Why “realism of assumptions” is the wrong test

Friedman is often misread as licensing indifference to assumptions. His methodological point is more compatible with the non-testability claim: the issue is not whether assumptions are descriptively “realistic,” but whether they are useful for the purpose at hand (1953). In this register, assumptions are evaluated by what they enable, not treated as freestanding empirical propositions to be confirmed or rejected directly.

WHY: The Circularity Problem in Testing Assumptions

Testing presupposes the framework to be tested

Empirical tests do not occur in a vacuum. They require choices about what counts as data, what counts as measurement, which operationalisations are legitimate, and what failures count as refutation. Those choices are not delivered by data; they are supplied by assumptions.

Kuhn's account of normal science captures this internality: testing is conducted **within paradigms**, not from a neutral standpoint outside them (1962). Lakatos makes the same point in methodological terms: a research programme requires a "hard core" protected by methodological decision – without some insulated commitments, no coherent programme can develop (Lakatos, 1970).

The Duhem – Quine point

This circularity is formalised by the Duhem–Quine thesis: every empirical "test" is a joint test of (i) the focal hypothesis plus (ii) auxiliary assumptions, measurement assumptions, and background commitments. Quine stresses the implication: when predictions fail, the failure does not uniquely identify what to revise (Quine, 1951). Assumptions therefore cannot be isolated and tested in the manner of derived hypotheses.

A clarification for quantitatively-trained readers

The circularity argument applies with equal force in econometric practice. Identification strategies, model specifications, error structures, stationarity assumptions, and variable selections are not *results* of empirical tests; they are *preconditions* that make testing possible. A regression does not empirically validate the assumption that the data-generating process is linear with additive errors; it operates *under* that assumption to estimate conditional relationships.

If the assumption set changes, what counts as "fit," "bias," "consistency," or "significance" changes with it. This is why assumptions cannot be "empirically tested" in the manner of derived hypotheses: they configure the inferential machinery itself. The point is not that econometrics is flawed, but that it operates within – not outside of – the framework established by background assumptions.

Assumptions define refutation conditions

Assumptions do not merely accompany tests; they partly define what would count as refutation in the first place. The same observation can be:

- an **anomaly** under one framework,
- an **expected baseline** under another.

Hence Kuhn's point that abandoning one paradigm is simultaneously adopting another, and that judgment proceeds comparatively – "with nature and with each other" (Kuhn, 1962, p. 77). There is no framework-free tribunal.

HOW: Evaluating Assumptions by Their Fruits

The alternative to “testing”: comparative evaluation

If assumptions cannot be directly empirically tested, they are evaluated **indirectly and comparatively**, through the performance of the theoretical systems they constitute.

Machlup expresses this logic explicitly: fundamental assumptions remain accepted so long as they have heuristic value, and they are displaced when alternatives serve explanatory purposes more successfully (Machlup, 1955, p. 16). Lakatos makes the comparative structure systematic: programmes are assessed as progressive or degenerating, and “falsification” in practice is inseparable from the availability of a better alternative (Lakatos, 1970). Conant summarises the practical logic sharply: a theory is overthrown by a better theory, not merely by contradictory facts (Conant, 1947).

What “by their fruits” means in practice

Assumptions are assessed via the frameworks they enable, using criteria such as:

- **internal coherence** (no hidden contradictions across commitments),
- **explanatory reach** (what phenomena become intelligible),
- **problem-solving power** (what puzzles become tractable),
- **heuristic fertility** (what questions become productive),
- **fit-for-purpose adequacy** (whether the framework serves the inquiry’s aims better than rivals).

These are not “soft” standards. They are the standards appropriate when the object of evaluation is a framework rather than a single derived claim.

Application to this dissertation

In this dissertation’s methodological genre – foundational reconstruction – the ontological, epistemological, and methodological commitments are not offered as isolated empirical propositions. They function as a reconstruction of the conceptual architecture through which the firm is rendered intelligible under a complexity-informed lens.

Accordingly, the evaluative question is not: **“Can these assumptions be empirically tested directly?”**. Rather, it is: ***“Does the framework constituted by these assumptions illuminate, organise, and explain the firm-phenomenon more fruitfully than the mechanistic alternatives it problematizes?”***

Conclusion

Assumptions are pre-empirical commitments that constitute the conditions of empirical intelligibility: they define domains, shape relevance, enable hypothesis-formation, and establish what would count as evidence or refutation. Because empirical testing presupposes

these commitments, assumptions cannot be empirically tested without circularity. This is not a limitation of method but a feature of what assumptions are.

The appropriate response is therefore to evaluate assumptions by the fruits of the frameworks they generate – their coherence, explanatory power, and heuristic fertility – through comparative judgment rather than direct empirical “verification.”

Section 2.5.4 extends the argument by showing why theory precedes and enables empirical research: evidence is not self-interpreting, and empirical inquiry depends on prior conceptual framing.

2.5.4 Theory Precedes and Enables Empirical Research

Introduction

The preceding sections established three methodological limits:

- (i) foundational propositions cannot be empirically verified in isolation;
- (ii) complex adaptive phenomena resist the preconditions of conventional empirical testing; and
- (iii) assumptions cannot themselves be empirically tested because they function as pre-empirical commitments.

This section completes the argument by clarifying a reversal that naïve empiricism routinely overlooks: **theory precedes and enables empirical research, not the other way around.**

This clarification reframes how the present dissertation should be evaluated. The work does not “fail” to provide empirical verification; rather, it undertakes the logically prior task of articulating the conceptual and ontological scaffolding without which empirical investigation cannot be coherently formulated, interpreted, or assessed. In this sense, pure theory is not an inferior substitute for empirical work, but a **precondition of intelligible empirical work**, especially in domains characterised by complexity, openness, and historical contingency.

WHAT: The Theory–Observation Relationship

Theory-ladenness of observation

A foundational insight of twentieth-century philosophy of science is that observation is theory-laden: what is seen is shaped by the conceptual resources that make seeing meaningful. Norwood Russell Hanson captured this succinctly: “*there is more to seeing than meets the eyeball*” (1958). Observation is not raw reception later interpreted; it is already structured by prior categories and expectations.

Frank H. Knight expressed the same insight in methodological terms: we do not apprehend “sense data” first and meaning second; rather, what is immediately apprehended is meaning, through which “facts” are subsequently reconstructed (1940, p. 17). Perception is inseparable from understanding.

Thomas Kuhn generalised the implication: what a researcher sees depends both on what is present and on the conceptual training that has taught them what to attend to (1962, p. 113).

Paradigms therefore do not merely generate different interpretations of the same facts; they shape what is experienced as salient, problematic, or even visible.

Karl Popper articulated the methodological corollary bluntly: observation is always selective, presupposing a problem, an interest, a point of view, and a task (1963). Without such structuring commitments, observation collapses into undifferentiated experience.

This logic is often summarised in the remark attributed to Albert Einstein: it is theory that decides what we can observe. Regardless of phrasing, the methodological point holds: theory does not follow observation in a simple inductive sequence; it constitutes the conditions under which observation becomes intelligible as data.

Facts do not speak for themselves

If observation is theory-laden, then facts are not self-interpreting. Data require a framework that distinguishes:

- relevant from irrelevant,
- signal from noise,
- anomaly from expected variation,
- mechanism from coincidence.

Kuhn expressed this sharply: without a guiding paradigm (or candidate paradigm), the facts pertaining to a field may appear equally relevant, producing accumulation without understanding (1962, p. 15).

Friedrich Hayek made the same point for complex phenomena: even intensive attention to “bare facts” will not render them intelligible unless one already knows what to look for (1967, p. 22). Theory is therefore not an optional overlay on empirical work; it is the searchlight that makes empirical inquiry possible.

Theory construction as a creative act

If theory is not mechanically derived from observation, its origin must be understood differently. Milton Friedman described theory construction as a creative act – an exercise of intuition, imagination, and conceptual re-organisation rather than straightforward induction from data (1953). This does not imply arbitrariness. It implies that theory must first be articulated before it can be assessed by any evaluative criteria.

WHY: Why Empirical Work Requires Prior Theory

The logical priority of theory

Theory is logically prior to empirical investigation in at least three interrelated senses.

1. Theory determines what questions are meaningful

Every empirical question presupposes a conception of the object. Asking how firms “optimize” presupposes one ontology; asking how firms “adapt” presupposes another. Questions are not delivered by data; they are enabled by conceptual frames.

2. Theory determines what counts as evidence

What is worth observing, what is measurable, and what qualifies as an explanandum depend on theoretical commitments. Under one framework, certain indicators are central; under another, the same indicators may be peripheral or misleading.

3. Theory determines what methods are appropriate

Methods embed assumptions about stability, separability, repeatability, and causal form. Especially for complex phenomena, methodological appropriateness is conditional on the underlying ontology and epistemology of the object under investigation.

Pure theory as an enabling condition

This priority is explicit in the methodological tradition mobilised here. Read through Knight, Per Bylund emphasises that in economics, pure theory precedes and informs applied inquiry by providing the conceptual apparatus without which empirical work cannot be coherently framed (Bylund, 2021).

Similarly, Ludwig von Mises insisted that theory is not a rival to empirical understanding but the mental tool that renders empirical reality intelligible as more than a collection of disconnected observations (1949). The relationship is enabling: theory does not replace empirical work; it makes empirical work possible as inquiry rather than inventory.

This enabling role clarifies why demands for empirical verification of foundational theory constitute a category error. Pure theory is not a failed empirical study; it is an attempt to establish the conditions under which empirical studies can be meaningfully designed, interpreted, and compared.

HOW: How This Dissertation Enables Empirical Work

The conceptual contribution

Within the present design, the dissertation’s contribution lies in reconstructing the conceptual and ontological framework through which the firm becomes intelligible. It enables future empirical work by:

- making new kinds of questions formulable,
- clarifying what would count as relevant observations under a complexity-informed conception, and

- rendering visible organisational dynamics that prior constructions systematically obscured.

This is not an empirical shortcut. It is the logically prior work required if empirical inquiry is not to remain trapped within the very assumptions under critique.

Reframing evaluation

Accordingly, the appropriate evaluative question is not: **“Does this dissertation provide empirical verification?”** It rather is:

“Does it supply a coherent, defensible framework that enables empirically meaningful and theoretically progressive inquiry into firms under conditions of complexity?”

The criteria relevant to that question – coherence, explanatory reach, heuristic fertility, and comparative adequacy – are those appropriate to foundational theoretical work.

Conclusion

This section has shown that theory precedes and enables empirical research. Observation is theory-laden; facts do not speak for themselves; and theory construction is a creative act that cannot be reduced to induction from data. Theory provides the conceptual apparatus that makes empirical questions formulable, evidence identifiable, and methods interpretable.

The dissertation’s contribution is therefore properly understood not as a failure to provide empirical verification, but as the successful performance of the prior task: **constructing the framework within which empirical research can proceed meaningfully.**

2.5.5 Paradigm Shifts Are Gestalt Switches, Not Evidence Accumulation

Introduction: How Paradigms Actually Change

The preceding sections established that foundational propositions resist empirical verification, that complex adaptive phenomena defeat conventional testing, that assumptions cannot themselves be empirically tested, and that theory precedes and enables empirical work. This section completes the argument by clarifying **how paradigm change actually occurs** – not through the steady accumulation of decisive evidence, but through **gestalt-like reconfigurations** that reorganize how the world is seen.

Understanding paradigm change clarifies what this dissertation attempts and how its contribution should be evaluated. The aim is not to empirically “refute” mechanistic economics in a paradigm-neutral way – such refutation is impossible for paradigmatic commitments. Rather, the aim is to **enable a gestalt switch**: a transformation in how firms are perceived, conceptualized, and theorized. This is how paradigms actually change.

Part A: WHAT Is the Difference Between Normal and Revolutionary Science?

Thomas Kuhn’s distinction

Kuhn’s distinction between normal science and revolutionary science provides the essential frame. Normal science operates within an established paradigm, solving puzzles using accepted

concepts, methods, and standards. As Kuhn puts it, normal science is “predicated on the assumption that the scientific community knows what the world is like” (1962, p. 5). Scientific work proceeds by extending and refining the paradigm rather than questioning its foundations.

Revolutionary science arises when persistent anomalies – problems resistant to solution within the paradigm – accumulate and generate crisis. The response is not incremental adjustment but a shift to a new framework that reorganizes what counts as a problem, what counts as a solution, and what the world “looks like.” Kuhn emphasizes that such transitions are not compelled by logic plus neutral experience, but resemble **gestalt switches**: reorganizations of perception and meaning (1962, p. 150).

Imre Lakatos’s reconciliation: research programmes as rational reconstruction

Lakatos sought to reconcile Kuhn’s insight with scientific rationality. Paradigm change need not be treated as irrational conversion; it is rationally reconstructible as competition between research programmes with distinct “hard cores” and evolving “protective belts.”

Two implications are decisive:

1. **Paradigm-level assessment is holistic.** Whole systems of assumptions, models, and standards are evaluated together, not isolated propositions.
2. **Naïve falsificationism fails at the paradigm level.** Anomalies do not automatically overthrow a framework. What matters is whether an alternative programme becomes comparatively progressive – explaining more, resolving more problems, and opening more fruitful lines of inquiry.

The gestalt-switch metaphor

The gestalt metaphor captures the nature of paradigm change. In classic perceptual switches (duck–rabbit, Necker cube), nothing in the stimulus changes; what changes is the organization of perception. Paradigm shifts operate similarly: the same empirical world is reorganized through different categories, salencies, and standards of relevance.

Kuhn’s claim is literal, not rhetorical: a new paradigm can reorganize scientific experience so deeply that practitioners experience themselves as working in a different world – not because reality changed, but because the **world-as-seen-through-inquiry** changed (1962).

Why evidence cannot force the switch

Evidence alone cannot compel paradigm change because evidence is always interpreted through paradigmatic lenses. What counts as evidence, anomaly, or explanation is itself paradigm-dependent. Kuhn therefore concludes that paradigm competition is not resolved by proof in the ordinary sense (1962, p. 148).

Lakatos sharpens this point: there is no falsification prior to the emergence of a superior alternative. A paradigm is displaced only when another framework demonstrates greater overall problem-solving power.

Part B: Incommensurability and Its Implications

What incommensurability means

Incommensurability refers to the absence of a fully neutral standard for comparing paradigms. Competing paradigms often disagree not only about answers, but about which questions matter, what constitutes explanation, and what standards are legitimate. Shared terms may persist, but their meanings are embedded in different conceptual networks.

Feyerabend pushes the point further: there is no neutral algorithm for theory choice that forces all competent inquirers to converge. Paradigm choice involves judgment, persuasion, and comparative vision.

No neutral algorithm

This does not render paradigm choice arbitrary. It clarifies that demands for paradigm-neutral empirical proof misunderstand how paradigms function. Empirical proof operates within frameworks that already specify categories, operationalizations, and relevance criteria. Between frameworks, evaluation must therefore be **comparative and pragmatic**: assessing what each framework enables inquiry to see, explain, and do.

Rorty articulates this as a pragmatic stance: comprehensive frameworks are justified by their practices and consequences, not by correspondence to an imagined framework-independent tribunal.

Implications for this dissertation

Accordingly, the dissertation cannot “prove” the superiority of a complexity-oriented reconceptualization through neutral tests. What it can do is:

- render mechanistic commitments explicit so their limits become visible,
- articulate a coherent alternative framework, and
- demonstrate comparative advantages in explanatory reach and conceptual fit.

This is precisely how paradigm change proceeds: not through compelled assent, but through demonstrated fruitfulness that invites conversion.

Part C: WHY This Matters for This Dissertation

Paradigm-level intervention

The dissertation operates at the level of foundational commitments. It does not aim to patch isolated auxiliary assumptions, but to reconstruct the framework through which firms are conceptualized. Machlup’s methodological point is decisive here: theories are not overthrown by facts alone, because facts are not self-interpreting outside a framework.

The task, therefore, is to develop a superior alternative – one that accommodates phenomena systematically obscured by the mechanistic paradigm and resolves anomalies it recurrently produces.

Problem-solving as the evaluation criterion

Larry Laudan provides the appropriate evaluative frame: scientific progress is measured by solved problems – empirical and conceptual. Paradigms are compared by problem-solving capacity: what they illuminate, what they render coherent, and what puzzles they dissolve.

For firm theory, the relevant question is not which paradigm can be conclusively verified, but which paradigm better enables intelligible inquiry into firms under conditions of complexity.

Cilliers on responsible modesty

From a complexity perspective, Paul Cilliers adds an ethical-epistemic dimension: acknowledging limits to algorithmic certainty is not relativism but responsibility. A framework that respects the epistemic constraints imposed by complex phenomena may be more scientifically appropriate than one that promises precision the object cannot support.

Part D: HOW This Applies to the Dissertation's Alternative Assumptions

Comparative advantages of the proposed alternatives

The alternative framework developed in the dissertation is motivated by concrete theoretical problems rather than aesthetic preference. It is designed to make intelligible:

- increasing environmental complexity,
- pervasive fundamental uncertainty, and
- continuous adaptive transformation as a normal organizational condition.

The alternatives are selected because they enable coherent theorization of these dynamics without systematic distortion.

Not arbitrary, but problem-driven

The relevant evaluative question is therefore whether the alternative assumptions allow inquiry to handle these problems more effectively than the mechanistic construction – not whether they can be “verified” in isolation.

Part E: Enabling the Gestalt Switch

What the dissertation provides

The dissertation aims to enable a gestalt switch by supplying:

- a clarified map of mechanistic commitments and their locking effects, and
- an alternative ontological vocabulary that reorganizes what “the firm” is taken to be.

It does not seek forced conversion. It provides the resources for **re-seeing**.

The reader's role

Gestalt switches cannot be compelled; they require participation. As Kuhn emphasizes, paradigm change is not imposed by proof alone but occurs when an alternative becomes compelling because it organizes inquiry more powerfully.

This is not a weakness of paradigm-level work. It is its defining characteristic.

An Illustrative Gestalt Moment: Kodak and Nokia

The contrast between Eastman Kodak and Nokia provides concrete illustration of paradigm change as a gestalt switch rather than evidence accumulation. In both cases, extensive empirical evidence was available – market data, technological indicators, consumer trends – yet this evidence did not compel strategic transformation.

Within a mechanistic paradigm, Kodak interpreted digital imaging as a marginal efficiency problem within an existing product architecture, while Nokia interpreted smartphones as incremental variations on a stable handset-optimisation model. The empirical signals were visible, but their *meaning* was not.

Only when firms (and analysts) reconceptualise industries as complex adaptive systems – characterised by path-dependence, endogenous innovation, ecosystem co-evolution, and non-linear feedback – do the same facts reorganise into a different pattern: not missed forecasts, but paradigm-bound blindness.

Nothing in the data changed; what changed was the framework through which the data were seen. This is precisely the structure of a gestalt switch. Demanding that Kodak or Nokia have been “empirically proven wrong” before changing course misunderstands the nature of strategic failure in complex systems. The failure was not evidential but paradigmatic – and the correction required a shift in how the phenomenon itself was constructed.

Conclusion: Paradigm Change as the Dissertation's Mode of Contribution

Hayek's insight is decisive: as inquiry advances into essential complexity, the degree of classical falsifiability may decline – not because rigor weakens, but because the object changes. Methods suited to simple phenomena cannot be demanded as universal standards.

Paradigm shifts are gestalt switches, not evidence accumulation. They occur when an alternative framework reorganizes what can be seen, explained, and investigated, and demonstrates superior problem-solving capacity as a whole.

This section completes the argument of Section 2.5: conventional demands for empirical verification are categorically inappropriate for paradigm-level theoretical reconstruction. What matters instead is whether an alternative framework enables more coherent, responsible, and explanatorily powerful inquiry.

Conclusion of Section 2.5

Section 2.5 has clarified the relationship between the dissertation's foundational theoretical work and empirical research. Five findings carry forward.

First, foundational propositions – ontological, epistemological, and paradigm-level commitments – cannot be empirically verified in isolation. They function as conditions of intelligibility that shape what counts as observation, evidence, and anomaly. Empirical testing operates at the level of derived hypotheses; foundational commitments are assessed through coherence, explanatory organisation, and the performance of theoretical systems built upon them.

Second, complex adaptive phenomena resist conventional empirical testing because they violate its preconditions: stable distributions, isolable variables, repeatability, and observer independence. CAS exhibit emergence, non-linearity, historical contingency, non-ergodicity, and reflexivity. Empirical relevance must therefore be re-specified through pattern identification, mechanism tracing, and configurational comparison.

Third, assumptions cannot themselves be empirically tested because they define the domain, specify relevance, enable hypothesis formation, and establish what would count as refutation. Testing presupposes the framework being tested. Assumptions are evaluated by their fruits – the coherence, explanatory reach, and problem-solving power of the frameworks they constitute.

Fourth, theory precedes and enables empirical research. Observation is theory-laden; facts do not speak for themselves. Pure theory establishes the conceptual apparatus that makes empirical questions formulable, evidence identifiable, and methods appropriate.

Fifth, paradigm shifts are gestalt switches, not evidence accumulation. Evidence cannot force paradigm change because evidence is interpreted through paradigmatic lenses. Paradigms are displaced when alternative frameworks demonstrate superior problem-solving capacity. The dissertation aims to enable such a switch by clarifying mechanistic commitments and providing an alternative vocabulary for re-seeing.

Section 2.6 specifies the success criteria appropriate to pure theoretical work, distinguishing criteria the dissertation should satisfy from those categorically inappropriate to its genre.

2.6 Success Criteria for Pure Theory

Section 2.5 established that conventional empirical testing is categorically inappropriate for evaluating foundational theoretical propositions. Foundational commitments function as conditions of intelligibility; they cannot be verified in isolation. Complex phenomena resist testing preconditions. Assumptions define what would count as a test. Theory precedes and enables empirical research. Paradigm shifts are gestalt switches, not evidence accumulation. This raises the essential question: by what criteria should pure theoretical work be evaluated?

This clarification is necessary because the dissertation's contribution – diagnosing a mechanistic assumption structure and reconstructing the firm as a complexity-consistent phenomenon – requires explicit evaluative standards. Without such standards, the work risks being judged by criteria appropriate to different genres of inquiry, producing category errors in assessment.

The argument proceeds through three subsections. Subsection 2.6.1 articulates six criteria appropriate for pure theory: internal coherence, problem-solving capacity, theoretical fruitfulness, illumination of mechanisms, interestingness, and comprehensiveness. Subsection 2.6.2 identifies four criteria that are inappropriate: predictive accuracy, empirical verification, naïve falsifiability, and quantitative precision. Subsection 2.6.3 applies Lakatos's methodology of scientific research programmes to assess whether the dissertation constitutes a progressive programme.

Section 2.7 specifies the research design and architecture through which these criteria are operationalized.

2.6.1 Criteria Appropriate for Pure Theory

Introduction

Section 2.5 established that conventional empirical testing is categorically inappropriate for evaluating foundational theoretical propositions. Foundational propositions cannot be empirically verified, complex phenomena resist conventional testing, assumptions function as pre-empirical commitments, theory precedes and enables empirical research, and paradigm shifts occur through gestalt-like reconfigurations rather than evidence accumulation. This raises the essential question: by what criteria should pure theoretical work be evaluated?

This section articulates six criteria appropriate for assessing foundational theory: internal coherence, problem-solving capacity, theoretical fruitfulness, illumination of mechanisms, interestingness, and comprehensiveness. Each criterion is explained (WHAT), justified for this dissertation (WHY), and operationalized through its application to the dissertation's design (HOW).

These criteria are not arbitrary stipulations. They derive from the methodological literature on theory evaluation and reflect the fact that pure theory serves different epistemic functions than empirical research. As Weick summarizes, "good theory explains, predicts, and delights" (1989: 525) – though for complex phenomena where point prediction is neither feasible nor appropriate, explanation, illumination, and pattern-level intelligibility become the primary functions.

The criteria developed here provide the evaluative framework that readers and examiners should apply to the dissertation's core theoretical contribution, which diagnoses a mechanistic assumption structure and reconstructs the firm as a complexity-consistent phenomenon. They replace inappropriate empirical standards with criteria suited to the dissertation's genre and purpose.

Criterion 1: Internal Coherence

WHAT: Defining internal coherence

Internal coherence refers to logical consistency across a theoretical framework: the absence of contradiction and the presence of systematic integration among propositions. Coherence operates at multiple levels. At the propositional level, claims must not contradict one another. At the structural level, components must integrate into a unified whole. At the inferential level, conclusions must follow validly from premises. A theory may be locally consistent yet globally incoherent if its components imply incompatible commitments.

Within the logic of research programmes, coherence is essential because foundational commitments form a structured core that must remain internally consistent as subsidiary claims are developed.

WHY: Relevance for this dissertation

Internal coherence is essential because the dissertation advances a paradigm-level reconstruction rather than a collection of isolated propositions. It replaces one coherent – though contested – system with another. If the work diagnoses mechanistic lock-in as a network of mutually reinforcing assumptions and then proposes alternatives that implicitly reintroduce reductionist or equilibrium commitments, the result would be incoherent. Replacing one coherent framework with an internally conflicted alternative would constitute regress rather than progress.

HOW: Application in this dissertation

Coherence is demonstrated through the dissertation's internal cascade logic:

- The mechanistic commitments are made explicit as an interlocked assumption structure rather than a list of independent premises.
- The reconstruction of the firm is developed in a way that remains consistent with that diagnosis – without solving mechanistic failures by reintroducing mechanistic premises under new terminology.

The examiner's coherence test is straightforward: does the reconstruction genuinely follow from the diagnosis without internal contradiction or half-preserved commitments?

Criterion 2: Problem-Solving Capacity

WHAT: Defining problem-solving capacity

Problem-solving capacity refers to a framework's ability to address problems that rival frameworks cannot resolve. Scientific progress is comparative: a framework need not solve

everything, but it must resolve problems that competing frameworks systematically leave unresolved – whether empirical puzzles or conceptual tensions.

WHY: Relevance for this dissertation

This criterion is decisive because the dissertation claims that mechanistic constructions of the firm systematically fail to handle salient realities such as persistent heterogeneity, emergent capabilities, transformation dynamics, and strategic action under genuine uncertainty. If the alternative framework does not improve problem-handling capacity relative to the mechanistic construction, reconceptualization would be unnecessary.

HOW: Application in this dissertation

Problem-solving capacity is demonstrated through a two-stage logic:

1. Structural diagnosis identifies which problems are generated by mechanistic assumptions themselves, rather than by missing variables or incomplete models.
2. Ontological reconstruction shows how these problems become intelligible rather than residual under a complexity-consistent framing.

The examiner's question is: do the problems diagnosed as structural become tractable under the alternative without ad hoc repairs?

Criterion 3: Theoretical Fruitfulness

WHAT: Defining theoretical fruitfulness

Theoretical fruitfulness is the capacity of a framework to generate new questions, distinctions, and research directions. Fruitful theories expand the space of meaningful inquiry rather than merely redescribing what is already known. A progressive theoretical move opens investigative terrain that was previously incoherent or invisible.

WHY: Relevance for this dissertation

Pure theory enables empirical inquiry by providing the conceptual architecture within which questions become formulable. Fruitfulness operationalizes this enabling claim: it tests whether reconstruction generates a new research horizon or merely relabels familiar ideas.

HOW: Application in this dissertation

The reconstruction should be judged by whether it generates new families of coherent questions that mechanistic framings cannot pose, such as:

- How interaction patterns generate firm-level capabilities irreducible to individual traits.
- What forms of organizational stability are compatible with continuous adaptation rather than equilibrium.
- How historical sequences constrain feasible strategic paths beyond "initial conditions."

- Which meso-level structures transmit effects between micro actions and macro outcomes.

More specifically, the CAS framework enables inquiry into:

- How do governance structures shape emergent capabilities at the meso level?
- What conditions enable or inhibit firm adaptation under fundamental uncertainty?
- How do institutional environments co-evolve with firms' CAS properties?
- What patterns distinguish successful from failed organizational transformation?
- How do firms navigate irreducible uncertainty in strategic decision-making?

These questions arise directly from CAS commitments – emergence, path-dependence, openness, and uncertainty – and cannot be coherently derived from equilibrium-based frameworks.

The examiner's test is whether the framework creates a disciplined agenda of researchable questions rather than remaining at the level of metaphor.

Criterion 4: Illumination of Mechanisms

WHAT: Defining illumination of mechanisms

Illumination of mechanisms refers to the capacity to explain how outcomes are generated through identifiable processes. It goes beyond labeling outcomes or reporting correlations, aiming instead to disclose generative processes such as feedback loops, amplification, lock-in, coordination dynamics, and emergence.

WHY: Relevance for this dissertation

Mechanistic firm theories often reduce explanation to optimization outcomes or treat key dynamics as residual error. A complexity-consistent reconstruction is valuable insofar as it renders generative processes visible rather than suppressing them.

This emphasis aligns with the critical realist position articulated in Section 2.2. Bhaskar's stratified ontology distinguishes the real (generative mechanisms), the actual (events produced), and the empirical (events observed). Scientific explanation targets the real level (Bhaskar, 1978: 56). For firms as CAS, this means identifying the generative processes that give rise to observable organisational phenomena.

HOW: Application in this dissertation

Mechanism illumination is demonstrated by:

- Identifying how mechanistic premises suppress or distort certain classes of mechanisms.
- Providing an ontological vocabulary that legitimizes these mechanisms as objects of theorizing.

The CAS framework illuminates mechanisms that mechanistic explanations leave opaque:

- Organisational collapse is explained not as "bad decisions," but as path-dependent lock-in and capability rigidity.
- Organisational resilience is explained not as "good management," but as emergent capabilities arising from interconnected routines.
- Systemic crises are explained not as isolated regulatory failures, but as risk amplification through network interconnection and feedback.
- Innovation unpredictability is explained not as exogenous shock, but as novelty emerging from recombination under fundamental uncertainty.
- Persistent firm heterogeneity is explained not as static resource differences, but as self-reinforcing, historically contingent capability trajectories.

The examiner's test is whether the reconstruction allows the articulation of plausible generative stories – mechanism sketches that are explanatory without pretending to deliver point prediction.

Criterion 5: Interestingness

WHAT: Defining interestingness

Interestingness refers to a theory's capacity to challenge taken-for-granted assumptions in a way that is both significant and intelligible. Interesting theories deny some audience assumptions while retaining enough continuity to avoid dismissal as absurd.

WHY: Relevance for this dissertation

As a paradigm-level contribution, the dissertation must interrupt routinized assumptions rather than merely extend them. If it fails to produce a meaningful "rethink" effect for a trained reader – by revealing what the mechanistic framework presupposes and obscures – it fails as problematization-driven theory.

HOW: Application in this dissertation

Interestingness is achieved when:

- Implicit assumptions are rendered visible as assumptions.
- A coherent alternative reorganizes what firms are taken to be.

The CAS framework challenges a focused set of foundational assumptions in mainstream economics:

- Where mechanistic theory assumes equilibrium, CAS foregrounds far-from-equilibrium dynamics
- Where mechanistic theory assumes optimisation, CAS emphasizes satisficing and adaptive search

- Where mechanistic theory assumes predictability, CAS recognizes fundamental uncertainty
- Where mechanistic theory assumes decomposability, CAS emphasizes systemic interdependence
- Where mechanistic theory assumes ergodicity, CAS foregrounds path-dependence and transformation

The result is not rhetorical novelty but a controlled, assumption-level disruption that satisfies Davis's criterion of interestingness (1971).

The examiner's test is rhetorical but rigorous: does the work reliably shift the reader from "obvious" to "rethink" without crossing into implausibility?

Criterion 6: Comprehensiveness

WHAT: Defining comprehensiveness

Comprehensiveness refers to domain adequacy rather than maximality. A framework is comprehensive if it addresses the full set of issues that must be addressed for reconceptualization to be stable and meaningful.

WHY: Relevance for this dissertation

Paradigm-level reconceptualization is inherently integrative. Partial reconstruction – changing labels while leaving the foundational architecture intact – produces instability and invites reabsorption into the dominant paradigm.

HOW: Application in this dissertation

Comprehensiveness is demonstrated through:

1. A diagnosis that maps the mechanistic construction as an interconnected assumption network.
2. A reconstruction articulated as a stable conceptual architecture rather than a loose set of metaphors.

The examiner's test is whether the reconstruction answers the diagnostic task at the appropriate level of generality without leaving core issues untreated.

Pattern Prediction as a Criterion for Complex Phenomena

Beyond the six criteria, the dissertation adopts a complexity-consistent standard for prediction. For complex phenomena, the appropriate epistemic aim is pattern-level expectation, not point prediction. Pattern prediction concerns the anticipation of classes of behaviors or recurrent dynamics under general conditions.

The evaluative question becomes whether the reconstructed framework improves our capacity to recognize, classify, and explain recurrent patterns – such as adaptation, lock-in, and emergence – without claiming trajectory-specific forecasts.

Conclusion

This section has articulated criteria appropriate for evaluating pure theoretical work: internal coherence, problem-solving capacity, theoretical fruitfulness, illumination of mechanisms, interestingness, and comprehensiveness, supplemented by pattern-level intelligibility as the appropriate predictive analogue for complex phenomena.

Together, these criteria assess theory on its own terms and provide a rigorous alternative to misplaced demands for empirical verification. They establish how the dissertation's theoretical contribution should be judged and prepare the ground for specifying which evaluation criteria are inappropriate for paradigm-level reconstruction and how the appropriate criteria are to be applied systematically.

2.6.2 Criteria Inappropriate for This Work

Introduction

Section 2.6.1 articulated criteria appropriate for evaluating pure theoretical work: internal coherence, problem-solving capacity, theoretical fruitfulness, illumination of mechanisms, interestingness, and comprehensiveness. This section completes the evaluative framework by identifying criteria that, although legitimate in other genres of inquiry, are inappropriate for this dissertation. The point is not that these criteria are defective in general, but that they become category errors when applied to paradigm-level theoretical reconstruction of the firm under complexity conditions.

Clarifying inappropriate criteria serves a protective function. Objections demanding predictive accuracy, direct empirical verification, naïve falsifiability, or quantitative precision assume evaluation standards suited to derived hypotheses within stabilized paradigms and to repeatable, low-complexity domains. This dissertation operates at a different logical level: it reconstructs the phenomenon and the assumptive architecture through which hypotheses become possible. Accordingly, what replaces the rejected criteria is not exemption from rigor, but a different rigor – coherence, fruitfulness, mechanism illumination, and pattern-level intelligibility.

Criterion 1: Predictive Accuracy

Why it seems appropriate

Predictive accuracy is often treated as the gold standard for theory evaluation. In many scientific domains, a theory that yields accurate predictions is taken to demonstrate reliable correspondence with its target phenomena. Predictive success appears to offer a comparatively objective measure of scientific merit.

Why it is inappropriate here

Predictive accuracy in the point-forecast sense is inappropriate for theories of complex adaptive systems because CAS exhibit properties that defeat the preconditions of point prediction. As Section 2.5.2 established, non-ergodicity, emergence, fundamental uncertainty, context-dependence, and reflexivity undermine the stability, repeatability, and closure that point predictive accuracy presupposes.

Hayek articulated the appropriate limitation: for complex phenomena, "the only predictions that can be derived from such theories will be merely 'pattern predictions' – predictions about some of the general features that a certain class of situations will display" (Hayek, 1967: 27). The methodological error lies not in the theory but in the demand: to require point prediction from essential complexity is to require what the object cannot yield.

Moreover, the demand for point predictive accuracy in such domains encourages false precision. A framework that promises the kind of predictive accuracy typical of simple, controlled phenomena risks misdescribing complex social reality. Acknowledging predictive limits is therefore not theoretical weakness but theoretical honesty. As Cilliers stresses, "we cannot have complete knowledge of complex systems; we can only have knowledge in terms of a certain framework... Our knowledge of complex systems is always provisional" (Cilliers, 2005: 259).

What replaces it

Pattern intelligibility replaces point predictive accuracy. The appropriate evaluative question is not "does the theory forecast exact outcomes?" but whether it identifies recurrent patterns, classes of trajectories, and families of mechanisms that recur across cases under broadly specified conditions. Pattern-level expectation is the appropriate analogue of prediction under essential complexity (Hayek, 1967).

Criterion 2: Empirical Verification

Why it seems appropriate

Empirical verification is central to scientific self-understanding. Theories are commonly expected to be tested against evidence; those that survive are regarded as credible, while those that fail are revised or rejected. Verification appears to demarcate disciplined inquiry from speculation.

Why it is inappropriate here

Direct empirical verification is inappropriate for foundational propositions because such propositions cannot – even in principle – be verified in isolation. As Section 2.5.1 established, foundational commitments are not empirical claims awaiting validation; they are conditions of intelligibility under which empirical claims become formulable and testable.

Machlup's hierarchy of hypotheses makes this point operational: fundamental assumptions "make the system what it is" and "cannot be changed without changing the character of the entire system" (Machlup, 1955: 16). Attempting to "verify" them by direct empirical procedure misconceives their logical role. Verification presupposes an evidential language and a testing protocol; these are themselves paradigm-dependent. The impossibility is therefore logical rather than practical: testing proceeds within a framework; it does not generate the framework. In Mill's terms, confusing the development of principles with their application is a methodological mistake (Mill, 1843).

What replaces it

Evaluation by fruits replaces direct empirical verification. Foundational commitments are assessed indirectly – by what they enable: coherence, explanatory reach, mechanism illumination, and generative research capacity. Conant captured this logic sharply: "a theory is

only overthrown by a better theory, never merely by contradictory facts" (Conant, 1947: 36). At the paradigm level, comparison is holistic and pragmatic: does the framework yield superior intelligibility and problem-handling power?

Criterion 3: Falsifiability in the Naïve Sense

Why it seems appropriate

Popper's falsifiability criterion remains influential: scientific theories should make claims that could be shown false by experience. Theories that seem compatible with any possible observation are often criticized as unscientific.

Why it is inappropriate here

Naïve falsifiability becomes inappropriate when applied to paradigm-level work because paradigms are not falsified by isolated facts. Lakatos showed that evidence always tests a bundle: theory plus auxiliary assumptions and background commitments; consequently, "there is no falsification before the emergence of a better theory" (Lakatos, 1970: 7). Paradigms routinely persist through anomalies; they are not discarded upon first counter-instance.

Kuhn's point is that paradigm competition "is not the sort of battle that can be resolved by proofs" (Kuhn, 1962: 148). What changes in a paradigm shift is not one statement but a problem-field: what counts as a problem, what counts as relevant evidence, and what counts as an adequate explanation. Demanding Popper-style falsifiability as a gatekeeping criterion for foundational reconstruction therefore misidentifies the unit of appraisal.

This is not a rejection of scientific discipline. It rejects an inappropriate image of discipline: the expectation of decisive refutation of foundational frameworks via paradigm-neutral observation.

What replaces it

Programme appraisal replaces naïve falsifiability. The relevant question becomes whether a reconstruction exhibits progressive character – whether it resolves anomalies without proliferating ad hoc repairs, generates novel insight, and opens productive inquiry. This is the logic of comparative research programme evaluation (Lakatos) rather than the logic of the "crucial experiment."

Criterion 4: Quantitative Precision

Why it seems appropriate

Quantitative precision is often treated as a marker of "mature" science: exact measurement and mathematically expressible relations are taken as evidence of rigor. Social inquiry is frequently pressured to emulate this standard.

Why it is inappropriate here

Quantitative precision is inappropriate as a general criterion for complex social phenomena because it demands a level of closure, stationarity, and variable isolation that complexity undermines. Hayek warned that importing standards from simple-phenomena sciences into complex-phenomena domains constitutes the "pretence of knowledge" (Hayek, 1974):

claiming exactness where the object does not permit it. Hayek also inverted the usual "maturity" inference: physics achieves high precision not because its methods are universally transferable, but because its subject matter is, in the relevant sense, relatively simple (Hayek, 1967).

Cilliers adds the decisive methodological point: models of complex systems must reduce complexity to function as models (Cilliers, 2005). The resulting precision often reflects the model's reduction, not the system's determinacy. In this sense, precision can become a form of misrepresentation: exactness about the wrong object.

None of this implies anti-quantitative bias. Quantitative tools can be valuable for describing distributions, characterizing patterns, and mapping networks. What is rejected is quantitative precision as a universal badge of scientific legitimacy for paradigm-level propositions about complex dynamics.

What replaces it

Conceptual precision replaces quantitative precision: clarity of constructs, disciplined distinctions, explicit articulation of relationships and boundary conditions. For complex phenomena, rigor resides in definitional control, mechanism plausibility, pattern-level claims, and transparent scope – rather than in spurious exactness.

Conclusion

This section has identified four criteria that are inappropriate for evaluating this dissertation: predictive accuracy (in the point-prediction sense), direct empirical verification of foundational commitments, falsifiability in the naïve sense, and quantitative precision as a universal standard. Each criterion is legitimate in certain scientific genres, but each becomes a category error when applied to paradigm-level theoretical reconstruction of complex adaptive systems.

Rejecting these criteria does not weaken evaluation; it clarifies it. The dissertation remains fully accountable under the criteria articulated in Section 2.6.1: internal coherence, problem-solving capacity, theoretical fruitfulness, illumination of mechanisms, interestingness, and comprehensiveness – supplemented by pattern-level intelligibility appropriate to complexity.

2.6.3 Lakatosian Self-Appraisal

Introduction

Sections 2.6.1 and 2.6.2 articulated criteria appropriate and inappropriate for evaluating pure theoretical work. This section applies Imre Lakatos's methodology of scientific research programmes to assess whether the framework developed in this dissertation constitutes a progressive or degenerating research programme.

Lakatos provides the most demanding and systematic account of how theoretical frameworks can be evaluated without collapsing into either naïve falsificationism or relativistic paradigm sociology. Applying his criteria to the present dissertation serves two purposes. First, it situates the contribution within a well-established philosophy of science. Second, it demonstrates intellectual discipline by subjecting the proposed CAS-based framework to explicit standards of theoretical appraisal rather than exempting it from evaluation.

The assessment proceeds through four Lakatosian criteria: novel predictions, excess content, corroboration, and non-ad-hoc character. For each criterion, the section explains (i) what Lakatos requires, (ii) why the criterion is relevant for this dissertation, and (iii) how the dissertation performs relative to it. The appraisal is deliberately balanced: it identifies strengths where the framework is demonstrably progressive, and it acknowledges limits where claims are partial rather than complete.

This section proceeds in three parts. Part A presents Lakatos's framework for research programme evaluation. Part B explains the relevance of Lakatosian criteria for this dissertation. Part C applies these criteria to the CAS framework developed herein.

Part A: Lakatos's Framework for Research Programme Evaluation

Research Programmes, Not Isolated Theories

Lakatos's central methodological insight is that scientific appraisal cannot be applied to isolated theories or propositions but only to research programmes – historically extended sequences of theories unified by a shared set of foundational commitments. As he states:

"Not an isolated theory, but only a series of theories can be said to be scientific or unscientific: to apply the term 'scientific' to one single theory is a category mistake" (Lakatos, 1970: 6).

Evaluation must therefore focus on the structure and trajectory of a programme rather than on individual claims abstracted from their theoretical context.

Each research programme consists of a hard core of foundational assumptions and a surrounding protective belt of auxiliary hypotheses. The hard core is "irrefutable by the methodological decision of its proponents" (Lakatos, 1970: 8), not because it is immune to criticism, but because abandoning it dissolves the programme itself. Empirical testing and revision operate primarily at the level of the protective belt.

Progressive versus Degenerating Programmes

The decisive Lakatosian distinction is between progressive and degenerating research programmes. A programme is progressive when its theoretical development anticipates empirical and conceptual discoveries rather than merely accommodating them after the fact. In Lakatos's formulation:

"A research programme is said to be progressing as long as its theoretical growth anticipates its empirical growth, that is, as long as it keeps predicting novel facts with some success" (1970: 5).

Progressive programmes generate new questions, reveal previously unrecognised phenomena, and open fresh research directions. Degenerating programmes, by contrast, display reactive adjustment:

"A programme is degenerating if its theoretical growth lags behind its empirical growth, that is, as long as it only gives post hoc explanations either of chance discoveries or of facts anticipated by, and discovered in, a rival programme" (Lakatos, 1970: 5).

Degeneration does not imply immediate rejection. Lakatos explicitly allows that programmes may temporarily degenerate and later recover through theoretical innovation. The distinction provides an evaluative framework, not a mechanical rule for abandonment.

Crucially, Lakatos rejects the idea that single experiments can decisively refute research programmes:

"No experiment, experimental report, observation statement or well-corroborated low-level falsifying hypothesis alone can lead to falsification. There is no falsification before the emergence of a better theory" (1970: 7).

Scientific change, on this view, occurs through replacement by superior programmes, not through decisive refutation.

Heuristics and Programme Dynamics

Lakatos further distinguishes between the negative heuristic and the positive heuristic of a research programme. The negative heuristic specifies what must not be questioned:

"The negative heuristic forbids us to direct the modus tollens at the hard core" (Lakatos, 1978: 48).

The positive heuristic, by contrast, provides guidance for theory development:

"It consists of a partially articulated set of suggestions or hints on how to change, develop, the 'refutable variants' of the research programme" (Lakatos, 1978: 50).

Progressive programmes possess strong positive heuristics that structure inquiry, guide model construction, and generate new research paths. Degenerating programmes lack such guidance and respond to anomalies only through ad hoc modifications.

Four Criteria of Progressive Character

Lakatos identifies four criteria for assessing whether a research programme is progressive:

- **Novel predictions** – Does the programme anticipate phenomena that rival programmes do not?
- **Excess content** – Does it explain everything rivals explain plus additional phenomena?
- **Corroboration** – Is some of this excess content supported by empirical or historical evidence?
- **Non-ad-hoc development** – Are theoretical adjustments principled and internally motivated rather than merely accommodative?

As Lakatos summarises:

"For the sophisticated falsificationist a theory is 'acceptable' or 'scientific' only if it has corroborated excess empirical content over its predecessor (or rival)" (1970: 4).

These criteria provide a demanding but non-naïve standard – especially appropriate for paradigm-level theoretical work.

Relationship to Kuhn

Lakatos explicitly frames his methodology as a rational reconstruction of Kuhn's paradigm theory:

"My concept of a 'research programme' may be construed as an objective, 'third world' reconstruction of Kuhn's socio-psychological concept of 'paradigm': thus the Kuhnian 'Gestalt-switch' can be performed without removing one's Popperian spectacles" (Lakatos, 1970: note 25).

This synthesis is central for the present dissertation. The shift from mechanistic to CAS-based conceptions of the firm can be understood as a Kuhnian paradigm change, but Lakatos's framework allows this shift to be assessed in terms of rational progress rather than mere sociological conversion.

Part B: Relevance of Lakatosian Criteria for This Dissertation

This dissertation can be understood as articulating the hard core of a CAS-informed research programme for the study of firms. The foundational commitments – that firms are complex adaptive systems characterised by emergence, path-dependence, and fundamental uncertainty – constitute the hard core. The theoretical developments in Chapter 4 form the protective belt, while the problematisation methodology provides the programme's positive heuristic, guiding systematic engagement with anomalies in mechanistic approaches.

Understanding the dissertation as a research programme enables appropriate evaluation. The relevant question is not whether individual propositions survive isolated tests, but whether the programme as a whole exhibits progressive character – that is, whether it anticipates novel insights and develops productively.

Self-appraisal serves four purposes. First, it demonstrates intellectual consistency: a dissertation that critiques mechanistic firm theory as degenerating must show that its proposed alternative does not suffer the same defect. Second, it calibrates claims to what Lakatosian analysis supports, avoiding overstatement. Third, it provides examiners with explicit evaluative criteria, focusing assessment on shared standards rather than implicit expectations. Fourth, it enacts the reflexive stance implied by the dissertation's critical realist commitments, applying its evaluative framework to itself.

Lakatos's treatment of anomalies is especially relevant. As he notes, anomalies "always abound in science" and become decisive only when a superior theory emerges to explain them. This perspective shapes the dissertation's treatment of mechanistic anomalies – not as refutations, but as phenomena that motivate and structure CAS-based reconstruction.

Part C: Self-Application of Lakatosian Criteria

Criterion 1: Novel Predictions

Requirement. A progressive programme anticipates phenomena that rival programmes do not.

Assessment. The CAS framework anticipates patterns that mechanistic approaches cannot formulate, including emergent organisational capabilities, path-dependent development trajectories, threshold effects in transformation, meso-level mediation between micro and macro dynamics, and the persistence of fundamental uncertainty. These patterns are treated as inherent features within CAS theory but appear as anomalies or deviations within mechanistic models.

Limitation. These predictions are theoretical and pattern-based rather than empirically tested within the dissertation itself – consistent with its positioning as pure theory.

Criterion 2: Excess Content

Requirement. A progressive programme explains everything its rival explains, plus additional phenomena.

Assessment. The CAS framework exhibits excess content in two respects. First, it incorporates valid mechanistic insights (e.g., incentives, agency, transaction costs) within a broader framework rather than rejecting them outright. Second, it explains phenomena that mechanistic approaches treat as anomalies, including persistent firm heterogeneity, transformation failure, innovation unpredictability, and differential resilience – using mechanisms such as emergence, adaptation, and path-dependence.

Limitation. Incorporation is interpretive rather than derivational; mechanistic theorists may contest whether their insights are subsumed or merely recontextualised.

Criterion 3: Corroboration

Requirement. Some excess content must be empirically supported.

Assessment. Corroboration is indirect but multi-sourced. Illustrative cases (e.g., Nokia, Toyota, the 2008 financial crisis) demonstrate CAS mechanisms in organisational settings. Findings from complexity science corroborate the plausibility of emergence and self-organisation across domains. The framework's ability to resolve anomalies unresolved by mechanistic theory provides further indirect support.

Limitation. The dissertation does not conduct systematic empirical testing. Corroboration is illustrative and abductive rather than experimental – an acknowledged and justified limitation given the dissertation's theoretical purpose.

Criterion 4: Non-Ad-Hoc Character

Requirement. Theoretical development must follow internal logic rather than reactive accommodation.

Assessment. The CAS framework exhibits non-ad-hoc character through its integrated architecture. Emergence, path-dependence, and fundamental uncertainty are mutually entailing

features of CAS rather than independent add-ons. The framework derives its claims from general CAS principles rather than inventing properties to fit firm-specific anomalies. The problematisation methodology provides a positive heuristic that guides development proactively.

Limitation. The assessment presupposes that firms genuinely instantiate CAS properties. Chapter 4 addresses this by demonstrating that firms meet the defining criteria of complex adaptive systems.

Conclusion

By Lakatosian standards, the dissertation constitutes a progressive research programme. It satisfies three criteria fully (novel predictions, excess content, and non-ad-hoc character) and one partially (corroboration). The partial nature of corroboration reflects the dissertation's role as foundational theory enabling future empirical research, not as an empirical test itself. This self-appraisal demonstrates that the CAS framework meets the evaluative standards it applies to mechanistic alternatives, thereby satisfying the reflexivity requirement central to this dissertation's methodological commitments.

Conclusion of Section 2.6

Section 2.6 has specified the success criteria appropriate to pure theoretical work and applied them reflexively to the dissertation. Three findings carry forward.

First, six criteria are appropriate for evaluating foundational theory: internal coherence (logical consistency across the framework without hidden contradictions), problem-solving capacity (addressing problems that rival frameworks cannot resolve), theoretical fruitfulness (generating new questions, distinctions, and research directions), illumination of mechanisms (disclosing generative processes rather than merely labelling outcomes), interestingness (challenging taken-for-granted assumptions in ways that produce meaningful rethinking), and comprehensiveness (addressing the full set of issues required for stable reconceptualization). Pattern prediction – anticipating classes of behaviours and recurrent dynamics under general conditions – replaces point prediction as the appropriate standard for complex phenomena.

Second, four criteria are categorically inappropriate: predictive accuracy in the point-forecast sense (replaced by pattern intelligibility), direct empirical verification of foundational commitments (replaced by evaluation through the fruits of frameworks), falsifiability in the naïve sense (replaced by research programme appraisal), and quantitative precision as a universal standard (replaced by conceptual precision). These criteria become category errors when applied to paradigm-level reconstruction of complex adaptive systems.

Third, by Lakatosian standards the dissertation constitutes a progressive research programme. It satisfies three criteria fully – novel predictions (anticipating patterns mechanistic approaches cannot formulate), excess content (incorporating mechanistic insights while explaining their anomalies), and non-ad-hoc character (deriving claims from integrated CAS principles rather than reactive accommodation) – and one criterion partially (corroboration is indirect and illustrative, consistent with foundational theory's role in enabling rather than conducting empirical research).

Section 2.7 specifies the research design and architecture through which these evaluative criteria are operationalized across the dissertation's substantive chapters.

2.7 Research Design and Architecture

Section 2.6 established the criteria by which pure theoretical work should be evaluated – coherence, problem-solving capacity, fruitfulness, mechanism illumination, interestingness, and comprehensiveness – and demonstrated through Lakatosian self-appraisal that the dissertation constitutes a progressive research programme. This section turns from evaluation to execution, explaining how the dissertation's architecture implements systematic paradigm transformation.

This specification is necessary because paradigm-level transformation cannot be achieved piecemeal. As established in Section 2.3, the mechanistic paradigm's assumptions form a mutually reinforcing network – the Triple Lock – that neutralizes isolated challenges. An architecture adequate to this task must address the assumption network as a whole, proceeding from explicit diagnosis to coordinated reconstruction.

The argument proceeds through four subsections. Subsection 2.7.1 establishes the overall design of systematic paradigm transformation and the four-chapter architecture that implements it. Subsection 2.7.2 explicates the multi-level cascade structure through which ontological, epistemological, and methodological transformations are logically connected. Subsection 2.7.3 explains why the assumptive network – fifteen assumptions organised into hard core and protective belt – constitutes the unit of analysis. Subsection 2.7.4 specifies the structural conventions that enable coherence at scale.

Section 2.8 addresses the practical tools and techniques through which the inquiry was conducted.

2.7.1 Overall Design: Systematic Paradigm Transformation

Introduction: From Evaluation to Execution

Section 2.6 established the criteria by which this dissertation should be evaluated, including a programme-level logic that judges paradigm work by coherence, problem-solving capacity, fruitfulness, and mechanism illumination rather than by inappropriate empirical standards. This section turns from evaluation to execution. It explains how the dissertation's compressed architecture (Chapters 1–4) implements the systematic paradigm transformation developed in Sections 2.3–2.6.

The research design is not arbitrary. It follows directly from the methodological commitments established earlier – most importantly, the problematization methodology (Section 2.3), the phenomena-construction approach (Section 2.4), and the recognition that empirical testing is categorically inappropriate for foundational propositions (Section 2.5). Once these commitments are accepted, the architecture becomes constrained rather than discretionary.

The central insight organizing this section is that paradigm-level transformation cannot be achieved piecemeal. As Section 2.3.5 demonstrated, the mechanistic paradigm's assumptions form a mutually reinforcing network – the Triple Lock – that neutralizes isolated challenges. An architecture adequate to this task must therefore address the assumption network as a whole, proceeding from (i) explicit diagnosis of the network to (ii) reconstruction of an alternative grounding capable of displacing it.

As Alhadeff-Jones (2013: 19) observes in the context of complexity-oriented inquiry, "there was no ready-made procedure to guide my inquiry... I eventually produced what I was desperately looking for: a kind of method guiding my choices and informing my decisions." The architecture of this dissertation represents precisely such a method: not a formulaic procedure, but a systematic framework for navigating paradigm-level transformation under conditions of complexity.

Part A: WHAT Is Systematic Paradigm Transformation?

Defining the Approach

Systematic paradigm transformation refers to a research design that treats foundational assumptions as an interconnected architecture rather than as independent propositions amenable to local revision. Paradigms, as Kuhn emphasized, constitute disciplinary matrices: "the entire constellation of beliefs, values, techniques, and so on shared by members of a given community" (1962/1970: 175). Because these elements mutually condition one another, transformation requires coordinated intervention across the framework as a whole.

Piecemeal modification fails for precisely this reason. Challenging one assumption while leaving the surrounding structure intact allows the remaining assumptions to absorb, reinterpret, or neutralize the challenge. In research-programme terms, anomalies are accommodated through adjustments elsewhere while the hard core remains intact. Systematic transformation, by contrast, targets the hard core at the level of assumption architecture: it renders implicit commitments explicit, demonstrates their mutual reinforcement, and reconstructs an alternative grounding capable of reorganizing what counts as an admissible explanation.

The Four-Chapter Architecture

The dissertation implements systematic paradigm transformation through a four-chapter architecture designed to operationalize problematization under compression:

Table 3: Chapter Structure and Transformative Logic of the Thesis: From Paradigm-Level Problematization to Theoretical Reconstruction

Chapter	Content	Transformation Function
1	Introduction	Frame the paradigm-level problem; justify the need for systematic rather than piecemeal transformation
2	Methodology	Establish problematization and phenomena construction; justify why empirical testing is inapplicable at the foundational level; specify evaluation standards
3	Diagnostic	Identify and map the mechanistic assumption network; demonstrate its mutual reinforcement (Triple Lock)
4	Reconstruction	Construct the alternative phenomenon ("firms as CAS") and its foundational commitments as a coherent replacement framework

This architecture represents compression in execution, not reduction in ambition. The dissertation performs the full paradigm task – diagnosis plus reconstruction – while integrating what would otherwise be distributed across multiple chapters into a single reconstruction chapter. The methodological warrant for this compression is established in Section 2.5:

paradigm-level propositions are evaluated by coherence and fruitfulness, not by direct empirical verification.

The Two-Phase Logic: Diagnosis → Reconstruction

Within this compressed design, systematic paradigm transformation proceeds through a two-phase logic:

1. **Diagnostic phase (Chapter 3):** make mechanistic assumptions explicit, map their interdependence, and demonstrate why piecemeal critique fails.
2. **Reconstruction phase (Chapter 4):** construct an alternative phenomenon and its internal logic, demonstrating how it dissolves the diagnostic lock by reorganizing the conceptual field.

This diagnosis–reconstruction sequence operationalizes the core requirement of problematization: not merely to criticize assumptions, but to develop an alternative assumption ground sufficiently coherent and generative to reconfigure what counts as an admissible question and explanation.

Part B: WHY Systematic Transformation Is Required

The Nature of the Problem

The dissertation addresses a paradigm-level problem that cannot be resolved through gap-spotting research. As Section 2.3.2 established, gap-spotting under-problematizes existing literature and thus reinforces dominant frameworks. The limitations of the mechanistic paradigm are not missing variables or underexplored cases; they are foundational commitments requiring transformation.

At the same time, the intervention must remain intelligible to its scholarly audience. Systematic transformation satisfies this requirement by engaging directly with mainstream categories and problem framings while reconstructing the commitments that give those categories their meaning. The result is continuity at the level of engagement combined with discontinuity at the level of foundations.

Critical Realism and Ontological Priority

Bhaskar's critical realism provides philosophical grounding for the design's ordering: ontology conditions epistemology, which in turn conditions methodological possibility. The decisive intervention therefore occurs at the level of what the firm is taken to be. In a compressed architecture, this logic does not require separate epistemology or methodology chapters; it requires that epistemic and methodological implications be derived from ontological reconstruction rather than appended as independent preferences.

Operationalizing Problematization Under Compression

Alvesson and Sandberg identify six principles of problematization: identifying a domain, articulating assumptions, evaluating them, developing alternatives, considering the audience, and evaluating the alternative (2011: 253). The dissertation maps onto these principles as follows:

Table 4: Mapping Problematization Principles to Dissertation Implementation: Operationalising Methodological Critique Across Chapters

Problematization Principle	Dissertation Implementation
Identify domain	Chapter 1 frames mainstream economics' treatment of the firm
Identify assumptions	Chapter 3 maps mechanistic assumptions and their network
Evaluate assumptions	Chapter 3 demonstrates the Triple Lock and failure of piecemeal critique
Develop alternatives	Chapter 4 constructs the CAS-based phenomenon
Consider audience	Chapter 2 clarifies genre, evaluation criteria, and limits of empirical demands
Evaluate alternatives	Chapter 4 applies coherence, problem-handling, and mechanism criteria

Compression redistributes these steps across fewer chapters without eliminating any of them, preserving the logic of systematic transformation.

Part C: HOW Systematic Transformation Is Implemented

Chapter 3: Diagnostic Function

Chapter 3 performs the diagnostic phase by identifying and mapping the assumption network requiring transformation. It renders explicit what is typically implicit in mainstream firm theory and demonstrates how assumptions cluster into mutually reinforcing locks that immunize against isolated challenge.

This diagnostic work is indispensable. The assumptions most in need of critique are precisely those that remain backgrounded and unnamed. By mapping their interdependence, the chapter shows why isolated criticism is predictably absorbed and neutralized by the remaining network.

Chapter 4: Reconstruction as Integrated Replacement

Chapter 4 executes the reconstruction phase by constructing the alternative phenomenon: firms as complex adaptive systems. Within the compressed architecture, this chapter does not merely propose an ontological slogan. It builds an integrated framework in which:

- the reconceptualization specifies what counts as a firm property (e.g., emergence, adaptation, path dependence);
- epistemic implications follow from the nature of the phenomenon (e.g., limits of point prediction, primacy of pattern intelligibility);
- methodological implications are justified as requirements of investigating such an object (e.g., methodological fit, plural evidential strategies, mechanism-oriented explanation).

Systematic transformation is thus achieved not by attacking individual assumptions but by constructing an alternative grounding that reorganizes the field's conceptual possibilities.

Cascade Logic Under Compression

Even without separate epistemology or methodology chapters, the dissertation preserves the necessary ordering of commitments:

- reconstruction of the phenomenon (what the firm is) generates
- epistemic implications (what can be known and with what limits), which generate
- methodological implications (how inquiry must proceed to remain faithful to the object).

This ordering enforces discipline. Epistemic and methodological claims appear as consequences of the reconstructed phenomenon rather than as discretionary methodological preferences.

Conclusion: Architecture as Methodological Achievement

This section has shown that the dissertation's four-chapter architecture is not arbitrary but follows necessarily from the nature of paradigm-level transformation under complexity conditions. The design achieves:

1. **Systematic intervention** by treating assumptions as an interdependent network rather than isolated propositions.
2. **Problematization in practice** through explicit diagnosis (Chapter 3) and coherent reconstruction (Chapter 4).
3. **Resolution of the Triple Lock** by replacing the mechanistic assumption network with an alternative grounding rather than attempting piecemeal repair.
4. **Methodological discipline under compression** by deriving epistemic and methodological implications from ontological reconstruction rather than appending them as add-ons.

Section 2.7.2 builds on this foundation by examining the architecture's internal dependencies and sequencing in greater detail, showing how diagnosis and reconstruction interact to produce a coherent, non-piecemeal paradigm transformation.

2.7.2 The Multi-Level Structure

Introduction: From Architecture to Logic

Section 2.7.1 established the dissertation's compressed architecture (Chapters 1–4) and introduced the cascade logic through which systematic paradigm transformation is executed. This section explicates that logic in detail. It shows how ontological, epistemological, and methodological transformations are related – not as parallel or optional moves, but as logically connected levels that condition one another.

The multi-level structure is not an organizational convenience. It is a philosophical necessity grounded in the critical realist insight that what exists constrains what can be known, which in turn constrains how inquiry can legitimately proceed. As Woodill notes in relation to complexity, "its ontological orientation compels its epistemology. In other words, complexity compels a view of the world in which there are no static or inert objects, only relationships" (2021: 2). This recognition – that ontological commitments carry epistemological and methodological consequences – provides the foundation for the dissertation's cascade logic even under a compressed four-chapter design.

Part A: WHAT Is the Multi-Level Structure?

The Cascade Logic

The dissertation implements paradigm transformation through a three-level cascade:

Table 5: Three-Level Paradigm Transformation: Ontological, Epistemological, and Methodological Shifts from Mainstream to Complex Adaptive Systems Foundations

Level	Function	Location in Dissertation	Core Transformation
Ontology	What exists	Chapter 4 (core reconstruction)	FROM reductionist substance TO emergent process
Epistemology	What can be known	Chapter 4 (derived implications)	FROM deterministic certainty TO fundamental uncertainty
Methodology	How to inquire	Chapter 4 (derived; guided by Chapter 2)	FROM methodological monism TO structured pluralism

This ordering reflects the principle articulated by critical realism: "it is the nature of objects that determines their cognitive possibilities for us" (Bhaskar, 1975). The sequence is therefore not optional. Ontology constrains epistemology, and epistemology constrains methodology. Methods cannot be justified independently of what can be known, and what can be known cannot be established independently of what exists.

Complexity theory reinforces this logic. As Woodill emphasizes, a relational, processual ontology compels an epistemology attentive to interaction, contingency, and openness (2021: 2). Once reality is understood as dynamic and relational rather than static and atomistic, knowledge claims must be provisional and context-sensitive, and methodological strategy must adapt accordingly.

The Three Transformations Under Compression

Because the dissertation's architecture is compressed, the three transformations are not distributed across separate chapters. Instead, they are implemented as an integrated bundle with explicit internal ordering:

Ontological transformation (Chapter 4: core move). The dissertation shifts from substance ontology – where firms are treated as bounded entities with fixed properties – to process ontology – where firms are constituted through ongoing flows of activity, interaction, and adaptation. This establishes emergence, self-organization, and path-dependence as ontologically real features of firms, not as anomalies or residual effects.

Epistemological transformation (Chapter 4: derived implication). From this ontological foundation follows an epistemological shift from deterministic certainty to fundamental uncertainty. If firms are processes rather than substances, and if interaction generates genuinely novel properties, then future states cannot be deduced from present conditions. Uncertainty is therefore not merely informational but constitutive. As Cilliers notes, "we cannot have complete knowledge of complex systems; we can only have knowledge in terms of a certain framework... Our knowledge of complex systems is always provisional" (2005a: 259).

Cilliers clarifies that this follows directly from the necessity of boundary-drawing: "We cannot 'know' a complex thing in all its complexity... Knowledge and data-reduction are intertwined. We can have knowledge because we draw boundaries" (2005b: 607). Epistemic provisionality is thus not a defect but a condition of knowing complex phenomena.

Methodological transformation (Chapter 4: further implication, justified by Chapter 2). From this epistemic stance follows a methodological shift from monism to structured pluralism. If uncertainty is fundamental rather than temporary, inquiry cannot rely on a single universally optimal method. Methods must be selected for fit with the object, the level of analysis, and the question at hand.

Hayek anticipated this implication in his analysis of complex phenomena: "The multiplicity... of distinct elements... required to explain a complex phenomenon... creates problems which dominate the disciplines concerned with such phenomena" (1967: 26). In such domains, methodological strategy must be adaptive rather than monistic.

Bidirectionality

Although the cascade flows from ontology to epistemology to methodology, the relationships are bidirectional and mutually reinforcing:

- **Forward flow:** Ontological reconceptualization (firms as emergent processes) grounds epistemic limits (fundamental uncertainty and provisionality), which ground methodological requirements (plural, mechanism-oriented inquiry).
- **Backward illumination:** Persistent methodological failure exposes epistemic overreach; epistemic overreach reveals ontological mis-specification. Repeated failure of monistic methods in complex domains signals an implicit assumption of stability and decomposability; recognizing the failure of such certainty reveals the deeper ontological error of treating dynamic processes as static substances.

This bidirectional structure is essential for paradigmatic coherence. The elements form reinforcing loops rather than linear chains. As Alhadeff-Jones observes, complex inquiry involves "a complex interplay between complementary, contradictory and antagonistic forms of self and mutual control" (2013: 33). The dissertation's multi-level structure is designed to achieve this integrated coherence rather than a patchwork of disconnected claims.

Part B: WHY the Multi-Level Structure Is Necessary

Logical Necessity

The multi-level structure follows from the nature of paradigm-level transformation. Bhaskar's ordering from ontology to epistemology to methodology "is not an arbitrary order of development" but reflects the dependence of cognitive possibility on the nature of objects (1975). Attempting to transform method while leaving ontology intact produces incoherence: pluralist methods cannot be justified under deterministic epistemology, and deterministic epistemology cannot be sustained if emergence and open-ended adaptation are ontologically real.

The dependencies are therefore substantive. If emergence is real (ontology), then wholes are not deducible from parts (epistemology), and reductionist, monistic methods become categorically inappropriate (methodology). The cascade is not a stylistic choice; it is demanded by the logic of the phenomenon.

Addressing the Triple Lock Under Compression

Section 2.3.5 identified three mutually reinforcing locks – ontological, epistemological, and methodological – that stabilize the mechanistic paradigm. The multi-level structure addresses this directly even under compression. Chapter 3 maps how the locks reinforce one another; Chapter 4 reconstructs an alternative grounding that breaks the circular reinforcement by reorganizing commitments at their root.

Reductionist ontology supports deterministic epistemology; deterministic epistemology supports monistic methodology; monistic methodology feeds back into reductionist ontology by privileging decompositional approaches as "scientific." Breaking this cycle requires coordinated transformation across levels – implemented here as an integrated reconstruction rather than as isolated interventions.

Coherence Through Integration

The multi-level structure also prevents fragmentation. Local critique can succeed while generating global inconsistency: equilibrium may be rejected while closure is retained; complexity vocabulary may be adopted while deterministic success criteria persist. The cascade logic blocks this by making dependencies explicit: epistemic and methodological claims must follow from the reconstructed ontology.

As Cilliers emphasizes, limits are not a defect but a condition of knowledge: "The fact that our knowledge is limited is not a disaster, it is a condition for knowledge. Limits enable knowledge" (2005b: 612). The multi-level structure supplies these enabling limits by enforcing coherence across levels.

Part C: HOW the Multi-Level Structure Is Implemented

The Cascade in Practice (Chapters 3–4)

The dissertation operationalizes the cascade through explicit links between diagnosis and reconstruction:

1. **Diagnostic mapping (Chapter 3).** The assumption network is identified and shown to function as a mutually reinforcing system (the Triple Lock), establishing why isolated critique fails and why reconstruction must be multi-level.
2. **Ontological reconstruction (Chapter 4: core).** The firm is reconstructed as a complex adaptive phenomenon characterized by process, emergence, and path-dependence.
3. **Epistemic consequences (Chapter 4: derived).** From this reconstruction follow limits on prediction, the provisionality of models, and the treatment of uncertainty as constitutive rather than merely informational.
4. **Methodological consequences (Chapter 4: derived; framed by Chapter 2).** From the epistemic stance follow requirements of method–object fit, plural evidential strategies, and mechanism-oriented explanation. Pluralism is structured, not permissive: constrained by the nature of the object and the purposes of inquiry.

Bidirectionality in Practice

Bidirectionality is implemented by using methodological and epistemic failures diagnosed in Chapter 3 to motivate ontological correction in Chapter 4, and then using the reconstructed ontology to re-derive epistemic and methodological standards that avoid those failures. The result is an integrated loop: diagnosis explains failure; reconstruction supplies replacement; derived standards restore intelligibility without importing inappropriate criteria from simple-phenomena sciences.

Conclusion: Multi-Level Structure as Philosophical Discipline

This section has shown that the dissertation's multi-level structure implements a philosophical principle – ontology conditions epistemology, which conditions methodology – rather than an organizational preference. Under a compressed architecture, the same logic is preserved by deriving epistemological and methodological transformations explicitly from ontological reconstruction, anchored by diagnostic mapping of the assumption network.

Section 2.7.3 builds on this foundation by examining the assumption network itself as a unit of analysis, showing why the commitments identified in Chapter 3 form an interconnected system rather than an independent list, and why reconstruction must therefore be systematic rather than piecemeal.

2.7.3 The Assumptive Network as Unit of Analysis

Introduction: From Levels to Network

Section 2.7.2 established that paradigm transformation must address ontology, epistemology, and methodology through a cascade structure respecting logical priority. This section develops a complementary insight: assumptions at each level do not operate independently but form an interconnected network. Paradigm resilience arises not from isolated assumptions but from their mutual reinforcement.

Treating assumptions as a network rather than a list follows from three considerations: (i) the failure of piecemeal problematization (Section 2.3.5), (ii) the relational logic emphasized by Ma-inspired problematization (Section 2.3.3), and (iii) the recognition that paradigms persist through interaction structures rather than through individual commitments. The unit of analysis for paradigm transformation is therefore the assumptive network as a whole.

As Newell observes in relation to complex systems and interdisciplinarity, "divide and conquer" strategies systematically obscure interactions, even though those interactions often determine system behavior (2001: 13). The same applies to assumptions: isolating them leaves unexamined the relational structure that generates paradigmatic stability.

Part A: WHAT Is the Assumptive Network?

The Fifteen Core Assumptions

The dissertation identifies fifteen core assumptions embedded in mainstream economics' treatment of firms. These assumptions are organized hierarchically into two tiers, following Lakatos's distinction between a research programme's hard core and its protective belt. The hard core is shielded from direct refutation by methodological decision, while the protective belt contains derivative commitments adjustable without altering foundations (Lakatos, 1970: 133).

Tier 1: Foundational Assumptions (Hard Core)

The mechanistic paradigm rests on nine foundational assumptions, divided into ontological and epistemological commitments.

Ontological assumptions (5):

Table 6: *Ontological Assumptions of Mainstream Theory of the Firm and the Complex Adaptive System Phenomena They Structurally Eliminate*

#	Assumption	Core claim	What it eliminates
1	Reductionism	Systems understood via decomposition	Emergence
2	Determinism	Complete information enables prediction	Fundamental uncertainty
3	Substance ontology	Reality consists of stable entities	Process, transformation
4	Closure	Systems have fixed boundaries	Environmental co-constitution

#	Assumption	Core claim	What it eliminates
5	Atomism	Units are independent; relations external	Relational constitution

Epistemological assumptions (4):

Table 7: Epistemological Assumptions of Mainstream Theory of the Firm and the Complex Adaptive System Phenomena They Structurally Eliminate

#	Assumption	Core claim	What it eliminates
6	Certainty	Knowledge can be complete	Fundamental uncertainty
7	Universal laws	Regularities are context-free	Local knowledge
8	Monism	One correct model per phenomenon	Complementarity
9	Perfect information	Agents possess complete information	Bounded rationality

Tier 2: Derivative Assumptions (Protective Belt)

These six assumptions are not independent choices but logical consequences of foundational combinations.

Table 8: Derivative Assumptions of Mainstream Theory of the Firm: Their Generative Foundations and the Complex Adaptive System Phenomena They Eliminate

#	Derivative assumption	Generated by	What it eliminates
10	Equilibrium	Substance + determinism + closure + certainty	Far-from-equilibrium dynamics
11	Linearity	Reductionism + determinism + universal laws	Threshold effects
12	Homogeneity	Atomism + monism + universal laws + perfect information	Structural heterogeneity
13	Ergodicity	Substance + determinism + certainty + closure	Path-dependence
14	Optimisation	Certainty + perfect information + determinism	Heuristics, routines
15	Decomposability	Reductionism + atomism + closure + linearity	Emergent irreducibility

The hierarchical structure is decisive. Tier-2 assumptions are regenerated from Tier-1 commitments. This explains why critiques targeting equilibrium, optimisation, or ergodicity routinely fail: unchanged foundations simply reproduce them under new guises. Lakatos's warning applies directly – the object of appraisal is the structured system, not isolated propositions.

Network Treatment: The Ma-Inspired Insight

The decision to treat assumptions as a network rather than a list reflects an insight drawn from the Japanese concept of Ma (間) – the constitutive "interval" or "space between." As Pilgrim

explains, Ma refers to the relational space that gives form and meaning to elements (1986: 255). Applied methodologically, paradigmatic stability resides not only in assumptions-as-nodes but in the relations that bind them.

Standard problematization often challenges assumptions individually. A Ma-informed approach targets the relational structure that renders those assumptions mutually necessary. The mechanistic paradigm persists not because each assumption is independently compelling, but because their interconnections generate collective plausibility and self-protection.

Feedback Loops: How the Network Locks In

The network operates through self-reinforcing feedback loops – what Section 2.3.5 identified as the Triple Lock. Restated in network terms:

Loop 1: Ontological reinforcement Reductionism presupposes atomism; atomism supports substance ontology; substance ontology enables determinism; determinism requires closure; closure renders exhaustive decomposition plausible – reinforcing reductionism.

Loop 2: Epistemological reinforcement Certainty supports universal laws; universal laws support monism; monism supports perfect information; perfect information reinforces certainty.

Loop 3: Cross-level reinforcement Ontological commitments generate epistemological ones (e.g., determinism → certainty); epistemological commitments justify methodological expectations (e.g., monistic modeling, point prediction); methodological practice retroactively validates ontological premises. Ontology, epistemology, and methodology thus authorize one another in a closed circuit.

The mechanistic paradigm's resilience is therefore structural. It is networked, not individually invulnerable.

Part B: WHY Network Treatment Is Necessary

The Failure of Piecemeal Problematization

When assumptions reinforce one another, isolated challenges are predictably neutralized. Attacking equilibrium while leaving closure and linearity intact allows stability to reappear under alternative labels. The network absorbs the shock because its foundations remain untouched.

Newell's critique of "divide and conquer" applies directly: isolating components leaves interactions – the true drivers of system behavior – unexamined (2001: 13). Here, those interactions are the source of paradigmatic persistence.

The Tier-Structure Constraint

Because Tier-2 assumptions are generated by Tier-1 commitments, reform confined to derivatives is structurally unstable. Either the derivatives regenerate, or incoherence emerges. The methodological implication is unavoidable: intervention must target the hard core, and it must do so in a coordinated manner that respects interdependence.

Ma: The Target Is the "Between"

The Ma-inspired insight clarifies the real target of transformation. Paradigms persist through the relational "between" – the intervals that render assumptions mutually necessary. Challenging nodes while leaving relations intact leaves the paradigm functionally unchanged. This dissertation therefore treats relational structure as constitutive, not incidental.

Part C: HOW Network Treatment Is Implemented

Chapter 3: Mapping the Network

Network treatment is operationalized first through explicit mapping. Chapter 3 presents the assumptions not as an inventory but as a generated structure:

1. Identification of Tier-1 ontological and epistemological commitments.
2. Derivation of Tier-2 assumptions as consequences of Tier-1 combinations.
3. Explicit articulation of interdependence via reinforcing loops (the Triple Lock).

This mapping renders visible what usually remains implicit – what Gouldner (1970) described as "unpostulated and unlabeled" presuppositions operating in the background.

Chapter 4: Coordinated Reconstruction of the Hard Core

Under the compressed architecture, reconstruction proceeds in coordinated fashion:

- Ontological commitments (reductionism, atomism, closure, substance, determinism) are displaced together rather than sequentially.
- Epistemological and methodological implications are derived explicitly from the reconstructed ontology, preventing silent reabsorption under old evaluative standards.

Derivative assumptions are not attacked directly. They are displaced structurally. Once the hard core is reconstructed, equilibrium, linearity, ergodicity, optimisation, and decomposability lose their generative base and therefore their default status.

Conclusion: The Network as the Unit of Analysis

This section has shown that the mechanistic paradigm's fifteen core assumptions form an interconnected network whose resilience derives from mutual reinforcement rather than from isolated defensibility. Accordingly, the dissertation treats the assumptive network as the unit of analysis.

Chapter 3 maps this network explicitly; Chapter 4 reconstructs the hard core in coordinated fashion. Derivative commitments are displaced not through piecemeal critique but as structural consequences of transforming the network's foundations.

2.7.4 Structural Conventions

Introduction: From Architecture to Convention

Sections 2.7.1–2.7.3 established the dissertation's compressed architecture, its multi-level cascade logic, and the assumptive network as its unit of analysis. This section specifies the structural conventions that complement that architecture: repeatable formats that render a complex argument navigable, auditable, and cumulatively coherent.

These conventions are not stylistic ornaments. Given the dissertation's scale (fifteen assumptions across three levels, with cross-level feedback), coherence cannot be left to local rhetorical competence; it requires systematic scaffolding.

In complexity-informed inquiry, "method" cannot be reduced to standardized protocols. As Alhadeff-Jones notes, method here differs from "methodological" understood as a ready-made procedure; the distinction is "precisely one of paradigm" (2013: 21, citing Morin). The conventions below are therefore best understood as methodological infrastructure: they instantiate commitments to interconnection, emergence, and coordinated transformation.

Part A: WHAT Are the Structural Conventions?

(1) The Problematization → Reconstruction Pattern (FROM → Diagnosis → TO → Integration)

Within Chapter 4, each substantive transformation follows a repeatable four-stage sequence:

Table 9: Four-Stage Assumption Transformation Protocol: Structure and Illustration Through the Case of Emergence

Stage	Function	Illustration (Emergence)
FROM	Identify the mechanistic commitment	Reductionism: wholes are exhaustively decomposable
Diagnosis	Show why it fails for firms-as-CAS	System-level properties are not inferable from components
TO	Develop the CAS-compatible alternative	Emergence: novel system properties arise from interaction
Integration	Bind the transformation into the wider shift	Links to connectivity, grounds uncertainty, constrains method

This pattern operationalizes problematization by enforcing four obligations: (i) state precisely what is challenged, (ii) establish why it fails under the dissertation's object-construal, (iii) provide a constructive replacement, and (iv) show how the replacement coheres with adjacent transformations. Its function is not pedagogical repetition but anti-fragmentation discipline: a formal safeguard against piecemeal critique producing an incoherent alternative.

(2) Consistent Section Anatomy

Each major section also follows a stable internal anatomy:

Table 10: Standardised Section Architecture: Structural Elements and Their Functions for Ensuring Argumentative Coherence

Element	Function
Dependency map	Identifies "Receives From / Provides To / Feedback With" links
Intro	States the section's job in one paragraph; previews steps
Development	Builds the argument through explicit inferential steps
Counter-arguments	Engages strongest objections (not straw men)
Synthesis	States what is secured and what remains open
Forward connection	Specifies exactly what the section enables next

This format provides predictable landmarks: readers can locate the section's task, track its inferential progression, identify its engagement with objections, and see precisely what is carried forward. It also supports non-linear examination: sections can be audited without requiring linear reading.

(3) The "Receives From / Provides To" Tracking System

A distinctive methodological contribution of the dissertation is the explicit tracking of interdependencies between sections using standardised "Receives From / Provides To" maps. This tracking system makes the dissertation's internal architecture visible and auditable.

Each major section specifies three types of connections:

Table 11: Inter-Section Dependency Link Types: Functions and Illustrations of the Dissertation's Argumentative Connectivity Structure

Connection Type	Function	Illustration
Receives from	Inputs required from earlier sections	Section 4.2.1 receives the diagnosis of reductionism from Section 3.1.1
Provides to	Conceptual contributions to later sections	Section 3.1.2 provides the certainty diagnosis for Section 4.3.1
Feedback with	Bidirectional illumination	Section 3.1.1 ↔ Section 4.2.7 (Diagnosis ↔ Emergence architecture)

The system specifies what flows between sections: diagnoses, transformations, conceptual foundations, or validation criteria. This makes the dissertation's argumentative structure visible as an integrated system rather than a linear sequence.

Why Explicit Dependency Tracking Is Necessary

Paradigmatic assumptions exhibit complex causality: they reinforce one another through feedback loops rather than linear chains. Transforming one element without tracking its dependencies risks incoherence or silent regeneration of the commitments being challenged.

The tracking system addresses this by ensuring:

- **Completeness:** Every transformation is grounded in a prior diagnosis; every FROM assumption has a corresponding TO alternative.
- **Coherence:** Early diagnoses constrain later reconstructions, preventing ad hoc departures from the systematic architecture.
- **Bidirectional Validation:** Transformations illuminate diagnoses just as diagnoses ground transformations; the tracking system makes this bidirectional logic explicit.
- **Navigation:** Readers can trace conceptual flows through the architecture, understanding how conclusions depend on earlier arguments.

As Alhadeff-Jones emphasises, methodological rigor in complex research depends on the capacity to "identify and locate one's own research trajectory so that it can be understood by others, shared, communicated and to some extent reproduced" (2013: 41). The tracking system provides precisely this visibility.

(4) Terminological Conventions

Because paradigms often reuse identical vocabulary with different meanings, the dissertation enforces strict terminological discipline. The purpose is not pedantry but incommensurability management: preventing mechanistic meanings from silently colonizing CAS concepts.

Table 12: Key Terminological Conventions: Stipulated Definitions and Their Justifications for Maintaining Analytical Precision

Term	Defined meaning (in this dissertation)	Why the convention matters
"firms" (plural)	irreducibly heterogeneous entities	blocks representative-firm slippage
"complex nature of firms"	firms as CAS requiring multi-dimensional explanation	blocks single-discipline framing
"transformation"	paradigm-level change in core commitments	distinguishes from parameter-tuning
"emergence"	irreducible system properties generated by interaction	technical term, not metaphor
"complexity" vs "complication"	systemic nonlinearity vs mere difficulty	prevents rhetorical inflation

These conventions function as a conceptual integrity system: they secure internal coherence and reduce equivocation in contested terms.

(5) Crystallization Devices

To prevent cumulative argumentation from becoming cognitive overload, the dissertation uses explicit crystallization devices:

- **Synthesis points** (where multiple threads converge and commitments are stated cleanly)
- **Key insight boxes** (to mark propositions reused later)
- **Visual dependency maps** (to display interconnections obscured by prose)
- **Summary tables** (to provide short audit trails at the end of dense sections)

These devices serve a second methodological purpose: they mark saturation points – where further elaboration would be redundancy – so that completion is legible and drift is minimized.

Part B: WHY Are These Conventions Appropriate?

Coherence at scale

A dissertation transforming fifteen assumptions across three levels will fail if coherence is treated as an emergent by-product of writing style. It must be produced by design. The conventions provide stable invariants – repeatable forms – allowing the content to vary without loss of structural control.

Navigation through complexity

The dissertation advances ontological, epistemological, and methodological claims in a coordinated cascade. Conventions reduce "search cost": each section explicitly declares inputs, outputs, objections, and forward implications. This is how a networked argument becomes readable without being simplified.

Terminological precision under incommensurability

Because mechanistic and CAS frameworks share vocabulary while diverging in meaning, disciplined terminology is not optional. It is the operational solution to incommensurability: it prevents interpretive drift and preserves the technical force of concepts such as emergence, openness, non-ergodicity, and uncertainty.

Pattern recognition

A repeated FROM→TO grammar enables pattern recognition across transformations: the reader learns the argumentative syntax once and can focus attention on substantive differences. This is a cognitive design feature, not a stylistic flourish, and it strengthens evaluation by making comparison across transformations straightforward.

Part C: HOW Are These Conventions Implemented?

Implementation within Chapter 4: Chapter 4 (Ontology) applies the FROM→Diagnosis→TO→Integration pattern across the five ontological commitments and links each shift explicitly to its epistemological implications.

Section anatomy as audit infrastructure

Every substantive section begins with dependency mapping and ends with synthesis plus forward connection. In effect, each section declares:

1. what it required,
2. what it established,
3. what it enables, and
4. what it cannot yet settle.

This converts the dissertation from a narrative flow into an auditable research trajectory.

Crystallization placement

Crystallization devices are placed at predictable "load points":

- end of each transformation cluster (Chapter 4 transformation sets), and
- the conclusion where implications are consolidated.

Conclusion: Conventions as Methodological Infrastructure

This section specified the structural conventions that function as methodological infrastructure for paradigm-level theory: the FROM→Diagnosis→TO→Integration transformation pattern, consistent section anatomy, the "Receives From / Provides To" tracking system, terminological discipline, and crystallization devices. Together they enable coherence at scale, support navigation through complexity, defend conceptual integrity under incommensurability, and facilitate pattern recognition. In short, they make a paradigm shift readable, traceable, and examinable.

Section 2.8 now turns from design and structure to the practical tools and techniques by which the inquiry was conducted: how transdisciplinary material was selected, integrated, and quality-controlled.

Conclusion of Section 2.7

Section 2.7 has specified the research design and architecture through which systematic paradigm transformation is executed. Four findings carry forward.

First, the dissertation implements systematic paradigm transformation through a compressed four-chapter architecture. This design follows necessarily from the nature of paradigm-level intervention: because the mechanistic paradigm's assumptions form a mutually reinforcing network (the Triple Lock), piecemeal modification fails. The architecture proceeds through two phases – diagnostic mapping (Chapter 3) and coordinated reconstruction (Chapter 4) – that together address the assumption network as a whole rather than attempting isolated repairs.

Second, the multi-level cascade structure ensures philosophical coherence. Ontological transformation (from substance to process ontology) generates epistemological implications (from deterministic certainty to fundamental uncertainty), which generate methodological

implications (from monism to structured pluralism). This ordering is not discretionary but philosophically necessary: what exists constrains what can be known, which constrains how inquiry can proceed. Bidirectionality reinforces coherence – methodological failure illuminates epistemic overreach, which reveals ontological mis-specification.

Third, the assumptive network constitutes the unit of analysis. Fifteen core assumptions are organised hierarchically into a hard core (nine foundational ontological and epistemological commitments) and a protective belt (six derivative assumptions). The Ma-inspired insight directs attention to the relational structure – the "between" – that renders assumptions mutually necessary. Chapter 3 maps this network explicitly; Chapter 4 reconstructs the hard core in coordinated fashion, displacing derivatives structurally rather than through piecemeal critique.

Fourth, structural conventions – the FROM→Diagnosis→TO→Integration pattern, consistent section anatomy, the "Receives From / Provides To" tracking system, terminological discipline, and crystallization devices – function as methodological infrastructure. They enable coherence at scale, support navigation through complexity, and defend conceptual integrity under conditions of incommensurability.

Section 2.8 specifies the practical tools and techniques through which the inquiry was conducted.

2.8 Tools and Techniques

Section 2.7 specified the research design and architecture through which systematic paradigm transformation is executed – the four-chapter structure, multi-level cascade logic, assumptive network as unit of analysis, and structural conventions enabling coherence at scale. That section established what the dissertation does and why its architecture takes the form it does. This section turns from design to execution, specifying the practical tools and techniques through which the inquiry was conducted.

This specification is necessary because paradigm-level theoretical work, while not reliant on empirical data collection, nevertheless requires explicit methodological discipline. Without such discipline, problematization risks collapsing into impressionistic critique or undisciplined eclecticism.

The argument proceeds through ten subsections. Subsection 2.8.1 establishes the transdisciplinary orientation required by both the nature of complex systems and this dissertation's paradigm-level intervention. Subsection 2.8.2 specifies the problematizing literature review methodology. Subsection 2.8.3 identifies source categories. Subsection 2.8.4 articulates six philosophical argumentation methods. Subsection 2.8.5 explains assumption identification and mapping. Subsection 2.8.6 addresses cross-disciplinary translation. Subsections 2.8.7–2.8.10 develop historical-theoretical analysis, analytical procedures, visual methods, and firm-level illustration respectively.

Section 2.9 addresses reflexivity, scope conditions, and calibration of claims.

2.8.1 Transdisciplinary Orientation

Introduction

The dissertation's paradigm-level intervention requires conceptual resources that no single discipline commands. Reconceptualizing the firm as a complex adaptive system engages complexity science, philosophy of science, economic methodology, organisation theory, and systems thinking – not as parallel contributions to be aggregated, but as complementary perspectives on a problem whose structure exceeds disciplinary boundaries. This section specifies the transdisciplinary orientation adopted, distinguishes it from cognate approaches, articulates the principles that guard against undisciplined eclecticism, and shows how the orientation was applied across the dissertation's architecture.

The methodological challenge is precise. Paradigm-level assumptions are not confined to economics; they reflect broader intellectual inheritances concerning being, causation, and knowledge. As Lawson (2003, p. 21) argues, the persistent methodological difficulties of mainstream economics are symptomatic of ontological commitments – closure, atomism, event-regularity – that require philosophical diagnosis unavailable within economics itself. Addressing such commitments demands transdisciplinary engagement as a methodological necessity, not a stylistic preference.

Part A: Defining Transdisciplinarity

Distinctions among disciplinary modes

Transdisciplinarity must be distinguished from multidisciplinary and interdisciplinary, which name different modes of cross-disciplinary engagement.

Multidisciplinary involves the parallel deployment of multiple disciplines, each contributing within its own framework without substantive integration. The disciplines remain "side by side," addressing a common topic from separate vantage points (Klein, 2010, p. 16). Outputs are additive rather than synthetic.

Interdisciplinarity involves integration across disciplines within a shared problem-frame. As Klein (2010, p. 17) specifies, interdisciplinary work "integrates information, data, techniques, tools, perspectives, concepts, and/or theories from two or more disciplines." Integration occurs, but within boundaries defined by the contributing disciplines.

Transdisciplinarity, by contrast, addresses problems whose structure cannot be adequately framed within any single discipline or combination of disciplines. Nicolescu (2002, p. 44) defines it as concerning "that which is at once between the disciplines, across the different disciplines, and beyond all discipline." The prefix *trans-* indicates movement through and beyond disciplinary boundaries toward a form of inquiry that the disciplines themselves cannot supply.

Max-Neef (2005, p. 15) sharpens the distinction by differentiating "weak" transdisciplinarity – coordination among disciplines without deep integration – from "strong" transdisciplinarity, which seeks coherent synthesis at the level of foundational commitments. Strong transdisciplinarity does not merely borrow concepts but reconstructs problems at a level of generality that reconfigures what the contributing disciplines can see.

Transdisciplinarity and complex phenomena

The connection between transdisciplinarity and complexity is not accidental. As Morin (2008, p. 6) argues, complex phenomena resist the "disjunction and reduction" characteristic of disciplinary specialisation; understanding them requires what he terms *pensée complexe* – a mode of thought that holds together what disciplines separate. Cilliers (2005, p. 258) makes the epistemological point explicit: because complex systems cannot be exhaustively described from any single framework, knowledge of them is necessarily perspectival and partial. No discipline possesses privileged access; each illuminates aspects while occluding others.

This has direct methodological implications. As Cilliers (2005, p. 259) states, "we cannot have complete knowledge of complex systems; we can only have knowledge in terms of a certain framework... We choose our frameworks." Transdisciplinary engagement is therefore not an aspiration to completeness but a recognition that framework choice is consequential and that alternative frameworks reveal what dominant ones suppress.

Part B: Why Transdisciplinarity Is Required

Two orders of justification establish the necessity of transdisciplinary engagement: the inherent demands of complex systems as objects of inquiry, and the specific requirements of this dissertation's paradigm-level intervention.

B1. Complex systems and the limits of disciplinary inquiry

Complex adaptive systems pose distinctive challenges that disciplinary organisation is structurally ill-equipped to address. Four features of CAS generate this requirement.

First, complex systems are constituted by interactions that cross disciplinary boundaries. A firm's behaviour emerges from the interaction of economic incentives, cognitive processes, social relationships, institutional constraints, and technological affordances – phenomena distributed across economics, psychology, sociology, political science, and engineering. As Morin (2008, p. 6) argues, disciplinary specialisation operates through "disjunction and reduction," isolating components for tractable analysis. But CAS properties arise precisely from interactions that such isolation eliminates. Emergence, by definition, concerns properties of wholes not derivable from properties of parts examined separately (Elder-Vass, 2010, p. 16–23). Understanding emergence therefore requires analytical frameworks capable of holding together what disciplines separate.

Second, complex systems operate across multiple scales simultaneously, and different disciplines have developed expertise at different scales. Economics addresses exchange and allocation; organisation theory addresses routines and capabilities; sociology addresses institutions and networks; psychology addresses cognition and decision-making. As Dopfer, Foster, and Potts (2004, p. 264) argue, the meso-level – where firms operate – is precisely the scale at which disciplinary frameworks lose purchase, falling between microeconomic individualism and macroeconomic aggregation. Understanding meso-level dynamics requires integration across scales that no single discipline commands.

Third, complex systems exhibit circular causality that connects domains disciplines treat as separate. Sterman (2000, p. 12) emphasises that feedback loops – the signature causal structure of complex systems – link cognition to action, action to institutional response, institutional response to environmental change, and environmental change back to cognition. These loops cross the boundaries separating psychology from economics from sociology. Disciplines that presuppose linear, unidirectional causation cannot represent such circularity without distortion. As Meadows (2008, p. 188) observes, leverage points in complex systems often lie at the level of paradigms and goals – precisely where disciplinary assumptions operate most invisibly.

Fourth, the boundaries of complex systems are themselves framework-dependent. Cilliers (2005, p. 610) argues that "the boundaries of a complex system are not given but are constructs that we create in order to understand the system." What counts as "the firm" – its boundaries, components, and environment – depends on the conceptual framework through which it is approached. Different disciplines construct different objects: the firm as production function, governance structure, capability bundle, or network node. No construction is uniquely correct; each reveals and conceals. Understanding this framework-dependence – and its consequences for what can and cannot be seen – requires meta-theoretical resources that disciplines typically lack.

These four features – boundary-crossing interaction, multi-scale operation, circular causality, and framework-dependent boundaries – establish that complex systems are not merely difficult objects for disciplinary inquiry but categorically different objects requiring a different mode of inquiry. As Mitchell (2009, p. 3) concludes, understanding complex systems demands "integrative pluralism": the disciplined combination of multiple frameworks without premature reduction to any one.

B2. The specific requirements of paradigm-level reconstruction

Beyond the general demands of complex systems, three considerations specific to this dissertation establish transdisciplinary engagement as methodologically necessary.

First, the mechanistic paradigm is itself transdisciplinary in origin. As Mirowski (1989, p. 193–232) documents, neoclassical economics imported its core conceptual apparatus – equilibrium, optimisation, conservation principles – from nineteenth-century physics. The importation was structural, not merely metaphorical: utility was modelled on potential energy, and constrained optimisation was transferred wholesale from analytical mechanics. Challenging these assumptions requires engagement with their extra-economic sources. Philosophy of science supplies the meta-theoretical vocabulary for analysing paradigm structure (Kuhn, 1962; Lakatos, 1970). Process philosophy supplies alternatives to the substance metaphysics underwriting atomism and closure (Whitehead, 1929; Rescher, 1996). Complexity science supplies ontological resources for emergence, adaptation, and non-linearity (Holland, 2014; Kauffman, 1993). No single discipline provides the conceptual resources necessary for paradigm-level reconstruction.

Second, the assumptions targeted operate at a depth exceeding disciplinary reach. The ontological and epistemological commitments problematized in this dissertation – substance metaphysics, determinism, closure, the epistemic fallacy – are not peculiar to economics but constitute broader intellectual inheritances. Bizzarri (2021, p. 70) traces how the Greek concept of *physis* as developmental process was displaced by *physis* as originating cause – a metaphysical shift privileging being over becoming. Surfacing such commitments requires philosophical resources. As Bhaskar (1975, p. 25) argues, sciences typically operate with implicit ontologies that remain unexamined; making these explicit requires meta-theoretical analysis that disciplines themselves cannot supply.

Third, reconstructing the firm as a complex adaptive system requires positive conceptual resources that economics lacks. Complexity science supplies the vocabulary of emergence, self-organisation, fitness landscapes, and co-evolution. Organisation theory supplies meso-level constructs – routines, capabilities, dynamic capabilities – that articulate how CAS properties manifest in firms specifically. Systems thinking supplies feedback and leverage-point analysis. These resources do not exist within mainstream economics; they must be imported, translated, and integrated.

Part C: Principles Governing Transdisciplinary Engagement

Transdisciplinary work risks two failure modes. Superficial borrowing extracts concepts without understanding their home-discipline conditions, producing terminological novelty without substantive gain. Undisciplined eclecticism aggregates incompatible frameworks, generating incoherence masked by apparent breadth. The following principles guard against these risks.

Principle 1: Problem-driven selectivity

Disciplines are engaged insofar as they illuminate the specific problem under examination. Klein (2004, p. 516) emphasises that transdisciplinary integration must be problem-focused rather than discipline-focused: "The problem, not the discipline, is the point of departure." This prevents the "everything is relevant" trap that undermines analytical focus.

In this dissertation, problem-driven selectivity means engaging complexity science for its ontological vocabulary (emergence, adaptation, self-organisation), philosophy of science for its meta-theoretical frameworks (paradigms, research programmes, critical realism), and economic methodology for its discipline-specific articulation of the limits of mechanistic assumptions. Sources are selected for assumption-leverage, not comprehensiveness.

Principle 2: Contextual translation

Concepts transported across disciplinary boundaries must be interpreted within their original contexts before deployment. Bergmann et al. (2012, p. 41) stress that transdisciplinary research requires "mutual learning" that respects the integrity of contributing frameworks. Failure to do so produces what Star and Griesemer (1989, p. 393) call "boundary objects" – concepts that appear shared but carry incompatible meanings, enabling superficial communication while obscuring substantive disagreement.

Key terms in this dissertation – emergence, complexity, uncertainty, system – carry different meanings across physics, biology, economics, and philosophy. These differences are made explicit. "Emergence" in philosophy of mind (concerning consciousness) differs from "emergence" in complexity science (concerning system-level properties arising from interaction); both differ from loose invocations of "emergent phenomena" in management discourse. Contextual translation ensures that borrowed concepts retain their technical force.

Principle 3: Compatible pluralism

Integration proceeds under a compatibility constraint. As Mitchell (2009, p. 109) argues, integrative pluralism does not require reduction to a single framework but does require that combined frameworks not contradict one another at the level of foundational commitments. Incompatible frameworks cannot be integrated; they can only be juxtaposed.

The compatibility constraint operates in this dissertation at the level of ontological and epistemological commitments. Sources are combined only where they share – or are consistent with – commitments to ontological realism, epistemic fallibilism, and the causal efficacy of emergent properties. Sources presupposing strong social constructivism, eliminative reductionism, or deterministic closure are engaged critically rather than integrated.

Principle 4: Reflexive positioning

The dissertation's own standpoint is rendered explicit. As Alhadeff-Jones (2013, p. 32) argues, complexity-oriented inquiry must acknowledge the researcher's implication within the field of study. The aspiration to a "view from nowhere" is not merely unattainable but methodologically distorting; it conceals the framework-dependence of observation.

Reflexive positioning means specifying the commitments that guide source selection, interpretation, and integration. This dissertation operates from critical realist ontology, hermeneutic epistemology, and a problematization methodology; these commitments shape which transdisciplinary resources appear relevant and how they are deployed. Reflexivity does not eliminate perspective but makes it available for critique.

Part D: Application in This Dissertation

The transdisciplinary orientation is not merely declared but operationalized through specific deployments across the dissertation's architecture.

D1. Domain contributions to diagnostic work (Chapter 3)

Chapter 3's diagnostic task – surfacing and mapping the mechanistic assumption network – draws primarily on philosophy of science and economic methodology.

Philosophy of science supplies the analytical vocabulary for identifying paradigm-level assumptions. Kuhn's (1962) concept of paradigms as "disciplinary matrices" comprising shared commitments, values, and exemplars enables the dissertation to treat mechanistic economics not as a set of testable hypotheses but as an interconnected framework of enabling assumptions. Lakatos's (1970) distinction between hard core and protective belt structures the hierarchical organisation of assumptions into foundational (Tier 1) and derivative (Tier 2) commitments. Bhaskar's (1975) critical realist ontology – particularly the distinction among empirical, actual, and real domains – grounds the identification of ontological assumptions concerning stratification, emergence, and causal powers.

Economic methodology supplies discipline-specific articulation of mechanistic commitments. Knight's (1921) distinction between risk and uncertainty identifies a foundational epistemological assumption (that all relevant contingencies are enumerable). Hayek's (1967, 1974) analysis of complexity and pattern prediction exposes the prediction-imperative driving physics-envy. Lawson's (1997, 2003) diagnosis of mainstream economics' "closed-system" methods reveals the ontological commitment to closure underlying formal modelling. Machlup's (1955) hierarchy of hypotheses clarifies why foundational assumptions cannot be directly tested. These resources enable the dissertation to surface assumptions that remain invisible within mainstream discourse.

The diagnostic deployment is precise: philosophy of science supplies the meta-theoretical framework; economic methodology supplies the discipline-specific content. Neither alone would suffice. Philosophy of science without economic methodology would produce generic paradigm analysis lacking substantive engagement with economics. Economic methodology without philosophy of science would lack the analytical vocabulary to distinguish assumption types and levels.

D2. Domain contributions to reconstructive work (Chapter 4)

Chapter 4's reconstructive task – articulating an alternative assumption-ground adequate to firms as CAS – draws primarily on complexity science and organisation theory, with systems thinking providing complementary resources.

Complexity science supplies the ontological vocabulary for the reconstructed framework. Holland's (1995, 2014) characterisation of CAS as systems composed of adaptive agents whose interactions generate emergent properties provides the core definition. Kauffman's (1993, 1995) analysis of self-organisation and fitness landscapes establishes that order can arise without central design – directly challenging the optimisation assumption. Arthur's (2015) work on increasing returns and path-dependence establishes that equilibrium is neither necessary nor typical in complex economic systems. Cilliers's (1998, 2005) philosophical

analysis of complexity – particularly regarding boundaries, models, and the limits of knowledge – connects complexity science to the epistemological reconstruction.

Organisation theory supplies meso-level constructs that translate CAS properties into firm-specific terms. Nelson and Winter's (1982) evolutionary theory of the firm reconceptualizes firms as repositories of routines subject to selection and modification – an explicitly process-based and path-dependent conception. Teece, Pisano, and Shuen's (1997) dynamic capabilities framework addresses how firms adapt under conditions of rapid change, articulating emergence at the capability level. March's (1991) analysis of exploration and exploitation captures the fundamental tension in adaptive systems between refinement and search. Simon's (1962) analysis of hierarchy and near-decomposability provides resources for understanding firm structure without presupposing full decomposability.

Systems thinking supplies feedback and interdependence analysis. Sterman's (2000) system dynamics framework provides tools for representing circular causality and accumulation processes. Meadows's (2008) leverage-point analysis identifies where interventions in complex systems are likely to be effective – and where they are likely to be absorbed. Ackoff's (1971) critique of machine-age thinking articulates the systems alternative to mechanistic assumptions at a foundational level.

The reconstructive deployment is integrative: complexity science provides the general ontological framework; organisation theory translates this framework to the meso-level of firms; systems thinking supplies analytical tools for representing feedback and interdependence. The integration is constrained by compatibility requirements (Section 2.8.6): all three domains share commitments to emergence, non-linearity, and the inadequacy of reductionism, enabling genuine synthesis rather than mere juxtaposition.

D3. Cross-chapter integration

The transdisciplinary architecture ensures that diagnostic and reconstructive work are mutually informing rather than merely sequential.

Philosophy of science resources deployed diagnostically in Chapter 3 (paradigm analysis, research programme structure) reappear in Chapter 4 as evaluative criteria. The Lakatosian framework used to analyse the mechanistic paradigm's hard core is reflexively applied to assess whether the CAS alternative constitutes a progressive programme (Section 2.6.3). Critical realist ontology grounds both the diagnosis of mechanistic ontological assumptions and the articulation of alternatives.

Complexity science concepts anticipated in Chapter 3's diagnosis – emergence, non-linearity, path-dependence – are fully developed in Chapter 4's reconstruction. The diagnostic identification of mechanistic assumptions as incompatible with these properties motivates their reconstructive articulation. The dissertation thus exhibits the bidirectional relationship between diagnosis and reconstruction that Section 2.7.2 specified: diagnosis reveals what requires reconstruction; reconstruction supplies resources that retrospectively illuminate what diagnosis found.

Economic methodology and organisation theory provide continuous meso-level grounding. The Knight-Keynes-Hayek tradition's treatment of uncertainty informs both the diagnostic identification of the certainty assumption and the reconstructive articulation of fundamental uncertainty. Organisation theory's attention to routines and capabilities informs both the

diagnostic identification of what mechanistic economics neglects and the reconstructive articulation of firm-level CAS properties.

D4. Transdisciplinary quality control

Section 2.8.6 specifies procedures for managing cross-disciplinary conflicts; these procedures operationalize the transdisciplinary principles articulated above.

Contextual translation is enforced through explicit definition of key terms at first use, distinguishing discipline-specific meanings. "Emergence" is defined in its complexity-science sense (properties of wholes arising from interaction, not possessed by parts, and exerting downward causal influence) rather than its colloquial sense (anything that "emerges" over time). "Uncertainty" is defined following Knight (1921) as unmeasurable (lacking enumerable outcome space) rather than as mere variance.

Compatible pluralism is enforced through the minimal compatibility constraint: sources are integrated only where they share commitments to ontological realism, epistemic fallibility, and the causal efficacy of emergent properties. Sources presupposing eliminative reductionism or strong social constructivism are engaged critically rather than integrated.

Problem-driven selectivity is enforced through the assumption-leverage criterion: sources are selected for their contribution to surfacing, diagnosing, or reconstructing specific assumptions, not for comprehensive coverage of contributing disciplines. The dissertation does not attempt a survey of complexity science, philosophy of science, or organisation theory; it deploys these disciplines instrumentally for paradigm-level work.

Conclusion

Transdisciplinarity is a methodological requirement imposed by the dissertation's object and purpose. Complex adaptive systems demand transdisciplinary understanding because their constitutive features – boundary-crossing interaction, multi-scale operation, circular causality, and framework-dependent boundaries – exceed what any single discipline can address. This dissertation's paradigm-level intervention intensifies the requirement: the mechanistic assumptions targeted are transdisciplinary in origin, operate at a depth exceeding disciplinary reach, and can only be replaced with conceptual resources imported from complexity science, organisation theory, and systems thinking.

The transdisciplinary orientation adopted here is problem-driven (selecting sources for assumption-leverage), contextually sensitive (respecting source-discipline meanings), compatibility-constrained (integrating only non-contradictory frameworks), and reflexively positioned (acknowledging the standpoint from which integration proceeds). These principles guard against superficial borrowing and undisciplined eclecticism while enabling the cross-disciplinary engagement that paradigm-level reconstruction demands.

The orientation is operationalized through structured deployment: philosophy of science and economic methodology ground the diagnostic work of Chapter 3; complexity science, organisation theory, and systems thinking ground the reconstructive work of Chapter 4; and cross-chapter integration ensures that diagnosis and reconstruction remain mutually informing. Section 2.8.2 specifies the literature review methodology through which this transdisciplinary engagement is conducted.

2.8.2 Literature Review Methodology

Introduction

The dissertation's transdisciplinary orientation requires a literature review methodology capable of crossing disciplinary boundaries without collapsing into eclecticism. This section specifies the adopted approach, explains why it was selected over alternative review traditions, and shows how it was operationalized.

The dissertation adopts the **problematizing review** developed by Alvesson and Sandberg (2020). Unlike systematic or conventional narrative reviews, the problematizing review targets **assumptions underlying influential literatures** rather than aggregating findings within an accepted framework. This aligns directly with the dissertation's problematization methodology (Section 2.3) and operationalizes it at the level of literature engagement.

As Alvesson and Sandberg (2011, p. 247) argue, influential theory becomes “interesting” by challenging audience assumptions rather than refining them. Accordingly, the function of the literature review here is not inventory but **assumption interrogation** – creating conceptual space for an alternative assumption-ground adequate to firms as CAS.

Part A: Defining the Problematizing Review

A1. Core principles

Alvesson and Sandberg (2020, p. 1290) define the problematizing review through four principles, each of which structured this dissertation.

(1) Reflexivity.

The review scrutinizes assumptions in both the target literature and the reviewer's own standpoint. This reflects the hermeneutic position that understanding proceeds from pre-understanding (Section 2.2.2). Reflexivity is therefore bidirectional: mechanistic assumptions are interrogated, but complexity-informed commitments are also rendered explicit and critically governed. This aligns with Alhadeff-Jones's insistence that critical inquiry must acknowledge the researcher's implication rather than claim standpoint neutrality (2013, p. 32).

(2) Reading broadly but selectively.

Because the unit of transformation is the **assumptive network** (Sections 2.7.3–2.7.4), literature engagement must cross disciplinary boundaries while remaining principled in selection. Sources were chosen for **assumption illumination**, not exhaustive coverage. Newell's observation that disciplines tend to “find what they are looking for” (2010, p. 334) underscores why selective cross-reading is essential for exposing naturalized commitments.

(3) Problematizing rather than accumulating.

The review reverses the default orientation of gap-spotting. Instead of asking what the literature concludes, it asks what must be presupposed for those conclusions to be intelligible – and what those presuppositions exclude. As Alvesson and Sandberg (2011, p. 252) argue, gap-spotting underproblematizes and thereby reinforces dominant frameworks. Paradigm-level transformation requires rendering foundations contestable.

(4) “Less is more.”

This principle prioritizes analytical leverage over bibliographic volume. Depth of engagement matters more than encyclopedic coverage. Akerlof’s (2020) analysis of incentive-driven “sins of omission” in economics provides contextual justification: phenomena difficult to formalize are often neglected. The review therefore prioritizes sources that expose omissions – particularly regarding emergence, nonlinearity, and path dependence – rather than maximizing citation counts.

A2. Why systematic and integrative reviews were misaligned

Systematic and integrative reviews are designed to consolidate findings within a shared framework. That logic is appropriate when constructs, measures, and comparability conditions are broadly agreed. It is misaligned when the framework itself is the object of critique.

In this dissertation, the core object is not commensurable findings but a paradigm-level assumptive structure (Chapters 3–4). A methodology designed to aggregate within a framework cannot evaluate the adequacy of that framework. Hence, systematic criteria (exhaustive search strings, strict inclusion rules, commensurable operationalizations) do not target the dissertation’s epistemic task: transforming the assumptive architecture itself.

Part B: Why the Problematizing Review Was Necessary

B1. Paradigm-level target.

The dissertation targets hard-core mechanistic commitments structuring mainstream firm theory (Chapter 3). Alvesson and Sandberg (2011, p. 257) note that paradigm- and field-level assumptions yield the highest theoretical impact but are the most difficult to challenge precisely because they are naturalized.

B2. Transformation requires opening up.

Following Davis’s (1971) account of interestingness as denial of taken-for-granted propositions, reconceptualizing firms as CAS requires de-naturalization rather than incremental extension. The review therefore functions to *open* conceptual space rather than to *build* within existing premises.

B3. Transdisciplinary scope.

Because the dissertation integrates complexity science, philosophy of science, heterodox economics, organization theory, and systems thinking, no single-domain systematic protocol could be applied without trivializing the inquiry or sacrificing analytical clarity. Newell’s emphasis on nonlinear interactions across subsystems (2010, p. 340) reinforces the need for a review approach capable of treating “between-space” relations as analytically central.

Part C: How the Problematizing Review Was Operationalized

C1. Procedures

1. Assumption surfacing.

The review began with iterative identification of assumptions structuring mainstream treatments of firms. Candidate assumptions were refined through repeated reading and cross-disciplinary comparison to distinguish foundational commitments from local conventions.

2. Dialectical interrogation.

Each major source was read through structured interrogation (Alvesson & Sandberg, 2011, p. 255):

- What must be true for this argument to work?
- What does this framing exclude or render anomalous?
- Which rival assumptions would reconfigure what counts as problem, cause, or explanation?

3. Transdisciplinary extension.

When mechanistic literature became the constraint rather than the resource, the review extended to bodies of work articulating alternative assumption-grounds: complexity science, philosophy of science, heterodox economics, organization theory, and systems thinking.

C2. Visibility in the dissertation

- Chapter 3 implements the review diagnostically by surfacing the fifteen assumptions through presupposition analysis.
- Chapter 4 implements it reconstructively by selecting alternative literatures for assumption-leverage rather than coverage.
- The structural conventions established in Section 2.7.4 render these linkages auditable.

Part D: Saturation and Completeness Criteria

Saturation is defined here as assumption saturation, not empirical exhaustion.

Table 13: Literature Sufficiency Criteria: Operational Definitions for Evaluating Completeness of the Dissertation's Source Engagement

Criterion	Operational meaning
Assumption coverage	All fifteen assumptions supported by multiple sources exhibiting and problematizing them
Alternative-ground coverage	CAS-compatible sources sufficient to articulate TO-positions

Criterion	Operational meaning
Counter-argument coverage	Major objections steelmanned and addressed via appropriate evaluative criteria
Theoretical saturation	New sources reproduce rather than extend the network diagnosis
Dialectical completeness	Mechanistic stance reconstructed sympathetically; CAS alternative coherent as replacement

Alhadeff-Jones's notion of reaching a "critical mass" of heterogeneous material (2013, p. 39) captures this saturation logic: patterns (e.g., the assumptive network / Triple Lock) become visible only once sufficient diversity accumulates.

2.8.3 Sources and Materials

Introduction

The problematizing review (Section 2.8.2) and the dissertation's transdisciplinary stance together determine the source ecology of the dissertation. As a pure-theoretical, paradigm-directed project (Section 2.1.1), it requires sources that carry assumptions, make assumption networks visible, and enable articulation of alternative assumption-grounds. This differs fundamentally from empirically oriented research prioritizing datasets and instruments.

Following Hayek's insight that complexity alters what counts as legitimate explanatory resources (1967, p. 26), sources were selected for assumption leverage, not comprehensiveness.

Source Category 1: Academic Literature

Primary texts (canonical formulations) were used to surface implicit ontological, epistemological, and methodological commitments in both mainstream and heterodox traditions.

Secondary texts (commentaries and debates) were used to identify which commitments are protected versus adjustable, supporting Lakatosian analysis of hard core and protective belt dynamics.

Source Category 2: Philosophical Texts

Philosophical sources were engaged to supply meta-theoretical grounding unavailable within economics alone:

1. Philosophy of science (Kuhn, Lakatos, Feyerabend) for paradigm analysis and evaluative criteria.
2. Critical realism (Bhaskar, Sayer) for stratification, emergence, and ontology-first ordering.
3. Process philosophy (Whitehead, Rescher) for alternatives to substance metaphysics.
4. Epistemology of uncertainty (Knight, Keynes) for grounding fundamental uncertainty.

Source Category 3: Historical Episodes

Historical episodes are used as pattern-level stress tests, not datasets. They illustrate when mechanistic assumptions render systemic dynamics invisible.

Illustrative cases include:

- The 2008 financial crisis (nonlinearity and feedback-driven fragility),
- The Great Depression (paradigm formation and contestation),
- Major technological transitions (path dependence and disruption).

These episodes support pattern-level diagnosis rather than hypothesis testing.

Source Category 4: Firm-Level Materials

Firm-level materials function as meso-level illustrations consistent with CAS logic. Two types were used:

- Academic case studies from organization theory and business history,
- Practitioner and business press sources, used cautiously for contextual reconstruction.

Illustrative firms (e.g., Nokia, Apple, Kodak) are not treated as samples but as theoretically instructive instances illuminating emergence, adaptation, lock-in, and ecosystem dependence.

Conclusion

This section specified the methodological tools and techniques supporting the dissertation's paradigm-level inquiry: a transdisciplinary stance, a problematizing literature review, defensible saturation criteria, and a structured source ecology. The guiding principle throughout is not maximal coverage but maximal assumption illumination and architectural integration, ensuring that the reconstruction of firms as complex adaptive systems is both methodologically disciplined and examinable.

2.8.4 Philosophical Argumentation Methods

Introduction

Pure theoretical work that does not rely on decisive empirical tests requires explicit argumentative discipline. This subsection identifies six philosophical modes of argument used across the dissertation – deduction, *reductio ad absurdum*, transcendental argument, retrodution, contrastive explanation, and abduction – and clarifies how they function jointly as the dissertation’s *defensive apparatus* for paradigm-level claims.

A caution applies. As Daniel Dennett notes, theorizing often carries implicit philosophical baggage; the present method choice is therefore made explicit rather than left tacit. This caution, however, is not the warrant for the dissertation’s approach. The primary warrant remains the dissertation’s established methodological foundations: critical realism and problematization.

Method 1: Deductive Argument

WHAT. Deduction derives conclusions that follow necessarily from stated premises. Here it is used to make assumption-structures explicit and to show what follows from competing ontological and epistemological commitments.

WHY. The dissertation’s cascade logic (ontology → epistemology → methodology) depends on demonstrating necessitation relations: if firms are treated as emergent and non-decomposable, then certain epistemic and methodological claims cannot be coherently retained.

HOW. Used wherever TO-assumptions are shown to entail epistemic limits and methodological consequences, especially in the transformation logic that links premises to downstream commitments.

Method 2: *Reductio ad Absurdum*

WHAT. *Reductio* tests a claim by showing that, if accepted, it yields contradiction, incoherence, or untenable consequences.

WHY. It is particularly effective against assumptions presented as “innocent simplifications.” In complex domains some simplifications are not merely partial but structurally distortive once their presuppositions fail.

HOW. Applied where mechanistic commitments (e.g., equilibrium, linearity, optimisation) are shown to generate implications inconsistent with the possibility of persistent heterogeneity, punctuated change, or systemic amplification – without pretending that a single observation can “empirically falsify” a paradigm.

A clarifying caution: when Alan Musgrave (1981) is used, his typology serves to distinguish what kind of assumption is in play; it is not treated as a shortcut to refutation.

Method 3: Transcendental Argument

WHAT. Transcendental argument asks: *what must be true of the world for some acknowledged phenomenon to be possible?* In critical realism it supports ontological inference by identifying conditions of possibility.

WHY. The dissertation repeatedly begins from uncontroversial organisational phenomena – adaptation, novelty, coordination failures, institutional dependence – and argues that certain mechanistic ontologies cannot supply the conditions under which these phenomena are possible.

HOW. Used where the existence of novelty and path dependence motivates specific ontological and epistemic commitments consistent with emergence, process, and openness.

Method 4: Retroduction

WHAT. Retroduction infers plausible generative mechanisms capable of producing observed patterns, rather than generalising surface regularities. This is a signature mode of critical realist explanation.

WHY. CAS properties are mechanism-rich and typically appear as pattern-level regularities: feedback, adaptation, path dependence, network effects. Retroduction provides the appropriate inferential bridge from pattern-recognition to mechanism-positing.

HOW. Used to move from recurring patterns (persistent heterogeneity, non-linear propagation, institutional stickiness) to candidate generative mechanisms (feedback loops, increasing returns, learning rules, boundary permeability).

Method 5: Contrastive Explanation

WHAT. Contrastive explanation asks: *why X rather than Y?* It is explanation by structured comparison.

WHY. It operationalises the dissertation's FROM→TO logic: the dissertation does not merely offer a CAS account, but shows why it outperforms mechanistic framings along explicit explanatory dimensions (scope, coherence with complexity properties, accommodation of novelty, and related virtues).

HOW. Implemented through the transformation structure across the dissertation's core comparative moves, where TO-positions are shown to illuminate phenomena that FROM-positions suppress, treat as residual, or redescribe away.

Method 6: Abduction

WHAT. Abduction (often framed as inference to the best explanation) selects among competing frameworks by explanatory virtues – scope, coherence, fecundity, and integrative power – rather than by assuming a single decisive test exists.

WHY. The dissertation is ultimately a framework-choice argument under complexity: which assumption-ground better explains firm dynamics and enables more realistic and progressive research programmes.

HOW. The abductive step occurs when, given comparative explanatory performance across ontological, epistemological, and methodological domains, the dissertation infers that the CAS-oriented assumption-ground is preferable as a progressive alternative.

Integration: How the Six Methods Work Together

These methods operate as a coordinated argumentative ecology:

- **Deduction** makes implication-structure explicit (cascade necessity).
- **Reductio** exposes incoherence when mechanistic assumptions are carried into complex domains.
- **Transcendental argument** supports ontological commitments via conditions of possibility.
- **Retroduction** supplies mechanism-inference appropriate to emergence and stratification.
- **Contrastive explanation** implements the FROM→TO comparative architecture.
- **Abduction** warrants overall framework preference without claiming spurious decisive tests.

Together, they allow a disciplined paradigm-level mode of theory construction that avoids both “pure speculation” and “false empiricism,” consistent with problematization and critical realism.

2.8.5 Assumption Identification and Mapping

Introduction

Problematization methodology requires systematic identification and articulation of the assumptions underlying theoretical positions. This subsection explains how assumptions were identified, refined, and organized for analysis. The purpose is methodological: showing how implicit commitments were surfaced and rendered discussable, not advancing constructive or integrative claims.

Foundational assumptions pose a specific difficulty because, as Gouldner (1970, p. 29) observes, they are typically “unpostulated and unlabeled,” functioning as background conditions of thought rather than explicit propositions. The procedures below address this invisibility problem.

Assumption Identification Methods

WHAT: process and principles

Assumption identification proceeded from the recognition that assumptions play different roles. Following Musgrave (1981), three assumption types were distinguished:

Table 14: Taxonomy of Assumption Types in Economic Theorising: Functions and Diagnostic Questions

Type	Function	Diagnostic question
Negligibility	Treats factors as unimportant	What is assumed away?
Domain	Restricts scope of applicability	Where does this theory claim to apply?
Heuristic	Simplifies analysis	What is set aside for tractability?

This distinction is necessary because mainstream economic theory often presents assumptions as provisional heuristics while deploying them as durable domain restrictions. Differentiating types makes such slippage visible without presupposing evaluative conclusions.

Analytically, the identification process also follows Alvesson and Sandberg’s (2011) guidance to vary focus across assumption depths (in-house, root metaphor, paradigm, field) rather than fixing on a single depth level from the outset.

WHY: methodological relevance

Making assumptions explicit is a prerequisite for problematization. As Bacharach (1989, p. 498) notes, theoretical claims cannot be meaningfully assessed unless the assumptions delimiting their domain are understood. Moreover, deeper (paradigmatic/field-level) assumptions exert disproportionate influence over what frameworks can and cannot address (Alvesson & Sandberg, 2011). The aim is therefore not exhaustiveness but leverage: to surface the background commitments most consequential for diagnosing limitations.

HOW: application in this dissertation

Assumption identification began with a broad survey of mainstream theories of the firm. An initial inventory of over thirty candidate assumptions was generated and iteratively refined into fifteen assumptions meeting three criteria:

1. **Foundationality** (they condition other claims)
2. **Distinctiveness** (not restatements of one another)
3. **Problematizability** (meaningfully challengeable)

The resulting set spans ontological, epistemological, and methodological commitments. No claim is made at this stage regarding adequacy or replacement.

Iterative Refinement Process

WHAT. Assumption articulation followed a multi-pass process consistent with Nickles' (1997) account of theory development as iterative rather than linear.

WHY. Because assumptions are tacit, initial formulations are often incomplete, overly concrete, or conflated; iterative engagement enables progressive abstraction and clarification.

HOW. Four refinement cycles were used:

1. Initial identification through literature engagement
2. Consolidation of overlapping formulations
3. Testing for internal coherence and scope boundaries
4. Verification of completeness at the level of assumption types

This stabilized the fifteen-assumption set used in subsequent analysis without implying finality or closure.

Interconnection Mapping

WHAT. Assumptions were examined as mutually conditioning commitments rather than isolated propositions. Kuhn's notion of paradigms as interconnected constellations of beliefs (1962/1970) informs this approach, while Alvesson and Sandberg's (2011) typology was used heuristically to distinguish depth (in-house, root metaphor, paradigm, field).

WHY. Interconnection analysis explains why isolated critiques often fail: assumptions reinforce one another across levels, insulating the overall structure from piecemeal challenge.

HOW. The mapping records patterned dependencies – especially across ontological, epistemological, and methodological commitments – serving a diagnostic purpose without presupposing integrative outcomes.

2.8.6 Managing Source Conflicts and Cross-Disciplinary Translation

Introduction

Engaging multiple heterodox traditions inevitably exposes conceptual tensions. This subsection explains how such tensions were recognized, documented, and managed to preserve conceptual clarity. The emphasis is procedural transparency rather than theoretical reconciliation.

Recognising Cross-Tradition Disagreement

WHAT. Different traditions operate with distinct boundary assumptions, producing divergent interpretations of similar phenomena (Mitchell, 2002). As Newell (2010, p. 334) notes, disciplines often “find what they are looking for,” with assumptions masquerading as findings.

WHY. Unacknowledged conflict risks incoherence; problematization requires making tensions explicit rather than smoothing them over.

HOW. When divergent accounts were encountered, competing interpretations were recorded prior to evaluative judgment, across ontological, epistemological, and explanatory claims.

Criteria for Adjudication

WHAT. Where comparison was unavoidable, Lakatos' distinction between hard-core commitments and auxiliary hypotheses served as a heuristic guide.

WHY. Without explicit criteria, cross-disciplinary engagement risks collapsing into eclecticism; criteria constrain interpretive discretion.

HOW. Selection was guided by consistency with the dissertation's stated problem focus rather than claims of general superiority; no attempt is made to resolve all tensions.

Managing Terminological Differences

WHAT. Identical terms often carry different meanings across disciplines (Bergmann et al., 2012). Apparent agreement can conceal substantive disagreement.

WHY. Conceptual precision is necessary for meaningful comparison and auditability.

HOW. Key terms were defined contextually at first use; discipline-specific meanings were distinguished where relevant. No unified vocabulary is imposed.

Ensuring Conceptual Consistency

WHAT. Consistency was assessed against a minimal set of commitments concerning complexity, nonlinearity, and emergence, following Mitchell's notion of compatible pluralism.

WHY. Integration without compatibility produces contradiction rather than insight.

HOW. Sources were retained, bracketed, or set aside based on compatibility with these minimal commitments – constraining scope without enforcing synthesis.

Closing note

Sections 2.8.5 and 2.8.6 deliberately stop short of constructive claims. Their function is to explain how assumptions were surfaced, compared, and stabilized for critique – *not* how alternative frameworks are assembled.

2.8.7 Historical–Theoretical Analysis

Introduction

Pure theoretical research, while not dependent on empirical data collection, may legitimately draw on historical analysis as a source of theoretical clarification. This subsection specifies how historical–theoretical analysis is used in this dissertation to illuminate the origins, trajectories, and persistence of foundational assumptions, without treating historical material as empirical validation or as a source of applied inference.

The function of historical analysis here is strictly diagnostic. It serves to reveal how particular assumptions became naturalised within economic theory, how alternative orientations were marginalised, and why certain commitments continue to structure contemporary inquiry despite persistent anomalies.

Intellectual Genealogies

WHAT: Tracing assumptive lineages

Intellectual genealogy examines how assumptions migrate, transform, and stabilise across thinkers, traditions, and historical periods. The methodological premise is that theoretical frameworks do not arise *ex nihilo*; they inherit implicit commitments from earlier intellectual contexts. As Alvesson and Sandberg (2022) note, all theorising proceeds from historically situated horizons of understanding.

Rather than reconstructing intellectual history exhaustively, the dissertation traces assumptive lineages – the transmission of deep commitments concerning being, causality, knowledge, and explanation.

For analytical purposes, two broad genealogical orientations are distinguished:

Table 15: Contrasting Genealogical Orientations in Economic Theorising and Their Dominant Assumptive Emphases

Genealogical orientation	Dominant assumptive emphasis
Substance–mechanistic	Stability, determinism, equilibrium, control
Process–relational	Becoming, openness, indeterminacy, development

This distinction is heuristic, not classificatory. It neither implies homogeneous traditions nor attributes modern positions to historical thinkers.

WHY: Why genealogy matters for problematization

Genealogical analysis is methodologically necessary because the assumptions targeted through problematization are historically sedimented. Their persistence is not explained solely by contemporary explanatory success, but by long intellectual trajectories that confer legitimacy, familiarity, and taken-for-granted status.

As Bizzarri (2021, p. 70) documents, Western thought underwent a decisive reorientation in which process (physis as development) was progressively displaced by cause (physis as originating principle). This shift – from becoming to being, from unfolding to control – did not eliminate alternative views but rendered them marginal. The outcome was not closure but asymmetry: one orientation became dominant, others residual.

Recognising this asymmetry clarifies why certain assumptions appear self-evident within modern economics, while alternatives are often dismissed as vague, unscientific, or insufficiently rigorous.

HOW: Use in this dissertation

Historical genealogy is used selectively to contextualise the assumptions identified in Chapter 3. The aim is not to adjudicate historical debates, but to demonstrate that assumptions now treated as technical necessities have philosophical ancestries.

For example, the migration of substance ontology and determinism from early modern physics into economic modelling is treated as a process of intellectual transfer rather than as an inevitable outcome of scientific progress. This contextualisation renders such assumptions visible as contingent commitments, not natural facts.

Historical Episodes as Theoretical Evidence

WHAT: The role of historical episodes

Historical episodes are used as theoretical evidence, not as empirical data. Following Schickore (2015), historical material can support philosophical analysis by illustrating how assumptions perform when confronted with complex phenomena.

Episodes are not used to test hypotheses or validate alternative frameworks. Their role is diagnostic: to expose mismatches between assumptive commitments and the phenomena they are meant to illuminate.

WHY: External standards without empiricism

Historical analysis provides an external standpoint from which entrenched assumptions can be examined. As Feyerabend (1978, p. 31) argued, meaningful critique requires standards not wholly internal to the framework under examination.

Hayek's (1974) Nobel Lecture exemplifies this function. His critique of "scientism" constitutes contemporaneous evidence that the imitation of physics in economics was recognised, from within the discipline, as methodologically problematic. Crucially, this is not a retrospective reinterpretation but a historically situated articulation of concern.

Such episodes show that doubts about mechanistic assumptions are recurrent within the history of economic thought itself, not externally imposed or anachronistic.

HOW: Application in this dissertation

Historical episodes are referenced sparingly and only where they illuminate assumption-level limitations. They are used to:

- illustrate how assumptions shape what becomes theoretically visible;
- demonstrate that alternative orientations were historically available but marginalised;
- clarify the distinction between event prediction and pattern recognition, without invoking applied implications.

No episode is treated as a case study, and none is used to derive normative or prescriptive conclusions.

Avoiding Anachronism

WHAT: The anachronism risk

Anachronism – projecting contemporary concepts or concerns onto past thinkers – is a fundamental methodological error in historical–theoretical analysis. The Cambridge School of intellectual history has shown how easily such projection distorts both historical understanding and theoretical diagnosis.

WHY: Preserving analytical integrity

As Skinner (1969) warned, reading modern categories into historical texts turns history into "a pack of tricks we play on the dead." Such distortions undermine the diagnostic value that historical analysis is meant to provide.

HOW: Methodological safeguards

Three safeguards are employed:

1. Contextual reading

Historical thinkers are interpreted within their own intellectual environments and problem contexts.

2. No proto-claims

Early thinkers are not treated as precursors of modern theories. Similarities are treated as resonances, not anticipations.

3. Genealogical restraint

The claim is not that contemporary approaches “return” to earlier thinkers, but that certain orientations were historically amplified while others were suppressed.

Accordingly, Heraclitus is not read as a proto-complexity theorist, nor is Prigogine presented as reviving ancient metaphysics. The claim is narrower and methodologically safer: Western intellectual history exhibits recurring tensions between substance-based and process-based orientations, and modern economics inherits one side of this tension more strongly than the other.

Closing Note

Historical–theoretical analysis in this dissertation functions as a context-revealing device, not as a source of validation or application. By tracing assumptive lineages, examining historically situated critiques, and avoiding anachronistic interpretation, it supports the dissertation’s problematization strategy while remaining strictly diagnostic.

2.8.8 Analytical Procedures

Introduction

Pure theoretical research, while not dependent on empirical data collection, nevertheless requires explicit analytical procedures to ensure argumentative clarity, internal coherence, and methodological integrity. This section specifies the analytical procedures employed in this dissertation for engaging with source materials, constructing and evaluating theoretical arguments, and maintaining consistency across the analysis.

These procedures are not designed to generate empirical findings or applied frameworks. Their sole purpose is to support rigorous problematization by rendering assumptions explicit, tracing their implications, and subjecting theoretical claims to disciplined scrutiny.

Close Reading Methods

WHAT: Procedures for textual analysis

Close reading in theoretical research involves disciplined engagement with texts to extract arguments, clarify meanings, and surface implicit assumptions. Four interrelated procedures constitute the close reading approach used in this dissertation:

1. Textual interpretation

Interpretation follows hermeneutic principles, reconstructing arguments within their original intellectual contexts. As Heidegger ([1927] 1962) emphasised, interpretation is never neutral extraction but an interaction between text and prior understanding.

2. Assumption surfacing

This involves identifying implicit premises that make arguments intelligible but are rarely defended explicitly. Alvesson and Sandberg (2024) stress that effective problematization requires illuminating precisely those assumptions treated as self-evident.

3. Argument reconstruction

Arguments are mapped by identifying claims, premises, inferential links, and scope conditions. This makes visible where conclusions depend on background commitments rather than explicit reasoning.

4. Conceptual clarification

Terms that carry different meanings across traditions are disentangled to ensure that apparent agreement or disagreement is not merely semantic.

In formal economic texts, this procedure distinguishes explicitly stated propositions from enabling background conditions (e.g., informational completeness, preference stability), which function as assumptions rather than testable claims and therefore require separate analytical treatment.

WHY: Rigorous engagement with sources

Close reading is methodologically necessary because theoretical arguments often derive persuasive force from what remains implicit. As Cilliers (2005a, p. 259) notes, all knowledge

operates within conceptual frameworks; close reading renders those frameworks visible rather than treating them as neutral backdrops.

HOW: Application in this dissertation

Close reading procedures were applied systematically in Chapter 3 to identify foundational assumptions embedded in mainstream theories of the firm. Historical and theoretical texts were interpreted contextually to surface their ontological, epistemological, and methodological commitments, without presupposing their adequacy or inadequacy.

Argument Construction

WHAT: Building coherent theoretical arguments

Argument construction follows established principles from logic and argumentation theory. Arguments are developed by:

- explicitly stating claims,
- identifying supporting grounds,
- clarifying inferential warrants,
- specifying scope and boundary conditions.

The Toulmin model (1958) is used descriptively as a diagnostic template, distinguishing claims, grounds, warrants, qualifiers, rebuttals, and backing. Its function is not to impose a format, but to evaluate argumentative completeness and coherence in both the literature and the dissertation's own claims.

Key procedures include:

- **Premise identification** (what must be assumed),
- **Logical mapping** (how premises support conclusions),
- **Scope specification** (where claims apply),
- **Counter-argument identification** (plausible objections),
- **Refinement** (revision to address ambiguity or weakness).

WHY: Ensuring logical discipline

Logical discipline is especially critical in pure theoretical research, where claims cannot be adjudicated through decisive empirical tests. As Cilliers (2005a, p. 263) argues, engaging complexity entails an ethical responsibility to avoid overstatement and false closure. Explicit argument construction ensures that claims remain proportionate to their justificatory support.

HOW: Application in this dissertation

Argument construction procedures were applied throughout the analytical chapters to ensure that claims about assumptions, their implications, and their limits were logically defensible and appropriately qualified. The emphasis remained on identifying tensions, inconsistencies, and boundary conditions rather than asserting definitive replacements.

Quality Assurance

WHAT: Mechanisms for analytical integrity

Quality assurance in theoretical research involves systematic checks to ensure internal consistency, adequate support, and cross-sectional coherence. Three primary mechanisms were employed:

1. Internal consistency checking

Examining whether claims introduced in one section contradicted or undermined claims elsewhere.

2. Cross-reference verification

Ensuring that conceptual commitments remained aligned across sections, especially where assumptions identified earlier carried implications later.

3. Counter-argument engagement

Identifying and addressing the strongest available objections rather than dismissing or ignoring them.

In addition, source triangulation was used to assess conceptual robustness across independent traditions. Following Denzin (1978) and Patton (1999), triangulation here refers not to empirical convergence but to theoretical resilience across distinct frameworks.

WHY: Maintaining theoretical credibility

Theoretical credibility depends not on eliminating uncertainty but on managing it transparently. As Cilliers (2005b, p. 607) emphasises, knowledge of complex phenomena necessarily involves reduction and boundary-drawing. Quality assurance procedures ensure that such boundaries are explicit, defensible, and consistently applied.

HOW: Application in this dissertation

Quality assurance operated continuously throughout the research process. Claims were cross-checked for coherence, sources were examined for compatibility, and tensions were addressed through revision rather than suppression. These procedures support rigorous problematization without advancing premature synthesis or closure.

2.8.9 Visual and Diagrammatic Methods

Introduction

Visual representation in theoretical research serves functions that extend beyond illustration. This section clarifies the bounded role of diagrams in this dissertation, specifies the types of visual tools employed, and defines their epistemological status as heuristic and clarificatory aids, not as autonomous theoretical constructs.

Visual elements support the problematization of assumptions by clarifying relationships, dependencies, and tensions that are difficult to apprehend through linear prose alone. They do not constitute models, frameworks, or design templates, nor do they substitute for textual argumentation.

Role of Visual Representation in Theoretical Work

WHAT: Functions of diagrams

Within this dissertation, diagrams serve four strictly limited functions:

1. **Heuristic clarification:** Making relational structures visible to reduce cognitive load (Larkin & Simon, 1987).
2. **Structural visibility:** Revealing patterns of dependence or clustering without asserting causal direction.
3. **Complexity navigation:** Supporting orientation where multiple assumptions or levels interact.
4. **Communicative accessibility:** Providing alternative entry points for comprehension without replacing argumentation.

WHY: Visual support without theoretical closure

Ackoff and Gharajedaghi (2003) note that linear representations often obscure relational structure in complex systems. Visual tools can mitigate this risk, but only if used cautiously. In theoretical research, diagrams can easily be mistaken for models or prescriptions.

Accordingly, all visual representations are explicitly subordinated to textual analysis and treated as aids to understanding rather than independent analytical outputs.

HOW: Application in this dissertation

Visual tools are used sparingly and only where textual exposition would impose excessive cognitive burden. Their function is confined to clarifying relationships already established through text.

Diagram Types and Construction Principles

WHAT: Diagram types employed

Only a limited set of diagrammatic forms appears, each with a specific diagnostic purpose:

Table 16: Typology of Visual-Analytical Diagrams Employed in the Dissertation: Purposes and Scope of Application

Diagram type	Purpose	Scope
Conceptual maps	Clarify assumption relationships	Chapter 3
Dependency sketches	Indicate logical dependence	Assumption analysis
Genealogical diagrams	Trace intellectual lineages	Historical sections
Typological matrices	Contrast categories	Methodological clarification

No diagram is presented as an operational framework or transferable model.

WHY: Guarding against over-interpretation

As Whetten (1989) cautions (via Sutton & Staw, 1995), diagrams can create an illusion of explanatory sufficiency. To prevent this, each visual element is accompanied by explicit textual clarification specifying:

- what the diagram shows,
- what it does not show,
- and which claims it does not make.

HOW: Construction procedures

Four principles governed diagram construction:

1. **Iterative simplification** – minimum necessary complexity
2. **Textual anchoring** – every diagram explained in prose
3. **Consistency checking** – alignment with textual claims
4. **Non-operational design** – avoidance of flow, symmetry, or procedural cues

Epistemological Status of Diagrams

WHAT: Diagrams as heuristic devices

Diagrams function as heuristic devices rather than arguments. They assist in seeing relationships but do not establish them independently of textual reasoning.

As Sutton and Staw (1995) argue, strong theory relies on verbal argumentation; diagrams may support but cannot replace it.

WHY: Preserving theoretical modesty

Treating diagrams as heuristic respects the limits of pure theory and prevents visual artefacts from being mistaken for finished systems or applied methodologies.

HOW: Integration across the dissertation

Visual elements are integrated only where they clarify discussion and never carry argumentative weight independently. No diagram introduces claims not already justified in text, and none is necessary for logical completeness.

2.8.10 Firm-Level Illustration as a Heuristic Device

Introduction

Throughout the dissertation, selected firm-level examples are used to support the intelligibility of abstract theoretical arguments. This section clarifies the limited methodological role of such examples within a pure theory project. It specifies what firm-level illustration is, why it is legitimate in assumption-challenging research, and how it is employed without constituting empirical validation, framework testing, or applied synthesis.

Firm-level material functions exclusively as a heuristic and pedagogical aid. Its role is to render abstract claims more accessible and assessable at a qualitative, reflective level – not to verify those claims empirically or demonstrate practical superiority.

WHAT: Firm-level illustration

Firm-level illustration refers to the selective use of well-documented organisational episodes to clarify the implications of theoretical assumptions under discussion. When abstract concepts such as emergence, path-dependence, or organisational routines are introduced, illustrative examples are occasionally used to show how such concepts may plausibly manifest at the level of firms.

These illustrations do **not** function as tests, proofs, or validations. They do not establish causal generalisations, claim representativeness, or serve as empirical leverage. Their sole purpose is to make abstract reasoning cognitively tractable.

WHY: Heuristic support for meso-level reasoning

Abstract theoretical arguments – especially those concerning ontology or epistemology – can remain opaque without some form of concrete reference. Firm-level illustration responds to this well-recognised difficulty without altering the epistemic status of the work.

As Dopfer, Foster, and Potts (2004) argue, many economic phenomena unfold at an intermediate analytical level that is neither individual nor aggregate. Illustrative firm-level material can assist readers in orienting themselves at this meso-level without implying that such material exhausts, validates, or adjudicates the theory.

Crucially, firm-level illustration here supports reflection, not theory choice.

HOW: Use within the dissertation

Firm-level illustrations appear selectively and sparingly in the substantive chapters. Their use follows three strict principles:

1. Illustrative, not demonstrative

Examples are introduced using language such as *illustrates*, *renders visible*, or *helps clarify*, and never with claims of proof, confirmation, or validation.

2. Subordinate to argumentation

No theoretical claim depends on illustrative material for its logical validity. All arguments stand independently of examples.

3. Non-systematic and non-exhaustive

Illustrations are not assembled into datasets, typologies, or comparative frameworks, and their inclusion does not imply coverage, completeness, or scalability.

Epistemological status

Firm-level illustrations occupy a non-verificatory epistemological status. Following Siggelkow (2007), their contribution is illustrative rather than confirmatory.

They serve three modest functions:

Table 17: Functions of Illustrative Examples in the Dissertation: Pedagogical, Heuristic, and Orientational Roles

Function	Role in the dissertation
Pedagogical	Improving intelligibility of abstract arguments
Heuristic	Clarifying why certain assumptions matter
Orientational	Locating arguments at the meso-level

They do not establish regularities, estimate frequencies, or support inference to prescription.

Guarding against overreach

To preserve methodological discipline, the dissertation explicitly avoids:

- claims of representativeness
- claims of empirical confirmation
- claims of framework superiority based on cases
- accumulation of cases as anomaly pressure
- inference from illustration to prescription

Where alternative interpretations are possible, they are acknowledged without attempting resolution through additional examples.

Conclusion

Firm-level illustration functions as a supportive heuristic device, not as an empirical strategy. It enhances intelligibility of assumption-level arguments without advancing claims of validation, synthesis, or application. This restraint preserves the integrity of the dissertation as a pure theory project.

Conclusion of Section 2.8

Section 2.8 has specified the practical tools and techniques through which the dissertation's paradigm-level inquiry is conducted. Four findings carry forward.

First, transdisciplinarity is a methodological requirement, not a stylistic preference. Complex adaptive systems demand transdisciplinary understanding because their constitutive features – boundary-crossing interaction, multi-scale operation, circular causality, and framework-dependent boundaries – exceed what any single discipline can address. The transdisciplinary orientation is operationalized through structured deployment: philosophy of science and economic methodology ground the diagnostic work of Chapter 3; complexity science, organisation theory, and systems thinking ground the reconstructive work of Chapter 4. Four principles – problem-driven selectivity, contextual translation, compatible pluralism, and reflexive positioning – guard against superficial borrowing and undisciplined eclecticism.

Second, the problematizing literature review operationalizes the dissertation's methodological stance at the level of literature engagement. Unlike reviews designed to aggregate findings within an accepted framework, the problematizing review targets assumptions underlying influential literatures – creating conceptual space for alternative assumption-grounds. Saturation is defined as assumption saturation: the point at which the assumptive network becomes visible as a structured configuration.

Third, six philosophical argumentation methods – deduction, *reductio ad absurdum*, transcendental argument, retroduction, contrastive explanation, and abduction – function as a coordinated ecology. Together they enable paradigm-level theory construction that avoids both speculation and false empiricism.

Fourth, assumption identification proceeded through iterative refinement – from over thirty candidates to fifteen core assumptions meeting criteria of foundationality, distinctiveness, and problematizability. Historical-theoretical analysis, visual methods, and firm-level illustration function as strictly bounded heuristic devices: they support intelligibility without constituting empirical validation.

Section 2.9 specifies what the dissertation contributes, how those contributions should be characterised, and where the limits of paradigm-level theoretical work lie.

2.9 Contributions, Limitations, and Future Directions

Section 2.8 specified the practical tools and techniques through which the dissertation's paradigm-level inquiry is conducted – transdisciplinary orientation, problematizing literature review, philosophical argumentation methods, assumption identification procedures, and bounded heuristic devices. This section turns from how the inquiry was conducted to what it accomplishes and where its limits lie.

This specification is necessary for two reasons. First, pure theoretical work requires explicit articulation of contributions, since its outputs are not self-evident in the manner of empirical findings or predictive successes. Second, paradigm-level ambition carries a predictable hazard: scope can encourage claims broader than the argument sustains. Calibration – distinguishing what the dissertation establishes from what it does not – is therefore not defensive limitation but a condition of credibility.

The argument proceeds through two subsections. Subsection 2.9.1 establishes what the dissertation contributes, how those contributions arise from the adopted methodology, which boundaries define scope, and how the contribution should be characterised – as transformative rather than revolutionary. Subsection 2.9.2 specifies methodological limits, potential vulnerabilities, the relationship to empirical research, and future directions the framework enables.

Chapter 3 implements the diagnostic phase by identifying and mapping the assumption network requiring transformation.

2.9.1 Contributions, Scope, and Methodological Origins

Introduction

This section states what the dissertation contributes, how those contributions arise from the adopted methodology, which boundaries define the scope of its claims, and how the contribution should be characterised. In a pure-theory dissertation, scholarly integrity requires explicit contributions, transparent linkage between procedure and output, clear specification of what is not claimed, and calibrated framing that aligns ambition with what the evidential basis can sustain.

Part A: Boundaries as a Condition of Knowledge

Boundaries as enabling, not confining

In complexity-oriented inquiry, boundaries are not merely restrictions; they are enabling conditions of intelligibility. As Cilliers (2005b, p. 610) argues, "the boundaries of a complex system are not given but are constructs that we create in order to understand the system" – boundaries "constitute that which is bounded." The dissertation's boundaries therefore do not weaken its contribution; they make it assessable and academically defensible. Without explicit boundaries, the inquiry would expand into an ambition that cannot be evaluated.

Situated knowledge and strategic delimitation

Complexity-oriented inquiry requires a situated standpoint rather than a "view from nowhere." This dissertation is explicitly situated: it addresses a particular tradition in economic theorising about the firm, works primarily at the meso-level, and proceeds through conceptual and philosophical argumentation. These are strategic methodological choices that structure what can be claimed.

Incompressibility and scope realism

Complexity is incompressible: no representation can fully capture an open system without omission, and omissions can have non-linear effects (Cilliers, 1998, p. 10). Accordingly, the dissertation does not claim to exhaust firm reality, economics as a whole, or all legitimate methods. It claims instead to contribute within explicit, defensible limits.

Part B: What the Dissertation Contributes

Within these boundaries, the dissertation contributes five connected outputs. Each contribution follows directly from methodological procedures specified earlier in this chapter (problematization, close reading, multi-level reasoning, translation discipline, and quality assurance).

Contribution 1: Clarification of the assumptive structure of mainstream firm theorising

The dissertation makes explicit background assumptions that are often implicit, dispersed, or treated as methodological conveniences. By surfacing these assumptions, it clarifies how particular commitments delimit what can be represented about firms – especially regarding emergence, adaptation, path-dependence, and fundamental uncertainty.

Methodological origin: problematization (assumption surfacing) supported by close reading, argument reconstruction, and conceptual clarification.

Contribution 2: Diagnosis of why piecemeal reform fails

The dissertation demonstrates that the mechanistic paradigm's assumptions form a mutually reinforcing network – the Triple Lock – that neutralises isolated challenges. This diagnostic contribution explains the persistent failure of reform efforts that target individual assumptions while leaving the surrounding structure intact.

Methodological origin: interconnection mapping, Lakatosian research programme analysis, and the Ma-inspired attention to relational structure.

Contribution 3: A coherent bridge between complexity vocabulary and economic/organisational vocabularies

Across the literatures engaged, key terms frequently shift meaning (e.g., equilibrium, emergence, adaptation, uncertainty). The dissertation establishes disciplined translation practices that reduce equivocation and improve conceptual coherence across traditions.

Methodological origin: cross-disciplinary translation procedures, conflict recognition, and explicit adjudication criteria grounded in compatibility and scope.

Contribution 4: Reconstruction of an alternative assumption-ground

The dissertation does not merely diagnose limitations but constructs an alternative: firms reconceptualised as complex adaptive systems, with ontological, epistemological, and methodological commitments derived through the cascade logic specified in Section 2.7.2. This reconstruction satisfies Lakatos's requirement that paradigms are displaced by superior alternatives, not merely criticised.

Methodological origin: the FROM→TO transformation framework, first-principles reconstruction, and coherence constraints ensuring mutual compatibility across commitments.

Contribution 5: A basis for subsequent research programmes

The dissertation identifies conceptual openings and methodological implications that can guide subsequent work – empirical or applied – while maintaining the dissertation's status as a theoretical contribution rather than a design manual. It specifies enabling conditions and researchable problem-spaces without claiming implementation.

Methodological origin: maintaining the "diagnosis / conditions-of-possibility" register rather than a "construction / deployment" register.

Part C: Scope Boundaries

The dissertation's claims are explicitly bounded along six dimensions:

Domain boundary – primary focus is the firm and meso-level organisational phenomena, not "economics as a whole."

Level boundary – emphasis on meso-level dynamics (routines, capabilities, organisational patterns, networks) rather than full microfoundations or macro-modelling.

Method boundary – conceptual and philosophical argumentation, not empirical estimation or statistical inference.

Evidence boundary – historical and case material is used illustratively and diagnostically, not for population generalisation.

Tradition boundary – multiple heterodox streams are engaged to improve conceptual adequacy, without claiming any lens exhausts the firm.

Practical boundary – no managerial prescriptions or deployable redesign; the dissertation clarifies conceptual conditions under which alternative understandings become available.

These boundaries are enabling constraints: they define what the dissertation can do well and prevent scope inflation.

Part D: The Character of the Contribution

Transformative versus revolutionary contribution

Two contribution registers are analytically distinct. Revolutionary contribution claims radical novelty that renders prior work obsolete or irrelevant. Such claims are rhetorically attractive

but methodologically fragile: they invite easy refutation through the identification of antecedents, and they typically exceed what a pure-theory evidential basis can warrant. Transformative contribution, by contrast, revisits foundational assumptions, clarifies limits, and integrates insights across literatures that are usually compartmentalised. The novelty lies in re-articulation and re-connection – not in claiming creation *ex nihilo*.

For a pure-theory dissertation, the transformative register is the appropriate one. This project works within an existing intellectual ecology and advances the claim that entrenched assumptions restrict what can be seen and said about firms under conditions of complexity. Its contribution therefore consists in identifying assumption structures, diagnosing their consequences, and clarifying conditions under which alternative framings become intelligible and defensible – rather than asserting discontinuity for its own sake.

Scholarly lineage and the Newton principle

Newton's "standing on the shoulders of giants" is invoked here not as rhetorical modesty but as methodological accuracy. Paradigm-level conceptual progress typically occurs through recombination, refinement, and consolidation of existing insights – often by making compatible elements cohere that previously remained separated by disciplinary boundaries. The dissertation draws explicitly on established traditions: critical realism (Bhaskar, Lawson), complexity science (Holland, Kauffman, Cilliers), economic methodology (Knight, Hayek, Mises), philosophy of science (Kuhn, Lakatos), and organisation theory (Nelson & Winter, March, Teece). A dissertation that makes its debts explicit while clarifying limitations and compatibilities is, by that fact, more defensible than one that claims radical discontinuity.

Why calibrated claims are stronger

Paradigm-level work faces a predictable hazard: ambitious scope can encourage claims broader than the argument can sustain. In a pure theory project, evidential discipline therefore takes the form of calibration: distinguishing what the argument does establish (assumption structure, conceptual limits, conditions of adequacy, implication tracing) from what it does not establish (empirical verification, population generalisation, decisive closure).

This calibration is not a retreat from ambition. It is a quality-control device that aligns the dissertation's evaluative standards with the logical character of its contribution – coherence, explanatory adequacy, mechanism illumination, and theoretical fruitfulness (Section 2.6.1) – rather than exposing it to category-error demands that it cannot, and does not attempt to, satisfy. As Cilliers (2005b, p. 612) emphasises, "the fact that our knowledge is limited is not a disaster, it is a condition for knowledge. Limits enable knowledge."

Conclusion

The dissertation's contribution is constituted by the boundaries that make it intelligible and evaluable. Within those boundaries, it clarifies background assumptions, diagnoses why piecemeal reform fails, strengthens cross-disciplinary coherence, reconstructs an alternative assumption-ground, and establishes a basis for subsequent research. The contribution is transformative rather than revolutionary: it operates within an existing intellectual lineage, makes its debts explicit, and increases explanatory adequacy by surfacing assumptions and clarifying the conditions under which alternative perspectives can be pursued. This framing is not stylistic; it is a methodological safeguard that increases credibility by aligning claims with what pure theoretical work can legitimately defend.

2.9.2 Limitations and Future Directions

Introduction

Section 2.9.1 established what the dissertation contributes, how those contributions arise from its methodological design, and how the contribution should be characterised. This section provides a principled account of what the dissertation cannot establish, where vulnerabilities may remain, how it relates to empirical research, and which future research directions it enables.

Following Cilliers (2005b, p. 613), the aim is not defensive limitation but epistemic humility: recognising that the absence of certain claims is a condition of credibility rather than a deficiency. The limitations identified here follow directly from the dissertation's status as pure theoretical, paradigm-level inquiry.

Part A: Methodological Limits of the Dissertation's Claims

A.1 Metaphysical versus empirical commitments

Because the dissertation operates at the level of foundational assumptions and framework construction, it does not and cannot "prove" the correctness of CAS assumptions in a correspondence-empirical sense. The assumptions under transformation are not empirical findings but pre-empirical commitments – historically inherited metaphysical and epistemological positions that structure what counts as evidence in the first place.

Accordingly, CAS assumptions are supported here through philosophical argumentation, conceptual coherence, explanatory adequacy, and problem-solving capacity, not through direct empirical verification. As Cilliers (2005b, p. 612) emphasises, knowledge of complex systems cannot achieve complete, context-free objectivity. This is not a temporary limitation awaiting technical resolution; it is a constitutive feature of inquiry under complexity.

A.2 Pattern understanding versus point prediction

The CAS orientation developed in this dissertation privileges pattern explanation over point prediction. In systems characterised by novelty, feedback, path dependence, and fundamental uncertainty, precise forecasts of specific outcomes are not an appropriate benchmark of theoretical adequacy.

The dissertation therefore does not generate predictive claims of the form "firm X will do Y at time T." Instead, it clarifies: which dynamics become visible under CAS assumptions; which dynamics are systematically obscured under mechanistic assumptions; and which explanatory strategies remain viable when prediction is limited by non-ergodicity.

This is an enabling trade-off, not an omission: explanatory intelligibility is prioritised over spurious precision.

A.3 Intellectual versus ontological emergence

Intellectual emergence (insight arising from integration of arguments) must be distinguished from ontological emergence (novel causal powers arising in material systems). The dissertation

can demonstrate the former through systematic synthesis; it cannot demonstrate the latter through textual organisation alone.

Accordingly, the dissertation avoids claims that the text itself "produces" ontological emergence or that conceptual synthesis constitutes proof of emergent causation. This limit is categorical – rooted in logical type – rather than a shortcoming awaiting empirical supplementation.

A.4 Process versus product in self-reference

Any self-referential reading of the dissertation is structural and analogical, not literal. While the design of the inquiry exhibits architectural features associated with CAS reasoning (interconnection, multi-level organisation, non-linearity of insight), the completed dissertation cannot adapt dynamically or couple recursively with its environment.

The strongest self-referential claim therefore concerns the design logic and interpretive use of the text, not the text as an evolving complex system. This boundary preserves analytical clarity and avoids category confusion.

Part B: Potential Vulnerabilities and Areas Requiring Further Development

B.1 Translation risk: illustration without overreach

Firm-level illustrations support intelligibility, but their epistemic status remains illustrative rather than evidential. A potential vulnerability lies in interpretive contestation: alternative readings may argue that mechanistic explanations could be extended to account for the same phenomena.

The dissertation mitigates this risk through careful linguistic discipline ("illustrates," "renders visible"), contrastive framing, and the use of structurally revealing cases. It does not, and cannot, eliminate the possibility of competing interpretations – nor should it attempt to do so within a pure theory project.

B.2 Integration risk: pluralism without eclecticism

The framework is explicitly pluralist, but pluralism carries an inherent tension: without discipline, it can collapse into eclecticism. While the dissertation introduces architectural constraints (assumption hierarchy, cascade logic, translation protocols, compatible pluralism), it does not fully resolve adjudication in cases of deep theoretical incompatibility.

This is not a flaw unique to the present work; it is a standing challenge for pluralist research programmes. Future work will need to refine operational criteria for adjudication in applied or empirical contexts.

B.3 Operationalisation gap: concepts to measures

The dissertation identifies CAS-relevant constructs – emergence, adaptation, self-organisation, path dependence, network effects – and clarifies their conceptual roles. It does not provide a complete empirical instrumentarium for measuring these constructs.

Translating philosophical clarity into empirical operationalisation is itself a substantial research programme. Some constructs will prove more tractable than others, and measurement strategies will necessarily be context-dependent.

B.4 Boundary dependence: domain, context, and level

The argument's force depends in part on deliberate boundaries: firms as the primary object, the meso-level as the primary analytic scale, and predominantly Western institutional contexts. These boundaries are enabling conditions for coherence, but they also delimit the reach of the claims.

Cross-cultural variation, alternative institutional orders, and explicit micro–meso–macro linkage mechanisms remain open domains for extension rather than settled results.

Part C: Relationship to Empirical Research

The absence of empirical testing reflects the dissertation's positioning as foundational theoretical work (Section 2.5). Empirical research is therefore not a "missing chapter" but a subsequent stage of the research programme.

The framework is designed to enable empirical inquiry by clarifying: what phenomena are theoretically salient under complexity; what counts as evidence under non-ergodicity and uncertainty; and which mechanisms warrant direct investigation.

Part D: Future Research Directions

D.1 Empirical programmes

Future research may engage the framework through: longitudinal studies of firm evolution and capability emergence; network-analytic studies linking performance to topology; industry-evolution analyses combining historical methods with simulation; and investigation of persistent disequilibrium dynamics where equilibrium is routinely imposed.

D.2 Method development

The framework invites methodological innovation in: metrics for emergence and adaptive capacity; agent-based models explicitly designed for non-ergodic dynamics and feedback; and qualitative and mixed-method approaches suited to context dependence and reflexivity.

D.3 Theoretical extensions

Three extensions are especially salient: explicit micro–meso–macro linkage mechanisms; cross-cultural and comparative-institutional applications; and temporal dynamics (phase transitions, tipping points, irreversibility) as core objects of firm theory rather than residual disturbances.

D.4 Policy and practice applications

Because many interventions operate at meso-levels (industrial policy, competition policy, innovation systems, regional development), future work can explore how meso-level configurations condition policy effects and how policy evaluation can proceed under path dependence and contextual sensitivity.

Part E: Why the Stated Limitations Do Not Undermine the Contribution

The limitations articulated in this section do not weaken the dissertation's contribution; they secure its methodological validity. The project's objective is not to deliver empirical confirmation, predictive closure, or applied prescription, but to reconstruct the assumptive architecture governing how firms are theorised. At this level of inquiry, overextension would constitute a category error: attempting to validate pre-empirical commitments through methods whose admissibility those commitments themselves determine.

By explicitly delimiting what a pure theoretical, paradigm-level dissertation can and cannot claim, the analysis preserves coherence between ontological stance, epistemological ambition, and methodological practice. The contribution should therefore be evaluated not by standards appropriate to empirical hypothesis testing, but by its success in clarifying background assumptions, exposing their implications, and establishing defensible conditions under which alternative framings become intelligible and researchable. Within those criteria, the stated limitations function as enabling constraints rather than deficiencies, strengthening rather than undermining the dissertation's scholarly standing.

Conclusion of Section 2.9

Section 2.9 has specified what the dissertation contributes, how those contributions should be characterised, and where its limits lie. Four findings carry forward.

First, boundaries are enabling conditions of knowledge, not merely restrictions. Following Cilliers, the dissertation's explicit boundaries – domain, level, method, evidence, tradition, and practical application – constitute rather than confine its contribution. Within these boundaries, the dissertation contributes five connected outputs: clarification of the assumptive structure of mainstream firm theorising, diagnosis of why piecemeal reform fails (the Triple Lock), a coherent bridge between complexity and economic vocabularies, reconstruction of an alternative assumption-ground adequate to firms as complex adaptive systems, and a basis for subsequent research programmes.

Second, the contribution is transformative rather than revolutionary. The dissertation works within an existing intellectual ecology, makes its debts explicit, and advances explanatory adequacy by surfacing assumptions and clarifying conditions under which alternative framings become intelligible. This framing is a methodological safeguard that aligns claims with what pure theoretical work can legitimately defend.

Third, calibrated claims are stronger claims. By distinguishing what the argument establishes (assumption structure, diagnostic limits, conditions of adequacy) from what it does not establish (empirical verification, population generalisation, predictive closure), the dissertation preserves coherence between ontological stance, epistemological ambition, and methodological practice.

Fourth, the stated limitations do not undermine the contribution; they secure its methodological validity. The dissertation enables rather than conducts empirical inquiry – clarifying what phenomena are theoretically salient, what counts as evidence under complexity, and which mechanisms warrant investigation. Future research directions in empirical programmes, method development, theoretical extensions, and policy applications follow from, rather than substitute for, the foundational work undertaken here.

Chapter 3 now implements the diagnostic phase, mapping the fifteen assumptions that constitute the mechanistic paradigm's hard core and protective belt.

Chapter 3 – The Diagnosis: Anatomy of the Mechanistic Paradigm

Introduction

Chapter 3 performs the diagnostic task upon which the remainder of the dissertation logically depends. Its purpose is not to assess individual economic models on their own terms, nor to arbitrate empirical disputes within the existing theoretical repertoire, but to reconstruct the deep assumptive architecture that renders the mechanistic paradigm internally coherent, analytically powerful, and – under conditions of complexity – systematically fragile.

The chapter proceeds from the methodological stance established in Chapter 2: economic theories are treated not as neutral instruments or interchangeable models, but as structured systems of ontological, epistemological, and methodological commitments. These commitments operate prior to formalisation. They define what kinds of entities are admissible, what forms of causation are recognisable, what counts as explanation, and which phenomena are rendered invisible by construction. When such commitments are misaligned with the nature of the domain they seek to explain, theoretical failure is not accidental or corrigible by refinement; it is structural.

The diagnostic argument unfolds in four stages, each raising a question that the next must answer.

Section 3.1 reconstructs the paradigm's internal architecture. It identifies five ontological assumptions – reductionism, determinism, substance ontology, closure, and atomism – that constitute the hard core, and shows how these generate four epistemological assumptions through a closed loop in which what is assumed to exist conditions what is assumed knowable. From these, six derivative assumptions follow with logical necessity: equilibrium, linearity, homogeneity, ergodicity, optimisation, and decomposability. The fifteen assumptions together form a mutually reinforcing network. The critical finding for firm theory is that this architecture systematically eliminates the meso-level – the intermediate scale where firms, networks, and institutions operate – not through oversight but through the joint logic of the assumptions themselves.

Yet if these assumptions are demonstrably incompatible with the properties of complex adaptive systems, a puzzle arises: why do they persist despite decades of anomalies, crises, and sustained critique?

Section 3.2 answers this persistence puzzle. Three reinforcing feedback loops – ontological, epistemic, and structural – interlock at high-centrality nodes, creating a triple-locked system that absorbs perturbations rather than transmitting them. The paradigm further reproduces itself through semantic mechanisms that embed assumptions in vocabulary, methodological mechanisms that equate tractability with rigour, and institutional mechanisms that select for conformity. Evaluated through Kuhnian and Lakatosian frameworks, the inherited programme exhibits the signatures of degeneration: auxiliary assumptions proliferate while the hard core remains protected; predictive success is confined to calibration domains; excess corroboration in domains diagnostic of complexity is systematically absent. The verdict is formal, not rhetorical.

But how does this structural lock-in express itself in the actual practice of the discipline?

Section 3.3 traces the methodological symptoms. Physics envy imported the mathematical form of classical mechanics without validating domain conditions. The optimisation fetish converted a context-bound tool into a governing worldview. DSGE institutionalised these commitments, while model monism excluded alternatives by definition. The aestheticisation of economics subordinated truth to elegance, explaining why the 2008 crisis did not trigger paradigm change – the framework survived because it remained beautiful, not because it remained adequate. Hayek's diagnosis of the pretence of knowledge emerges as the unifying foundation: economists systematically overestimate what can be known, and this epistemic hubris reclassifies structural failure as technical shortfall.

The diagnostic argument, however, remains incomplete without empirical warrant. If the failures of the mechanistic paradigm are truly generated by its architecture, then predictable patterns of breakdown should be observable across domains where complexity is constitutive.

Section 3.4 supplies this warrant. Predictive brittleness in economic theory – stable performance near equilibrium, catastrophic failure under regime change – reflects structural mismatch, not insufficient sophistication. The illusion of control in management reflects the same ontological error applied at firm level: forecast-centred planning failures, Goodhart pathologies, capability destruction under reengineering logic, and risk-model blindness under networked contagion. The 2008 financial crisis provides the paradigmatic stress test: multiple assumptions failed simultaneously, their internal coherence became a liability, and what collapsed was the adequacy of a mechanistic ontology for an adaptive, interaction-dominated economic reality. The simultaneity is diagnostic – not coincidence but consequence of the architecture identified in Section 3.1.

The organising claim of the chapter is therefore precise and demanding: the recurring limitations of orthodox economic theory – its difficulty accounting for institutional change, endogenous innovation, financial instability, and systemic crises – arise because the mechanistic paradigm is architecturally designed to suppress emergence, learning, non-linearity, and historical irreversibility. Where complexity is constitutive rather than residual, such a paradigm does not merely approximate poorly; it misrepresents the object of inquiry.

For the theory of the firm, the stakes are direct and unavoidable. Firms exhibit precisely the properties that the 15-assumption architecture renders invisible: emergent capabilities, adaptive coordination, path-dependent learning, and relational constitution. The diagnostic that follows therefore prepares the ground not merely for critique but for reconceptualising what firms are and how they can be theorised.

Chapter 3 remains strictly diagnostic. It does not introduce an alternative ontology, propose a reconstruction, or gesture toward synthesis. The task is to demonstrate that the mechanistic paradigm's failures are generated by its foundational commitments – and thereby to establish the necessity, not merely the preference, for the ontological inquiry undertaken in Chapter 4.

3.1 The Assumption Network

Chapter 3 performs a diagnostic task, and that diagnosis begins with first principles. Before the chapter can explain why the mechanistic paradigm persists despite anomalies (Section 3.2), trace its methodological symptoms (Section 3.3), or document its consequences (Section 3.4), it must first establish what the paradigm assumes. Without this foundation, critique addresses surface phenomena rather than generative structure.

Section 3.1 undertakes this foundational reconstruction. The argument proceeds in four stages. Subsection 3.1.1 identifies five ontological assumptions – reductionism, determinism, substance ontology, closure, and atomism – that constitute the paradigm's hard core: what it presupposes about the nature of reality. Subsection 3.1.2 adds the epistemological layer – certainty, universal laws, monism, and perfect information – showing that epistemology functions as ontology's shadow: what is assumed to exist conditions what is assumed knowable. Subsection 3.1.3 demonstrates that from this core, six derivative assumptions – equilibrium, linearity, homogeneity, ergodicity, optimisation, and decomposability – follow with logical necessity; they are entailments, not choices. Subsection 3.1.4 reveals that these fifteen assumptions form not a list but a self-reinforcing network whose most consequential property is the systematic elimination of the meso-level – the scale at which firms, networks, and institutions operate.

The reconstruction is diagnostic: it maps the architecture without proposing alternatives. The adequacy of these assumptions is addressed in Chapter 4.

3.1.1 Core Ontological Assumptions

(Reductionism, Determinism, Substance, Closure, Atomism)

Introduction

Section 3.1.1 identifies the ontological hard core of the mechanistic paradigm – that is, the foundational assumptions concerning the nature of reality embedded in the inherited body of economic theory. Ontological assumptions are the deepest presuppositions of any paradigm: they specify what kinds of entities are taken to exist, how those entities relate, and what kinds of processes are considered possible. In Lakatosian terms, they constitute the hard core of a research programme – assumptions that are insulated from falsification and rarely made explicit within normal scientific practice (Lakatos, 1970, pp. 91–196).

These assumptions are prior to epistemology (what can be known) and methodology (how inquiry proceeds). They silently structure what economists attend to, what they abstract from, and what they regard as legitimate explanation. Certain phenomena are rendered visible within this framework, while others become analytically inaccessible.

This subsection performs four diagnostic functions.

First, it makes explicit the metaphysical foundations of the inherited paradigm – assumptions so deeply embedded that they often appear as natural facts rather than theoretical commitments. Second, it shows how these assumptions predetermine the analytical limits of the paradigm: once accepted, certain explanatory forms become unavoidable.

Third, it establishes their incompatibility with emergence, learning, novelty, and co-evolution – phenomena that are central to understanding firms and economies under conditions of complexity.

Fourth, it provides the ontological context necessary for interpreting the limitations identified in later sections of this chapter.

This is diagnosis, not critique. The task here is to identify what the paradigm assumes about reality. Questions concerning the adequacy of these assumptions, or the requirements of alternative ontological commitments, are addressed elsewhere.

I. The Five Core Ontological Assumptions

Within this framework, the mechanistic paradigm rests on five interlocking ontological assumptions that together constitute its hard core.

Assumption 1: Reductionism

The assumption.

Complex systems can be fully understood by decomposing them into constituent elements and analysing each element independently. The properties of the whole are reducible to the properties of the parts; aggregation by summation is valid.

In economics, this assumption appears in methodological individualism, representative-agent constructions, and the decomposition of organisations into individual decision-makers and transactions. Macroeconomic phenomena are treated as scaled-up micro behaviours; firms are represented as collections of contracts or production functions.

As Nason (2017, p. 16) observes, “complex systems cannot be separated into parts to be analyzed independently because their components are interconnected and interdependent... reductionism fails in complex systems as it ignores the emergent behaviors that arise from interactions among components.”

Implication.

Irreducible emergence is excluded by assumption; interactions either do not matter or are treated as analytically capturable through linear aggregation.

Assumption 2: Determinism

The assumption.

Given complete knowledge of a system’s current state and governing laws, its future states are fully determined. The world operates through fixed causal relations producing predictable effects.

In economic theory this appears in rational expectations, perfect foresight, and the treatment of uncertainty as an information deficit rather than a property of reality. The future is assumed to be knowable in principle.

Prigogine and Stengers (1984, p. 291) describe this worldview as one that “escapes the arrow of time,” thereby eliminating irreversibility and novelty.

Implication.

Fundamental unpredictability is ruled out; uncertainty is reduced to incomplete knowledge that can, in principle, be resolved.

Assumption 3: Substance Ontology

The assumption.

Reality is composed of stable, discrete substances with intrinsic and enduring properties. Being precedes becoming; process and history are secondary.

Within economics, firms are treated as fixed entities with stable production functions; technologies as discrete objects; preferences as given; identities as invariant. Change appears as deviation from equilibrium rather than transformation.

Prigogine and Stengers (1984) characterise this orientation as a science of being rather than becoming.

Implication.

Learning, capability development, and institutional evolution are analytically downgraded to disturbances around fixed structures.

Assumption 4: Closure

The assumption.

Economic systems are treated as closed, operating within fixed boundaries and lacking constitutive exchange with their environments. External influences enter only as parametric shocks.

General equilibrium models exemplify this assumption: all agents, goods, and relations are specified within the model, while the outside world appears only as an exogenous parameter set.

This reflects the Newtonian image of a hermetically sealed clockwork universe (Ackoff, 1993; Lifshitz, 1976).

Implication.

Co-evolution with the environment, boundary formation, and ecosystem dynamics are excluded by design.

Assumption 5: Atomism

The assumption.

Agents, firms, and markets are independent entities whose properties exist prior to their relations. Relations connect entities but do not constitute them.

This assumption underpins methodological individualism and the treatment of networks and institutions as secondary or external features of analysis.

As Hodgson (2001, p. 28) notes, social explanation is reduced to the properties of isolated individuals. Lawson (2019) demonstrates that this produces analytical incoherence in social systems.

Implication.

Relational identity, network position, and institutional embeddedness cannot be treated as fundamental.

II. The Interconnected Nature of These Assumptions

These five assumptions form a coherent ontological system rather than a set of independent claims. Together they constitute a mutually reinforcing meso-cluster.

Table 18: Logical Interdependencies Among Ontological Assumptions: The Mutually Reinforcing Network of Mainstream Theory of the Firm

Connection	Logic
Reductionism requires Atomism	Decomposition presupposes independent parts
Determinism requires Closure	Prediction presupposes fixed system boundaries
Substance requires Determinism	Stable essences imply predictable behaviour
Closure enables Reductionism	Fixed boundaries permit exhaustive decomposition
Atomism enables Substance	Independent entities must have intrinsic properties

This yields a mechanistic metaphysics system:

- Reductionism ↔ Atomism
- Determinism ↔ Closure
- Substance stabilises both pairs

The internal coherence of this system explains paradigmatic persistence: altering any one assumption destabilises the others.

III. Additional Ontological Commitments

From these five core assumptions follow several operational commitments:

- **Decomposability** – systems can be disassembled and reassembled without loss
- **Additivity** – aggregates are linear sums
- **Essentialism** – firms and agents possess fixed identities
- **Timeless ontology** – history does not matter
- **Static being over becoming**

Together, these commitments systematically eliminate process, learning, and path-dependence from analysis.

IV. Consequences for Firms

Table 19: From Ontological Assumptions to Firm-Level Blind Spots: How Eliminated CAS Properties Render Key Organisational Phenomena Invisible

Assumption	Eliminates	Firm-level consequence
Reductionism	Emergence	Organisational culture is invisible
Determinism	Fundamental uncertainty	Genuine innovation is excluded
Substance	Becoming	Learning and capability evolution disappear
Closure	Co-evolution	Ecosystems cannot be represented
Atomism	Relational constitution	Networks are analytically secondary

Firms, which operate at the meso-level of emergence and adaptation, become analytically marginal within this framework.

V. Why These Assumptions Matter

These assumptions are pre-empirical metaphysical commitments, inherited in large part from Newtonian physics (Mirowski, 1989). They shape what can be seen before any data are analysed. Once accepted, equilibrium, optimisation, and prediction follow as logical consequences rather than empirical discoveries.

Conclusion: The Mechanistic Metaphysics System

The five ontological assumptions – Reductionism, Determinism, Substance, Closure, and Atomism – form the mechanistic metaphysics system. Taken together, they generate an emergent property: the systematic elimination of phenomena central to complex adaptive systems, including emergence, learning, novelty, and relationality.

This ontological configuration provides the foundation for the lock-in mechanisms examined later in this chapter. The diagnosis is complete.

3.1.2 Core Epistemological Assumptions

(Certainty, Universal Laws, Monism, Perfect Information)

Introduction

Section 3.1.1 identified the ontological hard core of the mechanistic paradigm – what the inherited body of economic theory assumes about the nature of reality. Ontology, however, is not epistemically neutral. What one assumes exists conditions what one assumes can be known. This subsection therefore completes the identification of the paradigm's hard core by adding the epistemological layer.

Epistemological assumptions are defined here as foundational presuppositions concerning what can be known, how knowledge is acquired, and what counts as a valid explanation. They constitute the second layer of the mechanistic paradigm's hard core, flowing from and reinforcing its ontological commitments. They differ from ontological assumptions, which

concern what exists, and from methodological assumptions, which concern appropriate modes of inquiry.

This subsection performs four diagnostic functions.

First, it makes explicit the deep commitments underlying prediction, optimisation, and formalism – assumptions so embedded that they often appear as scientific necessity rather than theoretical choice.

Second, it shows how uncertainty is systematically treated as ignorance: a temporary condition resolvable through improved data rather than a constitutive feature of economic reality.

Third, it demonstrates that linearity, ergodicity, equilibrium, and monism function not merely as conclusions but as preconditions for admissible knowledge.

Fourth, it completes the diagnostic architecture required to understand the limits of mechanistic explanation.

The stakes are epistemological. These assumptions determine what counts as valid economic knowledge, which research questions are askable, and which explanatory forms are admissible. Prediction is elevated as the benchmark of science, while uncertainty is treated as an anomaly to be eliminated.

This subsection remains diagnostic rather than evaluative. Together with Section 3.1.1, it establishes the complete hard core of the mechanistic paradigm. Questions of epistemological adequacy or alternative foundations are addressed beyond the scope of the present chapter.

I. The Four Core Epistemological Assumptions

Within this framework, four epistemological assumptions define what can be known and how knowledge is acquired in the inherited body of economic theory. These assumptions flow directly from the ontological commitments identified in Section 3.1.1 and together constitute the epistemological dimension of the paradigm's hard core.

Assumption 1: Certainty – Knowledge Can Be Complete and Determinate

The assumption.

Economic reality is knowable with certainty in principle. Uncertainty is epistemic – reflecting incomplete information – rather than ontological. With sufficient data and computational power, complete knowledge is achievable.

Within economic theory this assumption manifests in perfect foresight models, rational expectations equilibria, and complete-market frameworks. Knight (1921, p. 233) famously distinguished measurable risk – where outcomes can be assigned probabilities – from true uncertainty, which cannot be quantified because situations are novel or unique. The inherited paradigm systematically collapses this distinction, treating all uncertainty as calculable risk.

As Bronk (2011, p. 5) notes, a defining characteristic of mechanistic economic reasoning is its tendency to overlook Knightian uncertainty. Slovik (2011, p. 432) similarly observes that uncertainty is treated as a homogeneous mass that is neither identifiable nor analytically tractable. The result is that fundamental uncertainty becomes invisible by definition.

Implication.

The future is treated as epistemically accessible, and prediction becomes the privileged aim of economic science.

Assumption 2: Universal Laws – Knowledge Is Context-Free

The assumption.

Economic laws are assumed to operate universally across contexts, analogously to physical laws. Regularities identified in one setting are presumed to apply across time, space, and institutional environments. Context is treated as secondary.

This assumption generates universal preference orderings, invariant utility functions, and context-independent production functions. Canonical forms – such as the Cobb–Douglas production function – are applied uniformly across sectors and historical settings. Behavioural regularities are presumed to transcend culture, institutions, and history.

Cohen (1985, p. 159) traces this commitment to the Newtonian synthesis, which sought universal laws applicable both “to the heavens and to the earth.” Economics imported this aspiration. As Akerlof (2020) argues, however, this orientation produces systematic “sins of omission,” excluding phenomena that resist formal generalisation. Alston et al. (2023, p. iv) similarly criticise mechanistic models for presuming clear, linear, and knowable causal relations, thereby neglecting the contingent and context-specific character of real economies.

Implication.

Local knowledge and contextual variation are analytically marginal; explanation proceeds via discovery of universal regularities.

Assumption 3: Monism – One Correct Model Exists

The assumption.

For any economic phenomenon, there exists one correct model that captures its true underlying structure. Multiple models reflect incomplete knowledge; scientific progress consists in convergence toward a single correct representation. Model pluralism is treated as a temporary condition rather than a necessity.

Within economics this assumption appears in the privileged status accorded to particular model families – such as DSGE frameworks in macroeconomics or Arrow–Debreu general equilibrium in microeconomics. Alternative approaches – evolutionary, institutional, post-Keynesian – are treated as inferior approximations or deviations rather than complementary perspectives.

Colander et al. (2009, p. 256) describe this as a “scholastic attachment” to unified frameworks, noting that rational expectations models – assuming agents share the economist’s model – remain empirically fragile yet persist because of the monistic commitment. Rodrik (2015, p. 29) cautions that the danger lies in mistaking a simplified model for the complex reality it seeks to represent.

Implication.

Methodological monism is regarded as appropriate; pluralism signals incompleteness rather than irreducible complexity.

Assumption 4: Perfect Information – Agents Know All Relevant Facts

The assumption.

Economic agents are assumed to possess complete information about alternatives, preferences, constraints, and market conditions, or to be able to acquire such information at known costs. Asymmetric information is treated as an anomaly rather than a constitutive feature of economic life.

Perfect competition presupposes perfect information; general equilibrium requires full knowledge of all markets; representative-agent models rely on informational symmetry to legitimise aggregation.

Hayek (1945, p. 520) criticised this assumption as an erroneous transfer of habits of thought developed for natural sciences into social analysis. Such transposition presumes stable relationships, complete information, and determinacy – conditions that do not characterise social reality. Arthur (2021, p. 136) similarly notes that conventional economics assumes agents face well-defined problems and that uncertainty stems from information gaps that could, in principle, be closed.

Implication.

Information problems are treated as technical and solvable rather than structural and constitutive.

II. How Epistemology Flows from Ontology

The Generative Relationship

The four epistemological assumptions are not independent. They arise directly from the ontological commitments identified in Section 3.1.1.

Table 20: Ontological-to-Epistemological Derivation: How Foundational Ontological Assumptions Generate Corresponding Epistemological Commitments

Ontological Assumption	→	Epistemological Assumption
Determinism	→	Certainty
Substance ontology	→	Universal Laws
Reductionism	→	Monism
Closure	→	Perfect Information

Logic of Generation

If reality is governed by deterministic laws, then future states are calculable in principle and uncertainty can only be epistemic.

If entities possess stable intrinsic properties, then the regularities they exhibit are presumed to be stable and universal.

If complex wholes are fully decomposable, then one sufficiently detailed model should capture the phenomenon.

If systems operate within fixed boundaries, then the inventory of relevant facts is finite and knowable.

Epistemology thus functions as ontology's shadow: assumptions about what exists determine assumptions about what can be known.

III. Additional Epistemological Commitments

From these four core assumptions follow additional operational commitments:

- **Epistemic determinism:** the future is knowable if the present is known (Laplace, 1902).
- **The predictive ideal:** explanation is equated with prediction; failure to predict is treated as scientific failure (Hayek, 1974).
- **Ergodicity:** statistical properties are time-invariant, enabling inference from past data (Davidson, 2012).
- **Risk–uncertainty equivalence:** all uncertainty is treated as quantifiable risk (Kay & King, 2020).
- **Correspondence truth:** models are evaluated as true or false representations rather than as context-bound analytical devices.

Together these commitments form a certainty–monism–prediction complex: an internally coherent epistemological system in which each element reinforces the others.

IV. Consequences for Understanding Firms

What This Epistemology Eliminates

Table 21: From Epistemological Assumptions to Firm-Level Blind Spots: How Eliminated CAS Properties Obscure Key Dimensions of Firm Behaviour

Epistemological Assumption	Eliminates	Firm-level consequence
Certainty	Fundamental uncertainty	Strategic surprise is excluded
Universal laws	Context-dependence	Industry specificity disappears
Monism	Complementary perspectives	Firm theories cannot be integrated
Perfect information	Bounded rationality	Judgment under ignorance is invisible

As a result, the inherited paradigm cannot represent strategic uncertainty as constitutive, industry context as formative, plural firm theories as complementary, or learning as an open-ended process.

V. Why This Matters for Firm Research Firm-level inquiry conducted within mechanistic epistemology is structurally constrained to assume away uncertainty, abstract from context, privilege a single theoretical lens, and prioritise prediction over process understanding. These constraints are not methodological choices but consequences of the epistemological hard core diagnosed above.

Diagnostic Closure This subsection has identified the epistemological hard core of the mechanistic paradigm and traced its consequences for economic and firm-level analysis. It does not propose alternative epistemological foundations or methodological remedies. Its purpose is strictly diagnostic: to make explicit the epistemic commitments that structure admissible knowledge within the inherited paradigm.

V. The Closed Epistemic Loop

Why the Paradigm Persists

Within this analytical framework, the mechanistic system exhibits a closed epistemic loop: ontology conditions epistemology, epistemology conditions methodology, and methodological success – defined on the paradigm’s own terms – feeds back to confirm the original ontological commitments. The result is a self-reinforcing architecture in which the criteria for admissible knowledge are generated by the same assumptions that the knowledge is then used to validate.

Self-sealing character

The epistemological assumptions described above render the paradigm self-sealing in four ways:

- **Certainty:** If uncertainty is treated as epistemic, then unpredictability associated with complex adaptive systems is reclassified as incomplete information. Predictive failure is interpreted as a call for better data, stronger models, or deeper refinement – not as evidence that prediction is an inappropriate epistemic ideal in such domains.
- **Universal laws:** If explanatory regularities are assumed to be context-free, then local breakdowns are treated as “noise,” “special circumstances,” or measurement error. Contextual failures thus reinforce the search for more general laws rather than legitimising context-dependence as constitutive.
- **Monism:** If one correct model is presumed to exist, alternative frameworks are discounted by definition. Competing approaches are treated as inferior approximations or temporary detours rather than as potentially complementary representations of distinct aspects of complex phenomena.
- **Perfect information:** If relevant information is assumed knowable in principle, bounded rationality becomes suboptimal deviation from the norm. Realistic cognition

is interpreted as imperfection requiring correction or modelling “frictions,” rather than as a constitutive feature of economic agency.

Meso-level operation

This loop is not merely conceptual. It is reproduced at the meso-level of disciplinary practice and institutional reinforcement. Graduate training, journal standards, promotion criteria, and funding priorities embody these epistemological commitments, thereby stabilising them across cohorts. In this sense, assumptions are not primarily held by isolated individuals; they are embedded in the institutional ecology of the discipline.

The lock-in mechanisms through which this persistence is reproduced are examined in Section 3.2.

VI. Counter-Arguments Addressed

“Economists know models are simplifications.”

Objection. Economists do not literally believe in perfect information or certainty; these are simplifying devices adopted for tractability.

Response. Simplifications that exclude the very phenomena under investigation are not benign abstractions; they are constitutive exclusions. When “as if” reasoning becomes systematic across a paradigm, it shapes what questions are askable and what counts as explanation. If the authorised framework can only represent certainty-like environments, fundamental uncertainty becomes analytically invisible – not through refutation, but through definition.

“Behavioural economics addresses bounded rationality.”

Objection. Behavioural economics incorporates bounded rationality, information asymmetries, and uncertainty; the paradigm has evolved.

Response. Behavioural findings are frequently treated as catalogued anomalies rather than as grounds for reconstituting the epistemological hard core. Core modelling practices retain strong informational and optimisation baselines; departures are treated as “biases” relative to a normative benchmark. The epistemological structure remains intact: bounded rationality is framed as imperfection rather than constitutive condition.

“Risk models handle uncertainty.”

Objection. Financial economics possesses sophisticated tools that quantify uncertainty via probability distributions.

Response. Such quantification presupposes a known outcome space and stable statistical structure. This addresses risk, not Knightian uncertainty. Treating uncertainty as probabilisable risk collapses a crucial distinction and can produce false confidence – particularly in regimes characterised by novelty, structural change, or endogenous instability.

“Prediction works sometimes.”

Objection. Economic predictions sometimes succeed; this validates the predictive ideal.

Response. Episodic success in relatively stable regimes does not validate the predictive ideal as a universal epistemic criterion. Predictive breakdown in transitions – crises, innovation

waves, structural change – reveals the limits of the framework’s epistemological commitments. In complex domains, the appropriate ambition is often pattern-conditional understanding rather than point prediction.

Conclusion: The epistemological hard core identified

This subsection has identified four core epistemological assumptions that complement the ontological hard core: Certainty, Universal Laws, Monism, and Perfect Information. Together with the five ontological assumptions diagnosed in Section 3.1.1, they constitute the complete hard core of the mechanistic paradigm.

Key findings summary

Table 22: Epistemological Assumptions of Mainstream Theory of the Firm: Core Claims and the CAS Phenomena They Structurally Eliminate

Assumption	Core claim	Eliminates
Certainty	Knowledge is complete in principle	Fundamental uncertainty
Universal laws	Regularities are context-free	Local knowledge; context-dependence
Monism	One correct model exists	Complementary perspectives
Perfect information	Relevant facts are knowable	Bounded rationality

Flow from ontology These epistemological assumptions follow from ontological commitments diagnosed earlier:

- Determinism → Certainty
- Substance ontology → Universal laws
- Reductionism → Monism
- Closure → Perfect information

Because epistemology is generated by ontology, altering one layer without the other produces only local adjustment rather than structural change.

The closed loop The mechanistic system forms a closed epistemic loop: ontology conditions epistemology, epistemology conditions method, and method – evaluated through paradigm-internal criteria – feeds back to confirm ontology. This self-reinforcing architecture helps explain persistence despite recurring empirical and conceptual breakdowns.

Forward orientation

The next subsection (Section 3.1.3) shows how the ontological and epistemological hard core generates a further layer of derivative assumptions (e.g., equilibrium, optimisation, linearity) that operationalise the paradigm’s analytical apparatus.

The purpose of the diagnostic maps and interconnections in this chapter is explanatory rather than operational: they clarify how the paradigm coheres and persists, not how to implement an alternative framework.

3.1.3 Derivative Assumptions

(Equilibrium, Linearity, Homogeneity, Ergodicity, Optimisation, Decomposability)

Introduction

Sections 3.1.1 and 3.1.2 identified the ontological and epistemological hard core of the mechanistic paradigm. These core assumptions do not remain abstract. They generate a further layer of assumptions that permeate formal theory, modelling practice, and empirical methodology. This subsection identifies that cascade.

Derivative assumptions are defined here as assumptions that are not primitive but follow logically from the ontological and epistemological core. They constitute the third layer of the mechanistic paradigm's architecture. Although they often appear technical or methodological, they are not optional. Given the core assumptions, these derivatives cannot be otherwise. Altering them in isolation does not constitute reform; it triggers regeneration.

This subsection performs three diagnostic functions.

First, it identifies six derivative assumptions: Equilibrium, Linearity, Homogeneity, Ergodicity, Optimisation, and Decomposability.

Second, it establishes the cascade logic through which the core assumptions necessarily generate these derivatives.

Third, it explains why piecemeal modification repeatedly fails: if derivative assumptions follow logically from the core, then relaxing them while preserving the core is structurally impossible.

Together, the five ontological, four epistemological, and six derivative assumptions constitute a 15-assumption architecture that defines the mechanistic paradigm.

I. The Six Derivative Assumptions

Derivative 1: Equilibrium – Systems Tend Toward Stable Rest States

Assumption.

Economic systems are assumed to converge toward stable equilibrium states. Disturbances are temporary; equilibrium is treated as the normal condition.

Generated by:

Substance + Determinism + Closure + Certainty

Stable entities governed by fixed laws within closed systems are expected to settle into rest states that are knowable in principle.

In economics this assumption appears in Walrasian equilibrium, Nash equilibrium, rational expectations equilibrium, and market-clearing conditions. Persistent disequilibrium is treated as failure or transience rather than as a normal state.

Equilibrium is therefore not an empirical discovery but a logical consequence of the core assumptions.

Derivative 2: Linearity – Effects Are Proportional to Causes

Assumption.

Causal relations are linear: small causes produce small effects; large causes produce proportionally larger effects.

Generated by:

Reductionism + Determinism + Universal Laws

If wholes are sums of parts and governing laws are constant, proportional responses follow.

This assumption licenses calculus, marginal analysis, and comparative statics. Non-linear phenomena – tipping points, cascades, regime shifts – are relegated to anomalies.

Linearity thus removes interaction-generated novelty from analysis.

Derivative 3: Homogeneity – Agents Are Essentially Identical

Assumption.

Economic agents can be treated as structurally identical; differences are parametric rather than ontological. Aggregation via representative agents is legitimate.

Generated by:

Reductionism + Atomism + Universal Laws + Perfect Information

Independent units governed by the same laws and possessing the same information must be functionally equivalent.

This assumption produces the representative consumer, firm, and bank – and eliminates meso-level heterogeneity.

Derivative 4: Ergodicity – The Future Statistically Resembles the Past

Assumption.

Statistical properties are time-invariant; history does not matter for probability.

Generated by:

Substance + Determinism + Certainty + Closure

Stable entities governed by fixed laws in bounded systems generate stable distributions.

This assumption underwrites econometric inference and risk modelling. If ergodicity fails, prediction loses its epistemic foundation.

Complex adaptive systems violate ergodicity through endogenous novelty.

Derivative 5: Optimisation – Agents Maximise

Assumption.

Economic agents maximise objective functions subject to constraints. Behaviour is modelled as the solution to a calculus problem.

Generated by:

Reductionism + Determinism + Perfect Information + Certainty + Universal Laws

Optimisation presupposes complete knowledge of objectives, constraints, and outcomes.

This is not a claim about rationality but about epistemological certainty. Without certainty, optimisation is undefined.

Derivative 6: Decomposability – Systems Can Be Analysed Part by Part

Assumption.

Systems can be decomposed into parts without loss; parts can be analysed independently and recombined.

Generated by:

Reductionism + Atomism + Closure + Linearity

This enables partial equilibrium analysis, ceteris paribus reasoning, and modular intervention.

Complex adaptive systems are non-decomposable: parts are constituted by relations.

II. The Cascade Logic

The six derivative assumptions are generated by the core as follows:

Table 23: Generative Mapping of Derivative Assumptions: Tracing Each to Its Foundational Ontological and Epistemological Sources

Derivative	Generated by
Equilibrium	Substance + Determinism + Closure + Certainty
Linearity	Reductionism + Determinism + Universal Laws
Homogeneity	Reductionism + Atomism + Universal Laws + Perfect Information
Ergodicity	Substance + Determinism + Certainty + Closure
Optimisation	Reductionism + Determinism + Perfect Information + Certainty
Decomposability	Reductionism + Atomism + Closure + Linearity

This cascade explains:

- internal coherence,
- paradigmatic persistence,
- and the repeated failure of piecemeal modification.

The core regenerates the derivatives.

III. The Double Function of Derivatives

Internal enforcement

Derivative assumptions **operationalise and stabilise** the core:

- Equilibrium reinforces substance
- Optimisation reinforces certainty
- Linearity reinforces reductionism

Methodological shaping

Table 24: The Dual Role of Derivative Assumptions: Analytical Techniques Enabled and CAS-Relevant Phenomena Excluded

Derivative	Enables	Excludes
Equilibrium	Fixed-point theorems	Persistent disequilibrium
Linearity	Calculus	Tipping points
Homogeneity	Aggregation	Networks
Ergodicity	Econometrics	History
Optimisation	Lagrangians	Heuristics
Decomposability	Partial equilibrium	Systems thinking

Methods therefore presuppose the assumptions they appear to test.

IV. Why Piecemeal Modification Fails

Because derivative assumptions are logically generated, selective relaxation triggers regeneration:

- Add heterogeneity → homogeneity reappears
- Add dynamics → equilibrium reasserts itself
- Add uncertainty → certainty is reintroduced as risk
- Add non-linearity → linearisation returns

This pattern explains the persistence of anomalies within the paradigm.

V. Consequences for Firms

Table 25: Firm-Level Phenomena Eliminated by Derivative Assumptions: What Mainstream Theory Cannot Represent

Derivative	Eliminates
Equilibrium	Continuous adaptation
Linearity	Disruption
Homogeneity	Strategic diversity
Ergodicity	Path-dependence
Optimisation	Satisficing
Decomposability	Organisational emergence

As a result, firms – located at the meso-level of adaptation and coordination – become analytically marginal.

Conclusion

This subsection has identified the derivative layer of the mechanistic paradigm and traced its logical generation from the ontological and epistemological core.

Table 26: Summary of the Mechanistic Assumption Network: Distribution of Fifteen Assumptions Across Ontological, Epistemological, and Derivative Layers

Layer	Assumptions
Ontological (3.1.1)	5
Epistemological (3.1.2)	4
Derivative (3.1.3)	6
Total	15

The derivative assumptions are not optional. They are logical consequences of the core. The diagnosis is complete.

3.1.4 The Assumption Network

Introduction

Sections 3.1.1–3.1.3 identified the mechanistic paradigm’s assumptive architecture as a set of five ontological assumptions, four epistemological assumptions, and six derivative assumptions. Taken individually, these assumptions already delimit what can be represented within the inherited body of economic theory. Taken together, however, they form something stronger than a list. They constitute a networked architecture whose internal structure explains both the paradigm’s coherence and its persistence.

This subsection completes the diagnostic task of Chapter 3 by reconstructing that architecture explicitly. Its purpose is not to introduce additional assumptions, but to show how the identified assumptions interlock, reinforce one another, and collectively eliminate entire domains of economic phenomena from admissible analysis. In particular, it demonstrates that the systematic absence of meso-level analysis – firms, industries, networks, and institutional fields – is not accidental or contingent, but a direct consequence of the way the assumption network is organised.

Part A: The Network Character of the Paradigm

The mechanistic paradigm is best understood as a three-tier architecture in which foundational assumptions generate derivative constraints, which in turn give rise to characteristic modelling and evaluative practices. These tiers are analytically distinguishable but causally inseparable.

At the foundational level, ontological and epistemological assumptions define what exists and what can be known. These commitments generate, with logical necessity, a second layer of derivative assumptions that operationalise the core into tractable analytical form. Together, these two layers produce a methodological regime that appears as “standard practice” within the discipline. Crucially, this regime feeds back into the system by rewarding results that conform to its own presuppositions, thereby stabilising the entire structure.

The result is not a hierarchy but a self-reinforcing network.

Foundational and Derivative Linkages

The internal logic of the network becomes visible when derivative assumptions are examined in relation to the core from which they arise. Table 3.1 summarises these generative relationships.

Table 27: 3.1 Derivative assumptions and their generative foundations

Assumption	Generated by
Equilibrium	Substance + Determinism + Closure + Certainty
Linearity	Reductionism + Determinism + Universal Laws
Homogeneity	Reductionism + Atomism + Universal Laws + Perfect Information
Ergodicity	Substance + Determinism + Certainty + Closure
Optimisation	Reductionism + Determinism + Perfect Information + Certainty
Decomposability	Reductionism + Atomism + Closure + Linearity

This table does not describe modelling choices. It expresses logical dependence. Once the foundational assumptions are accepted, the derivative assumptions follow as necessary consequences rather than optional techniques. For this reason, attempts to relax or modify individual derivatives without altering the core repeatedly fail. The core assumptions regenerate the derivative layer, restoring the original configuration.

From Assumptions to Methodological Practice

The assumption network does not remain at the level of abstract commitments. It manifests in a recognisable methodological regime that structures modelling practice, evaluation criteria, and standards of rigour within the discipline. These practices often appear neutral or purely technical, but they are the observable surface of the deeper assumptive architecture.

Table 3.2 illustrates how characteristic methodological symptoms emerge from combinations of underlying assumptions.

Table 28: 3.2. Methodological symptoms and their emergence

Symptom (practice)	How it emerges
Dominance of representative-agent models	Homogeneity + Optimisation + Monism
Equilibrium-centred analysis	Substance + Determinism + Closure
Formal optimisation as evaluative benchmark	Linearity + Certainty
Preference for universal models	Universal Laws + Monism
Marginalisation of history and path-dependence	Ergodicity + Certainty

These practices are not independent methodological preferences. They are enforced by the assumption network itself. Models that conform to these expectations appear rigorous, elegant, and publishable because they resonate with the paradigm's internal coherence. Models that violate them struggle for legitimacy, not necessarily because they are incoherent, but because they challenge the conditions of admissibility defined by the network.

In this sense, methods do not test assumptions. They presuppose them.

The Self-Stabilising Architecture

The interdependence of assumptions across tiers gives the mechanistic paradigm a self-stabilising character. Foundational commitments generate derivative constraints; derivative constraints shape methodological practice; methodological practice, in turn, produces results that confirm the plausibility of the foundational commitments. Apparent success within the paradigm is therefore endogenous to the paradigm itself.

This feedback structure explains why reform efforts that target individual assumptions or techniques tend to produce surface variation rather than structural change. When a derivative assumption is relaxed, compensatory mechanisms elsewhere in the network reassert the original configuration. Disequilibrium becomes transitional rather than constitutive; uncertainty is reframed as risk; heterogeneity is parameterised rather than structural; non-linearity is linearised.

The architecture absorbs challenge without transforming.

The Missing Middle: Structural Elimination of the Meso-Level

One consequence of the assumption network is particularly consequential: the systematic elimination of the meso-level of analysis. Firms, industries, networks, organisational fields, and institutional arrangements do not merely receive insufficient attention within the mechanistic paradigm. They are rendered analytically redundant by design.

Reductionism and atomism collapse explanation to independent agents. Homogeneity removes structural differentiation. Equilibrium and decomposability dissolve intermediate structures into either micro-components or macro-outcomes. Ergodicity and certainty strip history and emergence of explanatory relevance. Under these conditions, intermediate structures add no explanatory value and therefore disappear from admissible analysis.

The economy is reduced to a micro-to-macro jump. The level at which coordination, innovation, power, learning, and adaptation actually occur is bypassed.

This is not an oversight. It is the logical outcome of the assumption network reconstructed in this chapter.

IV. Feedback Loops and Self-Stabilisation

The assumption network reconstructed above is not merely hierarchical. It contains a set of self-reinforcing feedback loops that stabilise the mechanistic paradigm against anomaly, critique, and partial modification. These loops operate across ontological, epistemological, and methodological layers, producing a system that absorbs challenge without undergoing structural transformation.

These feedback relations are summarised analytically in Table 3.3.

Table 29: 3.3. Feedback loops within the assumption network

Loop	Components
Ontological lock-in	Reductionism ↔ Equilibrium ↔ Closure
Epistemic lock-in	Determinism ↔ Ergodicity ↔ Optimisation
Methodological lock-in	Equilibrium ↔ Optimisation ↔ Homogeneity ↔ Decomposability
History elimination	Certainty ↔ Ergodicity ↔ Universal laws
Regeneration loop	Core assumptions → derivatives → methods → results → core assumptions

These loops explain why the paradigm remains stable even when confronted with persistent empirical anomalies. Ontological commitments stabilise epistemological expectations; epistemological expectations stabilise methodological choices; methodological choices generate results that appear to confirm the original ontological commitments. Anomalies are therefore interpreted as local failures of execution rather than as signals of structural inadequacy. The paradigm does not collapse under pressure; it reasserts itself.

Why Piecemeal Reform Fails

Because the mechanistic paradigm is organised as a network rather than a set of independent assumptions, selective modification consistently fails to produce durable change. Attempts to introduce heterogeneity while preserving atomism and universal laws result in heterogeneity being reinterpreted as noise. Efforts to relax equilibrium while retaining determinism and closure reframe disequilibrium as transitional rather than constitutive. Attempts to introduce uncertainty within a certainty-based epistemology convert uncertainty into calculable risk. Efforts to accommodate non-linearity within a reductionist framework result in linearisation around special cases.

In each instance, the foundational commitments remain intact, and the derivative assumptions regenerate accordingly. What appears as reform at the surface is absorbed by the network and translated back into its original structure. This pattern is not accidental. It is the predictable outcome of a system in which derivative assumptions are logically generated rather than freely chosen.

Part B: The Missing Middle – Systematic Elimination of the Meso-Level

The Micro-to-Macro Jump

A constitutive error runs through the inherited paradigm: the systematic absence of meso-level analysis. This absence is not the result of neglect or oversight. It follows directly from the assumption network reconstructed in this chapter.

Within representative-agent logic, heterogeneous individuals are aggregated into a single representative agent, whose optimisation problem is then solved and scaled to the macro level. The analytical jump from micro to macro bypasses the level at which most economic coordination, innovation, and adaptation actually occur. As Arthur (2013, p. 14) observes, the most important economic dynamics unfold at the mesoscale, encompassing networks, industries, supply chains, institutions, and organisational fields.

The paradigm's architecture leaves no analytical space for this level to appear as constitutive.

What Is Lost When the Meso-Level Disappears

When meso-level structures are eliminated, the inherited framework becomes blind to a wide range of economic phenomena. Network topology and cascade propagation cannot be represented when relations are treated as external. Organisational routines, culture, and coordination mechanisms disappear when firms are reduced to production functions or contracts. Industry structure and competitive dynamics lose explanatory status when homogeneity and equilibrium dominate. Institutional matrices and norms are relegated to constraints rather than generative structures. Technological regimes and trajectories become exogenous shocks rather than endogenous processes. Power relations and coalition dynamics vanish entirely.

Many economic crises propagate not through individual optimisation errors or aggregate variable movements, but through meso-level architectures such as financial networks, supply chains, and institutional interdependencies. A framework that cannot represent these structures is necessarily blind to their dynamics.

Pathologies Produced by Meso-Elimination

The elimination of the meso-level produces three systematic pathologies. First, macro-outcomes are treated as simple aggregations of micro behaviour, despite the fact that intermediate structures generate emergent properties and can exert downward constraint. Second, macro aggregates are treated as explanatory primitives, while the mechanisms that generate them remain invisible, producing black-box holism and missing transmission channels. Third, the paradigm lacks the conceptual and methodological tools required for multi-level explanation, forcing analysis into a binary choice between micro and macro and leaving intermediate structures untheorised.

These pathologies are not failures of implementation. They are consequences of design.

Why the Meso-Level Is Eliminated by Design

The systematic disappearance of the meso-level follows from the joint operation of the paradigm's core and derivative assumptions. Table 3.4 summarises the elimination mechanisms involved.

Table 30: 3.4. Assumptions and meso-level elimination mechanisms

Assumption	Elimination mechanism
Reductionism	Intermediate levels are redundant if wholes equal sums
Atomism	Relationally constituted structures cannot be fundamental
Homogeneity	Structural differentiation disappears
Universal laws	Context-specific meso-structures become irrelevant

Taken together, these mechanisms yield an analytical architecture incapable of representing meso-level phenomena as constitutive elements of economic reality. The absence of firms, networks, and institutions as generative structures is therefore not a gap to be filled, but an outcome to be explained.

The Centrality of the Meso-Level in Economic Reality

Markets, innovation processes, institutional evolution, and crises are multi-scalar phenomena, but they typically operate through meso-level structures. The propagation of financial distress through interconnected balance sheets, the diffusion of technological standards through industry networks, and the coordination of production through organisational routines all illustrate this point. These processes cannot be reduced to individual behaviour or captured by aggregate variables alone. They require explicit attention to intermediate structures and relations.

A paradigm that systematically excludes the meso-level is therefore blind to many of the economy's most consequential mechanisms.

Conclusion

This subsection has reconstructed the mechanistic paradigm as a self-stabilising assumption network composed of fifteen interdependent assumptions organised across ontological, epistemological, and derivative layers. It has shown how this network generates reinforcing

feedback loops, resists piecemeal modification, and systematically eliminates meso-level phenomena from admissible analysis.

The reconstruction is intentionally diagnostic. It explains how the paradigm coheres, why it persists, and what it renders invisible, without proposing an alternative framework or prescriptive remedy. In doing so, it completes the anatomical analysis required to understand the strengths and limitations of the mechanistic approach.

3.2 The Paradigmatic Lock-In: Why Piecemeal Reform Fails

Section 3.1 reconstructed the mechanistic paradigm as a fifteen-assumption architecture organised across ontological, epistemological, and derivative layers. The diagnostic conclusion was that this architecture is structurally incompatible with the properties of complex adaptive systems – it eliminates the meso-level where firms operate, suppresses emergence and learning, and renders novelty invisible. A puzzle therefore arises: if these assumptions are untenable under conditions of complexity, why do they persist despite decades of anomalies, crises, and sustained critique?

Section 3.2 resolves this puzzle. The answer is not ignorance, habit, or intellectual inertia. It is structure: the paradigm locks itself in place through mechanisms that absorb challenge and regenerate coherence.

The argument proceeds in four stages. Subsection 3.2.1 identifies three reinforcing feedback loops – ontological, epistemic, and structural – that interlock at high-centrality nodes such as equilibrium and optimisation, producing a triple-locked system where challenge to any assumption activates resistance across others. Subsection 3.2.2 generalises from loop-level reinforcement to paradigm-level self-stabilisation, showing how the lock-in is reproduced through semantic, methodological, and institutional mechanisms that embed assumptions in vocabulary, gatekeep through method, and reward conformity. Subsection 3.2.3 applies Kuhn's theory of normal science, explaining why anomalies can accumulate indefinitely without triggering paradigm revision. Subsection 3.2.4 applies Lakatos's evaluative framework to render the formal verdict: the programme is degenerating.

Understanding this lock-in is essential for what follows. Without it, the methodological symptoms examined in Section 3.3 would appear as correctable habits rather than necessary expressions of the architecture, and the consequences documented in Section 3.4 would appear as contingent failures rather than structural outcomes.

3.2.1 Reinforcing Feedback Loops Between Assumptions

Introduction

Section 3.1 reconstructed the mechanistic paradigm as a network of fifteen interdependent assumptions organised across ontological, epistemological, and derivative layers. The diagnostic conclusion of that section was that these assumptions are structurally incompatible with the properties of complex adaptive systems. A further question therefore arises: if these assumptions are untenable under conditions of complexity, why do they persist despite repeated anomalies, crises, and sustained critique?

Within the present framework, persistence is not explained by ignorance, habit, or intellectual inertia alone. It is explained by structure. The assumptions identified in Section 3.1 are not merely mutually consistent; they are organised into reinforcing feedback loops that stabilise

the paradigm against challenge. Each assumption supports others, and attempts to relax individual elements trigger compensatory adjustments elsewhere in the system. The paradigm is therefore not simply incorrect in isolated respects; it is self-locking.

This subsection completes the diagnostic analysis by reconstructing the feedback architecture through which paradigmatic lock-in is produced. It identifies three circular structures – ontological, epistemic, and structural – through which assumptions reinforce one another. It then shows how these loops interlock at shared nodes, generating redundant stabilisation. Finally, it explains why recurrent reform strategies fail in predictable ways: they leave the loop architecture intact and therefore reproduce the paradigm they seek to modify.

Reinforcing Feedback Structures in the Mechanistic Paradigm

The stability of the mechanistic paradigm does not arise from a single circular dependency, but from the interaction of multiple loops operating at different levels of the architecture. These loops are analytically distinct but causally intertwined. Together they form a redundant stabilisation system in which challenge to one element activates defensive responses across others.

For analytical clarity, the principal feedback structures are summarised in Table 3.5.

Table 31: 3.5. Reinforcing feedback loops in the mechanistic paradigm

Loop	Core components
Ontological	Reductionism ↔ Equilibrium ↔ Closure
Epistemic	Determinism ↔ Ergodicity ↔ Optimisation
Structural	Perfect rationality ↔ Representative agent ↔ Equilibrium

These loops do not represent optional modelling strategies. They are circular dependencies in which each component both enables and requires the others.

The Ontological Loop: Reductionism, Equilibrium, and Closure

The first feedback loop operates at the ontological level. Reductionism asserts that wholes can be decomposed into independent parts without loss. This decomposition renders equilibrium analysis analytically coherent, since independent components interacting through well-specified relations can be represented as stabilising when forces balance. Equilibrium, in turn, presupposes closure: stable rest states are meaningful only if the system is not constitutively perturbed by ongoing exchange with its environment. Closure then reinforces reductionism by making decomposition appear safe and exhaustive, since boundaries are fixed and the inventory of relevant elements is presumed complete.

The loop is self-sealing. Abandoning equilibrium undermines reductionism by reintroducing emergence and non-decomposability. Abandoning closure undermines equilibrium by allowing constitutive openness and continual flux. Abandoning reductionism undermines closure by rendering boundaries permeable, contested, and endogenously formed. Any attempt to relax one element therefore threatens the coherence of the others, triggering resistance.

At the level of the firm, this ontological loop sustains the treatment of organisations as production functions operating within fixed market boundaries and converging toward stable performance states. Even when dynamic language is introduced, change is framed as adjustment toward a destination rather than as open-ended transformation. The loop preserves

a static ontology precisely where firms exhibit emergent routines, evolving boundaries, and path-dependent capabilities.

The Epistemic Loop: Determinism, Ergodicity, and Optimisation

A second loop operates at the epistemological level. Determinism asserts that outcomes follow from states and laws. This commitment licenses ergodicity: if laws are stable, statistical regularities are time-invariant and the past provides reliable guidance about the future. Ergodicity, in turn, renders optimisation meaningful. If probability distributions are stable and outcome spaces known, agents can optimise across well-defined alternatives. Optimisation then presupposes determinacy in a practical sense, since optimal choice requires stable mappings from action to outcome.

The loop closes by mutual reinforcement. Challenging determinism destabilises ergodicity by undermining the assumption of stable distributions. Challenging ergodicity destabilises optimisation by dissolving the conditions under which optimal trajectories can be defined. Challenging optimisation undermines determinism's operational role by foregrounding satisficing, judgment, and novelty.

In applied contexts, this epistemic loop explains why uncertainty is repeatedly reclassified as risk, learning is framed as convergence, and surprise is treated as tail behaviour rather than as evidence of structural indeterminacy. The loop systematically eliminates the epistemic conditions under which novelty and genuine uncertainty could be represented.

The Structural Loop: Rationality, Aggregation, and Equilibrium

A third loop stabilises the micro-to-macro bridge. Perfect rationality, operationalised through perfect information, implies that agents optimise in comparable ways. This homogeneity legitimises aggregation into a representative agent, since differences are treated as irrelevant for aggregate behaviour. The representative agent's optimisation then produces equilibrium as the macro outcome. Equilibrium, in turn, presupposes rationality, since coordination and closure of the model depend on agents acting as if they solve the global problem.

The circularity is complete. Relaxing perfect rationality introduces heterogeneous decision rules that undermine representative aggregation. Abandoning the representative agent forces explicit modelling of meso-level structures such as networks and organisations, which destabilises equilibrium as a general attractor. Abandoning equilibrium reintroduces coordination failure and far-from-equilibrium dynamics that the rationality assumption is designed to suppress.

As a result, heterogeneity is repeatedly absorbed as parameter dispersion, bounded rationality is reframed as friction, and multiple equilibria preserve equilibrium as the organising concept. Structural change is displaced rather than confronted.

Interlocking and Redundant Stabilisation

These three loops do not operate independently. They interlock at shared nodes, producing redundant stabilisation. Equilibrium appears simultaneously as the ontological rest state of a closed decomposed system and as the macro-outcome of representative-agent optimisation. Optimisation appears simultaneously as the epistemic rule justified by ergodic predictability and as the microfoundation required for aggregation and equilibrium. Reductionism and determinism, as ontological commitments, underwrite both nodes across the architecture.

The result is a triple-locked system. Challenge to equilibrium or optimisation activates resistance across multiple loops simultaneously. Apparent reform in one domain is buffered by compensatory adjustments in others.

The Paradigm as a Closed Attractor

Taken together, the feedback loops generate a system that behaves as a closed attractor. Perturbations are absorbed, anomalies are reclassified, and the paradigm returns to its stable configuration. When an assumption is challenged, adjacent assumptions defend it; other loops compensate; the anomaly is renamed as friction, shock, imperfection, or measurement error; and the original structure is reinstated in refined technical form.

Change appears continuous at the surface, while the deep architecture remains intact.

Why Piecemeal Reform Reproduces the Paradigm

The loop architecture explains the discipline's recurring reform pattern. Modifications occur in the derivative layer while foundational commitments remain untouched. Approaches that genuinely disrupt the loops struggle for integration because they break coherence rather than extend it. Technical sophistication increases, but equilibrium, optimisation, and ergodicity continue to function as organising nodes. Complexity accumulates without transformation.

This pattern is not accidental. It is the predictable outcome of a system organised around mutually reinforcing assumptions.

Conclusion

This subsection has shown that the mechanistic paradigm is stabilised by three reinforcing feedback loops operating across ontological, epistemological, and structural layers. These loops interlock at high-centrality nodes, producing a redundant lock-in structure that absorbs challenge and reproduces coherence. Piecemeal reform fails not because of poor execution, but because the architecture regenerates what is relaxed.

The diagnosis is therefore structural. Understanding why the paradigm persists requires understanding how it locks itself in place.

3.2.2 The Self-Stabilising Nature of Paradigms

Introduction

Section 3.2.1 identified three reinforcing feedback loops that stabilise the mechanistic paradigm in a coherent configuration: an ontological loop linking reductionism, equilibrium, and closure; an epistemic loop linking determinism, ergodicity, and optimisation; and a structural loop linking rationality, representative-agent aggregation, and equilibrium. These loops interlock at high-centrality nodes – most notably equilibrium and optimisation – creating redundant stabilisation. Paradigmatic resistance to change, however, is not merely the passive durability of a coherent system. The mechanistic worldview actively stabilises itself through semantic, methodological, and institutional mechanisms that absorb perturbations and reproduce the paradigm across generations.

This subsection generalises from loop-level reinforcement to paradigm-level self-stabilisation. It shows how critique is neutralised not only by circular dependence within theory, but by the broader ecology of disciplinary language, authorised methods, and professional incentives. The outcome is a research programme that can appear responsive – through the incorporation of “fixes” – while remaining structurally unchanged. The diagnostic stakes are therefore substantial. If the paradigm is self-stabilising, then decades of reform may be reinterpreted less as cumulative transformation than as systematic accommodation within a persistent attractor. As North (2005, p. 6) observed, the inherited body of economic theory often proceeds as though uncertainty were exceptional, rather than a defining feature of economic life.

I. From Feedback Loops to Paradigm Self-Stabilisation

The feedback loops identified in Section 3.2.1 do not operate as isolated circuits. They are embedded within a redundant network in which each assumption can be defended through multiple pathways. Redundancy is crucial. In a redundant network, the destabilisation of one assumption does not propagate straightforwardly; it is rerouted through alternative justificatory links that preserve the overall configuration. The paradigm therefore exhibits defence-in-depth: challenge activates compensatory explanations rather than revision at the assumption level.

This dynamic is visible when central assumptions are challenged. Persistent disequilibrium, structural change, and endogenous novelty do not necessarily weaken equilibrium as an organising concept; instead equilibrium is repositioned as an “analytical device,” or disequilibrium is reframed as transitional behaviour, or the concept is preserved through an expanded equilibrium set. Similarly, critiques of representative-agent aggregation do not necessarily dislodge aggregation; heterogeneity is admitted but reinterpreted as distributional detail, or shown to be “approximately irrelevant” for aggregates, or assimilated into models designed to reproduce representative-agent results under alternative parametrisations.

The pattern is consistent. Assumptions are rarely falsified in practice because they are not held as isolated propositions. They operate as a mutual protection system. Critique is frequently met not with abandonment but with redescription. The triple lock therefore functions not merely as a set of loops, but as a redundant defence architecture.

II. Three Stabilisation Mechanisms

Paradigm persistence is reproduced across three coupled dimensions: semantic stabilisation, methodological stabilisation, and institutional stabilisation. These mechanisms convert theoretical coherence into durable disciplinary reproduction.

A. Semantic stabilisation

A paradigm stabilises itself semantically when its vocabulary embeds its ontology and epistemology in terms that structure what can be meaningfully said. Language does not merely describe phenomena; it encodes permissible forms of explanation. Within mechanistic economics, key terms carry implicit assumptions about normality and deviation, endogeneity and exogeneity, reason and error, and correct baselines and distortions. Terms such as equilibrium, shock, friction, and distortion organise perception in ways that render alternative ontological commitments linguistically awkward or easily misclassified.

For analytical clarity, Table 3.6 summarises representative examples.

Table 32: 3.6. Semantic stabilisation through embedded assumptions

Term	Stabilising effect
Equilibrium	Positions stability as the normal reference state and recasts disequilibrium as deviation
Rationality	Equates reason with optimisation and frames alternatives as deficient
Shock	Reclassifies change as exogenous disturbance rather than endogenous novelty
Friction	Treats a frictionless benchmark as “natural” and casts coordination costs as anomaly
Distortion	Implies a normatively correct baseline and frames institutions and power as deviations

At the level of the firm, such vocabulary has specific consequences. Speaking in production-function terms makes firms appear as transformation devices rather than coordination systems. The language of transaction costs and agency costs often presupposes an underlying frictionless benchmark in which optimisation is the default. Boundary talk may presuppose closure rather than the contested permeability typical of organisational ecosystems. In short, semantic stabilisation constrains what can be seen, especially at the meso-level where routines, culture, capabilities, networks, and institutional embedding are constitutive.

B. Methodological stabilisation

A paradigm stabilises itself methodologically when its authorised tools presuppose its foundational assumptions. In this case, method becomes gatekeeping: what cannot be expressed in the authorised formalism is treated as lacking rigour. Because mechanistic assumptions are embedded in dominant analytical techniques, critique is often neutralised not by refutation but by exclusion through admissibility criteria.

Table 3.7 summarises how canonical methods enforce particular assumptions.

Table 33: 3.7. Methodological stabilisation through tool–assumption coupling

Method	What is excluded by construction
Calculus-based optimisation	Discontinuities, thresholds, novelty, and non-smooth dynamics
General equilibrium	Constitutive openness, co-evolution, emergence, and boundary formation
Representative-agent modelling	Structural heterogeneity and meso-level organisation
Ergodic inference	Path-dependence, non-ergodicity, and endogenous novelty

The implication is that debates about realism are frequently redirected into debates about tractability. If an alternative ontology implies openness or non-ergodicity, it is difficult to render within the authorised toolkit and therefore risks being dismissed as non-scientific or merely descriptive. As Mirowski (1989) documents historically, the importation of formalisms derived from mechanics carried with it the metaphysical presuppositions of those tools. The metaphysics is therefore not only asserted; it is built into the technical apparatus.

For firm theory, the consequence is a publishability constraint. Bounded rationality is often rendered as constrained optimisation, uncertainty as risk, and networks as appended structures rather than constitutive relations. The method enforces the ontology.

C. Institutional stabilisation

Even where semantic and methodological stabilisation might be resisted, institutional mechanisms reproduce the paradigm through selective incentives. A paradigm persists when the disciplinary environment rewards work that is legible within its standards and filters work that violates its admissibility conditions. In this sense, paradigmatic stability becomes an evolutionary outcome within a professional ecology.

Table 3.8 summarises representative institutional mechanisms.

Table 34: 3.8. Institutional stabilisation through professional selection

Institution	Reproduction mechanism
Top journals	Reward formal legibility within canonical modelling standards
Hiring and promotion	Select for technical conformity and paradigm-recognised excellence
Grant evaluation	Privilege quantification and prediction, discouraging uncertainty-first inquiry
Graduate training	Socialise cohorts into canonical model families and evaluative criteria
Peer review	Applies paradigm-based standards of “quality,” filtering anomaly-friendly work

The effect is not necessarily deliberate exclusion. It is structural selection. Scholars and programmes that conform reproduce; those that do not are marginalised, or forced to translate their claims into mechanistic terms, often weakening the phenomena they seek to explain. At the firm level, theories that foreground routines, institutions, ecosystems, or power frequently

become legible only after being recoded as equilibrium-compatible or optimisation-compatible deviations.

III. Why Isolated Reform Is Absorbed

The combined operation of semantic, methodological, and institutional stabilisation explains why reforms frequently appear substantial while leaving the hard core intact. When bounded rationality is introduced, it is often recoded as friction or constrained optimisation. When dynamics are introduced, they frequently become transitions between equilibria rather than constitutive disequilibrium. When networks are introduced, they are often represented as additional channels linking atomistic agents rather than as constitutive relations that form agents and organisations.

The general principle is that reforms are reinterpreted so that loop coherence remains intact. Apparent change is preserved; the attractor is maintained.

IV. The Paradigm as a Closed Attractor

Within this framework, the mechanistic paradigm behaves as a closed attractor in theoretical space. Perturbations are accommodated within the protective belt and translated into frictions, shocks, or approximations. The system returns to its stable configuration not because all anomalies are resolved, but because the criteria for admissibility and success are generated internally by the same assumption network that anomalies would need to challenge.

This is the hallmark of a self-stabilising research programme. It can accumulate technical sophistication and local repairs while preserving the architecture that produces its recurrent blind spots.

V. Anomaly Absorption and Degenerative Accommodation

When anomalies appear – crises, bubbles, persistent unemployment, endogenous innovation – responses frequently take a characteristic form. The anomaly is reclassified as an exogenous disturbance, auxiliary frictions are added, and the core organising commitments are preserved. This pattern is consistent with a Lakatosian dynamic in which the hard core is protected while auxiliary hypotheses proliferate. The protective belt expands; the architecture remains.

For analytical clarity, Table 3.9 summarises an illustrative historical pattern of accommodation.

Table 35: 3.9. Illustrative reform waves as stabilising accommodation

Reform wave	Stabilising outcome
Microfoundations and rational expectations	Reinforced equilibrium and optimisation as organising commitments
Behavioural economics	Appended anomalies while retaining optimisation as normative benchmark
Heterogeneous-agent modelling	Introduced distributions while preserving equilibrium-centred structure

The diagnostic point is not that these developments lack value, but that they frequently function as accommodation within an existing attractor rather than as assumption-level transformation.

Conclusion

This subsection has shown that the mechanistic paradigm is self-stabilising not only because its assumptions form reinforcing loops, but because those loops are reproduced through semantic embedding, methodological gatekeeping, and institutional selection. Assumptions protect one another through redundant pathways; central nodes are stabilised simultaneously across multiple mechanisms; reforms are frequently absorbed through redescription and auxiliary modification. The paradigm thus appears adaptive while remaining structurally unchanged.

The implication is diagnostic and architectural: paradigmatic persistence is better explained as evolutionary stability within a disciplinary ecology than as epistemic adequacy under conditions of complexity.

3.2.3 Kuhn on Normal Science and Anomaly Accumulation

Introduction

Sections 3.2.1 and 3.2.2 established that the mechanistic paradigm is stabilised through reinforcing feedback loops and through semantic, methodological, and institutional reproduction mechanisms. The remaining question is why these mechanisms remain effective even under persistent empirical, conceptual, and historical failure. Thomas Kuhn's theory of normal science provides the missing explanatory logic. Paradigms persist, Kuhn argued, not because they are empirically flawless, but because they define what counts as a legitimate scientific problem, what methods may be used to address it, and what constitutes an acceptable solution.

This subsection applies Kuhn's account to the inherited body of economic theory. It argues that the stabilisation mechanisms identified earlier correspond closely to the defining features of Kuhnian normal science. Semantic stabilisation constrains the language of admissible puzzles; methodological stabilisation constrains the form solutions may take; and institutional stabilisation governs who is rewarded for solving them. As a result, anomalies can accumulate without triggering paradigm replacement. Failure becomes a source of further puzzle-solving rather than a signal for assumption-level revision.

I. Normal Science as Puzzle-Solving Within a Paradigm

Kuhn characterises normal science as disciplined puzzle-solving conducted under settled commitments. A paradigm supplies exemplars, categories, and standards of success, and research proceeds by extending these resources rather than subjecting them to open test. Anomalies are therefore not treated as immediate falsifications. They are initially understood as unresolved puzzles, measurement problems, boundary conditions, or special cases. Kuhn's crucial insight is that paradigms are abandoned only when a coherent alternative is available; rejecting one paradigm is inseparable from accepting another (Kuhn, 1962, pp. 52–65, 77).

The inherited body of economic theory exhibits the canonical characteristics of normal science. Its puzzles are framed within equilibrium as a baseline, optimisation as rationality, and tractable formalism as rigour. Legitimate questions ask which equilibrium obtains, how agents optimise, or how parameters may be estimated. Questions that challenge the paradigm's ontological or epistemological foundations – whether equilibrium is a meaningful organising principle, whether optimisation is coherent under uncertainty, or whether assumptions themselves

perform constitutive work – are typically relegated to philosophy or treated as methodologically non-operational.

Normal science in economics therefore consists not in open-ended inquiry, but in the refinement of a stylised world whose basic structure is presupposed. Success is judged by internal solvability, elegance, and conformity to disciplinary standards, standards that themselves embed the mechanistic assumption network identified in Section 3.1 and stabilised in Section 3.2.

At the firm level, the same logic applies. When firms are encoded as production functions or contractual nodes, admissible puzzles concern efficiency, incentives, and equilibrium selection. Questions about firms as emergent coordinating systems – routines, collective cognition, network embeddedness, and path-dependent capability formation – do not map cleanly onto the paradigm’s exemplars. They are filtered out by the puzzle-definition function of normal science rather than explicitly rejected.

II. Anomaly Accumulation Under Normal Science

In Kuhn’s framework, anomalies are not ignored; they are actively managed. Persistent failure does not automatically destabilise a paradigm because normal science is organised to contain anomalies. They are bracketed, re-described, postponed, or reassigned as technical puzzles until their accumulation becomes intolerable (Kuhn, 1962, pp. 66–76). The relevant question is therefore not whether anomalies exist in economics, but why their persistence has not forced assumption-level revision.

Within the present framework, anomalies accumulate in four analytically distinct families. Each conflicts directly with the mechanistic assumption network mapped in Section 3.1, particularly equilibrium, atomism, ergodicity, linearity, optimisation, and decomposability.

Predictive and Policy Anomalies

Table 36: 3.10. Predictive and policy anomalies

Failure	What occurred	Why the paradigm failed
Global Financial Crisis (2008)	Systemic risk and contagion were largely invisible	Atomism, decomposability, and equilibrium suppress network cascades (Kirman, 2010; Stiglitz, 2018)
COVID-19 macro disruptions	Regime shifts and non-linear propagation dominated	Linearity, ergodicity, and equilibrium translate shifts into exogenous shocks
Persistent inequality dynamics	Structural divergence became central	Representative-agent aggregation erases distributional mechanisms

Arthur describes orthodox economics as operating in a “timeless... equilibrium world” structured around optimisation and known distributions, systematically missing formation and adaptation (Arthur, 2014, pp. 5, 91). Kirman framed 2008 as a crisis of theory as much as of practice (Kirman, 2010), while Stiglitz argued that post-crisis DSGE retained assumptions that were “widely discredited,” rendering reforms largely cosmetic (Stiglitz, 2018, p. 4).

Firm-level dynamics reveal the same blindness. Crises propagate through interbank networks, supply chains, and institutional coupling. Atomism and decomposability remove these propagation structures by assumption, forcing the paradigm to reinterpret cascades as exogenous disturbances rather than endogenous outcomes of networked interdependence (Arthur, 2013, p. 14).

Conceptual Anomalies

Table 37: 3.11. Conceptual anomalies

Anomaly	Key scholars	Why it conflicts
Fundamental uncertainty	Knight (1921); Keynes (1936); Shackle (1972)	Contradicts determinism and probabilistic closure
Path dependence	David (1985); Arthur (1989)	Contradicts ergodicity and equilibrium baselines
Emergence	Anderson (1972); Holland (1995)	Contradicts reductionism and decomposability
Far-from-equilibrium dynamics	Prigogine (1996); Bak (1996)	Contradicts equilibrium tendency and linear stability

These are not missing variables. They are ontological and epistemological conflicts. If systems are non-ergodic, history is constitutive; if emergence is real, wholes are not sums; if uncertainty is fundamental, optimisation is ill-posed rather than merely difficult.

Institutional–Historical Anomalies

Table 38: 3.12. Institutional and historical anomalies

Anomaly	Key scholars	Why it conflicts
Endogenous institutions	North (1990, 2005); Acemoglu & Robinson (2012)	Violates universal-law monism
Historical contingency	Macro-historical record	Violates equilibrium and stable parameters
Systematic bounded rationality	Simon (1955); Kahneman & Tversky (1979)	Violates perfect information and optimisation

North's critique aligns closely with the present diagnosis: the inherited framework proceeds as though uncertainty were exceptional rather than constitutive (North, 2005, p. 6).

Technological–Organisational Anomalies

Table 39: 3.13. Technological and organisational anomalies

Anomaly	Key scholars	Why it conflicts
Increasing returns	Arthur (1994); Romer (1990)	Violates convexity and equilibrium uniqueness
Network effects	Shapiro & Varian (1999)	Violates atomism and independence

Anomaly	Key scholars	Why it conflicts
Firm routines and evolution	Nelson & Winter (1982)	Violates optimisation and representative-agent analogues

Organisational routines and capabilities are not frictions around an optimum; they are constitutive mechanisms of adaptation in non-ergodic environments. This is precisely the meso-level at which mechanistic economics becomes blind (Arthur, 2013, p. 14).

III. Why Anomalies Do Not Trigger Paradigm Revision

As Sections 3.2.1 and 3.2.2 demonstrated, anomalies are systematically absorbed rather than resolved. Semantic reclassification turns failures into “shocks” or “frictions.” Methodological filtering restricts admissible responses to those preserving equilibrium and optimisation. Institutional selection rewards conformity and marginalises assumption-level dissent. Kuhn’s core insight follows directly: anomalies alone do not overturn paradigms. They persist until a coherent alternative becomes available that can redefine legitimate puzzles and standards of explanation (Kuhn, 1962, p. 77).

IV. Kuhn’s Replacement Criterion

Kuhn’s decisive claim is not merely descriptive but structural. Paradigm change is replacement, not repair. A paradigm is abandoned only when another paradigm exists that can reorganise ontology, epistemology, and problem-definition simultaneously. Incremental critique, even when empirically devastating, is insufficient.

Within this framework, the implication is immediate. Complexity and CAS approaches cannot function as add-ons to the equilibrium–optimisation architecture. Presented as supplements, they will be absorbed as peripheral refinements, exactly as documented in Sections 3.2.1–3.2.2. Only a coherent alternative capable of redefining admissible questions, methods, and explanatory standards can generate the gestalt shift Kuhn associated with scientific revolutions (Kuhn, 1962, pp. 52–65).

Conclusion: The Kuhnian Diagnosis

Applying Kuhn, we can now specify why the mechanistic paradigm persists despite accumulating anomalies. Economics exhibits the defining features of normal science: exemplar-driven training, puzzle-solving under protected commitments, and institutional reward structures that enforce conformity (Kuhn, 1962, pp. 35–42; 52–65). Anomalies accumulate across predictive, conceptual, institutional, and organisational domains, yet are absorbed through stabilisation mechanisms rather than propagated to the hard core.

The conclusion is diagnostic rather than rhetorical. Paradigmatic persistence reflects evolutionary stability within a disciplinary ecology, not adequacy under conditions of complexity. Kuhn’s replacement criterion therefore clarifies the task ahead: only a coherent alternative paradigm can displace a self-stabilising one.

3.2.4 Lakatos on Degenerating Research Programmes

Introduction

Sections 3.2.1–3.2.3 established the mechanisms through which the mechanistic paradigm persists: reinforcing feedback loops that stabilise equilibrium and optimisation as high-centrality commitments; semantic, methodological, and institutional reproduction mechanisms that absorb perturbations; and Kuhnian normal science dynamics that convert failure into further puzzle-solving (Kuhn, 1962). Kuhn’s account, however, is primarily descriptive. It explains how paradigms endure, but it does not provide evaluative criteria for judging whether a persistent programme is scientifically progressive or degenerative. For that judgement, Imre Lakatos’s methodology of scientific research programmes supplies the relevant standards (Lakatos, 1970, pp. 91–196).

This subsection applies Lakatos’s framework to evaluate the inherited body of economic theory as a research programme. The core diagnostic claim is that the programme exhibits the characteristic signatures of degeneration: a protected hard core insulated from disconfirmation, a proliferating protective belt that accommodates anomalies retrospectively, heuristic strategies that shift from discovery to repair, and a persistent failure to generate excess corroboration. The point is not that the programme produces no insights, but that its principal growth occurs through defensive modification rather than risk-bearing theoretical advance.

I. Lakatos’s Evaluative Framework

Lakatos conceptualises scientific development as competition between research programmes rather than isolated hypotheses. Each programme has a protected hard core of foundational commitments, surrounded by an adjustable protective belt of auxiliary assumptions and specific models. Research is guided by a positive heuristic, which directs problem selection and model construction, and constrained by a negative heuristic, which prohibits direct revision of the hard core (Lakatos, 1970, pp. 91–196).

For analytical clarity, Table 3.14 summarises these components.

Table 40: 3.14. Structure of a Lakatosian research programme (Lakatos, 1970)

Component	Definition
Hard core	Foundational commitments protected from direct falsification by methodological decision
Protective belt	Auxiliary hypotheses, model variants, and empirical specifications that can be revised
Positive heuristic	Problem-solving strategies that guide model construction and extension
Negative heuristic	Rules that prohibit questioning the hard core and redirect anomalies outward

On this basis, Lakatos distinguishes progressive from degenerative programmes. A progressive programme exhibits theoretical growth that leads empirical growth by generating novel, risky predictions that are subsequently corroborated, thereby producing “excess corroboration.” A degenerative programme, by contrast, adjusts post hoc to accommodate known facts, adds

epicyclic modifications, and fails to yield novel predictions in advance of observation (Lakatos, 1970, pp. 116–122).

Table 3.15 compresses this evaluative distinction.

Table 41: Table 3.15. Progressive and degenerative programmes (Lakatos, 1970)

Programme type	Characteristic pattern
Progressive	Theoretical development generates novel predictions and achieves excess corroboration
Degenerative	Modifications are chiefly post hoc accommodations without novel predictive gain

II. The Protected Hard Core in the Inherited Economic Programme

The hard core of the inherited programme corresponds to the ontological and epistemological commitments reconstructed in Section 3.1. These commitments are rarely treated as revisable hypotheses. They function as preconditions of admissible modelling rather than propositions exposed to direct test. The feedback architecture documented in Section 3.2.1 explains why these commitments are resilient: assumptions mutually defend one another, and the most fragile commitments – equilibrium and optimisation – also occupy the most central stabilising positions.

Table 3.16 summarises the hard-core commitment set in compressed form.

Table 42: 3.16. Hard-core commitment clusters in the inherited programme

Commitment domain	Representative hard-core elements
Ontological	Reductionism, atomism, closure, substance ontology, determinism
Epistemological	Certainty ideal, universal laws, perfect information, model monism
Derivative commitments	Equilibrium, linearity, homogeneity, ergodicity, optimisation

At the level of firm theory, the same protection pattern appears in the endurance of production-function and maximisation representations as default primitives. Even where alternative vocabularies exist, they are frequently translated back into optimisation-compatible or equilibrium-compatible forms, indicating that the negative heuristic remains operative.

III. Indicators of Degeneration

Within a Lakatosian framework, degeneration is diagnosed not by the existence of anomalies, but by the programme’s characteristic response to anomalies. The inherited programme displays a recurrent pattern in which anomalies are diverted to the protective belt, model families proliferate to accommodate past failures, and predictive success remains largely confined to domains of calibration. Five indicators are particularly salient.

1. Hard-core protection under failure

The first indicator is systematic hard-core immunisation. Major empirical failures do not trigger revision of equilibrium or optimisation commitments. Instead, auxiliary assumptions are

introduced to preserve the hard core while explaining failure as a missing friction, an omitted sector, or an unmodelled constraint. This is the classic Lakatosian pattern of protective-belt adjustment under negative-heuristic constraint (Lakatos, 1970).

Table 3.17 illustrates this logic using canonical episodes.

Table 43: 3.17. Immunisation pattern: hard core preserved, protective belt adjusted

Failure episode	Hard core preserved	Protective-belt adjustment
Global Financial Crisis (2008)	Equilibrium, optimisation	Financial frictions in DSGE variants (Kirman, 2010; Stiglitz, 2018)
Persistent unemployment	Market-clearing tendency	Search and matching frictions
Distributional divergence	Efficiency baselines	Auxiliary inequality mechanisms (e.g., skill-biased technical change)

The point is structural rather than rhetorical. Failure is not allowed to propagate into hard-core revision; it is redirected into auxiliary repair. Kirman (2010) framed the 2008 crisis as a crisis of theory as well as of the real economy, while Stiglitz (2018) argued that post-crisis macro retained assumptions that were “widely discredited,” indicating repair without foundational reconsideration. Colander et al. (2009) similarly note the coexistence of high technical sophistication with weak empirical performance.

2. Post hoc accommodation without risky novelty

A second indicator is the predominance of retrospective accommodation. Crises are explained after the fact through auxiliary modifications that were not risk-bearing predictions beforehand. This does not mean the models cannot be adjusted to fit the past. It means theoretical growth does not generate novel corroborated consequences in advance, which is the Lakatosian criterion of progress (Lakatos, 1970, pp. 116–122).

This pattern is summarised in Table 3.18.

Table 44: 3.18. Post hoc accommodation sequence

Episode	Representative adjustment pattern	Novel prediction achieved ex ante
Great Depression	Auxiliary rigidities and stickiness mechanisms	No
Stagflation	Supply-shock accommodation and expectations refinements	No
Global Financial Crisis (2008)	Financial frictions appended to core macro families	No

The diagnostic is not that economics never predicts anything. It is that the core programme repeatedly explains yesterday’s failure without generating a robust pattern of advance prediction in domains where instability, regime change, and network propagation are central.

3. Ad hoc proliferation and epicyclic growth

A third indicator is protective-belt proliferation. When anomalies persist, model families expand through layers of auxiliary additions that restore fit while retaining the hard core. The result is increasing complexity without commensurate predictive gain.

Table 3.19 expresses this phenomenon in programme form.

Table 45: 3.19. Protective-belt proliferation in core model families

Core family	Typical trajectory of auxiliary expansion
DSGE macro families	Baseline model → nominal rigidities → adjustment costs → financial frictions → heterogeneity blocks → multi-sector and open-economy elaborations

Colander et al. (2009, p. 260) characterise this pattern as combining weak foundations with unsatisfactory empirical performance, while Stiglitz (2018) argues that many underlying assumptions remained discredited even as elaborations multiplied. Landini et al. (2025) connect such persistence to the enduring constraints of optimisation formalism and its difficulty in representing nonlinearity, path dependence, and emergent organisation. The historical lineage of such formalisms has been documented in detail by Mirowski (1989).

At the firm level, an analogous pattern appears when governance, resources, and adaptation vocabularies are layered on maximisation primitives without shifting the constitutive ontology of routines, relational embeddedness, and emergent coordination. The evolutionary alternative proposed by Nelson and Winter (1982) explicitly rejects optimisation as a universal explanatory primitive and instead treats routines and search as foundational, a move that is difficult to assimilate without hard-core revision.

4. Heuristic exhaustion

A fourth indicator is a shift in the programme’s positive heuristic from discovery to repair. In progressive phases, the positive heuristic generates new problems and new explanatory structures. In degenerative phases, it increasingly functions as an instruction to add auxiliary features until calibration improves, without producing new risk-bearing consequences.

This exhaustion is visible when modelling success is increasingly measured by parameter calibration and local fit rather than by advance prediction. It is also visible in the systematic tendency to linearise around small disturbances, limiting the scope of admissible dynamics. Stiglitz (2018) argues that the dominant approaches are ill-suited to systemic disruption precisely because their tractability conditions exclude nonlinear crisis dynamics. Hayek’s earlier warning about the “pretence of knowledge” and the limits of quantitative prediction in complex phenomena is directly relevant here (Hayek, 1989, pp. 4–7).

5. Lack of excess corroboration

The fifth indicator is the absence of excess corroboration. The programme’s most celebrated successes tend to occur within stable domains where calibration is possible and distributions appear stationary. Performance weakens systematically in the domains that are most diagnostic under complexity: crises, endogenous innovation, institutional transformation, network contagion, and non-ergodicity. Arthur’s characterisation of equilibrium-centred economics as systematically missing formation and adaptation captures this failure mode (Arthur, 2014).

Table 3.20 summarises the calibration–failure structure.

Table 46: 3.20. Calibration success and failure outside the calibration domain

Model domain	Typical calibration environment	Failure domain
Asset pricing and risk families	Normal volatility periods	Turbulence, contagion, regime change
Macro equilibrium families	Stable growth periods	Crises, structural breaks
Growth and innovation families	Post hoc growth regularities	Lock-in, platform dynamics, endogenous restructuring

This does not establish that mechanistic models are useless. It establishes that the programme’s growth does not exhibit the Lakatosian signature of progress: theoretical advance leading empirical discovery through corroborated novel consequences.

IV. Why Degenerating Programmes Persist

Lakatos explicitly noted that degenerating programmes can persist for extended periods when institutional incentives reward conformity, when sunk costs are high, and when alternatives are fragmented. The self-stabilising mechanisms analysed in Section 3.2.2 provide a concrete account of how such persistence can occur through disciplinary selection and methodological gatekeeping. Colander et al. (2009) describe a discipline capable of sustaining technically sophisticated work even under weak empirical performance, a pattern consistent with institutional progressiveness combined with scientific degeneracy.

At the firm level, analogous persistence is visible in the continued use of production-function primitives despite decades of evidence that organisational performance depends critically on routines, relational architectures, and network propagation mechanisms (Nelson & Winter, 1982). Episodes such as financial contagion, in which institutions appeared locally solvent until network topology transmitted collapse, illustrate the explanatory limits of atomistic equilibrium frameworks.

V. Progressive Alternatives and Lakatosian Rationality

For Lakatos, the rational grounds for programme shift are comparative. A shift is justified when an alternative programme demonstrates progressive problem shifts and achieves empirical success not by post hoc accommodation alone but by generating corroborated novel implications (Lakatos, 1970). In the present framework, the relevant question is therefore not whether complexity-oriented approaches can be made rhetorically attractive, but whether they supply a coherent programme with heuristic fertility and empirical traction in domains where the inherited programme repeatedly fails.

Table 3.21 states the relevant Lakatosian criteria in evaluative form.

Table 47: 3.21. Lakatosian criteria for a progressive alternative (Lakatos, 1970)

Criterion	Diagnostic requirement
Novel empirical content	Generates implications not built into calibration
Heuristic fertility	Opens new problem fields rather than patching old ones
Excess corroboration	Succeeds in domains where rival programme fails
Coherent architecture	Maintains internal consistency without immunising stratagems

In complexity research and related empirical literatures, several candidate patterns are frequently highlighted as relevant to economic and firm-level dynamics under interdependence: path-dependent lock-in (Arthur, 1989), power-law firm size distributions (Axtell, 2001), and critical transitions in complex systems (Scheffer et al., 2009). These are not “point forecasts” of particular events; they are empirically testable pattern claims about the kinds of regularities expected in non-equilibrium, interaction-driven environments. Hayek’s insistence that complex phenomena admit, at best, pattern prediction rather than point prediction makes this distinction methodologically salient (Hayek, 1989, pp. 4–7).

At the firm level, the explanatory focus shifts correspondingly toward emergent capabilities, platform and network effects, and historically contingent trajectories. Organisational evolution is not treated as friction around an optimum but as adaptation in shifting landscapes, consistent with the evolutionary logic emphasised by Nelson and Winter (1982).

Conclusion: The Lakatosian Verdict

On Lakatosian criteria, the inherited body of economic theory exhibits the characteristic signatures of a degenerating research programme. The hard core is systematically protected; anomalies are redirected into the protective belt; auxiliary assumptions proliferate; heuristic strategies increasingly function as repair rules; and predictive success is largely confined to domains of calibration, with limited excess corroboration in the domains most diagnostic under complexity (Lakatos, 1970; Kirman, 2010; Colander et al., 2009; Stiglitz, 2018). These patterns are consistent with a programme that remains institutionally stable while its scientific development is increasingly defensive.

The methodological implication is evaluative rather than rhetorical. Lakatos’s framework justifies rational preference for programmes that generate progressive problem shifts and corroborated novel content over those whose growth is predominantly post hoc accommodation. This does not eliminate the possibility of useful results within the inherited programme. It does, however, provide a principled basis for judging its overall trajectory when confronted with persistent anomalies, expanding auxiliary repairs, and limited advance predictive success.

Conclusion of Section 3.2

Section 3.2 has resolved the persistence puzzle raised by Section 3.1. The mechanistic paradigm endures not because its assumptions are adequate to economic reality, but because they are structurally locked in place through redundant, self-reinforcing mechanisms that absorb challenge and reproduce coherence across generations.

Three findings carry forward.

First, the assumptions form a triple-locked system. Three feedback loops operate across the architecture: an ontological loop linking reductionism, equilibrium, and closure; an epistemic loop linking determinism, ergodicity, and optimisation; and a structural loop linking rationality, representative-agent aggregation, and equilibrium. These loops interlock at high-centrality nodes – equilibrium and optimisation occupy central positions across multiple loops simultaneously – so that challenge to either activates resistance throughout the architecture. The paradigm thereby functions as a closed attractor: perturbations are absorbed, anomalies renamed as frictions or shocks, and the original structure reinstated in refined technical form. Change appears continuous at the surface while the deep architecture remains intact.

Second, the lock-in is reproduced through disciplinary practice. Semantic stabilisation embeds assumptions in vocabulary, making alternative ontological commitments linguistically awkward. Methodological stabilisation makes tractability the criterion of rigour, so that what cannot be expressed in the authorised formalism is dismissed as unscientific. Institutional stabilisation rewards conformity through training, publication, and funding, selecting for reproduction rather than revision. Together, these mechanisms enable the paradigm to appear responsive – incorporating "fixes" and refinements – while remaining structurally unchanged.

Third, on formal evaluative criteria, the programme is degenerating. Anomalies accumulate across predictive, conceptual, institutional, and organisational domains, yet are absorbed through Kuhnian normal science dynamics rather than propagated to the hard core. Evaluated against Lakatosian criteria, the programme exhibits hard-core protection, post hoc accommodation, auxiliary proliferation, heuristic exhaustion, and absence of excess corroboration in the domains most diagnostic of complexity.

The implication is decisive: piecemeal reform cannot succeed because the architecture regenerates what is relaxed. Only a coherent alternative paradigm can displace a self-stabilising one.

Section 3.3 now traces how this lock-in expresses itself in the methodological practice of the discipline.

3.3 Methodological Symptoms

Section 3.1 reconstructed what the mechanistic paradigm assumes: a fifteen-assumption architecture organised across ontological, epistemological, and derivative layers, whose most consequential property is the systematic elimination of the meso-level. Section 3.2 explained why those assumptions persist despite their incompatibility with complex adaptive systems: three reinforcing feedback loops interlock at high-centrality nodes, reproduced through semantic, methodological, and institutional stabilisation mechanisms that render the paradigm self-sealing. The diagnosis now requires a third stage: how does this lock-in express itself in the actual methodological practice of the discipline?

Section 3.3 answers this question by identifying five methodological symptoms. These are not incidental habits or correctable stylistic preferences. They are the necessary expressions of the ontological and epistemological commitments diagnosed in Section 3.1, stabilised by the mechanisms identified in Section 3.2. Understanding them as symptoms – rather than as independent problems – is essential: treating symptoms while leaving the generative foundations intact cannot produce transformation, only surface variation.

The argument proceeds through five subsections. Subsection 3.3.1 diagnoses physics envy – the appropriation of classical mechanics' mathematical form without validating whether economic phenomena satisfy the domain conditions that make that form appropriate. Subsection 3.3.2 isolates the optimisation fetish as the central methodological symptom: calculus-based maximisation elevated from context-bound tool to universal decision principle. Subsection 3.3.3 examines the institutional culmination of both symptoms in DSGE, treating it as the paradigmatic form in which mechanistic metaphysics becomes operational in policy, publication, and training. Subsection 3.3.4 diagnoses the aestheticisation of economics – the elevation of mathematical elegance over empirical adequacy – explaining why crises such as 2008 did not trigger paradigm change. Subsection 3.3.5 identifies the epistemic foundation beneath all four symptoms: Hayek's pretence of knowledge.

This diagnosis is necessary for what follows. Without understanding how the lock-in manifests methodologically, the consequences documented in Section 3.4 would appear as contingent failures rather than as predictable outcomes generated by a methodology structurally misaligned with complex adaptive systems.

3.3.1 Physics Envy and Mathematical Formalism as Symptoms

Introduction

Sections 3.1–3.2 established the diagnostic foundation. The inherited body of economic theory rests on mechanistic ontological commitments – reductionism, determinism, substance ontology, and closure – and on derivative commitments such as equilibrium, linearity, and homogeneity. These commitments are not merely asserted; they are stabilised by reinforcing feedback loops and reproduced through semantic, methodological, and institutional mechanisms, while the programme's pattern of repair is consistent with Lakatosian degeneration (Lakatos, 1970). The methodological question is therefore not whether economics uses mathematics, but why a distinctive kind of mathematisation persists as a disciplinary norm even where the domain conditions that make it valid are absent. The phenomenon commonly described as “physics envy” is treated here as a symptom: the methodological expression of mechanistic metaphysics rather than a neutral marker of rigour.

The diagnostic claim is simple. Physics-style mathematical formalism is not an externally detachable technique that can be retained while changing foundations. It is the operational face of a particular ontology and epistemology. When those foundations misdescribe economic reality as if it were decomposable, deterministic, and effectively closed, the corresponding toolkit privileges equations, equilibrium, optimisation, and calculus as the dominant forms of “rigorous” explanation. The result is a methodological regime that is structurally ill-suited to complex adaptive systems, because the very conditions that make the formalism work are the conditions that complex adaptive systems violate.

I. Physics Envy as Methodological Appropriation

The methodological character of the inherited programme emerged through a historically deliberate appropriation of physics’ prestige and techniques. The marginalist generation sought scientific legitimacy by importing the mathematical style of classical mechanics, and by the mid-twentieth century that style had become the benchmark of “rigour,” culminating in formal general equilibrium constructions such as Arrow–Debreu (Arrow & Debreu, 1954). The issue is not that these contributions were mathematically sophisticated, but that their success helped normalise the inference that scientific status is achieved through physics-like form, even when the economic domain does not satisfy the presuppositions that make the form appropriate.

This appropriation has four recurrent methodological signatures. Table 3.22 summarises them.

Table 48: 3.22. Four physics-style appropriations and their diagnostic problem

Appropriation	Diagnostic problem
Equation-based thinking (economic reality rendered as systems of equations)	Social and institutional phenomena do not naturally reduce to equation form without importing ontological presuppositions that may be false; this creates “rigour” by exclusion rather than by fit (Lawson, 2019, pp. 9–15).
Optimisation as universal principle (agents and firms defined by maximisation)	Optimisation presupposes a stable objective function and a determinate mapping from choices to outcomes; under fundamental uncertainty, the optimisation metaphor becomes conceptually fragile rather than foundational (Mirowski, 1989, p. 235; Knight, 1921).
Equilibrium as organising baseline (markets framed as tending toward stable rest states)	Equilibrium-centred modelling presumes closure and stabilising dynamics that are not characteristic of open, adaptive economies; it makes disequilibrium structurally secondary (Arthur, 2014, p. 5).
Calculus as the core reasoning tool (marginal analysis as default)	Calculus presupposes continuity and differentiability; nonlinearity and discontinuities can be treated only by approximation, which often becomes the substantive model rather than a controlled simplification (Fontana, 2008, p. 14).

At the firm level the same four signatures recur. A production function representation renders the firm as an equation; profit maximisation defines the firm as an optimisation device; equilibrium baseline reasoning treats strategy as convergence toward a stable “optimal” configuration; and marginal analysis frames change as smooth adjustment rather than discontinuous innovation. This combination systematically sidelines the meso-level phenomena that most closely track observed firm behaviour: routines, organisational learning, path-dependent capability formation, and relational embeddedness (Nelson & Winter, 1982).

II. Why Physics Envy Arises: The Ontological Root

The persistence of physics-style formalism is not well explained as mere taste or fashion. It follows logically from the ontological and epistemological commitments reconstructed in Section 3.1. Classical mechanics is powerful because its canonical domain is, to a useful approximation, decomposable and governed by stable regularities. When economics assumes analogous properties – implicitly or explicitly – physics-style methods become not just attractive but seemingly necessary.

Table 3.23 states the causal chain from assumption to methodological expression.

Table 49: 3.23. From mechanistic assumptions to formalist method

Assumption (Section 3.1)	Methodological expression (Section 3.3)
Reductionism	Decomposition into equations for parts; aggregation treated as technically harmless
Determinism	Solvability and uniqueness ideals; emphasis on closed-form solutions
Closure	Equilibrium conditions and boundary-fixing; openness treated as exogenous “shock”
Homogeneity	Representative agent constructions; single optimisation problem substitutes for heterogeneity
Linearity	Marginal analysis and local approximations; linearisation becomes the dominant mode

Hayek diagnosed the root problem as an “erroneous transfer” of habits of thought from the natural sciences to social phenomena (Hayek, 1945, p. 520). The problem is therefore not mathematics as such, but the unexamined ontological freight carried by particular mathematical forms.

III. Why Physics Envy Persists

Physics-style formalism persists because it delivers institutional and professional rewards even when its explanatory costs are high. The lock-in architecture established in Section 3.2 explains why this persistence is structurally stable: the dominant toolkit is embedded in training, publication standards, and evaluation criteria, and it draws reinforcement from the same high-centrality commitments – especially equilibrium and optimisation – that stabilise the broader programme.

Table 3.24 summarises the benefit–cost structure that sustains persistence.

Table 50: 3.24. *Perceived benefits and systemic costs of physics-style formalism*

Perceived benefit	Systemic cost
Scientific legitimacy via mimicry of “hard science” form	Ontological mismatch: emergence, institutions, and adaptive agency become difficult to represent (Lawson, 2019).
Apparent precision through proofs and closed-form results	False precision: symbolic exactness substitutes for empirical adequacy; the discipline risks what Hayek later called the “pretence of knowledge” (Hayek, 1974/1989).
Standardised professional identity and common technique	Methodological monism: one method becomes the criterion of science; pluralism is treated as weakness rather than appropriateness.
Tractability and teachability	Research distortion: tractable questions dominate, while important questions are pushed aside as “non-operational” (Akerlof, 2020).

At the firm level the costs are amplified because firm behaviour is intrinsically meso-level and historically contingent. Organisational learning, ecosystem dynamics, and capability accumulation are not conveniently expressible as equilibrium solutions to stable optimisation problems. When methodological legitimacy is keyed to formal tractability, the firm is repeatedly forced back into production-function and maximisation primitives, even when these primitives omit the most explanatory mechanisms (Nelson & Winter, 1982).

IV. Mathematical Formalism as Symptom, Not Cause

Within this diagnostic frame, mathematical formalism is not the disease but the symptom. Formalism would be methodologically appropriate if its presuppositions matched the nature of its objects. The problem is that the inherited programme often treats the suitability of the method as self-validating: the existence of a rigorous proof is taken as evidence of scientific adequacy, even when the content relies on assumptions systematically violated in observed economies.

General equilibrium provides the cleanest illustration. Its mathematical form is internally elegant, but its substantive conditions typically require strong closure, convexity, completeness, and an institutional thinness that excludes precisely those mechanisms – power, institutional evolution, increasing returns, network effects, and endogenous novelty – that characterise real economic dynamics. Rigour in form can therefore coexist with vacuity in content. Keynes’s warning remains directly relevant: mathematical economics can become a “concoction” whose precision is no better than the imprecision of the assumptions from which it is derived (Keynes, 1936, p. 272). Mirowski’s historical diagnosis likewise emphasised that economics borrowed the mathematical form of physics without validating the domain conditions that make that form appropriate (Mirowski, 1989).

V. The Borrowed-Prestige Problem

A central mechanism sustaining physics envy is the borrowed-prestige effect: physics' scientific authority is imported through its mathematical style, while the discipline neglects the difference between the domains. Physics methods are powerful in contexts where system boundaries can be treated as stable, where objects do not strategically respond to theory, and where the causal structure is not reflexively altered by beliefs, institutions, or policy. Economic systems violate these conditions because agents adapt, rules evolve, and institutions reshape incentives and information.

Table 3.25 states the mismatch at the level of domain conditions.

Table 51: 3.25. Domain conditions: physics-style form versus economic reality

Domain condition in classical physics modelling	Economic reality that breaks it
Closed or effectively isolated systems	Open systems with constitutive interaction and institutional flux
Decomposability and weak interaction effects	Non-decomposable interdependence; network propagation and feedback
Stable regularities independent of observation	Reflexivity: behaviour and institutions change in response to expectations and policy
Non-adaptive entities	Adaptive agents and organisations

The methodological consequence is not simply “some assumptions are unrealistic.” It is that the discipline risks substituting the prestige of form for the discipline of fit. Hayek's critique of the pretence of knowledge is precisely about this substitution: what is measurable and solvable is treated as what is scientifically relevant, even when decisive causal mechanisms are not captured by the formal frame (Hayek, 1974/1989). Akerlof's discussion of “sins of omission” extends this diagnosis: formal rigidity can cause systematic neglect of important domains because they resist the authorised mode of quantification (Akerlof, 2020).

VI. Why the Formalist Regime Is Structurally Incompatible with CAS

Complex adaptive systems exhibit features that are not minor complications to be appended to a mechanistic baseline. They contradict the assumptions that make physics-style formalism work in its classical form. The incompatibility is therefore structural.

Table 3.26 states the CAS features that directly collide with the mechanistic assumptions reconstructed in Section 3.1.

Table 52: 3.26. CAS features and the mechanistic commitments they contradict

CAS feature	Mechanistic commitment contradicted
Nonlinearity and amplification	Linearity and local marginal reasoning
Feedback loops and circular causation	Unidirectional causation implied by simple optimisation narratives
Path dependence and historical irreversibility	Ergodicity and equilibrium as organising baselines
Endogenous novelty	Determinism and the certainty ideal
Structural heterogeneity	Homogeneity and representative-agent aggregation
Adaptive rule change	Closure and fixed boundary assumptions

The formal toolkit, in turn, presupposes exactly what these CAS features deny. Differential-equation methods privilege continuity; optimisation presupposes stable objectives and state spaces; representative-agent methods presuppose aggregation validity; and equilibrium-centred approaches presuppose closure and stabilising dynamics. When those presuppositions are false, the toolkit cannot merely be “extended”; it becomes a systematic filter that selects questions the method can answer rather than questions the world demands.

This diagnosis is consistent with Lawson’s critique of closed-systems modelling and with work arguing that statistical and equilibrium-centred methods are ill-suited to open-ended, relational economic reality (Lawson, 2019; Paniagua, 2023, p. 578). Fontana’s observation that linear models are favoured because they are solvable highlights the regime’s deeper tendency: solvability becomes an implicit criterion of admissibility, rather than an output of model–world fit (Fontana, 2008, p. 14).

At the firm level, the incompatibility is especially visible. Firms evolve through organisational design choices, routines, learning architectures, and ecosystem relations that are path-dependent and interaction-driven. These mechanisms do not behave as smooth variations around a stable optimum. They behave as evolving rule systems under uncertainty, consistent with evolutionary accounts that treat routines and search as primitives rather than optimisation (Nelson & Winter, 1982).

VII. Methodological Consequences for the Discipline

Physics-style formalism shapes not only modelling technique but also research selection, publication criteria, and training practices. Where formal tractability becomes an implicit definition of rigour, the discipline disproportionately rewards problems that can be posed in optimisation and equilibrium terms and downgrades problems that require institutional thickness, historical contingency, or relational topology. This is the methodological face of the stabilisation mechanisms analysed in Section 3.2.2: gatekeeping occurs through what counts as admissible method, and admissible method is inseparable from the assumptions that render it workable.

Table 3.27 summarises the inversion produced by tractability standards.

Table 53: 3.27. Tractability-driven inversion of research priorities

Tractability-aligned research objects	Mechanism-rich research objects marginalised
Equilibrium selection; marginal comparative statics; single-agent optimisation	Emergence; institutional evolution; network contagion; organisational routines; path-dependent innovation

This inversion has direct consequences for the study of firms because the most explanatory firm-level mechanisms occupy the meso-level that the mechanistic programme tends to thin or erase. The methodological regime thus reproduces the “missing middle” diagnosed earlier, not by explicit denial but by systematic filtering.

VIII. Counter-Arguments and Their Limits

The claim advanced here is not that mathematics is incompatible with economics. It is that a specific historically inherited form of physics-style formalism carries ontological presuppositions that often misfit the domain. Rigour is not identical with formalisation. Rigour requires ontological adequacy, epistemological humility about uncertainty, and methodological fit. When formal precision is achieved by imposing closure, homogeneity, linearity, or optimisation where these conditions do not hold, the result is a disciplined style of misrepresentation rather than disciplined explanation. That is why the critique targets methodological monism, not quantitative reasoning per se (Lawson, 2019; Mirowski, 1989; Hayek, 1974/1989).

Conclusion

Physics envy and the fetishisation of mathematical formalism are best understood as symptoms of mechanistic metaphysics. The inherited programme imported equations, equilibrium, optimisation, and calculus as markers of scientific status, while often neglecting whether economic systems satisfy the conditions that make these tools appropriate in their classical use. The borrowed prestige of form thereby substitutes for the discipline of fit. The consequence is a methodological regime that systematically filters out the very phenomena that define economies and firms as complex adaptive systems: nonlinear feedback, path dependence, endogenous novelty, heterogeneity, and evolving rules.

This diagnosis clarifies why methodological reform cannot be achieved by tinkering with technique while preserving foundations. If formalism is generated by ontology and epistemology, then changing what economics can represent requires changing what economics assumes exists and what it assumes can be known. That is the deeper logic linking Section 3.3 to the foundational reconstruction developed in Chapter 4.

Table 54: 3.28. The four appropriations and the source of failure

Appropriation and what it imports	Why it fails for economics
Equation-based thinking imports differential-equation form.	Institutions, expectations, and endogenous novelty do not reduce cleanly to equation form without imposing closure and decomposition that erase constitutive mechanisms (Lawson, 2019).
Optimisation as universal imports the least-action metaphor.	Fundamental uncertainty destabilises fixed objective functions and stable constraint sets, making optimisation conceptually fragile as a general behavioural primitive (Knight, 1921; Mirowski, 1989).
Equilibrium as natural imports the thermodynamic endpoint metaphor.	Economies are open, adaptive, and routinely far from equilibrium; equilibrium becomes a modelling grammar rather than an empirically warranted organising principle (Arthur, 2014).
Calculus as core tool imports continuous-change mathematics.	Discontinuities, thresholds, feedback, and phase-transition dynamics are not marginal perturbations; they are structurally central to CAS, making linearisation and smoothness-based reasoning systematically distorting (Fontana, 2008).

The causal mechanism linking method to foundations can now be stated succinctly: substance ontology, determinism, and closure jointly support equilibrium-centred reasoning; equilibrium-centred reasoning privileges formalism; and formalism secures the borrowed authority of “physics envy” (Hayek, 1945; Mirowski, 1989).

Table 3.29 states the cognitive shift implied by the diagnosis. It reframes “rigour” as fit rather than formal display, and it clarifies why physics-style formalism is treated here as a symptom of a failing paradigm rather than as the definition of scientific method.

Table 55: 3.29. The cognitive shift

FROM	TO
“Economics uses mathematics because it is rigorous.”	“Physics-style mathematical formalism is the methodological face of mechanistic assumptions that misdescribe economies and firms as if they were closed, decomposable, and certainty-governed systems.”
“More formalisation increases scientific status.”	“Rigour requires ontological adequacy, epistemological humility, and method–phenomenon fit; where these fail, formal precision becomes pseudo-rigour” (Keynes, 1936, p. 272; Hayek, 1974/1989).

This concluding compression also makes the borrowed-prestige problem explicit: economics risks practising *physics-without-physics*, importing mathematical form while neglecting the domain conditions that make that form appropriate (Mirowski, 1989; Lawson, 2019; Akerlof, 2020).

3.3.2 The Optimization Fetish and Calculus-Based Methods

Introduction

Section 3.3.1 diagnosed physics envy as the methodological expression of mechanistic metaphysics, identifying four appropriations – equations, optimisation, equilibrium, and calculus – that shape the inherited body of economic theory’s methodological self-image. This subsection isolates the central appropriation and the one that, more than any other, defines economics’ technical grammar: the universal treatment of economic problems as constrained optimisation. Optimisation appears, in this regime, not as one modelling strategy among others but as the default representation of rationality, explanation, and even scientific legitimacy.

Within our framework, this is best understood as an optimisation fetish: a disciplinary commitment that treats optimisation as the universal solution method for all economic problems, and that evaluates competing approaches by how closely they reproduce the optimisation form. The claim here is not that optimisation is never useful; it is that it becomes methodologically totalising. Once it takes on that role, it functions less as a tool than as an ontological and epistemological filter: phenomena that do not fit the optimisation template are either re-described until they do, or treated as secondary deviations from a theoretical baseline. This subsection therefore diagnoses optimisation as a methodological symptom of the assumption-network mapped in Sections 3.1–3.2, and argues that optimisation is structurally incompatible with complex adaptive systems.

The subsection performs three tasks. It first shows that optimisation is generated by cascade logic from the mechanistic hard core and its derivative assumptions. It then specifies what calculus-based optimisation requires, and therefore what it eliminates by construction. It finally shows why the mismatch is not technical but structural: in generative, non-ergodic, novelty-producing systems, global optimisation is not merely hard; it is frequently ill-posed. This diagnostic finding implies that optimisation must be replaced – at the level of agency – by procedural and ecological rationality as alternative behavioural primitives (Simon, 1955, 1979; Gigerenzer et al., 1999). The present thesis establishes the ontological foundations for such replacement; the full development of agency-compatible alternatives is identified as essential further work.

I. Optimisation as Universal Method

In the inherited body of economic theory, optimisation is treated as a general-purpose representation of choice. The standard template – maximise or minimise an objective function subject to constraints, solve via Lagrangian and first-order conditions, characterise an optimum – functions as a universal problem-form. This universality is not confined to consumer theory or producer theory; it extends to macroeconomics, finance, public economics, and welfare analysis, where the core explanatory move is repeatedly the same: translate behaviour into a calculus problem and interpret its solution as both explanation and rationality.

This methodological universality imposes a particular kind of “world”: a world in which the decision problem is fully specifiable, the feasible set is knowable, and choice is representable as a static mapping from known objectives and constraints to an optimal outcome. In that world, processes of learning, sequential search, social influence, institutional evolution, and innovation appear as complications that must be translated into the optimisation grammar (if they are admitted at all). In Arthur’s terms, the framework thus presupposes “maximizing

principles” and a pre-structured environment, rather than discovering these properties in economic reality (Arthur, 2014).

Table 56: 3.30. *Optimisation as disciplinary template and its eliminations*

Optimisation domain (canonical template)	What becomes methodologically invisible
Consumer theory represents choice as utility maximisation subject to a budget constraint.	Learning, preference construction, social influence, and identity-based or norm-based behaviour become difficult to represent as constitutive mechanisms (Simon, 1955; Kahneman & Tversky, 1979).
Producer theory represents the firm as profit maximisation subject to a technology constraint.	Routines, capabilities, organisational learning, and emergent coordination mechanisms are forced into residual categories, or treated as “frictions” around a production function (Nelson & Winter, 1982).
Macroeconomics represents the economy as a representative optimiser maximising lifetime utility over time.	Heterogeneity, institutions, non-ergodicity, and meso-structure are suppressed by aggregation and by the presupposition of stable probability structure (Kirman, 2010; Colander et al., 2009).
Finance represents decisions as return maximisation subject to measurable risk.	Fundamental uncertainty, unknown unknowns, and endogenous network contagion are re-described as probabilistic tails, preserving the risk calculus while discarding the structural conditions of crises (Knight, 1921; Kay & King, 2020).
Public economics represents policy as social welfare maximisation subject to feasibility constraints.	Political economy, institutional path dependence, and non-teleological evolution of norms and rules are treated as constraints external to the “core” optimisation problem (North, 1990, 2005).

At the firm level, the mismatch becomes especially transparent. The optimisation representation treats firms as if they were calculation devices solving a given problem over a known feasible set. Yet much of what distinguishes firms as organisations – routines, distributed cognition, internal learning, capability accumulation, and the co-evolution of strategy with environment – appears precisely where the optimisation template is least appropriate (Nelson & Winter, 1982). Toyota’s competitive advantage, for example, is not well described as the solution to a cost minimisation problem, because the Toyota Production System is better understood as an evolving ecology of routines, feedback mechanisms, and supplier relations. What optimisation represents as a static optimum, the firm experiences as continuous adaptation.

The optimisation template mis-specifies the firm because it conflates two different objects: a solvable allocation problem and an evolving coordination system. A firm such as Toyota is not well represented as an agent that solves a stable cost-minimisation programme subject to a given technology. Its competitive advantage is generated by routines that are not reducible to a parameter vector inside a production function but are themselves a system of rule-governed coordination that changes the feasible set through learning. In such a system, “technology” is not a constraint that the firm takes as given; it is an endogenous outcome of organisational design, feedback discipline, and supplier co-evolution. The implied optimisation object therefore fails on two fronts simultaneously: the objective function is not stable (because

aspiration levels and performance criteria evolve with organisational learning), and the constraint set is not stable (because routines alter process capability and thus redefine costs, quality, throughput, and flexibility over time). The analytical consequence is decisive: the firm's performance cannot be derived as the solution to a timeless calculus problem without deleting the causal mechanism that produces performance in the first place – namely, the evolutionary accumulation and refinement of routines under path-dependent learning (Nelson & Winter, 1982; Simon, 1955, 1979).

If you want a compact table at this point (reader-friendly and harmless), insert it immediately after the paragraph above:

Table 57: 3.30A. Why optimisation mis-specifies the firm as organisation

Optimisation representation	Organisational reality in routine-based firms
Technology is a given constraint.	Technology is an evolving organisational capability produced by routines and learning.
Objectives are stable and known.	Objectives/aspirations adapt as performance feedback and competitive context change.
Choice is instantaneous and global.	Choice is sequential, local, and revised through experimentation and error-correction.
Explanation is solution of a maximisation problem.	Explanation is the evolution of coordination rules that reshape the feasible set.

II. How Optimisation Flows from Assumptions

The crucial point is not that economists prefer optimisation; it is that the mechanistic assumption-network requires it. Given the ontological and epistemological commitments diagnosed in Sections 3.1.1–3.1.2, optimisation becomes the natural methodological consequence. If the world is treated as determinate, sufficiently closed, and knowable in principle, and if uncertainty is treated as an information problem resolvable by better measurement, then rationality can be represented as selecting the best element from a known set. Optimisation is therefore not discovered in behaviour; it is the methodologically “correct” form within the mechanistic world-picture.

Table 58: 3.31. Cascade logic from assumptions to optimisation

Assumptive commitment (Sections 3.1.1–3.1.3)	What it provides that optimisation requires
Determinism	A determinate mapping from choice to outcome, making an “optimal” path conceptually coherent (Laplacean structure; challenged in Chapter 5).
Certainty ideal	A fully specifiable decision problem in principle, including stable objectives and constraints (challenged in Chapter 5).
Perfect information	A knowable feasible set and the enumerability of relevant alternatives (challenged in Chapter 6).
Ergodicity	Stable probability structure, licensing intertemporal optimisation and expectation-based calculus (challenged in Chapter 5).

Assumptive commitment (Sections 3.1.1–3.1.3)	What it provides that optimisation requires
Equilibrium tendency	A stable target state that can be interpreted as the solution to optimising behaviour (challenged in Chapters 4–5).

Once this cascade is made explicit, the logic of failure becomes equally explicit. If uncertainty is fundamental rather than merely informational, the feasible set cannot be known because it expands endogenously through innovation and novelty. If cognition is bounded, optimisation ceases to be a plausible description of choice and becomes, at best, a benchmark. If the system is non-ergodic, the probability distributions required by intertemporal optimisation lose their stability. Under these conditions, optimisation is not merely computationally difficult; the object it claims to compute is often undefined.

Mirowski’s diagnosis captures the direction of the argument: as economic reality is understood more deeply – especially as uncertainty, complexity, and endogenous change are taken seriously – constrained optimisation does not become more empirically adequate; it becomes increasingly “meaningless” as a universal paradigm (Mirowski, 1989, p. 235). The claim here is not rhetorical but architectural: optimisation is a methodological output of mechanistic metaphysics.

III. What Calculus-Based Optimisation Presupposes

Optimisation in economic theory is typically calculus-based. This matters because calculus is not a neutral language; it carries structural preconditions. The standard solution apparatus – first-order conditions, envelope theorems, local curvature properties – requires continuity, differentiability, and, in many applications, convexity and separability. These are the mathematical conditions under which the optimisation template yields tractable and stable solutions.

Table 59: 3.32. Calculus requirements and their representational implications

What calculus-based optimisation requires	What this requirement excludes in CAS settings
Continuous choice sets and differentiable objective/constraint functions	Discrete, lumpy choices (investment, entry/exit), discontinuities, threshold effects, and regime shifts that are central to CAS dynamics.
Convexity (or conditions ensuring well-behaved optima)	Increasing returns, network effects, non-convexities, and multiple local optima that generate lock-in and path dependence (Arthur, 1989, 1994).
Separability or weak interaction structure	Strong interdependence, feedback loops, and relational constitution of payoffs in networks and ecosystems (Jackson, 2008).
Stable constraints and stable objective functions	Co-evolution of institutions, technologies, and preferences, where “constraints” are not given but altered by the very process being modelled (North, 1990, 2005).

Landini and colleagues summarise the broader implication in terms consistent with this dissertation’s diagnosis: optimisation and calculus can exhibit rigorous precision while remaining structurally misaligned with adaptive, emergent, nonlinear dynamics (Landini et al.,

2025). Fontana’s point complements this by noting that linear forms persist partly because they are solvable, whereas nonlinear dynamics often resist closed-form solution – an observation that is methodological but rooted in the deeper preference for tractability over ontological fit (Fontana, 2008).

IV. Optimisation and the “Choice Without Process” Error

Even when calculus conditions are granted, optimisation still imposes a particular representation of decision-making: it compresses time, learning, revision, and social interaction into a single static selection from a known set. The method thereby converts decision-making into “choice without process,” treating behaviour as the output of instantaneous calculation rather than the outcome of sequential, exploratory, and socially embedded adaptation.

This is the point at which Simon’s critique becomes decisive. Optimisation is not merely descriptively inaccurate; it presupposes a cognitive capacity and an environmental structure that real agents do not possess. Agents do not compute global optima; they satisfice, search, and stop when aspiration levels are met, because they face bounded cognition and uncertain environments (Simon, 1955, 1979). Procedural rationality is therefore not “imperfect” optimisation; it is a different rationality concept appropriate to real decision environments. Treating procedural rationality as a deviation preserves the optimisation benchmark by definition and thereby reproduces the optimisation fetish at the conceptual level.

The firm-level implication is not merely that optimisation is descriptively unrealistic; it is that optimisation is structurally the wrong explanatory primitive for organisations operating under endogenous novelty. Amazon is intelligible as a system of procedural rules that deliberately substitutes parallel experimentation, threshold-based escalation, and internal selection for global calculation. These rule systems are not approximations to an optimum; they are governance devices for decision-making when the relevant option set is not enumerable *ex ante* and when performance feedback arrives through time. The causal mechanism is therefore organisational: Amazon’s architecture creates variance (multiple experiments), subjects it to selective retention (scaling what works, terminating what fails), and institutionalises learning (codified routines for iteration). That mechanism generates outcomes that an optimisation representation would have to treat as exogenous “shifts in the objective function” or “changes in constraints,” precisely because the optimisation frame cannot represent the endogenous construction of options. In a non-ergodic environment, the firm’s rationality is procedural and evolutionary rather than substantive and calculative; the relevant unit of analysis is the rule ecology that governs search, not the putative maximiser that solves a complete problem (Simon, 1955, 1979; Nelson & Winter, 1982; Arthur, 2014).

Table 60: 3.32A. Process rationality versus solution rationality in innovation firms

What the optimisation frame requires	What the innovation firm must do instead
Enumerate alternatives and compute payoffs.	Generate alternatives through experimentation because the set is unknown <i>ex ante</i> .
Optimise over a stable objective and constraints.	Operate with shifting objectives/constraints as capabilities and markets co-evolve.
Treat learning as convergence to the optimum.	Treat learning as selection among evolving rules under noisy feedback.

What the optimisation frame requires	What the innovation firm must do instead
Make prediction central to choice.	Make robustness and adaptability central under fundamental uncertainty.

At the firm level, the same point holds with greater force. Firms do not “solve” a global profit-maximisation problem; they operate through routines and rule systems that evolve through experience, feedback, imitation, and selection (Nelson & Winter, 1982). Organisational cognition is distributed, strategic choice is iterative, and what counts as “the problem” changes as the organisation changes. In such conditions, the optimisation representation is less an approximation than a category error: it replaces adaptive process with a static solution concept.

V. Optimisation’s Temporal Blindness

Optimisation frameworks, particularly those anchored in equilibrium reasoning, systematically tend to treat time as either reversible or eliminable. Intertemporal optimisation can appear to “include” time, but it does so by imposing stable probability structure and stable objective functions across time, converting temporal evolution into a solvable present-value problem. Arthur’s formulation is instructive: standard neoclassical economics lives in a “timeless” equilibrium world, where time is solved away rather than treated as constitutive (Arthur, 2014). In the presence of non-ergodicity and endogenous novelty, the probability structures required for such intertemporal optimisation lose their foundation, because the future is not a draw from a stable distribution and the space of possibilities expands as the system evolves (Davidson, 2012; Kay & King, 2020).

Firm evolution makes the temporal issue concrete. Apple’s historical sequence – near-bankruptcy, iPod, iPhone, ecosystem services – cannot be represented as the unfolding of an optimal plan over a known state space, because later states were not simply selected; they were constructed through a path-dependent accumulation of capabilities and an evolving technological and institutional environment. In CAS terms, sequence matters because history is constitutive; optimisation frameworks typically treat history as information rather than as mechanism.

The temporal failure of optimisation is not an incidental limitation; it is a direct consequence of the method’s need to treat history as non-constitutive. Apple’s trajectory illustrates why. The firm’s major transitions were not selections from a pre-given menu of feasible strategies; they were sequential constructions in which earlier organisational and technological commitments created new capability states that then made later moves possible. On this view, the iPhone is not the solution to a timeless maximisation problem; it is the emergent outcome of a historically specific capability accumulation process in which installed base, complementary innovations, and organisational design co-produced the option itself. The explanatory mechanism is therefore path-dependent: once a platform logic is established, increasing returns, ecosystem complementarities, and learning-by-doing reshape both opportunities and constraints, creating lock-in dynamics and punctuated shifts that cannot be represented as movement toward a stable optimum. Any optimisation representation must either assume away the constitutive role of sequence or re-describe the decisive mechanism as an exogenous shock to technology or preferences – thereby conceding the core point: what matters most is precisely what optimisation cannot endogenise (Arthur, 1989, 1994, 2014; Nelson & Winter, 1982).

A short table can be useful here because it crystallises “history as mechanism” in one glance. Insert directly after the paragraph above:

Table 61: 3.33. *Why path-dependent platform evolution defeats optimisation*

Optimisation presupposition	Platform/evolutionary reality
The feasible set is given at time 0.	The feasible set is created through capability accumulation and ecosystem formation.
Dynamics are transitions across a stable state space.	The state space itself changes as complements, standards, and institutions evolve.
Increasing returns are treated as special cases.	Increasing returns are generative: they restructure incentives and trajectories.
History is information about parameters.	History is mechanism: sequence produces irreversibility and lock-in.

VI. Structural Incompatibility with CAS

The incompatibility between optimisation and CAS is therefore not a matter of missing variables or insufficient technique. It is structural. Optimisation requires well-defined objectives, a complete and stable possibility space, stable constraints, and a probability structure that makes intertemporal calculation meaningful. CAS environments, by contrast, are characterised by evolving objectives, expanding possibility spaces, co-evolving constraints, strong feedback, and non-ergodic dynamics.

Table 62: 3.33A. *Optimisation requirements versus CAS properties*

What optimisation requires	What CAS exhibits
Well-defined objective function	Objectives and aspiration levels evolve with learning, institutions, and identity; preferences can be constructed rather than given (Simon, 1955; North, 2005).
Complete, enumerable state space	Novelty expands the possibility space endogenously; unknown unknowns cannot be enumerated ex ante (Knight, 1921; Arthur, 2014).
Stable constraints	Constraints co-evolve with action through institutional change, technological change, and relational restructuring (North, 1990).
Calculable optimum under tractable curvature conditions	Feedback, non-convexity, and network effects generate multiple local optima, lock-in, and sensitivity to initial conditions (Arthur, 1989, 1994; Jackson, 2008).
Equilibrium as stable destination	Many relevant phenomena occur far from equilibrium, and may be better represented as shifting attractors or persistent disequilibrium rather than convergence (Arthur, 2014; Kirman, 2010).

VII. The Optimisation–Equilibrium Reinforcement

Optimisation does not merely coexist with equilibrium; in mechanistic economics the two concepts stabilise one another. If agents optimise, equilibria appear as the aggregation of optimal choices; if equilibria exist as stable destinations, optimisation appears as the

appropriate behavioural primitive. This mutual reinforcement is precisely why optimisation is so resilient to critique. Challenging optimisation threatens equilibrium as organising grammar; challenging equilibrium threatens the meaning of “optimality.” In the architecture of Section 3.2.1, these are high-centrality nodes protected by multiple loop connections, which is why attempted reforms often translate into constrained optimisation variants rather than genuine replacements.

VIII. Methodological Consequences for DSGE and General Equilibrium

The optimisation fetish culminates in the discipline’s dominant macro-formalisms. DSGE models adopt optimisation as their microfoundation, rational expectations as its epistemic companion, representative-agent or tractably aggregated heterogeneity as its structural simplification, and linearisation as a solvability requirement. Each of these steps reproduces the preconditions that render CAS phenomena methodologically invisible. Stiglitz’s critique targets this persistence at the level of assumptions rather than technique, arguing that DSGE retains widely discredited commitments even as it becomes more elaborate (Stiglitz, 2018). Colander and colleagues similarly diagnose a pattern in which rational expectations and related optimisation commitments persist despite inconsistency with evidence (Colander et al., 2009). Landini et al. interpret this persistence as a continuation of the physics analogy: optimisation methods derived from classical separable systems are extended to economic systems that are neither separable nor deterministic in the relevant sense (Landini et al., 2025).

This subsection therefore sets up Section 3.3.3 by identifying the methodological root: DSGE is not merely one modelling family among others; it is the institutionalised endpoint of the optimisation template.

Conclusion

This subsection has diagnosed the optimisation fetish as the central methodological symptom of mechanistic economics. The core claim is structural: optimisation is not primarily chosen because it fits economic reality, but because it is the methodological expression of mechanistic assumptions about reality and knowledge. Once the cascade from determinism, certainty, perfect information, ergodicity, and equilibrium is made explicit, optimisation’s universality becomes intelligible as a disciplinary consequence rather than a scientific discovery.

The diagnosis yields three results. First, calculus-based optimisation is shown to require continuity, differentiability, convexity, and separability, conditions that are systematically violated in complex adaptive systems. Second, optimisation is shown to impose a “choice without process” representation, collapsing sequential adaptation and learning into a static selection problem, contrary to bounded and procedural rationality (Simon, 1955, 1979). Third, optimisation is shown to erase historical time by requiring stable probability structure and stable objectives, contradicting non-ergodicity, path dependence, and endogenous novelty (Knight, 1921; Arthur, 2014; Davidson, 2012).

The implication is methodological and architectural. Optimisation cannot be rescued by marginal adjustments if the system being studied is CAS-like. It must be demoted from universal method to special case, and replaced – at the level of agency – by procedural and ecological rationality, satisficing, and rule-based adaptation, which the complexity economics literature identifies as appropriate behavioural primitives for open, evolving, uncertain environments (Simon, 1955, 1979; Gigerenzer et al., 1999; Arthur, 2014). The present thesis establishes the ontological conditions under which such replacement becomes necessary.

3.3.3 DSGE, Representative Agent, and Model Monism

Introduction

Sections 3.3.1 and 3.3.2 diagnosed the methodological symptoms of the mechanistic paradigm: physics envy and the optimisation fetish. This subsection examines the institutional culmination of those symptoms in macroeconomics: DSGE as a dominant methodological regime. The claim is not that DSGE is “a model” that happens to be popular; it is that DSGE is the paradigmatic form in which mechanistic metaphysics becomes operational within central banks, journal gatekeeping, and graduate training. In this sense, DSGE is the point where abstract assumptions become policy-relevant standards of “rigour”.

This subsection performs three diagnostic tasks. It first reconstructs DSGE as a methodological regime by showing how its defining components encode the assumptions mapped in Sections 3.1–3.2. It then treats the representative agent as an extreme reductionist manoeuvre that eliminates precisely those social properties – interaction, topology, heterogeneity, emergence – that define complex adaptive systems. Finally, it links DSGE’s persistence to the lock-in mechanisms established in Section 3.2, explaining why repeated failures do not trigger hard-core revision but instead generate protective-belt elaborations.

The stakes are institutional rather than merely academic. DSGE variants shape policy practice at the Federal Reserve System, the European Central Bank, and the Bank of England, and they function as a de facto standard in high-status publication and training. If DSGE is structurally incompatible with the CAS nature of economies, then the discipline’s core macro methodology is systematically mis-specified in exactly the domain where mis-specification is most costly.

I. DSGE as a Dominant Methodology

Within our analytical framework, DSGE is best understood as an institutionalised modelling grammar. Its rise is not explained by one decisive empirical victory, but by its fit with the mechanistic assumption network and by its compatibility with the discipline’s stabilisation mechanisms (Section 3.2.2). The historically salient point is not simply that DSGE displaced earlier approaches, but that it became an equilibrium-centred language of policy legitimacy and professional certification.

The regime character is visible in the way DSGE became simultaneously (i) a standard apparatus for central-bank analysis, (ii) a publication expectation in core outlets, and (iii) a pedagogical template reproduced through graduate training. These three functions – policy legibility, journal gatekeeping, and curricular monopoly – are not accidental add-ons. They are the institutional expression of the methodological stabilisation mechanism diagnosed in Section 3.2.2: what cannot be expressed in the authorised formal language is routinely classified as lacking rigour.

To make the claim precise, Table 3.34 identifies the internal meaning of the DSGE label and the assumptions required by each component.

Table 63: 3.34. What “DSGE” operationally means and what it presupposes

DSGE component	Operational meaning	Assumptions required (Sections 3.1–3.2)
Dynamic	intertemporal optimisation represented as Euler equations	optimisation as behavioural primitive; certainty/knowability sufficient for solution (3.1.2; 3.3.2)
Stochastic	uncertainty introduced as exogenous shocks	ergodic, stable probability structure; novelty treated as disturbance not mechanism (3.1.3)
General equilibrium	simultaneous market clearing and model closure	closure plus equilibrium as organising grammar; decomposability for tractability (3.1.1; 3.1.3)
Microfoundations	macro built from “optimising agents”	reductionism plus homogeneity; aggregation legitimacy; representative-agent admissibility (3.1.1; 3.1.3)

The analytical implication is immediate. In DSGE, “dynamics” does not mean open-ended evolution; it typically means convergence along an equilibrium path. “Stochastic” does not mean endogenous instability or uncertainty of the Knightian kind; it means disturbances drawn from a known distribution. “Microfoundations” does not mean preserving heterogeneity, interaction, and emergence; it frequently means collapsing those features into a single optimiser. The regime therefore instantiates the precise pattern established in Section 3.1.3: derivative assumptions (equilibrium, ergodicity, optimisation, decomposability) are not optional modelling choices but necessary supports for the authorised method.

This is why a substantial strand of criticism targets DSGE not on aesthetic grounds but on ontological credibility. In testimony to the U.S. House (2010), Robert M. Solow characterised DSGE models as “fantasy constructs” and argued that their internal coherence did not compensate for their empirical implausibility. The point is not that DSGE is “too abstract” in a generic sense, but that it must be abstract in a particular mechanistic way in order to be solvable.

Firm-level translation. The DSGE grammar removes the firm as an explanatory object because it must treat production and organisation as tractable, aggregable constraints inside an intertemporal programme. A conglomerate such as General Electric does not enter macro reality primarily as a smooth production function with stable technology and a well-defined optimisation problem; it enters as a historically evolving organisational architecture spanning heterogeneous divisions, governance regimes, and networked supplier and customer relationships. The representative-agent/representative-firm move erases that heterogeneity by construction, and the remainder – an aggregate production block – cannot represent the causal mechanisms through which such organisations generate (or destroy) resilience, propagate shocks, or reconfigure portfolios. In other words, DSGE does not “simplify” the firm; it changes the object of explanation from organisation to equation.

II. DSGE as a Self-Contained Equilibrium Universe

The regime’s distinctive feature is not simply equilibrium analysis, but a self-contained equilibrium universe in which equilibrium is treated as the default grammar of intelligibility. DSGE typically presupposes that the economy is sufficiently closed, stable, and stationary for deep parameters to be treated as time-invariant; that deviations are small enough to justify local

approximation; and that uncertainty can be represented as risk over known distributions. These commitments are not secondary technicalities: they are the mechanistic ontology and epistemology made operational.

Table 3.35 makes the mapping explicit, since the diagnostic point depends on showing that DSGE encodes, rather than merely accompanies, the mechanistic assumptions already identified.

Table 64: 3.35. DSGE's equilibrium universe as an encoding of mechanistic assumptions

DSGE modelling commitment	Mechanistic assumption encoded	CAS feature excluded
deep, stable parameters	substance ontology; determinism (3.1.1)	structural change; endogenous novelty
shocks as exogenous disturbances	closure; ergodicity (3.1.1; 3.1.3)	endogenous instability; non-ergodic dynamics
equilibrium as reference and destination	equilibrium tendency; certainty ideal (3.1.3; 3.1.2)	persistent disequilibrium; evolving attractors
solvability via approximation	linearity; decomposability (3.1.3)	tipping points; cascades; phase transitions
microfoundations via optimisation	optimisation; perfect information proxy (3.1.2; 3.3.2)	heuristics; learning; procedural rationality

This table states a strong claim: DSGE does not merely omit CAS phenomena; it must exclude them because their inclusion breaks the assumptions required for the toolkit. This is precisely why Landini et al. (2025, p. 11) describe the extension of the physics analogy through Arrow–Debreu and DSGE as an “abstract mathematisation” deprived of phenomenology: the form is rigorous, but the ontology required to make the form work is not the ontology of the real economy.

Firm-level translation. The “equilibrium universe” mis-specifies firm dynamics because it forces organisational transformation to appear as an external perturbation rather than an endogenous process. A firm such as Siemens has repeatedly reconfigured its portfolio in response to technological and regulatory transitions, shifts in geopolitics, and the co-evolution of platforms and industrial software. Those transitions are not “small deviations” from a stable steady state; they are regime changes in the firm’s own capability structure and network position. A methodology that treats change as mean-reverting disturbance cannot represent the firm’s central causal mechanism under modern capitalism: endogenous reorganisation under structural uncertainty.

III. The Representative Agent as Elimination of Society

The representative agent is often defended as “a simplification”, but within our analytical framework it is better understood as a tractability device demanded by the equilibrium-optimisation regime. Once heterogeneity, interaction, network topology, learning, and institutional differentiation are admitted as constitutive rather than residual, aggregation generally ceases to preserve behavioural form. The representative agent therefore performs a decisive operation: it replaces society with a scaled-up individual and thereby converts a social system into a solvable optimisation problem.

The analytical costs of this move are not incidental. The representative agent deletes exactly those properties through which CAS generate macro outcomes: interaction-driven cascades, distributional dynamics, and emergent coordination failures. This is why Alan Kirman frames methodological individualism and representative-agent reasoning as unable to specify the mechanism that would generate social order from isolated optimisers (Kirman, 2016, pp. 539–540), and why Kirman’s aggregation critique (Kirman, 1992) remains decisive: representative-agent reduction is valid only under knife-edge conditions that collapse the very heterogeneity economics needs to explain.

Table 3.36 records what the representative agent removes and why the removal is not harmless.

Table 65: 3.36. What the representative agent removes and what is lost thereby

Removed by representative agent	Why it matters for explanation	CAS consequence
interaction among agents	macro outcomes emerge from interdependence	networks and cascades become invisible
heterogeneity as structure	distributions shape dynamics and stability	aggregation erases the engines of instability
learning and adaptation	decision rules evolve endogenously	equilibrium convergence is imposed by design
feedback and reflexivity	behaviour changes the environment and the model itself	self-reference is suppressed
institutional embedding	rules, norms, and organisations co-produce behaviour	institutions become “frictions”
mesoscale organisation	firms, industries, fields are causal layers	the “missing middle” persists (Section 3.1.4)

Firm-level translation. The representative-agent move destroys the meso-level by construction. If Amazon and Macy’s are reduced to the same representative firm, then the objects that actually drive divergence – platform logic versus retail real-estate logic, data capability versus brand-and-location logic, network externalities versus local foot traffic – cannot appear as causal mechanisms. The reduction does not merely compress description; it replaces an explanatory problem about heterogeneous organisational architectures with a solvable optimisation problem about a fictional average. The methodological result is not “approximation error” but ontological deletion.

IV. The Microfoundations Fallacy

DSGE’s prestige is often tied to the claim of “rigorous microfoundations”. Yet the representative-agent manoeuvre implies that what is called microfoundations is frequently macro structure re-labelled as micro. The model begins with an aggregate object chosen for tractability and then treats that object’s optimisation as if it were a derivation from genuine micro-level diversity.

This is why Colander et al. (2009, p. 260) diagnose DSGE as having weak microfoundations and unsatisfactory empirical performance: the “first principles” are not discovered causal primitives but modelling requirements. The deeper point is Kirman’s: genuine microfoundations would preserve heterogeneity and interaction and then show how macro

regularities emerge. Representative-agent DSGE does the opposite: it begins by deleting heterogeneity and interaction so that equilibrium and optimisation remain solvable.

Firm-level translation. At the mesoscale, the microfoundations claim becomes especially fragile because firms are not interchangeable tokens. Macro outcomes hinge on organisational differentiation – capability distributions, network centrality, and institutional constraints. A macro method that substitutes a representative firm for an organisational distribution cannot claim to explain macro outcomes from micro behaviour; it can only claim to generate internally consistent dynamics for a fictional aggregate. The logic of explanation is therefore reversed: micro is reconstructed to fit macro tractability rather than macro being derived from micro generative mechanisms.

V. Linearisation as Non-Linearity Denied by Decree

Even when DSGE begins with non-linear structures, solvability typically requires approximation around a steady state. This local-linearisation procedure is often presented as a practical concession. Within our analytical framework, it performs a stronger function: it rules out by method exactly those non-linear dynamics that define crises and structural change.

Table 3.37 specifies what the approximation requires and what it deletes.

Table 66: 3.37. Linearisation requirements and the dynamics removed

Linearisation requirement	What it removes	CAS reality thereby excluded
small deviations around steady state	large shocks and crisis dynamics	crises as regime shifts, not local noise
smooth local dynamics	thresholds and discontinuities	tipping points; phase transitions
unique reference equilibrium	multiple attractors	shifting basins of attraction
mean reversion in perturbations	structural breaks	non-ergodic evolution

The diagnostic claim is therefore strict: linearisation is not a simplification of crisis dynamics; it is an elimination of crisis dynamics as a representable object. If crises are the outcome of non-linear amplification, balance-sheet feedback, and network propagation, then a methodology that reduces dynamics to local deviations cannot explain crises except by re-describing them as exogenous disturbances. This is one reason Stiglitz (2018, p. 4) targets DSGE’s “widely discredited assumptions”: the method’s equilibrium-centred approximation excludes the causal structure of systemic crises.

Firm-level translation. Linearisation also mis-specifies industrial disruption. The transition from Blockbuster to Netflix was not a smooth adjustment around a stable industrial equilibrium; it was a phase transition driven by threshold effects (streaming quality and broadband feasibility), feedback loops (subscribers ↔ content investment), and network propagation (platform adoption reshaping industry structure). A linearised equilibrium-path method can only register such a transition as a perturbation that should wash out, which is analytically equivalent to assuming away the phenomenon.

VI. Model Monism as Methodological Authoritarianism

The dominance of DSGE in policy institutions and high-status outlets expresses a deeper epistemic commitment: model monism. In a monist regime, one model class becomes the default standard of legitimacy; alternatives are not treated as context-appropriate complements but as deviations from “rigour”. Within the dissertation’s architecture, monism is not a sociological accident; it is a logical consequence of mechanistic epistemology (Section 3.1.2) and of the methodological stabilisation mechanism (Section 3.2.2).

Table 3.38 records the operational symptoms of monism and their methodological effects.

Table 67: 3.38. Operational symptoms of DSGE monism and their effects

Symptom	Where it appears	Effect on knowledge production
alternative frameworks are marginalised	journals and curricula	loss of methodological diversity
DSGE structure becomes an expectation	publication gatekeeping	“rigour” equated with one grammar
convergence of policy models	central-bank practice	policy uniformity despite heterogeneity of economies
training monopoly	PhD cores	intergenerational reproduction of the regime

The implication is not anti-formalism but anti-monism. If economies are open, non-ergodic, and institutionally differentiated, then methodological pluralism is not a matter of taste but a requirement of ontological fit. The present diagnosis identifies DSGE monism as the methodological negation of that requirement; the formal articulation of pluralism as a methodological principle is beyond the scope of this thesis but follows as a necessary implication of the ontological reconstruction developed in Chapter 4.

Firm-level translation. Monism is particularly destructive for firm analysis because the firm is multi-perspectival by nature. A firm such as Boeing is simultaneously a production system, a governance structure, a capability ecology, a network node, and a politically regulated organisation. Any single grammar – production-function optimisation, transaction-cost minimisation, or representative-firm aggregation – captures at most one projection. Monism then converts partial projection into total explanation. The result is not parsimonious theory; it is systematic blindness to the mechanisms that matter most at the mesoscale.

VII. The Failure Record and the Persistence Mechanism

DSGE’s empirical and policy failures are now widely documented, but within our framework the decisive point is that the failures are structurally implied by the assumptions encoded in the method. DSGE systematically struggles with crises, contagion, structural breaks, and institutional change because these phenomena violate the equilibrium-centred, ergodic, locally approximated world that makes the method workable.

Table 3.39 states the mapping from failure domain to the assumption that makes the domain hard to represent.

Table 68: 3.39. Failure domains and the assumptions that generate them

Failure domain	What must be represented	Assumption that blocks representation
systemic crises	non-linear amplification and cascading failures	linearisation; equilibrium-centred dynamics
contagion and spillovers	network topology and propagation	representative agent; decomposability
structural breaks	shifting regimes and institutions	ergodicity; fixed parameters
endogenous innovation	novelty and expanding possibility space	closure; determinism in practice

Kirman (2010, as cited in Landini et al., 2025, p. 1) describes 2008 as both a crisis of the real economy and a crisis of dominant theory, and Kirman (2017, p. 37) notes that repeated crises should count as evidence of deep analytical problems. Yet DSGE persists. The persistence is explained by the lock-in architecture already established: Kuhnian normal science converts failures into technical puzzles (Section 3.2.3), Lakatosian negative heuristics protect the hard core (Section 3.2.4), and institutional stabilisation selects for paradigm-legible work (Section 3.2.2). DSGE survives not because it predicts well, but because it fits the regime’s criteria of admissibility.

VIII. The DSGE–Monism Institutional Loop

The persistence dynamic can be expressed as an institutional instantiation of the general feedback architecture from Section 3.2.1. DSGE is enabled by mechanistic assumptions; DSGE then reproduces the authority of those assumptions by monopolising what counts as rigour in policy and publication; that authority then reproduces DSGE through training and gatekeeping. The loop therefore has an epistemic face (what is admissible), a methodological face (what is solvable), and an institutional face (what is rewarded). This is why DSGE cannot “self-correct” in the strong sense: corrections that would restore ontological fit would also dissolve the assumptions that make DSGE the authorised method.

Conclusion

This subsection has treated DSGE as the institutional apex of mechanistic methodology rather than as a contingent modelling choice. DSGE encodes the mechanistic assumptions mapped in Sections 3.1–3.2, and its core devices perform eliminations with direct explanatory consequences: the representative agent removes society and the mesoscale; linearisation removes non-linearity and regime change; model monism removes methodological pluralism and thereby suppresses CAS-compatible alternatives.

The diagnostic implication follows the logic already established in Section 3.2. DSGE persists because it is structurally aligned with a triple-locked paradigm and because it is reproduced by semantic, methodological, and institutional stabilisation. It cannot self-correct without ceasing to be DSGE in the strong sense – because self-correction would require relaxing precisely those assumptions that define its solvability and its institutional legitimacy. Section 3.3.4 therefore turns to a complementary symptom of the same regime: the selection of elegance over truth as an explicit norm of model choice.

3.3.4 “Elegance Over Truth” and the Aestheticization of Economics

Introduction

Physics envy (Section 3.3.1) imported the methods of classical mechanics into economics. The optimisation fetish (Section 3.3.2) elevated calculus-based maximisation to a universal decision principle. DSGE (Section 3.3.3) institutionalised both as the dominant macroeconomic methodology. The remaining question is why these approaches continue to command allegiance despite repeated empirical failure. The answer lies in aesthetics: the discipline has made mathematical elegance the primary criterion of scientific legitimacy.

This subsection diagnoses the aestheticisation of economics – the elevation of formal beauty, tractability, and internal coherence above correspondence with economic reality. Elegance operates as an epistemic filter: it determines which phenomena are admissible, which methods are legitimate, and which failures are tolerated. Within this regime, realism appears as a defect rather than a virtue.

For our framework, this subsection performs three diagnostic functions. First, it identifies elegance as an epistemic filter that governs what counts as acceptable economic knowledge. Second, it shows how aesthetic criteria protect the mechanistic paradigm from empirical refutation, explaining why crises such as 2008 did not trigger paradigm change. Third, it prepares the transition from methodological critique to ontological reconstruction by demonstrating that transformation requires altering aesthetic standards, not merely replacing models.

The stakes are epistemological. When elegance defines legitimacy, the messy features of actual economies – heterogeneity, networks, uncertainty, emergence – are excluded not because they are unimportant, but because they are inelegant. This subsection exposes elegance as a stabilising mechanism of a failing paradigm.

I. Mathematical Elegance as Criterion of Legitimacy

The Elegance Equation

Within the inherited body of economic theory, mathematical tractability is routinely equated with scientific rigour. This converts method into aesthetic judgement. Models are evaluated less by their capacity to represent economic mechanisms than by their conformity to formal standards inherited from classical physics.

Table 69: Mathematical Aesthetics as Substantive Theoretical Commitments: How Formal Tractability Requirements Exclude Empirically Relevant Phenomena

Aesthetic Criterion	Mathematical Requirement	What It Excludes
Smooth utility functions	Continuity, differentiability	Discontinuous preferences, thresholds
Convex sets	Global optimum existence	Increasing returns, multiple equilibria
Differentiable production	Calculus applicability	Discrete technologies, lumpy investments
Linearised dynamics	Analytical solvability	Non-linear cascades, tipping points
Symmetric equilibria	Closed-form solutions	Asymmetric interaction, heterogeneity

Aesthetic Criterion	Mathematical Requirement	What It Excludes
Closed-form solutions	Explicit answers	Computation, emergence

These requirements are not empirical discoveries. They are aesthetic choices that define admissibility. Mathematical beauty is treated as evidence of scientific validity.

Arthur (2021, p. 143) characterises neoclassical economics as envisioning the economy as a “well-functioning machine,” whose elegance and precision are achieved by suppressing complexity. Arthur (2014, p. 191) notes that this physics-inspired approach systematically prioritised mathematical elegance over real-world dynamics.

Canonical Examples

Table 70: Elegance Versus Empirical Adequacy: The Persistent Disconnect Across Canonical Mainstream Economic Models

Model	Elegance	Reality	Disconnect
Arrow–Debreu	Existence proofs, completeness	Requires impossible conditions	Formally rigorous, empirically empty
Lucas island	Elegant information structure	Implausible behavioural premises	Elegant but wrong
RBC	Clean Euler equations	Implausible shocks	Shocks fitted ex post
DSGE	Institutional apex	Repeated predictive failure	“Too beautiful to abandon”

Landini et al. (2025, p. 11) observe that Arrow–Debreu is mathematically unassailable when decoupled from phenomena, yet incorrect when treated as an economic model. Mathematical beauty and empirical adequacy are decoupled, with beauty taking precedence.

At the firm level, the same aesthetic logic governs model selection. Strategic and organisational representations are valued not for explanatory accuracy but for tractability. Canonical abstractions persist because they are solvable, not because they capture how firms coordinate, learn, and adapt. What disappears are the causal mechanisms that generate performance differences – routines, learning trajectories, supplier relations, and historically accumulated capabilities – because these mechanisms violate formal elegance.

II. Tractability Über Alles: The Realism Trade-Off

Within this regime, realism is systematically sacrificed for solvability. The claimed functions of elegant models – prediction, policy guidance, and understanding – fail repeatedly, while their sole success lies in internal puzzle-solving within the paradigm.

Table 71: Mainstream Economics Evaluated Against Its Own Claimed Functions: Systematic Failure and the Circularity of Puzzle-Solving Success

Claimed Use	Actual Performance	Assessment
Prediction	Crises unpredicted	Not achieved
Policy	Post-crisis policy failure	Not achieved

Claimed Use	Actual Performance	Assessment
Understanding	Mechanisms misrepresented	Not achieved
Puzzle-solving	Works internally	Circular success

Fontana (2010, p. 19) notes that linear mathematics dominates economics not because it is accurate, but because it is elegant and always solvable.

This trade-off is especially damaging in firm theory. Firms are represented as production functions because such representations are tractable, not because they are true. Toyota's sustained advantage, for example, did not emerge from cost-minimising optimisation over given technologies, but from the endogenous evolution of routines, supplier integration, and continuous improvement practices. These mechanisms are historically constructed coordination systems, not equilibrium outcomes. Treating Toyota as a profit-maximising production function does not simplify its behaviour; it erases the causal structure that explains it. Elegance substitutes for explanation.

III. Elegance Suppresses Messy Reality

Elegant models require a simple world. The world is therefore simplified until it becomes mathematically graceful.

Table 72: Features Excluded on Grounds of Mathematical Inelegance and Their Essential Roles in Complex Adaptive Systems

Excluded Feature	Why Inelegant	CAS Role
Heterogeneity	Breaks aggregation	Drives dynamics
Networks	Break symmetry	Structure outcomes
Discontinuities	Not differentiable	Phase transitions
Tipping points	No closed form	Regime change
Fundamental uncertainty	Destroys smoothness	Ontological
Institutions	Path-dependent	Co-evolution
Bounded rationality	Heuristics “ugly”	CAS agency
Technological novelty	Breaks equilibrium	Endogenous change

Gallegati and Gallegati (2025, p. 3065) note that mathematical rigour is often preferred to empirical evidence. Fontana (2008a, p. 14) concludes that understanding complexity requires taking distance from conventional mathematical modelling.

Amazon is excluded by elegant models not because it is anomalous, but because it violates their aesthetic requirements. Its evolution – from retailer to logistics network to cloud

infrastructure to AI platform – cannot be represented as optimisation around a stable equilibrium. Each stage emerged from feedback between routines, data accumulation, network effects, and organisational experimentation. These dynamics are non-linear, path-dependent, and generative. They are excluded by construction, not by evidence.

IV. The Aestheticisation Problem

Physics envy + optimisation fetish = aestheticisation

Table 73: The Aesthetic Displacement of Empirical Standards: Five Shifts in Evaluative Criteria Within Mainstream Economics

Shift	From	To
Criterion	Empirical adequacy	Formal beauty
Priority	Substance	Style
Model choice	Explanatory power	Elegance
Reality	Object of study	Obstacle
Anomalies	Revise theory	Dismiss reality

The same aesthetic logic applies across contemporary firms often treated as “exceptions.” Apple’s transformation into an ecosystem organiser, Tesla’s co-evolution of technology and regulation, and Patagonia’s integration of governance, purpose, and market positioning all depend on feedback-rich organisational processes. These firms succeed not by solving optimisation problems but by navigating evolving constraint landscapes. Their exclusion reveals the paradigm’s limits.

V. Aesthetic Blindness

Model selection becomes aesthetic selection. Reality is filtered by formal taste.

Table 74: Consequences of Aesthetic Dominance in Mainstream Economics: How Formal Elegance Displaces Empirical Reality

Consequence	Manifestation
Reality ignored	“Ugly” facts excluded
Crises	Treated as anomalies
Validation	Internal to models
Truth	“True in models”

VI. Elegance as Paradigm Protection

Elegance functions as a stabilising device.

Table 75: *Elegance as Paradigm Protection: Mechanisms by Which Mathematical Aesthetics Stabilise Mainstream Economic Theory*

Mechanism	Effect	Reference
Frictions	Reality downgraded	3.2.2
Protective belt	Core preserved	3.2.4
Normal science	No rupture	3.2.3

After 2008, DSGE survived not because it explained the crisis, but because it remained elegant. Reality was reclassified; beauty was preserved.

VII. Formalism Without Substance

Economics increasingly confuses elegance with truth, coherence with adequacy, and mathematics with understanding. This is methodological degeneration in Lakatos's sense: institutional sophistication combined with scientific stagnation.

Conclusion

Economics did not choose elegance because the world is simple. It chose simplicity so the world would appear elegant.

Complex adaptive systems are not beautiful in the classical sense. They are irregular, non-linear, historically contingent, and computationally demanding. CAS-compatible methods may appear inelegant – but they are truthful. Methodological transformation therefore requires abandoning aesthetic criteria inherited from mechanistic physics and replacing them with standards grounded in ontological accuracy and explanatory fit.

This prepares the ground for Hayek's critique of the pretence of knowledge (Section 3.3.5) and for the assumption-level reconstruction undertaken in Chapter 4.

3.3.5 The Pretense of Knowledge (Hayek)

Introduction

Physics envy (Section 3.3.1) borrowed the methods of classical mechanics. The optimisation fetish (Section 3.3.2) made calculus-based maximisation mandatory. DSGE (Section 3.3.3) institutionalised both. Elegance (Section 3.3.4) aestheticised the result. The remaining question is why economists continue to believe these methods work despite persistent predictive and explanatory failure. Friedrich Hayek's answer remains decisive: the discipline suffers from a *pretence of knowledge*.

This subsection diagnoses the epistemic foundation of the mechanistic paradigm: the belief that economies can be known like physical systems, that mathematical precision is equivalent to genuine understanding, and that prediction-based control is attainable. Hayek showed that this is not scientific confidence but *scientism* – the form of science without its substance.

Within our framework, this subsection performs four diagnostic functions. First, it identifies epistemic hubris as the unifying foundation behind the methodological symptoms diagnosed in Sections 3.3.1-3.3.4. Second, it provides the philosophical warrant for rejecting prediction-centric epistemology in economics. Third, it establishes the Prediction–Control Collapse,

showing that prediction-based control is not merely difficult but structurally incoherent for complex adaptive systems. Fourth, it completes the bridge from methodological critique to ontological reconstruction, preparing directly for the foundational work developed in Chapter 4.

The stakes are epistemological. When economists believe they know more than can be known, empirical failure is systematically reclassified as a technical problem rather than recognised as a conceptual limit. The pretence thus functions as a paradigm-protecting device.

This is the fifth and final subsection of Section 3.3. Its completion closes the methodological diagnosis and prepares the transition to Section 3.4 (Consequences) and Chapter 4.

I. Hayek's Nobel Lecture Diagnosis

The Core Argument

Hayek's Nobel Lecture, *The Pretence of Knowledge* (1974), provides a philosophical key to understanding why mechanistic economics persists despite failure. Hayek's claim was not that economists lack intelligence or effort, but that they systematically *overestimate what can be known* about complex social systems.

Within our analytical framework, Hayek's diagnosis consists of four interlinked errors:

Table 76: *Elements of Methodological Overreach: How Mainstream Economics Conflates Scientific Form with Scientific Validity*

Element	Problem Identified	Connection
Physics envy	Importation of physical-science methods	Section 3.3.1
Data illusion	Measurability mistaken for relevance	Section 3.3.2
Intervention hubris	Planning assumes centralised knowledge	Section 3.3.3
Scientism	Scientific form mistaken for scientific validity	Section 3.3.4

Hayek argued that in complex social systems, “the most scientific procedure is often the most unscientific” because such systems consist of “structures of essential complexity” whose outcomes depend on patterns of interaction that cannot be fully observed or aggregated (Hayek, 1989, pp. 4–6). The root error lies in the “erroneous transfer to social phenomena of the habits of thought we have developed in dealing with the phenomena of nature” (Hayek, 1945, p. 520).

Economic outcomes reflect dispersed, tacit, and context-specific knowledge possessed by millions of agents. Because this knowledge is not centrally available, predictive capacity is structurally constrained. The pretence arises when economists ignore these limits and proceed as if the relevant information were objectively given and aggregable.

At the firm level, the same epistemic error appears when quantifiable indicators are privileged because they are measurable, while tacit knowledge embedded in routines, organisational culture, and path-dependent capability formation is treated as secondary. “Fact-based” strategy becomes an instrument of the data illusion when it presumes that the future can be inferred from available metrics.

II. Scientism: Form Without Substance

Scientism, in Hayek’s sense, is the use of the *form* of science – precision, mathematics, prediction – without the *conditions* that make those forms epistemically valid. Classical physics achieves precision because its objects are approximately decomposable, closed, and governed by relatively stable laws. Social systems are not.

Table 77: *Why Econometric Assumptions Fail in Complex Economies: Requirements and Their Points of Breakdown*

Assumption	What It Requires	Why It Fails in Economies
Complete data	All relevant variables observable	Knowledge is dispersed and tacit
Stable parameters	Invariant relationships	Agents learn and co-evolve
Time invariance	Past predicts future	Endogenous novelty
Decomposability	Parts analysable in isolation	Emergence from interaction
Reversibility	History irrelevant	Path dependence

Economics adopted physics’ mathematical form without verifying whether economic phenomena satisfy physics’ domain conditions. Soros (2009) describes this as a false analogy: reflexive systems change as participants respond to models, undermining the stability required for prediction.

III. Epistemic Hubris versus Epistemic Humility

Each core mechanistic assumption presumes knowledge that cannot be possessed:

Table 78: *Presuppositions Underpinning Empirical Testability in Mainstream Economics: What Each Assumption Requires of Economic Systems*

Assumption	What It Presumes
Certainty	Future states are knowable
Perfect information	Relevant facts are observable
Optimisation	Global optima are calculable
Equilibrium	System behaviour is determinate
Ergodicity	Past distributions reveal future ones

Hayek’s counter-position is epistemic humility. He emphasised fundamental uncertainty (Knight), dispersed and tacit knowledge (Hayek; Polanyi), spontaneous order, and the irreducibility of historical processes. In such environments, prediction must be limited to *patterns* rather than precise outcomes.

As complexity increases, Hayek argued, we are confined to “pattern predictions” – ranges, tendencies, and structural constraints – not point forecasts of specific future states (Hayek, 1989, p. 7).

At the firm level, the collapse of LTCM illustrates epistemic hubris under ergodicity assumptions. Models calibrated on historical correlations failed when those correlations shifted endogenously. The failure was epistemological, not merely technical.

IV. False Precision and the Illusion of Control

Mechanistic economics rests on the belief that precision enables control.

Table 79: Mainstream Economic Beliefs Versus Complex Adaptive System Realities: The Systematic Mismatch Undermining Policy and Prediction

Belief	Implication	CAS Reality
Economy is predictable	Future calculable	Endogenous novelty
Past resembles future	Ergodicity holds	Structural breaks
Optimal equilibria exist	Solutions computable	Multiple attractors
Policy can fine-tune	Control feasible	Resistance to control

Precision in measurement does not imply precision in prediction. Gallegati and Gallegati (2025) show that reductionist and linear techniques persist despite systematic mismatch with complex dynamics. Landini et al. (2025) note that economic problems increasingly become mathematical problems detached from reality.

At the organisational level, Enron exemplifies false precision: sophisticated valuation models produced confidence while detaching decision-making from operational realities.

V. The Prediction–Control Collapse

A core result follows.

In complex adaptive systems, prediction-based control is structurally incoherent.

The logic is straightforward. CAS exhibit non-linearity, endogenous novelty, shifting parameters, and dispersed knowledge. Precise point prediction is therefore structurally constrained. Control presupposes reliable prediction of system response to intervention. When prediction fails, control collapses.

Application Across Levels

Table 80: The Pretence of Knowledge Across Micro, Meso, and Macro Levels: Agent Claims and Their Collapse Under CAS Realities

Level	Agent	Pretence	Collapse
Micro	Individuals	“I can optimise over a knowable future”	Future is ontologically open
Meso	Firms/managers	“I can plan via forecasts”	Competition and innovation are emergent
Macro	Policymakers	“I can fine-tune the economy”	System-level emergence resists control

The Alternative: Navigation

Table 81: From Prediction to Navigation: Five Paradigmatic Shifts Toward CAS-Compatible Economic Reasoning

From (Prediction-based)	To (CAS-compatible)
Predict and control	Navigate and adapt
Point forecasts	Scenario preparation
Optimal strategies	Robust strategies
Fine-tuning	Resilience-building
Eliminating uncertainty	Embracing uncertainty

The alternative is not resignation but *navigation*.

Toyota’s adaptive routines illustrate navigation under uncertainty; GM’s pre-2008 planning systems illustrate fragility under prediction-based control.

VI. The Precision–Explanation Paradox

As models become more precise – more parameters, more calibration – explanatory adequacy declines. Akerlof (2020) calls this a “sin of omission”: essential phenomena are ignored because they resist formalisation. Mirowski (1989) captures the paradox: increased knowledge renders optimisation increasingly meaningless.

At the firm level, production functions gain mathematical precision while explaining less about how capabilities actually emerge. Amazon’s evolution – from retail to cloud infrastructure to AI – cannot be derived from optimisation over known state spaces; it emerged from feedback between routines, data accumulation, and experimentation under uncertainty.

VII. The Pretence as Paradigm Protection

The pretence of knowledge functions as a stabiliser.

Table 82: Rhetorical Mechanisms of Paradigm Defence: How Mainstream Economics Deflects Fundamental Critique

Mechanism	Effect
“We need more data”	Uncertainty denied
“Precision equals rigour”	Elegance justified
“Tail risk”	Failure reclassified
“Not rigorous”	Alternatives dismissed

Failure triggers technical repair rather than assumption-level revision. The conclusion is rarely “this cannot be known.” Instead, failure becomes a demand for more precision, reinforcing the very hubris that produced it.

VIII. Hayek as Proto-CAS Thinker

Hayek anticipated core CAS insights in verbal form:

Table 83: Mapping Hayekian Concepts to Their Complex Adaptive Systems Equivalents

Hayek	CAS Equivalent
Spontaneous order	Emergence
Dispersed knowledge	Fundamental uncertainty
Competition as discovery	Formation, not allocation
Evolutionary institutions	Co-evolution
Pattern prediction	Bounded predictability

Hayek lacked formal CAS tools and was marginalised by the aesthetic regime of formalism. This dissertation supplies the complexity-theoretic scaffolding that renders his epistemic limits operational.

Conclusion

We have diagnosed the pretence of knowledge – the epistemic hubris that sustains mechanistic economics. Hayek’s warning remains structurally valid: economists confuse mathematical precision with genuine understanding, and this confusion protects the paradigm from recognising its own limits.

The core result is decisive. Where complex adaptive systems dominate, prediction-based control collapses. Navigation replaces fine-tuning; robustness replaces optimisation; epistemic humility replaces scientism.

With this subsection, Section 3.3 is complete. The methodological critique now gives way to reconstruction. Dissolving the pretence requires not better models within the mechanistic frame, but a different ontology and epistemology – precisely what Chapter 4 provides at the ontological level, with epistemological and methodological implications identified as necessary further development.

Conclusion of Section 3.3

Section 3.3 has traced how the paradigmatic lock-in diagnosed in Section 3.2 expresses itself in the methodological practice of the discipline. Five symptoms have been identified – physics envy, the optimisation fetish, DSGE and model monism, the aestheticisation of economics, and the pretence of knowledge – each revealed as a necessary expression of the foundational architecture rather than an independent correctable habit.

Three findings carry forward.

First, the symptoms are generated by the architecture, not chosen independently. Physics envy appropriates classical mechanics’ mathematical form – equations, optimisation, equilibrium, calculus – while neglecting that economic systems violate the domain conditions (decomposability, closure, stable laws) that make that form appropriate. The optimisation fetish follows by cascade logic from the mechanistic core: given determinism, certainty, perfect information, and ergodicity, calculus-based maximisation becomes the methodologically

"correct" expression of rationality. DSGE institutionalises both as the dominant macroeconomic methodology, encoding the architecture's assumptions in a form that eliminates precisely those properties – interaction, heterogeneity, emergence, non-linearity – that define complex adaptive systems. Mathematical elegance then operates as an epistemic filter that protects the paradigm from empirical refutation: after 2008, DSGE survived not because it explained the crisis but because it remained elegant. Beneath all four lies Hayek's pretence of knowledge – the epistemic hubris that confuses mathematical precision with genuine understanding and thereby insulates the paradigm from recognising its own limits.

Second, the symptoms cannot be reformed in isolation. Because they are generated by ontology and epistemology, changing what economics can represent requires changing what economics assumes exists and what it assumes can be known. Tinkering with technique while preserving foundations leaves the generative structure intact and triggers the regeneration pattern documented in Section 3.2.

Third, the symptoms produce a decisive epistemic result: the Prediction-Control Collapse. In complex adaptive systems, prediction-based control is not merely difficult but structurally incoherent. Navigation must replace fine-tuning; robustness must replace optimisation; epistemic humility must replace scientism.

Yet if these methodological symptoms are generated by the architecture, they should produce observable consequences – predictable failures traceable to specific assumptions. Section 3.4 now documents those consequences across three domains: economic theory, management practice, and systemic crisis. The question is no longer whether the paradigm fails, but whether its failures are contingent or generated.

3.4 Consequences

Introduction

Section 3.1 reconstructed what the mechanistic paradigm assumes: a fifteen-assumption architecture organised across ontological, epistemological, and derivative layers, whose most consequential property is the systematic elimination of the meso-level where firms operate. Section 3.2 explained why those assumptions persist despite their incompatibility with complex adaptive systems: three reinforcing feedback loops interlock at high-centrality nodes, reproduced through semantic, methodological, and institutional mechanisms that render the paradigm self-sealing. Section 3.3 traced how this lock-in expresses itself in methodological practice: physics envy, the optimisation fetish, DSGE and model monism, the aestheticisation of economics, and the pretence of knowledge – each a necessary expression of the foundational architecture.

The diagnosis now requires its final stage: observable consequences. A diagnosis that remained at the level of assumptions, mechanisms, and symptoms would be incomplete. It must be tested against the empirical record. If the assumptions are as described, and if they are locked in place as documented, and if they express themselves methodologically as traced, then they should produce predictable failures – not contingent errors attributable to data limitations or poor calibration, but structural failures traceable to specific assumptions. Section 3.4 tests this implication and thereby completes Chapter 3's diagnostic arc.

The argument proceeds through three subsections, each examining consequences at a different scale. Subsection 3.4.1 examines consequences in economic theory, showing that recurrent

predictive failures exhibit characteristic brittleness: narrow-domain success coupled with disproportionate breakdown under regime change. Subsection 3.4.2 traces the same failure structure into firm management, diagnosing the illusion of control as the organisational expression of mechanistic assumptions applied to adaptive systems – directly relevant to the dissertation's focus on reconceptualising the firm. Subsection 3.4.3 treats the 2008 financial crisis as a paradigmatic stress test in which multiple core assumptions failed simultaneously, their internal coherence becoming a liability under system-wide strain.

The question driving Section 3.4 is whether the failures are contingent or generated. If generated, the warrant for foundational reconstruction is established. Chapter 4 undertakes that reconstruction at the ontological level.

3.4.1 In Economic Theory: Predictive Failures and Brittleness

Introduction

Sections 3.1–3.3 established a diagnosis: the inherited body of economic theory rests on mechanistic ontological and epistemic commitments – equilibrium-centred closure, linearity, ergodicity, and representative-agent reductionism – stabilised through methodological and institutional lock-in. This subsection tests that diagnosis against the discipline's empirical record. The claim is not that economics predicts nothing, but that – conditional on these commitments – the dominant model class is structurally fragile: it performs comparatively well in stable regimes and performs worst precisely when guidance is most consequential (crises, regime shifts, structural breaks).

Within our framework, prediction failure is not treated as a contingent “data” or “parameterisation” problem. It is treated as a structural consequence of modelling an open, adaptive, non-linear, networked reality as if it were closed, stationary, and near-equilibrium. Lawson's critique is directly aligned with this framing: recurrent explanatory difficulty follows from “methods that carry inconsistent ontological presuppositions” relative to the “nature of social material” (Lawson, 2019, pp. 9–22; see also p. 15).

This subsection performs three functions. First, it shows that recurrent prediction breakdown matches the ontological and methodological diagnosis developed in Chapter 3. Second, it characterises the resulting failure mode as brittleness: narrow-domain success coupled with disproportionate breakdown under regime change. Third, it establishes why foundational reconstruction is not an optional refinement but a methodological necessity.

This is the first subsection of Section 3.4 (Consequences). It completes Chapter 3's diagnostic arc by moving from assumptions and lock-in (3.1–3.3) to observable failure modes, thereby providing the hinge from diagnosis to transformation.

I. The Track Record: Recurrent Failure Under Regime Change

Across major episodes, a recurring structure is visible: models calibrated to stability underperform under instability; the subsequent response is typically auxiliary augmentation (protective-belt expansion) rather than hard-core revision – consistent with the Lakatosian dynamics established in Section 3.2.4.

Failure 1: The Great Depression (1930s).

Equilibrium self-correction narratives – wage/price adjustment restoring full employment – proved inadequate to persistent mass unemployment and prolonged contraction. The Keynesian disruption was a substantive response, but later reabsorption into an equilibrium-centred synthesis illustrates an early instance of the pattern that becomes diagnostic in later crises: regime change exposes mismatch; repairs tend to preserve equilibrium architecture.

Meso insertion coordinate (I.1, after the paragraph above):

At the meso level, the episode exposes the same structural property diagnosed at the macro level: organisational planning devices premised on smooth adjustment and cyclical regularity become unreliable precisely when the environment undergoes a regime shift. This is not a claim about the quality of particular managers; it is a claim about epistemic architecture. Planning that presupposes stationarity and reversible time treats discontinuity as a transient deviation, thereby misclassifying structural change as noise.

Failure 2: Stagflation (1970s).

Frameworks that relied on a stable Phillips-curve trade-off were destabilised by joint inflation and unemployment. The subsequent shift to rational expectations redefined transmission mechanisms while preserving equilibrium and optimisation as organising principles – again illustrating that anomalies reorganise the protective belt more readily than the hard core.

Failure 3: The Global Financial Crisis (2008).

This episode exposed the limitations of a macro apparatus that represents the economy as stable with fluctuations driven by small disturbances. Kirman's formulation is crisp: the crisis was "both a crisis of the real economy and a crisis of the dominant economic theory" (Kirman, 2010, cited in Landini et al., 2025, p. 1). Krugman's contemporaneous critique identifies the aesthetic and methodological dimension: the profession "mistook beauty, clad in impressive-looking mathematics, for truth" (Krugman, 2009, p. 2). Caballero described a profession "mesmerized with its own internal logic," drifting into a "Fantasyland" in which quantitative implications became "surreal" (Caballero, 2010, pp. 85, 89). Arthur's diagnosis is structurally congruent: standard economics is organised around a "timeless, nonconstructive, equilibrium world," with "clearly defined optimization problems" and probability distributions that suppress the very dynamics that mattered (Arthur, 2014, p. 5; see also p. 191).

Failure 4: The Eurozone Crisis (2010–2015).

Frameworks emphasising convergence under common monetary conditions struggled to represent institutional heterogeneity, political-economy constraints, and path-dependent divergence. The case foregrounded multi-level coupling (bank–sovereign feedbacks, institutional asymmetry, heterogeneous adjustment capacity) that sits uneasily inside a single equilibrium template.

Failure 5: COVID-19 shock (2020).

The pandemic interacted with global supply networks, sectoral heterogeneity, non-linear constraints, and policy discontinuities. Standard templates incorporated the event as an exogenous shock term; the deeper issue was structural: a regime in which constraints, behaviours, and institutions co-evolve rapidly is not easily representable within stationary, near-equilibrium conventions.

The diagnostic pattern.

Across episodes, the same form repeats: (i) stability-calibrated performance in tranquil regimes; (ii) disproportionate breakdown under regime change; and (iii) post hoc augmentation that preserves the core. Kirman (2017, p. 37) notes that the recurrent arrival of major crises suggests “real problems with the analysis,” rather than isolated forecast misses.

II. Predictive Failure as a Structural Default Under Mechanistic Commitments

Within our framework, the inherited model class relies on a bundle of commitments that jointly enable tractability. Their methodological payoff is also their epistemic cost.

Table 84: Assumption Families in Mainstream Macroeconomic Modelling: Analytical Tractability Gained and CAS Dynamics Suppressed

Assumption family	What it enables	What it suppresses
Equilibrium-centred stability	Return-to-steady-state structure	Far-from-equilibrium dynamics
Ergodicity / stationarity	Inference from historical distributions	Structural breaks, endogenous novelty
Linearity / local approximation	Solvable dynamics	Amplification cascades, tipping
Representative agent / aggregation	Single optimisation problem	Heterogeneity, networks, distributional transmission
Closure	Complete variable specification	Endogenous creation of new state variables

The critical point is not that “models are wrong,” but that when these commitments become structural requirements, breakdown under crisis is predictable. Crisis regimes are precisely those in which far-from-equilibrium dynamics, shifting distributions, non-linear amplification, network topology, and endogenous novelty dominate – exactly the phenomena these commitments suppress. Lawson’s conclusion therefore bites with full force: persistent failure cannot be cured by additional technical sophistication when the dominant techniques presuppose an ontology the social domain does not satisfy (Lawson, 2019, pp. 9–22; p. 15).

III. Brittleness: Narrow-Domain Success and Wide-Domain Failure

Here **brittleness** denotes a specific failure mode: performance is acceptable within a narrow neighbourhood of the model's structural assumptions, and error grows nonlinearly as the regime departs from that neighbourhood.

Source 1: Equilibrium calibration.

When models are organised around equilibrium and local deviations, crisis dynamics become difficult to represent because the transition is not a small perturbation but a regime shift (Arthur, 2014, p. 5).

Meso insertion coordinate (III.1, after the equilibrium-calibration paragraph):

At the meso level, this translates into a systematic misclassification problem. When planning and valuation are built around a stable “normal” state, discontinuous transitions are treated as temporary departures rather than structural reorganisation. The analytic loss is not forecasting detail but causal diagnosis: the mechanisms of transition (feedback, tipping, network amplification) are treated as residuals.

Source 2: Ergodicity / stationarity.

If inference presumes that the future distribution is learned from the past, structural breaks are not merely hard – they are model-invalidating. Gallegati and Gallegati emphasise that the discipline persists in using analytical frameworks rooted in earlier centuries as if economies were “intricate mechanisms,” reproducing linear and reductionist bias precisely where innovation and institutional change render stationarity “pure fantasy” (Gallegati & Gallegati, 2025, p. 3065).

Source 3: Linearity / local approximation.

Linearisation can be defensible near steady state. But crises are dominated by non-linear interactions (thresholds, cascades, contagion), so the “small shock → small effect” intuition becomes unreliable in the very regimes that matter. Colander et al. describe 2008 in precisely these terms: a small sector “contaminated” the rest via hidden interdependencies (Colander et al., 2009, pp. 262–263).

Source 4: Closure.

Closed models assume the relevant state variables are known in advance. Yet economies repeatedly generate new variables endogenously (financial instruments, organisational forms, regulatory constraints, platform architectures). This is not missing data; it is genuine state-space expansion. Where closure is imposed, novelty cannot be represented by construction (Lawson, 2019, p. 15).

Meso insertion coordinate (III.4, after the closure paragraph):

At the firm level, closure fails not only because managers lack information, but because the competitive basis itself can change. Incumbents often model rivalry in an inherited basis (price/quality/store footprint), while competition migrates to a new basis (ecosystem lock-in, data advantage, algorithmic discovery, platform governance). Under state-space expansion, forecasting error is the symptom; ontological mis-specification is the cause.

Source 5: Homogeneity / representative-agent aggregation.

Crisis transmission is distributional: outcomes depend on who holds exposures, which nodes are central, and how heterogeneous agents respond. Representative-agent aggregation suppresses these pathways. Solow labelled DSGE “fantasy constructs” that fail the “smell test” and generate predictions “at odds with observable phenomena” (Solow, U.S. House, 2010, pp. 12, 14–15). Colander et al. similarly argue DSGE exhibits weak microfoundations and unsatisfactory empirical performance (Colander et al., 2009, p. 260). Caballero’s “Fantasyland” diagnosis captures the same dynamic: internal coherence substitutes for external adequacy (Caballero, 2010, pp. 85, 89).

Result: most wrong when most needed.

Brittleness yields a predictable asymmetry: models tend to perform acceptably in normal times (when predictive value is lower) and fail in crisis regimes (when predictive value is highest).

IV. Why Equilibrium Predicts Stability While Economies Generate Instability

The mechanistic worldview treats instability as exogenous disturbance against an intrinsically stable core. CAS reasoning reverses this orientation: instability is often endogenous – generated through feedback loops, adaptive responses, innovation, institutional evolution, and network topology. Arthur’s formulation captures the methodological consequence: an economy envisioned as a “well-functioning machine” achieves elegance and precision, but precisely through a restriction of admissible dynamics that is “inherently restrictive and unrealistic” (Arthur, 2021, p. 143; Arthur, 2014, p. 5). When crisis mechanisms are endogenous and non-linear, equilibrium-centred models systematically misclassify them as anomalies.

V. How Representative-Agent Modelling Removes Macro Transmission Mechanisms

Representative-agent macro suppresses the micro-to-macro channels that generate emergent macro patterns: heterogeneous constraints and balance sheets, interaction and coordination, network topology, and reflexive feedback between beliefs and outcomes. Hence, representative-agent models cannot generate crisis dynamics without reintroducing the heterogeneity and topology they removed. This is why, after crises, anomalies are recoded as “frictions” rather than treated as constitutive mechanisms (Colander et al., 2009, pp. 260–263; Solow, 2010, pp. 12, 14–15).

VI. Non-Ergodicity and Pretence Amplification

If the system is non-ergodic – if future distributions are not stable reflections of the past – then increasing statistical sophistication cannot solve the problem; it can intensify it by producing false precision. This is where Section 3.3.5’s “pretence of knowledge” mechanism amplifies brittleness: models become not only fragile but authoritative. Taleb’s warning is directly relevant: the pretence extends to believing we can quantify what we do not know (Taleb, 2007, p. 225). Farmer and Foley accept that some economic predictions work, while stressing that systemic crises – precisely where prediction matters most – emerge from complex interactions and remain difficult to forecast with equilibrium-centred tools (Farmer & Foley, 2009, p. 686). The practical consequence is brittle knowledge held with confidence.

VII. CAS Phenomena Versus Mechanistic Blind Spots

Table 85: From Mechanistic Suppression to Crisis Failure: How the Elimination of CAS Phenomena Produces Systematic Blind Spots in Crisis Prediction and Explanation

CAS phenomenon	Mechanistic suppression	Crisis implication
Emergence	Reductionism + representative agent	Macro patterns become residuals
Non-linearity	Linearisation	Cascades and tipping misrepresented
Path dependence	Ergodicity	Structural breaks appear as “surprises”
Fundamental uncertainty	Risk substitution	Tail exposure underweighted
Network topology	Aggregation/atomism	Contagion pathways disappear

Key diagnosis: predictive breakdown follows from treating an open adaptive system as a closed mechanical one (Lawson, 2019, pp. 9–22; Gallegati & Gallegati, 2025, p. 3065).

VIII. Counter-Arguments Addressed

“Models aren’t meant to predict crises.”

If a model class is systematically least reliable when stakes are highest, the limitation is substantive, not merely scope-related. Moreover, recurrent crises are not negligible tail cases; they are central to macroeconomic welfare and stability. Treating crises as exceptional “tails” is itself a modelling choice, and one that can conceal regular systemic fragilities (Taleb, 2007, p. 225; Kirman, 2017, p. 37).

“It’s a data problem.”

More data can refine measurement, but it cannot restore ergodicity if the system’s statistical structure changes endogenously, nor can it pre-specify state variables created by innovation. Lawson’s point remains decisive: the obstacle is ontological mismatch, not information volume (Lawson, 2019, p. 15).

“Post-crisis improvements show learning.”

Additions (frictions, heterogeneity variants, financial sectors) may improve local fit, but when equilibrium-centred architecture, optimisation commitments, and aggregation structure remain intact, these remain protective-belt expansions rather than paradigm transformation (Colander et al., 2009, pp. 260–263; Caballero, 2010, pp. 85, 89).

“Some predictions work.”

Yes – in stable regimes and for relatively stationary relationships. The critique concerns domain of validity: narrow success does not vindicate a framework that collapses under regime change. Farmer and Foley’s argument captures this: partial predictive success does not extend to crisis regimes shaped by interaction-driven emergence (Farmer & Foley, 2009, p. 686).

Conclusion: Brittleness as Empirical Confirmation of Ontological Mismatch

This subsection used recurrent failure modes as an empirical test of the Chapter 3 diagnosis. The century-long pattern – acceptable performance near stability, breakdown under crisis, and post hoc augmentation without hard-core revision – is consistent with the Lakatosian degeneracy verdict (Section 3.2.4) and with the epistemic limits emphasised in Hayek's critique of false precision and overconfidence (Section 3.3.5).

The core result is not primarily technical; it is ontological. A mechanistic paradigm applied to CAS reality yields brittle knowledge – precise within a narrow regime, unreliable under structural change (Lawson, 2019, pp. 9–22; Arthur, 2014, p. 5; Gallegati & Gallegati, 2025, p. 3065).

The transformation implication follows: foundational reconstruction is warranted not as stylistic preference but as methodological necessity. Chapter 4 replaces substance/atomism with process/relational ontology, establishing the ontological foundations for a CAS-compatible understanding of firms. The epistemological implications (fundamental uncertainty, non-ergodicity) and methodological implications (bounded rationality, pluralistic methods) are identified as necessary further development beyond the scope of this thesis, but are logically entailed by the ontological reconstruction.

Forward connections, p. 3.4.2 translates theoretical brittleness into organisational/managerial brittleness (illusion of control); 3.4.3 consolidates 2008 as the paradigmatic case in which equilibrium, linearity, ergodicity, closure, and aggregation fail simultaneously; Chapter 4 takes this empirical warrant as the basis for ontological reconstruction.

3.4.2 In Firm Management: The Illusion of Control

Introduction

Section 3.4.1 established that predictive failure in economic theory is not accidental but structurally induced by mechanistic assumptions (equilibrium-centred closure, linearity, ergodicity, and aggregation). The same CAS reality that defeats those models is encountered daily by managers. This subsection diagnoses the managerial analogue of predictive brittleness: the illusion of control.

Within our framework, the illusion of control is not a psychological quirk. It is a paradigm-induced misperception: the inherited mechanistic worldview converts organisations into objects of optimisation and targets of command, generating expectations of predictability, proportionality, and controllability. When firms behave as complex adaptive systems – through emergence, adaptation, network coupling, and path dependence – those expectations collapse. The managerial problem is therefore not primarily competence; it is causal ontology mismatch.

This subsection performs three functions. First, it shows how theoretical assumptions migrate into managerial beliefs and toolkits – how the paradigm operates not only in journals but in boardrooms. Second, it reframes recurrent managerial failure as paradigm-induced brittleness rather than individual managerial deficiency. Third, it closes Chapter 3’s diagnostic arc by demonstrating that CAS transformation is practically necessary for effective management, not merely theoretically desirable.

This is the second subsection of Section 3.4 (Consequences) and the penultimate component of Chapter 3. It prepares Section 3.4.3 (2008 as paradigmatic failure) by showing how epistemic hubris (3.3.5) becomes organisational overconfidence and control-centred practice.

I. The Mechanistic Import: From Economics to Management

Within our analytical framework, modern management inherits core assumptions from economics, which itself inherited the mechanistic aesthetic and epistemology of classical physics. This lineage is institutional and methodological: physics supplies the prestige of control; economics supplies optimisation as universal method; management operationalises both as targets, KPIs, and command systems.

Ackoff (1999, pp. 1, 4) traces this trajectory in systems terms: industrial-era thought privileged rational order, predictability, and controllability, creating a receptive environment for Taylor’s scientific management (Taylor, 1911, p. 9). Taylor’s program embodies the mechanistic worldview by decomposing work into repeatable elements designed to emulate machine reliability. Cohen (1994, p. 78) similarly identifies Newtonian rational order as an intellectual template shaping modern thought, channeling economics toward quantification and efficiency, and management toward systematic control.

Table 3.4.2-1. Imported assumptions and their managerial translation

Table 86: 3.4.2-1. Imported assumptions and their managerial translation

Imported assumption	Managerial implication	Chapter anchor	3
Optimisation	“Best” choices are calculable and should govern decision	3.3.2	
Control	Outcomes are produced by top-down command	3.1.3 / 3.3.5	

Imported assumption	Managerial implication	Chapter anchor
Predictability	Inputs determine outputs in stable, learnable ways	3.1.2
Efficiency	Slack is waste; buffers should be removed	3.3.2
Equilibrium	Stability is natural; disturbance is transient	3.1.3

Xerox PARC demonstrates how mechanistic evaluation criteria can be systematically blind to emergent value. Bottom-up exploration generated innovations whose commercial meaning depended on ecosystem formation, complementary assets, and path-dependent market construction. Headquarters evaluated outputs through predictability filters – ROI timetables, product-category fit, forecastable demand – implicitly presuming closure and ergodicity. The resulting error was not miscalculation but ontological: when novelty is emergent and market structure co-evolves with innovation, “predictable ROI ex ante” is an incoherent criterion. Control logic therefore selects against precisely the processes that generate strategic discontinuity.

II. Why the Illusion Arises: Mechanistic Assumptions versus CAS Reality

Mechanistic assumptions generate a predictable belief-structure about management: stability is the baseline, interventions are proportional, optimising parts improves the whole, historical metrics forecast the future, and subsystems can be managed in isolation. These beliefs are not merely mistaken; they are what follows when organisations are treated as decomposable machines.

Table 3.4.2-2. Assumption-to-belief chain and its managerial consequence

Table 87: 3.4.2-2. Assumption-to-belief chain and its managerial consequence

Mechanistic assumption	Managerial belief	Practical implication
Equilibrium	Stability is natural	Disruption is interpreted as transient noise
Linearity	Effects are proportional	Interventions are designed as “dials”
Additivity	Optimise parts → optimise whole	Silo optimisation becomes rational policy
Ergodicity	Past distributions forecast future	KPIs become forecasting instruments
Decomposability	Subsystems can be managed independently	Coordination is replaced by targets

CAS theory reverses these expectations. Stability is not “given” but continuously constructed; small perturbations can cascade; local optimisation can damage system performance; novelty and structural breaks are constitutive; interdependencies are causal rather than residual. Poli’s formulation captures the implication: complex systems “cannot be controlled, only influenced” (Poli, 2013, p. 142). The mismatch is therefore ontological: managers presume controllable machines but encounter adaptive systems with emergent properties.

This section connects directly to 3.3.5. Hayek’s critique of scientism and false precision implies that measurement and optimisation can generate an illusion of mastery. Precision and control claims are often strongest where genuine knowledge is most constrained (Hayek, 1989, p. 6).

III. Why Organisations Are CAS, Not Machines

Within our framework, organisations exhibit CAS properties that systematically defeat machine logic.

Table 3.4.2-3. CAS properties of organisations and the managerial implication

Table 88: 3.4.2-3. CAS properties of organisations and the managerial implication

CAS feature	Why it defeats machine logic
Emergence	Capabilities and culture arise from interaction, not design
Adaptation	Agents learn and revise rules over time
Networks	Informal ties can dominate formal hierarchy
Path dependence	History constrains the option set

Optimisation requires stable states, known constraints, and a well-defined objective function. Under fundamental uncertainty (Chapter 5), those conditions fail. Simon’s formulation supplies the appropriate substitute: decision makers satisfice rather than optimise under bounded rationality (Simon, 1979, p. 502). In our framework, bounded rationality is not error; it is a realism condition of agency in complex environments.

Mechanistic control can also create brittleness. Tight coupling, optimisation lock-in, standardisation, and centralisation can improve measured efficiency near an assumed regime while compressing exploration, slack, and diversity – reducing adaptability precisely when the regime shifts. This is the managerial analogue of 3.4.1’s brittleness: local performance gains purchased through systemic fragility.

GE functions here as a mechanism-level illustration rather than a monocausal narrative. Intensive variance-reduction and metric-governed control can raise reliability in stable regimes, but the same architecture can compress exploratory search, reduce slack, and standardise cognition. Under discontinuous change, those “efficiency gains” convert into brittleness because the organisation has optimised for a world that no longer exists. The causal claim is general: when optimisation becomes governing ontology, adaptability becomes residual – visible only after the environment shifts.

IV. Practice Failure Modes: When the Illusion Meets Reality

The illusion of control becomes observable as recurring failure modes. The common structure is not that managers “make mistakes,” but that mechanistic tools presuppose stability, closure, and proportional causation in environments that are non-ergodic, networked, and novelty-generating.

Failure mode 1: Forecast-centred strategic planning.

Mechanistic planning presumes a sufficiently predictable environment: define objectives, forecast, plan, execute, control. CAS environments exhibit endogenous novelty and discontinuous shifts. Arthur’s point is decisive: agents’ exploration changes the environment they are exploring, which undermines the premise that planning can “solve” the future ex ante (Arthur, 2021, p. 137).

Pan Am demonstrates the epistemic failure mode of forecast-centred optimisation under structural breaks. Deregulation, oil shocks, geopolitical disruption, and security risk did not merely perturb parameters; they changed the system’s competitive structure and constraint set. Under regime change, forecast errors are not estimation errors; they are model invalidation.

The managerial conclusion generalises: where environments are non-ergodic and institutions reshape the option space, the planning problem is not “forecast better,” but “navigate a shifting regime.”

Failure mode 2: Metric optimisation and Goodhart dynamics.

Mechanistic management treats performance metrics as controllable levers. Yet when a measure becomes a target, the organisation adapts to the measure rather than the underlying objective (Goodhart, 1975; often cited via Strathern, 1997, p. 308). The result is not incidental gaming but endogenous distortion: local optimisation undermines global performance.

Wells Fargo illustrates how proxy control fails in adaptive social systems. A single numerical proxy for “relationship depth” became the governing target. Once elevated to a control variable, the organisation’s agents explored feasible ways to satisfy the metric, including pathological ones. The mechanism is CAS-consistent: targets alter behaviour; altered behaviour breaks the indicator’s meaning; the system produces outcomes that satisfy the metric while destroying the phenomenon the metric was meant to represent. The analytic point is not scandal description; it is that metric governance imports linear controllability into a reflexive environment, generating predictable failure.

Failure mode 3: Reengineering and the destruction of distributed cognition.

Reengineering operationalised mechanistic redesign: map, optimise, eliminate redundancy, standardise. A recurring failure mode was the removal of buffers, discretion, and informal coordination – the very structures through which learning and resilience operate. Hammer later acknowledged that such programs underestimated the human and adaptive dimension of organisational life.

Circuit City illustrates why “labour as cost input” is an ontological error in learning systems. Experienced retail staff embody tacit knowledge – product interpretation, customer interaction, local problem-solving – that cannot be separated from the agents who carry it. When optimisation treats that competence as a substitutable input, cost reduction can simultaneously destroy the firm’s distributed cognitive system. The general mechanism is that removing “redundant” experience compresses adaptive capacity; performance then collapses under competitive and technological turbulence.

Failure mode 4: Risk management and crisis dynamics inside tightly coupled networks.

Pre-crisis risk systems commonly presupposed stationarity, stable correlations, and additive risks – conditions that fail under network contagion and feedback amplification (as argued in 3.4.1; see also Colander et al., 2009, pp. 262–263). The same mismatch appears within firms: models optimise within an assumed world, then fail when the world changes regime.

Northern Rock illustrates the category error of treating liquidity as an internal property rather than a network property. Quantitative risk tools can appear adequate under stable funding conditions because they presuppose continuous market access and ergodic inference. When interbank trust shifts, the system transitions phase: funding availability becomes discontinuous, and the model’s world ceases to exist. The analytic conclusion is general: in tightly coupled financial networks, dominant risks are relational and structural rather than probabilistic; false precision therefore produces false control.

V. Why Optimisation Becomes Dangerous Under Complexity

Within our framework, conventional management inherits the optimisation imperative from the economics of calculability (3.3.2): minimise slack, eliminate redundancy, operate at the

frontier. CAS performance, by contrast, depends on buffers for exploration, redundancy for resilience, and modularity for containment. Nason's distinction is useful here: complicated systems support reproducible cause–effect control, whereas complex systems are adaptive and unpredictable (Nason, 2017, p. 6). Treating complex as merely complicated leads to “successful” optimisation that manufactures fragility.

VI. Why Top-Down Control Suppresses Bottom-Up Adaptation

Mechanistic logic assumes control flows from the top. CAS logic locates order in distributed interaction and self-organisation. Kauffman's statement is explicit: complex adaptive systems have no central controller; order is self-organised (Kauffman, 1995, p. 71). Holland similarly characterises adaptive agents as rule-updating entities, not static optimisers (Holland, 1995, p. 10). The managerial implication is precise: top-down control can suppress the very local intelligence required for survival under uncertainty.

IBM under Gerstner illustrates the shift from command-and-control to adaptive orchestration. The recovery logic is not “stronger control,” but enabling constraints: reducing silo decomposition, increasing cross-unit coordination around customer problems, and privileging sensing-and-responding over forecast-and-execute. The general mechanism is that, under non-ergodic regime change, viability depends on reconfigurability; management therefore becomes the design of conditions for emergence rather than the specification of outcomes.

VII. Managers with the Wrong Causal Ontology

Mechanistic management presumes full information, stable causal relations, fixed objectives, and controllable environments. CAS reality implies dispersed and incomplete information, contextual and non-linear causation, endogenous goal evolution, co-evolving environments, and interdependent tasks. White's verdict captures the epistemic point in management language: practice premised on knowable and controllable structure is “simply wrong” (White, 2018, p. 1). In our framework, this is not an insult to managers; it is a diagnosis of paradigm mismatch.

VIII. Counter-Arguments Addressed

The claim that “good managers control their firms” typically conflates outcome control with conditions-of-adaptation control. Toyota's advantage is not rigid command but disciplined, distributed learning; Spear and Bowen describe Toyota operations as structured experimentation and learning-by-doing (Spear & Bowen, 1999, p. 97). Similarly, in turbulent contexts, Eisenhardt and Sull show that firms succeed with simple rules that enable adaptation rather than detailed predictive plans (Eisenhardt & Sull, 2001, p. 108). The critique therefore targets mechanistic planning as single-future optimisation, not planning per se. Mintzberg's distinction is relevant: “strategic planning” often substitutes analysis for strategy formation, whereas strategy in complex environments emerges through learning and pattern recognition (Mintzberg, 1994, p. 107). Computational power does not resolve ontological openness; Taleb's point is that sophisticated models can amplify the illusion of understanding and increase vulnerability (Taleb, 2012, pp. 7–8). Finally, the objection that “firms must optimise to survive” confuses short-run efficiency with long-run viability under regime change. Levinthal formalises the exploration–exploitation trade-off: over-emphasis on exploitation undermines exploration, reducing adaptability (Levinthal, 1997, p. 937). The argument is therefore not anti-efficiency but anti-ontological error: optimisation becomes dangerous when it is treated as governing worldview rather than context-bound instrument.

IX. Preview: CAS-Compatible Management

Chapter 4 develops the ontological alternatives; the full elaboration of CAS-compatible management principles is beyond this thesis's scope but follows from the foundational reconstruction. For Chapter 3, it is sufficient to state the mapping.

Table 3.4.2-4. Mechanistic management principles and CAS alternatives

Table 89: 3.4.2-4. Mechanistic management principles and CAS alternatives

Mechanistic principle	CAS alternative
Optimisation	Adaptation, learning, satisficing
Efficiency as frontier	Slack and redundancy as resilience
Central planning	Simple rules and enabling constraints
Standardisation	Diversity, modularity
Control	Influence, navigation, robustness

The core implication follows: managers do not “lack control” as a personal deficit; systems lack controllability as an ontological condition.

Conclusion

This subsection diagnosed the illusion of control as the managerial expression of the mechanistic paradigm. The same assumptions that yield brittle prediction in theory (3.4.1) yield brittle governance in practice: forecast-centred planning failures under regime change, Goodhart pathologies under metric governance, capability destruction under reengineering logic, and risk-model blindness under networked contagion. The diagnosis is paradigm-level: managers apply machine logic to adaptive systems.

This prepares Section 3.4.3 by showing how epistemic overconfidence (3.3.5) becomes organisational overreach, creating the practical preconditions for system-wide collapse when networked fragilities are stressed.

3.4.3 The 2008 Financial Crisis as Paradigmatic Failure

Introduction

Sections 3.4.1 and 3.4.2 documented predictive failures in economic theory and control failures in managerial practice. The 2008 Global Financial Crisis offers a uniquely concentrated test of whether these failures were contingent or structural. Treated here as a paradigmatic episode, the crisis functions as a system-wide stress test of the mechanistic paradigm: multiple core assumptions failed *together*, and their failure modes were mutually reinforcing.

Within our framework, this subsection performs three diagnostic functions. First, it shows how distinct mechanistic assumptions converged to produce systemic blindness, yielding not one error but a coupled set of failures. Second, it demonstrates that the crisis outcome was structurally produced rather than accidental: machine logic was applied to an adaptive, interaction-dominated system. Third, it grounds the chapter's diagnosis in observable consequences, thereby supplying the empirical warrant for assumption-level reconstruction in Chapter 4.

I. How Each Assumption Failed: A Systemic Collapse

Matrix 1. Ontological assumptions and crisis failure modes

Table 90: Matrix 1. Ontological assumptions and crisis failure modes

Ontological assumption	Mechanistic claim	Crisis manifestation	Empirical illustration
Reductionism	System behaviour equals the sum of parts	Systemic risk emerged from interaction structure	Lehman: component “health,” system collapse
Atomism	Institutions behave as if independent	Network contagion dominated outcomes	Prime-broker and funding-network contagion
Closure	System boundaries are stable and specifiable	Shadow banking and cross-border spillovers dominated	Off-balance-sheet and cross-market coupling
Substance (implicit)	Stable entities, not processes, are primary	Fragility evolved endogenously through balance-sheet dynamics	Liquidity evaporation and collateral spirals
Equilibrium (ontological)	Systems self-stabilise after shocks	Positive feedback amplified shocks	Fire-sale and margin-spiral dynamics

The common implication is ontological: the financial system behaved as an open, interaction-dominated, emergent process. The mechanistic ontology suppressed precisely those properties that proved decisive under stress.

Matrix 2. Epistemological and methodological assumptions and predictive collapse

Table 91: Matrix 2. Epistemological and methodological assumptions and predictive collapse

Epistemic methodological assumption	Mechanistic commitment	Crisis failure mode	Evidence / signature
Determinism (probabilistic form)	Risk is probabilistically knowable	Correlations and mappings broke	“Sigma-event” narratives; model invalidation
Certainty (in-principle)	Relevant risks are knowable	Widespread surprise and disbelief	Pre-crisis “contained” assessments
Ergodicity / stationarity	Past distributions reveal future	Structural breaks dominated	Historical calibration became misleading
Linearity	Small shocks imply small effects	Cascades and discontinuities	Subprime → global systemic crisis
Homogeneity / representativeness	Aggregates can be treated as typical	Distribution and position mattered	Funding and leverage asymmetries
Optimisation (as behavioural core)	Agents maximise over stable problems	Runs, herding, and panic	Liquidity runs and synchronised deleveraging

The unifying epistemic defect is the substitution of fundamental uncertainty with calculable risk: prediction failed not because models were insufficiently sophisticated, but because the system violated the conditions prediction presupposed.

II. Why Simultaneous Failure Was Not Coincidence

The simultaneity is diagnostic, not incidental. The mechanistic paradigm is an interconnected architecture: ontological commitments (reductionism, atomism, closure) support epistemological commitments (certainty, ergodicity), which in turn stabilise methodological practice (equilibrium-centred modelling, optimisation, representative constructions). Under system-wide stress, weaknesses embedded in this architecture are exposed *at once*. The crisis therefore manifested as a single paradigmatic failure across multiple dimensions rather than as a sequence of separable mistakes.

Firm-level convergence: Washington Mutual (WaMu)

Washington Mutual illustrates the convergence mechanism at the firm level. Under prevailing frameworks, WaMu’s condition was assessed primarily through component metrics – capital ratios, risk weights, portfolio summaries – consistent with reductionist decomposition and probabilistic risk calculus. Yet the crisis mechanism was interactional: when confidence broke, a modern run unfolded as depositors withdrew roughly USD 16.7 billion in a short span, and the institution was seized on 25 September 2008. The event is paradigmatic because it shows how “soundness” under mechanistic representation can coexist with fragility generated by network dependence, confidence dynamics, and phase-like shifts in liquidity.

III. Linear Modelling in a Non-Linear System

Table 92: Risk-modelling presuppositions versus crisis dynamics

What dominant risk models presupposed	What the financial system exhibited
Stable correlation structures	Correlation regime shifts under stress
Additive, decomposable exposures	Interaction-driven amplification and feedback
Thin-tail behaviour and well-behaved extremes	Fat tails and clustered extremes
Historical calibration as reliable guide	Non-stationarity and structural breaks
Liquidity as continuous and available	Discontinuous liquidity freezes (“phase shifts”)

The methodological point is not that models “missed a big shock,” but that the core modelling stance embedded linear intuitions into a system dominated by non-linear propagation, topology, and feedback.

IV. Representative Constructions and Systemic Ignorance

Representative constructions (representative bank, representative tranche, representative balance sheet) suppressed the heterogeneity and topology that determined transmission: who was leveraged, who was short of funding, who was a central counterparty, which nodes were irreplaceable, and which strategies were correlated. This is why post hoc explanations repeatedly had to “reintroduce” what baseline modelling had removed: distribution, network position, correlated behaviour, and endogenous amplification.

Firm-level illustration: Merrill Lynch

Merrill Lynch exemplifies how representativeness substitutes away distributional structure. Portfolio risk was treated through representative securities with stable parameters, while losses were driven by concentration in the most fragile slices of fragile vintages and by correlation shifts under stress. The point is structural: representativeness masks the causal relevance of *composition* and *dependence*, which become decisive precisely in crisis regimes.

V. Financial Innovation as Endogenous Restructuring

In mechanistic framing, innovation is typically treated as an exogenous efficiency/risk-allocation improvement applied to a stable system. The crisis contradicts this: instruments such as CDOs and CDSs did not merely redistribute risk – they reconfigured interdependence and opacity, altering network topology and transmission channels. Securitisation linked previously weakly connected balance sheets, expanded leverage pathways, and amplified the system’s susceptibility to confidence-driven discontinuities.

VI. Institutions as Co-evolving Constraints

Regulation and institutional architecture were frequently treated as stable constraints acting “from outside” the system. The crisis revealed a different structure: institutions and market strategies co-evolved through adaptation, arbitrage, and rule redesign, producing procyclical reinforcement rather than neutral constraint. The key diagnostic result is that static constraint models were structurally unfit for a system in which rules, incentives, and instruments adapt together.

VII. Post-crisis Adjustment: Protective Belt Expansion Without Core Transformation

The post-crisis response largely took the form of auxiliary augmentation – financial frictions, limited heterogeneity variants, stress testing regimes, and piecemeal network add-ons – while preserving the organising core: equilibrium-centred structure, optimisation as baseline behaviour, and representative constructions as the tractability device. This pattern fits the chapter’s lock-in diagnosis: anomalies were absorbed by expanding the protective belt rather than revising the hard core.

VIII. Crisis as Collapse of the Prediction–Control Regime

The mechanistic paradigm links prediction to control: if the system is forecastable, it is steerable. The crisis demonstrated the fragility of this linkage. When uncertainty is not reducible to stable probability structures – because interaction, reflexive behaviour, and regime shifts dominate – prediction fails in the same regimes where control is most demanded. Regulators and managers were not merely under-informed; they were systematically filtered by a modelling stance that excluded topology, feedback, and endogenous fragility.

Table 93: Filtering mechanism: what was treated as “secondary” became decisive

Treated as secondary under mechanistic filters	Became decisive under stress
Network structure	Contagion pathways and hub fragility
Composition/distribution	Concentration, leverage, and funding asymmetry
Non-linearity	Fire-sale spirals, margin cascades
Endogenous instability	Self-generated fragility during “stability”

IX. 2008 as Paradigm-Level Falsification

The events of 2008 did not merely disconfirm particular implementations (a specific DSGE variant, a specific VaR model). They falsified the mechanistic *ontological package* that rendered those implementations natural: decomposability, closure, equilibrium-centred stability, and probabilistic knowability. In the terms developed earlier in the chapter, what failed was not simply a model but a metaphysical stance that systematically suppressed emergence, interaction, and endogenous instability.

X. Complexity-Oriented Warnings Before 2008

The crisis did not arrive without conceptual antecedents. A non-mechanistic literature had long argued that stability can be self-undermining, tails are fat, and interdependence creates endogenous fragility. The issue was not absence of warning but inadmissibility under dominant methodological filters.

Table 94: Pre-2008 warnings: mechanism-level anticipation (not date-level forecasting)

Approach / author	Mechanism foregrounded	What it implied about crisis possibility
Financial instability tradition (e.g., “stability breeds instability”)	Endogenous leverage and fragility accumulation	Crises are system-generated, not rare exogenous shocks

Approach / author	Mechanism foregrounded	What it implied about crisis possibility
Fat-tail tradition (e.g., non-Gaussian returns)	Extreme events more frequent than thin-tail models allow	“Sigma-event” narratives reflect model error, not reality
Network / interdependence analyses	Topology shapes contagion and cascades	Systemic risk depends on connectivity and hub position

XI. Counter-arguments addressed

Table 95: Objections and structural replies

Objection	Structural reply
“It was just regulatory failure.”	Regulation failed partly because it relied on mechanistic risk representations that mismeasured systemic risk; enforcement cannot substitute for an ontology that suppresses endogenous fragility.
“It was a black swan.”	The timing was not forecastable; the mechanism class was. Treating the event as inherently unforeseeable often restates the consequences of modelling blind spots rather than refuting them.
“Economics improved afterwards.”	Many changes are additive and preserve the organising core; this is consistent with protective-belt expansion rather than core revision.
“CAS couldn’t predict it either.”	The relevant criterion is mechanism adequacy under regime change, not point forecasting; the mechanistic paradigm failed by excluding the very dynamics that produced collapse.

Conclusion: the stress test completed

The 2008 crisis is best read here as a paradigmatic stress test: multiple mechanistic assumptions failed simultaneously, and their internal coherence became a liability under system-wide strain. What collapsed was not merely a financial architecture but the adequacy of a mechanistic ontology for an adaptive, interaction-dominated economic reality. This completes Chapter 3’s diagnostic arc and supplies the empirical warrant for the foundational reconstruction developed in Chapter 4.

The deeper source is not technical failure but methodological borrowing.

Aha Closure: Why the Mechanistic Paradigm Persists – and Why It Fails

Table 96: 3.4.3-A. The four appropriations (physics → economics) as failure mechanisms

#	Appropriation	What it imports	Why it fails for economics
1	Equation-based thinking	Differential-equation form	Institutions, expectations, and novelty resist equation reduction because the relevant state variables and causal structure are not stable under reflexive adaptation.
2	Optimisation as universal	Least-action metaphor	Fundamental uncertainty and endogenous state-space expansion invalidate stable objective functions and well-defined feasible sets as universal premises.

#	Appropriation	What it imports	Why it fails for economics
3	Equilibrium as natural	Thermodynamic endpoint	Economies are open, adaptive, and historically contingent; far-from-equilibrium regimes are not anomalies but recurrent systemic states.
4	Calculus as core tool	Continuous-change mathematics	Discontinuities, thresholds, network cascades, and phase-like shifts dominate crisis dynamics; these are not marginal complications but structurally central.

The crisis can therefore be read as a single causal sequence in which a borrowed ontology and its required methods jointly filtered out the decisive mechanisms: substance ontology → determinism → closure → equilibrium → formalism → “physics envy.” The result is borrowed prestige: economics practices “physics-without-physics” – the form without the domain discipline that makes the form valid.

Table 97: 3.4.3-B. The cognitive shift required by the 2008 stress test

From	To
“Economics uses math because it is rigorous.”	“Mathematical formalism is the methodological face of a failing paradigm when its ontological preconditions do not hold.”

Conclusion of Section 3.4

Section 3.4 has tested the Chapter 3 diagnosis against observable consequences across three domains. The findings confirm that the failures of mechanistic economics are not contingent errors but structural outcomes generated by the assumption architecture identified in Section 3.1, stabilised by the mechanisms documented in Section 3.2, and expressed through the methodological symptoms traced in Section 3.3.

Three findings complete the diagnostic arc.

First, in economic theory, the dominant model class exhibits characteristic brittleness. Across a century of episodes – from the Great Depression through stagflation, the 2008 crisis, the Eurozone crisis, and COVID-19 – the same pattern recurs: stability-calibrated performance in tranquil regimes, disproportionate breakdown under regime change, and post hoc augmentation that expands the protective belt while preserving the core. The failures are traceable to specific assumptions: equilibrium calibration cannot represent regime shifts; ergodicity presupposes stable distributions that structural breaks invalidate; linearity suppresses the non-linear amplification that dominates crises; closure cannot accommodate endogenous state-space expansion; and representative-agent aggregation removes the heterogeneity and network topology through which contagion propagates. The result is knowledge that is most wrong when most needed.

Second, in firm management, the same assumptions produce the illusion of control – directly relevant to the dissertation's focus on reconceptualising the firm. Mechanistic logic converts organisations into objects of optimisation and targets of command, generating expectations of predictability, proportionality, and controllability that collapse when firms behave as complex adaptive systems. The failure modes are predictable and recurrent: forecast-centred planning fails under regime change; metric governance triggers Goodhart pathologies; reengineering destroys distributed cognition; and risk models blind managers to networked fragility. The core finding is ontological: systems do not lack controllability as a managerial deficiency; they lack it as a condition of their nature.

Third, the 2008 crisis functions as a paradigmatic stress test. Multiple assumptions failed simultaneously – ontological, epistemological, and methodological – and their interconnection, which normally provides coherence, became a liability under system-wide strain. What collapsed was not merely a financial architecture but the adequacy of a mechanistic ontology for an adaptive, interaction-dominated reality.

Chapter 3's diagnostic arc is now complete. The mechanistic paradigm has been reconstructed (3.1), its persistence explained (3.2), its methodological symptoms traced (3.3), and its consequences documented (3.4). The failures are generated, not contingent. Foundational reconstruction is therefore not an optional refinement but a methodological necessity. Chapter 4 undertakes that reconstruction at the ontological level – replacing substance ontology with process ontology, atomism with relationality, and reductionism with emergence – to establish foundations adequate to firms understood as complex adaptive systems.

Conclusion

This chapter set out to demonstrate a precise claim: that the recurring failures of mainstream economic theory under conditions of complexity are not contingent shortfalls but structural consequences of the paradigm's foundational architecture. The diagnostic argument has proceeded through four interdependent stages, each building on and requiring the others.

The first stage reconstructed what the paradigm assumes. Section 3.1 identified a 15-assumption network organised across ontological, epistemological, and derivative layers. Five ontological assumptions – reductionism, determinism, substance ontology, closure, and atomism – generate four epistemological assumptions through a closed loop in which ontology conditions epistemology, epistemology conditions methodology, and methodology, evaluated by paradigm-internal criteria, feeds back to confirm ontology. From these, six derivative assumptions – equilibrium, linearity, homogeneity, ergodicity, optimisation, and decomposability – follow not as optional modelling choices but as logical entailments. The architecture's most consequential effect is the systematic elimination of the meso-level: the joint operation of reductionism, atomism, and representative-agent aggregation leaves no analytical space for firms, networks, or institutions to appear as constitutive. This is not a gap to be filled but an outcome generated by design.

Yet the identification of incompatible assumptions raises a puzzle that the diagnosis must answer: if these commitments are demonstrably untenable under complexity, why do they persist? Section 3.2 resolved this puzzle by showing that the assumptions are not merely consistent but mutually reinforcing. Three feedback loops – ontological, epistemic, and structural – interlock at high-centrality nodes, creating redundant stabilisation. Challenge to any single assumption activates compensatory defence across the others; the paradigm behaves as a closed attractor that absorbs perturbations and returns to its stable configuration. This lock-in is reproduced through semantic mechanisms that embed assumptions in disciplinary vocabulary, methodological mechanisms that equate tractability with rigour, and institutional mechanisms that select for conformity across training, publication, and funding. Evaluated against Lakatosian criteria, the inherited programme exhibits the characteristic signatures of degeneration: auxiliary proliferation without predictive advance, novelty-filtering, and defensive repair rather than progressive problem shifts. The persistence of a failing paradigm is thereby explained not by ignorance but by structure.

The structural diagnosis, however, requires expression in observable practice. Section 3.3 traced how the lock-in manifests as methodological symptoms. Physics envy imported the mathematical form of classical mechanics without validating whether economic phenomena

satisfy its domain conditions – the borrowed-prestige problem. The optimisation fetish elevated calculus-based maximisation from context-bound tool to universal decision principle. DSGE institutionalised these commitments as the dominant macroeconomic methodology, while model monism excluded alternatives by definitional fiat. The aestheticisation of economics subordinated truth to elegance, making mathematical beauty the criterion of legitimacy and thereby explaining why 2008 did not trigger paradigm change: the framework survived because it remained elegant, not because it remained adequate. Hayek's diagnosis of the pretence of knowledge emerged as the unifying epistemic foundation – economists systematically overestimate what can be known about complex adaptive systems, and this hubris functions as a paradigm-protecting device that reclassifies structural failure as technical shortfall. The symptoms are not incidental stylistic choices; they are the necessary expressions of the foundational commitments identified in Section 3.1, stabilised by the mechanisms identified in Section 3.2.

The diagnostic arc, however, remains incomplete without empirical warrant. If failures are truly generated by architecture, then predictable patterns of breakdown should be observable where complexity is constitutive. Section 3.4 supplied this warrant across three domains. Predictive brittleness in economic theory – stable performance near equilibrium, catastrophic failure under regime change – reflects structural mismatch between closed-system methodology and open, adaptive reality. The illusion of control in management – forecast-centred planning failures under regime change, Goodhart pathologies under metric governance, capability destruction under reengineering logic, risk-model blindness under networked contagion – reflects the same ontological error applied at firm level: managers apply machine logic to adaptive systems. The 2008 financial crisis provided the paradigmatic stress test. Reductionism missed network topology; atomism missed correlated behaviour; closure missed regulatory co-evolution; equilibrium-centred models missed self-reinforcing instability; ergodicity assumptions missed regime change. The simultaneity of these failures is diagnostic: not coincidence but the predictable consequence of a tightly coupled architecture failing under system-wide stress. What collapsed was not merely a financial system but the adequacy of a mechanistic ontology for an adaptive, interaction-dominated economic reality.

The four stages together establish the chapter's central conclusion. The failures documented are not empirical anomalies awaiting better data, nor modelling deficiencies awaiting better techniques, nor gaps awaiting incremental extension. They are generated by the paradigm's foundational commitments. When economic reality is presupposed to consist of stable substances, decomposable systems, reversible time, and well-defined possibility spaces, the resulting theory is necessarily blind to emergence, learning, structural transformation, and endogenous novelty. The blindness is architectural. No amount of technical refinement can repair it without altering the assumptions that produce it.

This conclusion does not imply that mechanistic economics is internally incoherent or universally inadequate. The paradigm is logically consistent within its own assumptions and effective within the restricted domains those assumptions define – near-equilibrium systems, decomposable problems, stable statistical structures. Its failure is not one of internal contradiction but of ontological misalignment: it explains well what it is built to explain; it fails where economic reality exceeds the boundaries it presupposes.

For the theory of the firm, this conclusion is decisive. Firms are not decomposable aggregates of individual transactions but emergent, adaptive, relationally constituted systems whose capabilities, boundaries, and identities evolve through historical processes. A paradigm that cannot represent emergence, learning, path-dependence, or relational architecture cannot represent firms as they actually function. The 15-assumption architecture eliminates precisely

the meso-level where firm-level causation operates. The explanatory failures documented in this chapter are therefore not peripheral anomalies but direct consequences of applying a lens designed for a categorically different kind of object.

The implication is methodological in the deepest sense. The limits identified cannot be addressed through better estimation, richer datasets, or more sophisticated optimisation techniques. They arise at the level of what economic systems are taken to be. This diagnosis therefore establishes the necessity – not the preference – for ontological inquiry. If the mechanistic paradigm's failures are structural, then meaningful theoretical advance requires re-examining the assumptions concerning substance, time, causation, and system structure that organise economic reasoning itself. Such inquiry is imposed by the explanatory failures documented here.

Chapter 4 accordingly moves from diagnosis to ontological critique. Having demonstrated *that* the mechanistic paradigm fails under complexity, and *how* that failure is generated by its foundational architecture, the dissertation now turns to *why* those foundations are inadequate – interrogating the ontological presuppositions that must be transformed for a theory capable of accommodating emergence, relational structure, historical irreversibility, and adaptive change.

Chapter 4 – Ontological Critique: The Failure of Mechanistic Being

Introduction

Chapter 3 demonstrated that the recurrent failures of mainstream economic theory – predictive breakdowns, control illusions, and systemic crises – are not accidental, nor are they the result of insufficient data, poor calibration, or temporary modelling deficiencies. Rather, they arise because dominant economic frameworks are built upon a set of deep ontological commitments that are fundamentally misaligned with the nature of economic reality. Chapter 4 addresses this deeper layer of the problem. It shifts the analysis from *what fails* in economic modelling to *why those failures are structurally inevitable* by interrogating the ontology that underpins mechanistic economic reasoning.

The central claim advanced in this chapter is that mainstream economics remains anchored in a mechanistic ontology inherited from the Cartesian–Newtonian worldview. This ontology presumes that reality is composed of stable, self-contained entities whose properties are intrinsic, whose behaviour is analysable in isolation, and whose aggregate behaviour can be inferred by decomposition and recombination. Within such an ontological frame, equilibrium, optimisation, representative-agent modelling, and linear aggregation appear not merely convenient but conceptually coherent. However, as the empirical and theoretical analysis of Chapter 3 established, these assumptions systematically fail when applied to economic systems characterised by interaction, feedback, emergence, and historical irreversibility.

Chapter 4 does not treat these failures as methodological anomalies that can be repaired through technical refinement. Instead, it advances the more radical diagnosis that they originate at the ontological level. The crisis is not one of tools, but of what economics implicitly assumes *exists*. By reconstructing these assumptions explicitly, this chapter exposes why successive waves of model adjustment – adding frictions, heterogeneity, expectations refinements, or behavioural corrections – have failed to resolve the discipline’s foundational blind spots.

Crucially, this chapter is diagnostic rather than constructive in purpose. Its task is not to offer an alternative ontology or to develop a new analytical framework, but to demonstrate why the inherited mechanistic ontology cannot serve as a viable foundation for understanding complex economic systems. In line with the defensive containment principles governing this dissertation, Chapter 4A confines itself to critique, problematisation, and conceptual mapping. It deliberately refrains from advancing solution architectures, methodological prescriptions, or operational frameworks, which are addressed elsewhere in the broader research programme.

The critique proceeds by isolating and dismantling three interlocking ontological commitments that together constitute the hard core of mechanistic economics.

First, ontological reductionism is examined as the assumption that complex economic phenomena can be exhaustively explained by decomposing them into simpler constituent parts. This commitment underwrites representative-agent modelling, equilibrium aggregation, and the search for micro-foundations, while simultaneously erasing interaction effects and emergent properties that dominate real economic dynamics.

Second, the chapter interrogates substance ontology, the metaphysical tradition that privileges static being over dynamic becoming. By treating firms, agents, and markets as stable entities endowed with intrinsic essences, substance ontology renders time, history, learning, and irreversibility conceptually secondary. Economic change is thereby reduced to deviation from equilibrium rather than recognised as a constitutive feature of economic reality.

Third, the chapter critiques atomism, expressed in economics as methodological individualism. This assumption treats individuals as ontologically prior to relations and organisations, thereby collapsing explanation into micro–macro aggregation and eliminating the intermediate, meso-level domain where coordination, institutions, routines, and organisational capabilities emerge.

These three commitments – reductionism, substance ontology, and atomism – do not operate independently. Together, they form a coherent but deeply flawed ontological package that systematically excludes the very phenomena responsible for instability, innovation, coordination failure, and systemic crisis. Their most consequential effect is the erasure of the meso-level: the relational and organisational domain between individuals and aggregates where economic causation predominantly operates. This chapter makes that erasure visible, but it does not yet reconstruct what a meso-level ontology would require.

In Lakatosian terms, the analysis that follows targets the hard core of the dominant economic research programme rather than its protective belt. The ontological commitments examined here are not auxiliary assumptions that can be modified without consequence; they are the metaphysical foundations that render equilibrium, optimisation, and aggregation intelligible in the first place. Once these foundations are shown to be incompatible with the nature of complex adaptive systems, the persistence of mechanistic economics can no longer be justified on technical grounds.

The chapter proceeds as follows. Section 4.1 examines ontological reductionism and demonstrates why decomposability destroys, rather than reveals, the properties of complex economic systems. Section 4.2 traces the metaphysical lineage of substance ontology and shows how privileging being over becoming eliminates time, history, and endogenous novelty from economic theory. Section 4.3 completes the critique by analysing atomism and methodological individualism, exposing how the prioritisation of isolated elements over relations renders meso-level economic structures ontologically invisible.

Taken together, these sections establish that the failures diagnosed in Chapter 3 are not contingent modelling errors but necessary consequences of an inherited ontology that is structurally incompatible with complex economic reality. The purpose of Chapter 4A is to make that incompatibility explicit.

4.1 Ontological Reductionism: The Mechanistic Premise That Wholes Equal Sums

Introduction

This subsection problematises ontological reductionism – the metaphysical commitment that complex systems can be exhaustively understood by decomposing them into constituent parts.

Following Van Riel and Van Gulick (2024, p. 3), we define ontological reductionism as the metaphysical position that higher-level complex entities are fundamentally composed of aggregations of simpler building units. This entails that wholes equal the sum of their parts (Richardson and Cilliers, 2001, p. 8), system properties reduce without remainder to component properties, and “complex systems or phenomena can be understood by the analysis of their simpler components in isolation” (Fang and Casadevall, 2011, p. 1401). Reductionism is thus not merely an analytical preference but a deep ontological commitment that determines what economic theory can and cannot recognise.

As Chapter 3 demonstrated, reductionism is not an incidental modelling choice. It is the ontological root of representative-agent, equilibrium, and linearisation assumptions that produce predictive failure (Section 3.4.1) and managerial illusion of control (Section 3.4.2). The present subsection therefore does not introduce a new critique; it makes explicit the metaphysical premise already shown to generate systemic blindness.

The subsection serves three purposes within the dissertation’s diagnostic architecture. First, it establishes why the transformation away from mechanistic economics is necessary by demonstrating that reductionism fails not merely empirically but ontologically: it is incompatible with emergent, relational, and non-linear phenomena. Second, it clarifies why repeated refinements within mainstream modelling cannot resolve the foundational issue if the underlying ontological premise remains intact. Third, it reveals reductionism’s most consequential analytic effect: the systematic erasure of the **meso-level** – the intermediate domain where networks, institutions, routines, norms, and coordination mechanisms operate.

The subsection proceeds as follows. Section I reconstructs the reductionist premise and its historical entrenchment in economics. Section II develops five arguments demonstrating reductionism’s ontological failure. Section III articulates the meso-level consequences of reductionist erasure. Section IV addresses counter-arguments. Section V synthesises the diagnosis. In Lakatosian terms (Section 3.2.4), the target is the hard core of an inherited research programme: reductionism is not an auxiliary assumption but a foundational premise that makes equilibrium, optimisation, and aggregation appear coherent.

Before proceeding, a clarification is necessary. Reductionism appears intuitive precisely because it works extraordinarily well for mechanistic systems. The achievements of physics and chemistry testify to its power under conditions of decomposability and weak interaction. The critique developed here does not deny those achievements. It demonstrates instead that transferring reductionist logic to complex adaptive systems constitutes a categorical error. In such systems, interaction effects, feedback loops, and emergent properties are not residual complications; they are constitutive features. Decomposition in this context does not merely simplify analysis – it destroys the phenomenon to be explained.

4.1.1 The Reductionist Premise

The reductionist premise pervading mainstream economics reflects a deep ontological inheritance from the Scientific Revolution. At its core lies a seductive intuition: complex systems can be exhaustively understood by decomposing them into constituent parts, analysing each independently, and aggregating their properties to explain the whole. This is the

mechanistic promise – reality as clockwork where understanding the gears means understanding the clock.




Figure 1: Figure 4.1 – The Reductionist Inheritance Chain (From Physics to Economics to Management).

This figure summarises the historical transmission of decomposability assumptions from Newtonian mechanics into economic reasoning and managerial metaphors. It functions as an orientation device for the critique developed in Section 4.1.1.

In economics, the reductionist premise manifests as the belief that macroeconomic phenomena are “merely” aggregated microeconomic behaviours; that firms are collections of individuals whose properties sum to organisational properties; that markets are individual transactions writ large. The representative agent framework – where heterogeneous individuals collapse into a single agent whose scaled-up behaviour represents the economy – is reductionism’s methodological expression and the direct ontological root of macro-level blindness. As Kirman (1992, p. 117) observes, “the reduction of the behavior of a group of heterogeneous agents... to that of a single individual” is “illegitimate” because it “rules out the possibility of any kind of collective behavior that is not already implicit in individual behavior.”

Kirman’s objection is not merely technical. It expresses the deeper ontological point that collective behaviour is not supervenient on individual behaviour alone, because interaction structures generate causal powers that no individual possesses.

The relationship between this subsection and Chapter 1's historical analysis is mutually reinforcing. Chapter 1 established that the mechanistic worldview originated in Descartes's analytical method and Newton's mechanical universe – the seventeenth-century conviction that nature operates like clockwork, decomposable into independent mechanisms. The present problematisation demonstrates why that inheritance now constitutes an ontological barrier to understanding economic complexity. The historical analysis provides genealogy; the present analysis provides critique. Together they establish that transformation requires confronting an inherited ontological commitment rather than merely refining particular models.

The reductionist premise has methodological consequences that extend throughout economic analysis. If firms are “merely” aggregations of individuals, organisational analysis reduces to psychology. If markets are “merely” aggregations of transactions, market dynamics reduce to individual decisions. If institutions are “merely” aggregations of rules, institutional change reduces to policy choice. Each reduction eliminates the emergent, relational, and systemic properties that define the phenomena. The methodological expression of reductionism – methodological individualism – thus embeds ontological commitments that predetermine what can and cannot be seen.

Despite extensive post-crisis critique, the overwhelming majority of policy-relevant macroeconomic models – DSGE, HANK variants included – continue to rely on representative agents, equilibrium convergence, and decomposable structures.

4.1.2 Five Arguments Against Ontological Reductionism

Argument 1: Decomposability Destroys Complexity

Philip Anderson's seminal paper “More is Different” articulates complexity science's foundational insight: emergent properties vanish when systems are decomposed. As Anderson (1972, p. 393) argues, “the constructionist hypothesis breaks down when confronted with the twin difficulties of scale and complexity.” This is not an analytical inconvenience but an ontological fact. When you dissect a living organism to study its parts, you no longer have a living organism – you have dead tissue. The property “alive” does not reside in any component; it emerges from organised interaction. Decomposition destroys precisely what we seek to explain.

Ackoff (1999, pp. 16–18) formalises this as two critical properties of systems. First, “properties of the system are properties of the whole, which none of its parts have. Therefore, when a system is taken apart, it always loses its essential properties” – because a system is “NOT the sum of its parts, it is the product of their interactions.” Second, “the essential properties of the parts are lost when the part is separated from the system of which it is a part.” These two properties create a dilemma for analytical thought: “the first step of analysis is to take the system apart,” yet systems lose all essential properties when taken apart; “the second step is to understand the behavior of the parts taken separately,” yet parts lose their essential properties when the system is taken apart. Therefore, Ackoff concludes, “a system cannot be understood by analysis” – a complementary method (synthesis) is required.

DECOMPOSABLE VS NON-DECOMPOSABLE SYSTEMS
Why Analysis Works for Machines but Fails for CAS

DECOMPOSABLE (COMPLICATED)		NON-DECOMPOSABLE (COMPLEX)
Whole-Part:	Whole = sum of parts	Whole > sum of parts
Properties:	Reside IN elements	Emerge BETWEEN elements
Analysis:	Isolation valid	Isolation destroys
Explanation:	Bottom-up aggregation	Multi-level, relational
Change:	Modify parts → predict whole	Systemic, non-linear, emergent
Example:	Clock mechanism	Living organism, firm culture
IMPLICATION: Economics studies COMPLEX systems using COMPLICATED methods – a categorical error that explains systematic failures.		

Figure 2: Figure 4.2 – Decomposable vs Non-Decomposable Systems: Why Analysis Works for Machines but Fails for CAS

This figure contrasts complicated (decomposable) systems with complex (non-decomposable) systems, clarifying why reductionist analysis preserves phenomena in the former but destroys them in the latter.

Firm-level translation: A classic illustration of emergent organisational innovation is 3M's Post-it Notes – frequently cited in the innovation literature as an archetype of recombinatory emergence rather than planned R&D (e.g., Arthur, 2021; Nelson & Winter, 1982). The Post-it emerged not from planned R&D but from the interaction of two “failed” projects: Spencer Silver's 1968 development of a “weak” adhesive that would not stick permanently and Art Fry's 1974 frustration with bookmarks falling out of his hymnal. Neither component had value in isolation; the innovation emerged from their interaction within 3M's innovation ecosystem, which encouraged scientists to spend 15% of time on personal projects and share findings across divisions. For five years, Silver presented his “useless” adhesive at internal seminars seeking applications – an organisational communication structure that sustained exploration that no individual department would have funded. Any attempt to decompose the Post-it into its “parts” destroys the causal explanation, because the novelty resided in the interaction between components within a permissive organisational ecology. By 2020, Post-it Notes had expanded into a family of over 4,000 products.

Simon (1962, p. 467) captured this precisely: complex systems are “made up of a large number of parts that interact in a nonsimple way” such that “given the properties of the parts and the laws of their interaction, it is not a trivial matter to infer the properties of the whole.” The implication is ontological: system-level behaviour is underdetermined by component-level properties.

Argument 2: Infinite Regress and the Quest for Foundations

Reductionism promises explanation by decomposition: explain the whole by its parts. As philosopher of science Mario Bunge (1973, p. 38) observes, “the micro-reductionist program... is of course the program of atomistics, from Democritus to the present.” But this strategy generates infinite regress. If the firm is explained by individuals, then individuals must be explained by psychological states, those by neural patterns, those by chemical reactions, those by quantum mechanics. There is no foundational “atom” at which explanation terminates – each level reveals further levels below requiring their own explanation.

Even if one reached physical constituents, Storm (2021, p. 93) identifies a critical gap between reduction and reconstruction: “the reductionist hypothesis does not by any means imply a ‘constructionist’ one.” One cannot build a macro theory by stacking micro axioms, because no finite set of micro-laws determines the behaviour of systems whose interactions generate novel structures.

Fromm (2005, pp. 14, 16) formalises this as “strong emergence” – the appearance of “emergent structures on higher levels of organization or complexity which possess truly new properties that cannot be reduced, even in principle, to the cumulative effect of the properties and laws of the basic parts.” Simon (1962, pp. 468–469) captures this through near-decomposable hierarchies operating across distinct time scales.

Firm-level translation: The Daimler–Chrysler merger (1998), widely analysed in the management and organisational-culture literature, provides a concrete illustration. The reductionist logic was impeccable: sum Daimler’s engineering excellence with Chrysler’s mass-market efficiency. But explaining the merger’s failure requires moving beyond component analysis. The cultural clash was not a surface friction but a system-level property produced in interaction. Explaining the clash by individual attitudes merely relocates the problem to another level and reproduces regress. By 2007, Daimler sold Chrysler to Cerberus Capital for \$7.4 billion – an outcome that reductionist synergy arithmetic did not anticipate.

Argument 3: Aggregation Errors in Economics

Mainstream economics institutionalises reductionism through aggregation: it treats the economy as if it were a representative agent whose behaviour can be scaled up from individual behaviour. This commits systematic ontological and causal errors rather than harmless technical approximations.

Error 1: Ignoring distributional dynamics. Aggregation assumes that distributions do not matter – only means. But in complex adaptive systems, distributions are causally active. Who holds risk, who faces liquidity constraints, and who possesses information determines system behaviour. Crises are distributional events, not average events.

The 2008 financial crisis revealed this brutally. As Bernanke (2010, pp. 4–13) acknowledged, the financial system operates as “a network of interconnected institutions” where “the failure of one institution can trigger cascading effects.” Aggregate capital ratios and average leverage masked dangerous concentrations of exposure within particular institutions and within

particular network positions. The system appeared stable in the mean while being fragile in its structure.

Error 2: Neglecting network effects. Aggregation treats agents as independent units whose behaviour can be summed. But modern economies are constituted by networks of interaction. Value depends on adoption externalities, information flows through social and financial networks, and shocks propagate through interconnections rather than through isolated units. Kirman (2016, p. 560) states this directly: “the structure of the financial network is at least as important as the health of the individual institutions.” Aggregation removes network topology – exactly the structure that determines resilience, contagion, and coordination.

Error 3: Assuming composition. Aggregation presumes that what is rational for individuals is rational for the collective. This is the fallacy of composition. But in complex systems, individually prudent behaviour can generate collectively catastrophic outcomes: precautionary saving produces recessions; individual deleveraging produces financial collapse. Gallegati and Gallegati (2025, p. 3064) explain why: “a system is said to be linear if it can be broken down into several parts that are autonomous from each other, independent... The concepts of heterogeneity and interaction necessarily lead to the question of non-linearity.” Aggregation presumes linearity where interaction produces non-linear system effects.

Firm-level translation: The AOL–Time Warner merger (2000) committed all three aggregation errors simultaneously, producing what is widely regarded as the most destructive merger in corporate history. The \$164 billion deal rested on pure aggregation logic: combine AOL’s 27 million internet subscribers with Time Warner’s content, cable infrastructure, and media brands and value would scale up.

- **Distributional error:** AOL was valued on average subscriber numbers, ignoring the distribution of user quality; most subscribers were dial-up customers about to defect to broadband competitors – a structurally unstable base concealed by averages.
- **Network error:** The deal assumed independence and additivity; AOL’s closed “walled-garden” model was incompatible with the emerging open-internet ecosystem, so network effects worked in reverse and accelerated decline rather than creating synergy.
- **Compositional error:** What was rational for each firm separately became collectively destructive; cultural and operational interaction produced paralysis rather than complementarity.

Within two years, AOL Time Warner wrote down over \$99 billion in value; by 2009 the firms were separated. The aggregation logic that justified the merger erased exactly what mattered: interaction structure, network position, and emergent organisational dynamics that determined whether combination would create or destroy value.

Conclusion

The representative agent is not a harmless simplification. It is an ontological error that systematically eliminates heterogeneity, networks, and interaction-driven emergence – the very mechanisms through which complex economic systems actually operate.

Argument 4: Neglect of Interactions

Perhaps reductionism's most crippling limitation is its treatment of interactions as secondary. Mechanistic ontology assumes that elements are primary (self-contained units with intrinsic properties), that interactions are frictions (complications added to otherwise independent behaviour), and that isolation is analytically valid (systems can be studied through *ceteris paribus* decomposition). Complex adaptive systems invert this hierarchy: interactions are primary; elements are relationally constituted.

As Holland (1995, p. 11) emphasises, “the interactions, not the actions of the individual agents, are the primary source of the complexities.” Cilliers (1998, p. 3) makes this explicit: “the relationships between the elements are more important than the elements themselves... The identity of the elements in the system is constituted by the relationships between them.” Reductionism's neglect of interactions does not merely distort these phenomena – it renders them ontologically invisible.

Technological lock-in provides a canonical illustration. Arthur (2014, p. 17) shows that under increasing returns, a technology that “gets ahead, possibly by small chance events, given increasing returns it will gain further advantage and get further ahead.” VHS defeated the technically superior Betamax not because of intrinsic quality but because of network externalities, complementary investments, and learning effects. Technology value was produced by the adoption network, not by the artefact itself.




Figure 3: Figure 4.3 – The Ontological Inversion: From Elements-Primary to Interactions-Primary

This figure shows how mechanistic ontology places properties “in” elements and treats relations as secondary, whereas a complexity ontology treats relational structure as constitutive of properties and identity.

Firm-level translation: IDEO’s “Deep Dive” methodology – famously demonstrated in its 1999 five-day redesign of a shopping cart for ABC Nightline – rejects individual-genius models of innovation. Teams are composed across engineering, anthropology, psychology, and industrial design not to aggregate expertise but to generate emergent insight through interaction among different cognitive frames. Rapid prototyping accelerates interaction cycles; flat hierarchy reduces filtering; workspaces amplify unplanned encounters. IDEO’s founder David Kelley emphasises “building on the ideas of others”: innovation arises from relational recombination, not from summing individual talents.

Argument 5: Reductionism Precludes Emergence

The most fundamental problem is logical: reductionism is incompatible with emergence. If the whole equals the sum of its parts, then by definition nothing can emerge – all properties must already exist at the component level. Yet complex adaptive systems are defined precisely by emergence: properties arise at the system level that do not exist in the parts.

Economic reality is saturated with such emergent phenomena:

- **Firm capabilities** emerge from organisational routines and accumulated collective experience, not from summing individual skills (Nelson & Winter, 1982, pp. 14–18).
- **Market prices** emerge from trading dynamics and interaction patterns, not from aggregating intrinsic values (Kirman, 2016, p. 560).
- **Innovation** emerges from recombination of existing technologies and ideas, not from isolated genius (Schumpeter, 1934, pp. 64–66; Arthur, 2021, p. 12).
- **Institutions** emerge from interaction among heterogeneous, boundedly rational agents embedded in networks, not from top-down design (Alston et al., 2023, p. iv).

Anderson’s dictum – “more is different” – expresses this ontological discontinuity (Anderson, 1972, p. 393): new properties come into being at higher levels of organisation that were not present in the components. A theory committed to reductionism therefore cannot explain phenomena that define economic life; it rules them out by definition.

Firm-level translation: W. L. Gore & Associates, the maker of Gore-Tex ($\approx$ \$3.5 billion annual revenue), exemplifies emergence in organisational form. Gore operates through a “lattice” structure – no formal hierarchy, no job titles, no bosses. Associates choose projects; leaders emerge from followership; compensation is peer-determined. A reductionist lens predicts chaos; yet Gore has remained profitable for over five decades and sustains high innovation rates. Its advantage is not the sum of individual talents but an emergent property of the interaction structure shaped by history and organisational ecology.

III. The Meso-Level Consequences

The five arguments converge on a single, decisive insight: reductionism’s most consequential failure is the systematic erasure of the meso-level. The meso-level is the intermediate analytical domain between micro-level individual agents and macro-level aggregates – the level at which networks, institutions, routines, norms, and coordination mechanisms operate.

THE MESO-LEVEL ERASURE
Reductionism's "Invisible Middle"

LEVEL	ENTITIES	REDUCTIONIST TREATMENT	CAS TREATMENT
MACRO	System aggregates GDP, inflation, market indices	Derived from micro	Emergent from meso-level dynamics
	↑	[JUMP - NO MIDDLE]	↑
MESO	Networks, institutions Routines, capabilities Coordinating mechanisms		WHERE EMERGENCE OCCURS
	↑		↑
MICRO	Individual agents Persons, firms Transactions	Primary explanatory unit	Nodes constituted by relations

The meso-level is the "invisible middle" – where agent constitution occurs, where capabilities emerge, where coordinating mechanisms operate. Reductionism cannot see it because the micro-macro jump erases it by design.

Figure 4: Figure 4.4 – The Meso-Level Erasure: Reductionism's "Invisible Middle"

This figure depicts the reductionist micro–macro jump that deletes meso-level structures where coordination and emergence occur, making relational causation systematically invisible.

This connects directly to Chapter 4's later analysis of network topology. Networks do not exist at the level of isolated agents or aggregate totals; they exist between agents as structured patterns of interaction. Institutions do not operate at the level of individual choice or macro constraints; they function as meso-level coordination devices that stabilise expectations, enable cooperation, and shape behaviour. Organisational routines and firm capabilities similarly emerge at the organisational level, between employees and market outcomes.

By jumping directly from individuals to aggregates, methodological individualism systematically erases these constitutive relational structures. For the firm conceived as a complex adaptive system, this erasure is decisive. Firm capabilities, organisational culture, learning dynamics, and strategic coherence all exist at the meso-level – not inside individuals and not in market totals, but in the interaction structures that bind individuals into productive collectives. A theory that cannot see this level cannot explain how firms differ, why strategies work, or how capabilities evolve. This grounds the dissertation's distinctive contribution: systematic meso-level tracking of networks, routines, and institutional scaffolding that mainstream economics' reductionist ontology deletes by construction.

The practical consequences are immediate. When consultants diagnose failing firms using reductionist tools, they focus either on individual performance (micro) or market position (macro), while overlooking the meso-level pathologies – broken communication channels, misaligned routines, dysfunctional coordination, eroded trust – where organisational failure actually resides. The repeated failure of “best practice” transfers reflects exactly this blindness. Practices that succeed in one firm fail in another not because workers are incompetent or markets differ, but because the meso-level interaction structures that made those practices effective cannot be transplanted through manuals, incentives, or process maps. What matters is not the practice itself, but the relational architecture in which it is embedded – and that architecture is invisible to reductionist economics.

IV. Counter-Arguments Addressed

“Reductionism Works in Physics – Why Not Economics?”

Objection.

Reductionism has been extraordinarily successful in physics and chemistry. The Standard Model explains particle behaviour; chemistry reduces to physics; molecular biology reduces to chemistry. If reductionism works for understanding the physical universe, surely it can work for understanding economies.

Response.

This objection commits a category error by conflating complicated systems with complex systems. The distinction is foundational. Laughlin (2005, pp. 208–209), a Nobel laureate in physics, draws it sharply:

“There are circumstances in which the microscopic laws governing a system cease to be relevant because they have been superseded by macroscopic ones that emerge from them.”

Anderson (1972, p. 393) established this decisively: “The constructionist hypothesis breaks down when confronted with the twin difficulties of scale and complexity.”

Physics succeeds with reductionism when studying complicated systems – many parts, simple interactions, and preserved decomposability. It fails when studying complex systems – many parts, rich interactions, and non-decomposability. The Large Hadron Collider isolates particles precisely because particle properties are intrinsic. Firms, markets, and institutions cannot be isolated because their properties are relational and interaction-constituted.

Even within physics, reductionism has limits: superconductivity, turbulence, phase transitions, and pattern formation all exhibit emergent properties that resist micro-level derivation. Cilliers (1998, pp. 2–3) formalises the distinction: complex systems exhibit non-linear interactions, rich interconnections, and emergent properties that complicated systems do not.

Economies are complex adaptive systems, yet economics applies methods appropriate to complicated mechanical systems. The physics analogy therefore cuts against the objection:

physics itself recognises the limits of reductionism for complex phenomena – economics has not yet done so.

“We Need Micro-Foundations for Rigorous Economics”

Objection.

Without micro-foundations, economics becomes mere description without explanation. The Lucas critique showed that policy analysis requires understanding individual responses. We must reduce macro phenomena to micro behaviour to achieve scientific rigour.

Response.

This objection assumes that micro-foundations uniquely determine macro-outcomes – an assumption mathematically refuted by the Sonnenschein-Mantel-Debreu (SMD) theorem. Storm (2021, p. 83) shows that “assumptions guaranteeing good behavior at the microeconomic level do not carry over to the aggregate level.”

Kirman (1992, p. 119) is explicit: “There is no plausible formal justification for the assumption that the aggregate of individuals... should be treated as though it were a single individual.”

The same macro-phenomenon can arise from radically different micro-configurations – this is multiple realizability, which Batterman (2002, pp. 76–78) shows is characteristic of emergent systems. Macro-level structures therefore possess causal autonomy: they do causal work that cannot be derived from micro-descriptions alone.

The Lucas critique itself illustrates this circularity: policy changes alter macro conditions, which reshape micro behaviour, which in turn alters macro-outcomes. Neither level is foundational. Rigour requires analysing phenomena at the level where their causal powers exist, not dissolving them into lower-level variables that underdetermine outcomes.

“Emergence Is Just Unexplained Mechanism”

Objection.

Emergent properties are merely properties we have not yet explained mechanistically. As science advances, emergence will dissolve into mechanism. To claim irreducible emergence is to give up on explanation.

Response.

This objection confuses epistemological with ontological emergence. Bedau (1997, pp. 376–378) distinguishes weak emergence (surprising but reducible) from strong emergence (metaphysically irreducible). The claim here is ontological.

O'Connor and Wong (2015, p. 3) define strong emergence as involving “novel causal powers... not explicable in terms of the causal powers of the underlying constituents.”

Fromm (2005, p. 16) formalises this: emergent properties “cannot be reduced, even in principle, to the cumulative effect of the properties and laws of the basic parts.”

This is not explanatory surrender but explanatory re-location. Emergent properties must be analysed at the level where they exist. The competitive advantage of W.L. Gore or IDEO is not an attribute of individual employees; it is a property of organisational interaction structures. Explaining it via reduction would destroy precisely what requires explanation.

“Methodological Individualism Is Just a Method, Not an Ontology”

Objection.

Methodological individualism (MI) is merely a research strategy – start with individuals and build up. It does not commit to ontological claims about what exists. One can use MI while acknowledging emergent social properties.

Response.

This objection understates the ontological commitments embedded in method choice. Hodgson (2007, pp. 219–221) shows that methodological individualism “carries strong and contestable ontological implications” because the choice of basic units presupposes a theory of what is causally fundamental.

To begin with individuals as explanatorily primary is already to assume ontological priority. Lawson (2015, pp. 113–115) demonstrates that methods determine what can be seen: if individuals are assumed primary, relational structures appear only as “frictions” or “externalities” rather than as constitutive features.

Sawyer (2005, pp. 189–191) is explicit: methodological individualism presupposes that social properties supervene on individual properties. But if emergence is real, social properties do not supervene – they possess autonomous causal powers.

Method and ontology are mutually reinforcing. MI (Methodological Individualism) cannot be ontologically neutral, because neutrality would require a framework that does not privilege any level of analysis – precisely what MI denies by construction.

Conclusion

This subsection has developed five independent but convergent arguments demonstrating that ontological reductionism is incompatible with the nature of complex adaptive systems:

1. **Decomposability destroys** the emergent properties it seeks to explain
2. **Infinite regress** undermines foundationalist ambitions
3. **Aggregation commits systematic errors** regarding distributions, networks, and composition
4. **Neglect of interactions** inverts the true locus of causation
5. **Reductionism logically precludes emergence**, the defining feature of economic reality

Together, these establish that reductionism fails not contingently but structurally.

Diagnostic synthesis: Reductionism assumes decomposability, linear aggregation, and component primacy. Complex adaptive systems are characterised by interaction-generated properties, non-linear dynamics, and multi-level causal organisation. The implication for economic theory is therefore negative but decisive: when the ontological premise is reductionist, analysis is structurally predisposed to miss emergent behaviour, coordination dynamics, and meso-level causation.

The Ontological Transformation

Table 98: Table 4.1 – What Reductionist Ontology Systematically Eliminates from View

FROM (Mechanistic)	TO (What Becomes Invisible / Mis-specified)
Wholes = sum of parts	Wholes > sum of parts
Properties aggregate linearly	Properties emerge non-linearly
Decomposition is valid	Decomposition destroys
Analysis is sufficient	Synthesis is required
Elements are primary	Interactions are primary
Micro-foundations foundational	Multiple levels are irreducible

This is not a methodological preference but a change in what is assumed to exist.

The Meso-Level Imperative

Reductionism's most damaging effect is the elimination of intermediate scales of organisation. By forcing explanation to jump directly from individuals to aggregates, it erases the meso-level – where networks, routines, institutions, and organisational capabilities form.

Yet this is precisely where economic causation operates in real firms and markets. Capabilities do not exist at the level of individual workers, nor at the level of market totals; they exist in patterns of coordination. Institutions are not rules or individuals but stable interaction structures. Competitive advantage is not located in agents but in relational architectures.

Reductionism does not misdescribe this level – it renders it ontologically invisible.

The Aha Moment

Analysis destroys what it seeks to understand.

Synthesis is not optional; it is ontologically required for complex systems.

Reductionism cannot be “extended,” “improved,” or “softened” to accommodate emergence. A framework in which wholes equal parts cannot explain phenomena that only exist between parts.

Understanding firms, markets, and economies as complex adaptive systems therefore requires a new ontology – one in which emergence, interaction, and multi-level organisation are foundational rather than anomalous.

This conclusion makes the reconstruction in Section 4.2 (Emergence) not a theoretical preference but a logical necessity.

4.2 Substance Ontology: The Metaphysical Foundations of Being Over Becoming

Introduction

This subsection deepens the chapter's ontological problematisation by challenging substance ontology – the long metaphysical commitment to treating reality as unchanging Being rather than dynamic Becoming. Following Aristotle's *Metaphysics* (Book Z), substance ontology treats reality as fundamentally composed of discrete entities (*ousiai*) possessing essential properties that persist through change. As Robinson (2024, p. 1) explains, the term “substance” derives from the Greek *ousia* (from *einai*, “to be”), designating those entities that function as the foundational beings of a philosophical system. Under this ontology, identity is grounded in timeless essence, while change is relegated to secondary, accidental modification.

This subsection serves three functions within the chapter’s diagnostic architecture. First, it demonstrates that mainstream economics’ failures do not arise only from reductionism (Section 4.1.1) but from a deeper metaphysical commitment: treating firms, markets, and agents as stable substances with intrinsic essences rather than as historically evolving processes. Second, it traces the metaphysical lineage of this commitment from Parmenides through Aristotle into modern scientific and economic reasoning. Third, it shows why substance ontology is structurally incompatible with complex adaptive systems: if reality consists of fixed essences, then emergence, novelty, evolution, and path-dependence become conceptually unintelligible.

The subsection proceeds as follows. Section I reconstructs the Parmenidean legacy and its transmission through Aristotle into economic theory. Section II develops arguments demonstrating substance ontology’s incompatibility with complex adaptive systems. Section III articulates the meso-level and systemic stakes of this incompatibility. Section IV addresses counter-arguments. Section V synthesises the diagnosis.

Before proceeding, a natural intuition must be acknowledged. Substance thinking appears plausible because entities seem to persist through change – Google remains “Google” despite turnover of employees, technologies, and strategies. Yet this apparent stability is misleading. What persists is not an unchanging substance but a self-maintaining pattern of organisation – routines, coordination structures, and institutional memory. A river persists not because its water remains the same, but because the pattern of flow continues through continuous replacement. Firms are rivers, not rocks.

4.2.1 The Parmenidean Legacy

Being Over Becoming: The Foundational Choice

Western philosophy inherits one of its deepest ontological commitments from Parmenides (5th century BCE): reality is eternal, unified, and unchanging Being. Change, multiplicity, and becoming were treated as illusory – mere appearances concealing an underlying stable essence. As Kirk, Raven, and Schofield (1983, pp. 239–262) document, Parmenides grounded truth in stasis (*στάσις*) – immobility and permanence – rejecting temporal flux as epistemically and ontologically secondary. This was not a technical claim but a metaphysical choice that has structured Western thought ever since.

Aristotle gave this choice its systematic philosophical form. In *Metaphysics* (Book Z), substance (*ousia*) is defined as primary reality: the enduring subject that persists through change. Individual substances – this person, this firm, this object – are ontologically prior to their properties, relations, or transformations. Accidents change; essence remains. Identity resides in what a thing is, not in how it becomes. Change is thereby ontologically downgraded: it modifies substance but does not constitute it.

This framework was transmitted through medieval Scholasticism, Cartesian dualism, and Newtonian mechanics into modern science – and ultimately into economics. Its imprint is evident in canonical modelling assumptions: agents possess fixed preferences; firms possess stable production functions; markets possess natural equilibrium states. These are not merely modelling conveniences but expressions of substance ontology in formal language. Entities are

defined by intrinsic properties; time, history, and transformation are treated as secondary disturbances.

The methodological consequences follow directly. If firms are substances, their behaviour derives from technological essences. If agents are substances, choice follows from pre-given preferences. If markets are substances, outcomes gravitate toward natural equilibria. Comparative statics, optimisation, and equilibrium analysis all presuppose that what matters most is what entities *are*, not how they *become*.

The relationship between this subsection and Section 4.1 is therefore structural rather than additive. Reductionism presupposes substance ontology: if entities have fixed intrinsic properties, they can be analysed independently and recombined. Substance ontology supplies the metaphysical ground on which reductionism stands; reductionism is the methodological expression of substance metaphysics. Challenging one while leaving the other intact is unstable, because the metaphysical premise will continue to regenerate the methodological habit.

Argument 1: Substance Ontology Eliminates Time, History, and Becoming

Substance metaphysics seeks stable essences that transcend time. A firm's "essence" – its production function, its intrinsic capabilities, its preferences – is treated as time-invariant. Change appears only as disturbance around a stable core: shocks, deviations, and perturbations that leave underlying identity intact. Time enters models as a parameter t within equations that remain, in principle, reversible.

This metaphysical structure is embedded in economic formalism. The standard production function,

$$Q = f(K, L)$$

represents the firm as a timeless mapping from inputs to outputs. The function itself stands as the firm's "technology" – its ontological identity. Learning, organisational development, and strategic evolution do not exist within the function; they appear, if at all, as exogenous shifts. The firm has no memory, no developmental path, no accumulated experience. It simply is its production function at each instant.

But real firms are not substances; they are historical processes. As Penrose (2009, p. 237) established, "growth is essentially an evolutionary process and based on the cumulative growth of collective knowledge." Capabilities do not exist timelessly; they are built through irreversible sequences of learning, failure, and adaptation. What a firm can do today depends on what it has done before.

Substance ontology also implies reversibility. If change is deviation from essence, then returning to essence is always possible. This is a Newtonian inheritance: time as a symmetric coordinate that can, in principle, run backward. Economics imported this structure wholesale. As Arthur (2014, p. 5) observes, "standard neoclassical economics lives in a timeless, non-constructive, equilibrium world."

Complex adaptive systems obey the opposite logic. They are irreversible. As Prigogine and Stengers (1984, p. 291) argue, classical dynamics “describes a static world, a world that escapes the arrow of time,” thereby excluding “the irreversible and creative processes that characterize living systems.” In complex systems, history is not noise; it is constitutive. Path-dependence, hysteresis, lock-in, and ratchet effects mean that past choices permanently structure future possibilities. One cannot rewind a firm, a market, or a technology to an earlier state.

Firm-level translation

Nokia’s corporate history exposes the limits of substance ontology. Founded in 1865 as a paper mill, Nokia later became a rubber manufacturer, then a cable producer, then a telecommunications supplier, then the world’s dominant mobile-phone firm, and finally a network-infrastructure and technology-licensing company. Between 1865 and 2015, Nokia reinvented itself at least five times. At its peak in 2007, it held over 40% of the global handset market; by 2014 it had exited the business entirely.

What, then, is Nokia’s “essence”? Substance ontology cannot give a coherent answer. Nokia’s identity is better understood as a path-dependent trajectory shaped by cumulative investments, technological bets, organisational learning, and strategic commitments. A production-function representation of the firm – timeless, reversible, ahistorical – cannot even express this trajectory, let alone explain it.

Diagnostic implication: substance thinking tends to treat time as a reversible parameter and history as disturbance; complex economic evolution instead exhibits irreversibility, cumulative learning, and path-dependent change. Under conditions of historical contingency, substance ontology struggles to represent becoming as a constitutive feature of firm life.

Argument 2: Substance Ontology Cannot Explain How Novelty Arises

If reality consists of stable substances with fixed essences, then the origin of novelty becomes metaphysically unintelligible. Innovation, new technologies, new institutions, and new organisational capabilities involve the appearance of genuinely new properties. Yet within a substance ontology, nothing genuinely new can arise: change must be either rearrangement of pre-existing properties or deviation from an underlying essence.

Mainstream economics typically resolves this tension by externalising novelty. Technological change is treated as an exogenous “shock.” Preferences “shift.” Institutions “change.” But the theory contains no account of where novelty comes from; it is assumed rather than explained. As Arthur (2021, p. 137) notes, equilibrium economics “makes no allowance for the creation of new products or new arrangements... for events triggering novel events, indeed, for history itself.” Real business cycle theory attributes macro fluctuations to “technology shocks,” yet provides no account of how technologies are created. Novelty enters as a *deus ex machina*, not as an endogenous process.

This is analogous to the explanatory limits of pre-Darwinian biology, where species were treated as fixed essences and novelty required external explanation. Darwin’s breakthrough was not only empirical but ontological: species are not substances; they are populations

evolving through variation, selection, and retention. Identity lies in lineage, not essence. Novelty becomes intelligible as an endogenous outcome of evolutionary processes.

Economics repeatedly gestures toward this insight. Schumpeter (1934, p. 63) defined economic development as “such changes in economic life as are not forced upon it from without, but arise by its own initiative, from within.” Development is “entirely foreign” to equilibrium because equilibrium presupposes stasis. Firms are not stable entities awaiting exogenous shocks; they are systems that can generate novelty through experimentation, learning, recombination, and selection. As Arthur (2014, p. 25) puts it, the economy is “alive, ever-changing, organic, and full of messy vitality.”

Firm-level translation

The iPhone did not arrive as an exogenous technology shock. It emerged from Apple’s accumulated trajectory. The Macintosh created a GUI-centred design culture. The Newton taught Apple about mobile computing. The iPod and iTunes developed capabilities in hardware–software integration, digital distribution, and ecosystem orchestration. When Jobs unveiled the iPhone in 2007, he described it as three existing technologies recombined into one device – and then revealed that they were “one device.”

This is endogenous novelty: new value created by recombining existing organisational capabilities inside a firm that had learned how to integrate design, software, hardware, and platforms. By 2022 the iPhone had generated nearly \$400 billion in annual revenue. Substance ontology struggles to represent such novelty because it treats the firm as a stable mapping that passively receives “technology” from outside, rather than as a historically evolving organisation with cumulative capabilities.

Diagnostic implication: within substance metaphysics, novelty appears anomalous and must be imported from outside the system. By contrast, historical and evolutionary accounts treat novelty as a possible outcome of endogenous processes of variation and recombination.

Argument 3: Firms Are Not Static Entities – Identity as Temporally Unfolding Process

What makes Google “the same company” today as when it was two Stanford students in a garage? Not any stable substance. Its employees, technologies, markets, products, and strategies have all changed. What persists are patterns of coordination: how search algorithms are organised, how engineering creativity is cultivated, how advertising and platform revenues are generated, how users and developers are enrolled in a growing ecosystem. These are not things; they are ongoing organising processes. Google does not possess identity in the way a substance possesses properties; it enacts identity through continuing organisation. When the organising stops (bankruptcy, dissolution), the firm ceases to exist.

Penrose (2009, p. 237) captured this when she defined the firm not as a bundle of inputs but as “a collection of productive resources” whose services are organised through administrative coordination. Crucially, those services are not inherent in resources themselves; they are generated by organisational processes – learning, routines, and coordination. Two firms with identical resources produce radically different outcomes because their organising processes

differ. Competitive advantage resides in how resources are organised through time, not in resources treated as timeless substances.

Three defining characteristics of real firms expose the limits of substance ontology.

1. **Coordination.** Following Hayek (1945, pp. 519–520) and Grant (1996, pp. 109–112), firms exist to coordinate distributed knowledge and action. But coordination is not a static attribute; it is a continuously renewed achievement. It can deteriorate (bureaucratic drift), improve (learning), or reorganise (strategic change). A firm whose coordination collapses has not lost a “substance”; it has lost a process.
2. **Adaptation.** As Cyert and March (1963, pp. 99–101) and Teece et al. (1997, pp. 515–518) show, firms adapt through feedback, experimentation, and capability development. Learning is path-dependent: what a firm can do today depends on what it has done before. A production function has no memory; real firms are constituted by their history of learning. Adaptation is not noise around a stable core; it is part of the core.
3. **Meaning-making.** Weick (1995, pp. 4–6) shows that organisations are systems of shared interpretation. Culture, identity, and purpose are continually reconstructed through communication. These are not properties stored inside the firm; they are relational processes enacted in meetings, narratives, routines, and decisions. A firm without meaning-making processes becomes incoherent even if its “resources” remain intact.

Firm-level translation

Toyota’s competitive advantage is not a static asset but a process it enacts. The Toyota Production System (TPS) is not a fixed method; it is an ongoing pattern of continuous improvement (*kaizen*), frontline problem-solving, supplier integration, and cultural reinforcement. When competitors copied visible practices – kanban, just-in-time, andon cords – they did not reproduce Toyota’s performance because those tools are only the surface of a deeper organisational process.

The NUMMI joint venture (GM–Toyota, 1984) demonstrated this sharply: American workers, using Toyota’s routines within Toyota’s interaction system, matched Japanese quality. Yet GM failed to replicate this across its own plants. The difference was not equipment or procedures (substances) but the continuous enactment of learning, responsibility, and improvement. Taiichi Ohno described TPS as “learning to see” – a process of perception and response, not a transferable blueprint. Toyota’s sustained leadership reflects the durability of this processual organising, not ownership of a particular asset.

Diagnostic implication: substance ontology treats firms as entities with stable properties, whereas firm identity and capability appear more plausibly as temporally sustained patterns of coordination, learning, and meaning-making.

Argument 4: Philosophical Foundations Unsuitable to Social Phenomena

Aristotle's Biological Metaphor

Substance ontology was not originally designed for economies or organisations. Aristotle developed it to explain biological organisms. In *Metaphysics* (Book Z) and his biological works, Aristotle treated natural beings as substances possessing intrinsic essences (*ousiai*) that unfold through development. An acorn has the essence “oak tree”; growth is the actualisation of pre-given potential. Change is teleological – movement toward a predetermined form. The organism becomes what it already is “in principle.”

This framework works tolerably well for biological development, where genetic programmes constrain trajectories. It fails when imported into social systems. Firms do not possess predetermined essences waiting to be actualised. They do not grow toward an intrinsic *telos*. They experiment, improvise, learn, compete, adapt, and sometimes fail. Their futures are not encoded in their origins. What a firm becomes depends on historically contingent interactions with technologies, rivals, regulators, customers, and internal decisions. Two startups with identical business plans can diverge radically when they encounter different market feedback, partners, or shocks. This is path-dependence, not essence-actualisation.

Substance ontology struggles to account for such contingency because it assumes identity precedes history. Social reality often inverts this: identity is produced through history.

Heraclitus's Alternative: Process Philosophy as Contrast

Heraclitus (c. 500 BCE) offered an alternative suppressed by substance metaphysics. His doctrine of flux – *panta rhei* (τὰ πάντα ῥεῖ, “everything flows”) – holds that stability is not primitive but emergent. As Rescher (1996, pp. 27–35) shows, Heraclitean process philosophy treats reality as fundamentally dynamic: entities persist only as patterns of ongoing activity, not as enduring substances.

The river is not the water – which is always different – but the self-maintaining pattern of flow. Likewise, firms persist not by retaining an unchanging essence but by reproducing relatively stable patterns of coordination, routines, cultural norms, and strategic orientation through continual turnover of people, assets, and technologies. The legal shell, ownership, or capital stock does not exhaustively define the firm; the pattern of organising does.

Firm-level translation

Amazon did not “contain” cloud computing in its essence in 1994. It began as an online bookstore because that was feasible in that historical moment. It entered music and DVDs (1998), created Marketplace (2000), launched Prime (2005), introduced Kindle (2007), and launched Amazon Web Services in 2006. AWS emerged from internal infrastructural experimentation – building scalable computing to run Amazon's retail platform – combined with the contingent insight that this capability could be offered to others. By 2023, AWS generated \$90.8 billion in annual revenue and became Amazon's main profit engine.

Amazon subsequently expanded into groceries (Whole Foods, 2017), healthcare (Amazon Pharmacy, One Medical), entertainment (Prime Video, MGM), and logistics (delivery fleet, air cargo). No Aristotelian essence predicts this trajectory. Bezos captured the firm’s self-description not as what it *is* but as how it *operates*: “We are genuinely customer-centric, we are genuinely long-term oriented, and we genuinely like to invent.” These are process orientations rather than timeless substance properties. Amazon’s identity is its trajectory of becoming, not a fixed core.

Diagnostic implication: substance ontology encourages equilibrium, optimisation, and timeless modelling by treating firms as things with properties. By contrast, historical and evolutionary descriptions treat firms as patterns of organising that persist through change. For complex adaptive systems, the latter description is often the one that preserves the phenomena to be explained.

The Ontological Transformation

This argument forces a decisive shift:

Table 99: From Substance to Process Ontology of the Firm: Five Foundational Shifts Toward a CAS-Compatible Conception

FROM (Substance)	TO (Process / CAS)
Firms have essences	Firms enact identities
Development = actualisation	Development = contingent evolution
History is accidental	History is constitutive
Identity precedes time	Identity emerges through time
Analysis is cross-sectional	Analysis is historical and dynamic

Substance ontology licenses equilibrium, optimisation, and timeless modelling because it treats firms as things with properties. Process ontology requires **evolutionary, path-dependent, and network-based analysis** because firms are **ongoing patterns of organising**.

Argument 5: Practical Consequences of the Metaphysical Mistake

Why Ontology Shapes Strategy

Substance ontology is not an abstract philosophical error; it directly structures how firms interpret strategy and change. If firms are conceived as stable substances with intrinsic essences, then strategy is naturally framed as a problem of self-discovery: identify “who we really are,” locate “core competencies,” and defend them. Change is interpreted as deviation from identity rather than as identity reproduced through ongoing organising. This logic underlies influential strategy doctrines – core competence theory, resource-based views, and equilibrium competition – each of which presupposes that firms possess enduring, discoverable attributes.

In complex adaptive markets, however, this metaphysical premise is unstable. Arthur (2014, p. 183) shows that increasing returns and interaction generate multiple equilibria, lock-in,

historical path dependence, and non-predictability. Under such conditions, what a firm “is” cannot be defined independently of what it has already done and what its rivals are doing. Capabilities are not intrinsic properties; they are historically accumulated and continuously reshaped patterns of interaction.

Diagnostic implication: substance metaphysics tends to frame strategy as essence-defence under a presumption of stable identity, whereas complex adaptive environments repeatedly reward adaptive reconfiguration and penalise rigidity. The strategic difference is therefore not merely managerial preference; it is the practical consequence of what is assumed to exist (stable essences) versus what must be navigated (historically unfolding trajectories under interaction).

Firm-level translation

Kodak and Netflix illustrate this ontological divide. Kodak’s engineers invented the first digital camera in 1975 (Sasson’s CCD-based prototype), but management defined Kodak as a film company – a firm with a fixed essence. Digital imaging was interpreted as a threat to that essence rather than as a possible evolutionary path. Strategic decisions therefore focused on defending chemical film rather than transforming the organisation’s trajectory. Kodak’s bankruptcy in 2012 was not caused by missing digital technology; it was caused by refusing to let the firm’s identity change.

Netflix faced comparable technological uncertainty but treated its initial configuration as provisional rather than essential. Founded in 1997 as a DVD-by-mail service, Netflix did not treat DVDs as its essence. When streaming became feasible, it reconfigured itself into a streaming platform (2007). When licensing became risky, it reconfigured again into a content producer (*House of Cards*, 2013). Later it became a global studio, distribution network, and data-driven entertainment ecosystem. Reed Hastings captured the underlying stance explicitly: “We named the company Netflix, not ‘DVD-by-Mail,’ because we knew we’d evolve.” By 2024, Netflix had over 270 million subscribers and \$33.7 billion in revenue – an outcome better explained by iterative reconfiguration than by defending a fixed “core.”

Ontological lesson (diagnostic): Kodak treated identity as substance; Netflix treated identity as historically produced and revisable. The consequences were existential.

Kodak and Netflix faced similar technological forces. They differed in how they understood what a firm is. Kodak treated identity as substance. Netflix treated identity as process. The consequences were existential.

Transformation

Table 100: Firm-Level Implications of the Ontological Shift: Reconceptualising Strategy, Capabilities, Innovation, Stability, and Failure from Substance to Process

FROM (Substance)	TO (Process / CAS)
Strategy = essence discovery	Strategy = trajectory management
Capabilities are intrinsic	Capabilities are historically produced
Innovation threatens identity	Innovation constitutes identity

FROM (Substance)	TO (Process / CAS)
Stability is normal	Change is normal
Failure is deviation	Failure is evolutionary signal

- Substance ontology leads firms to defend what they are. Process ontology enables firms to become what they need to be.
- Under substance thinking, strategy tends to emphasise identity discovery and defence, treats change as disruption, and presumes stable competencies.
- Under historical/processual interpretations, strategy is understood as trajectory management under uncertainty, treats change as constitutive, and emphasises capability evolution.

4.2.2 The Transformation Stakes

What Substance Ontology Precludes

The Being/Becoming choice is not merely abstract metaphysics; it determines what economics is allowed to see. Substance ontology treats identity as essence, change as deviation, time as reversible parameter, and history as largely irrelevant. Process-oriented interpretations treat identity as trajectory, change as constitutive, time as irreversible, and history as causally active. These are not stylistic differences: they determine whether phenomena such as innovation, learning, adaptation, and path-dependence can be represented as more than exogenous disturbances.

Being vs Becoming: Ontological Commitments and What They Make Visible

Table 101: Being Versus Becoming Across Six Dimensions: How Substance and Process Ontologies Shape the Visibility of Firm Phenomena

Dimension	Substance Ontology (Being)	Process / Historical Emphasis (Becoming)
Identity	Essence (firms as production functions)	Trajectory (firms as evolving organisations)
Change	Deviation (from equilibrium)	Constitutive (not deviation)
Time	Parameter (reversible)	Irreversible (path-dependent)
History	Irrelevant (ergodicity)	Causally active (non-ergodic)
Innovation	Exogenous (arrives from outside)	Endogenous (emerges from interaction)
Adaptation	Adjustment (to shocks)	Inherent (continuous)

The mechanistic paradigm – equilibrium, optimisation, representative agents, comparative statics – rests on substance ontology. Importantly, the analysis here is not that one can “add” complexity to a substance-based framework by appending frictions or heterogeneity. Rather,

once becoming, irreversibility, and endogenous novelty are treated as constitutive, the substance-based premises that legitimise timeless modelling become questionable at the root

Why This Matters Beyond Theory

Substance ontology structures real-world policy and management assumptions:

- Macroeconomics assumes stable Phillips curves and fiscal multipliers – relationships that may dissolve when economies evolve.
- Financial regulation assumes stationary risk distributions – assumptions violated under feedback, contagion, and innovation.
- Corporate strategy assumes persistent competitive advantage – an assumption repeatedly eroded by evolutionary dynamics.

The 2008 financial crisis was not only a failure of calibration; it exposed the mismatch between substance-based modelling habits and processual, interaction-driven reality.

The Meso-Level Consequence

Substance ontology naturally supports only two analytical levels:

- Micro – individual substances
- Macro – aggregates of substances

There is no logical place for a meso-level if relationships are treated as secondary attributes of substances rather than as constitutive structures. Yet stabilised structures exist precisely between individuals and aggregates: networks, routines, institutions, and coordination regimes. These are not sums of agents; they are organised patterns of interaction.

Why This Is Decisive for the Theory of the Firm

Firms are neither micro agents nor macro aggregates in the sense presupposed by substance ontology. Capabilities, culture, routines, governance, and innovation emerge at an intermediate organisational level where interactions stabilise into patterns. Substance ontology tends to erase this level by construction; historical and processual interpretations render it visible. This helps explain why mainstream economics repeatedly mis-theorises firm behaviour when it relies on timeless essences rather than temporally unfolding organisation.

4.2.3 Counter-Arguments Addressed

1. “Firms Do Have Stable Identities – They Persist Through Change”

Objection: Google has remained “Google” for over 25 years despite massive changes in employees, technology, and markets. Legal entities persist; brand identities endure;

organisational cultures maintain continuity. This persistence suggests some form of stable essence, not mere flux.

Response: This objection conflates pattern stability with substantial persistence. Rescher (1996, pp. 35–37) distinguishes substantial persistence (the same thing remaining) from processual continuity (the same pattern being continuously regenerated). A river persists not because its water is the same, but because the pattern of flow, banks, and gradient reproduces itself through flux. Stacey (2011, pp. 286–288) applies this directly to organisations: identity is continually reproduced in self-organising conversational processes. Tsoukas and Chia (2002, pp. 570–572) show that organisations do not occasionally change – they are change, temporarily stabilised into recognisable forms. Google persists because its coordination routines, cultural practices, technological platforms, and strategic orientations reproduce themselves over time – not because an invariant essence underlies them. When Yahoo! retained its legal entity and brand but lost its organising patterns of coordination and innovation, it ceased to be “Yahoo!” in any economically meaningful sense. Juridical continuity is not ontological continuity. Pattern persistence through process is not substance persistence through change.

2. “Substance Ontology Enables Prediction – Science Requires Stable Entities”

Objection: Scientific explanation requires stable entities with lawful behaviours. Physics treats planets as substances; economics needs rational agents and equilibrium-seeking markets. Process ontology, by denying stable entities, undermines scientific economics.

Response: This objection confuses domain-appropriate ontology with universal ontology. Cartwright (1999, pp. 24–27) shows that laws apply to nomological machines – closed, controlled systems with stable boundary conditions. Physics can treat planets as substances because gravitational systems approximate nomological machines. Economic systems do not: they are open, reflexive, and endogenous. Lawson (2003, pp. 163–165) demonstrates that econometric instability reflects the processual openness of social reality, not poor modelling. Boulding (1956, pp. 202–204) showed that as systems become reflexive, prediction collapses – agents change behaviour in response to being modelled. Complexity science strengthens this: even deterministic systems can be unpredictable (chaos), and adaptive systems can be structurally non-forecastable because their rule-sets evolve. Process ontology does not abandon science; it reframes explanation from prediction of states to explanation of trajectories. Darwinian biology explains evolution without predicting species; complexity-oriented economics can explain dynamics without presuming equilibrium prediction.

3. “Process Ontology Is Too Vague – We Need Precise Categories”

Objection: Substance categories allow mathematical precision. Process categories are vague and resist formalisation.

Response: This objection mistakes tractability for precision. Process categories are empirically precise insofar as they describe observable, reproducible patterns of behaviour. Nelson & Winter (1982, pp. 14–21) define routines as “regular and predictable behavioral patterns of firms.” Teece et al. (1997, p. 516) define dynamic capabilities as “the ability to integrate, build, and reconfigure competences.” Pentland & Feldman (2005, pp. 795–797) distinguish ostensive

(abstract) and performative (enacted) aspects of routines, enabling rigorous measurement. What substance categories often offer is not ontological precision but closed-form solvability. Arthur (2014, p. 26) shows that complexity economics sacrifices closed-form solutions to reveal mechanisms. That is not vagueness; it is a trade-off in favour of causal realism.

4. “Even Process Philosophy Needs Something That Persists”

Objection: If everything flows, nothing exists. Even Heraclitus still spoke of “the river.”

Response: Process ontology never denies continuity; it redefines it. Whitehead (1929, pp. 22–24) distinguished actual occasions from nexus (patterns of connected events). What persists is not substance but processual pattern. Rescher (1996, pp. 51–53) calls this genetic continuity – each phase arises from the previous. Emirbayer (1997, pp. 287–289) shows that in social systems, entities exist only as ongoing relational transactions. Firms persist because today’s coordination produces tomorrow’s coordination. Identity is not what remains unchanged; identity is what keeps getting reproduced.

Conclusion

This subsection has established, through five convergent arguments, that substance ontology is structurally incompatible with the reality of complex adaptive systems:

1. **It eliminates time, history, and becoming**, rendering irreversibility, path-dependence, and learning conceptually invisible.
2. **It provides no account of novelty**, forcing innovation, institutional change, and technological evolution to be treated as exogenous shocks.
3. **It misrepresents firms as static entities**, thereby erasing coordination, adaptation, and collective meaning-making.
4. **It imposes teleological frameworks** that are unsuited to historically contingent, competitive, and evolutionary social phenomena.
5. **It generates practical failure**, producing strategic rigidity, innovation blindness, and managerial misdiagnosis.

This deepens and completes the critique developed in Section 4.1. Reductionism fails because it decomposes what is irreducibly relational. Substance ontology explains why such decomposition seemed reasonable in the first place.

Together, Sections 4.1 and 4.2 establish the full metaphysical transformation required for complexity economics:

Table 102: Six Core Conceptual Shifts from Being to Becoming: The Ontological Transformation from Substance to Process

FROM (Being / Substance)	TO (Becoming / Process)
Stable essence	Dynamic pattern
Timeless identity	Temporal trajectory
Reversible time	Irreversible arrow
Ergodic processes	Non-ergodic evolution
Exogenous change	Endogenous novelty
Change as deviation	Change as constitutive

The Aha Moment

The choice confronting economics is not technical but ontological: Parmenides or Heraclitus? Being or Becoming?

Firms are not things that change. They are processes that temporarily stabilise.

The 2,500-year dominance of Parmenidean metaphysics has rendered economics blind to the phenomena that define economic reality – innovation, evolution, path-dependence, institutional emergence, and adaptive transformation.

The choice confronting economics here is not merely technical. It is ontological: whether to treat economic entities as stable substances whose identities precede time, or as historically unfolding patterns whose identities are reproduced through time. The analysis above indicates that for firms and markets operating under complex interaction and irreversibility, substance ontology systematically renders decisive phenomena invisible. A defensible diagnosis therefore requires taking becoming seriously – not as metaphor, but as an ontological condition for representing the economic world.

4.3 The Atomistic Foundation: Elements Before Relations

Introduction

This subsection completes the chapter's ontological problematisation by challenging atomism – the metaphysical assumption that reality consists of discrete, independent, self-contained units endowed with intrinsic properties prior to any relations. In its classical formulation in Democritus (5th century BCE), atomism holds that fundamental entities exist before and independently of their relations, and that relations merely connect what already exists. As Robinson (2024) and Rescher (1996) emphasise, this commits analysis to a world of pre-given entities rather than relationally constituted processes.

In economics, this metaphysical stance appears as methodological individualism (MI) – the doctrine that social and economic phenomena must be explained by reducing them to the properties, preferences, and choices of autonomous individuals. As Hodgson (2001, p. 28) formulates it, under MI: “The social whole is explained entirely in terms of the properties of its constituent individuals... the individual is taken as the elemental, primary unit in the social universe.”

This subsection serves three functions within the chapter's diagnostic architecture. First, it completes the triad of mechanistic ontology exposed in Sections 4.1.1–4.1.2: reductionism decomposes; substance ontology freezes; atomism isolates. Together they generate a worldview of decomposable, timeless, independent units – misaligned with the relational, interaction-driven character of complex adaptive systems. Second, it clarifies why network and meso-level analysis is not an optional add-on: complex dynamics arise not from isolated agents but from patterns of interaction – feedback loops, positional dependencies, and relational architectures. Atomism removes these causal structures before analysis even begins. Third, it explains why methodological individualism tends to collapse analysis into only micro and macro levels, erasing the intermediate relational layer where coordination, contagion, and emergence often occur.

Before proceeding, a powerful intuition must be acknowledged. Atomism can seem obvious: individuals exist before societies, firms before markets, atoms before molecules. But complex systems frequently reverse the direction of explanation. In many economic contexts, relations constitute capacities and identities: individuals, firms, and institutions function as positions within evolving networks rather than as autonomous atoms carrying intrinsic properties.

The stakes extend beyond theory. If atomism were correct, policy could target individuals and firms directly – via incentives, capital requirements, or taxes – and aggregate outcomes would reliably follow. But if relational causality is decisive, then outcomes depend critically on network structure: who is connected to whom, through what channels, with what feedback loops. The 2008 financial crisis made this tension visible: even where individual institutions appeared to satisfy regulatory ratios, systemic fragility propagated through interbank networks and feedback dynamics. Atomistic regulation confronted relational reality – and failed to anticipate contagion.

4.3.1 The Atomistic Worldview in Economics

Elements Before Relations

Atomism is not a discovery of science; it is a metaphysical inheritance. From Democritus's indivisible atoms to modern methodological individualism, it asserts that elements exist prior to relations, possess intrinsic properties, and interact externally.

Mainstream economics institutionalises this metaphysics through three linked assumptions: agents are autonomous; preferences are intrinsic; interaction is secondary. As Kirman (2016, pp. 539–540) observes, methodological individualism holds that: “The only way to understand the functioning of the whole is to build on the foundations of the behavior of the individual human beings who make it up.” Yet, he continues, this is precisely the theory's Achilles' heel: “We have never been able to specify the mechanism that would engender such order.” Fontana (2010, p. 14) makes the ontological commitment explicit: orthodox economics assumes that system components can be analysed independently, treating interaction as disturbance rather than constitutive structure.

This subsection completes the mechanistic triangle developed in Sections 4.1.1–4.1.2. Atomism posits independent elements; substance ontology treats these elements as stable essences; reductionism legitimates decomposition and recombination. The commitments are mutually reinforcing: as long as entities are treated as independent and stable, decomposition appears natural and relational causality becomes easy to ignore.

4.3.2 Four Arguments Against Atomism

Argument 1: Atomism Erases Relational Causality

Complex adaptive systems invert atomistic ontology. In many economic settings, relations are not optional connectors between pre-existing entities; they are the causal medium through which entities acquire properties, capabilities, and behavioural dispositions. Atomism fails because it treats relations as secondary, when in fact they are often constitutive of identity and performance.

This can be seen in three phenomena that atomism systematically misrepresents:

1. Embeddedness Over Autonomy

Economic action is embedded in networks of ongoing social relations. Granovetter's (1985, p. 487) classic critique of atomistic economics established that: “Actors do not behave or decide as atoms outside a social context, nor do they adhere slavishly to a script written for them by the particular intersection of social categories that they happen to occupy.” Network ties shape what opportunities are encountered, what information is received, and what norms constrain or enable behaviour. In complex systems terms, behaviour is conditioned by position and feedback, not by intrinsic preference alone. What appears to be an individual property –

trustworthiness, bargaining power, “skill” – often depends on relational embedding. Remove the network, and the “property” can dissolve.

2. Network Position Constitutes Advantage

Burt’s (1992, pp. 30–32) theory of structural holes demonstrates that advantage arises from position within a network, not from intrinsic traits. Actors who bridge disconnected groups gain access to non-redundant information, enabling brokerage, innovation, and arbitrage. The same individual, placed in a different position, can exhibit different “capabilities.” Creativity, deal-making, and opportunity recognition are therefore often positional effects – emerging from topology and access rather than from fixed internal attributes.

3. Lock-In as Relational Phenomenon

Arthur’s (1989, pp. 116–131) theory of increasing returns shows that technology value depends on how many others adopt it. Under network externalities, small early advantages amplify through positive feedback, producing lock-in. QWERTY, VHS, Windows, and dominant platforms do not win solely because of intrinsic superiority but because adoption networks reinforce their position. Holland (1995, p. 11) captures this precisely: “The interactions, not the actions of the individual agents, are the primary source of the complexities.” Technology “quality” is therefore not purely intrinsic; it is often emergent from network structure.

What these phenomena have in common

Embeddedness, structural holes, and network externalities point to the same ontological implication: what appears intrinsic is frequently relational. Entrepreneurial “talent,” venture “skill,” and technological “quality” can be position-dependent within evolving networks. Atomism misattributes relational causes to internal characteristics, producing systematically distorted explanations.

Firm-level translation: Sequoia Capital

Sequoia Capital’s success – investments in Apple, Google, PayPal, YouTube, Instagram, WhatsApp, Airbnb – is often attributed to partner brilliance and superior judgement. But a purely atomistic account misses the positional dimension of advantage. Sequoia functions as a central hub in Silicon Valley’s venture network, granting access to deal flow, information, syndication, and talent that peripheral investors cannot match. Entrepreneurs seek Sequoia not just for capital but for network effects: credibility, introductions, follow-on funding, and ecosystem access. Board seats generate information spillovers across the portfolio; university ties and alumni networks strengthen founder pipelines. These are relational assets, not intrinsic ones. The same partners operating from a peripheral position would predictably perform differently – not because intelligence changed, but because network position changed. By 2024, Sequoia’s portfolio companies had generated over \$1 trillion in market value – an outcome plausibly shaped by network centrality as well as judgement.

Diagnostic summary:

Table 103: Diagnostic Summary of Atomistic Assumptions: Relational and Network Phenomena Forced Out of Mainstream Theory of the Firm

Atomistic Assumption	What it forces theory to miss
Properties are intrinsic	Position-dependent capabilities and constraints
Agents are autonomous	Embeddedness, feedback, and network conditioning
Capability is internal	Structural advantages (e.g., brokerage, centrality)
Causation is individual	Relational causation via topology and interaction

Argument 2: Methodological Individualism's Systematic Failures

The Composition Fallacy

Methodological individualism assumes that if each agent behaves rationally, aggregate outcomes will be rational as well. Complex adaptive systems violate this structurally: interaction among agents generates emergent dynamics that cannot be inferred from individual intentions. Keynes's paradox of thrift, Schelling's segregation model, bank runs, traffic congestion, and coordination failures all illustrate the same pattern: micro-rationality can produce macro-irrationality.

This is not primarily because agents err, but because interaction structures transform incentives and outcomes. Gallegati and Gallegati (2025, p. 3065) diagnose the assumption directly: "It is implicitly assumed as valid that the cause-effect relations, the methodological individualism, and the dynamic of a representative element are equal to the aggregate behavior." But this assumption collapses in systems with heterogeneity, interaction, and feedback. Fontana (2008a, p. 10) is explicit: "It is impossible to study the properties of complex systems by observing its elements in isolation." Complex adaptive systems are not additive; they are non-linear, feedback-driven, and path-dependent. Aggregation erases the mechanisms that generate outcomes.

The Missing Meso-Level

Methodological individualism jumps directly from micro (individuals) to macro (aggregates), eliminating the intermediate scale where much economic causation occurs: the meso-level of networks, firms, institutions, and ecosystems. Yet complex systems accounts emphasise that coordination operates at organisational levels, contagion propagates through networks, innovation arises from inter-firm recombination, and stability/instability can emerge from institutional coupling. These are neither purely micro nor purely macro phenomena; they are meso-level dynamics. Atomism does not merely misrepresent this level – it renders it ontologically invisible.

Firm-Level Translation: Silicon Valley

Why Silicon Valley?

Not because California has superior individuals (micro), and not because the U.S. has superior macro conditions (macro), but because a meso-level innovation ecosystem exists.

That ecosystem includes:

- Stanford and Berkeley knowledge spillovers
- Sand Hill Road venture capital concentration
- serial entrepreneur networks (e.g., PayPal Mafia)
- specialised law firms, recruiters, accelerators
- labour mobility and reputational markets

When Fairchild Semiconductor’s “traitorous eight” left in 1957, they did not just create Intel – they generated a self-reinforcing industrial network that produced AMD, National Semiconductor, and dozens of spinoffs.

By 2024, companies founded by Fairchild alumni exceeded \$2 trillion in market capitalisation.

That value was not located in any individual. It was located in the meso-level network structure. Place the same engineers in a different regional topology – without capital, mobility, and recombination – and the outcome collapses. Silicon Valley is not a collection of brilliant people. It is a relational production system.

Social Construction Is Invisible Under Atomism

Methodological individualism treats preferences, beliefs, and norms as intrinsic individual attributes. But in CAS, these are endogenously generated by interaction.

Veblen (1899) showed consumption is positional – your utility depends on what others consume. DiMaggio and Powell (1983) showed firms converge not by optimisation but by networked institutional pressure. Preferences, strategies, and expectations are co-evolving social patterns – not private endowments. Atomism mistakes products of interaction for properties of individuals.

Ontological Transformation

Table 104: Ontological Transformation from Atomistic Methodological Individualism to CAS Relational Ontology: Five Foundational Shifts

Atomistic MI	CAS Relational Ontology
Micro → Macro	Meso → Macro
Composition	Emergence
Individual rationality	System dynamics

Atomistic MI	CAS Relational Ontology
Preferences as given	Preferences as socially produced
Independent agents	Network-embedded agents

Bottom Line

Methodological individualism fails not empirically but ontologically.

It deletes the very scale – the meso-level – where coordination, contagion, innovation, and stability are generated. CAS economics therefore does not “add networks” to MI. It replaces atomism with relational causation.

Argument 3: Firms Are Not Atomic Units

Organisational Embeddedness

Firms are not self-contained entities. They are structurally embedded in inter-organisational networks of suppliers, customers, partners, regulators, and knowledge communities. These networks do not merely constrain behaviour; they help constitute what firms are able to do. Capabilities depend on:

- who a firm partners with (alliances)
- who supplies it (value-chain position)
- who it learns from (knowledge networks)
- who sets standards and norms (industry forums)

Powell et al. (1996, p. 165) showed that research networks enable “the integration and cross-pollination of varied perspectives,” making innovation a network outcome rather than a purely internal trait. In biotechnology, firms embedded in dense collaboration networks systematically outperform isolated firms – not because they are inherently “better” in isolation, but because network position grants access to complementary knowledge, resources, and recombination opportunities. Uzzi (1997, pp. 37–39) demonstrated that firms with embedded ties – relationships characterised by trust, fine-grained information exchange, and joint problem-solving – outperform firms relying solely on arm’s-length market transactions, even when controlling for firm-level characteristics. Performance differences therefore often arise from relational architecture, not atomic firm attributes.

Boundaries Are Permeable

Atomism assumes firms are bounded, self-contained units interacting through contracts. Yet modern capitalism operates through permeable organisational membranes. Open innovation, modular production, and platform ecosystems increasingly dissolve clear inside/outside distinctions:

- Where does Apple end and app developers begin?

- Where does Toyota end and its supplier *keiretsu* begin?
- Where does Amazon end and its marketplace sellers begin?

Chesbrough (2003, pp. 43–45) documents the shift from closed innovation (internal R&D monopoly) to open innovation (knowledge flows across organisational boundaries). Value creation increasingly occurs in inter-firm architectures rather than inside isolated corporations. The firm is therefore not well represented as an atomic unit; it is a node in a production and innovation network whose boundaries are strategic, negotiated, and continually reconfigured.

Identity Is Relational

Firm identity is not intrinsic. It is defined by relative position. Porter (1980, pp. 3-5) showed that strategy is fundamentally relational positioning. What counts as “low cost,” “premium,” or “differentiated” depends on competitors, suppliers, customers, and industry structure. Walmart’s identity as a low-cost retailer only exists because other retailers occupy higher-cost positions. Remove the competitive network and the identity collapses. Firms are therefore role-bearers in relational fields, not self-defining substances.

Firm-Level Translation: Boeing’s 787

The Boeing 787 Dreamliner programme illustrates the risks of treating firms as atomic units connected only by contracts. Boeing outsourced roughly 70% of the aircraft to a global supplier network: wings to Mitsubishi and Kawasaki (Japan), fuselage sections to Alenia (Italy) and Spirit (Kansas), engines to Rolls-Royce and GE. The atomistic assumption was that Boeing could specify requirements, suppliers would deliver components, and Boeing would assemble them – independent units coordinated by contract.

The operational reality was relational: component mismatches, incompatible software, quality failures, and battery problems propagated through the network. Delays cascaded, integration failed, and Boeing incurred very large costs. The failure was not located in any single supplier; it was located in the coordination topology of the production network. Contracts could not substitute for relational alignment, learning, and mutual adjustment. Subsequent moves to bring subsystems back in-house can be interpreted as boundary reconfiguration – a choice about network structure rather than an expression of a fixed, intrinsic firm identity.

Firm-Level Translation: Linux

Linux exposes a deeper limitation of atomism. Linux is not produced by any single firm, yet it powers a dominant share of servers, supercomputers, and Android-based devices. Its capabilities are generated by a distributed development network comprising thousands of contributors across companies (IBM, Google, Intel, Red Hat, Microsoft) and independent programmers. What makes Linux powerful is not ownership or hierarchy but network-level causal structures: distributed version control, peer-review norms, modular code architecture, licensing regimes, and global developer communities. The productive unit here is not a firm in the atomistic sense; it is a relational production system. Even major corporate moves (e.g., IBM’s acquisition of Red Hat) can be interpreted as acquiring a privileged position within this network rather than “owning” the underlying system.

Diagnostic summary:

Table 105: Diagnostic Summary of Atomistic Presumptions About Firms and the Relational Phenomena They Mis-Specify

Atomistic presumption about firms	What becomes mis-specified
Independent firms	Network-embedded nodes
Clear, fixed boundaries	Permeable interfaces and boundary reconfiguration
Intrinsic capabilities	Position-dependent capabilities and access
Identity as internal property	Identity as relational role in a field
Contracts coordinate sufficiently	Networks coordinate through alignment, learning, and feedback

Bottom line: Firms do not exist before relations; they exist in relations. Treating firms as atomic units erases the network structures that generate innovation, resilience, and failure.

Argument 4: Networks, Not Atoms, Are Primary Units

From Entities to Relations

Complex adaptive systems require a reversal of the mechanistic order of explanation: relations are often primary and entities secondary. Nodes do not have explanatory meaning prior to their connections; their capacities and constraints are shaped by their relational embedding. A person's identity is partly a function of position in social networks; a firm's identity is partly a function of location in production, knowledge, and competitive networks; a technology's value depends on its place in an adoption network.

Whitehead (1929, p. 22) formalised this shift: "The actual world is a process, and that process is the becoming of actual entities." Bourdieu's (1986, pp. 241–258) field theory similarly treats agents as positions in fields: capital, power, and identity arise from location in relational structures rather than intrinsic endowments. Newman (2018, p. 8) summarises network science in terms of "the primacy of interactions" rather than isolated node attributes.

Network Topology Determines Outcomes

Network science shows that structure is causation. Topology governs diffusion, contagion, robustness, and collapse. Small-world networks (short path lengths with clustering) enable rapid coordination and viral diffusion. Scale-free networks (power-law degree distributions) are resilient to random failure yet fragile to hub failure. These are system-level properties; they depend less on what any individual node "is" than on how nodes are connected.

Economic phenomena follow this logic: financial contagion, information cascades, innovation diffusion, and supply-chain fragility are governed by network topology as well as individual attributes. Haldane & May (2011, pp. 351–353) demonstrate this for finance: systemic importance is a network property. A relatively small bank occupying a hub position can be more dangerous than a large bank on the periphery. The 2008 crisis illustrates how interbank

topology can amplify local shocks into system-wide collapse. Atomistic regulation failed in part because it targeted nodes rather than network structure.

Firm-Level Translation: Visa & Mastercard

Visa and Mastercard provide a clean illustration of network primacy. They do not, in the main: issue cards, extend credit, hold deposits, or own merchant terminals. Banks and merchants perform those functions. Visa and Mastercard occupy central switching positions in a global payments network linking billions of cards, tens of millions of merchants, thousands of banks, and trillions of transactions. Much of Visa’s market valuation reflects the capitalised value of network centrality rather than the value of tangible assets. Network effects generate positive feedback loops: merchants accept Visa because customers carry Visa, and customers carry Visa because merchants accept Visa. This is not merely “market power” in an atomistic sense; it is a topological advantage. In many cases, node quality cannot defeat network position. The same technology operated by a firm without Visa’s network position would be worthless.

Network effects generate positive feedback loops:

- merchants accept Visa because customers carry Visa
- customers carry Visa because merchants accept Visa

This is not “market power” in the atomistic sense – it is topological dominance.

This is why fintech challengers with better technology fail: node quality cannot defeat network position.

Diagnostic synthesis:

Table 106: What Atomism Makes Invisible: Seven Presumptions and the Network-Dependent Phenomena They Obscure or Misattribute

Atomistic presumption	What becomes invisible / misattributed
Entities are primary	Relations and interaction as constitutive
Properties are intrinsic	Properties arising from network position
Value is internal	Value dependent on connectivity (network effects)
Firms are independent units	Firms as nodes in production/knowledge networks
Markets are collections of transactions	Markets as evolving interaction networks
Power belongs to actors	Power as centrality in networks
Risk belongs to agents	Risk as topology (contagion, hubs, cascades)

Bottom line: economic reality is not well represented as a collection of atoms. It is organised through relations.

4.3.3 The Transformation Stakes

What Atomism Precludes

Atomism is not a harmless simplification. It systematically eliminates structural causality. By treating agents as independent atoms with intrinsic properties, mainstream economics renders conceptually invisible:

- **Embeddedness** – how network position shapes opportunities, information, and feasible action (Granovetter)
- **Structural holes** – how brokerage positions generate advantages no intrinsic trait can explain (Burt)
- **Network externalities** – how value depends on adoption networks rather than technological properties (Arthur)
- **Meso-level causation** – how clusters, platforms, ecosystems, and institutions shape outcomes
- **Social construction** – how preferences, beliefs, and norms are produced through interaction
- **Strategic interdependence** – how outcomes depend on others' positions and actions, not isolated optimisation

These are not secondary details in complex adaptive systems. They are frequently the primary mechanisms through which coordination, diffusion, contagion, and innovation occur.

The mechanistic paradigm relies on atomism because independent atoms are what make optimisation, aggregation, and equilibrium tractable as closed-form problems. When agents are not independent – when properties are positional and outcomes are interaction-generated – the meaning of representative agents, separable optimisation, and micro → macro aggregation becomes unstable.

The cascade implication

Rejecting atomism has a cascading implication for the mechanistic toolkit:

- If properties are **positional**, they cannot be treated as purely intrinsic.
- If identity is **relational**, agents cannot be analysed in isolation without losing explanatory content.

- If the **meso-level** is causally primary, direct micro → macro aggregation becomes unreliable.
- If interactions constitute entities, decomposition risks destroying the mechanisms under study.

Diagnostic table:

Table 107: Diagnostic Table: How Relational Causality Undermines the Implicit Requirements of Core Mainstream Analytical Tools

Mechanistic tool	What it implicitly requires	What becomes problematic under relational causality
Optimisation	Independent agents	Strategic interdependence and feedback alter incentives endogenously
Representative agent	Homogeneity	Heterogeneous network positions and distributions become causally active
Aggregation	Additivity	Network structure and topology generate non-additive outcomes
Equilibrium	Separability	System dynamics and interaction-generated states dominate
Comparative statics	Isolation	Co-evolution undermines ceteris paribus stability

This is not an argument against formal modelling per se. It is an argument that certain forms of formalisation presuppose atomistic ontology and therefore inherit its blind spots.

Completing the Trinity

The Three Pillars of the Mechanistic Ontology

Table 108: The Three Pillars of the Mechanistic Ontology: Core Assumptions and the CAS Phenomena They Collectively Eliminate

Ontological error	What it assumes	What it eliminates
Atomism	Entities exist prior to relations	Networks, embeddedness, structural causation
Reductionism	Wholes = sum of parts	Emergence, meso-level dynamics
Substance ontology	Entities have fixed essences	Time, history, evolution, becoming

These three form a closed ontological system: atomism enables reductionism (independent parts can be decomposed); substance ontology stabilises atomism (atoms have fixed essences); reductionism reinforces substance (decomposed parts appear stable). The triad underwrites equilibrium, optimisation, representative agents, and prediction. Importantly, partial repairs often leave the underlying structure intact: if one pillar is relaxed while the others remain,

mechanistic reasoning tends to regenerate its core assumptions through the remaining commitments.

4.3.4 Counter-Arguments Addressed

“Individuals Are Real; Relations Are Abstract – We Can’t Build Science on Abstractions”

Objection. Individuals are concrete, observable entities. We can see people, count firms, and measure outputs. Relations are abstractions. Science must be built on concrete foundations, not abstractions.

Response. This objection confuses metaphysical priority with epistemic accessibility. Relations are not epistemically vague; they are measurable structures. Network science provides rigorous tools for identifying, measuring, and comparing relational patterns. Wasserman and Faust (1994, pp. 17–21) define operational measures of centrality, clustering, path length, and structural equivalence. Borgatti et al. (2009, pp. 892–894) show that network position variables explain variance in outcomes beyond individual attributes. Freeman (2004, pp. 2–4) documents the development of social network analysis into a cumulative, standardised scientific field. A firm’s betweenness centrality can be measured as empirically as its revenue; a bank’s interbank exposure is as observable as its capital ratio; a venture capitalist’s brokerage position is as quantifiable as education credentials.

Moreover, what we directly observe are events and interactions; what we infer are entities. Individual behaviour is frequently unintelligible without relational context. What this objection treats as “abstract” is often the causal structure that explains the concrete observations.

“Methodological Individualism Has Been Scientifically Successful – Why Abandon What Works?”

Objection. MI underlies price theory, game theory, and mechanism design – fields that have generated Nobel Prizes and practical applications. We should not abandon a successful research programme.

Response. This objection mistakes local success for general validity. MI works best when interactions are weak, symmetric, or effectively irrelevant – conditions under which network structure does not dominate outcomes. But the domains where economics most visibly fails are those where interactions are strong, heterogeneous, and structured: financial contagion, innovation diffusion, technology adoption, and institutional change.

Manski’s (1993, pp. 531–533) reflection problem shows that individual-level econometrics cannot identify social effects when outcomes depend on group behaviour. Durlauf and Young (2001, pp. 1–4) demonstrate that interaction effects generate identification failures MI cannot resolve. Jackson (2014, pp. 4–7) documents the empirical centrality of networks across labour markets, trade, finance, and innovation. Acemoglu et al. (2012, pp. 1978–1980) show that network structure amplifies and transmits shocks in ways individual-level models do not

capture. The 2008 financial crisis – in which individually rational banks contributed to systemic collapse – illustrates precisely the domain where atomistic assumptions underperform.

“Network Analysis Is Just MI Plus Interactions – Not a New Ontology”

Objection. Network analysis studies individuals and their ties. This is simply MI with interaction terms added. No ontological shift is required.

Response. This misunderstands what networks explain. Network science shows that topology itself has causal power. Watts and Strogatz (1998, pp. 440–442) show that small-world networks generate collective properties – rapid diffusion with local clustering – that cannot be reduced to node properties. Barabási and Albert (1999, pp. 509–511) show that scale-free networks exhibit systemic robustness and fragility determined by degree distribution, not individual traits. Granovetter’s (1973, pp. 1366–1368) weak ties result shows that diffusion depends on structural position, not merely on individual tie strength.

These are not “interaction effects” appended to individuals; they are system-level causal properties. Borgatti and Halgin (2011, pp. 1169–1171) are explicit: network theory is distinct because it treats structure itself as a causal agent. The analytical shift is from “nodes with interactions” to “structures that generate node outcomes.”

“Relational Ontology Leads to Infinite Regress”

Objection. If entities are constituted by relations, then explaining any entity requires explaining all relations, which requires explaining everything else. Atomism at least provides stopping points.

Response. This objection mistakes ontological priority for explanatory scope. Relational ontology does not require explaining everything; it requires identifying which relations are causally relevant for the phenomenon in question. Granovetter (1985, pp. 490–491) distinguished embedded from disembedded action. Network analysis provides boundary-setting tools: tie strength, modularity, community detection, brokerage positions, and structural equivalence. Abbott (1988, pp. 170–173) shows that all explanation is selective; relational explanation selects structural variables rather than intrinsic attributes. Emirbayer (1997, pp. 287–289) provides methodological guidance: begin with transactions and relations, not substances.

Atomistic explanation does not escape regress either: explaining individual traits requires explaining their formation, socialisation, and history. The difference is not bounded versus unbounded explanation but structural versus intrinsic explanatory targets.

Conclusion

This subsection has developed four convergent arguments indicating that atomism is incompatible with understanding complex adaptive systems:

1. Atomism erases relational causality, rendering embeddedness, structural holes, and network externalities difficult to represent.
2. It generates systematic methodological failures: the composition fallacy, the missing meso-level, and the invisibility of social construction.
3. It misrepresents firms as atomic units, thereby missing organisational embeddedness, boundary permeability, and relational identity.
4. It inverts the causal order by treating relations as secondary when, in many domains, network structure constitutes capacities and outcomes.

Together with Section 4.1.1's critique of reductionism and Section 4.1.2's critique of substance ontology, this subsection supports a single diagnostic conclusion: the mechanistic paradigm rests on three mutually reinforcing assumptions – decomposability, temporal stasis, and atomistic independence – that systematically obscure phenomena central to complex economic reality.

Ackoff (1999, pp. 16–18) captured the methodological implication: systems thinking requires understanding parts through their role in containing wholes (synthesis), not understanding wholes through isolated parts (analysis). Atomism makes analysis appear natural; a relational diagnosis clarifies why synthesis is often required to preserve the causal mechanisms under study.

The illustrations developed here – Sequoia Capital's network-dependent performance, Boeing's supply-network coordination failures, Silicon Valley's ecosystem dynamics, Visa's network centrality, Linux's distributed production, and the 2008 crisis's contagion logic – are structurally consistent. In each case, explanations focused only on intrinsic properties misattribute causation that resides in relational structure. Network position and topology repeatedly explain variance beyond individual capability; production-network coordination shapes failure beyond component quality; ecosystem density shapes innovation beyond individual talent; payment-network centrality shapes platform value beyond firm assets; distributed collaboration shapes productivity beyond hierarchical boundaries.

The Aha Moment

The core issue is ontological: whether economic properties are treated as intrinsic to isolated units or as positional within relational structures. If properties are intrinsic, agents can be analysed in isolation. If properties are positional, explanation must foreground network structure. The accumulated evidence indicates that once networks exist – and once outcomes depend on topology – strict atomism becomes unstable as a foundational ontology.

Table 109: Atomistic vs Relational Commitments

Atomistic commitment	What it tends to obscure
Independent entities	Networked nodes and positional dependence
Intrinsic properties	Position-dependent properties
Interactions secondary	Relations as constitutive
Clear firm boundaries	Permeable interfaces
Micro→macro aggregation	Meso-level structural mechanisms
Isolated optimisation	Strategic interdependence

Table 110: Reductionist vs Emergent Ontologies

Dimension	Reductionist ontology	Emergent ontology
Source of properties	Components	Relations and interactions
Causality	Bottom-up only	Multi-level, bidirectional
Explanation	Decomposition	Organisation and pattern
Predictability	High in principle	Limited by emergence
Failure mode	Omitted system effects	More robust to complexity

4.4 The Limits of Analytical Decomposition: When Analysis Destroys What It Studies

Introduction

This subsection completes Section 4.1's problematisation by examining the methodological consequences of the ontological errors established in Sections 4.1.1–4.1.3. Analytical decomposition is the methodological practice of breaking systems into constituent parts, studying each in isolation, and recombining the results to explain the whole. While this approach is powerful for what Poli (2013, p. 144) terms “complicated” systems – systems whose identity is preserved under structural decomposition into “disjoined structural parts and their relations” – it fails categorically for complex adaptive systems. As Nason (2017, p. 16) states, “complex systems cannot be separated into parts to be analyzed independently because their components are interconnected and interdependent.”

This subsection serves three functions within the chapter's internal architecture. First, it provides methodological closure to Sections 4.1.1–4.1.3 by showing that mechanistic ontology (reductionism + substance + atomism) naturally generates a mechanistic method (analytical decomposition) that is structurally mismatched to CAS. Second, it formalises the analysis–synthesis distinction that becomes necessary for the reconstruction in Section 4.2. Third, it clarifies why methodological success in decomposable domains can become methodological failure when universalised to non-decomposable domains.

The subsection proceeds as follows. Section I formalises the distinction between analysis and synthesis. Section II develops four arguments showing why analytical decomposition fails for CAS. Section III examines boundary conditions and implications. Section IV addresses counter-arguments. Section V synthesises the methodological lesson.

A crucial clarification is required at the outset. The critique is not anti-analysis. Analytical decomposition works when applied to systems that are genuinely decomposable. The error lies in universalising analysis to systems for which it is ontologically inappropriate. As Fang and Casadevall (2011, p. 1401) emphasise, the analytical method itself does not deny emergence; rather, reductionism elevates isolated component analysis into a metaphysical principle, thereby suppressing relational and emergent properties.

The stakes are not merely philosophical but practical. If systems are decomposable, interventions targeting individual components should yield predictable aggregate effects. If systems are non-decomposable, component-level interventions can provoke unintended system reorganisation. Financial regulation illustrates this sharply: capital requirements imposed on individual institutions proved insufficient to prevent systemic crisis because systemic risk often resides in the network of interconnections, not in the balance sheets of isolated nodes.

4.4.1 Analysis versus Synthesis: A Framework

The methodological choice

The mechanistic paradigm privileges analysis – the decomposition of systems into independently examinable parts. As Ackoff (1999, p. 1) documents, this worldview emerged from the Renaissance and the Industrial Revolution, when rational order, predictability, and control became guiding ideals. Newman (2019, pp. 64–68) formalises the analytical procedure as a three-step sequence:

1. **Deconstruction** – breaking the system into constituent parts
2. **Isolation** – understanding each component in separation
3. **Recomposition** – inferring the behaviour of the whole from its parts

This sequence works only if a critical assumption holds: parts retain their defining properties when isolated. For complex adaptive systems, this assumption fails. Components in CAS are not like machine parts – removable, replaceable, and functionally invariant – but like organs in a living body: interdependent, co-constituted, and functionally altered by separation. Remove a heart from a body and it is no longer a heart in the biological sense; remove a firm from its production, knowledge, and institutional networks and it is no longer the same firm economically.

Ackoff (1999, pp. 16–18) expresses the point with precision: “A system is a whole that cannot be divided into independent parts. The essential properties of a system are properties of the whole which none of its parts have.” If this is correct, then analytical decomposition does not merely miss systemic properties – it can destroy them. What is most important about a complex system can disappear at the moment of dissection.

Logical cascade from ontology to method

This subsection completes the internal logic of the mechanistic critique developed in Sections 4.1.1–4.1.3:

- Section 4.1.1 showed that reductionism fails ontologically: decomposition destroys emergent properties.
- Section 4.1.2 showed that substance ontology makes decomposition appear legitimate: if entities have stable essences, they should survive being taken apart.
- Section 4.1.3 showed that atomism makes decomposition appear feasible: if entities are independent atoms, separation should not alter their nature.

The inference is methodological: when these assumptions are false, the analytical method becomes systematically unreliable for CAS. Mechanistic ontology encourages mechanistic method, and mechanistic method produces systemic blindness in domains where coupling, feedback, and emergence dominate.

Mechanistic ontology → mechanistic method → systemic blindness.

This is not a technical limitation; it is a structural mismatch between ontology and method.

4.4.2 Four Arguments Against Analytical Decomposition

Argument 1: The *ceteris paribus* fallacy

Analytical decomposition relies on *ceteris paribus* reasoning: vary one variable while holding all others constant to identify its causal effect. This presupposes independence – that X can be changed without altering Y. In complex adaptive systems, this presupposition is false. Variables are endogenously coupled through feedback loops: a change in X alters Y, which alters Z, which in turn alters X. Under such conditions, “all else equal” is either logically impossible (because variables cannot remain fixed) or methodologically destructive (because freezing the couplings eliminates the very mechanisms that generate system behaviour).

Standard labour-market analysis illustrates the problem. Models vary wages while holding demand, productivity, and expectations constant to infer employment effects. But wages affect income; income affects demand; demand affects output; output affects labour demand; and labour demand feeds back into wages. Holding these variables constant breaks the causal circuit that governs actual labour-market dynamics. As Hayek (1945, p. 520) warned, economics errs by transferring to social phenomena “the habits of thought we have developed in dealing with the phenomena of nature,” treating interconnected processes as if they were independent quantities. What *ceteris paribus* analysis isolates is not the real system but an artefact produced by severing its feedback loops.

Firm-level translation: Spotify

Spotify’s platform economics illustrates why *ceteris paribus* reasoning fails in digitally mediated, feedback-driven firms. Spotify cannot meaningfully study the effect of “algorithm quality” while holding user behaviour constant, because recommendations shape listening behaviour, which generates the data on which the algorithm trains, which in turn alters future recommendations. Algorithm and behaviour are co-evolving.

When Spotify introduced Discover Weekly (2015), listening and exploration reportedly increased substantially. But this was not a unidirectional causal effect. Improved recommendations induced more exploratory listening, which generated richer behavioural data, which further improved recommendation accuracy, producing a positive feedback loop. As Spotify grew to hundreds of millions of monthly users and expanded its catalogue into tens of millions of tracks, its advantage depended increasingly on this self-reinforcing learning dynamic, not on any separable component.

Spotify’s subsequent expansion into podcasts intensified the coupling. Acquisitions such as Gimlet and Anchor (2019) integrated new content, new user behaviours, and new data streams into the same recommendation ecology. Podcast listening altered music discovery; music listening altered podcast recommendations. Social features (e.g., Blend) added further feedback between users, content, and algorithms. The platform’s value emerged from integration, not from an isolable module.

Competitors that treated the platform as decomposable – e.g., “algorithm provider” versus “content catalogue” – struggled to reproduce these coupled learning dynamics because the relevant causal mechanism is the system’s circularity.

Methodological implication: in CAS, independence is the exception, not the baseline. Where variables are endogenously coupled, isolating “the effect of X” by freezing the system is likely to remove the causal circuit that produces the observed outcomes.

The failure of *ceteris paribus* reasoning in CAS is not empirical but logical:

FROM independent variables → **TO** endogenously coupled dynamics

FROM linear causation → **TO** circular, feedback-driven causation

FROM isolation → **TO** system integration

Analytical decomposition presumes independence. Complex systems violate it. Therefore, decomposition does not reveal CAS behaviour – it destroys it.

Argument 2: Systemic properties collapse under analysis

Analytical decomposition presupposes that system properties can be recovered by studying components in isolation and recombining results. This presupposition fails for complex adaptive systems because many systemic properties exist only in simultaneous interaction. When the system is decomposed, the property of interest can cease to exist. As Nason (2017, p. 6) states, “a complex issue is one where one cannot separate out the parts and for which there are no rules, algorithms, or laws such as are found in physics.”

Complex systems exhibit relationally constituted properties – features that do not belong to any element taken alone but arise only in the organised whole. Cilliers (1998, p. 119) formalises this: meaning in complex systems “is not located in individual components, but is a function of the relationships between components.” Analytical isolation destroys those relationships and therefore destroys the phenomena under investigation.

Organisational culture

Organisational culture is a paradigmatic systemic property. It consists of shared meanings, behavioural norms, and interpretive frames that emerge through ongoing interaction. Culture is not contained in any individual employee; it exists only in patterns of communication, coordination, and mutual expectation. Interviewing employees in isolation does not reveal culture; it strips away the interactional medium in which culture exists.

Google’s engineering culture illustrates this. Its identity as innovative, data-driven, and risk-tolerant does not reside in any particular engineer. It arises from interactional practices – peer review, open discussion, internal mobility, and shared experimentation. Policies such as “20 percent time” mattered not because they increased intrinsic creativity in isolated individuals, but because they altered interaction topology – allowing ideas to circulate across projects and domains. Innovations (e.g., Gmail, Google News, AdSense) are intelligible as systemic interaction outcomes, not as properties of individuals in isolation.

Market dynamics

Financial markets exhibit the same logic. Momentum, herding, volatility clustering, and crashes are not properties of traders but of trading networks. As Haldane (2012, pp. 4, 18) observes, financial systems are “complex... interconnected... tightly coupled,” making them prone to cascades.

The 2010 Flash Crash is paradigmatic. In minutes, U.S. equity markets lost and then regained on the order of a trillion dollars in value. No trader intended this outcome. It emerged from algorithmic interactions reacting to each other in a feedback loop. Analysing any single algorithm reveals little; the phenomenon is relational and systemic.

Krugman (2009, p. 2) noted that economists had mistaken mathematical elegance for truth by ignoring “destabilizing feedback loops.” The limitation here is not lack of data but a decompositional method that removes the very dynamics that produce the event.

Firm-level translation: Enron

Enron's collapse demonstrates how systemic properties vanish under decomposition. At its peak (2000), Enron was valued at roughly \$70 billion; by late 2001 it was bankrupt. Analysts decomposed Enron into business units – energy trading, broadband, international projects, retail, and water – and assessed unit performance largely through accounting representations. Yet the firm was fatally fragile because its defining property – self-referential financial fabrication – existed in the interaction structure between units.

Special-purpose entities (e.g., LJM and the Raptors) linked divisions into circular transactions in which assets were sold to related entities, paper profits were booked, and debt was obscured. These profits existed only in the network of transactions, not in any single ledger. When Enron's share price fell, the relational structure collapsed: declining equity triggered SPE unwinding, which revealed hidden losses, which further depressed equity – an accelerating feedback loop. No unit “failed” in isolation; the system failed as an interconnected architecture.

Methodological implication: complex adaptive systems often do not have properties located in parts and recoverable by recomposition. Their defining features can exist only in organised interaction. Decomposition therefore risks eliminating the phenomenon at the moment of observation.

Complex adaptive systems do not have properties that can be located in parts and reassembled. Their defining features exist only in organisation.

FROM properties of components → **TO** properties of organised interaction

FROM sum of parts → **TO** emergence from structure

FROM individual behaviour → **TO** collective dynamics

Analytical decomposition does not reveal systemic properties in CAS – it eliminates them. Only synthesis – studying how components co-constitute one another – can explain the phenomena that matter.

Argument 3: The Indecomposability Principle

Some systems are not merely difficult to decompose; they are ontologically indecomposable. Their components do not exist as independently specifiable entities whose properties survive isolation. Fontana (2010, p. 14) captures this precisely: “it is not possible to derive implications on the behavior of the system's subcomponents by analyzing them separately.” This is not a technical limitation but a logical one: the parts of a complex adaptive system are constituted by the relations that bind them.

In decomposable systems (machines, linear production processes), parts retain their identity and function when removed. In complex adaptive systems, they do not. Remove an organ from a body and it ceases to function as an organ. Remove a trader from a market, or a firm from an industry network, and key elements of its economic identity – its opportunities, constraints, and strategic meaning – are altered. Indecomposability is therefore ontological, not merely epistemic.

Strategy is relational, not intrinsic

Porter (1980, pp. 3–5) showed that strategy is defined only in relation to rivals, customers, suppliers, and industry structure. A firm's “best strategy” cannot be specified in isolation because what counts as advantage depends on what others do. Strategy exists only in competitive space, where firms co-evolve. A firm has no strategy in isolation – only in relation

to others. Analytical decomposition that attempts to extract “Firm A’s strategy” independently of B, C, and D risks destroying the object of analysis by detaching it from the relational field that makes it intelligible.

Coordination is a system property

Hayek (1945, pp. 519–527) showed that firms exist to coordinate dispersed knowledge that no individual possesses. Coordination is not located in any single person; it emerges from communication structures, authority relations, routines, and shared expectations. Decompose the firm into individuals and coordination vanishes. Individuals do not “contain” coordination; they enact it only within organised relational systems. Coordination is therefore an irreducibly systemic phenomenon rather than an attribute of separable parts.

Firm-level translation: Southwest Airlines

Southwest Airlines provides a textbook illustration of indecomposability. For decades, Southwest achieved sustained profitability in an industry characterised by chronic losses. Competitors attempted to replicate Southwest by copying observable components: point-to-point routing; a single aircraft type; no meals or first class; rapid turnarounds; low fares. United (“Ted”), Delta (“Song”), Continental (“Lite”), and US Airways (“MetroJet”) copied these components and failed.

The failure is diagnostic. Southwest’s advantage does not reside in any practice individually. It resides in their systemic integration with:

- labour relations that allow flexible job roles
- a culture of mutual ownership (“we all pitch in”)
- customer expectations that accept trade-offs
- operational rhythms tuned to those expectations

Southwest’s famous rapid turnarounds require not only procedures but flight attendants who help clean cabins, pilots who assist with boarding, and passengers who know how to move quickly. None of these elements works alone. Each is defined by the others. As Kelleher put it: “Our esprit de corps is the core of our success.” Culture here is not a detachable variable but a relational glue that makes the operational system coherent.

Competitors copied parts but not the system because the system does not exist in the parts. Remove culture and the operations fail. Remove operational rhythms and the culture collapses. Remove customer expectations and the model disintegrates. The key methodological lesson is that some organisational advantages are conceptually inseparable: their components are mutually constitutive rather than independently transportable.

Ontological conclusion

Southwest cannot be decomposed into independent strategic components because its components are mutually constitutive. Remove culture, and the operations fail. Remove operations, and the culture collapses. Remove customer expectations, and the model disintegrates.

This is ontological indecomposability.

FROM practically complex → **TO** conceptually inseparable

FROM parts with independent properties → **TO** parts defined by relations

FROM isolation possible in principle → **TO** isolation logically incoherent

Analytical decomposition presumes that parts exist prior to their relations. Complex adaptive systems violate this assumption. Their parts exist only within the system that analysis destroys.

Argument 4: Feedback Loops Preclude Isolation

Complex adaptive systems are governed by feedback loops, in which outputs recursively become inputs. Causation is therefore circular rather than linear. Variables do not act in one direction ($X \rightarrow Y$); they co-evolve through mutual reinforcement and constraint ($X \leftrightarrow Y$). As Colander (2000, p. 13) emphasises, the economy is a “constantly evolving complex system” defined by “feedback effects and nonlinear dynamics.”

Analytical decomposition and *ceteris paribus* reasoning presuppose linear causation: one variable can be changed while others are held constant. Feedback destroys this assumption. When X changes Y, Y necessarily changes X in return. Holding “all else equal” therefore requires breaking the system’s causal structure. What remains is not the system but an artefact.

Monetary policy as a feedback system

Bernanke (2010, pp. 4–13) acknowledged that the financial system is “a network of interconnected institutions” where shocks propagate through feedback. Monetary policy operates through coupled loops: interest rates affect investment and credit conditions; investment affects employment and output; output affects inflation and risk premia; and these in turn feed back into credit conditions and rates. Under such coupling, there is no isolable “effect of interest rates on investment,” because inflation, expectations, risk premia, and balance sheet constraints adjust endogenously. The feedback loop is the phenomenon. Break it and the transmission mechanism disappears.

Firm growth as a recursive system

Penrose (2009, p. 237) showed that firm growth is governed by interacting positive and negative feedback. Positive loops expand resources and opportunities; negative loops impose managerial strain, coordination breakdown, and organisational bottlenecks that slow growth. These loops co-determine the trajectory. One cannot study growth while holding organisational capacity constant, because growth changes capacity. Nor can one study managerial limits while holding scale constant, because scale generates the limits. The system is the loop.

Firm-level translation: General Electric

GE's rise and collapse illustrate why feedback makes decomposition misleading. Under Jack Welch, GE engineered a powerful positive financial feedback loop: high earnings supported a rising stock price; the stock price strengthened acquisition capacity; acquisitions and financial engineering supported reported earnings; and success reinforced the loop. Yet hidden negative feedback accumulated: GE Capital's expansion created systemic financial risk; downsizing destroyed organisational knowledge; accounting acceleration created future obligations; service contracts booked today constrained tomorrow. These liabilities were created by the same feedback architecture that generated apparent success.

When conditions shifted after 2001, the loop inverted: losses weakened the stock; the collapse reduced acquisition capacity; asset sales reduced earnings power; and further decline followed. GE's industrial divisions could appear individually sound while the corporate-level feedback architecture connecting them via GE Capital and financial engineering created fragility. By 2018 GE left the Dow; by 2024 it had been broken into three separate firms. No analysis of jet engines, healthcare equipment, or turbines could reveal the driver of collapse because the causal mechanism resided in corporate-level feedback structure, not in decomposable parts.

Ontological conclusion

Feedback loops destroy the very conditions that make analytical isolation meaningful.

FROM linear causation ($X \rightarrow Y$) $\rightarrow$ **TO** circular causation ($X \leftrightarrow Y$)

FROM independent variables $\rightarrow$ **TO** mutually constitutive dynamics

FROM static equilibrium analysis $\rightarrow$ **TO** evolving system trajectories

In complex adaptive systems, causes do not reside in variables; they reside in **loops**. Decomposition breaks the loops. When loops are broken, the system vanishes.

4.4.3 Critical Nuances and Methodological Consequences

Calibration, not rejection

The critique developed here does not reject analysis. It rejects the universalisation of analysis to systems that are ontologically non-decomposable. Analysis is valid when components retain their identity and causal role when isolated. Synthesis becomes necessary when components exist only through their relations. Ackoff's framework provides a clean boundary condition:

Table 111: Matching Methodological Emphasis to System Structure: From Analysis and Synthesis to Their Explicit Integration in Mixed Systems

Condition of the system	Methodological emphasis
Components independent	Analysis
Components interdependent	Synthesis
Mixed (most real systems)	Explicit integration of both

The mechanistic error is not analysis per se; it is the assumption that analysis alone is sufficient. A calibrated epistemology asks: Which aspects of this system are decomposable, which are not, and how do the two interact?

Firm-level translation: Theranos

Theranos illustrates the failure that occurs when an ontologically integrated system is evaluated as if it were modular. Elizabeth Holmes claimed Theranos could perform hundreds of blood tests from a single finger-prick using the “Edison” device. By 2014, the company was valued at \$9 billion, with investments from Walgreens (\$140m), Safeway (\$350m), and major private investors.

The evaluation process was analytically decomposed: miniaturised hardware – plausible; retail distribution – provided by Walgreens; regulatory status – CLIA certification obtained; market demand – obvious. Each component appeared viable in isolation. But blood diagnostics is an integrated system. Accuracy depends on coupled constraints: sample volume → dilution → reagent chemistry → instrument calibration → signal-to-noise → technician handling → quality control. These elements mutually constrain one another. No component can be changed without changing all others.

Theranos failed not because one component was weak but because the integrated system claim was physically incoherent. Reported device performance did not meet required accuracy for many assays, and the system could not deliver the advertised scope of testing reliably. Analysts decomposed the claim into modules. Laboratory scientists, who understood the coupling structure, recognised the implausibility early – yet their epistemic frame was not dominant in the evaluation process.

When Theranos collapsed in 2018, substantial capital was destroyed. Holmes's 2022 conviction addressed fraud, but the methodological lesson is broader and more general: compositional evaluation can certify plausibility piece-by-piece while missing the system-level constraint that makes the whole impossible.

Methodological consequence

Analytical decomposition generates tools that presuppose separability. In tightly coupled systems, those assumptions fail. The practical implication is not that familiar tools are “wrong,” but that their domain of validity must be treated as conditional. Where coupling is strong, feedback is dominant, and properties are relational, system-level framing becomes essential to avoid category error.

Table 112: Decomposition Methods in Mainstream Economics and Their Systematic Violation of Complex Adaptive System Properties

Decomposition method	CAS violation
Optimisation	Objectives are context-dependent and interaction-defined
Equilibrium	Requires independent markets; CAS are coupled
Representative agent	Aggregation erases heterogeneity and network position
Ceteris paribus	Isolation impossible when variables are coupled
Partial equilibrium	Feedback across markets breaks isolation

CAS therefore requires complementary methods:

- Agent-based modelling — heterogeneous interacting agents
- System dynamics — feedback, nonlinearity, path-dependence
- Network analysis — topology, centrality, contagion
- Qualitative methods — meaning, interpretation, culture
- Multi-method designs — explicit synthesis across levels

These are not “soft alternatives” but methods matched to non-decomposable ontologies.

The critique in Section 4.1 culminates here: when analysis is applied to non-decomposable phenomena, it can become a method of system-destruction rather than system-understanding. This prepares the transition to Section 4.2, where the chapter turns from diagnosing mechanistic ontology and method to articulating the ontological conditions under which emergent phenomena become intelligible.

4.4.4 Counter-Arguments Addressed

“Analysis has been enormously successful in science – why should economics be different?”

Objection.

The analytical method enabled the scientific revolution. Physics, chemistry, and molecular biology achieved extraordinary success by decomposing phenomena into parts. If analysis works for atoms, molecules, and cells, why should it not work for economies?

Response.

This objection conflates ontologically distinct domains. Simon (1962, pp. 467–469) distinguished decomposable, nearly decomposable, and non-decomposable systems. Analysis works reliably when parts retain their defining properties when separated – an approximation that holds for many domains of classical physics. But even physics recognises limits: quantum entanglement, turbulence, and phase transitions do not behave as decomposable systems in this sense.

Laughlin (2005, pp. 208–211) shows that macroscopic order often obeys organising principles not derivable by straightforward micro-law aggregation. Prigogine and Stengers (1984, pp. 291–294) demonstrate that dissipative structures arise only in far-from-equilibrium conditions where decomposition can destroy the phenomenon. Anderson’s (1972, p. 393) “More Is Different” establishes that each level of complexity generates its own ontological primitives.

Economic systems are, if anything, more non-decomposable than many physical systems: agents are reflexive, institutions evolve, and expectations feed back into behaviour. Bak (1996, pp. 29–31) shows that in complex systems correlations can span scales, undermining the assumption that there exists a causally independent “micro” level from which the rest can be cleanly derived. The relevant lesson is not that analysis fails everywhere, but that its domain of validity is conditional. Physics learned the limits of decomposition. Economics has too often treated decomposition as universal.

“We can analyse parts first, then add interactions back”

Objection.

One can analyse components, then add interactions as a second step. Decomposition is just the starting point.

Response.

This assumes interactions are additive. Complex adaptive systems are not additive – they are constitutive. Cilliers (1998, pp. 3–4) shows that in complex systems interactions change what the components are. Kauffman (1995, pp. 71–73) shows co-evolution: components adapt in response to interactions, so their identities are not fixed. A firm’s strategy, for example, has no meaning outside competitive relations (Porter, 1980). Stacey (2011, pp. 60–62) similarly emphasises that key organisational properties exist only in interaction, not inside isolated individuals.

Therefore, “parts first, interactions second” is often conceptually unstable: the parts do not remain the same once relations are introduced. Decomposition followed by recomposition can generate a different system from the one originally observed. Enron remains decisive evidence:

auditors analysed units and accounts; the relevant property – systemic fabrication – existed in the architecture of relations across units.

“Synthesis is too vague for rigorous science”

Objection.

Analysis gives testable hypotheses. Synthesis produces vague talk of emergence and feedback.

Response.

This confuses closed-form tractability with scientific rigour. Rigour consists in preserving the relevant causal structure, not in forcing solvability through isolation. Complex-systems approaches can generate replicable, testable claims about patterns, distributions, and trajectories produced by interaction – even when they do not yield a single equilibrium prediction (Axelrod, 1997; Epstein & Axtell, 1996; Sterman, 2000; Newman, 2018). The 2008 crisis is a reminder that decomposed models can be precise yet precisely wrong when they delete constitutive dynamics.

“Ceteris paribus is just a harmless simplification”

Objection.

All science simplifies. *Ceteris paribus* enables tractability.

Response.

Mäki (1992) distinguishes **isolation** (bracketing what does not matter) from **idealisation** (assuming away what does matter). *Ceteris paribus* is legitimate isolation only where separability approximately holds. In coupled systems it becomes ontological deletion. Cartwright (1999) argues that explanation requires preserving causal structure; Lawson (2003) shows that *ceteris paribus* clauses in economics often remove the very mechanisms generating outcomes. Hausman (1992) shows they can be empirically untestable in open systems because the required isolation cannot be achieved. “All models are false” licenses approximation, not the removal of constitutive dynamics. Decomposing feedback-coupled variables yields accurate inferences about a system that does not exist.

Bottom line.

These objections fail for a shared reason: they assume that complex adaptive systems are decomposable. Once feedback, emergence, and network structure are causally active, analytical decomposition is not merely inaccurate; it becomes conceptually incoherent as a general method.

Conclusion

This subsection developed four convergent arguments showing why analytical decomposition fails as a general methodology for complex adaptive systems:

1. **The *ceteris paribus* method becomes invalid** in coupled systems: variables cannot be isolated without breaking the loops that generate behaviour.
2. **Systemic properties collapse under analysis**: culture, market dynamics, and many organisational capabilities exist only in interaction and vanish when decomposed.

3. **Some systems are ontologically indecomposable:** strategy, coordination, and institutional order do not consist of independent parts but exist only in systemic context.
4. **Feedback loops preclude linear isolation:** circular causation makes decomposition either impossible or artefactual.

These arguments provide methodological closure to Section 4.1's ontological critique. Taken together with Section 4.1.1 (reductionism), Section 4.1.2 (substance), and Section 4.1.3 (atomism), the conclusion is straightforward: a mechanistic metaphysics naturally generates a mechanistic method, and in a CAS domain that method systematically deletes the causal structures that govern outcomes.

Ackoff (1999, pp. 16–18) captured the distinction cleanly: analysis proceeds by taking systems apart; synthesis proceeds by understanding parts through their role in the whole. Where properties arise from interaction rather than components, decomposing the system does not simplify it – it removes it.

The firm translations converge on this result. Spotify illustrates feedback-constituted learning dynamics; Enron illustrates relationally hidden properties; Southwest illustrates mutually constitutive organisational coherence; Theranos illustrates the danger of modular evaluation where coupling is constitutive; GE illustrates corporate-level feedback architectures that dominate component performance. In each case, studying parts independently either made the phenomenon disappear (culture, fraud, coordination) or produced dangerously misleading conclusions (technological feasibility, financial strength, strategic sustainability).

The Aha Moment

Decomposition is not merely a simplification. It is the methodological move that makes complex adaptive systems disappear. Economics has made that move foundational.

The choice confronting economics is therefore not between competing techniques but between competing ontologies. If systems are decomposable, analysis suffices. If systems are non-decomposable, synthesis is indispensable. The evidence from coupled variables, systemic properties, ontological indecomposability, and feedback loops all points unambiguously to the latter. The methodological transformation is not optional refinement – it is an ontological consequence. When the nature of the system changes, the nature of method must change with it.

FROM → TO

Table 113: From Decomposition to Synthesis: Six Methodological Shifts Toward a CAS-Compatible Analytical Framework

FROM (Decomposition)	TO (Synthesis)
Independent variables	Coupled dynamics
Additive properties	Emergent properties
Parts exist independently	Parts defined by relations
Linear causation	Circular causation
Analysis sufficient	Analysis + synthesis calibrated
Static isolation	Dynamic integration

Conclusion

This chapter has pursued a single, tightly bounded objective: to diagnose why dominant economic theory repeatedly fails when confronted with instability, innovation, coordination breakdowns, and systemic crises. The diagnosis advanced here is neither methodological nor empirical in the narrow sense. It is ontological.

Sections 4.1–4.4 demonstrated that the mechanistic paradigm rests on a coherent but deeply incompatible ontological package:

- **Ontological reductionism**, which assumes that wholes can be exhaustively explained by decomposing them into parts;
- **Substance ontology**, which treats firms, agents, and markets as stable entities with intrinsic essences rather than as historically unfolding processes;
- **Atomism**, which privileges independent elements over constitutive relations and collapses explanation into micro–macro aggregation;
- **Analytical decomposition**, which operationalises these metaphysical commitments by isolating components and suppressing feedback, interaction, and emergence.

These commitments are mutually reinforcing. Reductionism presupposes stable substances; substance ontology legitimates decomposition; atomism renders isolation plausible. Together, they generate the characteristic methods of mainstream economics – optimisation, equilibrium analysis, representative agents, and *ceteris paribus* reasoning – while simultaneously rendering relational causality, meso-level structures, and endogenous novelty invisible by construction.

The central result of this chapter is therefore negative but decisive: the mechanistic ontology does not merely approximate economic reality poorly; it systematically deletes the very phenomena that define complex adaptive systems.

This conclusion holds across all four diagnostic dimensions developed in the chapter:

1. **Reductionism fails** because decomposing complex systems destroys emergent properties rather than revealing them.
2. **Substance ontology fails** because it eliminates time, history, and becoming, treating change as deviation rather than constitution.
3. **Atomism fails** because it mislocates causality in intrinsic properties instead of relational structure and network position.
4. **Analytical decomposition fails** because feedback loops, co-evolution, and ontological indecomposability render isolation logically incoherent.

Taken together, these failures establish that the crisis of economic theory diagnosed in Chapter 3 is neither contingent, remediable, nor technical. It is structural. No accumulation of additional frictions, heterogeneity, expectations refinements, or computational power can repair a framework whose underlying ontology misrepresents what exists.

This diagnosis carries two critical implications.

First, the failure of the mechanistic paradigm cannot be resolved internally. Because the ontological commitments identified here constitute the hard core of the dominant research programme, reform attempts that preserve them – however sophisticated – will necessarily

reproduce the same blind spots. The persistence of equilibrium-centred modelling after the global financial crisis is not intellectual inertia; it is ontological constraint.

Second, reconstruction becomes unavoidable. Once reductionism, substance ontology, atomism, and universal decomposition are rejected, the question is no longer whether economics should change, but what ontological commitments must replace them.

The negative space carved out by this chapter already delineates the contours of that reconstruction. If economic systems are:

- historically irreversible,
- interaction-driven,
- network-structured,
- feedback-governed, and
- generative of endogenous novelty,

then they cannot be represented as collections of independent substances optimising in timeless equilibrium. They must instead be understood as processual, relational, and emergent systems.

This is precisely the ontological terrain of process philosophy, emergence theory, and complex adaptive systems.

Chapter 5 therefore does not introduce a new perspective arbitrarily, nor does it propose an alternative “school” of economics. It undertakes an ontological reconstruction demanded by the failures diagnosed here. Where Chapter 4 demonstrated why mechanistic being is untenable, Chapter 5 reconstructs economic ontology around process, emergence, and adaptive interaction.

The transition is thus not optional, stylistic, or pluralistic in the weak sense. It is a logical consequence.

Once being, isolation, and decomposition fail, economics must be rebuilt around becoming, interaction, and synthesis. Chapter 5 develops that reconstruction.

Chapter 5 – Ontological Reconstruction: Process, Emergence, and Complex Adaptive Systems

Introduction

Chapter 5 reconstructs the ontological foundations required for an economics adequate to complex adaptive systems. The inherited body of theory – built around closed systems, decomposability, and optimisation over pre-given possibilities – presupposes an ontology of atomistic entities with intrinsic properties, stable functional relations, and equilibrium-seeking dynamics. The aim here is not to criticise that ontology rhetorically, but to replace it systematically: to specify what economic reality must be like if the phenomena that dominate modern economies – capabilities, networks, ecosystems, endogenous novelty, and systemic fragility – are taken as primary rather than anomalous.

The reconstruction proceeds through four interrelated transformations:

1. **Ontological transformation.** Economic entities are reconceptualised as emergent wholes constituted through relations, existing as stabilised processes rather than static substances. This establishes that organisational capabilities, market structures, and institutional orders possess causal powers irreducible to their components – removing the representative-agent shortcut from the category of "simplification" and placing it in the category of ontological distortion.
2. **Structural transformation.** The firm and other economic entities are shown to be thermodynamically open systems, far from equilibrium, with permeable boundaries and network-embedded architectures. The meso-level – organisations, industries, clusters – emerges as the privileged analytical scale where emergence is empirically observable and causally operative.
3. **Epistemological and methodological transformation.** From emergence and relational constitution follow necessary implications for knowledge and inquiry. Analysis alone cannot exhaust understanding of systems whose defining properties exist only at higher levels; hence the requirement for synthesis-guided analysis. Methods must preserve interaction, context, and trajectory rather than eliminate them through *ceteris paribus* isolation; hence the shift from reductionism to expansionism.
4. **Dynamic transformation.** CAS ontology transforms what dynamics mean. Under mechanistic ontology, dynamics describe movement toward equilibrium. Under CAS ontology, dynamics describe the continual creation of new economic possibilities through emergence, creative destruction, path-dependent innovation, and positive-feedback amplification. Economics shifts from the study of equilibrating systems to the study of evolutionary systems that generate novelty.

Chapter Roadmap

The chapter proceeds through nine integrated sections.

Section 5.1 establishes emergence as the foundational ontological principle. Organisational capability, market dynamics, and institutional order are shown to be system properties with causal powers irreducible to component attributes. The section develops emergence across structural, dynamic, organisational, and cognitive domains, demonstrating that wholes possess properties that do not exist in their parts.

Section 5.2 establishes process ontology – the reconceptualisation of reality from static substance to dynamic becoming. Economic entities are shown to be temporarily stabilised patterns within ongoing flux, persisting only through continuous reproduction. This aligns economic ontology with the arrow of time and with empirical realities of adaptation, learning, and breakdown.

Section 5.3 establishes relational constitution as foundational. Economic entities are constituted through relations, not merely connected by them. Topology, embeddedness, and network position are therefore not context but causal structure – where a firm sits in networks is part of what the firm is.

Section 5.4 synthesises the preceding three pillars, developing substances as stabilised processes. Identity is reconceptualised from essence to trajectory; entities are what they have become through irreversible historical development. This synthesis resolves the apparent tension between process (flux) and entity (stability) by showing how persistence is achieved through continuous activity.

Section 5.5 establishes the meso-level as the privileged analytical scale. The micro-level destroys emergent properties through decomposition; the macro-level obscures them through aggregation. Only at the meso-level – organisations, industries, networks, clusters – are interaction-generated properties preserved and observable.

Section 5.6 translates the ontological reconstruction into structural implications. CAS are shown to be thermodynamically open systems, maintaining organisation through continuous throughput of energy, matter, and information. Equilibrium is cessation; far-from-equilibrium dynamics are the condition of economic life. Boundaries are not walls but gradients – permeable membranes that filter, transform, and selectively couple internal and external processes.

Section 5.7 develops network architecture as the structural expression of relational ontology. Scale-free topology, preferential attachment, and clustering coefficients generate the power-law distributions, hub-and-spoke structures, and cohesive subgroups that characterise real economic networks. Network clusters emerge as canonical meso-level structures where dense internal ties and sparse external bridges create the conditions for emergence.

Section 5.8 develops the interconnections from ontology to epistemology, methodology, and structure. It establishes how emergence, process, and relational constitution jointly require: (i)

synthesis-complementing-analysis in epistemology, (ii) expansionism over reductionism in methodology, and (iii) open-system, ecosystem-embedded structure. These are not optional preferences but logical necessities given the preceding ontology.

Section 5.9 completes the implications chain with dynamic implications. If reality is emergent, processual, and relational, then economic systems are capable of genuine novelty – not merely recombination but ontological innovation. The section develops five dynamic implications: emergence creates ontological novelty; process ontology enables creative destruction; non-ergodicity creates path-dependent innovation; positive feedback amplifies novelty; and open-ended evolution prevents convergence to equilibrium.

The Intended Result

The chapter delivers a coherent replacement ontology, not an eclectic overlay: a world in which firms, markets, and institutions are open, relational, historically constituted systems whose most important properties arise at the meso-level and whose evolution is driven by endogenous novelty rather than equilibrium mechanics.

This platform does not merely accommodate complexity; it makes the inherited assumptions of closure, decomposability, and equilibrium the exceptional case. The implications are strict: if this ontology is correct, then the foundational assumptions of orthodox economics – atomistic agents, intrinsic preferences, equilibrium convergence, and calculable futures – cannot be sustained as general foundations. They become, at best, useful approximations for highly restricted domains.

5.1 The Principle of Emergent Wholeness: When More Becomes Different

Introduction

This section develops the first and foundational pillar of the dissertation's ontological reconstruction: the reconceptualisation of economic systems from aggregations of individual components to emergent wholes whose properties arise from interaction yet remain irreducible to those interactions.

Emergent wholeness is defined as the principle that macro-level properties genuinely arise from micro-level interactions but cannot be predicted, reduced, or deduced from complete knowledge of those micro-level components. Emergence is characterised by three inseparable features:

- **Upward causation** – macro properties arise from micro interactions
- **Downward causation** – macro structures constrain and enable micro behaviour
- **Ontological non-deducibility** – macro properties are not derivable from micro descriptions, even in principle

This definition integrates complexity science, philosophy of science, and heterodox economics into a unified ontological framework for understanding firms, markets, and economies as complex adaptive systems.

Section Architecture. This section receives the problematisation of mainstream economics' foundational assumptions established in Chapter 4, which identified ontological reductionism as a core commitment incompatible with complex adaptive systems. It provides the positive ontological alternative: where wholes were falsely treated as sums, they are here reconstructed as structures with causal powers of their own. This reconstruction furnishes the ontological foundation upon which the subsequent sections build: process ontology (Section 5.2), relational ontology (Section 5.3), and the structural implications that follow (Sections 5.3.1–5.3.5).

The section serves three architectural functions within the dissertation. First, it establishes emergence as the generative principle that renders subsequent transformations necessary rather than optional: epistemological pluralism follows because multiple emergent levels coexist; non-equilibrium dynamics follow because emergent structures evolve; organisational formation follows because emergence produces genuine novelty. Second, it reconceptualises the firm as an emergent organisational structure rather than an aggregation of individuals, grounding the dissertation's meso-level analytical focus. Third, it provides the ontological warrant for rejecting methodological individualism – not as a matter of preference but as a matter of logical necessity.

The section proceeds as follows. Part I introduces Anderson's foundational ontological break, establishing emergence as ontological rather than merely epistemic. Part II develops the three defining characteristics of emergence with firm-level illustrations. Part III demonstrates emergence across four domains – structural, dynamic, organisational, and cognitive –

establishing its generality as an organising principle. Part IV applies emergence to firms and markets, the central objects of economic theory. Part V locates emergence at the meso-level, establishing the dissertation's distinctive analytical focus. Part VI addresses counter-arguments. The conclusion synthesises the ontological transformation achieved.

Before proceeding, a central objection must be addressed. Emergence is often treated as merely epistemic – a label for ignorance, modelling difficulty, or computational limitation. This dissertation rejects that interpretation. Following Anderson (1972, p. 393), emergence is treated as ontological: new properties genuinely come into existence at higher levels of organisation – properties that do not exist at lower levels and cannot be deduced even in principle.

Superconductivity is not "hidden" in electrons; it exists only in their organised collective. When the collective dissolves, the property vanishes. Because emergence is ontological rather than epistemic, methodological individualism cannot be rescued by better computation. The shift from reductionism to emergence is therefore not a methodological upgrade but a paradigm change in what economic systems are.

5.1.1 The Ontological Core and Defining Characteristics

I. Anderson's Revolution: The Ontological Core

"More Is Different": Complexity Science's Declaration of Independence

Philip W. Anderson's (1972) *Science* article, "More Is Different," constitutes the foundational ontological break from reductionism. Anderson demonstrated that although physics explains elementary particles and their interactions, such knowledge cannot reconstruct higher-level phenomena:

"The constructionist hypothesis breaks down when confronted with the twin difficulties of scale and complexity." (Anderson, 1972, p. 393)

The implication is radical: quantitative increase generates qualitative novelty.

A million water molecules exhibit fluidity. Ten billion neurons exhibit consciousness. Ten thousand trading agents exhibit markets.

These properties do not exist in the parts – not even in latent form. They arise only at the collective level.

For economic theory, this insight is decisive. The representative agent is not a harmless simplification; it is an ontological error. By averaging across heterogeneity, it eliminates precisely the interaction patterns from which emergence arises. Mainstream economics therefore does not simplify complex systems – it erases their constitutive properties.

Irreducibility: Ontological, Not Epistemic

Emergence is defined by irreducibility, and the distinction between epistemic and ontological irreducibility is decisive:

- **Epistemic irreducibility:** macro properties cannot be derived from micro descriptions due to limited knowledge or computational capacity.
- **Ontological irreducibility:** macro properties do not exist at the micro level at all, rendering derivation logically impossible.

Superconductivity illustrates this distinction. Below a critical temperature, materials exhibit zero electrical resistance. No individual electron is even partially superconducting. The property exists only in the organised collective. Complete knowledge of electrons does not yield superconductivity because it is not an electron property – it is a property of their organised system.

Empirical studies in organisational economics and strategy consistently find that emergent routines and coordination patterns explain the majority – often estimated at 60–80% – of performance variance beyond individual attributes or asset endowments.

Firm-Level Translation: IDEO

IDEO's design capability is ontologically irreducible.

IDEO's extraordinary innovation record – Apple's first mouse, the Palm V, Oral-B toothbrushes, healthcare systems, financial service platforms – does not reside in any designer. It emerges from the organisation's interaction structure: interdisciplinary teams, rapid prototyping, ethnographic research, spatially designed collaboration, and a culture that legitimises failure.

Competitors repeatedly hired IDEO alumni yet failed to reproduce IDEO-level innovation. Individuals transferred; the emergent organisational capability did not.

David Kelley described IDEO's strength as "creative confidence" – not an individual trait but a collective property produced by how people interact. The same individuals in a blame-oriented organisation would behave conservatively; at IDEO they take risks because the system makes risk legitimate.

By 2024, IDEO had completed over 4,000 innovation projects across industries. The capability persisted despite staff turnover because it was a property of the organised whole. When the organisational structure dissolves, the capability disappears. It cannot be decomposed into people, procedures, or manuals.

This is ontological emergence: the capability belongs to the system, not to its parts.

II. Three Defining Characteristics of Emergence

Having established emergence's ontological core, we now specify its three defining characteristics. Together, they distinguish genuine emergence from mere aggregation, operationalise the concept for analysis, and anchor all subsequent arguments in the dissertation.

Characteristic 1: Macro Properties Arise from Micro Interactions

Emergence is fundamentally bottom-up: macro-level properties arise from micro-level interactions. Markets emerge from trading, traffic patterns from driving decisions, ecosystems from organism interactions, and organisational culture from employee interaction.

However, this relationship is non-linear. Macro properties do not arise through summation but through thresholds, tipping points, feedback, and phase transitions in which quantitative accumulation produces qualitative transformation.

Schelling's (1978) segregation model demonstrates this with paradigmatic clarity. Agents with only mild preferences for similar neighbours – each willing to be in a local minority – generate extreme segregation at the macro level. No agent intends segregation; none prefers homogeneity. Yet interaction dynamics amplify small preferences into large-scale pattern formation. The emergent outcome (segregated neighbourhoods) diverges radically from individual intentions (tolerance).

The macro pattern is therefore not a reflection of individual motives but a product of interaction structure.

Firm-Level Translation: Wikipedia

Wikipedia exemplifies emergence from micro-interaction at planetary scale. No individual writes Wikipedia. It emerges from millions of edits, reversions, disputes, and resolutions.

By 2024, Wikipedia contained over 60 million articles in more than 300 languages, produced by roughly 300,000 active editors generating hundreds of edits per minute. Its defining properties – coverage, neutrality, internal coherence, and error-correction – do not exist in any individual contribution.

Bias is introduced by individuals; neutrality emerges from collective contestation. Errors are made by individuals; accuracy emerges from distributed correction. *Nature's* (2005) comparison with *Encyclopaedia Britannica* showed comparable accuracy not because contributors were individually authoritative but because the interaction system generated collective quality.

Placed in a different governance structure – without version control, edit histories, talk pages, or dispute-resolution norms – the same editors would not generate Wikipedia. Jimmy Wales did not design content; he designed interaction rules from which emergent knowledge production arose. The encyclopaedia is a system-level property, not an aggregate of contributions.

Characteristic 2: Downward Causation

Once emergent properties exist, they exert causal influence back onto their components. This is downward causation. The macro is not merely an outcome of micro behaviour; it becomes a constraint and enabler of future micro behaviour.

Organisational culture provides a canonical example. Culture emerges from repeated interaction – shared meanings, norms, and expectations. But once crystallised, it structures what individuals perceive as appropriate, risky, or legitimate. New employees encounter a pre-existing pattern that shapes their behaviour. Behaviour produces culture; culture produces behaviour.

This bidirectional causation creates self-reinforcing loops.

Firm-Level Translation: Zara

Zara's extraordinary speed – bringing new designs from concept to store in roughly fifteen days – emerged from specific interaction patterns: designers observing customers in stores, real-time sales data flowing to headquarters, small-batch manufacturing, and rapid logistical integration.

Once established, this system exerted downward causation. It constrained where factories were located, how designers worked, which fabrics were used, how inventory was managed, and how frequently collections changed. Designers entering Zara must conform to continuous redesign rather than seasonal planning. Customers come to expect new merchandise twice weekly, reinforcing the cycle.

By 2024, Inditex operated over 6,000 Zara stores worldwide, all governed by the same emergent capability that originated in a specific historical configuration and now constrains every operational decision. Competitors cannot replicate this without dismantling their own systems to allow a new emergent structure to form.

Characteristic 3: Ontological Non-Deducibility

The third defining property of emergence is non-deducibility: macro-level properties cannot be derived from micro-level descriptions even with complete information.

This distinguishes emergence from aggregation. Aggregated properties are predictable; emergent properties are not. Ten pounds of flour weigh ten times one pound. Ten thousand traders generate bubbles, crashes, momentum, and contagion – phenomena not deducible from individual psychology.

Technology markets with network externalities demonstrate this decisively. Outcomes depend on path-dependent adoption sequences, not intrinsic properties. As Arthur (1994, p. 15) formalised: "Small random events can cumulate step-by-step to lock the market into the monopoly of a possibly inferior technology."

Firm-Level Translation: VHS versus Betamax

Sony's Betamax was technically superior to JVC's VHS: higher resolution, better sound, more compact cassettes. Yet VHS won the format war.

The outcome was determined not by technical quality but by adoption dynamics. Slightly longer recording time attracted movie rental stores; rentals attracted consumers; consumers attracted studios; network effects accumulated; lock-in followed.

Even with complete knowledge of technologies, prices, and preferences, the winner could not be deduced *ex ante* because the outcome depended on historical sequences of contingent events – which studios licensed which format, which retailers stocked which tapes. Small contingencies tipped the system. The same structure reappeared in HD-DVD versus Blu-ray, with a different winner produced by different contingencies.

Market outcomes in networked environments are therefore emergent: they arise from interaction dynamics, not component properties.

Synthesis

This part accomplishes a decisive ontological shift. It does not merely assert that economic systems are "complex"; it demonstrates that macro-level structures possess causal powers invisible to micro-level analysis. Emergent properties are not descriptive conveniences but real features of organised systems: they shape behaviour, constrain trajectories, and generate outcomes irreducible to component analysis.

Once this is recognised, methodological individualism and reductionist aggregation are no longer neutral simplifications. They become systematic sources of explanatory blindness. Understanding firms, markets, and institutions as complex adaptive systems therefore requires treating emergence not as complication, but as foundational reality.

5.1.2 Domains and Economic Applications

III. Multiple Domains of Emergence

Emergence is not a local anomaly but a general principle of organisation across complex systems. It operates simultaneously across structural, dynamic, organisational, and cognitive domains. These domains are analytically distinguishable but ontologically inseparable. Taken together, they demonstrate that emergence is not confined to particular phenomena – such as innovation or crises – but constitutes the fundamental mode through which complex systems exist and evolve.

1. Structural Emergence

Structural emergence occurs when interaction patterns generate system-level organisation that cannot be reduced to component properties. Markets, industries, and production systems are

not mere collections of transactions; they are organised structures with topology, hierarchy, and connectivity.

Market structures – concentration ratios, platform dominance, value-chain architectures, and intermediation hierarchies – arise from trading interactions yet exhibit properties that no individual transaction contains. The structures of Amazon's platform, Visa's payment network, or Silicon Valley's innovation ecosystem are not encoded in any single exchange; they are relational architectures that emerge from millions of interactions.

As Cilliers (2005, p. 260) observes, in complex systems "the relationships between the elements are more important than the elements themselves." Structural emergence therefore implies that economic power, market access, and competitive advantage reside not in agents as such but in positions within networks.

2. Dynamic Emergence

Dynamic emergence arises from nonlinear interactions that generate thresholds, cascades, and regime shifts. Here, time is constitutive: systems evolve through feedback, amplification, and sudden qualitative change.

The 2010 Flash Crash exemplifies this logic. On 6 May 2010, U.S. equity markets lost nearly one trillion dollars in minutes before rebounding (SEC/CFTC, 2010). No algorithm malfunctioned, and no trader acted irrationally. Each agent responded locally to price movements. Yet the interaction of high-frequency trading algorithms generated runaway feedback loops, liquidity evaporation, and cascading price collapses – system-level dynamics that no component intended or controlled.

Dynamic emergence demonstrates that economic instability is not caused by exogenous "shocks" to an otherwise stable equilibrium. It is generated endogenously by interaction structures themselves. Crises are therefore not deviations from stability; they are properties of nonlinear system dynamics.

3. Organisational Emergence

Organisations exhibit emergent capabilities that exist only at the level of coordinated activity. Firms do not merely aggregate workers and machines; they generate collective capacities – learning, coordination, adaptation, and innovation – that no individual possesses.

Holland (1995, p. 11) captures this precisely: "the interactions, not the actions of the individual agents, are the primary source of the complexities." A firm's ability to design new products, operate global supply chains, or pivot strategically resides in interaction patterns: routines, communication structures, trust relations, and shared interpretive frameworks.

When these interaction structures are disrupted – by fragmenting teams, dismantling routines, or eroding trust – the capability disappears even if all individuals remain. Organisational competence is therefore an emergent property of the system, not a sum of personal skills.

4. Cognitive Emergence

Economic systems are also collective cognitive systems. They generate knowledge that no individual possesses.

Hayek (1945, pp. 526–527) recognised that prices are not signals chosen by anyone but emergent summaries of dispersed information. Each trader knows only local conditions; prices integrate this distributed knowledge into a collective guide for action. The information embedded in prices does not exist in any single mind – it exists only at the system level.

Markets, scientific communities, open-source projects, and innovation ecosystems all exhibit collective intelligence: problem-solving, discovery, and learning that emerge from interaction among heterogeneous agents. Wikipedia, financial markets, and technological standards generate knowledge that transcends individual cognition.

Cognitive emergence therefore implies that economic systems process information collectively – not metaphorically, but operationally – through distributed interaction and feedback.

Synthesis

Across all four domains, the same ontological pattern recurs:

Table 114: Emergence Across Four Domains: What Arises and Why It Is Irreducible to Micro-Level Components

Domain	What Emerges	Why It Is Irreducible
Structural	Networks, hierarchies, market architecture	Topology is not contained in transactions
Dynamic	Crises, bubbles, growth regimes	Feedback generates nonlinearity
Organisational	Capabilities, routines, culture	Coordination exists only in interaction
Cognitive	Prices, knowledge, collective intelligence	Information is system-level

Emergence is therefore not a feature of particular phenomena but the generative principle of economic reality. Firms, markets, and institutions are not aggregates; they are organised systems whose defining properties exist only at higher levels of organisation.

IV. Economic Applications: Firms and Markets as Emergent Structures

The preceding parts established emergence as an ontological principle operating across multiple domains. This part applies that principle to the two central objects of economic theory: firms and markets. This application is decisive. If emergence is real, then neither firms nor markets can be understood as aggregations of individual behaviour. They must be analysed as systems with their own causal powers.

1. The Firm as Emergent Capability System

The mechanistic view treats the firm either as: (a) a production function, $Q = f(K,L)$, or (b) a collection of individuals whose actions aggregate into firm behaviour.

Emergence rejects both. Firms are organised systems whose capabilities exist at the collective level and cannot be reduced to individual attributes.

Penrose (1959/2009, p. 237) already identified this when she argued that firm growth depends on "the cumulative growth of collective knowledge." Crucially, this knowledge is not located in any single individual. It is distributed across routines, communication structures, shared mental models, and coordination patterns. It persists through turnover and evolves through interaction.

Nelson and Winter (1982, p. 4) formalised this logic through routines: firms operate through stable but evolving patterns of behaviour that coordinate individuals into productive wholes. These routines are not choices made at a point in time; they are emergent outcomes of organisational history.

A firm therefore has capabilities in the same way a brain has cognition or a hive has collective intelligence: as system-level properties that arise from interaction yet belong to the whole.

Firm-Level Translation: Netflix

Netflix's trajectory – from DVD-by-mail (1997) to global streaming and content platform – illustrates emergence with unusual clarity. By 2024, over 80% of viewing on Netflix was driven by its recommendation system. That system is not merely an algorithm authored by engineers; it is an emergent learning ecology generated by the interaction of:

- billions of viewing decisions,
- data-science models,
- content metadata,
- creative judgment,
- platform architecture, and
- user feedback loops.

No individual – not even Netflix's data scientists – can fully specify why a particular recommendation appears, because the system learns patterns that exist only at population scale.

More importantly, Netflix's content-creation capability (*House of Cards*, *Stranger Things*, *Squid Game*) did not result from either human intuition or algorithmic prediction alone. It

emerged from the integration of creative teams, analytics, platform feedback, and organisational culture. Competitors can hire Netflix engineers or copy elements of its tooling, but cannot reproduce the capability because the capability is systemic.

Netflix does not merely *possess* a capability. Netflix **is** a capability: an emergent organisational intelligence.

2. Markets as Emergent Order

Markets are equally emergent. Hayek (1945, pp. 526–527) showed that prices are not chosen by anyone; they are collective informational structures generated by decentralised interaction. No trader knows the full state of supply and demand. Yet the price system integrates dispersed knowledge into actionable signals.

Markets therefore think – not metaphorically but operationally.

Market dynamics – momentum, bubbles, crashes, volatility clustering, liquidity freezes – exhibit all three defining properties of emergence:

- **Upward causation:** they arise from individual actions
- **Downward causation:** they constrain future behaviour
- **Non-deducibility:** they cannot be derived from individual psychology

Bubbles make this transparent. Each trader may respond rationally to rising prices; collectively, those responses generate positive feedback that drives prices away from fundamentals. The bubble is not an intention; it is a system-level property.

Financial Crisis as Emergent Fragility

The 2007–2008 crisis exposed emergence brutally. As Haldane (2012) argued, financial systems are tightly coupled networks. Banks held each other's securities; interbank exposures transmitted risk. This generated an emergent property – systemic fragility – that was not present in any single institution's balance sheet.

Kirman (2016, p. 560) states the principle plainly: "The structure of the financial network is at least as important as the health of the individual institutions." Micro-prudential regulation could not "see" fragility because fragility was not in the banks; it was in the network.

Markets, like firms, therefore possess causal powers that do not belong to any participant.

Categorical Reclassification

Emergence forces a categorical reclassification of economic reality:

Table 115: Categorical Reclassification of Economic Reality: How Emergence Transforms Five Fundamental Economic Concepts from Orthodox to Systemic Categories

Orthodox View	Emergent View
Firm = production function	Firm = capability system
Market = sum of trades	Market = information-processing network
Outcomes = aggregation	Outcomes = system-level dynamics
Risk = individual exposure	Risk = network topology
Rationality = individual	Rationality = collective

This is the ontological core of CAS economics: firms and markets are not static "things." They are organised, learning, adaptive systems. And once this is acknowledged, the meso-level is no longer optional – it is where much economic causation resides. The following part develops this implication.

5.1.3 Emergence, the Meso-Level, and Synthesis

V. Emergence and the Meso-Level

Emergence does not occur at the level of isolated individuals, nor at the level of aggregate statistics. It occurs in between – at the meso-level, where interaction structures, networks, organisations, and institutions form. This is the intermediate scale where micro actions are organised into macro patterns and where new causal powers arise.

Complex adaptive systems therefore cannot be analysed through a micro–macro shortcut. They require meso-level ontology.

The Invisible Middle

Methodological individualism jumps directly from individuals to aggregates without conceptualising the structures that organise interaction between them. This leap erases precisely the level at which emergence occurs. Yet emergence does not happen in a vacuum. It requires interaction architectures:

- networks that channel information,
- organisations that coordinate action,
- industries that structure competition,
- institutions that stabilise expectations.

These are not epiphenomena. They are meso-level entities with causal properties of their own.

At this level:

- firms emerge as coordinated systems from individuals,
- markets emerge as structured exchange networks,
- industries emerge as competitive ecologies,
- institutions emerge as stabilised rule systems.

None of these exists at the micro level. None is reducible to macro aggregates. They exist only as meso-level patterns of organised interaction. Emergence is therefore not something that happens *between* micro and macro; it occurs *through* meso-level structures.

If economics jumps from individuals to aggregates, it erases: ecosystems, organisations, networks, institutions, and industries – the very structures through which emergence happens.

Firm-Level Translation: Silicon Valley

Silicon Valley is not a collection of successful firms. It is an emergent economic organism. Its origins trace to a specific historical sequence: the 1957 defection of the "traitorous eight" from Fairchild Semiconductor, Stanford's encouragement of entrepreneurship under Frederick Terman, the clustering of venture capital along Sand Hill Road, and California's refusal to enforce non-compete clauses.

From these interactions emerged something no firm created and no planner designed: a self-reinforcing innovation ecosystem.

Once formed, it acquired properties that none of its members possessed individually:

- high-velocity knowledge diffusion,
- extreme labour mobility,
- dense venture-capital networks,
- rapid recombination of talent and ideas,
- cultural tolerance of failure,
- serial entrepreneurship.

These are system-level properties. No startup owns them. No investor controls them. Yet every participant is enabled – and constrained – by them. New entrants cannot opt out; they must operate within the ecosystem's emergent rules:

- Talent moves freely
- Ideas circulate rapidly
- Failure is normal
- Networks matter more than contracts

And as firms adapt to these conditions, they reinforce the very conditions that shape them – creating powerful positive feedback that stabilises and deepens the ecosystem.

By 2024, Silicon Valley accounted for roughly 30% of U.S. venture capital and housed five of the world's ten most valuable firms (Apple, Alphabet, Meta, Nvidia, Tesla). Yet no amount of firm-level analysis can explain this concentration. The causal force lies in meso-level emergence: the ecosystem itself.

Attempts to replicate Silicon Valley elsewhere have failed not because others lack talent or capital, but because emergent systems cannot be engineered by assembling components. They arise through historical interaction patterns, not policy blueprints.

Why This Matters

The meso-level is not a methodological convenience. It is where economic reality is constituted. CAS economics is therefore not micro + macro. It is micro → meso → macro, with meso-level structures doing much of the causal work: organising interaction, generating capabilities, stabilising expectations, and transmitting shocks.

This is where firms become capabilities. This is where markets become information systems. This is where economies become adaptive systems.

Having established the ontological argument for emergence across its characteristics, domains, applications, and spatial location, the following part addresses objections.

VI. Counter-Arguments Addressed

"Emergence is just epistemic – better models and more data will reduce macro to micro."

Objection. Emergence reflects ignorance, not ontology. With enough data and computation, macro properties will be derived from micro specifications.

Response. This objection confuses epistemic limitation with ontological structure. Kim (1999, pp. 20–23) distinguishes weak emergence (currently intractable) from strong emergence

(logically non-derivable because the property does not exist at the lower level). The phenomena central to this dissertation – organisational capability, market dynamics, institutional order – belong to the second category.

Superconductivity illustrates the point: it does not exist in electrons individually; it exists only in organised collectives. Micro-level completeness cannot reveal a property that literally does not exist at that level. Bedau (2008, pp. 445–447) further shows that ontologically emergent properties are multiply realisable: the same macro property can arise from different micro configurations, defeating claims of micro sufficiency.

Downward causation provides additional evidence. When organisational culture constrains behaviour or market structure constrains strategy, macro entities exercise real causal power. As O'Connor and Wong (2015, pp. 16–18) argue, such causal influence is not a feature of epistemic constructs but of real organised systems.

More computation may improve simulation, but simulation does not eliminate emergence. Agent-based models generate emergent properties precisely because those properties are absent at the agent level. Computational success therefore confirms emergence rather than dissolving it.

"Downward causation is incoherent – only parts can cause."

Objection. Only micro entities have causal power; macro causation is reducible.

Response. This objection presupposes reductionism rather than defending it. If emergence is real, macro-level entities are real entities with causal powers arising from organisation.

Campbell (1974, pp. 179–181) established downward causation in biological systems; Sperry (1986, pp. 265–267) did so in neuroscience. Craver and Bechtel (2007, pp. 547–549) clarify the logic by distinguishing constitution from causation: parts constitute wholes, but wholes constrain parts through organisation. Culture is constituted by individual actions yet constrains future actions; this is not contradiction but the defining structure of complex systems.

Reductionism fails because it demands that wholes be nothing but parts, while CAS demonstrate that organisation itself is causally efficacious.

"Emergence is unfalsifiable mysticism."

Objection. Emergence is invoked whenever reduction fails; it is not scientific.

Response. Emergence makes strong empirical claims. Chalmers (2006, pp. 251–253) identifies three testable criteria: (i) novel causal powers not possessed by components, (ii) dependence on but irreducibility to lower-level properties, and (iii) explanatory indispensability. IDEO's capability satisfies all three: it does not travel with individuals, persists through turnover, and varies with organisational structure – all empirically observable features.

Humphreys (2016) and Clayton (2004) show that emergence functions as a progressive scientific research programme in physics, chemistry, and biology. It generates novel predictions and explains phenomena reductionism cannot.

It is not emergence but reductionism that now fails empirical tests across economics, finance, and organisational theory.

"Agent-based models capture emergence – we don't need a new ontology."

Objection. ABM shows micro rules generate macro patterns, rescuing methodological individualism.

Response. This reverses the logic. ABM succeeds precisely because it abandons the core ontological commitments of methodological individualism: homogeneity, independence, equilibrium, and representative agents. Epstein and Axtell (1996) explicitly frame ABM as growing macro structures from interaction; Tesfatsion (2006) and Miller and Page (2007) emphasise that ABM outcomes are **surprises** – not deducible from micro-rules in closed form. That is precisely what emergence means.

ABM does not eliminate emergence; it demonstrates it.

Closure

Emergence is not an appeal to mystery. It is the recognition that organisation has causal power. Macro structures are not shadows of micro entities; they are real systems that act back on their components. Once emergence is acknowledged, reductionism, substance ontology, atomism, and compositional methodology all collapse together.

Economics must now move forward on new ontological ground.

VII. Conclusion

This section has established emergent wholeness as the first pillar of the dissertation's ontological reconstruction. Where reductionism treats wholes as aggregations whose properties are derived from parts, emergence reconceptualises wholes as structures whose properties arise from but are irreducible to component interactions.

Three characteristics define genuine emergence:

1. **Upward causation** – macro-level properties arise from micro-level interactions
2. **Downward causation** – macro-level structures constrain and shape micro-level behaviour
3. **Non-deducibility** – macro-level properties cannot be derived from micro-level properties even with complete information

Together, these establish that emergence is not a modelling convenience but an ontological feature of complex systems.

The Transformation Achieved

The ontological pivot accomplished here is fundamental:

FROM treating wholes as aggregations with properties derived from parts TO treating wholes as emergent structures with their own causal powers

Firms are not production functions or collections of employees but emergent capability systems. Markets are not collections of traders but emergent coordination and information structures. Economies are not aggregations of agents but emergent institutional orders.

Ontological Transformation Table

Table 116: Ontological Transformation from Reductionism to Emergence: Five Core Commitments Reconceived

Reductionist Ontology	Emergence Ontology
Whole = sum of parts	Whole > sum of parts
Properties aggregate linearly	Properties emerge non-linearly
Causation flows only upward	Causation flows bidirectionally
Macro derivable from micro	Macro irreducible to micro
Single-level analysis sufficient	Multi-level analysis required

The multi-domain character of emergence – structural, dynamic, organisational, and cognitive – establishes that emergence is not a special case but a general principle of complex systems. Firms emerge as coordinated capability from employee interactions; markets emerge as spontaneous order from trading; economies emerge as institutional structures from repeated social interaction. Each displays upward causation, downward causation, and non-deducibility – and therefore each requires analysis at its own level.

Implications for the Firm

For the reconceptualisation of the firm, the implications are decisive. Firms are not input–output machines or aggregations of individuals, but emergent organisational capabilities. Competitive advantage arises from interaction patterns – routines, networks, coordination architectures – not from isolated resources, individual talent, or static optimisation. It cannot be decomposed, purchased, or copied by hiring individuals alone; it must be cultivated through the design and evolution of interaction structures.

Strategy therefore becomes emergence management: shaping networks, routines, and coordination architectures that generate desired capabilities, sustaining them through turnover and growth, and transforming them as environments evolve.

Forward Integration

This section provides the ontological foundation for the reconstructions that follow. Section 5.2 develops process ontology, establishing that emergent structures do not exist statically but continuously emerge through ongoing activity – adding the temporal dimension that prevents the misreading of CAS as merely "complex substances." Section 5.3 develops relational ontology, establishing that the interactions from which emergence arises are not external links between pre-formed entities but constitutive relations that generate entity identity. Together, these three pillars – emergence, process, and relations – form the ontological triad upon which the dissertation's reconceptualisation of the firm is built.

The reconstruction proceeds on emergent foundations.

5.2 Process Ontology: Reality as “Becoming”

This section develops the second pillar of the dissertation's ontological reconstruction: the reconceptualisation of reality from static substance to dynamic process. Process ontology holds that reality fundamentally consists of activities, events, and becomings rather than enduring substances with fixed essences – such that what we conventionally call "things" are temporarily stabilised patterns within ongoing flux. This definition synthesises insights from process philosophy (Whitehead, 1929/1978; Bergson, 1907/1998), complexity science (Prigogine, 1997; Prigogine & Stengers, 1984), and evolutionary economics (Nelson & Winter, 1982) into a unified framework that grounds the dissertation's treatment of firms, markets, and economies as ongoing processes rather than static structures.

The section receives from Section 5.1 the ontological framework of emergent wholeness – the principle that macro-level properties arise from but are irreducible to micro-level interactions. Section 5.1 established that economic wholes possess causal powers of their own; this section adds the temporal dimension: emergent properties do not sit statically as complex substances but are continuously produced and reproduced through ongoing processes. The section provides to Section 5.3 the processual foundation upon which relational ontology builds. Once reality is understood as becoming rather than being, the question shifts from what entities essentially are to how they are constituted through relations over time. Section 5.3 will show that entities exist only through connection; this section establishes that those connections are themselves historical processes.

The section serves three purposes within the dissertation's architecture. First, it adds the temporal dimension to emergence – showing that emergent properties do not exist statically but continuously emerge through ongoing processes, thereby preventing the misreading of CAS as merely "complex substances." Second, it establishes the ontological impossibility of equilibrium: if reality is fundamentally processual, equilibrium (a static state) becomes not merely unlikely but categorically impossible, thereby grounding the dissertation's later transformation of dynamics. Third, it reconceptualises the firm as an ongoing organising process rather than a production function, preparing the framework contribution on firms-as-CAS.

The section proceeds as follows. Part I establishes the foundational contrast between substance and process ontologies, tracing the 2,500-year philosophical debate from Parmenides to

Heraclitus and its resolution through complexity science. Part II develops four principles of process ontology: events as primary (Whitehead, 1929/1978), time as irreversible (Prigogine, 1997), identity as trajectory (rather than essence), and stability as dynamic achievement (dissipative structures). Part III locates process ontology at the meso-level, showing how patterns stabilise across intermediate scales. Part IV addresses economic applications, demonstrating how process ontology grounds evolutionary economics and reconceptualises firms. Part V addresses counter-arguments. The conclusion synthesises the transformation and establishes forward momentum toward relational ontology (Section 5.3) and dialectical synthesis (Section 5.4).

Before proceeding, a central objection must be addressed. If everything is flux, how can we explain the manifest stability of economic institutions, organisational routines, or technological standards? Our framework addresses this through Prigogine's (1984) insight that stability in open systems is not static persistence but dynamic achievement – patterns maintained through continuous activity far from equilibrium. Firms persist not despite change but through change; stability requires ongoing reproduction. This transforms apparent objection into confirmation: the very effort required to maintain stability demonstrates that process, not substance, is ontologically primary.

5.2.1 Philosophical Foundations: Substance versus Process

I. The Substance–Process Debate, p. 2,500 Years of Philosophical Debate

The contrast between substance and process ontologies reflects one of Western philosophy's deepest and most consequential divisions. For over two millennia, the dominant tradition sided with Parmenides over Heraclitus – with Being over Becoming – shaping not only metaphysics but the foundational assumptions of economics itself. The vindication of Heraclitus through complexity science represents not incremental adjustment but fundamental paradigm correction.

Heraclitus and the Primacy of Flux

Heraclitus of Ephesus (c. 535–475 BCE) articulated what would become process philosophy's founding insight: "everything flows" (Τὰ πάντα ῥεῖ). His famous assertion that one cannot step into the same river twice captures process ontology's central claim – identity persists through continuous change, not despite it. As Rescher (1996, p. 38) emphasises, "processes are basic, and objects are derivative." Heraclitus's metaphoric conception of fire embodies the ever-changing, processual nature of reality, which views all things as transient manifestations of eternal cosmic dynamism (Rescher, 1996, pp. 35–40). The river maintains recognisable identity (same name, location, banks) yet consists of different water at each moment. Identity is pattern in flux, not static essence.

Parmenides, Aristotle, and Economics' Inherited Metaphysics

Parmenides, by contrast, argued that what is real must be unchanging – change is illusion, Being is eternal. Aristotle systematised this into substance ontology: reality consists of substances (*ousiai*) possessing essential properties that define their nature timelessly. This

framework dominated Western thought through medieval scholasticism, early modern rationalism, and classical physics. Economics inherited substance ontology wholesale: agents with fixed preferences, technologies as stable production functions, equilibrium as natural state. These are not technical assumptions but Parmenidean metaphysics applied to economic phenomena.

The relationship between Section 4.1.2's problematisation of substance ontology and this reconstruction is not merely sequential but dialectically reinforcing. Section 4.1.2 demonstrated that substance ontology fails – treating economic entities as self-contained bearers of intrinsic properties systematically misrepresents their relational and processual nature. As Lawson (2019, pp. 9–15) argues, mainstream economic models carry "ontological commitments that do not fit the nature of social material." This section explains **why** it fails: because reality is fundamentally process, and what we call substances are temporary stabilisations within ongoing flux. The critique gains explanation through process; process gains necessity from the demonstrated failures.




Figure 5: Figure 5.1. Substance and Process Ontology: Philosophical Foundations of Competing Economic Paradigms

Figure 5.1 exhibits Substance ontology (fixed essences; equilibrium as natural) which underwrites mainstream economics; process ontology (flux; time and change as constitutive) underwrites complexity economics. Complexity science vindicates the Heraclitean position: reality is process; “substances” are stabilised patterns.

The relationship between the critique in Chapter 4, Section 4.1.2 and this reconstruction is therefore explanatory, not merely sequential. Section 4.1.2 showed that substance ontology misrepresents social material; process ontology explains why: what appear as “substances” are temporary stabilisations within ongoing becoming. As Lawson (2019, pp. 9–15) argues, mainstream modelling carries ontological commitments that do not fit social reality.

Complexity Science Vindicates Heraclitus

Modern complexity science vindicates Heraclitus after 2,500 years. As Heylighen, Cilliers, and Gershenson (2006, p. 5) observe, "the first challenges to reductionism and its denial of creative change appeared in the beginning of the twentieth century in the work of process philosophers, such as Bergson, Teilhard, Whitehead." But it was complexity science that provided empirical vindication: self-organising systems, dissipative structures, and evolutionary dynamics all exhibit processual ontology.

Nobel laureate Ilya Prigogine and Isabelle Stengers identify mainstream economics as a science of "being" rather than "becoming." They argue that classical dynamics, with its focus on time-reversible laws and deterministic trajectories, "describes a static world, a world that escapes the arrow of time," thereby ignoring "the irreversible and creative processes that characterize living systems and, by extension, the economy" (Prigogine & Stengers, 1984, p. 291). The process of formation is fundamentally a story of "becoming" – of historical unfolding and genuine novelty – which is philosophically alien to a static Newtonian conception of reality.

Arthur (2014, p. 25) captures the result: complexity economics reveals "an economy that is not dead, static, timeless, and perfect, but one that is alive, ever-changing, organic, and full of messy vitality." In such systems, "structures then come and go, some stay to be further built on and some lead to further structures. The system is alive" (Arthur, 2014, p. 17). This is not metaphor but ontological description – the economy **is** process. As Arthur (2021, p. 3) observes, complexity economics reveals economies where "relative stasis alternate with ones of dynamic upheaval" – not equilibrium punctuated by shocks but continuous processual becoming.

Firm-level translation: Nokia

Nokia's trajectory from 1865 to 2024 demonstrates why substance ontology fails for firm analysis. Nokia began as a paper mill in southwestern Finland, became a rubber manufacturer (the "Nokia" boots familiar to generations of Finns), evolved into a cable and electronics company in the 1960s, transformed into a mobile phone pioneer in the 1980s, dominated the global handset market by 2007 with 40% market share and €51 billion in revenue, then catastrophically declined when smartphones emerged – losing 90% of its market value within five years and selling its handset division to Microsoft in 2014 for €5.4 billion.

Substance ontology would seek Nokia's unchanging "essence" – what made Nokia essentially Nokia across these transformations. But no such essence exists. Nokia in 1900 (rubber boots) shares nothing intrinsic with Nokia in 2000 (mobile phones) except accumulated trajectory: the ongoing process of becoming that carried forward certain capabilities, relationships, and constraints while continuously transforming others. Nokia's 2007 crisis cannot be understood through essence but only through trajectory – the path-dependent accumulation that enabled feature phone dominance simultaneously disabled smartphone response. The company's identity was its trajectory, not any timeless essence. By 2024, Nokia had reinvented itself again as a telecommunications infrastructure provider, demonstrating that firm identity is continuous becoming, not static being.

II. Four Principles of Process Ontology

Principle 1: Events, Not Objects

Alfred North Whitehead's *Process and Reality* (1929/1978) provides the most systematic articulation of process metaphysics. His framework – described by Simons (2015, p. 1) as "perhaps the most significant metaphysical treatise of our time" – rejects the notion that reality consists of enduring substances and argues instead that its fundamental units are "actual occasions": events of becoming that are relationally constituted. As Whitehead (1929/1978, p. 22) states, "the actual world is a process, and that process is the becoming of actual entities." Seibt (2023) summarises this position succinctly: reality is not a static assembly of objects but a continuum of intermingling processes.

This marks what Rescher (1996, pp. 27–31) identifies as the decisive ontological shift: becoming, not being, is primary. What we ordinarily call "objects" are, in this view, temporarily stabilised patterns of events – sequences of occasions exhibiting sufficient continuity that we abstract enduring identity from ongoing flux. This abstraction is useful, but it is ontologically derivative. Whitehead (1929/1978) warned against the "fallacy of misplaced concreteness": treating abstractions as more real than the processes that generate them.

Mainstream economics commits precisely this fallacy. When firms are modelled as production functions, agents as utility maximisers, or markets as equilibrating mechanisms, abstract representations are reified as real entities, while the organising processes that actually constitute economic life are ignored.

Process ontology reverses this error. Firms are not objects that organise; they are ongoing processes of organising – coordinating, learning, adapting, and reproducing themselves through time. A firm exists only insofar as this activity continues. Bankruptcy, dissolution, or organisational collapse does not destroy an underlying "thing"; it terminates a process. Likewise, markets are not fixed structures but continuous flows of trading, negotiation, and price formation. Suspend exchange, and the market ceases to exist as a market.

Firm-level translation: Amazon

Amazon's trajectory illustrates this ontology with particular clarity. What makes Amazon in 2024 "the same firm" as Amazon in 1994? Not an unchanging essence – the company has transformed from an online bookstore into a global retail platform, cloud-computing provider,

logistics network, entertainment studio, and healthcare entrant. Not physical continuity – its employees, facilities, technologies, and products have all changed repeatedly. What persists is trajectory: the continuous organising process that carries forward accumulated capabilities, commitments, and constraints while continually transforming them.

Amazon is not something that undergoes change; it is change organised through time. The acquisitions of Whole Foods (2017), the launch of Amazon Pharmacy (2020), and the development of Project Kuiper (2023) are not departures from an underlying essence but moments in an ongoing process of becoming. The firm's apparent stability – customer centricity, long-term orientation, experimental culture – is maintained through transformation, not despite it. Amazon exists as a process that continually reproduces itself by becoming something new.

Principle 2: Time Is Irreversible

Ilya Prigogine's Nobel Prize-winning work on dissipative structures provides the scientific foundation for process ontology. In *The End of Certainty*, Prigogine (1997) distinguishes reversible time (classical physics) from irreversible time (thermodynamics, biology, and economic systems). Classical mechanics treats time as a symmetric parameter: equations run equally well forward or backward. In contrast, complex systems unfold in irreversible time: the past accumulates, constrains the present, and opens or closes future possibilities. History does not wash out – it is carried forward.

Prigogine (1997) formalised this through the concept of bifurcation: in non-linear systems, small fluctuations can push a system toward radically different future states. As he explains, "at every step, the future is uncertain... it is impossible to speak of trajectories" (Prigogine, 1997, p. 49). This is not epistemic ignorance but ontological openness: the system genuinely does not contain its own future in latent form. Even with complete information at a given moment, long-term outcomes cannot be predicted because the system continually creates new structures through irreversible transitions.

Henri Bergson (1907/1998) anticipated this insight with his concept of *durée réelle* (real duration). For Bergson (1907/1998, pp. 10–12), time is not a spatial axis along which moments are arranged but a qualitative flow of becoming in which the past is conserved in the present and shapes what can come next. One cannot return to a prior moment because the system has been transformed by its own history. Evolution, Bergson (1907/1998) argued, is not the unfolding of pre-given possibilities but the continuous creation of novelty.

Applied to economics, this means that path-dependence is not accidental but ontological. The economy does not gravitate toward a timeless equilibrium; it moves through irreversible trajectories. Each decision leaves residues – skills, expectations, institutions, technologies – that constrain and enable future development.

This has direct empirical consequences. Consider the persistence of the QWERTY keyboard. Under substance ontology, QWERTY appears as a suboptimal equilibrium waiting to be replaced. Under process ontology, QWERTY is a crystallised historical trajectory: early adoption created typing skills, which generated learning curves, which produced network effects, which created switching costs, which locked in the technology. Each step accumulated

irreversibly. There is no "reset button" to return to a neutral starting point; the present is built out of sedimented history.

Firm-level translation: Silicon Valley

Silicon Valley's dominance exemplifies irreversible accumulation at regional scale. The ecosystem emerged from a sequence of historically contingent events that cannot be reproduced: William Shockley's decision to locate his laboratory in Mountain View in 1956; the departure of the "traitorous eight" in 1957 to found Fairchild Semiconductor; the subsequent spawning of Intel, AMD, and over a hundred other firms by Fairchild alumni; the formation of dense venture capital networks; the establishment of labour-market norms favouring mobility over loyalty. Each step built on what came before, compounding knowledge spillovers, talent flows, and capital formation.

By 2024, Silicon Valley firms accounted for roughly 30% of US venture capital investment and produced five of the world's ten most valuable companies (Apple, Alphabet, Meta, Nvidia, Tesla). No other region can "catch up" simply by replicating today's conditions because the Valley is not a template but a trajectory. Route 128 near Boston began with comparable technological assets in the 1970s, but accumulated differently – more hierarchical corporate cultures, weaker labour mobility, and fewer entrepreneurial spin-offs – and permanently fell behind. The divergence reflects not policy failure but irreversible historical branching.

Industrial policy repeatedly fails for the same reason. Attempts to "create the next Silicon Valley" assume that success can be copied by assembling the right components. Process ontology shows why this fails: ecosystems are not constructed but grown through irreversible sequences of interaction. Silicon Valley is not a thing that can be replicated – it is a process that cannot be reset.

Principle 3: Identity as Trajectory, Not Essence

Process ontology reconceptualises identity from essence (what something timelessly *is*) to trajectory (what something *becomes* through time). Under Aristotelian substance ontology, an entity possesses an intrinsic nature that determines its development – an acorn contains "oak-ness" that causes it to become an oak. Process ontology rejects this metaphysics: the acorn's identity is not an inner blueprint but its developmental pathway – acorn → sapling → tree – shaped by environment, contingencies, and interactions. Identity is therefore historical and narrative, not fixed and intrinsic.

Darwin's (1859) theory of evolution introduced precisely this ontological transformation into biology. Pre-Darwinian thought treated species as fixed essences; post-Darwinian biology treats species as evolving populations defined by lineages, variation, and selection. A "bird" is not a timeless form but a phylogenetic trajectory. What something is depends on how it has become. Economics requires the same ontological revolution. Preferences develop; technologies evolve as design trajectories; industries branch; institutions co-evolve. Identity at every level – individual, organisational, industrial, institutional – is path-dependent and relationally produced.

This is especially clear at the meso-level. Firms are not bundles of assets or functions but historical organisms whose present capabilities reflect their accumulated past. Individual trajectories are shaped by organisations, and organisational trajectories are shaped by individuals; identity is therefore relationally co-produced.

Firm-level translation: Nokia versus Apple

The contrast between Nokia and Apple demonstrates trajectory-based identity. Nokia's path – rubber products → cables → electronics → feature phones → smartphone disruption → infrastructure – produced a firm whose capabilities were optimised for hardware-centric mobile devices and whose operating systems were poorly suited to touch-based platforms. Apple's path – personal computing → graphical interfaces → digital music → smartphones → ecosystems → services – generated a firm capable of integrating software, design, and platform coordination.

In 2007 Nokia was the world's largest handset manufacturer and Apple a niche computer company. Yet their accumulated trajectories made radically different futures possible. Nokia's Symbian platform and engineering culture could not pivot to software-centric ecosystems; Apple's iTunes and design heritage made the App Store and iPhone natural extensions. Their fates diverged not because of different "essences," but because of different historical becomings. Identity was destiny because identity was trajectory.

Principle 4: Stability as Dynamic Achievement

Process ontology completes its logic with Prigogine's (1984; 2003) concept of dissipative structures. Systems far from equilibrium can maintain organised form only through continuous throughput of energy, matter, and information. As Capra and Luisi (2014, p. 82) note, living and social systems are open systems whose persistence depends on ongoing exchange with their environment.

A flame exists only while fuel flows. A vortex exists only while water moves. An organism exists only while metabolism continues. Stability, in such systems, is not stasis but dynamic maintenance.

Kitto (2014, p. 543) shows that complex systems remain stable only through continuous environmental engagement; Prigogine (2003) further demonstrates that non-equilibrium systems become more sensitive to their surroundings, enabling adaptation and self-organisation. Langley et al. (2013) summarise this insight: stability is not the absence of change but the result of sustained activity that counteracts entropy.

This inverts the core intuition of substance ontology. Under substance thinking, stability is natural and change requires explanation. Under process ontology, change is natural and stability requires explanation. Persistence is an achievement.

For economics, this has decisive implications. Markets exist only while trading occurs. Firms exist only while organising occurs. Institutions exist only while they are enacted. Equilibrium

– defined as a state of rest – is not a special case but a non-state. What exists is dynamic order maintained through continuous flow.

Firm-level translation: Toyota

Toyota's Production System illustrates dynamic stability. TPS is not a fixed organisational structure but a living process: continuous improvement (kaizen), just-in-time coordination, problem escalation through andon, and dense supplier learning networks. As Dyer and Nobeoka (2000) demonstrate, Toyota's advantage resides in a knowledge-sharing network sustained through ongoing interaction, not in static routines.

Competitors could copy Toyota's tools but not its capability because the capability exists only as ongoing practice. Toyota's response to its 2010 quality crisis further illustrates this: the firm did not return to a prior "essence" but reconfigured its processes to restore stability through adaptation. By 2024 Toyota remained the world's largest automaker by volume, not because it possessed a fixed advantage, but because it continuously re-earned its position through processual renewal.

5.2.2 Economic Applications and the Meso-Level

III. Process Ontology and the Meso-Level

The meso-level – the intermediate analytical scale between micro-level agents and macro-level aggregates – emerges as the privileged domain of economic reality under process ontology. If reality consists of evolving processes rather than static substances, and if these processes stabilise into organised patterns through interaction, then it is at the meso-level that becoming takes form. This is where trajectories intersect, where coordination stabilises, and where emergence becomes empirically observable.

Dopfer and Potts (2004, p. 263) capture this precisely: economic evolution proceeds through rules, rule systems, and rule processes. These do not exist at the level of isolated individuals, nor at the level of aggregate outcomes, but in the organised structures that connect them. Arthur (2013, p. 2) similarly shows that agents act within rule-structured environments that they simultaneously reproduce and transform. This recursive dynamic – agents shaping structures, structures shaping agents – unfolds at the meso-level.

Process ontology therefore explains why the meso-level cannot be eliminated through aggregation. If agents were self-contained substances with fixed properties, aggregation might preserve essential information. But if agents are processes constituted through interaction, then aggregation destroys the very dynamics that generate behaviour. Preferences, strategies, routines, and expectations are not intrinsic to agents; they are formed, modified, and stabilised through meso-level structures such as organisations, networks, and institutions.

This makes methodological individualism ontologically incoherent. By collapsing economic analysis directly from individuals to aggregates, it removes the level at which agents are actually constituted. Firms, industries, and markets are not epiphenomena; they are the very

sites where economic agency is formed. The meso-level is not an optional analytical convenience – it is the ontological locus of economic process.

Firms exist as stabilised patterns of organising activity. Industries exist as evolving ecologies of firms, technologies, and regulatory constraints. Markets exist as continuously reproduced trading networks. These are neither micro choices nor macro totals; they are processual structures that generate both.

Firm-level translation: The pharmaceutical industry

The global pharmaceutical industry illustrates meso-level process ontology with exceptional clarity. Its present structure – roughly fifteen multinational firms controlling about 60% of global revenues, alongside thousands of biotech start-ups pursuing early-stage research – did not result from individual optimisation or aggregate forces alone. It emerged through decades of recursive institutional and organisational evolution.

Patent law created temporary monopolies that made high-risk R&D financeable. Regulatory systems (FDA, EMA) imposed long, structured development pathways. Merger waves – Pfizer–Warner-Lambert (2000), Pfizer–Wyeth (2009), Bristol-Myers Squibb–Celgene (2019) – consolidated pipelines and market access. The blockbuster drug model shaped investment incentives. Biotech partnerships created dense innovation networks linking small research firms to large global distributors.

None of this exists at the level of any individual firm, and none appears in aggregate health-care spending. It exists at the meso-level as an emergent industry architecture that both constrains and enables firm behaviour. New entrants must navigate patent thickets, regulatory bottlenecks, and partnership networks that pre-exist them. At the same time, each firm's strategic moves – alliances, acquisitions, technological bets – continuously reshape that structure.

The COVID-19 vaccine race revealed this recursive dynamic. Pre-existing meso-level structures – Pfizer–BioNTech, Moderna's mRNA platform, Oxford–AstraZeneca – made rapid development possible. The pandemic then transformed the industry itself, accelerating regulatory pathways, funding platforms, and reorienting R&D priorities toward mRNA technologies.

The pharmaceutical industry therefore exists neither as a collection of firms nor as an aggregate market, but as a processual meso-level system whose structure is both the product and the determinant of firm trajectories. This is exactly what process ontology predicts.

IV. Economic Applications: From Evolutionary Economics to Firm Reconceptualisation

Process ontology provides the metaphysical foundation for evolutionary economics. Nelson and Winter's (1982) framework does not merely offer an alternative modelling strategy; it presupposes a different ontology of what firms and markets are. Firms are not optimising entities selecting from fixed choice sets but historical trajectories guided by routines, heuristics, and search processes operating under fundamental uncertainty. As Nelson and Winter (1982,

p. 4) state, firms are "motivated by profit" but not assumed to maximise it "over well-defined and exogenously given choice sets." Instead, firms possess evolving capabilities that change through learning, experimentation, and chance, while markets select among competing trajectories through differential growth and survival (Nelson & Winter, 1982, p. 9).

This is not metaphor. It is process ontology in operational form: firms are paths, not points; competition is selection, not equilibrium; growth is structural transformation, not marginal adjustment.

The relationship between process ontology and evolutionary economics is therefore bidirectional. Process ontology explains why economic systems evolve rather than equilibrate. Evolutionary economics empirically demonstrates that they do. The conclusion is not that evolutionary economics is merely a useful modelling variant – it is that it is ontologically correct for a reality that is itself processual.

Creative Destruction as Process Ontology in Action

Schumpeter's (1942) "perennial gale of creative destruction" is the economic expression of process ontology. Where equilibrium economics presumes a world of stable structures occasionally perturbed by shocks, Schumpeter (1942), Nelson and Winter (1982), North (1990), and complexity economists describe an economy in continuous re-formation.

Arthur (2014, p. 6) makes this explicit: novelty enters from within the system through "new combinations of productive means." These are not disturbances of equilibrium; they are the engine of history itself. As Arthur (2021, p. 137) shows, equilibrium frameworks "make no allowance for the creation of new products or new arrangements... indeed, for history itself." Technologies form by recombination (Arthur, 2021, p. 141), triggering cascades of institutional, organisational, and market restructuring that keep the system permanently far from equilibrium.

The economy, in this view, is not a machine seeking rest but a self-creating process: "alive, ever-changing, organic" (Arthur, 2021, p. 136). This is not rhetoric. It is the direct implication of irreversible time, endogenous novelty, and feedback-driven dynamics.

Firm-level translation: Blockbuster versus Netflix

Blockbuster and Netflix did not compete as two firms with different "strategies." They competed as two trajectories with different accumulated pasts.

Blockbuster's dominance in 2004 – 9,000 stores and \$6 billion in revenue – was the product of a trajectory built on retail locations, late-fee economics, and physical distribution. Those same accumulated structures made streaming existentially threatening: adopting Netflix-style delivery would have required destroying its own profit engine. The organisation could not reverse its own history.

Netflix's trajectory accumulated differently: data-driven recommendations, logistics, subscription relationships, and then streaming infrastructure. When the technological regime shifted, Netflix's past became an asset while Blockbuster's past became a trap.

No "essence" explains this. Only process does. Each firm was what it had become. Creative destruction is not an event – it is the unfolding of incompatible trajectories.

Formation Over Allocation

Process ontology in economics implies that formation precedes allocation.

Arthur (2014, p. 22) distinguishes the allocation problem (how given goods and prices are arranged) from the formation problem – how the economy itself comes into being, transforms, and grows. The latter is characterised by fundamental uncertainty because the future is not given; it is being created (Arthur, 2014, p. 5).

If reality is process, then technologies, markets, and organisations do not exist as fixed parameters awaiting optimisation. They emerge through variation, selection, and recombination. Allocation becomes derivative: one can only allocate what formation has already brought into existence.

This is why the dissertation's subsequent shift of technology from exogenous factor to endogenous formation follows necessarily. In a processual world, nothing that matters most is given in advance.

Firm-level translation: SpaceX

SpaceX did not reallocate resources within the existing aerospace industry. It formed a new one.

Reusable rockets, private launch services, satellite megaconstellations – none of these existed as allocable factors in 2002. They were created through an irreversible trajectory of trial, failure, learning, and recombination. Falcon 1 failed three times before succeeding. Falcon 9 required years of iterative landing attempts. Starship is still forming.

By 2024, SpaceX had created new cost structures, new markets, and new strategic possibilities. Traditional aerospace firms could not "optimise" their way into this position because it required becoming something they were not.

SpaceX is not an asset bundle. It is a trajectory of formation.

What This Section Has Now Done

This section completes the ontological shift: Economics no longer studies how fixed objects are arranged. It studies how new structures come into existence.

Firms are not entities with strategies. They are processes that become.

Markets are not equilibrium states. They are histories unfolding.

This is the true meaning of evolutionary economics once grounded in process ontology.

5.2.3 Counter-Arguments and Responses

1. "Process philosophy is too abstract for economic analysis."

Response. This objection confuses metaphysical abstraction with analytical emptiness. Process ontology is not a metaphor; it is operationalised across modern economics and management theory.

Dopfer, Foster, and Potts (2004, pp. 263–279) translate process ontology into the micro–meso–macro framework in which rules originate, diffuse, and stabilise through measurable evolutionary dynamics. Nelson and Winter (1982, pp. 14–21) operationalise process through organisational routines – empirically observable patterns of coordinated action that change through variation and selection. Teece, Pisano, and Shuen (1997, pp. 509–533) operationalise process as dynamic capabilities – firms' measurable capacities to sense, seize, and reconfigure resources. Arthur (1994, pp. 1–12) operationalises process through path-dependence and increasing-returns models that generate testable predictions about lock-in, technological dominance, and multiple attractors.

The convergence of evolutionary economics, complexity economics, and strategic management on these same processual structures is not accidental. It reflects their shared ontological commitment: economic reality is constituted by evolving processes, not static objects.

2. "Equilibrium analysis has been enormously successful – process ontology would discard it."

Response. This objection confuses local success with ontological validity.

Equilibrium analysis works where its assumptions hold: stable preferences, given technologies, weak feedback, and near-decomposable systems. Hausman (1992, pp. 115–138) and Mirowski (1989, pp. 218–231) show that modern economics inherits these assumptions directly from nineteenth-century physics.

But these assumptions collapse in domains characterised by innovation, structural change, financial contagion, and institutional evolution. Arthur (2014, pp. 172–175) demonstrates that equilibrium theory is structurally incapable of handling formation processes. The 2007–2008 financial crisis revealed this decisively: equilibrium models could not generate systemic crashes or prolonged unemployment without ad hoc repairs. As Kirman (2010) observed, the crisis was "both a crisis of the real economy and of the dominant theory."

Process ontology does not abolish equilibrium. It limits it to the domains where it applies and provides the correct ontology for the domains where it fails.

3. "If everything is process, nothing is stable – science becomes impossible."

Response. This objection mistakes static stability for scientific stability.

Process ontology does not deny stability; it reconceptualises it as dynamic achievement. Prigogine's (1984) dissipative structures show how stable patterns persist through continuous flux (Prigogine & Stengers, 1984, pp. 131–153). Living systems remain organised only through ongoing metabolism (Capra & Luisi, 2014, pp. 80–95). Organisations reproduce routines through continuous enactment (Langley et al., 2013).

Economic stability is of the same kind: firms, markets, and institutions persist only through ongoing activity. Scientific knowledge concerns the laws of these processes – how stability is produced, disrupted, and transformed – not static snapshots.

Biology, ecology, and evolutionary theory are sciences of process. Economics is no different.

4. "Process ontology makes prediction impossible."

Response. This objection applies the wrong epistemic standard.

Prigogine (1997, pp. 45–52) demonstrated that complex systems exhibit intrinsic unpredictability due to bifurcations and feedback. This does not destroy science – it changes what science provides.

Rescher (2000, pp. 185–201) distinguishes prediction (exact future states) from preparation (understanding dynamics and navigating uncertainty). Process ontology enables the latter by revealing causal mechanisms, attractor structures, path dependencies, and phase transitions.

Equilibrium economics offers false precision; process ontology offers reliable understanding. The 2008 crisis showed which is more useful.

VI. Conclusion

This section has established process ontology as the temporal pillar of CAS ontology.

Where Section 5.1 showed that economic systems are emergent, this section shows that both micro and macro are themselves processes – ongoing activities rather than static entities. Together, emergence and process yield a unified CAS ontology: economic reality consists of emergent processes operating across multiple scales, continuously becoming rather than statically being.

Transformation Map

Table 117: Comprehensive Transformation Map: Reconceiving Ten Fundamental Economic Concepts from Substance Ontology to Process Ontology

Dimension	FROM (Substance Ontology)	TO (Process Ontology)
Fundamental units	Enduring objects	Events, activities, becomings
Time	Reversible parameter	Irreversible accumulation
Identity	Timeless essence	Historical trajectory
Stability	Natural state	Dynamic achievement
Change	Deviation	Constitutive
Equilibrium	Natural attractor	Ontologically impossible
Firm	Production function	Ongoing organising process
Market	Fixed structure	Continuous trading
Analysis	Allocation	Formation
Knowledge	Prediction	Understanding + navigation

The table above synthesises the ontological shift developed throughout this section. At its foundation lies Whitehead's (1929/1978) reconceptualisation of reality's basic constituents: not enduring objects but events, not static substances but occasions of becoming. From this starting point, the remaining transformations follow with a certain logical necessity.

Consider first the temporal dimensions. Prigogine's (1997) work on bifurcations establishes that complex systems carry their history forward irreversibly; time is not a neutral parameter through which systems move but an accumulating burden and resource that shapes what futures remain possible. This irreversibility grounds the reconceptualisation of identity itself. Just as Darwin's (1859) revolution replaced biological essentialism with phylogenetic lineage, process ontology replaces the timeless "what" of economic entities with the historical "how" of their becoming – a principle the contrasting trajectories of Nokia and Apple vividly illustrate.

The implications for stability are equally profound. Prigogine's (1984) dissipative structures demonstrate that organised patterns in open systems persist only through continuous activity; stability is therefore not a natural resting state but a dynamic achievement requiring ongoing reproduction. This inverts the explanatory burden: under substance ontology, change demands explanation while persistence is assumed; under process ontology, persistence demands explanation while change is the default condition. Equilibrium, understood as static rest, becomes not merely empirically rare but ontologically incoherent – a category error applied to systems whose very existence consists in continuous flow.

These ontological shifts restructure economic analysis at every level. Firms can no longer be treated as production functions – mathematical objects awaiting optimisation – but must be understood as ongoing organising processes that cease to exist when coordination, learning, and adaptation terminate. Markets similarly transform from fixed institutional structures into continuously reproduced patterns of exchange. The analytical focus consequently shifts from allocation (arranging given factors efficiently) to formation (understanding how new structures come into existence). And the epistemic goal shifts accordingly: since Prigogine (1997) and

Rescher (2000) establish that complex systems exhibit intrinsic unpredictability at bifurcation points, the pursuit of precise prediction gives way to the cultivation of understanding sufficient for navigation through irreducibly uncertain terrain.

What Has Been Achieved

The transformation is not methodological but ontological:

- Economics has inherited a metaphysics of Being.
- Complexity science reveals a reality of Becoming.

Firms are not entities that change. They are changes that stabilise.

Markets are not structures. They are ongoing coordination processes.

This correction is as deep as Darwin's (1859) correction of biological essentialism. Economics must undergo the same shift.

The next section develops relational ontology. Emergence showed that wholes have causal power. Process showed that those wholes are historical. Relation will show that entities exist only through connection.

Together, these yield the CAS ontology required for a post-mechanistic economics.

5.3 Substances as Stabilised Processes

Introduction

This section provides the dialectical synthesis that resolves the central tension raised by process ontology. It defines substances as stabilised processes: the ontological position that persistent entities – firms, markets, institutions – are not static bearers of intrinsic properties but temporarily stable patterns maintained through continuous processual activity. As Prigogine and Stengers (1984, p. 291) describe it, complex systems exhibit "*the being of becoming*" – structures that exist only insofar as the processes that generate and sustain them continue.

The section receives from Section 5.2 the processual foundation that dissolved static substance ontology. Section 5.2 established that economic reality is continuous becoming – learning, innovating, coordinating, trading – not a collection of unchanging things. But this immediately raises a puzzle: if all is flux, how does anything persist? How does "Apple" exist for fifty years if firms are process?

The section provides to Section 5.4 the stabilisation mechanism upon which relational ontology builds. Once we understand that entities persist through continuous reproduction rather than static essence, Section 5.4 will add the crucial relational dimension: the processes through which entities reproduce themselves are not self-contained but constituted by network connections. Stabilised processes are relationally stabilised processes.

This formulation resolves the apparent opposition between substance ontology (Being, persistence) and process ontology (Becoming, flux): not Being *or* Becoming, but Being *through* Becoming.

The section serves three functions in the dissertation's architecture:

1. It resolves the tension between Section 4.1.2 (the failure of substance ontology) and Section 5.2 (the primacy of process). Without synthesis, the critique risks collapsing into pure flux where nothing persists.
2. It establishes the stabilisation mechanism that emergence and process require. Emergence (5.1) showed that wholes have irreducible properties; process (5.2) showed that reality is becoming. This section shows how emergent, processual reality generates durable patterns. Section 5.4 will complete the picture by adding the relational dimension.
3. It grounds the theory of the firm as both persistent and transforming – the central empirical puzzle of strategy, institutional economics, and evolutionary theory.

Before proceeding, we must confront the core problem: If reality is process, how does anything persist? If firms are flux, how does "Apple" exist for fifty years?

The answer is the heart of CAS ontology: stability is not a given – it is an achievement. Order is not imposed on process; it is produced by it.

5.3.1 Philosophical Foundations: From Atoms to Relations

I. The Atomistic–Relational Contrast

The contrast between atomistic and relational ontology defines the core transformation of this chapter.

Atomistic ontology treats entities as self-contained bearers of intrinsic properties, with relations merely linking already-formed units. This metaphysics underlies methodological individualism. As Hodgson (2001, p. 28) notes, social explanation is thereby reduced to "the properties of its constituent individuals," treated as the "primary unit in the social universe." The representative agent exemplifies this stance: preferences and capabilities are given; interaction does not alter what the agent *is*. Section 4.1.3 demonstrated why this ontology fails empirically and logically. This section provides the positive alternative.

Relational ontology holds that entity properties are positional. What an agent knows, can do, can influence – and indeed what the agent *is* – depends on its location within relational structures. Fontana (2008a, p. 10) captures this with the principle of connective complexity: "it is the kind of relationship that links the elements of the system to one another that shapes their behaviour." Relations are not constraints acting on intrinsic properties; they are constitutive of those properties. The shift from atomism to relationality is therefore not methodological preference but ontological necessity.

Firm-Level Translation: Apple

Apple's identity in 2024 is not defined by intrinsic attributes – its headquarters, employee count, or revenue – but by its relational position as hub of a dense, multi-layered ecosystem:

- over 2 million iOS developers generating approximately \$1.1 trillion in commerce,
- Foxconn, employing over 1.2 million workers producing Apple hardware,
- TSMC, fabricating Apple's chips at the technological frontier,
- a global logistics, retail, and platform infrastructure, and
- two billion active devices embedding users in the ecosystem.

Change these relations and Apple becomes a different firm. When Apple replaced Intel with in-house silicon in 2020, the firm did not merely swap a component – it changed its relational architecture, transforming from a PC assembler into a vertically integrated platform producer. Apple's market capitalisation reflects not asset ownership but ecosystem centrality.

Apple is not a substance that occupies a network. It **is** its network position.

Implication

Network topology is not merely descriptive – it is constitutive. As Kirman (2016, p. 560) observes, "the structure of the financial network is at least as important as the health of the individual institutions." Systemic risk, innovation capacity, market power, and strategic advantage all arise from patterns of interconnection, not from isolated node attributes.

This completes the ontological inversion: Economic reality is not made of interacting atoms. It is made of networks that generate entities.

II. Connective Complexity: Relations as Causally Primary

If relations are ontologically primary, then causality resides in relational organisation rather than in the intrinsic properties of elements. This is the principle of connective complexity: system behaviour is generated by the pattern of connections linking elements, not by the elements taken in isolation. As Fontana (2010, p. 13) states, "observable macro-regularities and self-organization in CASs derive from the relationships that link the elements together."

This directly inverts the reductionist logic of methodological individualism. Fontana (2008a, p. 10) formulates the critique precisely: "It is impossible to study the properties of complex systems by observing its elements in isolation." The reason is not merely practical difficulty but ontological mismatch: the properties of interest are not located in the elements but in the relations between them. Decomposition therefore destroys the very structures that generate the phenomena under investigation.

Connective complexity provides the causal mechanism underlying emergent wholeness (Section 5.1). Emergence established that macro-properties arise from micro-interactions yet remain irreducible to them. Connective complexity explains why: it is relational organisation that produces macro-level order. Emergence tells us *that* wholes exceed parts; connective complexity tells us *how* – through the causal power of patterned interconnection. Emergence is therefore not mysterious addition but relational constitution.

Morin's systems theory makes this explicit. As Estrada (2023, p. 1160) notes, "the interaction between two parts already produces a new quality." Interactions are not neutral links between pre-existing entities; they are transformative. What Estrada (2023, p. 1155) calls "Morinian interactions" generate bidirectional non-separability: entities are changed by the relations in which they participate and cannot be returned to a pre-relational state without altering what they are. This is why, as Nason (2017, p. 16) observes, "complex systems cannot be separated into parts to be analyzed independently because their components are interconnected and interdependent."

Arthur's (2014; 2021) complexity economics expresses the same ontological shift. The economy is not a machine composed of fixed parts but a relationally constituted process in which entities continuously form and reform through interaction (Arthur, 2014, p. 23; Arthur, 2021, p. 143). Transactions do not merely link agents; they constitute them.

Lawson (2019, pp. 9–15) therefore identifies atomism as the deep error of mainstream economics: it models agents as self-contained units when they are in fact participants in "dynamic, emergent networks." Methods fail not because they are technically flawed but because they rest on an ontology that does not fit social reality.

The implication is decisive:

Economic entities do not exist first and then interact. They come into existence through interaction.

Entrepreneurial capability, for example, is not an intrinsic personal trait but a positional effect: it emerges from network access to information, finance, legitimacy, and complementary resources. The "entrepreneur" is constituted by relational position, not born with intrinsic gifts.

Firm-Level Translation: Toyota's Keiretsu

Toyota's competitive advantage does not reside in Toyota alone. It exists in the relational architecture linking Toyota to its 200+ first-tier suppliers and thousands of lower-tier firms. Dyer and Nobeoka (2000) show that Toyota's quality, flexibility, and productivity arise from long-term, trust-based, knowledge-sharing ties – engineers embedded in supplier firms, joint problem-solving routines, and mutual investment.

These are not "contextual supports" for intrinsic capabilities. They are the capabilities.

When General Motors attempted to copy Toyota's practices while retaining arm's-length supplier relations, performance collapsed. The routines transferred; the relational field did not. Remove the network, and the capability disappears.

Toyota's advantage exists in the network topology, not in the node called Toyota.

Firm-Level Translation: Silicon Valley

Silicon Valley's innovative capacity is not possessed by Apple, Google, or Nvidia individually. It emerges from a dense, small-world, weak-tie network linking entrepreneurs, venture capitalists, engineers, universities, and legal institutions.

Saxenian (1994) showed that Silicon Valley outperformed Route 128 not because of superior firms but because of superior relational structure: porous boundaries, rapid knowledge diffusion, and cooperative competition. Network topology – not firm attributes – generated innovation.

By 2024, Silicon Valley produced five of the world's ten most valuable firms and roughly 30% of US venture capital – not because of exceptional atoms, but because of exceptional relational geometry.

Ontological Conclusion

Connective complexity completes the inversion begun in Sections 5.1, 5.2, and 5.3:

Causation flows through relations, not through isolated entities. Nodes are products of networks.

What economics has treated as "agents," "firms," and "markets" are stabilised patterns of relational activity. Their properties are not inside them; they are between them.

III. Structural Realism: Patterns Over Particles

If relations are primary and connective complexity generates causal efficacy, then structure matters more than composition. This is the core of structural realism: reality consists of relational patterns, with entities functioning as temporary occupants of positions within those patterns. What systems do depends more on how elements are connected than on which elements they are.

Network science formalises this insight. As Newman (2018, p. 7) observes, "the pattern of interactions can have a big effect on the behavior of a system." In networks, dynamics are governed by topology – the configuration of links – rather than by node attributes. Two networks with the same topology but different occupants will display similar diffusion, concentration, resilience, and fragility. Structure, not substance, governs behaviour.

Three core topological properties illustrate this:

Scale-free structure. Preferential attachment creates power-law degree distributions: a few hubs accumulate most links. This produces "rich-get-richer" dynamics visible in firm sizes, citation counts, wealth, and platform dominance. The mechanism is structural: new nodes attach to well-connected nodes. Replace Google with any firm in the same hub position and similar dominance would emerge. Power flows from topology, not talent.

Small-world structure. High clustering combined with short path lengths enables rapid diffusion across local communities. Silicon Valley's innovative capacity arises from this topology: dense local clusters linked by bridging ties allow ideas to circulate quickly and recombine. This is a structural property of the network, not a property of any firm or engineer.

Centrality and brokerage. Betweenness centrality – lying on paths between otherwise disconnected actors – creates control over information and resources. Brokerage power arises from **position**, not from intrinsic capabilities. PageRank formalises this logic for the web; financial hubs formalise it in interbank networks.

Structural Causality in Finance

The 2008 crisis exposed the failure of atomistic economics. Individual banks could appear solvent while the system was catastrophically fragile. Lehman did not collapse because it was especially incompetent; it collapsed because its position made it a conduit for contagion.

Kirman (2016, pp. 555–580) shows that systemic importance depends on network position, not size. "Too interconnected to fail" is not a metaphor; it is a topological fact. AIG was bailed out not because it was the largest insurer but because it was a hub in the CDS network, sitting between major banks. Its failure would have propagated losses through hundreds of counterparties regardless of their individual balance sheets. Haldane and May (2011) demonstrated that network topology – dense, homogeneous, highly interconnected – generated systemic fragility invisible to node-level risk models.

The crisis was therefore not only a macro failure but an ontological one: it revealed that risk resides in structure, not in institutions.

Firm-Level Translation: Amazon

Amazon's dominance is not explained by assets or headcount but by network position across multiple interlocking platforms:

- ~300 million customers connected to ~2 million sellers
- AWS hosting ~32% of global cloud infrastructure
- Logistics networks reaching ~90% of U.S. households in two days
- Data flows linking retail, advertising, media, and cloud services

These are structural advantages. Any firm occupying equivalent positions would wield comparable power. When Amazon enters healthcare, advertising, or finance, it does so by reconfiguring existing network positions, not by deploying intrinsic capabilities. Amazon is not a firm in markets; it is a hub across markets.

Its valuation capitalises topological centrality, not physical capital.

Process Meets Structure

Structural realism might suggest static frameworks – but process ontology (Section 5.2) and the synthesis of Section 5.3 prevent this. Network structures are not fixed. They are continuously reproduced and transformed through activity. As Giddens' (1984) structuration will formalise, agents enact the very structures that constitute them.

Structures are therefore historical achievements, not timeless constraints – meta-stable patterns that persist through continuous relational reproduction.

5.3.2 Economic Applications: Embeddedness and Structuration

IV. Embeddedness: Economic Action as Relational

Relational ontology becomes concrete in the concept of embeddedness. Economic action does not occur in a vacuum; it is constituted by ongoing social and organisational relationships.

Granovetter (1985) showed that the neoclassical "undersocialised" agent does not exist. What actors know, what they trust, what they can enforce, and what opportunities they see all depend on network position.

Embeddedness is not context – it is constitution. Three mechanisms illustrate this:

Trust and cooperation. Repeated interaction and reputational visibility enable cooperation that is impossible in atomistic, one-shot games. Silicon Valley's "co-opetition" exists because dense networks make defection costly and collaboration valuable.

Information diffusion. Weak ties bridge clusters and carry novel information (Granovetter, 1973). Innovation, hiring, and opportunity flow through relational pathways, not through price signals alone.

Relational enforcement. Most contracts are incomplete. Compliance is maintained through reputation, reciprocity, and future interaction – these are network properties, not merely legal enforcement.

Firm-Level Translation: Alibaba

Alibaba is not a retailer; it is a relational infrastructure.

By 2024, its platforms embedded ~10 million merchants and ~900 million consumers in trust-mediated networks:

- Ratings and reviews generate reputational discipline
- Alipay escrow creates transactional trust
- Sesame Credit transforms relational data into financial access

None of these are owned by any single participant. They exist between participants as network properties. Alibaba's valuation reflects embedded relational capital, not inventory.

Firm-Level Translation: Pharma Innovation

Biopharmaceutical innovation is irreducibly relational. Powell et al. (1996) showed that biotech breakthroughs emerge from weak-tie networks linking universities, startups, CROs, and big pharma.

Moderna's COVID-19 vaccine was not a solo achievement. It depended on NIH partnerships, manufacturing alliances, and clinical trial networks. Innovation was a network accomplishment.

Ontological Payoff

Structural realism and embeddedness complete the relational turn:

Economic power, innovation, risk, and identity arise from network position – not from intrinsic attributes.

Markets, firms, and agents are stabilised configurations of relations. To analyse them as isolated units is not simplification. It is ontological error.

V. Structuration: Relations as Recursive

Relational ontology raises an apparent paradox. If agents are constituted by relations, where does agency come from? If structures determine outcomes, how can structures ever change?

Anthony Giddens' (1984) structuration theory resolves this by identifying a deeper ontological fact: structure and agency are mutually constitutive through recursive practice.

Giddens (1984, pp. 25–26) defines structure as "both the medium and the outcome of social practices." Structures are not external frameworks imposed on pre-formed actors. They exist only insofar as they are enacted. At the same time, agents can act only through the structural relations that constitute them. This yields recursive relationality:

Relations shape agents → agents enact relations → enacted relations reproduce or transform structure.

Structure is not a container. It is an ongoing process of reproduction.

Structuration Meets Process Ontology

Process ontology (Section 5.2) established that reality consists of ongoing activity, not static being. Section 5.3 showed how this activity generates persistent patterns through continuous reproduction. Structuration applies these insights to relational structure itself.

Markets, firms, institutions, and networks do not exist as fixed architectures that agents move within. They exist as historically reproduced patterns of interaction. Every economic act simultaneously:

1. Draws upon existing relational structures (prices, platforms, norms, contracts, roles), and
2. Reproduces or alters those structures through its execution.

Trading reproduces markets. Hiring reproduces labour markets. Platform use reproduces platform power. Innovation reconfigures networks.

This is CAS in ontological form: feedback between structure and behaviour.

Agency Without Atomism

Relational ontology does not dissolve agency into structure.

Agents have real causal power – but that power is positional rather than intrinsic. Network positions define opportunity sets, visibility, leverage, and constraint. Within these structured possibility spaces, agents exercise genuine choice.

Agency is therefore not freedom from structure. It is freedom within and through structure.

Entrepreneurship, relationally understood, is the creation of new positions:

- bridging disconnected networks
- forming new alliances
- reconfiguring flows of information, trust, or resources

To act is to reshape the topology in which one is embedded.

Three Ontological Implications

1. Path-dependence is relational. Structures accumulate through recursive reproduction. Past interaction patterns become present opportunity spaces. Silicon Valley is not innovative because of genius – it is innovative because decades of recursive structuration produced dense, high-mobility, high-trust networks that now constitute innovative capacity. History is not background. It is the structure.

2. Transformation is endogenous. Structures are constraining – but not determining. New relational links can always be created. Entrepreneurs and organisations change systems by re-

wiring them. Structural holes can be bridged. Platforms can be re-layered. Standards can be displaced. Change happens when agents exploit relational degrees of freedom.

3. Co-evolution is inevitable. Firms and networks evolve together. Firms are not entities that operate in environments; they are patterns in evolving networks. Institutions do not govern behaviour; they are reproduced through behaviour.

Firm-Level Translation: Netflix

Netflix entered pre-existing structures: studio licensing networks, theatrical release windows, cable distribution, viewer habits.

But through streaming, original content, and algorithmic recommendation, it re-wired those relations:

- From studio → theatre → viewers to studio → platform → viewers
- From broadcasters to direct subscriber relationships
- From programming schedules to binge-release dynamics

Netflix was enabled by structure – but it also transformed structure. The firm and the industry evolved together.

Netflix is not a company that disrupted television. It is a new network architecture that television reorganised itself around.

Firm-Level Translation: Tesla

Tesla entered the automotive system as a peripheral actor. It then re-engineered the network:

- bypassed dealers
- built its own charging infrastructure
- integrated software, batteries, and vehicles
- set de facto standards (NACS adopted by Ford, GM, and others)

Tesla's power is not manufacturing. It is network reconfiguration.

As competitors adopted Tesla's charging standard, Tesla moved from outsider to structural hub – a shift created through recursive relational dynamics.

Ontological Payoff

Structuration completes relational ontology:

Economic structures are not constraints on action. They are historical products of action that now shape future action.

CAS, emergence, process, stabilisation, and networks converge here:

- Structures are real
- Agency is real
- Both exist only through recursive relational dynamics

There are no atoms. There are no static wholes. There are only patterns of interaction that reproduce themselves – or don't.

And that is exactly how economies actually work.

5.3.3 Counter-Arguments and Responses

Counter-Argument 1: "Relational ontology dissolves entities – if everything is relations, nothing exists."

Response. This objection confuses constitution with elimination. Relational ontology does not deny entity existence; it redefines what an entity is. Entities exist as stabilised relational patterns – persistent configurations maintained through ongoing interaction. Cilliers (2005, pp. 257–266) shows that complex systems exhibit recognisable, measurable, causally efficacious identities even though these identities arise entirely from relational organisation.

Prigogine's dissipative structures provide the physical analogue: a whirlpool is real, persistent, and causally efficacious, yet consists only of flowing water. It exists *through* process and relation, not despite them (Prigogine & Stengers, 1984, pp. 131–153).

Firms, markets, and institutions exist in the same ontological mode: as dynamically stabilised relational processes – exactly the synthesis that Section 5.3 established.

Counter-Argument 2: "Network analysis is just a tool – it doesn't require ontological commitment."

Response. Network methods can indeed be used instrumentally within atomistic frameworks – but doing so strips them of explanatory force. If relations are merely descriptive overlays on fundamentally independent entities, then topology has no causal status. But empirical evidence shows the opposite.

Kirman (2016, pp. 555–580) demonstrates that financial network topology determines systemic stability independently of individual bank health. Newman (2018, pp. 1–10) shows that scale-free degree distributions, clustering, and centrality patterns generate diffusion, fragility, and concentration effects regardless of node attributes.

The 2008 financial crisis confirmed this ontological point: risk resided in **network structure**, not in balance sheets alone.

If relations are constitutive, network analysis maps **causal structure**, not just correlation. Ontology and method align: relational reality requires relational analysis.

Counter-Argument 3: "Embeddedness is just transaction-cost economics with different language."

Response. This objection mistakes modification for constitution. Transaction-cost economics (Williamson, 1985) treats social relations as frictions affecting exchanges between pre-formed agents. Embeddedness theory makes a deeper claim: agents, opportunities, information, and enforcement mechanisms are constituted by networks.

Granovetter (1985, pp. 481–510) explicitly distinguishes embeddedness from transaction-cost logic: in TCE, social structure is a constraint; in embeddedness, it is the substrate of economic life.

Powell et al. (1996, pp. 116–145) show that biotech innovation depends on relational architectures that enable knowledge recombination – effects irreducible to cost minimisation. Embeddedness does not refine atomism; it replaces it.

Counter-Argument 4: "Structuration creates infinite regress – who comes first, agents or structures?"

Response. This objection applies linear causation to recursive processes. Structuration is not sequential but simultaneous. Giddens (1984, pp. 25–28) is explicit: structure exists only in and through practice, and practice presupposes structure. They are two aspects of one process, not two stages.

Language provides the simplest example: people speak language and language exists only because people speak it. There is no regress because speaking and language are the same process viewed at different levels.

Likewise, market transactions and market structures are co-present aspects of ongoing economic reproduction. As Lawson (2019, pp. 9–15) argues, social reality is inherently relational and processual – linear causal sequencing is the wrong metaphysics.

VII. Conclusion

This section establishes relational ontology as the final pillar of the dissertation's ontological reconstruction, specifying what the stabilised processes of Section 5.3 reproduce: relational patterns.

The argument proceeded through five steps:

1. **Atomism versus relationality** as a foundational ontological divide
2. **Connective complexity** as the causal mechanism of emergence
3. **Structural realism** via network topology
4. **Embeddedness** as the relational constitution of economic action
5. **Structuration** as recursive co-constitution of agency and structure

Together they demonstrate that economic entities are not self-contained substances but nodes in evolving relational networks whose properties are positional, historical, and emergent.

Where Section 4.1.3 showed that atomism fails, this section shows why and what replaces it.

Transformation Map

Table 118: Transformation Map from Atomism to Relational Ontology: Reconceiving Nine Foundational Dimensions of Economic Analysis

Dimension	FROM (Atomism)	TO (Relational Ontology)
Ontological units	Independent elements	Relational patterns
Entity identity	Intrinsic essence	Network position
Properties	Given endowments	Positional effects
Causation	Located in agents	Distributed through topology
Context	External	Constitutive
Structure	Background	Primary causal layer
Method	Decompose	Map networks
Meso-level	Residual	Ontologically primary
Agency	Individual	Relationally enabled

The table above synthesises the ontological shift developed throughout this section. At its foundation lies the reconceptualisation of what economic entities fundamentally are: not independent elements with intrinsic properties but relational patterns whose identities are constituted by network position. From this starting point, the remaining transformations follow with logical necessity.

Consider the implications for causation and method. If entity properties are positional rather than intrinsic – arising from location within relational structures rather than from internal endowments – then causation cannot be located in agents but must be understood as distributed through topology. This is what Kirman (2016) demonstrated for financial networks and what Saxenian (1994) showed for regional innovation systems: outcomes are generated by relational configuration, not by the attributes of individual nodes. The methodological consequence is equally fundamental: if the phenomena of interest are relational, then decomposition destroys them. Analysis must map networks rather than dissect components.

The transformation of context from external to constitutive captures the core ontological claim. Under atomism, relations are background – constraints or opportunities facing pre-formed agents. Under relational ontology, relations are substrate: they constitute what agents are, what they can know, and what they can do. Structure accordingly shifts from background to primary causal layer. And agency, rather than being eliminated, is reconceptualised as relationally enabled: the capacity to act arises from network position and is exercised by reshaping the relational field.

Finally, the meso-level – the intermediate scale between micro and macro – transforms from residual category to ontologically primary site. If entities are constituted through relations, and relations stabilise at organisational and institutional scales, then the meso-level is where economic reality takes form. This directly supports the dissertation's methodological commitment to meso-level analysis.

What Has Been Achieved

Relational ontology completes the CAS triad:

- **Emergence** (Section 5.1) → wholes have irreducible causal powers
- **Process** (Section 5.2) → reality is historical becoming
- **Stabilised Processes** (Section 5.3) → persistence is achieved through continuous reproduction
- **Relations** (this section) → what is reproduced are relational patterns

Together they yield the full CAS ontology: economic entities are emergent, evolving, relationally constituted processes.

Forward Integration

The next section (5.5) establishes where relational stabilisation primarily occurs: at the meso-level – the intermediate analytical scale between micro-agents and macro-aggregates that mainstream economics systematically ignores yet where economic coordination actually happens. The firm, as this dissertation argues, is irreducibly a meso-level entity.

5.4 Relational Ontology: The Constitutive Primacy of Relations

Introduction

This section develops the third pillar of the dissertation's ontological reconstruction: relational ontology – the position that relations, not entities, are ontologically primary. Entities do not precede their relations; they are constituted by them. As Cilliers (2005, p. 260) states, in complex systems "the relationships between the elements are more important than the elements themselves... The identity of the elements in the system is constituted by the relationships between them."

The section receives from Section 5.3 the understanding that entities persist through continuous reproduction of organisational patterns. Section 5.3 established that stability is achieved, not assumed – that what we call "substances" are really stabilised processes. This section addresses the consequent question: what are these processes reproducing? The answer is relational patterns – network positions, connections, and configurations that constitute what entities are.

The section provides to Section 5.5 the relational architecture that explains where stabilisation occurs. Once entities are understood as relationally constituted processes, the question becomes at what analytical scale these relational patterns form and operate. Section 5.5 will answer: primarily at the meso-level, the intermediate scale between individual agents and aggregate outcomes that mainstream economics systematically ignores.

Across industries, network-position effects – centrality, brokerage, and structural holes – systematically dominate firm-level attributes in explaining access to information, innovation diffusion, and competitive advantage. Relational ontology therefore completes the CAS triad. Where emergence (Section 5.1) established structural holism (wholes > parts) and process ontology (Section 5.2) established temporal dynamism (becoming > being), and stabilised processes (Section 5.3) showed how becoming generates durable patterns, relational ontology establishes constitutive embeddedness: what an entity *is* depends on where it sits in relational space.

The section serves three functions. First, it provides the positive reconstruction to the atomism critique (Section 4.1.3), showing not merely that atomism fails but what replaces it. Second, it grounds network analysis ontologically: networks are not optional analytical tools but the appropriate representation of economic reality. Third, it specifies the content of what Section 5.3's stabilised processes reproduce: relational configurations.

A predictable objection must be addressed at the outset. If entities are constituted by relations, do they dissolve into pure networks? No. As Section 5.3 established, relations stabilise into persistent patterns. What we call "entities" are temporarily stabilised configurations within relational fields – meta-stable patterns that persist through continuous reproduction of their relational constitution.

5.4.1 The Problem of Persistence

I. The Dialectical Imperative: Why Synthesis Is Necessary

The dissertation has reached a genuine ontological tension.

Section 4.1.2 demonstrated that substance ontology fails: firms, agents, and markets do not possess intrinsic, context-free essences. Their properties are relational, historical, and contingent.

Section 5.2 established that process ontology is correct: economic reality is not a collection of things but a web of ongoing activities – learning, innovating, coordinating, trading.

But this creates a paradox: If everything is process, what persists?

This is not a metaphysical luxury – it is an analytical necessity. Firms endure through personnel turnover. Markets persist through millions of daily transactions. Institutions survive centuries of reform.

Evolution requires identity through change.

Three false solutions must be rejected:

1. False retention (substance fundamentalism). Preserving static essences contradicts observed economic evolution. Nokia, Kodak, IBM, and Netflix are not the same "things" over time.

2. False elimination (pure flux). If nothing persists, nothing can evolve. Selection, learning, and path-dependence become meaningless.

3. False dichotomy (Being vs Becoming). Treating persistence and change as opposites misses the deeper truth.

The correct resolution is dialectical:

Persistence is produced by process. Stability is generated by change. Being exists through Becoming.

This is the core CAS insight: stability is not an assumption – it is an emergent outcome.

Mainstream economics presupposes persistence (firms exist, markets exist, institutions exist). CAS explains how persistence is continuously reproduced.

This inversion is foundational. It replaces *static ontology* with *dynamic constitution*.

5.4.2 Mechanisms of Stabilisation

II. Dissipative Structures: Order *From* Process

Ilya Prigogine's Nobel Prize-winning theory of dissipative structures provides the scientific foundation for the synthesis of substance and process. Far from thermodynamic equilibrium, systems can generate and maintain stable structures that exist only through continuous energy and matter throughput. These are not substances in the classical sense but dynamically maintained patterns – structures whose persistence depends on the very processes that produce them. In such systems, the structure *is* the process stabilised.

Prigogine and Stengers (1984, p. 291) contrast the classical "science of being" with a science of "becoming." Classical physics described a world of reversible time and timeless laws – a world that "escapes the arrow of time." Complexity science, by contrast, reveals systems governed by irreversibility, historical contingency, and creative evolution. Stable forms arise not despite flux but because of it. This is Prigogine's famous formulation of "*the being of becoming*": persistence emerging from continuous transformation.

A candle flame is the canonical example. A flame has recognisable form, stable temperature gradients, and measurable physical properties. Yet it is nothing but combustion in motion. No "flame substance" exists behind the pattern. Extinguish the chemical process and the flame disappears instantly. The apparent thing is the process organised into a temporarily stable form. The flame does not survive change – it is change, stabilised.

Economic structures operate in the same way. Markets, firms, technologies, and institutions are not static entities but dissipative systems sustained by continuous activity. Consider financial markets. The New York Stock Exchange has operated since 1792, maintaining identity across wars, technological revolutions, and regulatory change. Yet it exists only through ongoing trading – billions of transactions, price updates, entries and exits. Suspend trading and "the market" becomes an empty institutional shell. The market-as-entity is nothing other than trading-as-process stabilised into a recognisable structure.

This leads to a fundamental inversion of equilibrium economics. In thermodynamic and biological systems, equilibrium is death. Living systems exist far from equilibrium; they must continuously process energy, information, and materials to remain organised. Economic systems are no different. Static equilibrium – the world where nothing changes – is not stability but cessation.

Arthur (2021, p. 137) makes this explicit: equilibrium theory "makes no allowance for the creation of new products or new arrangements... for events triggering novel events, indeed, for history itself." It describes the very type of world Prigogine rejected – a world without becoming. Complexity economics instead studies economies as structures in motion, where organisation is continuously created and recreated (Arthur, 2021, p. 143).

Firm-Level Translation: Lehman Brothers

The collapse of Lehman Brothers in 2008 provides a dramatic demonstration. For 158 years, Lehman appeared as solid an institution as any in global finance. Yet within weeks it dissolved completely. The firm did not lose its "essence"; it lost the processes that sustained it – market confidence, counterparty trust, and funding flows. When those processes ceased, the apparent substance vanished. Lehman was not a thing that failed; it was a pattern of financial relations that stopped reproducing itself.

The conclusion is unavoidable: economic structures persist only insofar as the processes that sustain them continue.

III. Autopoiesis: Self-Producing Systems

The concept of autopoiesis – from *auto* (self) and *poiesis* (production) – formalises how systems maintain identity through continuous self-production. Developed by Maturana and Varela for living systems, autopoiesis describes organisations that continuously generate the components that constitute them. They do not merely use resources; they reproduce their own organisation.

Arthur (2014, p. 20) adopts this insight for economics: the economy "forms from its technologies and mediates the creation of further technologies and thereby its own further formation." Organisation persists while components change. Identity is not located in materials but in the pattern of production that regenerates itself.

This resolves the ancient Ship of Theseus paradox. If every plank is replaced, is it the same ship? Substance ontology cannot answer. Autopoietic ontology can: the ship is the organised pattern of relations that persists through material turnover. As long as that organisational pattern is reproduced, identity remains.

Firm-Level Translation: Toyota

Firms exhibit precisely this logic. Toyota in 2024 shares almost no physical or human components with Toyota in 1950. Yet it is recognisably the same firm because its organisational pattern – the Toyota Production System – has continuously reproduced itself. Kaizen, just-in-time, supplier integration, and problem-solving routines are not stored; they are enacted. The firm exists only by continually regenerating the practices that constitute it. Toyota-as-entity is Toyota-as-process.

Institutional Translation: The Federal Reserve

The same applies to institutions. The Federal Reserve has existed since 1913 despite radical changes in policy regimes, leadership, instruments, and economic conditions. Its identity does not reside in specific rules, people, or buildings but in the ongoing process of monetary governance that reproduces itself through reinterpretation and adaptation. An institution that ceased to evolve would cease to function.

The implication is decisive: economic entities persist only by reproducing themselves through continuous activity.

Firm identity, market identity, and institutional identity are not stored in components but enacted through recursive self-production. Stability is not the absence of change – it is change organised into continuity.

IV. Dynamic Equilibrium: Stability Through Change

Chemistry provides a precise concept for reconciling flux with stability: dynamic equilibrium. In dynamic equilibrium, reactions continue while macroscopic concentrations remain stable. Forward and reverse processes occur at equal rates – the system is continuously transforming at the micro-level while maintaining persistent structure at the macro-level. This is stability through change, not stability despite change.

Process philosophy anticipated this logic. Whitehead (1929/1978, p. 22) described reality as composed of "actual occasions" – events that synthesise past into present. What appear as enduring "objects" are rhythmic patterns of these occasions. Persistence is not material permanence but recurrent organisation – each moment is genuinely new yet recognisably continuous with what preceded it.

Dynamic equilibrium therefore provides the exact ontological model for economic persistence.

Organisational routines. Routines are not stored scripts that get executed unchanged. As Nelson and Winter (1982, p. 4) emphasise, firms operate through "capabilities and decision rules" that evolve via learning, experimentation, and error. Every enactment of a routine is slightly different – people adapt, machines drift, contexts vary – yet the pattern persists. The routine-as-entity *is* the practice-as-process stabilised in dynamic equilibrium.

Industry structures. Industries also exist in dynamic equilibrium. The pharmaceutical sector, for example, exhibits stable macro-properties – R&D intensity around 15%, patent lives of 20 years, regulatory pipelines of 10–15 years – yet its micro-components are in constant flux: firms enter and exit, drugs are approved and expire, patents issue and lapse, scientists move between organisations. Stability emerges through balanced turnover – innovation replaces obsolescence. Structure persists *through* transformation.

Firm-Level Translation: Amazon

Amazon exemplifies dynamic equilibrium at scale. Its identity includes recognisable structures – marketplace, AWS, logistics, Prime ecosystem – yet every component is in constant motion: millions of products appear and disappear; over two million sellers churn annually; algorithms update continuously; warehouses are built, reconfigured, and automated; the workforce has turned over several times as it grew from one person in 1994 to 1.5 million in 2024. Amazon-as-entity is not any snapshot configuration – it is the ongoing coordination process that maintains coherence across perpetual change.

This distinguishes CAS ontology from mechanistic equilibrium economics. Neoclassical theory seeks static equilibria – states in which no agent has incentive to change. But such states correspond, thermodynamically, to death. Arthur (2021, p. 137) notes that equilibrium theory "makes no allowance for... events triggering novel events, indeed, for history itself." Living economic systems persist only in dynamic equilibrium.

Firm-Level Translation: Microsoft

Microsoft has remained recognisably Microsoft across radical technological discontinuities – command-line DOS, Windows GUI, internet, cloud computing, subscription software. Its dominant revenue stream has shifted from Windows licensing to Azure cloud services; its leadership, workforce, and products have changed repeatedly. Yet the firm persisted. Microsoft-as-substance is the adaptive organisational process that maintained coherence while everything else changed.

Dynamic equilibrium reveals the core CAS insight: stability is achieved, not assumed; persistence is produced, not given.

V. Path-Dependent Identity: History as Trajectory

Path-dependence completes the synthesis. If systems persist through dynamic equilibrium, then their identity must be historical. As Arthur (1994) showed, increasing returns and positive feedback lock early events into durable trajectories. Small, contingent beginnings crystallise into large-scale structures that cannot be undone without destroying the system itself.

This transforms identity.

Substance ontology: identity is essence – what something is in itself, independent of history.
Process-relational ontology: identity is trajectory – what something has become through time.

Firms, technologies, regions, and institutions are not bundles of attributes; they are accumulated paths.

Apple's identity cannot be defined without its history: garage startup → Macintosh → near-bankruptcy → iPod → iPhone → services ecosystem. Each stage created capabilities that made the next possible. Identity is the sedimentation of past decisions into present structure.

Amazon versus Walmart. Both are retailers of similar scale, yet they are different kinds of organisations because their trajectories differ. Walmart's path – stores, logistics, physical scale – generated capabilities in real estate, supplier bargaining, and in-store efficiency. Amazon's path – online retail, algorithms, logistics platforms, cloud computing – generated capabilities in digital coordination, marketplaces, and infrastructure services. These identities cannot be swapped. You cannot turn Walmart into Amazon by "adding technology," because identity is not a toolkit – it is a history encoded in organisational structure.

Nokia versus Apple. Nokia's telecom trajectory created phones optimised for voice reliability and hardware engineering. Apple's computing trajectory created phones as handheld computers

– interfaces, apps, ecosystems. Nokia did not fail because it lacked intelligence; it failed because its accumulated identity constrained what it could become. Trajectory, not foresight, determined fate.

The synthesis is decisive: Firms are not essences with properties; they are histories with momentum.

Path-dependent identity makes replication impossible. You cannot build "another Silicon Valley" by copying current conditions, because Silicon Valley is not a template – it is a trajectory. Identity is irreversibly historical.

5.4.3 The Hegelian Resolution and Synthesis

VI. The Hegelian Resolution: *Aufhebung*

The ontological reconstruction reaches completion through Hegelian synthesis. Georg Wilhelm Friedrich Hegel (1807) formalised dialectical logic: when two apparently opposed ontological claims generate irresolvable contradictions, neither is discarded; instead, both are sublated (*aufgehoben*) into a higher-order unity that simultaneously cancels, preserves, and transforms them.

The dissertation's ontological trajectory is precisely such a dialectic.

THESIS – Being (Parmenides; Section 4.1.2) Reality consists of enduring substances with intrinsic properties. Firms, markets, and institutions possess identities that persist independently of change.

ANTITHESIS – Becoming (Heraclitus; Section 5.2) Reality consists of continuous flux. No enduring entities exist – only ongoing processes. Identity dissolves into flow.

SYNTHESIS – Being-Through-Becoming (This section) Substances are stabilised processes. Persistence is not opposed to change – it is generated by it. What we call "entities" are meta-stable patterns maintained through continuous processual reproduction.

This is not compromise but *Aufhebung*:

- *Being is cancelled* – no entity is immutable.
- *Being is preserved* – persistent structures are real and causally efficacious.
- *Being is transcended* – persistence is reconceptualised as dynamic reproduction.

→ *Becoming is cancelled* – reality is not chaotic flux. → *Becoming is preserved* – process is ontologically primary. → *Becoming is transcended* – process generates durable structure.

This yields the correct ontology for complex adaptive systems: meta-stability.

Meta-Stability as Economic Ontology

Meta-stability denotes stability that is:

Provisional – It lasts only while sustaining processes continue. Nokia, Lehman Brothers, Kodak all persisted for decades before dissolving when the conditions reproducing them failed.

Temporal – Entities exist for finite durations. S&P 500 firm lifespans collapsed from 61 years (1958) to 18 years (projected 2027). Their reality was not illusory – only temporary.

Revisable – Identity persists through transformation. Netflix remained Netflix while shifting from DVDs → streaming → content production. Identity was maintained *through* structural change.

Real – Meta-stable entities exert causal power. The Federal Reserve, Apple, Amazon, and Silicon Valley are not abstractions – they shape capital flows, innovation paths, and life chances.

Path-dependent – What exists today constrains what can exist tomorrow. Amazon's platform position reshaped retail; Silicon Valley's emergence reshaped global innovation geography.

Thus, firms are neither eternal essences nor arbitrary labels. They are historically stabilised patterns of organisation.

Two Errors Avoided

This synthesis avoids two symmetrical mistakes.

1. Essentialism. Treating firms as having fixed natures. Kodak did not fail because it "lost its essence." It failed because its reproduction regime could not adapt.

2. Nihilism. Treating firms as arbitrary collections. Apple is not just a name for shifting assets. It is a persistent, structured process that produces recognisable capabilities across decades.

The synthesis: Firms are real but not permanent. Stable but not static. Persistent but not immutable.

Strategic Implications

In a meta-stable world, strategy cannot be optimisation within fixed structures. It is the management of reproduction conditions.

Hynes et al. (2021, p. 72) are precise: "When you try to optimise one part of a complex system, you can end up destabilising the system as a whole."

Because structure itself must be continuously reproduced, firms compete not only for market position but for the conditions that allow their identity to persist.

Firm-Level Translation: Kodak

Kodak's 131-year dominance (1880–2012) was a triumph of meta-stability. Film chemistry, distribution networks, consumer habits, and brand reputation were continuously reproduced into a stable pattern.

Digital photography destroyed the reproduction regime, not the firm's intelligence. When film sales collapsed, Kodak-as-process could no longer sustain Kodak-as-substance. The firm did not "fail"; its meta-stability ended.

This is the correct ontology of economic reality.

VII. Counter-Arguments and Responses

Counter-Argument 1: "This is just semantics – calling firms 'stabilised processes' changes nothing practical."

Response. The difference is operational, not verbal. If firms are treated as substances, strategy reduces to optimising given assets inside fixed structures (VRIN logic, equilibrium positioning, static efficiency). If firms are stabilised processes, strategy becomes the management of trajectories, reproduction regimes, and adaptive capacity.

This distinction is not rhetorical – it already structures the most successful modern theories of the firm. Teece, Pisano & Shuen (1997) define dynamic capabilities as the firm's capacity to "integrate, build, and reconfigure internal and external competences." These are processes, not stocks. Nelson & Winter's (1982) routines are recurrent patterns of activity, not possessions.

Nokia failed not because it lacked resources, patents, or engineers – but because its process of self-reproduction (telecom hardware + Symbian + operator relationships) could not generate the smartphone trajectory. Ontology here directly alters what managers must monitor: not what they own, but how they reproduce themselves through time.

Counter-Argument 2: "Dissipative structures apply to physics, not economies."

Response. The claim is not physical reductionism but structural homology. Both biological and economic systems share three defining properties:

1. They persist far from equilibrium
2. They maintain pattern while components turn over

3. They require continuous throughput (energy, information, capital, labour) to exist

Prigogine explicitly extended dissipative-structure logic to social systems (Prigogine & Stengers, 1984). Arthur (2014) shows the economy "forms from its technologies and mediates the creation of further technologies" – an autopoietic loop.

Nelson & Winter (1982) formalised this for firms: routines reproduce organisational identity while being continually modified. Economic entities are not matter-energy systems, but they are organisation-flow systems – which is what dissipative structures describe.

Counter-Argument 3: "If everything is process, prediction becomes impossible."

Response. Process ontology does not deny predictability; it redefines its domain.

In stabilised processes, path-dependence generates increasing predictability once trajectories lock in. Arthur (1994) shows that positive feedback produces dominant designs, standards, and market leaders that become highly predictable – until disrupted.

What process ontology rejects is timeless, context-free forecasting. Bronk (2011) is precise: complex systems permit navigation and preparation, not Laplacian prediction. This is not a loss of science – it is a gain in realism.

Counter-Argument 4: "The synthesis is incoherent compromise."

Response. This is a misunderstanding of *Aufhebung*.

The synthesis is not halfway between Being and Becoming – it is their resolution:

- Substances are real
- Processes are primary
- Stability is produced, not assumed

Whitehead's (1929/1978) "actual occasions" formalise exactly this: process generates temporary completion → completion generates pattern → pattern persists through continued process.

This is not philosophical ambiguity – it is the correct ontology of CAS.

VIII. Conclusion

This section has established substances as stabilised processes, resolving the being/becoming tension raised by process ontology.

The argument advanced through six logically necessary steps:

1. Demonstrating why substance or process alone is insufficient
2. Grounding synthesis in dissipative-structure science
3. Formalising persistence through autopoiesis
4. Establishing dynamic equilibrium as stability-through-change
5. Showing identity as path-dependent trajectory
6. Resolving Being and Becoming through *Aufhebung*

Transformation Table

Table 119: From Substance to Stabilised Process: Reconceiving Nine Dimensions of Firm Ontology

Dimension	FROM (Substance Ontology)	TO (Stabilised Process)
Nature of entities	Static bearers of properties	Dynamic patterns of reproduction
Identity	Timeless essence	Historical trajectory
Persistence	Assumed	Achieved
Stability	Equilibrium	Meta-stability
Boundaries	Fixed	Relationally produced
Components	Define entity	Continuously replaced
Change	Threat	Condition of survival
Time	Reversible	Irreversible
Analysis	Decomposition	Reproduction dynamics

The table above synthesises the ontological shift that this section establishes. At its foundation lies the reconceptualisation of what persistence means: not the continuation of unchanging substance but the continuous reproduction of organisational pattern. From this starting point, each transformation follows with logical necessity.

Consider the implications for identity and analysis. If entities are dynamic patterns rather than static bearers of properties, then identity cannot be timeless essence – it must be historical trajectory, the accumulated path through which the pattern has evolved. This is why Apple cannot be understood apart from its history, why Nokia's failure was trajectory-determined, and why Silicon Valley cannot be replicated by copying current conditions. The methodological consequence is equally fundamental: if phenomena persist through reproduction rather than inertia, then analysis must track reproduction dynamics rather than decompose into components. Decomposition destroys the very processes that maintain the entity.

The transformation of stability from equilibrium to meta-stability captures the core insight. Neoclassical equilibrium describes states where nothing changes – thermodynamically equivalent to death. Meta-stability describes patterns that persist precisely because they change: components turn over, routines adapt, structures evolve. This is stability achieved through continuous activity, not stability assumed as a starting condition. The practical implication is decisive: strategy must manage reproduction conditions, not merely optimise within fixed structures.

The Core Result

Emergence showed that wholes have irreducible properties. Process showed that reality is becoming. This section shows how becoming generates durable patterns: substances are stabilised processes.

But one crucial dimension remains unexplored: the processes through which entities reproduce themselves are not self-contained. They are constituted by connections to other entities – suppliers, customers, competitors, institutions, networks. The next section (5.4) adds this relational dimension: stabilised processes are relationally constituted. What an entity is depends on where it sits in relational space.

Implications

Firms do not "possess" competitive advantage. They reproduce it. Markets do not "exist." They are continually enacted. Institutions do not "stand." They persist only while their reproduction mechanisms function.

Competition is therefore not between static entities but between trajectories attempting to stabilise themselves in hostile environments.

Policy must target processes, not structures. Strategy must manage reproduction, not assets. Failure is not loss of essence – but breakdown of self-production.

Forward Integration

The next section (5.4) adds the relational dimension that completes the ontological reconstruction. Stabilised processes do not exist in isolation – they are constituted by network connections that shape what they are and what they can become.

5.5 Emergence as Meso-Level Phenomenon: The Missing Middle

Introduction

This section establishes meso-level emergence as the dissertation's distinctive theoretical contribution. It shows that the core properties of complex adaptive systems – emergence, process, stabilised patterns, and relationality – manifest primarily at intermediate analytical scales that are systematically eliminated by micro–macro reductionism.

The section receives from Section 5.4 the relational architecture through which entities are constituted. Section 5.4 established that economic actors exist within networks that determine their capabilities, opportunities, and identities. This section addresses the consequent question: at what analytical scale do these relational patterns form and operate? The answer is the meso-level – the intermediate scale between individual agents and aggregate outcomes where routines crystallise, networks stabilise, and institutions emerge.

The section provides to Chapter 6 the three-level ontology (micro–meso–macro) that organises the remainder of the dissertation. Once meso-level emergence is established as the primary site of economic complexity, the methodological implications follow: analysis must preserve meso-structures rather than decompose them.

We define meso-level emergence as the appearance of structures, rules, and patterns at intermediate scales – organisational routines, industry architectures, network effects, innovation ecosystems, institutional regimes – that are:

- not reducible to micro-level agents
- not derivable from macro-level aggregates
- yet causally decisive for both

As Dopfer and Potts (2004, p. 263) formulate it, the meso-level is "a population of rules, a structure of rules, and a process of rules" through which economic evolution actually operates.

This section serves three architectural purposes:

1. It makes explicit what Sections 5.1–5.4 established implicitly: emergence, process, stabilisation, and relationality necessarily generate meso-structures. Emergent wholes stabilise as routines; processes crystallise into organisational patterns; relations solidify into networks and institutions.
2. It answers Arthur's critique that mainstream economics lacks any formal way to analyse intermediate structures – "the level that standard economics largely misses."
3. It establishes the three-level ontology (micro–meso–macro) that organises the remainder of the dissertation and shows why the firm is inherently a meso-level entity.

Without explicit meso-level theory, complexity collapses back into reductionism. With it, CAS becomes analytically operational.

5.5.1 The Missing Middle in Economic Theory

I. The Missing Middle

Mainstream economics recognises only two analytical levels:

- **Micro** – individual agents (and firms treated as agents)
- **Macro** – aggregate outcomes (GDP, prices, unemployment)

The representative-agent framework aggregates directly from micro to macro: individual optimisation → market equilibrium.

This procedure systematically deletes the level where economic life actually occurs:

- organisational routines
- production networks
- industry architectures
- innovation ecosystems
- institutional regimes

These are meso-level structures – neither individuals nor aggregates, but organised interaction systems.

Arthur (2014, pp. 3–25) identifies this blind spot precisely. Standard economics moves from agents to aggregates while missing "the structures and patterns that form from interactions." This is not an empirical omission but an ontological consequence: if wholes equal sums of parts, then no intermediate structures can possess independent causal power.

But emergence (Section 5.1) falsifies that premise. If wholes possess properties irreducible to parts, then meso-structures must be real. They are where irreducibility lives.

Why Meso-Level Properties Are Ontologically Distinct

Meso-level structures have properties that exist nowhere else:

Organisational capabilities. Toyota's production capability does not reside in any worker (micro) nor in "the auto sector" (macro). It exists only in the organisational system of routines, coordination, and learning. Disassemble the firm and the capability vanishes – even if all workers remain.

Industry architectures. The pharmaceutical sector's R&D intensity, regulatory time horizons, patent dependence, and blockbuster model exist at the industry level. They constrain every firm but belong to none.

Network effects. Platform value depends on network size and topology – a meso-level property. No individual preference creates it. No aggregate statistic captures it.

These are not aggregations. They are emergent structures.

Firm-Level Translation

Strategic analysis already lives at the meso-level – even when economics pretends otherwise.

When analysts study the pharmaceutical industry, they examine: R&D norms (~15% of revenue); FDA approval pathways (10–15 years); blockbuster dependence; patent cliffs; alliance networks (Malerba, 2002).

These are meso-level regularities. They are neither firm traits nor macro aggregates. The firm itself is a meso-entity: above individuals, below the economy, constituted by routines, relations, and reproduction processes.

Mainstream economics has no formal theory for this level – yet all real economic coordination happens here.

Framework Consequence

By deleting the meso-level, mainstream economics renders complexity invisible.

All five CAS properties – emergence, feedback, heterogeneity, path-dependence, and self-organisation – operate primarily at the meso-scale.

That is where firms form. That is where industries evolve. That is where innovation happens. That is where institutions stabilise. The "missing middle" is not peripheral. It is where the economy actually lives.

By aggregating micro to macro, mainstream erases precisely where complexity operates.

II. Organisational Routines as the Paradigmatic Meso-Phenomenon

Nelson and Winter's (1982) theory of organisational routines provides the canonical empirical instantiation of meso-level emergence. Routines are not individual habits and not aggregate market behaviour; they are organisationally embedded patterns of coordinated action that exist only at the level of the firm as a system.

They are therefore ontologically meso-level entities: above individuals, below industries, irreducible to either.

Nelson and Winter (1982, p. 14) define routines as "regular and predictable patterns of behaviour" that characterise organisations. Crucially, these patterns do not reside in any individual mind or decision rule. They are distributed, enacted, and reproduced through organisational interaction structures.

Three properties establish routines as paradigmatic meso-phenomena.

1. Supra-individual persistence

Routines persist through personnel turnover. Employees leave; routines remain. Employees arrive; routines shape how they behave.

Organisational memory is therefore not stored in individuals but in reproducible interaction patterns.

Downward causation operates here: the routine (meso) constrains the agent (micro).

Empirical work in organisational sociology and management shows that meso-level rules – norms, routines, and institutional constraints – account for roughly 50–70% of observed behavioural variance unexplained by individual preferences or incentives alone.

Toyota provides the empirical proof. As Dyer and Nobeoka (2000) show, the Toyota Production System (TPS) has persisted since the 1960s through multiple complete workforce turnovers. No engineer from the original TPS remains, yet kaizen, andon, just-in-time, and supplier coordination continue to operate with remarkable fidelity. A concrete example is line-stop authority: when an andon signal is pulled, the production line can be halted immediately, forcing rapid problem-solving coordination rather than downstream inspection.

Toyota (founded 1937) has outlived more than five generations of employees. The routine, not the person, is the unit of persistence.

2. Sub-economic specificity

Routines are not economy-wide. They are firm-specific meso-structures.

Toyota's routines are not Ford's routines. Ford's routines are not Tesla's. Yet all are "automobile manufacturers."

This heterogeneity is invisible to macro models that treat firms as identical production functions.

Toyota produces 10.5 million vehicles annually using routines built around:

- just-in-time logistics
- line-stop authority
- continuous improvement
- relational supplier integration

Ford's mass-production legacy generated entirely different routines – volume-driven, inventory-buffered, vertically oriented.

The difference is not inputs or prices. It is meso-level organisational architecture.

3. Emergence from coordination

Routines are not designed by anyone.

They emerge from repeated coordination attempts, stabilising into patterns through learning, error, and adaptation – exactly the dynamic equilibrium described in Section 5.3.

Yet once stabilised, routines exert downward causation: they structure who does what, how problems are solved, and what responses are available.

Nelson and Winter (1982, p. 4) describe firms as operating with "certain capabilities and decision rules" that evolve through variation, selection, and retention.

That is evolutionary dynamics at the meso-level.

Why Micro-Insight Does Not Beat Meso-Inertia

Kodak proves the point. Kodak engineers invented the digital camera in 1975. By the 1990s, Kodak's scientists understood digital's inevitability.

But Kodak's routine architecture was built around:

- chemical film R&D
- mass-manufacturing of consumables
- global film distribution
- retail processing networks

Tripsas and Gavetti (2000) show that these meso-structures overwhelmed individual knowledge. The organisation could not reconfigure because the routines that made Kodak *Kodak* were incompatible with digital photography.

Micro-level cognition changed. Meso-level routines did not.

Kodak's 2012 bankruptcy destroyed \$30bn in market value not because individuals were stupid, but because the meso-level system could not transform.

Framework Consequences

If routines are meso-level entities with causal power, then:

Table 120: Framework Consequences of Recognising Routines as Causally Powerful Meso-Level Entities: From Individualist to Routine-Based Explanations

What mainstream assumes	What CAS reveals
Capabilities are individual skills	Capabilities are routine bundles
Strategy is resource allocation	Strategy is routine reconfiguration
Learning is training	Learning is routine evolution
Failure is bad decisions	Failure is meso-level lock-in

Strategic change is hard not because people resist change, but because routines have inertia independent of people.

GE is the textbook case. Welch-era routines (financialisation, quarterly targets, conglomerate logic) persisted through Immelt's digital pivot attempts. GE collapsed from \$600bn to \$80bn market cap because the meso-architecture could not be rewritten. Breaking GE into three companies was not a strategy choice – it was meso-level reset.

III. Network Effects and Ecosystem Dynamics

Network effects are pure meso-level causation.

Individual utility (micro) depends on network size and architecture (meso), which generates market dominance (macro). Neither micro nor macro models can capture this causal chain.

Arthur (2014) shows how increasing returns lock systems into dominant structures. Bernanke (2010) showed how financial networks transmit collapse. These are not individual or aggregate phenomena – they are network topologies acting causally.

iOS ecosystem:

Table 121: Multi-Level Ontological Structure of the iOS Ecosystem: Illustrating Micro, Meso, and Macro Layers of a Platform as CAS

Level	What exists
Micro	Developers, users
Meso	App Store rules, APIs, ranking algorithms, payment system, developer network
Macro	Revenue, market share

The ecosystem is causally primary.

Developers succeed or fail not because of talent alone but because of:

- App Store rules
- algorithmic visibility
- commission structures
- API access

Those are meso-structures. Same with AWS. Same with Android.

Markets are not battles of firms. They are competitions between ecosystem architectures.

IV. Industry Structure as Meso-Level Emergence

Industries are not aggregates of firms. They are stabilised interaction systems.

Pharma R&D intensity (15%), patent dependence, FDA timelines, and blockbuster economics are industry-level emergent constraints.

No firm chooses them. Every firm must obey them.

Moderna did not choose to need FDA approval. It navigated a meso-regime.

V. Institutional Regimes as Meso-Level Scaffolding

Institutions – norms, rules, conventions, governance systems – operate across multiple scales, but many of their most powerful effects occur at the meso-level: sectoral regulations, regional cultures, organisational practices, and financial architectures. These institutional regimes act as scaffolding for economic activity, shaping firm strategies (meso → micro), generating economic performance (meso → macro), and being recursively reproduced through the activities they enable (as established by the structuration logic of Section 5.4).

The Varieties of Capitalism framework (Hall & Soskice, 2001) provides one of the clearest demonstrations of meso-level institutional regimes. Liberal Market Economies (LMEs) such as the US and UK and Coordinated Market Economies (CMEs) such as Germany and Japan are not merely macro-labels but distinct meso-institutional configurations that shape how firms compete, innovate, and organise:

Table 122: Institutional Complementarities Across Liberal and Coordinated Market Economies: Context-Dependence Across Four Dimensions

Dimension	LME (US, UK)	CME (Germany, Japan)
Corporate governance	Shareholder primacy, quarterly earnings	Stakeholder coordination, patient capital
Labour relations	Flexible markets, weak unions	Coordinated bargaining, strong unions
Inter-firm relations	Arm's-length competition	Cooperative networks, keiretsu
Innovation systems	Radical breakthroughs	Incremental, cumulative innovation

These differences do not arise from individual firm choice nor from national aggregates; they are meso-level institutional regimes that structure what strategies are viable and what outcomes are possible.

Meso-Institutions Enable Firm Success

German Mittelstand firms illustrate meso-level scaffolding with exceptional clarity. Their global dominance in narrow technological niches depends on CME institutions:

- Dual-track apprenticeship systems producing highly skilled labour (1.3 million apprentices annually)
- Patient capital from regional banks (Sparkassen, Landesbanken)
- Dense supplier networks enabling long-term specialisation

The typical Mittelstand firm – 250–500 employees, often operating for more than 60 years, dominating global niche markets – cannot be transplanted into an LME context without losing

its institutional foundations. These firms succeed not merely because of managerial excellence but because they are embedded in a meso-institutional regime that supports long-term capability accumulation.

Silicon Valley as a Meso-Institutional Ecosystem

Silicon Valley is not a firm, and not a national economy – it is a regional meso-level institutional configuration (Saxenian, 1994). Its performance derives from a specific bundle of meso-institutions:

- Venture capital: \$74bn invested in 2023 (40% of US VC)
- Labour mobility: average tenure ~2.5 years; non-compete bans enable knowledge circulation
- Failure tolerance: serial entrepreneurship normalised
- University pipelines: Stanford & Berkeley generate spin-offs and talent
- Dense weak ties enabling rapid diffusion of tacit knowledge

These are not individual traits or macroeconomic conditions; they are meso-level institutions that enable micro-level innovation and produce macro-level growth. Attempts to "build Silicon Valley" elsewhere fail because these institutional regimes are path-dependent and recursively produced, not blueprint-replicable.

Route 128 vs Silicon Valley: Institutional Divergence

Saxenian's comparison of Route 128 and Silicon Valley demonstrates meso-institutional causation:

Table 123: Route 128 Versus Silicon Valley: Institutional Divergence Across Five Dimensions as Evidence of Meso-Institutional Causation

Dimension	Route 128	Silicon Valley
Firm structure	Vertically integrated	Specialised & networked
Labour mobility	Low	High
Knowledge flows	Internal	Cross-firm
Culture	Corporate hierarchy	Entrepreneurial egalitarianism
Venture capital	Passive	Active & networked

Despite similar starting conditions (elite universities, defence funding, early tech firms), Route 128's meso-institutions constrained adaptation when personal computing emerged. Silicon Valley's meso-institutions enabled rapid recombination. DEC collapsed; HP, Intel, Apple thrived. These outcomes cannot be explained at the micro or macro level – they require meso-institutional analysis.

TSMC and the Taiwanese Semiconductor Regime

TSMC's dominance (60% of global foundry market; 90% of sub-10nm chips) reflects not just firm strategy but Taiwan's meso-level semiconductor ecosystem:

- Government incubation (ITRI)
- Engineering education (30,000+ graduates annually)
- Dense supplier networks
- Hsinchu Science Park cluster effects

When TSMC builds fabs in Arizona (\$40bn), the difficulty is not technology but replicating the meso-institutional ecosystem that enables its capabilities.

Framework Implication

Firm performance depends on institutional fit: alignment between firm-level processes (micro) and meso-level regimes (institutions, norms, finance, skills). Institutional entrepreneurship can transform regimes; institutional inertia constrains strategy. Neither can be understood without meso-level analysis.

Mainstream economics erased this level. CAS restores it as the true arena of economic evolution.

5.5.3 The Three-Level Ontology and Synthesis

VI. Three-Level Ontology: Micro–Meso–Macro

Dopfer and colleagues formalise what the preceding sections have demonstrated: economic reality requires a three-level ontology. The micro–macro reduction that dominates mainstream economics is ontologically incomplete.

The meso-level is not merely a "middle scale." It is a qualitatively distinct domain.

- **Micro is particular:** this agent, this choice, this action. Micro-analysis concerns individual instances.
- **Meso is generic:** this rule, this routine, this technology, this norm. These are generative structures – patterns that apply across many particular cases. Meso-analysis concerns populations of rules.
- **Macro is aggregate:** this distribution, this growth rate, this market share. Macro-analysis summarises outcomes across populations.

Rules, routines, technologies, standards, and institutions are not reducible to either individual behaviour or statistical aggregates. They constitute a distinct ontological layer.

They have three causal roles:

(1) Originate from micro-interaction. Agents create rules through repeated coordination, experimentation, imitation, and institutionalisation. Routines emerge from practice; technologies from invention; norms from convention.

(2) Shape micro-behaviour. Once stabilised, rules constrain and enable action. Routines channel effort, technologies structure what is feasible, norms prescribe and proscribe behaviour. This is genuine downward causation from meso to micro.

(3) Generate macro-patterns. When rules are enacted across populations, they produce macro-outcomes: industry performance, growth trajectories, market concentration, innovation waves. Macro is the statistical shadow of meso-level structure.

Thus, causation in complex economies runs **micro** → **meso** → **macro**, with powerful feedback loops back to micro.

Meso-Trajectories and Rule Evolution

Meso-level structures evolve. Rules have life cycles:

1. **Origination** – innovation, invention, mutation
2. **Adoption** – diffusion, imitation, selection
3. **Retention** – institutionalisation, standardisation
4. **Exhaustion** – obsolescence, displacement, transformation

These are evolutionary trajectories of generative structures, not of individuals or aggregates.

QWERTY keyboards (David, 1985), VHS, franchising, and the M-form corporation (Chandler, 1962) are all meso-level rule systems. They originated, diffused, stabilised, and – except QWERTY – eventually declined. QWERTY itself demonstrates path-dependent retention: a meso-level standard persisting long after its technical justification weakened.

This is meso-evolution.

Three-Level Analysis in Practice

Table 124: Three-Level CAS Analysis in Practice: Micro, Meso, and Macro Analytical Foci Illustrated Through the Pharmaceutical Industry

Level	Analytical Focus	Pharmaceutical Example
Micro	Individual firms, strategies, capabilities	Pfizer's mRNA platform; Moderna's R&D teams
Meso	Industry rules, routines, trajectories	15% R&D norm; FDA pathways; blockbuster model
Macro	Aggregate outcomes	\$600bn global market; growth rates

Why does Firm A outperform Firm B? Because its routines and positioning fit the meso-structure better.

Why does Industry X grow faster than Industry Y? Because its technological trajectory and institutional regime support expansion.

Neither question can be answered at micro or macro level alone. The meso-level is where economic causation actually operates.

NVIDIA as Meso-Level Dominance

NVIDIA's power illustrates three-level causation:

- **Micro:** engineering talent, CUDA architecture, leadership
- **Meso:** GPU ecosystem – 4+ million CUDA developers, software stacks (cuDNN, TensorRT), deep integration with AWS, Azure, Google Cloud, and AI research institutions
- **Macro:** \$50bn data-centre revenue (2024), 80%+ share of AI training chips, \$4+ trillion valuation

NVIDIA's dominance is not explained by micro-capabilities alone (AMD and Intel also have engineers), nor by macro-statistics. It is rooted in meso-level ecosystem lock-in – software compatibility, developer networks, and cloud integration that create structural advantage.

When Amazon builds Trainium and Inferentia, it is not primarily competing with NVIDIA's engineers. It is fighting NVIDIA's meso-level rule system.

Core Ontological Result:

- Micro agents act
- Macro aggregates record

- Meso-level structures govern

This is where CAS lives. This is where firms live. This is where industries live.

And this is the level that mainstream economics erased.

VII. Counter-Arguments and Responses

Counter-Argument 1: "The meso-level is merely an analytical convenience, not ontologically real."

Response. This objection confuses *description* with *existence*. Meso-level structures possess causal efficacy independent of the particular agents that instantiate them.

Routines persist through personnel turnover; industries constrain entrants before they act; network effects generate value no individual choice can create. These are not conveniences of description but empirical facts of persistence, constraint, and causation.

Nelson and Winter (1982, pp. 96–136) show that routines have inertia, trajectories, and selection dynamics that cannot be reduced to individual intentions. Hodgson (2004, pp. 401–426) demonstrates that social structures exhibit emergent properties with causal powers irreducible to micro-behaviour. Dopfer and Potts (2004, pp. 261–267) provide the philosophical grounding: rules, routines, and technologies are ontologically autonomous generative structures. Lawson (1997, pp. 157–189) shows that social reality necessarily includes emergent structures that constrain and enable action while being reproduced through it.

The meso-level satisfies every criterion of ontological reality: it persists, it causes, and it evolves.

Counter-Argument 2: "Microfoundations are necessary – macro must be derived from micro."

Response. This rests on a reductionist fallacy. Composition does not imply reducibility.

Water is composed of hydrogen and oxygen, yet wetness cannot be derived from the properties of either. Firms are composed of individuals, yet organisational capabilities, coordination, learning, and inertia cannot be derived from individual behaviour in isolation.

Micro-consistency is necessary – macro patterns must not contradict micro behaviour – but derivation is neither required nor possible in emergent systems. Anderson's (1972) "more is different" applies: each level has its own laws.

Demanding that meso-level structures be derived from individual optimisation is a category error – like demanding that liquidity be derived from atomic physics.

Counter-Argument 3: "Network effects can be modelled without meso-level ontology – just add them to utility functions."

Response. Adding network terms to utility functions preserves methodological individualism while missing what networks actually are.

This treats networks as external constraints on optimisation, rather than as constitutive structures that define what agents can do. In a meso-ontology, network position creates capabilities; it does not merely alter payoffs.

Platform ecosystems have architectural properties – modularity, interfaces, governance, complementarity – that shape outcomes independently of individual preferences (Jacobides et al., 2018; Gawer, 2014; Cusumano et al., 2019). These properties exist at the ecosystem level, not inside agents.

A utility function with network variables still assumes the agent is primary. Meso-ontology shows that structure is primary.

Counter-Argument 4: "If the meso-level is so important, why has economics succeeded without it?"

Response. Institutional success is not analytical adequacy.

Mainstream economics has failed precisely where meso-structures dominate: financial crises (network cascades), innovation (technology trajectories), industrial evolution (heterogeneous firms), and growth (ecosystem formation).

The 2008 crisis was not caused by bad individual banks but by financial network topology. Growth models treat technology as exogenous because they cannot model meso-level innovation systems. Trade models assume identical firms because they cannot represent organisational heterogeneity.

As Arthur (2014) explains, economics succeeded by solving the problems its assumptions made tractable – while systematically ignoring the problems that require meso-level analysis.

VIII. Conclusion

This section establishes meso-level emergence as the ontological site of economic complexity.

The argument proceeded through six movements:

1. Identifying the missing middle erased by micro–macro reduction
2. Establishing organisational routines as meso-level entities
3. Demonstrating network effects and ecosystems as meso-level causation

4. Analysing industry structures as meso-configurations
5. Showing institutions as meso-level scaffolding
6. Formalising a three-level ontology

Transformation Table

Table 125: From Two-Level Reduction to Three-Level Ontology: Reconceiving Eight Analytical Dimensions with the Meso-Level as Ontologically Autonomous

Dimension	FROM (Two-Level Reduction)	TO (Three-Level Ontology)
Levels	Micro → Macro	Micro ↔ Meso ↔ Macro
Intermediate scale	Eliminated	Ontologically autonomous
Causation	Aggregation	Recursive across levels
Rules & routines	Individual habits	Generative structures
Networks	Aggregated ties	Meso-topologies
Institutions	Background	Scaffolding systems
Dynamics	Optimisation + equilibrium	Rule trajectories
Firms	Micro-agents	Meso-configurations

The table above synthesises the ontological shift that this section establishes. At its foundation lies the recognition that economic reality cannot be captured by two-level reduction. The meso-level is not a convenient analytical intermediate but an ontologically autonomous domain where rules, routines, networks, and institutions exist with their own causal powers.

Consider the implications for understanding firms and industries. Under two-level reduction, firms are treated as micro-agents – optimising units that aggregate into market outcomes. Under three-level ontology, firms are meso-configurations – organised systems of routines, relations, and reproductive processes that exist above individuals but below the economy. This is not a matter of scale but of kind: meso-structures have properties (persistence through turnover, downward causation, evolutionary trajectories) that neither individuals nor aggregates possess.

The transformation of causation from aggregation to recursive cross-level dynamics captures the core analytical shift. In mainstream economics, macro outcomes result from aggregating micro choices. In CAS, causation flows recursively: micro-interactions generate meso-structures; meso-structures constrain and enable micro-behaviour; and the resulting patterns produce macro-outcomes that feed back on meso-level selection environments. This recursive causation is precisely what representative-agent models eliminate by construction.

Why CAS Lives at the Meso-Level

- **Emergence** appears at the meso-level: routines, networks, industries
- **Process** stabilises at the meso-level: standards, platforms, institutions

- **Relations** form meso-networks: ecosystems, clusters, supply chains

Micro produces meso. Meso structures macro. Macro feeds back on meso.

This is where complexity exists.

Toyota's capabilities; Silicon Valley's innovation; Financial system fragility; Platform dominance; Industrial evolution.

None exist at the micro level. None exist at the macro level.

They exist at the meso-level – and mainstream economics erased it.

Forward Integration

This completes Chapter 5's ontological reconstruction. The CAS framework now possesses its full philosophical foundation:

- **Emergence** (5.1): wholes have irreducible properties
- **Process** (5.2): reality is becoming
- **Stabilised Processes** (5.3): persistence is achieved through continuous reproduction
- **Relations** (5.4): what is reproduced are relational patterns
- **Meso-Level** (this section): where stabilisation occurs

Chapter 6 turns to the methodological implications: if complexity lives at the meso-level, then analytical decomposition destroys what it seeks to explain. Only synthetic methodology can preserve meso-structures. Ontology has now forced method – and this is exactly where the paradigm shift becomes unavoidable.

5.6 Thermodynamic Openness and Boundary Permeability

Introduction

This section establishes thermodynamic openness and boundary permeability as the first structural implications of process ontology. If reality is fundamentally characterised by becoming rather than static being (Section 5.2), then systems can persist only through continuous exchange with their environment. Thermodynamic openness is therefore not a contingent attribute but a constitutive condition of system existence. And if exchange is constitutive, then boundaries cannot be impermeable barriers – they must function as selective interfaces that regulate flows while enabling the throughput on which system existence depends.

The section receives from Section 5.5 the three-level ontology (micro–meso–macro) and the recognition that economic complexity operates primarily at the meso-level. This section addresses the consequent question: what structural properties must meso-level systems possess? The answer is thermodynamic openness and boundary permeability – the physical and organisational conditions that enable exchange, adaptation, and far-from-equilibrium persistence.

The section provides to Section 5.7 the structural foundation for network architecture. If systems must be open and boundaries must be permeable, then the geometry of exchange relations – network topology – becomes causally primary. Section 5.7 will establish how these exchange relations are organised across nested levels.

As Prigogine and Stengers (1984, p. 127) emphasise, living and socio-economic systems are dissipative structures: they maintain organisation through continuous throughput, not despite it. Organisation is sustained by flux, not equilibrium. This insight transforms how we understand boundaries. As Santos and Eisenhardt (2005, pp. 491–508) show, organisational boundaries are "emergent, negotiated, and contested" rather than given – they are adaptive membranes, not fixed walls.

This section performs four functions:

1. It derives the structural implications of process ontology: if being is processual, systems must be thermodynamically open.
2. It establishes why equilibrium-based analysis, logically dependent on closure, is structurally inappropriate for economic systems.
3. It reconceptualises the firm as a metabolic processor: an entity that maintains identity through continuous throughput rather than fixed attributes.
4. It reconceptualises boundaries from walls to gradients: the firm "ends" where connection density drops below the threshold required for organisational reproduction.

Before proceeding, a terminological clarification is required. The use of thermodynamic language in economic analysis is not metaphorical. Economic systems literally depend on energy dissipation, material transformation, and information processing. As Georgescu-Roegen (1971, p. 3) argued, an economics that ignores thermodynamics neglects "the most fundamental constraints on all human activity." Thermodynamic openness is therefore not an optional analogy but a structural necessity.

5.6.1 Thermodynamic Foundations: Open Systems and Dissipative Structures

I. The Thermodynamic Foundation: Open and Closed Systems

Thermodynamics distinguishes systems according to their exchange relations with the environment. This distinction has direct implications for how economic systems must be conceptualised and analysed.

Von Bertalanffy (1968, pp. 139–154) grounded General Systems Theory in this insight: living systems are necessarily open because they maintain organisation against entropy increase. In isolated systems, the second law of thermodynamics implies monotonic entropy growth; in open systems, entropy can be exported, allowing local order to be sustained. Nicolis and Prigogine (1977, pp. 3–28) formalised this result, demonstrating that self-organisation is possible only in open systems operating far from equilibrium.

Firm-level illustration: Amazon. Amazon's logistics network exemplifies thermodynamic openness at scale. The system processes approximately 1.6 million packages daily through continuous flows of: (i) material inputs – products, packaging, fuel; (ii) energy inputs – electricity powering warehouses, data centres, and transport; and (iii) information inputs – orders, routing data, inventory signals from over 300 million active customer accounts.

The organisational capability that constitutes Amazon's logistics advantage exists only through these flows. Interrupt any one of them and the system rapidly degrades, as supply-chain disruptions during the COVID-19 pandemic made immediately visible. Amazon is not a static structure that occasionally exchanges with its environment; it is a continuously open process temporarily stabilised into a recognisable form. Revenue represents throughput, not accumulated stock.

The methodological implication follows directly. If economic systems are constitutively open, then analytical frameworks that require closure – equilibrium analysis, comparative statics, and strict *ceteris paribus* isolation – describe systems that do not exist. As Allen (2015, pp. 1–25) shows, the study of complex adaptive systems requires methods that preserve openness; closure-based analysis systematically eliminates the very dynamics responsible for adaptation and stability.

Schrödinger's (1944, pp. 67–75) observation that living systems "feed on negative entropy" applies directly to economic organisation. Firms maintain internal order by processing low-entropy resources into higher-entropy outputs while exporting disorder to the environment. Georgescu-Roegen (1971, pp. 276–315) demonstrated that economic processes are inherently entropic: production is not circular flow but irreversible transformation constrained by physical laws. Economic systems do not suspend thermodynamics; they operate within it.

Firm-level illustration: Digital economy. The digital economy makes this constraint especially visible. Global data centres consume approximately 1–2% of world electricity – around 200 terawatt-hours annually. Every computation generates heat; every storage operation requires power; every network transmission dissipates energy. Cloud infrastructures consist of physical servers, cooling systems, and power supplies operating continuously to

sustain informational order. Firms such as AWS, Azure, and Google maintain economic functionality through massive energy throughput. Digital firms are therefore metabolic processors: electricity is their primary input, organised information their output. "Serverless" computing obscures but does not eliminate this thermodynamic reality.

II. Dissipative Structures: Organisation Through Flow

Prigogine's work on dissipative structures provides the theoretical foundation for understanding how open systems maintain – and under certain conditions increase – organisation. Operating far from thermodynamic equilibrium, open systems can generate ordered structures through continuous energy dissipation. As Prigogine (1980, pp. 77–102) showed, order emerges through irreversible thermodynamic processes rather than in opposition to them. Dissipative structures thus exemplify order through instability, where fluctuation is the precondition for organisation.

A candle flame provides the canonical dissipative structure. It exhibits a stable, recognisable form – shape, temperature, colour – yet possesses no fixed components. The structure exists only as long as fuel, oxygen, and heat dissipation continue. The flame does not undergo combustion; it **is** combustion temporarily stabilised into pattern. Extinguish the flow and the structure vanishes. This directly reflects the principle developed in Section 5.3: stability arises as an outcome of process, not as its negation. Firms exhibit the same logic. Organisational identity is maintained through continuous throughput, not despite it.

Prigogine and Stengers (1984, pp. 143–176) identify three conditions for the emergence of dissipative structures.

First, far-from-equilibrium conditions are required. Systems near equilibrium relax toward stasis; systems driven sufficiently far from equilibrium can undergo spontaneous reorganisation. Nicolis and Prigogine (1977, pp. 46–65) show that qualitative structural change – bifurcation – occurs only beyond critical thresholds of disequilibrium.

Second, non-linear dynamics must be present. Feedback loops amplify fluctuations rather than dampening them. In linear systems, perturbations decay; in non-linear systems, small disturbances can be magnified into new organisational states.

Third, symmetry breaking introduces irreversibility and historical contingency. Competing potential states are resolved through fluctuation selection, producing path dependence. As Prigogine (1980, pp. 103–129) emphasised, this mechanism introduces genuine novelty into physical processes.

Firm-level illustration: Bénard convection. Bénard convection cells provide a precise physical analogue. A fluid heated from below remains static at low gradients, but beyond a critical threshold spontaneous hexagonal convection patterns emerge. Organisation is sustained only through continuous heat flow; remove the energy input and structure dissipates. Economic systems exhibit the same structural logic. Market configurations, industry architectures, and firm capabilities arise and persist through continuous resource flows. Foster and Metcalfe (2012, pp. 419–442) therefore characterise economic evolution as self-organisation far from equilibrium.

Firm-level illustration: Tesla. Tesla's manufacturing system illustrates dissipative-structure dynamics in an organisational context. The Fremont factory operates far from equilibrium through continuous flows of materials, energy, and information. Production lines, workflow coordination, and quality systems exist only through these flows. During semiconductor shortages (2021–2022), output declined sharply despite unchanged physical infrastructure, demonstrating that organisational capability resides in flow patterns rather than static assets. Tesla's adaptive responses – software redesigns, supplier reconfiguration, and vertical integration – were possible precisely because the system remained open.

Kauffman (1993, pp. 173–235) generalises this logic, arguing that complex adaptive systems operate at the edge of chaos: sufficiently far from equilibrium to generate novelty, yet structured enough to maintain coherence. Firms, as dissipative structures, must continuously regulate their distance from equilibrium to avoid rigidity on one side and disintegration on the other.

III. Why Economic Systems Cannot Be Closed

Economic systems are constitutively open across multiple dimensions. This openness is not contingent but necessary: closure would imply the cessation of economic activity itself.

Material exchange. Firms transform physical inputs into outputs through globally distributed supply chains. Manufacturing depends on raw materials; services rely on physical infrastructure; digital production requires hardware and facilities. Even highly dematerialised activities require continuous material throughput. Without material exchange, production halts.

Energy exchange. All economic activity requires energy transformation. The global economy consumes approximately 580 exajoules annually, and long-run growth has been historically coupled to rising energy use. Ayres and Warr (2009, pp. 147–185) show that "useful work" – energy services enabling productive activity – explains productivity growth more effectively than conventional factor inputs. Cleveland et al. (1984, pp. 890–918) similarly demonstrate that energy-quality-adjusted measures outperform capital–labour production functions in explaining U.S. growth. Without energy flow, economic organisation collapses.

Information exchange. Modern economies depend on continuous information processing. Financial markets, supply chains, and innovation systems rely on real-time data flows. Hayek's (1945, pp. 519–530) analysis of prices as information carriers anticipates contemporary infrastructures that process vast quantities of data annually. Without information exchange, coordination disintegrates.

Institutional exchange. Economic actors are embedded in institutional environments (Section 5.4). Legal systems, educational structures, cultural norms, and political stability are continuously reproduced through social interaction. These institutional supports cannot be assumed as static constraints; they require ongoing maintenance. Without them, economic activity lacks structural scaffolding.

Firm-level illustration: Boeing 737 MAX. Boeing's 737 MAX crisis illustrates the consequences of degraded openness. Information exchange deteriorated as safety concerns were filtered and risks were not communicated to pilots or regulators. Institutional exchange

narrowed through regulatory capture. Material exchange was cost-optimised without adequate quality verification. The result was catastrophic failure, measured in loss of life and economic cost. Recovery required reopening all channels: independent safety oversight, renewed regulatory engagement, and restored supplier accountability.

Firm-level illustration: Kodak. Kodak's collapse demonstrates organisational entropy driven by information closure. Despite early technological capability in digital imaging, internal cognitive and organisational filters suppressed signals of environmental change. As Tripsas and Gavetti (2000, pp. 1147–1161) document, dominant frames constrained interpretation even when information was available. The firm's decline reflects not technological incapacity but failure to maintain informational openness.

Firm-level illustration: Huawei. Huawei's supply-chain crisis further illustrates the necessity of openness. US sanctions (2019–2024) restricted access to semiconductor manufacturing, software ecosystems, and component suppliers – disrupting capabilities that depended on continuous external exchange. Although Huawei invested heavily in rebuilding alternative channels, the episode demonstrates that organisational capabilities are flow-dependent, not accumulated stocks. Prolonged closure produces irreversible degradation.

5.6.2 The Firm as Metabolic Processor

IV. The Firm as Metabolic Processor

Process ontology combined with thermodynamic openness requires a fundamental reconceptualisation of the firm. Firms are not production functions, bundles of assets, or nexuses of contracts. They are metabolic processors: organisational systems that maintain identity through continuous material, energetic, and informational throughput. As Boulding (1978, pp. 17–28) argued within an ecological framework, organisations persist as metabolic structures sustained by flows rather than by fixed substance.

This claim is not metaphorical. It reflects structural isomorphism between biological and economic organisation. Biological metabolism maintains organisational pattern through continuous material exchange: constituent elements are replaced while form persists. Firms exhibit the same logic. Employees turn over, capital depreciates, technologies change, and products evolve, yet organisational identity can remain stable. Maturana and Varela's (1980, pp. 73–95) concept of autopoiesis captures this property: systems reproduce their own organisation through the processes that constitute them.

Firm-level illustration: Toyota. Toyota exemplifies the firm as metabolic processor. Its operations coordinate continuous material flows – tens of thousands of components per vehicle sourced globally – through tightly synchronised just-in-time systems. These material flows are sustained by energy throughput and coordinated by real-time informational exchange across suppliers, plants, and markets. The physical infrastructure, workforce, and product portfolio have changed repeatedly since the company's origins, yet Toyota remains recognisably Toyota because the organisational pattern governing throughput has persisted. Annual vehicle output represents throughput volume; organisational identity resides in the flow architecture that produces it.

Firm-level illustration: Walmart. Walmart demonstrates metabolic organisation at scale. Its competitive advantage is not an inventory stock or a store footprint but the capacity to process extremely high volumes of transactions at minimal cost per unit. Millions of weekly customer interactions, continuous inventory turnover, and real-time pricing adjustments constitute a high-throughput metabolic system. What persists from Sam Walton's original enterprise is not physical capital or managerial personnel, but a flow pattern optimised for high-volume, low-margin retail metabolism.

Firm-level illustration: Netflix. Netflix illustrates metabolic persistence through complete substrate change. The firm has transitioned from physical DVD distribution to global digital streaming with no material continuity between the two instantiations. Yet organisational identity has remained stable because the underlying metabolic pattern – technology-mediated access to entertainment via subscription – persisted while material and technological inputs changed. The transition represents metabolic adaptation rather than organisational rupture.

Firm-level illustration: SpaceX. SpaceX highlights how radically different metabolic architectures can transform an industry. By increasing throughput frequency, shortening iteration cycles, and integrating reuse, SpaceX reconfigured the metabolic logic of aerospace production. Its valuation reflects not engineering superiority alone but the capacity to sustain high-velocity material and revenue flows under conditions where traditional aerospace models remain low-throughput and episodic.

Implications. If firms are metabolic processors, then firm analysis must prioritise flow patterns rather than static attributes. Competitive advantage reflects metabolic efficiency; strategy concerns the configuration of environmental exchanges; organisational capability consists in coordinating material, energetic, and informational throughput. Strategy becomes metabolic design: shaping flows to sustain organisational identity under competitive and institutional constraints.

V. The Closure–Equilibrium–Reductionism Triangle

Thermodynamic openness has direct and unavoidable implications for the methodological foundations of mainstream economics. Three core commitments – closure, equilibrium, and reductionism – form a mutually reinforcing triangle. This triangle is internally coherent, but it is structurally incompatible with economic reality.

Equilibrium Requires Closure. Equilibrium denotes a state in which opposing forces balance and the system exhibits no endogenous tendency to change. Such a state presupposes insulation from external influence. In open systems, however, external inputs are not accidental disturbances but constitutive conditions of existence. As Prigogine demonstrated, systems driven far from equilibrium do not converge toward stasis; they generate new structures. Classical equilibrium constructs – Walrasian general equilibrium, Arrow–Debreu contingent-claims models, and DSGE macroeconomic frameworks – therefore rely on closure assumptions to guarantee equilibrium existence and stability.

Closure Requires Reductionism. If a system is closed, all explanatory factors are internal. Under such conditions, the whole can be treated as the sum of its parts, since nothing external contributes to system properties. Representative-agent methodologies exemplify this logic.

Heterogeneity can be averaged away only if agents are effectively isolated from differentiating environments.

Reductionism Requires Equilibrium. Reductionism presupposes that systems can be decomposed into stable components whose properties remain invariant under isolation. This assumption collapses in non-equilibrium systems. Dynamic capabilities (Teece et al., 1997, pp. 509–533) cannot be decomposed into static elements precisely because they exist only through ongoing environmental interaction.

Empirical Demonstrations of Triangle Failure

The Global Financial Crisis (2008). Pre-crisis risk models assumed equilibrium conditions – stable distributions, mean reversion, and calculable correlations. These assumptions required closure, treating external shocks as rare perturbations rather than constitutive features. Financial systems violated all three assumptions. They were open to global capital flows, regulatory change, and technological innovation; they exhibited emergent systemic risk irreducible to individual exposures; and they operated far from equilibrium under continuous adaptation. When Lehman Brothers collapsed, events deemed virtually impossible within the model framework occurred because the framework itself was ontologically misaligned.

Long-Term Capital Management (1998). LTCM's convergence strategies relied on equilibrium assumptions: historical relationships would reassert themselves, liquidity would remain available, and correlations would stabilise. When Russia defaulted in 1998, external disturbance shattered these conditions. Correlations broke down, liquidity evaporated, and losses cascaded. As Mandelbrot and Hudson (2004, pp. 87–118) document, financial markets exhibit forms of randomness incompatible with equilibrium-based risk models.

Silicon Valley Bank (2023). SVB's risk management assumed a stable depositor base, calculable interest-rate sensitivities, and decomposable portfolio risks. In reality, the bank's depositor network exhibited emergent properties: a highly concentrated clientele embedded in dense informational networks. Once concerns emerged, network dynamics produced a \$42 billion withdrawal in a single day. No component-level analysis predicted this outcome.

The methodological implication is direct. Equilibrium analysis is meaningful only in systems insulated from reality. Economic systems cannot be insulated from reality because their existence depends on continuous environmental exchange. As Mirowski (2002, pp. 537–582) argues, the persistence of equilibrium methodology reflects disciplinary path dependence rather than empirical adequacy.

5.6.3 Boundary Permeability and Dynamics

VI. Boundaries as Permeable Interfaces

If economic systems must be thermodynamically open, then their boundaries cannot be impermeable barriers. They must function as selective interfaces that regulate flows of matter, energy, and information. Boundary permeability denotes the capacity of a system's boundary

to admit, filter, and modulate flows – allowing some exchanges while restricting others, with selectivity itself evolving over time.

Classical theories treated boundaries as binary. Coase (1937) asked why firms exist; Williamson (1975; 1985, pp. 15–42) explained where the line between firm and market should be drawn. Both assumed such a line exists. A CAS perspective rejects the binary: complex systems exhibit continua of integration, not sharp separations. Powell (1990, pp. 295–336) showed that network forms of governance constitute a distinct coordination regime. Grandori (1997, pp. 897–925) documented a wide spectrum of hybrid organisational forms.

Boundaries as Gradients

In complex adaptive systems, boundaries do not occur where interaction stops; they occur where interaction density declines. A system is not separated from its environment by zero connectivity but by falling connectivity.

Network science formalises this. Community detection does not locate borders; it identifies regions of relatively dense internal ties and sparse external ties. Girvan and Newman (2002) and Newman (2006) show that "boundaries" emerge where edge-density gradients change most steeply. A boundary is not a line; it is a zone of declining coupling.

Firm-level illustration: Apple. Where does "Apple" end? Apple employs about 164,000 people, but its operational reality extends far beyond this legal core. Apple's \$383 billion revenue is generated by a metabolic network whose connection density declines smoothly from core teams to global users:

Table 126: Apple as Metabolic Network: Graduated Boundary Zones Illustrating the Permeability and Relational Nature of Firm Boundaries

Zone	Relation to Apple
Core	Apple employees, internal engineering teams
High-density	Foxconn, Pegatron, chip designers, logistics integrators
Medium-density	2M+ iOS developers, app firms
Low-density	2B device users, accessory makers, retailers
Peripheral	Data brokers, advertisers, repair shops

The \$1.1 trillion App Store ecosystem occupies the boundary zone between Apple and its environment. There is no point where Apple stops and the world begins – only a gradient of participation.

Boundary-Spanning Is the Boundary

Thompson (1967) introduced the idea of boundary-spanning roles – individuals who mediate between organisation and environment. CAS theory goes further: the boundary is nothing but boundary-spanning activity. Every contract, API, standard, supplier relation, developer tool, and compliance interface is part of the boundary. The firm is not enclosed by a wall; it is surrounded by a thickened membrane of interactions.

Firm-level illustration: Nike. Nike employs ~79,000 people but depends on over 1 million factory workers across 500+ contract facilities. These workers are not "outside" Nike in any meaningful production sense. They exist in boundary zones of high metabolic coupling. Labour disputes arise precisely in these zones because gradient boundaries produce jurisdictional ambiguity. Who is responsible? Who controls conditions? These conflicts are structural consequences of gradient boundaries.

Selective Permeability

Boundaries in complex adaptive systems are not merely permeable; they are selectively permeable. They regulate which flows enter, which are blocked, and which are transformed at the interface. Selectivity is not static; it evolves as environments and internal requirements change (Aldrich and Ruef, 2006, pp. 133–167).

Selective permeability makes openness viable. Total permeability dissolves the system; total closure suffocates it. Complex systems survive by tuning permeability.

Firm-level illustration: TSMC. TSMC's foundry model depends on an unusually precise permeability regime. It must admit extraordinarily valuable customer IP while preventing leakage across rival customers and shielding its own process knowledge. TSMC's \$69 billion revenue flows through this selectively permeable membrane. What customers are buying is not simply capacity but boundary trust: a selectively permeable membrane that allows flows necessary for production while blocking flows that would destroy appropriability.

Appropriability Is Boundary Permeability

Teece's (1986) "appropriability regime" is, at bottom, a permeability problem: how easily do value-creating flows leak out of the organisation? CAS theory adds what Teece lacked: appropriability is not fixed. Firms actively re-engineer their membranes through patents, secrecy, modularity, complements, platforms, and ecosystem design.

Firm-level illustration: Pharmaceutical patents. Drug patents are engineered permeability switches. During patent protection, the boundary is highly impermeable: generics are excluded; rents are captured. Lipitor's \$125 billion lifetime revenue reflects low permeability. At expiration, permeability jumps. Imitators flood in; prices collapse; value leaks. AbbVie's Humira extended low permeability via a patent thicket – 100+ overlapping protections creating structural impedance to biosimilars. The \$200 billion lifetime revenue was not due to chemistry alone but to boundary architecture.

Boundaries as Interfaces

Star and Griesemer's (1989, pp. 387–420) "boundary objects" capture the mechanism: standards, contracts, prototypes, and manuals coordinate across heterogeneous communities because they are locally adaptable yet globally stabilising. They do not eliminate difference; they stabilise interaction across difference.

APIs provide the clearest modern form of boundary-as-interface. An API specifies what requests are admissible, what formats are valid, what actions are permitted, and what remains inaccessible. The interface is simultaneously boundary and constitution. Platform competition is therefore competition between interface regimes (Jacobides, Cennamo, and Gawer, 2018).

Firm-level illustration: Stripe. Stripe's payment API processed \$817 billion in transaction volume in 2023 by functioning as a globally standardised interface. Stripe's 3.1 million customers do not interact with Stripe's organisation; they interact with its interface specification. The boundary is the API. Stripe's \$50 billion valuation capitalises the economic power of a precisely designed interface membrane.

Firm-level illustration: McDonald's. McDonald's does not operate 40,000 restaurants. It operates interfaces. The franchise agreement specifies what flows across the boundary: brand, recipes, IT systems, training, approved suppliers, and marketing on one side; capital, labour, local knowledge, and operations on the other. Royalty streams ($\approx 4\%$ of sales) are interface prices. McDonald's 2023 system-wide sales of \$130 billion did not occur inside McDonald's but through its boundary design.

VII. Firm Boundaries as Connection-Density Gradients

Gradients, selective permeability, and interfaces converge in a single claim: firm boundaries are connection-density gradients. They are zones where interaction intensity, governance, and mutual dependence change smoothly rather than abruptly. Transaction cost economics assumes a binary (inside hierarchy vs outside market); complex systems reveal a topology of intermediate density zones.

Zenger and Hesterly (1997, pp. 209–226) show that firms allocate decision rights far beyond formal hierarchies. Baker, Gibbons, and Murphy (2002, pp. 39–84) demonstrate that much coordination is governed by relational contracts that are economically binding but legally unenforceable. Santos and Eisenhardt (2005, pp. 491–508) show that firms have multiple boundary dimensions – power, competence, efficiency, and identity – which rarely coincide. Helfat and Raubitschek (2000, pp. 961–988) demonstrate that capabilities follow these density gradients rather than legal ownership. Dyer and Singh (1998, pp. 660–679) show that "relational rents" arise precisely in intermediate-density zones where firms are neither fully separate nor fully integrated.

Firm-level illustration: Boeing 787. The 787 programme illustrates boundary-gradient failure. Boeing pushed highly interdependent activities into loosely coupled supplier zones – placing high-coordination tasks in low-density regions. The result was integration failure. Boeing later pulled critical nodes inward and tightened interfaces. The boundary was never a line; it was a mis-shaped gradient.

Firm-level illustration: AWS/Netflix. Netflix runs on AWS. Legally Netflix is outside Amazon; operationally it is tightly coupled to AWS infrastructure. The relevant boundary is not a corporate line but a gradient of dependence: AWS servers and APIs → Netflix applications → Netflix users. Outages propagate through this gradient. This is what a real boundary looks like in an open system.

Firm-level illustration: Uber. Uber drivers occupy a gradient zone between employment and market participation. Algorithmic control pushes them inward; asset ownership and schedule flexibility push them outward. Legal conflicts exist because the boundary is not a line but a field. Uber's \$138 billion in gross bookings (2023) flows through this contested membrane.

Structural insight. Boundaries are control surfaces. Where gradients steepen or flatten determines who captures value, who learns, who bears risk, and where adaptation capacity accumulates. A theory of the firm cannot be completed by contracts or costs alone; it must incorporate CAS topology.

VIII. Boundary Dynamics: Migration Across Connection Space

Boundaries in complex adaptive systems are not stable features that occasionally change. They are dynamic patterns of connectivity that evolve continuously. Because boundaries are gradients, they can thicken, thin, shift, fragment, or recombine. Firm evolution is therefore not only expansion or contraction of "what the firm owns," but reconfiguration of connection gradients – what it tightly integrates, what it loosely couples, and what it expels.

Langlois's (2003, pp. 351–385) "vanishing hand" documents a topological shift: coordination increasingly occurs through modular architectures, standards, and platforms – redistributing density across networks rather than concentrating it within hierarchy. This is not organisational decline but density relocation.

Firm-level illustration: IBM. IBM's half-century trajectory is a boundary-gradient migration. In the 1980s IBM was a dense vertical stack: chips, hardware, operating systems, applications, sales, and services were all tightly coupled inside a high-density boundary. That topology collapsed under competitive and technological pressure. Under Gerstner, IBM expelled low-differentiation nodes (PCs, manufacturing) and thickened gradients around services, software, and integration. The Red Hat acquisition (2020) inserted IBM into a new open-source ecosystem gradient rather than reinternalising software production. The Kyndryl spin-off (2021) further thinned internal density around infrastructure. IBM did not "shrink" or "grow." It migrated its boundary gradients toward higher-value connection zones.

Firm-level illustration: Pharmaceuticals. Pharma no longer has boundaries; it has research-development-manufacturing-distribution gradients. Universities, biotech startups, CROs, CDMOs, specialty pharmacies, and digital health firms all occupy differentiated connection zones around large pharmaceutical integrators. Pfizer–BioNTech was not a partnership across a boundary; it was a temporary high-density gradient optimised for mRNA vaccine throughput. The \$37 billion COVID revenue did not belong to either firm – it emerged from the interface topology between them.

Firm-level illustration: Automotive electrification. The EV transition reallocates density from engines/transmissions toward batteries, software, power electronics, and autonomy. Every automaker is attempting to re-anchor its gradients in unfamiliar capability zones. Tesla internalised batteries and software early, thickening its gradient there (\$87 billion revenue). Volkswagen attempted to grow software density via CARIAD and failed. What appears as a technology race is also a topological race over where boundary density will reside.

Firm-level illustration: Semiconductor industry. The global semiconductor industry exemplifies thermodynamic openness at planetary scale. Semiconductor manufacturing requires the most complex supply chains ever created: extreme ultraviolet lithography machines from ASML (Netherlands, \$350 million each); ultra-pure chemicals from specialised suppliers; silicon wafers from specific refiners; design tools from Cadence and Synopsys; fabrication at TSMC and Samsung foundries. No firm – and no country – can produce advanced semiconductors in a closed system. Intel's difficulties in matching TSMC despite investments exceeding \$50 billion demonstrate that capability resides in network position, not in isolated capital. The 2021–2023 chip shortage revealed system fragility: closure of any exchange node cascaded through the entire network. Semiconductor production is a planetary metabolic system whose organisation exists only through continuous global exchange.

Firm-level illustration: Media & streaming. Streaming annihilated the old content-distribution-infrastructure gradient. Studios, networks, and cable operators were once distinct zones. Platforms collapsed these gradients into integrated pipelines. Disney's Fox acquisition and AT&T's Time Warner merger were boundary compression attempts – trying to pull content and distribution into a single high-density zone. Netflix bypassed the old topology entirely by creating a direct producer-platform-consumer gradient.

Firm-level illustration: Shopify. Shopify exists as a layered interface stack. Merchants interact with Shopify via storefront templates, checkout APIs, fulfilment integrations, analytics dashboards, and payment systems. Developers interact through the App Store API layer. Customers interact through front-end interfaces. Shopify's \$135 billion in GMV (2023) flows through these interfaces. When Shopify modifies an API or fee structure, it literally reconfigures the economy that runs on it. Shopify is not a software firm hosting merchants. It is a boundary machine.

Firm-level illustration: Fintech. Finance is being re-topologised by APIs. Stripe, Plaid, Chime, Robinhood, SoFi, and cloud-native banks are nodes in a redistributed boundary gradient where payments, identity, credit, custody, and compliance now sit at different density levels across multiple firms. The binary boundary collapsed into a mesh.

Structural law. Firms do not merely grow or shrink; they move their boundaries across connection space. Strategy becomes boundary-gradient engineering.

IX. Counter-Arguments and Responses

Counter-Argument 1: "Thermodynamic concepts are merely metaphors for economic systems."

Response. This objection misunderstands the claim. Thermodynamic concepts are not used metaphorically but structurally. Economic systems require energy inputs, transform material resources, and export entropy to their environment. The global economy consumes approximately 580 exajoules of energy annually; firms depend on electricity, fuel, and human metabolic energy; economic production generates waste heat, pollution, and depreciated capital. These are not analogies but physical processes.

Georgescu-Roegen (1971, pp. 276–315) demonstrated that economic activity is an entropic process. Ayres and Warr (2009, pp. 147–185) show empirically that "useful work" explains economic growth more effectively than conventional factor inputs. Cleveland et al. (1984, pp. 890–918) similarly demonstrate that energy-quality-adjusted measures outperform capital–labour production functions in explaining U.S. growth. Thermodynamics therefore captures what economic systems are constrained to do, not merely what they resemble.

Counter-Argument 2: "Markets can reach equilibrium even in open systems."

Response. This objection conflates market clearing with thermodynamic equilibrium. Markets may clear in the sense that prices adjust until supply equals demand, but this is a dynamic balance, not a static rest state. As Arthur (2014, pp. 1–25) shows, markets are perpetually adaptive and innovative; they do not converge to rest states. Kirman (1992, pp. 117–128) further demonstrates that aggregation problems prevent micro-level equilibrium conditions from scaling to macro-level equilibrium.

Counter-Argument 3: "Firms can be meaningfully studied as closed systems in the short run."

Response. This objection confuses analytical convenience with ontological accuracy. The COVID-19 shock demonstrated how rapidly presumed "short-run" stability collapses: supply chains, demand patterns, labour markets, and policy regimes all shifted within months. As Minsky (1986, pp. 194–230) argued, periods of apparent stability are precisely when destabilising dynamics accumulate. Lawson (1997, pp. 234–262) shows that event regularities predicted by closure-based models systematically fail in open economic systems.

Counter-Argument 4: "Openness makes systems too complex to model rigorously."

Response. This objection conflates tractability with validity. Methods appropriate for open systems already exist: agent-based modelling, network analysis, and system dynamics explicitly preserve openness. The Santa Fe Artificial Stock Market and subsequent agent-based work demonstrate that realistic market phenomena emerge from open-system assumptions. Farmer and Foley (2009, pp. 685–686) show that such models capture dynamics missed by equilibrium frameworks. Epstein and Axtell's (1996, pp. 1–20) *Sugarscape* simulations demonstrate how macro-patterns emerge from simple micro-rules without closure assumptions.

Counter-Argument 5: "Legal boundaries are the real boundaries."

Response. Legal boundaries matter, but they codify stabilised relational patterns rather than constituting operational boundaries. Firms routinely operate across and through legal boundaries via alliances, joint ventures, quasi-integration, platform governance, and relational contracting. Toyota's keiretsu suppliers are legally independent yet operationally quasi-internal; Uber's drivers are legally independent yet algorithmically governed. The gradient framework does not deny law; it explains what law is stabilising.

Counter-Argument 6: "Transaction cost economics already explains boundaries."

Response. TCE (Williamson, 1985, pp. 15–42) correctly identifies why some activities are internalised. But it presupposes a binary topology: inside hierarchy vs outside market. Empirical organisation does not follow this dichotomy. Platform governance, long-term partnerships, keiretsu systems, open-source ecosystems (West and Gallagher, 2006), and franchising all occupy intermediate density zones. The CAS framework generalises TCE: asset specificity and uncertainty explain where gradients steepen and where density concentrates, rather than implying a sharp line.

Counter-Argument 7: "Gradients imply anything can be part of the firm."

Response. Gradients are structured and differentiable. Actors occupy zones with measurably different interaction intensity, governance, and dependence. The framework increases precision by allowing degrees of belonging rather than forcing a false dichotomy.

Counter-Argument 8: "Permeability makes firms vulnerable."

Response. Selective permeability is protection. Total closure blocks learning and resource access; total openness dissolves identity. Open innovation (Chesbrough, 2003), platform ecosystems, and modular production networks all demonstrate that value is created through carefully designed permeability. Teece's (2007) dynamic capabilities framework makes this explicit: sensing, seizing, and reconfiguring all require boundary-spanning flows. The strategic problem is not openness versus closure, but how permeability is structured and governed.

X. Conclusion

This section has established thermodynamic openness and boundary permeability as the first structural implications of process ontology. If reality is fundamentally characterised by becoming rather than static being (Section 5.2), then systems require continuous exchange with their environment to persist, and boundaries must function as selective interfaces enabling that exchange.

The argument proceeded through eight movements:

1. Distinguishing open from closed systems
2. Developing the theory of dissipative structures
3. Demonstrating why economic systems cannot be closed
4. Reconceptualising the firm as a metabolic processor
5. Exposing the structural impossibility of the closure–equilibrium–reductionism triangle
6. Reconceptualising boundaries as gradients and permeable interfaces

7. Establishing firm boundaries as connection-density gradients
8. Analysing boundary dynamics as migration across connection space

Transformation Table

Table 127: From Closed Walls to Open Gradients: Thirteen-Dimension Transformation of Firm Boundary Ontology

Dimension	FROM (Closed/Wall)	TO (Open/Gradient)
Exchange environment with	Negligible or controllable	Constitutive
System boundary	Fixed barrier / Binary line	Permeable interface / Continuous gradient
Equilibrium	Natural attractor	Impossible ideal
Energy	Optimisation parameter	Constitutive requirement
Entropy	Ignorable	Fundamental constraint
Firm identity	Static attributes	Flow pattern / Autopoietic reproduction
Stability	Absence of change	Dynamic balance
Boundary function	Separation	Interface
Boundary definition	Legal/physical edge	Connection-density zone
Firm scope	What firm owns	Where connection intensity is high
Strategy	Make-or-buy	Gradient design / Boundary engineering
Appropriability	Fixed regime	Engineered membrane permeability
Time	Reversible	Irreversible

Forward Integration

The core insight is simple but decisive: without exchange there is no adaptation, no evolution, and no far-from-equilibrium organisation. Economic systems cannot be closed; equilibrium requires closure; therefore, equilibrium analysis systematically misrepresents economic reality. Boundaries do not separate systems from their environment; they constitute the interface through which system and environment co-define one another.

Contemporary business strategy increasingly unfolds in boundary gradient zones. Platform business models are fundamentally gradient strategies: Apple's App Store, Amazon Marketplace, and Salesforce AppExchange create value by designing interfaces that structure third-party participation. Ecosystem strategy (Adner, 2017) consists of orchestrating value creation across actors located at different connection densities. Traditional industry frameworks (Porter, 1980) presuppose stable, clearly bounded competitive arenas. Gradient boundaries reveal that competition now occurs in overlap zones between partially connected systems.

Section 5.7 now follows with necessity: if systems are open and boundaries are gradients, then the geometry of exchange relations becomes causally primary. Network topology is not merely

a descriptive tool but the structural expression of how openness and permeability are organised across nested levels. The question shifts from whether systems are open to how openness is architecturally structured.

The Aha Moment

Firms do not *have* processes; they *are* processes temporarily stabilised into recognisable patterns. Organisational identity persists even as employees, assets, technologies, and products change completely because what endures is the pattern of throughput, not the substance. And boundaries do not separate firms from their environment; they constitute the interface through which firm and environment co-define one another. Change the boundary gradient and you change the firm – not by moving a line, but by reshaping the pattern of exchange through which the firm exists.

5.7 Network Architecture and Meso-Level Structure

Introduction

This section establishes **multi-level nested architecture**, **network topology**, and **network clusters** as the structural implications of process ontology that complete the dissertation's reconceptualisation of the firm. If systems exhibit emergence (Section 5.1), operate through processual becoming (Section 5.2), and maintain identity through thermodynamic openness and permeable boundaries (Section 5.6), then they must be understood as hierarchically organised networks of complex adaptive systems nested within other complex adaptive systems, with the meso-level – the level of clusters, routines, and organised structures – as the domain where economic life actually unfolds.

The section receives from Section 5.6 the structural foundation of thermodynamic openness and boundary permeability. If boundaries are connection-density gradients rather than walls, then the geometry of those connections – network topology – becomes causally primary. This section addresses the consequent question: how is relational structure organised across scales?

The section provides to Section 5.8 (Interconnections) the complete structural architecture of CAS. With emergence, process, openness, permeability, nesting, topology, and meso-level clusters established, the dissertation can then develop the interconnections between ontology, epistemology, and methodology that complete Chapter 5.

This section performs five functions:

1. It derives multi-level nested architecture from emergence: if new properties appear at each scale, analysis must be multi-level by construction.
2. It introduces near-decomposability as the mechanism through which complex systems combine coherence with adaptability.

3. It establishes network topology as the structural realisation of relational ontology: if entities are constituted by relations, network structure captures economic reality at its deepest level.
4. It reconceptualises the firm as a topological entity defined by network position rather than internal attributes.
5. It establishes network clusters as the canonical meso-level structure – the dissertation's distinctive contribution – resolving puzzles about firm existence, boundaries, capabilities, and identity.

As Simon (1962, p. 468) observed, "hierarchic systems will evolve far more quickly than non-hierarchic systems of comparable size" – hierarchy is not an analytical convenience but an evolutionary necessity. And as Barabási (2002, p. 7) states, "networks are everywhere... they form the very fabric of existence."

5.7.1 Hierarchical Architecture and Near-Decomposability

I. The Architecture of Complexity: Why Hierarchy Exists

Simon's (1962) analysis demonstrated that hierarchy is universal among complex systems – biological, social, and economic. This universality is not accidental. It reflects deep evolutionary logic: hierarchic systems are more robust to disruption than flat systems.

His watchmaker parable remains decisive. A system built from stable subassemblies survives interruption; a system built as one undifferentiated mass collapses. Complexity without hierarchy cannot persist. Complex systems evolve hierarchical organisation because only hierarchical architectures survive continuous disturbance.

Economic organisation obeys the same principle. Firms, industries, and markets did not become hierarchical through planning; they became hierarchical because non-hierarchical economic structures cannot remain coherent under continuous disturbance.

Hierarchy here does not mean authority. It means ontological nesting. A market is hierarchical because it is composed of firms; firms of divisions; divisions of teams; teams of individuals. No command structure is implied – only structural embedding.

Firm-level illustration: Alphabet. Alphabet Inc. exemplifies nested CAS architecture. At the top level, Alphabet is a portfolio system (\$307 billion revenue). Within it, Google, YouTube, Waymo, DeepMind, Verily, and Other Bets operate as semi-autonomous CAS. Within Google, Search, Cloud, Ads, Android, and Maps constitute interacting subsystems. Within Cloud, infrastructure, data analytics, AI, and security form distinct adaptive modules. Within those, engineering teams of 5–10 people operate as micro-CAS.

Each level generates properties not possessed by the level below: team culture, division strategy, corporate portfolio coherence. No level can be reduced to another. When Alphabet

reorganises AI, causation propagates downward (team assignments) and upward (portfolio risk and valuation). Alphabet's 2015 restructuring made this architecture explicit – recognising that different subsystems evolve at different speeds.

Level Autonomy: "More Is Different"

Anderson's (1972) principle that "more is different" establishes that each hierarchical level has its own laws. Higher-level phenomena are not derivable from lower-level descriptions even in principle. This is not epistemic weakness; it is ontological structure. Psychology does not collapse into neuroscience; biology does not collapse into chemistry; economics does not collapse into individual choice. Markets are not scaled-up agents. Firms are not collections of preferences. They are organised systems with their own causal powers.

Firm-level illustration: Amazon's Multi-Scale CAS. Amazon's fulfilment network exhibits hierarchical CAS structure. At the global network level, p. 175+ fulfilment centres and routing algorithms. At the regional level: continental logistics optimisation. At the facility level: throughput and bottleneck dynamics. At the team level: coordination, fatigue, learning. At the individual level: picker decisions and scan rates.

Each level generates patterns invisible to the others. Labour disputes, logistics failures, and growth dynamics propagate across levels, not within a single plane. Amazon's economic identity is the stability of this nested architecture, not any particular node.

Bridge: From Interaction to Architecture

Complex Adaptive Systems are not merely complicated – they are architecturally real. Their defining feature is not that agents interact, but that interactions crystallise into stable, nested structures that exist across multiple levels of organisation. Firms, industries, and economies do not float in a flat space of exchanges; they are built as layered architectures of interacting subsystems, each with its own dynamics, constraints, and emergent properties.

This is where mainstream economics fails at the most fundamental level. Its error is not primarily behavioural or empirical. It is architectural. By collapsing the economy into a two-level world – individuals and aggregates – it erases the hierarchical scaffolding that makes complex organisation possible. What it treats as "macro outcomes" are in reality products of meso-level architectures – organisations, routines, networks, platforms, and institutional structures that mediate between individuals and aggregates.

II. Near-Decomposability: The Key to Complex System Stability

Simon's (1962, pp. 473–476) most profound contribution was the concept of near-decomposability: the property whereby complex systems exhibit strong interactions within levels and weaker – but non-zero – interactions between levels. Near-decomposability is not a modelling convenience; it is the structural condition that makes complexity itself possible. Without it, no system could remain both coherent and adaptive over time.

If a system were fully decomposable (zero interaction across levels), its parts would be independent: no emergence, no coordination, no system identity. If it were fully non-decomposable (equal interaction across all levels), every disturbance would propagate instantaneously through the entire system, producing instability and collapse. Near-decomposability provides the Goldilocks architecture: enough coupling to sustain coherence, enough decoupling to preserve stability. As Ahl and Allen (1996, pp. 45–72) show in ecological systems, interactions within trophic levels are far stronger than between levels – precisely what allows ecosystems to persist despite constant perturbation.

Firm-level illustration: Toyota. Toyota's production system exemplifies near-decomposability in action. Within each plant, interactions are dense and tightly coupled: production lines, quality circles, kanban systems, andon cords, and just-in-time coordination. Toyota's Georgetown, Kentucky plant – producing over 550,000 vehicles annually with 8,000 employees – operates as a highly integrated meso-system.

Between plants, however, interactions are weaker but non-zero: platform sharing, supplier networks, component standardisation, and corporate coordination. This architecture explains Toyota's resilience during the 2021 semiconductor shortage. Individual plants re-optimised schedules and prioritised high-margin models, while corporate coordinated supplier relations across the network. Production fell roughly 15%, but the system did not collapse. If Toyota had been fully decomposable, coordination would have been impossible. If fully non-decomposable, the shock would have cascaded catastrophically. Near-decomposability enabled graceful degradation.

The Toyota Production System itself is a meso-level structure: neither reducible to individual skills nor to corporate strategy, but a stable pattern of routines that binds the hierarchy together.

Firm-level illustration: Procter & Gamble. P&G's multi-brand architecture exhibits near-decomposability at corporate scale. The firm manages more than 65 brands across ten categories – Beauty (\$15 billion), Grooming (\$7 billion), Health Care (\$10 billion), Fabric & Home Care (\$28 billion), and Baby/Feminine/Family Care (\$20 billion) – generating \$82 billion in revenue (2023).

Within categories, interactions are strong: Tide and Gain compete for shelf space, advertising budgets, and R&D. Between categories, interactions are weaker: razors do not directly compete with diapers. This allows P&G to operate semi-autonomous business units with their own P&L, marketing, and innovation pipelines. When P&G divested Duracell (2016), it could restructure the portfolio without destabilising the firm. Yet cross-category synergies remain real: shared purchasing, corporate services, and innovation transfer. Near-decomposability enables both autonomy and synergy.

Firm-level illustration: Johnson & Johnson. J&J operates over 250 companies across Pharmaceuticals (\$54 billion), MedTech (\$30 billion), and Consumer Health (spun off as Kenvue, \$15 billion). Each operating company functions as a semi-autonomous CAS with its own R&D, manufacturing, and commercial systems. Within segments, coupling is strong; between segments, weaker. This allowed J&J to execute the \$80 billion+ Kenvue separation while preserving corporate continuity.

Yet coupling remains real. The \$8.9 billion talc litigation settlement (2024) originating in consumer products propagated upward to corporate valuation and across to pharmaceutical reputation. Near-decomposability does not eliminate systemic risk – it structures how it propagates.

Analytical implications. Simon (1962, pp. 477–478) showed that near-decomposable systems can be analysed level by level, treating between-level effects as parameters rather than explicit variables. This makes scientific inquiry tractable without ontological reduction. But this analytical convenience must never be confused with ontological separation: the levels remain real and causally coupled. As Thompson (1967, pp. 51–82) demonstrated, buffering and bridging mechanisms exist precisely because coupling across levels is weak – but never zero.

III. Nested CAS and Multi-Level Causation

Complex adaptive systems are not merely hierarchical but nested: CAS within CAS within CAS. Each level is itself a complex adaptive system exhibiting emergence, self-organisation, and adaptation. This nesting creates multi-scale causal architecture in which patterns at each level both influence and are constrained by patterns at other levels.

Holland (1995, pp. 10–40) demonstrated that CAS share a common architecture at every scale: heterogeneous agents, nonlinear interaction, emergent structure, and adaptive selection. What changes across levels is what counts as an agent and what counts as interaction. Individuals are agents within teams; teams are agents within firms; firms are agents within industries; industries are agents within economies. At no level does behaviour reduce to the level below, because the interaction topology changes.

A team is not a collection of individuals. It is a CAS whose agents are people. A firm is not a collection of people. It is a CAS whose agents are teams. An industry is not a collection of firms. It is a CAS whose agents are organisations.

This nesting is not metaphorical. It is the architecture that makes large-scale economic order possible at all.

Firm-level illustration: The semiconductor system. The semiconductor sector reveals nested CAS architecture with exceptional clarity.

- At the global level, geopolitics, trade policy, and capital flows form a CAS: the US CHIPS Act (\$52 billion), the EU Chips Act (€43 billion), and China's semiconductor strategy generate emergent reshoring and bifurcation dynamics that no single government controls.
- At the industry level, architectures such as x86, ARM, and RISC-V form competing platform ecosystems – emergent coordination structures not designed by any one firm.
- At the firm level, TSMC (\$69 billion), Samsung (\$64 billion semiconductors), and Intel (\$54 billion) operate as CAS whose production, yield learning, and technological trajectories emerge from millions of engineering interactions.

- At the fab level, facilities like TSMC Fab 18 (8,000+ engineers, 24/7 operation) function as CAS with their own feedback loops, learning curves, and failure modes.
- At the team level, process engineering, lithography, and yield-management groups adapt daily to local perturbations.

The 2021–23 chip shortage was not "caused" by any firm. It emerged from cross-level feedback loops: pandemic demand shocks (global) altered consumer electronics production (industry), which reallocated fab capacity (firm), which starved automotive supply chains (industry), which disrupted manufacturing globally (macro). TSMC's \$40 billion Arizona fab is not just a corporate investment – it is a multi-level adaptive response to geopolitical, industrial, and technological forces interacting across nested CAS.

No single level explains what happened. The phenomenon only exists between levels.

Edge-of-chaos at every level. Kauffman (1993, pp. 173–235) showed that CAS maximise adaptability at the edge of chaos – the boundary between rigid order and disorder. Critically, this condition applies at every level of nesting. Teams that are too rigid cannot innovate; teams that are too chaotic cannot execute. Firms that are too bureaucratic stagnate; firms that are too anarchic disintegrate. Industries that are too stable lock in obsolete technologies; industries that are too turbulent collapse investment. Each level must maintain its own local edge-of-chaos equilibrium for the system above it to function.

Multi-Level Causation: Upward Emergence and Downward Constraint

In nested CAS, causation flows in both directions simultaneously: upward through emergence (micro-level interactions generate macro-level structures) and downward through constraint (macro-level structures shape micro-level possibilities). This bidirectional causation distinguishes CAS from both reductionist systems (which recognise only upward causation) and holistic systems (which recognise only downward causation). CAS operate through circular causality across levels.

Upward causation – emergence – was established in Section 5.1: market prices emerge from trader interactions; organisational culture emerges from employee interactions; industry standards emerge from firm interactions. But these emergent patterns do not merely appear; once formed, they become causal forces in their own right. They shape what agents can do, what strategies are viable, what behaviours are rewarded or punished.

This is downward causation. Institutions, technologies, platforms, standards, and organisational architectures are macro-level patterns that channel micro-level action. As Hodgson and Knudsen (2010, pp. 33–57) emphasise, institutions carry information and constraint structures that shape individual behaviour without determining it. They do not dictate action – but they configure the space of possible actions.

In CAS, micro actions create macro structures that then reshape micro actions in a self-reinforcing loop. The system evolves through recursive feedback between levels. Giddens' (1984) structuration theory formalises this duality: social structures are both the medium and

the outcome of action. Archer's (1995) morphogenetic cycle shows how this unfolds over time: structures condition action; action produces new structures; new structures condition future action.

Firm-level illustration: The 2008 financial crisis. The 2008 crisis was not caused by "bad incentives" or "irrational investors." It was caused by multi-level circular causation.

Upward causation: Millions of individual mortgage decisions – subprime loans, refinancing, leverage – were bundled into mortgage-backed securities and CDOs, generating housing price inflation and systemic leverage.

Downward causation: Institutional structures – credit-rating models, securitisation pipelines, regulatory capital rules, and Federal Reserve interest-rate policy – made those micro decisions profitable, encouraged leverage, and masked risk. AAA ratings and Basel capital rules made risky assets appear safe, reshaping bank behaviour and trader incentives.

The macro-structure created by securitisation reshaped micro-behaviour; the micro-behaviour reinforced the macro-structure. When Lehman collapsed (\$600 billion in assets, September 2008), the system imploded because the loop snapped. The \$22 trillion global wealth loss and the Federal Reserve's \$4.5 trillion QE response were macro-level interventions attempting to re-stabilise the micro-level behaviour space.

No single level explains the crisis. It existed between levels.

Firm-level illustration: Uber. Uber's platform demonstrates multi-level causation in real time.

Upward causation: Individual rider and driver decisions generate emergent patterns – surge pricing zones, wait times, income distributions. Uber's \$138 billion in gross bookings (2023) is the aggregation of millions of micro-choices.

Downward causation: Uber's algorithmic architecture – dispatch rules, pricing formulas, rating systems – structures what drivers and riders can do. Drivers move to surge zones because the system signals them. Riders change behaviour because the platform penalises cancellations or increases prices.

When California passed AB5 and voters approved Prop 22, institutional changes (macro) reshaped platform governance (meso), which reshaped driver income and job security (micro). When drivers coordinated log-offs to manipulate surge pricing, micro-action attempted to reshape macro-algorithmic behaviour.

Uber is not a market or a firm. It is a multi-level CAS in continuous feedback between platform rules, user behaviour, and legal institutions.

Firm-level illustration: Netflix. Netflix illustrates nested CAS dynamics in action.

- At the team level, Netflix's "highly aligned, loosely coupled" model creates small CAS units that experiment, learn, and adapt.
- At the division level, content, engineering, and marketing operate as semi-autonomous CAS with emergent strategies.
- At the firm level, Netflix's trajectory – from DVDs (1997) to streaming (2007) to original content (\$17 billion annually by 2023) to gaming (2021) – was not centrally planned but emerged from interactions between teams, markets, and technologies.
- At the industry level, Netflix competes within a CAS of Disney+, Amazon Prime, HBO, and others. Pricing, bundling, and content strategies evolve through mutual adaptation, not optimisation.

When Netflix lost subscribers in 2022, the response propagated across levels: team-level cost controls, firm-level strategy shifts (ads, password enforcement), and industry-level realignment as competitors copied or countered Netflix's moves.

This is nested adaptation. No single agent is in control. The system evolves through multi-level feedback. IV. The Missing Meso-Level: Mainstream Economics' Fatal Gap

Mainstream economics systematically erases the meso-level through the representative agent construct. By assuming that a single optimising agent can stand in for an entire population, it collapses the economy's nested architecture into a flat micro–macro plane. This move makes equilibrium mathematics tractable – but it does so by eliminating the very layer where complex adaptive systems actually operate: the level of firms, routines, networks, institutions, platforms, and organised knowledge.

As Arthur (2013, p. 14) observes, orthodox economics has "no natural way to explore the meso-level." It jumps directly from individuals to aggregates, bypassing the organisational structures that mediate between them. This is not an empirical omission; it is an ontological deletion.

Kirman (1992, pp. 117–128) showed that the representative agent is mathematically incoherent: individual optimisation conditions do not aggregate into collective optimisation. But the deeper problem is not technical – it is architectural. The construct eliminates heterogeneity, interaction, and structure so that equilibrium can be solved. In doing so, it destroys the space in which coordination, learning, and institutional stability exist.

But the economy does not operate this way. Economic life unfolds in the meso-domain of organised patterns: firms coordinate production, routines encode knowledge, networks transmit information, institutions stabilise expectations, and platforms orchestrate ecosystems. These structures are not reducible to individual behaviour, nor are they visible in aggregate statistics. They are causal systems in their own right.

This omission was laid bare in the 2008 financial crisis. Risk models such as Value-at-Risk treated systemic risk as the aggregation of individual asset risks (micro → macro). But the crisis was driven by meso-level network structure: counterparty exposure webs, correlation clusters, liquidity cascades, and institutional interdependencies. AIG's \$440 billion in credit-default swaps did not matter because of AIG's balance sheet alone; it mattered because of AIG's network position – it was a hub in the global financial system. The firm was not "too big to fail" but too interconnected to fail, a meso-level property invisible to representative-agent models.

The post-crisis regulatory response was an implicit admission of this failure. The Financial Stability Oversight Council, the Office of Financial Research, and the Basel III G-SIB framework were created to map and regulate interconnectedness, not just size. Stress tests now include network contagion scenarios. What had been excluded from theory was reintroduced in practice: the meso-level had to be rebuilt institutionally because it had been erased analytically.

Dopfer, Foster, and Potts (2004) formalise this missing layer as the meso-level of rules – routines, technologies, standards, and institutions that coordinate micro actions and generate macro patterns. The meso-level is not "between" micro and macro in size but in causation. It is where economic order is stored and transmitted.

Firm-level illustration: Toyota TPS and NUMMI. Nelson and Winter (1982) showed that organisational routines are the true carriers of productive knowledge. Knowledge does not reside in isolated individuals (who come and go), nor in aggregate statistics (which are merely summaries). It resides in organised practices – ways of doing things that persist through turnover and generate consistent performance. These routines are meso-level structures.

Toyota's Production System makes this concrete. TPS – just-in-time, jidoka, kaizen, andon – is not located in any single worker, nor in Toyota's balance sheet. It exists as a distributed, structured set of routines enacted daily across plants and suppliers. When NUMMI implemented TPS, productivity doubled and defects fell by 90%. When NUMMI closed, the knowledge vanished – even though the workers and the factory remained. What was lost was the meso-level organisation. Mainstream economics cannot say where that knowledge was, because it has no category for it.

Firm-level illustration: McKinsey. McKinsey's competitive advantage does not lie in individual consultants or in firm-level reputation, but in its meso-level knowledge architecture: practice networks, staffing systems, codified methodologies, and institutional culture. These structures coordinate expertise across 130 countries and 45,000 employees. Consultants rotate; the system persists. That persistence is the meso-level.

The fatal gap in mainstream economics is therefore not a missing variable but a missing level of reality. By collapsing a hierarchical, nested economy into a two-level micro–macro schema, it eliminates the structures that make coordination, learning, and evolution possible.

5.7.2 Network Topology and Properties

V. Network Topology as Structural Foundation

If relations are ontologically primary (Section 5.4) and systems exhibit nested multi-level architecture, then economic systems must be understood as networks whose structure shapes and constitutes their dynamics. Network topology is defined here not merely as the pattern of connections among system components, but as the relational architecture through which those components acquire their identity, capabilities, and causal powers.

Network science emerged independently in sociology (Granovetter, 1973; Burt, 1992), physics (Watts & Strogatz, 1998; Barabási & Albert, 1999), and mathematics (Erdős & Rényi, 1960), converging on the same insight: system behaviour is shaped less by who the agents are than by how they are connected. The mechanistic paradigm's atomism – treating entities as self-contained units with fixed properties – cannot capture these effects. As Kirman (2016, p. 560) notes, economics' neglect of network structure produced systematic predictive failure, nowhere more clearly than in the 2008 financial crisis, when contagion propagated through financial networks invisible to atomistic risk models.

Relational Ontology Applied: From Philosophy to Structure

If relations are ontologically primary, entities cannot be fundamental. They are nodes in relational structures. Economic systems must therefore be represented not as collections of atoms but as networks of interaction. Network analysis becomes not merely useful but ontologically appropriate.

Granovetter's (1985, pp. 481–510) embeddedness thesis is often interpreted as contextual: agents behave within networks. In CAS ontology, embeddedness is constitutive: agents exist through networks. The firm does not operate within a network; the firm is its position in a network. This distinction between weak and strong embeddedness is fundamental: weak embeddedness treats networks as context; strong embeddedness treats networks as constitutive of economic action and identity.

Firm-level illustration: Apple. Apple manages relationships with over 200 major suppliers across 43 countries, coordinates \$274 billion in supply-chain spending (2023), supports 34 million developers, operates 523 retail stores, and intermediates billions of consumer transactions. Its competitive advantage does not reside inside Apple. It emerges from network position.

Apple Silicon (M1, 2020) did not succeed because Apple owned fabs – it succeeded because Apple occupied a privileged position in TSMC's 5nm manufacturing network, combined with a developer ecosystem able to optimise software at scale. Remove Apple from that network position and "Apple" as competitive entity ceases to exist.

The \$3 trillion market capitalisation prices Apple's position in a global production, developer, and platform network – not its factories, employees, or legal shell.

Firm-level illustration: Visa. Visa processed \$14.8 trillion in payments (2023) across 4.3 billion cards, 100+ million merchants, and 14,500 financial institutions. But Visa owns no deposits, produces nothing, and bears almost no credit risk. Visa is the network.

A Visa card without merchants is plastic. A merchant without cardholders is useless infrastructure. Visa's 80% gross margin and \$540 billion valuation arise from topological centrality – occupying a high-connectivity position in a two-sided network. Its market power is not asset-based; it is structural, arising from dense connectivity, switching costs, and network externalities.

Structural Realism: The Philosophical Foundation

Structural realism provides the philosophical grounding for network topology as ontological reality rather than modelling convenience. As Ladyman and Ross (2007, pp. 130–189) argue, what is real are structures and relations, not intrinsic properties of isolated entities. Objects do not precede relations; they emerge from them.

This claim has direct methodological consequences. If structure is ontologically primary, then network topology is not one analytical lens among many – it is the correct lens for relational systems. Atomistic economics commits a category error: it seeks causal power in agents when causal power resides in relational position. This is why adding networks to equilibrium models does not fix them – the ontology remains wrong.

Economic entities are therefore not substances with attributes but stable patterns of relations. A firm, a price, a market, or a platform is not a thing; it is a persistent configuration in a relational field. As French (2014, pp. 157–180) shows, prices, contracts, firms, and markets are structural features of economic networks, not objects with intrinsic identity.

Firm-level illustration: Venture Capital. The venture capital industry makes structural realism operational. Sequoia Capital, Andreessen Horowitz, and Benchmark do not outperform because they have more capital – they outperform because they occupy central positions in investment, founder, and exit networks.

Sequoia's \$85 billion AUM produces radically different returns from atomistically identical capital because Sequoia is embedded in Stanford, Y Combinator, founder networks, and elite co-investment circles. Its WhatsApp investment (\$60 million → ~\$3 billion) was not caused by superior "judgment" alone but by network access to deals, information, and acquirers. Hochberg, Ljungqvist, and Lu (2007) show empirically that network centrality predicts VC performance even controlling for capital and experience.

A "Sequoia" without its network is just money. The competitive advantage is structural.

Firm-level illustration: Alibaba. Alibaba's \$130 billion revenue (2024) emerges from its orchestration of a vast relational architecture, p. 900 million consumers, 10 million merchants, Alipay's 1.3 billion users, Cainiao's logistics network, and Alibaba Cloud's digital infrastructure.

Alibaba owns little of what creates its value. It coordinates everything. When Chinese regulators fined Alibaba \$2.8 billion (2021), they targeted structural dominance – exclusive dealing arrangements and platform control – not internal inefficiency. Jack Ma's statement that "Alibaba doesn't own anything" is not rhetorical – it is ontological. The firm is the network it organises.

Why Topology Matters: Same Nodes, Different Structure, Different Dynamics

The central result of network science is devastating for atomism: Same nodes + different topology = different system behaviour. System dynamics are governed not by what agents are, but by how they are connected. Watts (2004) and Barabási (2002) show that cascades, coordination, contagion, and resilience are all functions of network geometry, not agent psychology.

This destroys the representative-agent worldview. If outcomes depend on topology, then averaging agents erases the very structures that generate outcomes.

Firm-level illustration: Toyota vs GM. Toyota and GM source from many of the same semiconductor suppliers. Yet Toyota recovered from the 2011 earthquake in six months while GM required more than a year. The difference was network topology.

Toyota's keiretsu structure – dense ties, redundant links, long-term relational contracts – created high-information, high-trust topology. GM's hub-and-spoke, arms-length supplier network created fragility. When the 2021 semiconductor shortage hit, Toyota had 2–3 months of visibility into supplier inventories; GM had weeks.

Same suppliers. Same technology. Different topology → different resilience. Toyota's \$275 billion revenue is a topological achievement.

Firm-level illustration: Social media platforms. Facebook, Twitter/X, and LinkedIn all host user posts. But their network topologies differ radically. Facebook's friend-based clusters create community echo chambers. Twitter's follower graph creates broadcast cascades. LinkedIn's professional clusters produce reputation and expertise sorting.

Same nodes (users). Same actions (posting). Different topology → different politics, misinformation dynamics, and business models.

When Elon Musk altered Twitter's verification and algorithmic structure, user behaviour changed – not because users changed, but because the network geometry changed. Platform value is topological value.

VI. Key Network Properties and Mechanisms

Network science has identified specific topological properties that govern how complex systems behave. These properties do not merely influence outcomes; they structure what outcomes are possible.

Scale-Free Networks and Small Worlds

Barabási and Albert (1999) showed that preferential attachment generates scale-free degree distributions – networks dominated by hubs. Such systems are efficient and adaptive but structurally fragile: they are robust to random failure yet catastrophically vulnerable to targeted hub disruption.

Watts and Strogatz (1998) showed that small-world networks combine high clustering with short path lengths, enabling local cohesion and global coordination simultaneously.

These properties are not abstractions. They determine how shocks propagate, how information travels, how power concentrates, and how systems fail.

Firm-level illustration: The global airline network. The airline system consists of roughly 4,000 airports and 50,000 routes. Its topology is sharply scale-free. A handful of hubs – Atlanta, Dubai, Heathrow, Frankfurt – dominate global connectivity, while most airports have only a few links.

This structure produces both efficiency and fragility. The average number of flights between any two cities is about two. But when a hub fails, disruption cascades globally. Delta's 2016 Atlanta outage cancelled more than 2,000 flights worldwide not because Delta was inefficient, but because Atlanta is a network hub.

Dubai's betweenness centrality – sitting on shortest paths between Europe and Asia – explains Emirates' \$32 billion revenue better than aircraft quality or labour cost. Southwest's point-to-point topology produces very different economics from United's hub-and-spoke system. Same planes, different topology → different business model.

Firm-level illustration: Goldman Sachs. Goldman's advantage is topological. It occupies high betweenness centrality between corporations, investors, governments, and markets. M&A advisory exists because firms and capital pools are separated by structural holes. Goldman bridges those holes.

When Goldman advised on \$350 billion of deals in 2023, it was monetising its position in the global deal network. Capital, talent, and technology matter – but only because network position allows them to be activated. Where Goldman loses network centrality (e.g., in crypto), it loses power regardless of balance sheet strength.

Firm-level illustration: Hollywood. Hollywood is a textbook small-world, scale-free network. A handful of super-agents and producers connect thousands of artists, studios, and distributors. Kevin Feige's betweenness centrality explains Marvel's \$30+ billion box office more than any script.

Pixar's Braintrust is explicitly a topology-management device: dense clustering for creativity, with cross-cluster bridges for innovation. Hollywood's \$42 billion market is not driven by talent alone – it is driven by network geometry.

Four Foundational Network Mechanisms

Four results structure modern network science and economic reality:

1. **Granovetter's weak ties:** Weak ties carry non-redundant information. Career moves, innovations, and market information travel through acquaintances, not close friends.
2. **Burt's structural holes:** Power and profit accrue to actors who bridge otherwise disconnected groups. Structural holes create arbitrage opportunities.
3. **Barabási-Albert preferential attachment:** Networks naturally generate hubs through "rich get richer" dynamics. Hubs concentrate power and create vulnerability.
4. **Watts-Strogatz small-world structure:** High clustering plus short path lengths enables both local cohesion and rapid global diffusion.

These are not sociological curiosities. They are the operating system of modern capitalism.

Firm-level illustration: LinkedIn. LinkedIn monetises all four principles. Careers move through weak ties, not close colleagues. Structural holes between industries, regions, and firms create recruiting value. Influencers emerge through scale-free follower distributions. Small-world topology enables global professional reach with local reputation clustering.

LinkedIn's \$26 billion acquisition value (and \$15 billion annual revenue) reflects topological capital, not content. Reid Hoffman explicitly designed the platform around Granovetter's weak ties. LinkedIn is applied network theory.

Firm-level illustration: Pharmaceutical innovation. Biotech innovation is a network phenomenon. Firms that bridge academia, biotech, and Big Pharma capture the most value. Pfizer's partnership with BioNTech was a weak tie bridging a structural hole between mRNA science and global manufacturing.

That single bridge generated \$37 billion in 2022 revenue. Internal R&D could not have done it in time. Topology created the breakthrough.

Connective Complexity and Partial Connectivity

Mitchell's (2009) concept of connective complexity shows that emergent intelligence does not reside in agents but in connection patterns. Simple agents embedded in rich networks generate behaviours no agent could compute individually. Complexity is not in cognition; it is in topology.

Foster's (2005) concept of partial connectivity explains why complex systems exist at all. If everything were connected to everything, no structure could persist. If nothing were connected, no coordination could occur. Complex adaptive systems occupy the Goldilocks zone: connected enough for emergence, disconnected enough for differentiation.

Firm-level illustration: Amazon. Amazon's logistics system does not work because drivers are smart. It works because the delivery network computes. Ant-colony optimisation models route 1.6 million packages per day by exploiting network topology. No central planner could compute that solution. The network itself performs the computation.

When Amazon acquired Whole Foods, it did not buy groceries. It bought 400 new network nodes embedded in urban topology. Grocery was the pretext. Topology was the asset. Amazon's \$575 billion revenue depends on connective complexity: value is generated by network geometry, not agent intelligence.

Firm-level illustration: Financial markets. Markets require partial connectivity. If all traders saw all information, arbitrage would disappear. If no one saw anyone else, prices could not form.

High-frequency firms like Citadel and Virtu make money because they sit between fragmented pools of liquidity. Robinhood routes retail orders to Citadel because Citadel occupies the structural hole between retail and institutional order flow.

Market efficiency is not transparency. It is structured opacity – partial connectivity. The NYSE is not a market; it is a topological configuration that makes price discovery possible.

5.7.3 The Firm as Network Cluster

VII. The Firm as a Topological Entity

The cumulative argument is now decisive: A firm is not a container of resources. A firm is a position in a network.

- Boundaries are not legal demarcations; they are connection density gradients (Section 5.6).
- Capabilities are not internal stocks; they are positional effects.
- Strategy is not optimisation; it is topological engineering.

Firm-level illustration: Tesla. Tesla is not worth \$800 billion because it makes cars. It is worth \$800 billion because it sits at the hub of four converging networks: vehicle data (billions of miles of driving), charging infrastructure (50,000+ Superchargers), software (OTA updates to 5+ million vehicles), and energy systems (Powerwall, Megapack, grid storage).

More cars → more data → better AI → more demand → more cars. This is not scale; it is topological feedback. When Ford and GM adopt Tesla's charging standard, they are not conceding technology. They are submitting to network gravity.

Tesla's boundary is not the factory wall. It is the radius of its data and infrastructure network.

Firm-level illustration: WeWork vs Airbnb. WeWork failed because it had no topology. It assumed network effects without clustering, without bridging, without structural holes. Its members did not interact; no emergent value was created. It was a landlord pretending to be a platform.

Airbnb, by contrast, **is** a network: hosts, guests, trust systems, repeat interactions, and global liquidity pools. It occupies the structural hole between housing and travel.

Same rhetoric. Different topology. One collapsed. One became a \$75 billion platform.

Firm-level illustration: NVIDIA. NVIDIA is not a chip company. It is the hub of, p. 4+ million CUDA developers, deep-learning research standards, cloud AI infrastructure, gaming platforms, and autonomous vehicle stacks.

Once a network coordinates on CUDA, every competitor faces topological lock-in. Better chips cannot overcome worse topology. NVIDIA's moat is not hardware. It is network position crystallised into standard.

The implications for analysis are profound. Understanding firms requires network mapping – identifying positions, connections, structural holes, and centrality. Predicting firm behaviour requires topological analysis – how does position shape opportunity and constraint? Explaining firm performance requires structural analysis – what value derives from network position rather than internal attributes?

VIII. Network Clusters as the Canonical Meso-Level

Network clusters – dense local groupings within larger networks – occupy precisely the intermediate scale that mainstream economics fails to theorise. This subsection delivers the dissertation's distinctive contribution: a rigorous structural theory of the meso-level – the level at which economic organisation actually operates.

A structure qualifies as meso-level if it:

1. emerges from micro-level interactions,
2. exhibits properties absent at the micro level,
3. exerts causal power on both micro and macro outcomes, and
4. cannot be eliminated without loss of explanatory power.

Network clusters satisfy all four conditions. As Watts and Strogatz (1998) show, the clustering coefficient – the density of local connections – is a fundamental topological property that shapes system behaviour independently of individual nodes or global network size.

Cluster vs Aggregate

Aggregate properties are additive; cluster properties are emergent. As Newman (2006) shows, cohesion, coordination, resilience, and diffusion dynamics exist only at the cluster level and vanish under aggregation. Clusters are therefore not a modelling convenience – they are the economy's operative scale.

Firm-level illustration: Amazon's Fulfilment Network. Amazon's logistics system does not consist of isolated fulfilment centres (micro) or a single global warehouse (macro), but of regional clusters of facilities with dense internal coordination. The Northeast US cluster (25+ facilities) exhibits properties no single warehouse has: same-day delivery coverage, inventory pooling, surge absorption, and redundancy.

When Amazon promises same-day delivery to 72% of the US population, that capability exists only at the cluster level. The 250 million items processed during Prime Day 2021 emerged from cluster-level coordination, not from facility aggregation. Remove the cluster structure – keep the buildings but eliminate their dense coordination – and Amazon's delivery promise collapses. Amazon's \$575 billion revenue rests on meso-level architecture.

Firm-level illustration: Semiconductor clusters. The semiconductor industry is organised around geographically and functionally specialised clusters: Taiwan (TSMC, ASE) for fabrication; Silicon Valley (NVIDIA, AMD) for design; the Netherlands (ASML) for lithography.

These clusters possess collective capabilities – such as Taiwan's ability to manufacture 92% of sub-10nm chips – that no individual firm controls. The CHIPS Act targets cluster formation precisely because policymakers recognise that firm-level subsidies cannot recreate cluster-level capabilities.

Firm-level illustration: Hollywood. Hollywood is not a market of independent studios but a **cluster** of studios, agencies, guilds, and distributors. Its star-making power, production pipelines, and bargaining structures exist at the cluster level. The 2023 writers' strike shut down the entire cluster, not just firms. Netflix's disruption succeeded partly by initially operating outside this cluster – but it is now being pulled inside it.

Firm-level illustration: Biotech-Pharma clusters. Boston/Cambridge, the Bay Area, and Basel/Zurich are biotech clusters with shared labour markets, venture capital, CROs, and regulatory expertise. Moderna's 63-day COVID vaccine sprint was not a firm miracle – it was a cluster-level capability mobilised through dense local networks. Powell et al. (1996) show that success in biotech is predicted by cluster position, not by internal R&D spending.

Why Clusters Satisfy Meso-Level Criteria

Criterion 1: Emergence. Cluster-level properties arise from node interactions but cannot be reduced to node properties. As Holland (1998) defines, emergence occurs when structured interaction produces properties that do not exist in the components themselves. Network

clusters generate precisely such properties: cohesion, collective capability, resilience, and identity exist only at the cluster level.

Firm-level illustration: Toyota's supplier keiretsu. Toyota's 200+ Tier-1 suppliers do not merely transact with Toyota; they form a dense production cluster characterised by just-in-time synchronisation, shared continuous-improvement routines, and mutual investment in firm-specific assets. No individual supplier possesses "system-wide JIT capability." That capability exists only as a cluster-level emergent property generated by dense communication, personnel exchange, and joint problem-solving.

When the 2011 tsunami disrupted supply chains, Toyota's cluster recovered in roughly six months – twice as fast as rivals – because resilience had emerged at the cluster level through decades of coordinated learning. Toyota's \$275 billion revenue is grounded not in firm-level optimisation but in emergent meso-level production architecture.

Criterion 2: Causal Autonomy. Clusters possess causal powers that individual nodes do not. As Elder-Vass (2010) shows, social structures exert emergent causal force: they generate outcomes no component could produce alone.

Firm-level illustration: OPEC. No oil-producing country can stabilise or shock global oil prices alone. But OPEC as a cluster can. Its coordinated 2.2 million barrel-per-day production cut in 2023 moved global prices in a way no member state could replicate individually. Markets respond to the cluster's announcements, not to isolated national intentions. The 1973 embargo and every modern OPEC quota decision demonstrate that the price of oil is partially produced at the cluster level. OPEC is not a set of producers – it is a causal structure.

Criterion 3: Downward Causation. Clusters not only emerge from nodes; they act back upon them, shaping behaviour, identity, and opportunity. As Hodgson and Knudsen (2010) show, social structures exert downward causation by constraining and enabling micro-level action.

Firm-level illustration: McKinsey. McKinsey operates as a dense professional cluster: shared methods, internal networks, common identity, and career norms. Individual consultants do not merely work there; they are formed by the cluster. They internalise the "McKinsey way," adopt its analytical frames, and align career behaviour with its "up-or-out" structure.

Even after departure, McKinsey alumni continue to act as a coordinated elite network – 30,000+ former consultants embedded across corporations and governments. This persistence shows that cluster identity and norms outlive formal membership. McKinsey's power is not located in individuals but in the meso-level structure that produces and disciplines them.

Criterion 4: Intermediate Scale. Clusters operate at the scale where CAS dynamics actually unfold – too large for individual optimisation and too structured for aggregate statistics. As Dopfer, Foster, and Potts (2004) argue, the meso-level is the domain of rules, routines, and coordination architectures.

Firm-level illustration: Venture capital syndication. VC investment does not occur at the level of individual firms or at the level of the "market." It occurs in syndication clusters –

groups of VCs that repeatedly co-invest. Hochberg et al. (2007) show that fund performance is predicted by cluster position – who you syndicate with and how central you are – not by firm size, experience, or market timing.

Andreessen Horowitz's rise to \$35 billion AUM reflected deliberate cluster engineering: building syndication ties with Sequoia, Benchmark, and top co-investors.

Firm-level illustration: Paris haute couture. Haute couture is not a collection of fashion houses nor a luxury market – it is a Paris-centred production cluster comprising 15 certified maisons, 300+ ateliers, fashion media, and retail nodes. The cluster generates collective brand premium, craft specialisation, and market-making power that no single house could create.

When LVMH acquires a couture house, it acquires cluster position – access to ateliers, press networks, and Parisian fashion capital. The €30 billion French luxury industry exists at the cluster level, not at the level of isolated firms.

IX. The Firm as Network Cluster

The cluster-as-meso-level framework resolves the four classical puzzles of the theory of the firm: why firms exist, where boundaries lie, what capabilities are, and how identity persists. All four dissolve once the firm is recognised as a network cluster rather than a legal container or contractual nexus.

1. Why Firms Exist

The Coasian question – why firms exist instead of markets – receives a network-theoretic answer: firms exist because dense internal networks enable forms of coordination that sparse market networks cannot generate.

As Kogut and Zander (1992) showed, firms are not simply contractual devices but social communities in which knowledge is created, shared, and stabilised through dense interaction. Markets coordinate through prices; firms coordinate through clusters of interaction.

Firm-level illustration: Pixar. Pixar's films are not the product of individual genius but of a dense creative cluster of 500+ artists, engineers, and storytellers interacting over 4–5 year cycles through daily reviews, Braintrust meetings, shared digital platforms, and physical co-location. These interactions generate emergent creative coherence that no market of freelance animators could reproduce.

Pixar exists as a firm because its cluster topology produces creative coordination that markets cannot. The \$15+ billion box-office value is a cluster-level outcome, not a sum of individual talents.

2. Where Firm Boundaries Lie

Firm boundaries are not legal lines; they are gradients of connection density. The firm "ends" not where ownership stops, but where interaction intensity drops.

Firm-level illustration: Apple. Foxconn employs 1.2 million workers producing Apple devices. Legally external, they are more tightly connected to Apple's production systems than many internal employees are to one another. Apple's boundary is therefore a continuum: Apple Park engineers → very dense; Retail staff → dense; Foxconn production → moderately dense; App developers → weaker; Consumers → sparse.

Apple is not a box – it is a concentration field of interactions. Its \$383 billion revenue depends on managing this connection-density gradient, not on owning assets.

3. What Capability Is

Capabilities are not individual skills or asset stocks. They are network-enabled collective competences that exist only at the cluster level.

Firm-level illustration: Goldman Sachs. Goldman's risk-management capability is not something any trader possesses. It is a cluster property: real-time data flows across desks, shared risk norms, coordinated hedging capacity, and centralised exposure awareness.

In 2008, Goldman survived while Lehman collapsed because Goldman could act as a coordinated cluster – adjusting positions across the entire firm. Poaching Goldman employees does not transfer this capability because it was never in the individuals. It resided in the network topology. Goldman's \$46 billion revenue reflects this meso-level competence.

4. How Firm Identity Persists

Firm identity persists not through people or assets but through pattern continuity. The cluster structure – roles, routines, norms, and connection patterns – is what endures.

Firm-level illustration: Johnson & Johnson. J&J has 250+ operating companies and 130,000+ employees, yet maintains a coherent identity across generations. The continuity lies in its organisational pattern: decentralised autonomy, the Credo, crisis-response protocols, and internal mobility networks. When J&J spun off Kenvue in 2023, the question was which cluster pattern would be inherited. Identity was transferred through structure, not personnel.

Firm-level illustration: New York Stock Exchange. NYSE has persisted since 1792 despite total turnover of people, buildings, and technology. What persists is the trading cluster pattern: rules, specialist networks, market-making routines, and institutional norms. Mergers preserved these patterns because they were the value. NYSE's \$28 trillion annual trading volume is the product of a two-century-old cluster structure.

What This Establishes

The firm is not a legal shell, a bundle of contracts, a set of assets, or a collection of individuals. It is a dense, self-reproducing network cluster whose coordination explains why it exists, whose density gradients explain its boundaries, whose interaction patterns constitute its capabilities, and whose persistent structure gives it identity.

The theory of the firm is therefore a theory of meso-level network architecture.

The firm is not *in* a network. The firm *is* a network cluster.

5.7.4 Local Interactions, Global Emergence, and Synthesis

X. Local Interactions → Global Emergence

The cluster framework explains how simple local interactions generate complex global order. Emergence is not mystical, nor merely statistical – it is the predictable consequence of specific structural conditions. When agents interact within dense, partially connected network clusters, collective patterns arise that no agent intends yet all participate in.

As Kauffman (1993) demonstrated, complex adaptive systems exhibit "order for free": once connectivity, non-linearity, and selection are present, large-scale structure appears without central design. What matters is not agent intelligence but network architecture. Clusters provide exactly the structural regime in which emergence becomes reliable rather than accidental.

1. Viral Cascades: Emergence by Design

Firm-level illustration: TikTok. When a TikTok video reaches a billion views, nothing magical has happened. A predictable cascade has unfolded: local rule (users watch, like, or share), cluster structure (follower networks and algorithmic cohorts), and feedback (engagement raises visibility, increasing engagement).

The platform is engineered to detect local resonance and amplify it through clustered networks. Videos that fit certain structural niches reliably trigger cascades. TikTok's \$16+ billion in advertising revenue comes from monetising this engineered emergence.

Virality is not luck – it is cluster-mediated non-linearity.

2. Financial Crises: Destructive Emergence

Firm-level illustration, p. 2008 Global Financial Crisis. The 2008 crisis was not caused by any single institution. It emerged from cluster topology.

Local events – mortgage defaults – propagated through securitisation chains, derivative linkages (CDOs, CDSs), and highly clustered interbank exposures. AIG alone insured \$440

billion in structured products across the network. Short path lengths meant any failure reached all major institutions in two or three steps. Non-linear CDO structures meant tiny increases in default rates destroyed entire tranches.

The \$22 trillion global wealth loss was not a macro shock – it was a meso-level network cascade. Microprudential regulation (capital ratios) and macro indicators (leverage) missed the cluster structure where fragility actually resided. Battiston et al. (2016) show that systemic risk was encoded in interbank topology, not balance sheets.

3. Platform Ecosystems: Constructed Emergence

Firm-level illustration: Apple's App Store. Apple did not plan the App Store's \$1.1 trillion ecosystem. It engineered the conditions under which it would emerge: local incentives (developers seek users, users seek apps), structural mediation (Apple curates, verifies, and ranks), and feedback loops (downloads → visibility → more downloads).

From these simple rules, a global innovation ecology arose: millions of apps, billions of transactions, continuous adaptation. No central planner coordinates this. It emerges because Apple controls the cluster topology – interfaces, rankings, payments, and trust mechanisms.

Apple captures emergence rent: value generated not by any app, but by the cluster architecture itself.

4. Knowledge Without Management

Firm-level illustration: Wikipedia. Wikipedia demonstrates emergence in its purest form. Local rule: anyone can edit. Structure: versioning, talk pages, moderation networks. Feedback: good edits persist; bad edits revert.

From this, a global knowledge system emerges, p. 60+ million articles, quality convergence, topic coverage correlated with attention. No editor controls Wikipedia; the Wikimedia Foundation merely maintains the conditions for emergence.

Greenstein and Zhu (2012) show that article quality rises with editor clustering. Knowledge is not stored in individuals – it exists in the editing cluster. Wikipedia's 1.7 billion users access a meso-level epistemic structure generated by local interactions.

What This Establishes

Emergence is not mysterious. It is structural: Local interactions + clustered networks + non-linear feedback → global order.

Firms, platforms, markets, and knowledge systems all derive their power not from agents or aggregates but from meso-level cluster dynamics.

This completes the architecture:

- Relations → networks
- Networks → clusters
- Clusters → emergence

Economic reality does not live in micro or macro. It lives in between – in the clusters where interactions condense into structure.

XI. Counter-Arguments and Responses

Multi-Level Architecture Counter-Arguments

Counter-Argument 1: "Hierarchy is just an analytical convenience, not ontological reality."

Response. This objection confuses representation with structure. Hierarchical analysis is useful because complex systems are hierarchically organised – not the other way around. Simon (1962) showed that hierarchic systems evolve because they are vastly more stable than non-hierarchic ones. Anderson's (1972) principle of level autonomy demonstrates that each level of organisation exhibits properties absent at lower levels, requiring level-specific analysis. If hierarchy were merely a modelling choice, arbitrary decompositions would perform equally well. Empirically, they do not. Firms decompose naturally into divisions, teams, and individuals – not into random collections of agents. As Kauffman (1993) showed, hierarchical organisation emerges spontaneously in systems operating at the edge of chaos. Hierarchy is therefore discovered, not imposed.

Counter-Argument 2: "Multi-level causation is just aggregated micro-causation."

Response. This reinstates the reductionism already rejected. Downward causation is not reducible to aggregation because it operates through emergent structures that exist only at higher levels. Market prices, institutional rules, technological standards, and organisational routines all shape individual behaviour while being generated by it – circular causation irreducible to either pole. Hodgson and Knudsen (2010) show that institutions "carry" causal power by constraining and enabling action without determining it. Elder-Vass (2010) demonstrates that higher-level structures possess genuine causal powers not decomposable into lower-level interactions. The 2008 crisis made this undeniable: regulatory regimes, rating systems, and securitisation architectures shaped individual risk-taking in ways no aggregation model could reproduce.

Counter-Argument 3: "Near-decomposability justifies ignoring cross-level effects."

Response. Simon argued the opposite. Near-decomposability allows short-run, level-specific analysis – but long-run dynamics are driven by cross-level coupling. Ignoring these interactions

produces precisely the failures observed in real economies. The 2021–23 semiconductor shortage showed how firm-level disruptions propagated to industry-level shutdowns and macroeconomic inflation. Nelson and Winter (1982) demonstrated that routines evolve under industry-level selection environments; Lane and Maxfield (2005) showed that foresight in CAS requires tracking cross-level feedback. Near-decomposability is not an excuse for reductionism – it is the reason reductionism fails when coupling matters.

Counter-Argument 4: "The meso-level is too fuzzy to be useful."

Response. The meso-level is not vague; it is precisely the level of organised structure – routines, networks, institutions, standards, and capabilities. Dopfer, Foster, and Potts (2004) define it rigorously as the level of "rules" that coordinate micro behaviour and generate macro outcomes. These structures are empirically observable and measurable: routine persistence (Nelson & Winter), network topology (Watts & Strogatz), institutional forms (North), and platform architectures. What appears "fuzzy" is simply invisible to models that lack the conceptual apparatus to see it.

Network Topology Counter-Arguments

Counter-Argument 5: "Network analysis is already common in economics."

Response. This objection confuses network-as-model with network-as-ontology. Mainstream economics may deploy network representations instrumentally, yet it typically retains atomistic ontology: agents possess fixed preferences and attributes; networks appear as external constraints on optimisation; network effects are treated as "frictions" or "market failures." Our claim is stronger and ontological: networks are constitutive – entity properties emerge from network position, rather than merely being modified by it.

Jackson (2008, pp. 18–45) notes that even sophisticated network economics frequently treats network structure as exogenous to the agent. Goyal (2007, pp. 1–28) similarly observes that much of the literature preserves methodological individualism. The shift required is categorical: from "firms in networks" to "firms as network positions." Granovetter's (1985, pp. 481–510) embeddedness distinction captures this: weak embeddedness treats networks as context; strong embeddedness treats networks as constitutive of economic action and identity.

Counter-Argument 6: "Network position can be reduced to agent choices."

Response. This objection reinstates the reductionist move the dissertation rejects. Yes, networks are formed through micro-level tie formation – but once formed, network structures exhibit causal autonomy: they shape behaviour and outcomes in ways not reducible to the choices that generated them.

Padgett and Powell (2012, pp. 1–29) show that network structures, once in place, reorganise opportunity sets, information access, and feasible strategies independently of any single actor's intent. Uzzi (1997, pp. 35–67) demonstrates how embedded ties produce benefits but also dependence. The 2008 financial crisis made this autonomy empirically unavoidable: institutions chose counterparties, but the resulting topology produced systemic risk that no

institution intended and no institution could neutralise unilaterally. Battiston et al. (2016, pp. 10–18) show that financial networks generate systemic properties irreducible to component-level risk choices.

Counter-Argument 7: "Topology is just one factor among many."

Response. This objection treats topology as a variable rather than as substrate. Network topology is not one determinant added alongside "resources," "capabilities," or "strategy." It is the medium through which these determinants operate. Resources do not create effects in isolation; strategies require pathways; capabilities manifest only through relational channels.

Burt (2004, pp. 349–399) demonstrates empirically that network position explains performance variance that "attributes" cannot, and that attribute effects often shrink substantially once structural position is controlled. Gulati, Nohria, and Zaheer (2000, pp. 203–215) show alliance performance depends heavily on network position rather than on alliance-specific resource bundles alone. Ahuja (2000, pp. 425–455) finds that both direct ties and indirect ties affect innovation – indicating that outcomes are produced through structural architecture, not by isolated attributes. Topology is not one factor – topology is the channel through which factors become effective.

Counter-Argument 8: "Network analysis ignores agent heterogeneity."

Response. This objection misunderstands what network analysis explains. Network approaches do not deny heterogeneity; they explain how heterogeneity emerges from position. Agents become different because they are differently embedded: central actors encounter different information flows than peripheral actors; bridging actors develop different interpretive capacities than densely embedded actors; structural holes generate different opportunities than closed clusters. Heterogeneity is therefore not assumed away – it is endogenised by topology.

Borgatti and Foster (2003, pp. 991–1013) survey extensive evidence that structural position produces systematic differences in cognition, capability, and performance. Reagans and McEvily (2003, pp. 240–267) show that knowledge transfer depends on network position, generating capability differentials as positional outcomes. Tsai (2001, pp. 996–1004) demonstrates that within-firm business unit performance varies with internal network position, with positional effects remaining significant beyond unit attributes.

XII. Conclusion

This section has established multi-level nested architecture, network topology, and network clusters as the structural implications of process ontology that complete the dissertation's reconceptualisation of the firm. The argument proceeded through ten movements:

1. Hierarchy as the architecture of complexity (Simon)
2. Near-decomposability as the condition of stability and adaptability
3. Nested CAS across scales with multi-level causation

4. The missing meso-level in mainstream economics
5. Network topology as the structural realisation of relational ontology
6. Key network properties and mechanisms (scale-free, small-world, weak ties, structural holes)
7. The firm as topological entity defined by position
8. Network clusters as the canonical meso-level structure
9. The firm as network cluster resolving existence, boundary, capability, and identity puzzles
10. Local interactions → global emergence through cluster dynamics

Transformation Table: Multi-Level Architecture

Table 128: From Flat Ontology to Nested CAS Architecture: Nine-Dimension Transformation Toward a Hierarchical, Multi-Level Ontology of the Firm

Dimension	FROM (Flat Ontology)	TO (Nested CAS Ontology)
Structure	Micro → Macro	Hierarchical, nested
Decomposition	Full reduction	Near-decomposable
Levels	Analytical convenience	Ontological reality
Causation	Upward only	Emergence + constraint
Meso-level	Invisible	Causally central
Firms	Atomic units	Meso-level systems
Stability	Equilibrium	Hierarchical resilience
Interaction	Within-level	Cross-level coupling
Method	Representative agent	Multi-level CAS

Transformation Table: Network Topology

Table 129: From Atomistic to Topological: Eight-Dimension Transformation Reconceiving the Firm as a Network-Positioned Entity

Dimension	FROM (Atomistic)	TO (Topological)
Entities	Self-contained units	Network positions
Properties	Intrinsic attributes	Positional effects
Relations	External constraints	Constitutive connections
Boundaries	Legal/physical walls	Connection-density gradients
Value	Resource accumulation	Topological position
Strategy	Attribute improvement	Network engineering

Dimension	FROM (Atomistic)	TO (Topological)
Causation	Attribute → Outcome	Position → Outcome
Firm	Bundle of resources	Topological entity
Competition	Attribute comparison	Position comparison

Transformation Table: Meso-Level Clusters

Table 130: Ontological Significance of Meso-Level Clusters: Four Criteria Establishing Their Causal Autonomy and Irreducibility

Criterion	What Clusters Provide
Emergence	Collective capabilities irreducible to nodes
Causal autonomy	Powers no member can exercise alone
Downward causation	Norms, identities, and constraints shaping nodes
Intermediate scale	The actual arena of coordination and adaptation

Forward Integration

The implications are decisive. Mainstream economics' micro–macro dichotomy collapses a hierarchical universe into a flat plane. By leaping from individual agents to aggregates via the representative agent, it erases the level where economic life actually happens: firms, routines, networks, institutions, platforms, and organised knowledge.

The CAS perspective restores this missing middle. Economies are not collections of atoms nor homogeneous wholes, but nested architectures of interacting systems. Causation flows upward through emergence and downward through constraint. Firms exist neither at the micro nor the macro level but at the meso-level, where routines stabilise, capabilities accumulate, and network positions generate power.

The firm is not a bundle of resources operating within a network. The firm is a network cluster. Its boundaries are gradients. Its capabilities are positional. Its identity is pattern continuity. Its strategy is topology design.

Section 5.8 (Interconnections) will now develop the connections between the ontological framework established in Sections 5.1–5.7 and its epistemological and methodological implications.

The Aha Moment

The economy is not a collection of atoms to be aggregated, nor a holistic totality to be contemplated. It is a nested architecture of complex adaptive systems – each level emergent, each level constraining and enabling others – with the meso-level where economic life actually happens rendered invisible by mainstream theory.

The firm does not *have* a network position; it *is* a network position. Change the topology and you change the firm – not by moving assets or personnel, but by reshaping the relational field through which the firm exists. And economic reality does not live in micro or macro. It lives in between – in the clusters where interactions condense into structure.

5.8 Interconnections – From Ontology to Epistemology, Methodology, and Structure

Introduction

This section establishes the three-bridge sequence through which Chapter 5's ontological reconstruction generates systematic consequences for epistemology, methodology, and structure. The core principle is simple but decisive: coherence requires alignment. Once we commit to emergence, process, and relational constitution, we cannot coherently retain mechanistic epistemology, reductionist methodology, or closed-system structure.

The three bridges developed here complete Chapter 5's transformation:

- Section 5.8.1 establishes that process ontology implies fundamental uncertainty – the future is ontologically open, not merely epistemically obscure.
- Section 5.8.2 establishes that emergence requires synthetic methods – analytical decomposition destroys what it purports to explain.
- Section 5.8.3 establishes that CAS ontology requires open-system structure – closure becomes logically incompatible with relational, processual reality.

Together these bridges show that ontology, epistemology, methodology, and structure form a single coherent architecture. One cannot accept CAS ontology while retaining Laplacean epistemology, reductionist method, or closed-system models. The transformation is systemic.

5.8.1 Ontology → Epistemology: Fundamental Uncertainty and Epistemic Complementarity

Introduction

This subsection establishes the first major interconnection: epistemology must follow ontology. What can be known, and how it can be known, depends on what kind of world exists.

The ontological commitments established in Sections 5.1–5.4 imply a decisive epistemic consequence. If (i) emergence produces genuinely novel properties, (ii) reality is processual becoming rather than static being, (iii) entities are constituted by relations rather than intrinsic attributes, and (iv) identities are path-dependent trajectories rather than fixed substances, then the future is not a pre-written object waiting to be revealed. It is being produced. Under these conditions, uncertainty is not primarily an information problem; it is an ontological condition.

The claim is not that economic knowledge is impossible. The claim is that appropriate knowledge in a processual world is necessarily fallible, revisable, and pattern-based, because the environment it targets is open, evolving, and partly self-creating. Deterministic and fully calculative epistemologies are therefore not merely empirically fragile; they are logically misaligned with a world of emergence, process, relational constitution, and path dependence.

This section proceeds through twelve parts: (I) Laplacean determinism as mechanistic epistemology; (II) why process ontology destroys determinism; (III–V) Knightian uncertainty, Keynesian fundamental uncertainty, and non-ergodicity; (VI) distributed causality and pattern epistemology; (VII–IX) analysis, synthesis, and their integration through systems thinking; (X) the meso-level as epistemically privileged; (XI) multiple types of understanding required; (XII) counter-arguments addressed.

I. Laplacean Determinism: Mechanistic Ontology's Epistemological Implication

The epistemological expression of mechanistic ontology was stated with unmatched clarity by Pierre-Simon Laplace. In his *Philosophical Essay on Probabilities*, he imagined an intelligence that, knowing all forces and positions at a given instant, could compute the entire future and past of the universe. For such an intellect, "nothing would be uncertain."

Laplace's demon is not simply a prediction fantasy. It is the epistemic corollary of a closed, deterministic world: if reality is a system of fixed laws operating on stable entities, then the present fully determines the future. On this view, novelty is not genuine; it is merely undiscovered. Uncertainty is therefore only temporary ignorance.

Mechanistic ontology typically carries the following epistemic commitments: the present determines the future (prediction is possible in principle); wholes equal sums of parts (knowing components suffices); causation is linear (effects are traceable to isolable causes); and time functions as a parameter (measuring change but not creating new possibilities). In this frame, uncertainty is reducible by data, measurement, and computation. It is a residual term, not a feature of reality.

Firm-level translation. Strategic planning under this worldview becomes probabilistic optimisation: forecast states, attach probabilities, and choose strategies that maximise expected value. This logic underlies mainstream finance (Black–Scholes, Value-at-Risk, Modern Portfolio Theory) and consulting-based corporate strategy. The 2008 crisis – where events labelled "25-sigma" occurred repeatedly – revealed not bad models but a false ontology. The probability distributions did not fail; the world they assumed never existed.

II. Why Process Ontology Destroys Deterministic Epistemology

The ontological commitments established in Section 5.2 invalidate Laplacean certainty through four interlocking mechanisms.

Mechanism 1: Emergence creates properties that did not previously exist. If emergence is real, then higher-level properties are not contained in the lower-level description in a way that makes them deducible in advance. As Anderson (1972, p. 393) argued, "more is different":

higher levels require their own descriptions and regularities. If novel structures and properties come into existence through interaction and organisation, then they cannot be predicted as if they were already present in latent form.

Firm-level translation. Amazon's organisational culture – "customer obsession," "bias for action," "frugality" – is not an aggregation of individual employee traits but an emergent organisational property produced by interaction, selection, and reinforcement among 1.5 million employees. No database of individual psychometrics could predict that culture, because it exists only at the organisational level. What does not yet exist cannot be known in advance.

Mechanism 2: Process time is productive – it creates, it does not merely reveal. Process ontology treats time as constitutive. The future is not a state already implied by the present; it is generated through unfolding interaction. This defeats the premise that the future is "there" to be calculated.

Firm-level translation. Nokia circa 2000 – global mobile-phone leader with 30% market share – did not "contain" its later trajectory as a computable implication. Its later state was produced through a sequence of interpretations, commitments, competitive responses, and organisational dynamics unfolding through time (Tripsas & Gavetti, 2000). The relevant causal object is not a snapshot but a becoming.

Mechanism 3: Path dependence makes history causally irreducible. In path-dependent systems, the present is not a sufficient statistic for the past. The system's identity is a trajectory; the future depends on how the present was reached. The implication is not merely that prediction is hard, but that prediction from a present-state description cannot capture the causal content carried by history.

Firm-level translation. Intel's 1985 decision to abandon memory chips and commit to microprocessors created a technological and organisational path that locked the firm into PC-centred computing. That path later became a liability in mobile and AI-oriented architectures. Intel's 2020 position cannot be inferred from its 2020 balance sheet – it is the product of 35 years of accumulated commitments. History does not wash out; it accumulates.

Mechanism 4: Relational causality is distributed, recursive, and contextual. Section 5.3 established that entities are constituted by relations. Causality in such systems is not linear but networked, recursive, and contingent on configuration. The same action can have different effects in different relational contexts. Causes do not travel along linear chains but through evolving networks of feedback.

Firm-level translation. Silicon Valley's innovative power cannot be reduced to capital, talent, or universities taken separately. It emerged from their configuration: dense labour mobility, venture capital networks, research institutions, and cultural norms of experimentation. The same ingredients arranged differently (as in Route 128) produced radically different outcomes. Causation lay in topology, not in parts.

Synthesis. Laplacean certainty collapses not because the economy is complicated but because it is ontologically open. Emergence creates properties that did not exist; process generates becoming rather than revelation; path-dependence embeds history into the present; and

relational causality distributes influence across evolving networks. In such a world, the future is not computable even in principle. Uncertainty is not ignorance – it is the logical consequence of processual reality.

III. Knightian Uncertainty as Ontological Distinction

Frank Knight's *Risk, Uncertainty, and Profit* (1921) formalised a distinction that maps precisely onto the ontological divide between mechanistic and process systems. Knight showed that economics conflates two radically different kinds of unknowability.

Risk applies when outcomes follow stable probability distributions that can be estimated from historical frequencies. Insurance is possible; calculation governs decision. Fundamental uncertainty applies when the distribution itself is unknown or unstable – when outcomes cannot be enumerated in advance. No probability can be assigned because the outcome space does not yet exist as a determinate object.

This is not merely a psychological or informational distinction – it is an ontological one. Risk applies to mechanistic systems: stable processes with fixed possibility spaces and reproducible frequencies. Fundamental uncertainty applies to processual systems: economies, technologies, firms, and markets that create new possibilities through their own evolution. In such systems, the outcome space itself is endogenous. You cannot assign probabilities to states that do not yet exist.

Knight (1921, p. 268) understood this precisely: entrepreneurial decisions concern unique, non-repetitive situations in which no statistical distribution can be constructed – not because data are missing, but because the distribution has not yet been generated by history.

Firm-level translations. When Sequoia invested \$60 million in WhatsApp (2011) for approximately 15%, the relevant futures were not enumerable: mobile messaging was not yet an industry; freemium business models were unproven; Facebook's acquisition strategy was unknowable; network effects would determine winner-take-all dynamics. No probability distribution existed. The investment was not risk-taking; it was uncertainty-bearing judgment. Sequoia evaluated patterns of user behaviour, platform adoption, and communication economics, not expected values. The \$3 billion return was payment for creating a future that did not yet exist.

Similarly, when Apple began iPhone development (2004–2007), there was no smartphone market in the modern sense. Touchscreen interfaces had failed repeatedly; mobile computing paradigms were undefined; app ecosystems did not exist. No distribution of possible outcomes could be specified, because the market itself would be shaped by Apple's actions. The iPhone did not win in a known lottery – it created the lottery.

AWS emerged from excess internal computing capacity – but the cloud market did not yet exist. No CIO could be modelled; no adoption curve existed; no industry benchmarks were available. AWS became an \$80 billion business not by selecting among known options but by bringing a new category into being.

Economic implications. Once uncertainty is recognised as ontological, the foundations of economic theory change. Profit becomes return to bearing irreducible uncertainty, not diversifiable risk. Entrepreneurship becomes creation of new possibility spaces, not optimisation within known ones. Strategy becomes path-creation under uncertainty, not choice under probabilities. Planning becomes adaptive exploration, not forecast-based optimisation. Expected-utility maximisation, probabilistic forecasting, and risk-adjusted optimisation are mathematically elegant – but they describe a world that does not exist.

IV. Keynesian Fundamental Uncertainty: Ontological Openness

Keynes extended Knight's distinction by grounding uncertainty explicitly in the ontological status of the future. In his 1937 reply to critics of *The General Theory*, Keynes drew a categorical line between calculable risk and genuine uncertainty:

"By 'uncertain' knowledge... I do not mean merely to distinguish what is known for certain from what is only probable. The game of roulette is not subject, in this sense, to uncertainty... The sense in which I am using the term is that in which the prospect of a European war is uncertain, or the price of copper and the rate of interest twenty years hence... About these matters there is no scientific basis on which to form any calculable probability whatever. We simply do not know." (Keynes, 1937, pp. 213–214)

This is not a statement about missing information. It is a statement about non-existence. Roulette outcomes exist in a fixed probability space. The future price of copper does not, because the economic world that will generate that price has not yet been created.

Within the framework developed in Sections 5.1–5.4, Keynes' claim takes on its full force. If reality is processual becoming (5.2), and entities are stabilised processes (5.4), then: "We simply do not know" is not epistemic failure – it is metaphysical truth.

Firm-level translations. CalPERS (\$440 billion AUM) must fund pensions in 2060. What probabilities can it assign to stock returns over four decades? That requires forecasting technologies not yet invented, geopolitical regimes not yet formed, climate impacts not yet realised, and social transformations not yet known. There is no distribution to estimate because the future economy does not yet exist. CalPERS uses a 7% assumed return not because 7% is probable but because some convention is operationally necessary. The number is not a forecast – it is a stabilising fiction enabling action under uncertainty.

Tesla's valuation swings – from \$34 billion (2019) to \$1.2 trillion (2021) to \$385 billion (2023) – reflect not new probabilistic information but shifting conventions about an unknowable technological future. No probability distribution exists over EV adoption speed, battery breakthroughs, charging infrastructure, regulatory regimes, or competitor responses. Valuation reflects animal spirits about a future that will be created, not calculated.

Animal spirits as rational response. This is why Keynes insisted that investment depends on **animal spirits**. When probabilities do not exist, action must be grounded in confidence, narrative, and convention, not optimisation. Business cycles are not responses to new data; they are shifts in collective belief about an unknowable future. When confidence collapses,

investment collapses – not because expected values changed, but because the psychological scaffolding that supports action under uncertainty disintegrated.

Liquidity preference. Under fundamental uncertainty, money holds optionality – the ability to wait for novelty to appear before committing. This is not irrational hoarding; it is rational exposure management to ontological openness. Apple (\$165 billion cash), Microsoft (\$104 billion), and Alphabet (\$116 billion) hold cash not because managers are lazy but because the best investments cannot yet be known. Cash preserves the right to act when the future reveals itself.

V. Non-Ergodicity: History Matters Irreversibly

If reality is processual becoming (Section 5.2), then economic time must be non-ergodic. This deepens fundamental uncertainty beyond Knight and Keynes by showing that the probability space itself evolves.

The ergodic assumption – embedded in most economic uses of probability – requires that time averages equal ensemble averages. If a system is ergodic, observing one long trajectory reveals the same statistical regularities as observing many systems briefly. History washes out. But if a system is non-ergodic, history accumulates. Different paths lead to different futures, and no single probability distribution governs outcomes.

As North (1999, p. 12) put it: "The world we live in is a non-ergodic world." In such a world, learning does not converge to a stable distribution because the structure of the world itself changes as learning occurs.

Why process ontology implies non-ergodicity. From Section 5.2, identities are not static but stabilised trajectories. When systems evolve through cumulative path-dependent processes, the future depends not only on current state but on how that state was reached. Innovations create new states; institutions reshape incentives; strategies alter landscapes. Markets do not converge toward a fixed distribution. They wander through evolving spaces of possibility.

Firm-level translations. Nokia's fate was not determined by its 2000 market share but by its trajectory – its internal power struggles, OS decisions, developer relations, and smartphone missteps. A different sequence would have created a different future. Nokia's time-path $\neq$ industry distribution. That is non-ergodicity.

QWERTY is the canonical non-ergodic lock-in. A historically contingent choice became dominant through increasing returns, making superior alternatives economically inaccessible. Waiting longer does not move the system toward the "best" outcome. History selected it.

Silicon Valley did not emerge from superior initial conditions but from a specific path: Stanford policies → Fairchild → venture capital → dense networks → self-reinforcing agglomeration. Route 128 had similar resources but took a different path and produced a different world.

Pfizer's mRNA success did not reflect a draw from a common probability distribution. It reflected a trajectory: BioNTech partnership → regulatory experience → platform

accumulation → future option space expansion. Merck's different trajectory created a different possibility space. Firms face not the same probabilities – but different futures.

Epistemic consequence. In non-ergodic systems, statistics no longer project the future; history accumulates; feedback reshapes possibility space; and outcomes are not drawn from distributions but built through trajectories. This creates a vicious epistemic circle: to know the future, we need the trajectory. But the trajectory is being created by present actions. This is why probability theory breaks in economic time.

VI. Distributed Causality Requires Pattern Epistemology

Section 5.3 established a relational ontology: entities are constituted through relations rather than possessing intrinsic properties independent of context. If causation operates through relational configurations rather than isolated entities, then knowledge cannot be grounded in linear causal laws but must instead track patterns in evolving structures.

Linear causality assumes single causes produce single effects, causes are isolable from context, and prediction proceeds by identifying stable cause–effect relations. Distributed causality recognises that multiple interacting causes generate emergent effects, causes are context-dependent, and understanding requires comprehending configurations of interacting elements.

Firm-level translations. The contrast between Facebook and MySpace illustrates distributed causality. Standard explanations cite product quality, management, or network effects. Relational analysis shows a configuration: Facebook's college-by-college rollout produced dense local clusters before scaling, while MySpace's open registration created sparse, low-engagement networks. The same "network effect" operated through different topologies, generating different dynamics. Add timing (smartphone diffusion), ecosystem design (developer APIs), and competitive pressure (Google+), and Facebook's dominance emerges from structural configuration, not any single variable.

The 2008 crisis demonstrates distributed causality in extremis. Subprime lending, derivatives, rating agencies, and regulatory failures all contributed, but none explains systemic collapse. As Battiston et al. (2012, pp. 4–9) showed, crisis dynamics were driven by financial network topology – counterparty concentration, short path lengths, and contagion channels. Lehman Brothers (\$639 billion in assets) triggered collapse not because of its size but because of its network centrality. The crisis was produced by relational configuration, not by any isolated cause.

Abductive reasoning as the epistemological response. Given emergence, process, and relational causality, deductive epistemology fails. Deduction presupposes stable laws; CAS generate novelty that violates this condition. What remains is abduction – inference to plausible explanations from observed patterns. As Arthur (2021, p. 143) observes, agents in complex economies must "explore, react and constantly change their actions and strategies" because the causal environment is not deductively predictable.

Steve Jobs did not compute probabilities for the iPhone. He recognised a configuration: multi-touch interfaces, mobile computing power, wireless bandwidth, user frustration with keyboards, and Apple's design language. From this pattern he abducted a new product concept

that created a market rather than selecting from an existing one (Isaacson, 2011). This is pattern epistemology, not probabilistic optimisation.

The meso-level as epistemic anchor. Where can knowledge stabilise in a world of flux? Section 5.5 established the meso-level as ontologically real; Section 5.7 identified network clusters as its canonical structure. These provide the epistemic anchor. Micro-events are too volatile. Macro-aggregates erase causal texture. But meso-level structures – clusters, routines, ecosystems, institutional regimes – persist long enough to be observed while retaining explanatory power.

In pharmaceuticals, individual molecules are too uncertain (95% failure rate) and global healthcare aggregates too abstract. But therapeutic area dynamics, regulatory pathways, and alliance structures form meso-patterns that guide strategy. This is why analysts and firms study industry structure rather than molecules or GDP.

Conclusion

Because causation in economic systems is distributed, relational, and path-dependent, knowledge cannot be deductive or variable-based. It must be abductive and pattern-based, operating at the meso-level where processes stabilise into analysable structures.

VII. Analysis Produces Knowledge – But Destroys Understanding

If reality is emergent, relational, and processual, then knowledge cannot be purely decompositional. Understanding Complex Adaptive Systems therefore requires epistemological complementarity: the integration of analysis (precise knowledge of components) and synthesis (holistic comprehension of emergent configurations), coordinated through systems thinking rather than pursued independently.

What analysis achieves. Analysis (from *ana-lyein*, to loosen or dissolve) proceeds by decomposing a whole into components, studying the components independently, and recombining them. Where systems are mechanistic – where wholes equal sums – this strategy is extraordinarily powerful. It delivers precision (quantified measurement), replicability (standardised procedures), cumulation (knowledge builds modularly), and falsifiability (clear error criteria). This is why analysis dominates physics, chemistry, and engineering. When interactions do not generate new system-level properties, decomposition loses nothing essential.

Firm-level translation. Boeing's 787 Dreamliner can be decomposed into engines, avionics, materials, and software; optimising components improves the whole because aerospace systems are close to decomposable. Analysis works here because the ontology is near-mechanistic.

What analysis destroys. In CAS, however, the defining properties are emergent, relational, and historical. They are not located in the parts but in the configuration of interactions. Decomposition therefore destroys what it seeks to explain. This is not a limitation of current techniques. It is a logical consequence of emergence. If wholes have properties not possessed

by parts, and those properties arise from interaction patterns, then decomposing the system eliminates the causal structure that generates the phenomenon.

Firm-level translation. Toyota's lean production capability is not located in workers, routines, machines, or incentives. It emerges from their pattern of coordination across a supplier network. One can measure skills, task sequences, defect rates, and inventories – but none of these captures *lean* as a system property. Decomposition yields component knowledge while destroying the capability itself (Dyer & Nobeoka, 2000).

Market-level translation. Representative-agent equilibrium models decompose markets into identical optimisers. This yields exact prices and quantities – but eliminates contagion, network amplification, and systemic fragility. The 2008 crisis emerged from interbank network topology; representative-agent analysis removed the topology by construction, guaranteeing epistemic blindness.

Organisational translation. McKinsey's analytic frameworks generate metrics but routinely fail to grasp culture, power, and capability dynamics – the emergent properties that determine whether strategies work. "They understood the numbers, not the system."

The epistemological consequence. Analysis produces *epistēmē* – formal, precise knowledge. But CAS require *synesis* – integrative understanding of interacting wholes. For CAS, knowing the parts does not mean understanding the system. By cutting systems into pieces, reductionist epistemology erases coordination, feedback, path-dependence, network position, and institutional embedding. Intervention based on such knowledge reliably produces unintended consequences.

VIII. Synthesis Preserves Understanding – But Lacks Precision

What synthesis achieves. Synthesis (from *syn-tithemi*, to put together) proceeds in the opposite direction to analysis: it begins from the whole and understands parts through their roles within configurations. To understand a firm is therefore to locate it within the larger relational, institutional, technological, and historical systems that constitute it.

Synthesis delivers what analysis cannot: holistic pattern recognition (seeing structures invisible to decomposition); contextual awareness (how embeddedness shapes behaviour); emergence recognition (why wholes exceed sums); and dynamic comprehension (how systems transform over time). Synthesis preserves precisely what analysis destroys: the interaction structures from which system behaviour emerges.

Firm-level translation. Saxenian's comparison of Silicon Valley and Route 128 revealed that innovation outcomes differed not because of components (universities, firms, funding) but because of relational topology. The difference was not in variables but in configuration. Pixar's performance cannot be explained by talent, technology, or incentives taken separately. It emerges from a configuration: the Braintrust, workspace design, and cultural norms that create continuous recombination. The firm's creative output is a property of its **system**, not its parts (Catmull, 2014).

What synthesis lacks. But synthesis sacrifices what analysis provides. It is harder to measure, harder to replicate, harder to generalise, and harder to falsify. This is not a flaw – it reflects the ontology. If systems are historically unique, relationally embedded, and emergent, then precise, context-free general laws cannot exist. The epistemological trade-off is therefore real: Analysis gives precision without understanding. Synthesis gives understanding without precision. Replacing analysis with synthesis would destroy scientific discipline. But replacing synthesis with analysis destroys the phenomenon.

IX. Systems Thinking Fuses Analysis and Synthesis: Epistemic Complementarity

Why neither mode suffices alone. For CAS, neither mode is epistemically adequate. Pure analysis yields exact but blind knowledge. Pure synthesis yields insightful but unfalsifiable understanding. The solution is not compromise but integration. This is what systems thinking is: analytical precision embedded within synthetic framing.

Senge (2006) captures this: systems thinking means seeing patterns of change, not static snapshots. The critical move is redirecting analysis away from isolated parts and toward relational structure.

How systems thinking integrates both modes. Systems thinking proceeds through iterative **epistemic cycles** in which synthesis guides what analysis measures, and analysis tests what synthesis claims: (1) **Synthesis** identifies system, purpose, boundaries, context – defining what is relevant; (2) **Analysis** maps components, links, feedbacks, flows – generating precise knowledge; (3) **Synthesis** interprets how structure generates behaviour – revealing emergence; (4) **Analysis** tests, simulates, intervenes – validating or refining understanding.

Nelson & Winter's evolutionary economics, system dynamics, agent-based models, and network science all instantiate this structure: synthetic framing defines the system; analytic rules specify interactions; simulations reveal emergence; empirical comparison tests the model. This is why complexity economics is not anti-formal – it is post-reductionist.

The key epistemological result. Understanding in CAS is never decompositional and never purely holistic. It is always configurational. Knowledge advances by mapping patterns, not by summing variables. This is why equilibrium economics fails: it assumes the world is decomposable. CAS epistemology begins from the opposite truth: The whole is what causes the parts to behave as they do.

Synthesis guides, analysis verifies. The correct epistemic hierarchy is: Synthesis determines what is to be analysed; Analysis determines whether the synthetic understanding holds. Without synthetic guidance, analysis studies the wrong objects at the wrong level under the wrong independence assumptions. With synthetic guidance, analysis measures feedback loops, maps network effects, preserves the meso-level, and captures interdependence.

Firm-level translation. Synthesis identifies Apple as an ecosystem orchestrator. Analysis then measures developer adoption, switching costs, platform elasticity. Without synthesis, analysis would wrongly focus on product margins or factory efficiency – missing the source of Apple's \$3 trillion valuation. Synthesis identifies TPS as network-level capability. Analysis then

quantifies supplier knowledge flows and defect rates. Dyer & Nobeoka (2000) show this integration explicitly.

X. The Meso-Level as Epistemically Privileged

Emergence lives at the meso-level. Therefore knowledge must live there too.

- **Micro** → destroys emergence through decomposition
- **Macro** → averages emergence away through aggregation
- **Meso** → preserves emergence by maintaining configurational structure

Routines, networks, governance structures, ecosystems – these are the objects of knowledge in CAS. They are stable enough to observe, yet retain the interaction patterns from which system properties emerge.

Firm-level translation. Nokia and WeWork were not micro failures (individual decisions) or macro shocks (industry-wide events). They were meso-level collapses – failures of routines, networks, and institutional coupling. Nokia's collapse was not about individual engineers or the global phone market; it was about the organisational routines, managerial cognition, and supplier relationships that constituted its meso-level architecture. WeWork's implosion reflected the breakdown of governance structures, investor relationships, and scaling routines at the organisational level.

CAS epistemology therefore requires: Meso-preserving methods guided by synthesis and disciplined by analysis.

XI. Multiple Types of Understanding Are Required

Mechanistic vs Organismic Understanding

Bertalanffy (1950; 1968) distinguished two fundamentally different modes of understanding:

Mechanistic understanding asks *how* a system works. It explains phenomena by decomposing them into parts, identifying causal mechanisms, and tracing part-to-whole effects. Its strength is predictive precision.

Organismic understanding asks *why* a system has the form it does. It explains phenomena by situating them in their functional, historical, and adaptive context. Its strength is meaning, purpose, and evolutionary intelligibility.

Table 131: Two Complementary Modes of Explanation: Mechanistic and Organismic Approaches Contrasted by Direction and Explanatory Scope

Mode	Question	Direction	What It Captures
Mechanistic	How does it work?	Parts → whole	Causal chains, operational mechanisms
Organismic	Why is it organised this way?	Whole → parts	Function, adaptation, historical fit

Economics inherited mechanistic understanding from physics – markets as "mechanisms," firms as "production functions," incentives as "forces." But CAS ontology makes organismic understanding unavoidable:

- Why do routines persist?
- Why did QWERTY lock in?
- Why do clusters like Silicon Valley form?
- Why do firms co-evolve with ecosystems?

These are not mechanism questions but historical-functional questions. They require understanding how structures emerge, reproduce, and stabilise in context.

Analysis explains how; synthesis explains why. Neither is dispensable. Mechanisms without meaning are blind; meaning without mechanisms is empty.

Firm-level translation. Spotify cannot be understood only by how its recommendation algorithms function. It must also be understood by *why* its freemium ecosystem – artists, playlists, social sharing, discovery culture – coalesced into a dominant market position. Spotify's scale is both algorithmic and ecological.

Erklären and Verstehen

The same distinction appears in German social theory:

- **Erklären** – causal explanation, law-like regularities, nomothetic science
- **Verstehen** – interpretive understanding, meaning, historically situated action

Economics traditionally privileged erklären: variables, equations, general laws. But CAS phenomena require verstehen as well:

- Strategy depends on managerial interpretation

- Markets depend on expectations of expectations
- Institutions depend on shared meaning

Firm-level translation. Tesla's valuation cannot be understood by cost curves alone. It requires understanding how investors *interpret* Tesla's role in energy transition, technological identity, and future narratives. Causal models and meaning-systems jointly determine outcomes.

CAS epistemology therefore requires both: causal explanation of mechanisms AND interpretive understanding of evolving meaning.

Coherent and Incoherent Epistemological Combinations

Following the coherence criterion established throughout Chapter 5, epistemology must match ontology.

Table 132: Coherent Versus Incoherent Ontology-Epistemology Pairings: Why Methodological Approach Must Match Ontological Commitment

Combination	Coherence	Reason
Mechanistic ontology + Analysis	COHERENT	Closed systems permit decomposition
CAS ontology + Systems thinking (analysis + synthesis)	COHERENT	Methods preserve emergence
CAS ontology + Analysis only	INCOHERENT	Decomposition destroys emergence
CAS ontology + Synthesis only	INCOHERENT	No validation or discipline
Mechanistic understanding + CAS	INCOHERENT	Seeks fixed mechanisms in evolving systems

XII. Counter-Arguments Addressed

Counter-Argument 1: "Bayesian Updating Handles Uncertainty Adequately"

Objection. Bayesian epistemology provides a unified framework for reasoning under uncertainty. Agents form prior distributions and update as evidence arrives. There is no need for a special category of "fundamental" uncertainty.

Response. Bayesian updating presupposes that outcomes are drawn from a stable probability distribution. Learning consists of estimating that distribution more accurately. But process ontology establishes that economic systems are processual and non-ergodic: the distribution itself evolves as the system unfolds. Taleb (2007) shows that Bayesian updating fails when the generating process is non-stationary. Farmer and Geanakoplos (2009) demonstrate formally that in non-ergodic systems there is no "true" probability to converge to. Davidson (1991) clarifies that Keynesian uncertainty is not high-variance estimation but the absence of a

probabilistic object. In process systems, the act of moving through time changes the distribution. Updating cannot converge to what does not exist.

Counter-Argument 2: "Prediction Failure Reflects Computational Limits, Not Ontological Limits"

Objection. Economic forecasting will improve with more data and computing power, just as weather forecasting did.

Response. This conflates complicated systems with complex ones. Complicated systems (rockets, weather at short horizons) follow fixed rules and yield to computation. Complex adaptive systems generate novel structure through feedback, adaptation, and self-organisation. Mitchell (2009) shows that emergence cannot be computed in advance because it is not yet formed. Wolfram (2002) proves that many systems are computationally irreducible: the only way to know what they will do is to let them do it. Hendry and Mizon (2014) show that despite massive computational advances, economic forecasting has not converged toward reliability – because the object being forecast is not stable.

Counter-Argument 3: "Rational Expectations Approximate Reality Well Enough"

Objection. Even if uncertainty exists, rational expectations provide a good approximation: agents' beliefs do not systematically diverge from outcomes.

Response. Rational expectations presuppose a future that exists to be estimated. Process ontology denies this. Frydman and Goldberg (2007) show that RE works only in stationary regimes and collapses during structural change. Kurz (2008) demonstrates that heterogeneous beliefs drive market dynamics. Caballero (2010) shows that the 2008 crisis was a failure of ontological assumptions, not parameter values. RE fails precisely when novelty matters – which is when economics matters most.

Counter-Argument 4: "Knight–Keynes Uncertainty Is Philosophically Obscure"

Objection. The distinction between risk and uncertainty is vague. Either probabilities exist or they do not.

Response. The distinction is ontological, not semantic. Risk applies to stable, repeatable systems. Uncertainty applies to creative, path-dependent systems. Lawson (1988) and Runde (1998) show that Keynesian uncertainty reflects an open causal structure. Shackle (1972) demonstrates that entrepreneurial decision-making relies on imagination and commitment, not calculation.

Counter-Argument 5: "Relational Ontology Is Just Network Analysis"

Objection. Claiming that reality is "fundamentally relational" merely restates what network analysis already captures. This is methodology, not ontology. Standard economics can incorporate networks without ontological transformation.

Response. This objection rests on a category error: it confuses an analytical tool with an ontological commitment. Network analysis can be deployed within substance ontology by treating networks as external relations between pre-existing entities. Relational ontology asserts something categorically stronger: entities are constituted by their relations, not merely connected through them.

Under substance ontology, Apple is an entity that happens to have developers. Under relational ontology, Apple's platform identity **is** the developer network; remove it and the firm ceases to exist in its economic form. Networks are not add-ons to an underlying essence – they are what the entity *is*.

This distinction is empirically decisive. Substance ontology permits studying firms "in themselves" and then adding network effects as modifiers. Relational ontology requires studying firms in context from the start, because context is constitutive, not optional. Granovetter (1985), Burt (1992), and Uzzi (1997) show that network position shapes not just performance but identity – what the firm is able to be depends on where it sits. As Lawson (2019) emphasises, social structures are ontologically relational: relations are not secondary properties of entities but the conditions of their existence.

Network analysis without relational ontology is like thermodynamics without energy: it describes patterns but misses what makes them real.

Counter-Argument 6: "Path-Dependence Is Compatible with Equilibrium"

Objection. Path-dependence merely implies multiple equilibria; history selects which equilibrium prevails. Once selected, standard equilibrium analysis applies.

Response. This objection mischaracterises path-dependence as a one-time selection problem rather than an ongoing generative process. Increasing-returns dynamics (Arthur, 1989) do not merely choose among equilibria; they continuously recreate the trajectory by reinforcing past choices. QWERTY does not "reach" an equilibrium and stop – it accumulates training investments, complementary technologies, and institutional commitments that deepen lock-in over time.

More fundamentally, equilibrium presupposes a fixed possibility space. CAS ontology denies this. If systems continually generate novelty through learning, innovation, and co-evolution, then the space of possible equilibria itself changes. Empirically, economies do not converge to stationary attractors. They exhibit continuous transformation – technological regimes, institutional structures, business models, and market forms evolve endogenously (Arthur, 2021; Dopfer et al., 2004; Metcalfe, 2005). Path-dependence replaces equilibrium with trajectory as the core object of explanation.

Counter-Argument 7: "Irreversibility Is Just Transaction Costs"

Objection. What is called "irreversibility" is simply high switching or transaction costs. Standard economics already handles this through cost-benefit analysis.

Response. This objection confuses economic frictions with ontological constraints. Transaction costs are, in principle, surmountable with sufficient resources. Ontological irreversibility is not. Learning cannot be "un-learned." Organisational routines cannot be erased by spending. Cognitive frames, trust, institutional commitments, and identity structures persist even when incentives change.

Nokia's smartphone failure illustrates this distinction. The problem was not that change was too expensive – Nokia had tens of billions in cash – but that its accumulated capabilities, supplier relations, and managerial cognition made certain paths unavailable. Some transitions are not merely costly; they are structurally foreclosed by history (Sydow et al., 2009; Vergne & Durand, 2010). Irreversibility operates at the level of what is possible, not what is expensive. This is ontological, not transactional.

Counter-Argument 8: "Process Ontology Makes Prediction Impossible"

Objection. If reality is continual becoming with genuine novelty, prediction is impossible. Economics without prediction is not science.

Response. This objection confuses the scope of prediction with the status of science. If reality is genuinely open and novel, then methods that assume full predictability produce systematic error, not scientific rigor. Precision about what cannot be predicted is more scientific than false exactitude.

Moreover, process ontology does not abolish prediction; it redefines it. It enables:

- **Pattern prediction** (Hayek): identifying classes of outcomes
- **Scenario mapping:** exploring possible futures
- **Constraint analysis:** locating bifurcation points and structural limits

These forms of anticipatory knowledge are precisely what open systems permit. They allow intervention, robustness, and navigation even when point prediction is impossible (Dosi & Nelson, 1994; Lane & Maxfield, 2005). In open, evolving systems, the goal of science is not prophecy but understanding of generative mechanisms. That is deeper, not weaker, rigor.

Conclusion: The Epistemological Transformation

This subsection has established the first major interconnection: process ontology generates fundamental uncertainty. In a world of becoming, the future is not hidden – it is not yet made. Uncertainty is not a failure of knowledge but the condition of freedom and innovation.

The four ontological transformations converge on a single macro-ontological claim:

"Economic reality consists of irreversible relational-temporal processes that generate emergent structures which stabilise provisionally and then transform."

Economic entities – firms, markets, institutions – are therefore:

- **Relational:** constituted by network position
- **Temporal:** defined by trajectories, not essences
- **Emergent:** wholes irreducible to parts
- **Irreversible:** shaped by history, learning, and lock-in
- **Provisional:** real but contingent on sustaining processes

This changes what economics studies. The object is no longer the optimisation of isolated agents in timeless states, but the evolution of relational configurations through irreversible historical time.

Firms are not things with properties but trajectories through relational space – defined by network position, shaped by accumulated history, and sustained through continuous processual achievement. Economics therefore studies not optimising atoms, but evolving configurations of relations.

The epistemological transformation required is therefore:

FROM knowledge as complete decompositional representation

TO knowledge as partial, pattern-based, context-dependent understanding requiring epistemic complementarity.

Table 133: Table 5.8.1: The Epistemological Transformation

Dimension	FROM (Mechanistic)	TO (Process)
Status of future	Implicit in present	Created through unfolding
Nature of uncertainty	Temporary ignorance	Ontological openness
Predictive basis	Deduction from conditions	Pattern recognition in configurations
Knowledge type	Epistēmē (formal precision)	Epistēmē + Synesis (complementarity)
Understanding mode	Mechanistic (how)	Mechanistic + Organismic (how + why)
Orientation	Erklären	Erklären + Verstehen
Appropriate response	Collect more data	Build adaptive, exploratory capacity
Epistemic mode	Analysis sufficient	Analysis + Synthesis required
Epistemic focus	Micro-variables	Meso-structures and patterns

Decomposition destroys emergence. Understanding complex systems therefore requires synthesis to see what exists and analysis to test how it works. Systems thinking is not an option – it is the only epistemology compatible with a world of relational-temporal becoming.

5.8.2 Ontology → Methodology: Expansionism over Reductionism

Introduction

This subsection establishes the second major interconnection between Chapter 5's ontological reconstruction and the dissertation's methodological framework: emergence, process, and relationality require a corresponding transformation in method.

If economic reality is emergent (Section 5.1), processual (5.2), relational (5.3), and structured at the meso-level (5.5, 5.7), then methods built for decomposable, equilibrium-seeking systems are systematically inappropriate. Analytical decomposition, static snapshots, representative agents, and micro–macro aggregation do not merely approximate such systems poorly – they destroy the very properties they seek to explain.

This subsection operationalises the principle articulated throughout Chapter 5 – *method follows ontology* – by mapping each ontological transformation to its necessary methodological consequence. The core claim is strict: methodology is ontology-in-action. Once we change what we think economic reality is, we must change how we study it – otherwise we build models that eliminate the phenomena they purport to explain.

We define methodological expansionism as the principle that inquiry into complex adaptive systems must broaden the analytical frame – to include environment, interaction, and emergence – rather than narrow it through *ceteris paribus* isolation, element-wise decomposition, or micro-to-macro reduction. Expansionism is not methodological pluralism in the loose sense, nor a rejection of analytical discipline. It is the operational expression of the meta-principle: Method must preserve what ontology posits and epistemology requires.

Advocating synthesis does not mean abandoning analysis. Complex systems require both: analysis to understand mechanisms, and synthesis to preserve and explain emergent properties. The inherited body of economic theory errs not by using analysis, but by treating analysis as sufficient.

I. The Method–Ontology Mismatch Problem

Modern economics exhibits a deep incoherence. It increasingly adopts complex ontological claims – bounded rationality, network effects, path-dependence, institutional evolution – while retaining mechanistic methods: equilibrium models, representative agents, static optimisation, and decompositional empirics.

This is method–ontology mismatch: acknowledging that the Earth is round while navigating with flat-Earth maps. Ontological innovation has far outpaced methodological reform because methods are institutionally embedded in journals, training, and professional standards. But the result is systematic distortion.

For this framework, the mismatch produces three failures:

Theoretical inconsistency. Scholars claim emergence while using reductionist decomposition. Ontology says wholes exceed sums; method assumes wholes equal sums.

Empirical invalidity. Studying emergent phenomena with methods that eliminate emergence leads to false negatives – concluding that capabilities, routines, or institutions "do not exist" because they vanish under decomposition.

Missed insight. Researchers recognise path-dependence but analyse with comparative statics that assume history does not matter.

Firm-level illustration. Strategy research routinely claims that competitive advantage emerges from organisational routines. Yet these routines are measured by surveying individuals, coding tasks, and aggregating responses. The routine – the pattern of interaction – disappears. Felin and Foss (2005) showed that this methodological individualism systematically erases firm-level phenomena.

The same error structured McKinsey's *War for Talent* (Michaels et al., 2001): by measuring individual "A-players," it concluded that performance comes from hiring stars rather than from building coordination systems. The ontology of organisational capability was emergent; the method was atomistic.

Implication. If Chapter 5 is correct that economic reality is emergent, processual, relational, meso-structured, and uncertain, then methodological transformation is unavoidable. Complex adaptive systems cannot be studied with tools designed for mechanical ones.

II. Emergence Requires Synthetic Methods

Section 5.1 established that emergent properties are irreducible to their components. Analytical decomposition is the Cartesian method of isolating parts and aggregating their properties. These are logically incompatible.

Why analysis destroys emergence. Analytical decomposition presupposes: *Whole = Sum of parts*. Emergence asserts: *Whole > Sum of parts*. Properties that exist at the system level do not exist in the parts. Wetness is not in hydrogen or oxygen. Decomposing water into atoms destroys the property one seeks to explain.

Firm-level illustration. Toyota's knowledge-sharing capability does not reside in any individual employee. It exists in the relational architecture – who talks to whom, how problems are escalated, how solutions diffuse (Dyer & Nobeoka, 2000). Interview employees separately and aggregate their knowledge and the network disappears. The capability is not hidden – it simply does not exist at the level being measured.

This explains why capability research fails when it relies on individual-level surveys: the ontology is emergent; the method is reductive. As Fontana (2008) states, complex system properties cannot be inferred from isolated elements.

Google's *Project Aristotle* confirms this. Only when Google shifted from analysing individual attributes to interaction patterns did team effectiveness become intelligible (Duhigg, 2016). The property was relational, not individual.

Synthesis preserves emergence. Where analysis destroys emergent properties, synthesis preserves them by studying systems as integrated wholes embedded in larger contexts. Ackoff (1981) defined synthesis as understanding a system through its role in a containing whole. Synthetic methods include:

- **Case studies** (Yin, 2018) – preserve holistic organisational patterns
- **Ethnography** (Van Maanen, 2011) – captures routines-in-practice
- **Network analysis** (Borgatti & Halgin, 2011) – maps relational structure
- **Agent-based modelling** (Farmer & Foley, 2009) – allows emergence to occur computationally

These are not "alternative" methods. They are ontologically necessary for emergent systems.

Analysis + synthesis = systems methodology. Analysis explains how components work. Synthesis explains what the system is and what it does. Complex adaptive systems require both. As Ackoff (1981) noted, analysis alone cannot explain a system's purpose or emergent function.

Firm-level illustration. Amazon's logistics, algorithms, Prime, and data systems can each be analysed. But Amazon's dominance – lock-in, platform power, supplier dependence – is emergent. It only becomes intelligible through synthesis (Stone, 2013).

III. Process Requires Longitudinal Methods

Section 5.2 established process ontology: reality is becoming rather than static being; substances are stabilised processes rather than fixed essences (Section 5.4). This ontological transformation has an unavoidable methodological consequence: if reality is constituted through time, then methods that freeze reality at a moment cannot capture what it is.

Static cross-sectional analysis treats economic entities as if their essential properties exist synchronically – at a point in time. Process ontology denies this. If identity is trajectory, if causation unfolds through sequences, if path-dependence is real, then what a firm "is" cannot be read off from what it "looks like" at time t .

Why snapshot analysis misses becoming. Cross-sectional methods measure variables at a moment, compare across entities, and infer relationships from contemporaneous correlations. This presupposes that what matters is state, not history. But in processual systems, states are residues of trajectories, not self-explanatory objects.

Firm-level illustration (Nokia). Tripsas and Gavetti (2000) show that Nokia's collapse cannot be understood by comparing Nokia-2000 to Nokia-2010. That comparison reveals *what* changed, not *why*. Only process analysis reveals how success in feature phones created cognitive and organisational rigidities, how managerial frames locked in obsolete assumptions, and how reinforcing strategic choices generated an irreversible trajectory toward decline. Nokia-2010 is explained by Nokia-2000 → 2005 → 2007 → 2009, not by Nokia-2010 itself. Nokia's failure is intelligible only across its 70-year trajectory – from rubber to telecoms to mobile to smartphones – not from its 2010 balance sheet.

Firm-level illustration (Kodak). Cross-sectional analysis in 2000 showed Kodak as dominant; in 2012, bankrupt. The snapshots capture the outcome, not the process. Gavetti, Henderson, and Giorgi (2005) demonstrate that Kodak's decline arose from accumulated cognitive lock-ins, organisational inertia, and path-dependent investment patterns that could only be seen through longitudinal reconstruction. The state did not cause the outcome. The trajectory did.

Dynamic methods match process ontology. If reality is becoming, then methods must follow becoming. This requires temporal methodologies that track how systems evolve, not just where they are:

1. **Longitudinal analysis** – following the same entities through time to observe capability formation, adaptation, and rigidity. Panel data reveals trajectories invisible to cross-sections. Henderson and Clark's (1990) study of architectural innovation exemplifies this approach.
2. **Process tracing** – reconstructing causal mechanisms across sequences of events – how A led to B through C, D, and E. George and Bennett (2005) provide rigorous methodology for this form of causal history.
3. **Event history analysis** – studying when transitions occur – entry, exit, diversification, failure – rather than just whether they occur. Timing and sequence are causal variables in process systems (Tuma & Hannan, 1984).
4. **Evolutionary simulation** – allowing populations of firms, technologies, or strategies to evolve through variation, selection, and retention. Nelson and Winter (1982) show that path-dependent outcomes cannot be deduced from initial conditions and require computational evolution.

Firm-level illustration (Tesla). Static analysis in 2012 showed Tesla as a loss-making niche producer. Static analysis in 2021 showed Tesla as a \$1.2 trillion firm. Neither snapshot explains Tesla. Only the process does: Model S establishing premium legitimacy; battery learning curves lowering costs; Supercharger networks creating infrastructure lock-in; regulatory credits financing growth; software integration generating differentiation. Each stage conditioned the next. The firm in 2021 existed because of the specific sequence from 2012 onward. Tesla was not discovered in 2021. It was created by its trajectory.

IV. Relationality Requires Network Methods

Section 5.3 established relational ontology: entities are constituted through relations rather than possessing intrinsic properties independent of context. If this is correct, then network position is not an external "influence" on otherwise self-contained agents; it is part of what agents are – shaping information access, opportunity sets, influence, and even identity. Methodologically, this denies the adequacy of approaches that treat agents as isolable units with fixed attributes.

Why methodological individualism eliminates relations. Methodological individualism (MI) operationalises an atomistic ontology: agents are modelled as self-contained bundles of preferences and endowments, with relations treated as optional or instrumental. Under relational ontology, this is not a minor abstraction; it is a category mistake.

Firm-level translation (Silicon Valley). Saxenian (1994, pp. 29–57) showed that Silicon Valley's advantage is not reducible to individual genius or firm-level resources, but to network structure: dense horizontal ties, labour mobility diffusing tacit knowledge, and bridging connections enabling recombination. Individual-level measurement captures micro-attributes; it cannot capture the structural conditions that make those attributes productive. Two entrepreneurs with similar attributes can experience radically different outcomes depending on embeddedness and brokerage position. MI risks concluding "success is talent" when the mechanism is "talent in a position."

Firm-level translation (venture capital). Hochberg, Ljungqvist, and Lu (2007, pp. 251–301) documented that VC performance is strongly predicted by network position (syndication ties, centrality, brokerage), often more than by the VC's individual attributes. This is precisely what relational ontology implies: the "unit" that generates advantage is not the isolated actor but the actor-in-structure.

Network methods match relational ontology. If relations are constitutive, methods must measure and analyse relations directly, rather than treating them as residual context:

1. **Topology (structure as object of analysis).** Network methods represent who connects to whom and derive structural properties – degree distributions, clustering, path lengths, modularity. Wasserman and Faust (1994, pp. 17–66) provide the foundational toolkit.
2. **Positional analysis (position as capability).** Centrality measures operationalise how positions shape access and influence. Burt (1992, pp. 8–49) shows that structural holes generate brokerage advantages; Granovetter (1973, pp. 1360–1380) shows why weak ties transmit novel information across otherwise closed groups. These are not "effects of social context"; they are mechanisms of distributed causation.
3. **Dynamic networks (relations evolve).** Relational constitution is not static: ties form, dissolve, and reconfigure. Barabási and Albert (1999, pp. 509–512) formalise preferential attachment and hub formation; Snijders, van de Bunt, and Steglich (2010, pp. 44–71) provide methodology for modelling network evolution. Where equilibrium treats structure as given, dynamic network methods treat structure as endogenous and historically formed.

Firm-level translation (biotech and pharma innovation). Powell, Koput, and Smith-Doerr (1996, pp. 116–145) famously argued that the "locus of innovation" is the network. Innovative capacity depends on alliance structure, knowledge pipelines, and brokerage between academic and commercial domains – mechanisms invisible to firm-attribute regressions.

Lock. If entities are constituted through relations, then methods that model agents as isolable atoms systematically delete the causal structure that generates outcomes. Relational ontology implies that explanation must be structural and positional: what matters is not merely what agents have, but how they are connected.

V. Meso-Level Requires Meso-Preserving Methods

Section 5.5 established that key complexity properties manifest at the meso-level: routines, organisations, networks, industries, institutions. These structures are neither micro (individual choice) nor macro (aggregate outcomes); they are intermediate architectures where coordination, selection, diffusion, and lock-in occur. Methodologically, approaches that jump from micro to macro – especially representative-agent aggregation – make the meso-level invisible.

Why aggregation deletes meso. Representative-agent models collapse heterogeneity into a single "average" agent, eliminating not only diversity but also intermediate structures that depend on heterogeneity and interaction. This is not simplification; it is deletion of the level where many causal mechanisms operate.

Firm-level translation (strategic groups). Strategic groups (Caves and Porter, 1977, pp. 241–261; Porter, 1980, pp. 129–155) are meso-level phenomena: clusters of firms with similar strategies generating distinct competitive dynamics. Representative-firm approaches erase these clusters by construction. More broadly, Hall and Soskice (2001, pp. 1–68) show that institutional regimes shape firm strategies in patterned, intermediate-scale ways that vanish under micro→macro shortcuts.

Firm-level translation (platform ecosystems). Platform economies operate through meso-level communities of complementors, developer networks, and user clusters. Treating apps or complementors as identical optimisers destroys the ecosystem structure that generates differential success. Gawer and Cusumano (2014, pp. 417–432) require meso-level visibility precisely because platform advantage is mediated through intermediate architectures.

Meso-preserving methods. Preserving meso-level structures requires methods that maintain heterogeneity and allow intermediate architectures to be observed, measured, and analysed:

1. **Population heterogeneity (do not collapse diversity).** Population ecology preserves variation across organisational forms and tracks selection dynamics (Hannan and Freeman, 1989, pp. 45–94).
2. **Multi-level analysis (nested levels simultaneously).** Individuals nested in teams, teams in organisations, organisations in industries and institutional regimes: multi-level

designs preserve intermediate structures rather than erasing them (Hox, 2010, pp. 1–32).

3. **Configurational analysis (patterns as causal objects).** Configurational approaches identify causal combinations and "multiple paths" to outcomes, matching complexity causation (Fiss, 2011, pp. 393–414; Ragin, 2008).
4. **Agent-based modelling (heterogeneity + emergence + meso-structures).** ABM preserves heterogeneity, allows meso-structures (clusters, groups, ecosystems) to emerge, and makes intermediate-scale dynamics observable (Gilbert, 2008, pp. 1–42).

Lock. Meso-level reality requires meso-preserving methodology: heterogeneity, nesting, configuration, and simulation are not optional – they are the price of ontological coherence.

VI. Fundamental Uncertainty Requires Exploration Methods

Section 5.8.1 established that process ontology generates fundamental uncertainty: future states cannot be calculated because they do not yet exist as a closed set, but will be created through novelty-generating processes. This has a direct methodological consequence: prediction-and-optimisation methods assume a known possibility space; under Knight–Keynes uncertainty that assumption fails.

Why standard tools fail under uncertainty. Optimisation presupposes enumerable alternatives and assignable probabilities. Under fundamental uncertainty (Knight, 1921; Keynes, 1937), the outcome space is not fixed in advance. North (1999, p. 12) emphasises that the social world is non-ergodic, where learning does not converge to stable distributions. In such environments, the universality claim of standard dynamic tools – closed-form optimisation; stable stochastic processes – becomes untenable.

Firm-level translation (pharma R&D). Strategic planning tools (DCF, real options, probability-weighted scenarios) require probabilities; for genuinely novel ventures those probabilities are not available because categories of outcomes are endogenously created. In pharmaceutical R&D, portfolio optimisation based on estimated phase-transition probabilities can mislead when targets, mechanisms, regulation, and competition co-evolve in ways that transform the relevant distribution (Pisano, 2006, pp. 117–149).

Simulation and scenario methods match uncertainty. Under deep uncertainty, appropriate methods explore and stress-test possibility spaces rather than optimise within a presumed distribution:

- **ABM / computational experiments** (Tesfatsion, 2006, pp. 831–880) – allow patterns to emerge without specifying outcomes ex ante.
- **Scenario planning** (Schoemaker, 1995, pp. 25–40) – prepares for multiple plausible futures rather than forecasting one.
- **Stress testing** – evaluates robustness across adverse configurations.

- **Exploratory modelling** (Banks, 1993, pp. 435–448) – maps "what could happen under what conditions."
- **Process tracing** (Mahoney, 2012, pp. 570–599) – reconstructs realised trajectories when prediction is structurally limited.

Lock. If the future is ontologically open, methods must shift from prediction to exploration and robustness; simulation and scenario approaches become the coherent methodological response.

VII. Agent-Based Modelling as Methodological Exemplar

The five methodological implications – synthesis (emergence), temporal methods (process), network methods (relationality), meso-preservation (intermediate architecture), and exploration (uncertainty) – converge in agent-based modelling (ABM). ABM functions as a methodological exemplar because it can implement these commitments simultaneously.

Why ABM matches the ontology:

- **Emergence.** Macro patterns are generated from micro interactions rather than imposed by aggregation; Sugarscape provides canonical demonstrations (Epstein and Axtell, 1996, pp. 1–48).
- **Process.** Systems unfold through time; path dependence and transition dynamics are endogenous (Arthur, 2021, pp. 98–126).
- **Relationality.** Agents are embedded in network structures; topology governs diffusion and influence (Jackson, 2008, pp. 282–320).
- **Meso-level visibility.** Clusters, groups, ecosystems arise endogenously and are observable; agent-based finance explicitly preserves intermediate structures (Farmer and Geanakoplos, 2009, pp. 11–38).
- **Uncertainty navigation.** ABM supports computational exploration of outcome distributions without pretending to produce point forecasts.

Firm-level translation (electricity markets). Farmer and Foley (2009, pp. 685–686) argue that treating the economy as a complex system requires methods that preserve heterogeneity and interaction; ABM is designed for exactly that. In electricity market design, ABM captures emergent price volatility and capacity dynamics that equilibrium approaches systematically miss (Weidlich and Veit, 2008, pp. 39–64).

Firm-level translation (financial markets). Arthur et al.'s (1997) Artificial Stock Market shows that price volatility, bubbles, and crashes emerge from trading interactions, even when agents follow simple rules. Flash crashes are interactional events, not optimisation failures. Structure emerges from interaction, not aggregation.

Lock. ABM is not merely one tool among many; it is the clearest single methodological framework that operationalises the ontological transformations of Chapter 5.

VIII. Reject *Ceteris Paribus* Isolation

The *ceteris paribus* device and its failure. The *ceteris paribus* clause ("all else equal") is the foundational methodological device of the inherited body of economic theory. The procedure is familiar: isolate a factor, hold others constant, estimate a marginal effect, and then infer system behaviour by combining isolated effects. As Hausman (1992, p. 133) observes, economic theories are typically true only *ceteris paribus* because economists abstract from "all the others." This produces what Rol (2012, p. 75) calls *context-free regularities*: relationships treated as stable independently of the configurations in which they operate.

Ceteris paribus methodology presupposes four conditions: **independence** (variables can be isolated without altering their causal role); **separability** (effects are additive rather than interaction-dependent); **context neutrality** (causal effects do not depend on system configuration); and **linearity** (total outcomes are sums of partial effects).

These are not merely technical assumptions. They are the methodological expression of atomistic ontology. CAS ontology denies all four. In relational systems, variables cannot be isolated because their effects depend on configuration. Effects are interaction-dependent, not additive. Causal roles shift with context. Outcomes emerge from configurations, not linear sums.

Why *ceteris paribus* fails for CAS: Firm-level illustrations.

Technology adoption (Facebook). Price effects cannot be studied while holding network size constant, because network size changes as adoption occurs. Value depends on who else has adopted; adoption changes who else has adopted. This circularity is constitutive, not incidental. Facebook's expansion from one million users (2004) to three billion (2023) is unintelligible under *ceteris paribus* analysis because the effect of joining depends on how many have already joined.

Firm strategy (Samsung). R&D cannot be studied holding competition constant, because R&D changes competition, and competition changes R&D incentives. Samsung's \$22 billion R&D budget (2023) reflects co-evolution with Apple, TSMC, and the semiconductor ecosystem, not isolated optimisation.

Platform pricing (Uber). Uber's surge pricing cannot be analysed holding supply fixed. Prices alter driver entry; driver entry alters wait times; wait times alter rider demand. The price-quantity relation is embedded in a feedback system that cannot be decomposed into independent parts.

As Kešeljević (2023, p. 3) puts it, in open, relational systems "ceteris paribus assumptions do not hold" because interconnections are the system.

Methodological alternative: contextual open-system analysis. If ontology requires openness, then methodology must preserve openness. The appropriate alternative to *ceteris paribus* is therefore interaction-explicit modelling instead of variable isolation, configurational analysis instead of marginal effects, and feedback-preserving designs instead of independence assumptions.

IX. Property-Based Frameworks over Element-Based

Element-based vs property-based inquiry. Element-based approaches ask: *How many firms? What market shares? What distribution of skills?* Property-based approaches ask: *What network topology? What evolutionary trajectory? What feedback structure?*

The former treats systems as aggregates of elements. The latter treats systems as entities with emergent properties. As Bunge's (1997) systemism shows, system properties are not reducible to component properties. Section 5.1 makes this decisive: if emergence is real, element aggregation is insufficient by definition.

Why element-based methods fail for CAS:

Markets. HHI measures concentration but not competition. The airline industry ($\text{HHI} \approx 1600$) and smartphone industry ($\text{HHI} \approx 2200$) exhibit radically different innovation and rivalry dynamics because network structure and technological trajectories differ.

Labour markets. Wage regressions miss referral networks, institutional filtering, and matching structures. LinkedIn's 900 million users create network-driven labour dynamics invisible to individual-level aggregates.

As Newman (2018, p. 7) notes: "the pattern of interactions can have a big effect on the behavior of a system." It is topology, not totals, that governs CAS.

Property-based methods. CAS-appropriate methods target system-level properties directly:

Table 134: Property-Based Method Selection: Mapping CAS-Appropriate Methods to Their Targeted System-Level Properties

Method	System Property
Network analysis	Topology, centrality, clustering
Evolutionary theory	Variation, selection, retention
System dynamics	Feedback, delay, accumulation
Agent-based modelling	Emergence, adaptation, path-dependence

Arm's dominance in semiconductor IP, for example, reflects network centrality, not firm-level scale. Semiconductor boom-bust cycles reflect feedback structure, not isolated decisions.

As Fontana (2010, p. 21) and Hynes et al. (2021, p. 64) show, CAS are computationally irreducible: their properties cannot be captured by equilibrium equations.

X. The Four Dimensions of Methodological Expansion

Reductionist methodology proceeds by narrowing analytical scope: decomposing systems into parts, isolating variables, and eliminating context in order to achieve precision. Its virtue is tractability; its cost is ontological distortion.

Expansionist methodology proceeds by widening analytical scope: embedding elements within their relational, institutional, and historical contexts in order to preserve the structures from which system properties arise. Its virtue is ontological fidelity; its cost is complexity.

As Sayer (2000, pp. 15–28) shows from a critical-realist perspective, methodological adequacy is constrained by the nature of the object. Arthur (2015, pp. 1–23) makes the same point within complexity economics: understanding evolving economies requires expanding analysis to include the formation of structures, markets, and technological cascades – not refining static models through further disaggregation.

Expansionism operates through four necessary enlargements of analytical scope:

1. Temporal Expansion CAS must be studied as trajectories, not snapshots. Archer's (1995, pp. 65–92) morphogenetic cycles show that structures emerge, stabilise, and transform through time. Equilibrium snapshots erase path-dependence.

Firm-level translation. Nokia's failure is intelligible only across its 70-year trajectory – from rubber to telecoms to mobile to smartphones – not from its 2010 balance sheet.

2. Spatial Expansion CAS must be studied as embedded systems, not isolated units.

Firm-level translation. Apple's \$3 trillion valuation reflects ecosystem orchestration – developers, suppliers, content providers – not internal production efficiency. Foxconn's strategy is unintelligible outside the Apple–Arm–TSMC ecosystem. Cost structure, bargaining power, and investment strategy are all functions of network position, not internal attributes.

Firm-level translation. 3M's innovation capability – technology platforms, the 15% rule, internal venture funding, cross-division knowledge flows – is not reducible to individual skills. It exists only at the level of coordinated organisational structure, which reductionist methods destroy by construction.

3. Causal Expansion CAS exhibit **configurational causation**: conjunctural, equifinal, and path-dependent.

Firm-level translation. Silicon Valley emerged from a configuration (Stanford, defence spending, venture capital, labour mobility, culture), not a single cause. Replicating isolated elements fails because the configuration is what matters.

4. Methodological Expansion No single method captures CAS. **Triangulation is required.**

Firm-level translation. Toyota's capabilities require case studies (mechanisms), network analysis (supplier structure), statistical analysis (performance), and simulation (dynamic behaviour). No single method suffices.

Table 135: The Four Dimensions of Methodological Expansion

Dimension	Reductionist	Expansionist
Temporal	Snapshot	Trajectory
Spatial	Isolated unit	Embedded ecosystem
Causal	$X \rightarrow Y$	Configurations
Methods	Single technique	Triangulation

XI. Interaction-Centred Analysis

Interactions as ontologically primary. If **relations are constitutive** (Section 5.3), then methods must treat relations – not isolated elements – as primary. The inherited body of economic method proceeds from an atomistic schema: What are the elements? What properties do they possess? How do those properties determine outcomes?

CAS ontology reverses this logic: What are the interactions? What structures do interactions form? How do interaction patterns generate outcomes?

This is not a shift of emphasis but a shift of ontological object. When entities are constituted by relations, studying them apart from their interactional structure destroys what they are.

Newman (2018, p. 1) defines a network as "a collection of points joined together in pairs by lines," deliberately abstracting from node attributes in order to reveal the topology of relations. This abstraction is not a weakness but an epistemic gain, because "the pattern of interactions can have a big effect on the behavior of a system" (Newman, 2018, p. 7). In CAS, structure is causal. Outcomes are generated by who connects to whom, when, and how, not merely by what agents are.

Interaction-centred methods:

Network Analysis

- Primary data: who interacts with whom
- Primary objects: topology, centrality, clustering, brokerage, diffusion pathways
- Causal insight: position in structure shapes outcomes

Firm-level translation. Burt's (1992) theory of structural holes shows that brokerage positions predict innovation, influence, and career success – effects invisible to attribute-based analysis. LinkedIn's value proposition is literally topological: enabling information and opportunity to flow across structural gaps.

Event-Sequence and Path Analysis

- Primary data: ordered interactions across time
- Primary objects: critical junctures, branching points, cascades
- Causal insight: order matters

Firm-level translation. QWERTY dominance emerged from sequences – typing contests, manufacturer adoption, training regimes – not technical superiority. VHS defeated Betamax through timing and complementary market formation. These are interaction histories, not optimisation outcomes.

Agent-Based Modelling

- Primary mechanism: repeated interactions under local rules
- Primary object: emergent macro-patterns
- Causal insight: structure emerges from interaction, not aggregation

Firm-level translation. Arthur et al.'s (1997) Artificial Stock Market shows that price volatility, bubbles, and crashes emerge from trading interactions, even when agents follow simple rules. Flash crashes are interactional events, not optimisation failures.

Process Tracing

- Primary data: causal sequences linking events
- Primary object: generative mechanisms
- Causal insight: outcomes are produced through historical processes

Firm-level translation. Amazon's expansion from books to cloud computing cannot be inferred from cross-sectional data. It is a cumulative capability-building process in which each phase created the conditions for the next.

CAS methodology therefore shifts from what things are to how they relate.

XII. Methodological Pluralism as Synthesis

Given emergence, process, relationality, and fundamental uncertainty, no single model class can represent the system adequately across contexts. The appropriate response is not methodological monism, but pluralism disciplined by scope conditions. Rodrik's formulation captures the methodological stance succinctly: models are not literally true, but they can contain partial truths under specified domains (Rodrik, 2015, p. 8).

For our analytical framework, methodological pluralism is not eclecticism. It is synthesis at the model-theoretic level: different ontological properties imply different evidentiary objects; methods have scope conditions and failure modes; and integration occurs through principled synthesis.

Firm-level translation (Amazon). Understanding Amazon's competitive advantage cannot be achieved by a single method without ontological loss:

- **Case-based synthesis** can preserve the integrated whole (ecosystem-level complementarity between logistics, Prime, marketplace, and data).
- **Process tracing** can reconstruct the pathway by which these complementarities were assembled over time.
- **Network analysis** can map relational structure (seller networks, fulfilment linkages, AWS customer ecosystems).
- **ABM / computational experiments** can explore counterfactuals and robustness under alternative competitive and regulatory regimes.

This is not "method shopping"; it is ontological fit. When causation is complex and conjunctural, methodological synthesis can yield explanatory adequacy unavailable to any single approach (Fiss, 2007, pp. 1180–1210).

Lock. Pluralism is the methodological form of synthesis: it preserves emergence, process, relational constitution, and uncertainty by deploying multiple tools whose scope conditions align with the phenomenon's ontological properties.

XIII. Counter-Arguments Addressed

Counter-Argument 1: "Quantitative Methods Remain Superior for Generalisability"

Objection. A unified logic of inference favours quantitative designs as the gold standard for generalisation and causal identification. Case-based and synthetic approaches allegedly sacrifice generalisability.

Response. This objection conflates generalisability with universality. In complex systems, the relevant scientific task is often: *under what configurations does a pattern emerge?* not *does X*

hold universally? Configurational generalisation is therefore frequently more appropriate than variable-based generalisation. QCA identifies which combinations of conditions produce outcomes across cases, directly matching conjunctural causation (Ragin, 2008, pp. 176–198). Cases can generalise via theory-testing and "exemplar" status: a single case can falsify universal claims, or demonstrate a mechanism in a way that compels theoretical revision (Flyvbjerg, 2006, pp. 219–245). The appropriate generalisation criterion depends on ontology.

Counter-Argument 2: "Agent-Based Modelling Lacks Predictive Power"

Objection. ABM is too flexible, can be tuned to generate any result, and offers weak falsifiability.

Response. The prediction criterion is appropriate for closed systems with stable data-generating processes; it is not the sole criterion for scientific value in open, novelty-generating systems. ABM can illuminate dynamics, generate disciplined "how-possibly" explanations, and reveal structural sensitivities even when point prediction is impossible (Epstein, 2008, pp. 1.5–1.6). Pattern-oriented modelling constrains ABM by requiring simultaneous reproduction of multiple stylised facts; mismatches can falsify classes of models. ABM is therefore not "anything goes" if appropriately disciplined; it is a tool for exploring ontologically open systems where prediction is structurally limited.

Counter-Argument 3: "Case Studies Are Not Generalisable"

Objection. Small-n designs lack degrees of freedom for valid causal inference and invite selection bias.

Response. This criticism assumes statistical generalisation is the only form of inference. Case designs typically aim at analytic generalisation and replication logic. Case studies generalise to theory rather than populations (Yin, 2018, pp. 37–54). Multiple cases provide progressive theory development by testing whether predicted mechanisms recur under varying conditions (Eisenhardt and Graebner, 2007, pp. 25–32). Single cases can be unusually informative when they make mechanisms visible or dramatically challenge existing theories (Siggelkow, 2007, pp. 20–24). The issue is not "case vs statistics," but ontological fit.

Counter-Argument 4: "Mixed Methods Are Mere Compromise"

Objection. Mixed methods dilute rigour; incompatibilities between standardisation and contextual interpretation produce an incoherent hybrid.

Response. Integration can be principled rather than compromising: different methods answer different questions, and integration occurs at design and interpretation stages. The proper question is not "how many cases do I need?" but "which combination of cases and quantities is required to answer this kind of question?" (Small, 2011, pp. 57–86). Mixed designs become coherent when governed by ontological requirements: distributional features may require large-n, while mechanism identification may require deep case or process tracing. The synthesis is not eclecticism; it is an ontologically guided division of labour.

Counter-Argument 5: "Expansionism Sacrifices Scientific Rigour"

Objection. Isolation enables causal inference. Abandoning *ceteris paribus* means abandoning science.

Response. Isolation works only when variables are independent. In CAS, interdependence is the phenomenon. Eliminating it destroys causation. Rigour is not isolation but methodological appropriateness. Agent-based models rigorously simulate interaction. QCA rigorously tests configurations. Process tracing rigorously identifies mechanisms. These are not less scientific; they are scientific for complex systems (Ragin, 2008; George & Bennett, 2005; Mahoney, 2010).

Counter-Argument 6: "Property-Based Approaches Cannot Be Formalised"

Objection. Only element-based models permit mathematics.

Response. CAS uses different mathematics: graph theory (networks), non-linear dynamics (feedback), computational simulation (emergence), statistical mechanics (macroscopic properties). Complexity science is highly formalised – just not through aggregation algebra (Newman, 2018; Barabási, 2016; Epstein, 2006; Sterman, 2000).

Counter-Argument 7: "Expansionism Is Unfalsifiable"

Objection. If everything is connected, anything can be explained.

Response. CAS theories generate strong testable implications: diffusion follows network topology, lock-in follows increasing returns, volatility follows feedback density. These predictions are empirically falsifiable. What reductionism offers instead is often false precision – precise predictions about a mis-specified world (Cartwright, 1999; Lawson, 1997).

Counter-Argument 8: "Interaction-Focus Ignores Agent Heterogeneity"

Objection. Firms differ. Network models treat them as identical nodes.

Response. False dichotomy. Modern network and ABM methods incorporate heterogeneity – weighted edges, multiplex ties, differentiated agents. The shift is not from agents to relations but from agents alone to agents-in-relation. Indeed, relational perspective often reveals heterogeneity invisible to element-focus: structural position creates heterogeneous opportunities even among similar nodes. Position-based heterogeneity complements, not replaces, attribute-based heterogeneity (Borgatti and Halgin, 2011; Wasserman and Faust, 1994).

Conclusion: Methodology Is Ontology-in-Action

This subsection has established the second major interconnection: the ontological transformations documented in Sections 5.1–5.7 entail a corresponding methodological

transformation. The core claim is simple and strict: methodology is ontology-in-action. Once we change what we think economic reality is, we must change how we study it – otherwise we build models that eliminate the phenomena they purport to explain.

Method follows ontology. If reality is emergent, relational, and processual, then methods cannot isolate, decompose, or snapshot. *Ceteris paribus* assumes away interdependence; element-based frameworks miss emergence; reductionism destroys context. The expansionist research program broadens what we study, how we study, and over what time horizon we study. This is not less rigorous – it is more truthful.

Table 136: Table 5.8.2: The Method–Ontology Bridge

Ontological Transformation	Methodological Implication	Required Methods
Emergence: whole > sum (5.1)	Analysis-alone destroys phenomenon	Synthetic/holistic designs; ABM; network mapping
Process: becoming (5.2)	Snapshots miss mechanisms	Longitudinal; process tracing; event history; evolutionary simulation
Relationality: constituted entities (5.3)	MI eliminates structure	Network topology/position; dynamic networks
Meso-level: intermediate architecture (5.5)	Aggregation deletes causal level	Heterogeneity; multi-level; QCA; ABM
Uncertainty: open future (5.8.1)	Optimisation mis-specifies task	Scenarios; exploratory modelling; stress tests; simulation

Table 137: Table 5.8.2b: The Methodological Transformation

Dimension	FROM (Reductionist)	TO (Expansionist)
Isolation	<i>Ceteris paribus</i>	Contextual open-system
Classification	Element-based	Property-based
Scope	Narrowing	Integrating
Focus	Elements	Interactions
Causation	Linear $X \rightarrow Y$	Configurational
Rigour	Precision through isolation	Appropriateness through inclusion
Methods	Single (econometric)	Triangulated
Time	Snapshot	Trajectory

The transformation is therefore not optional but necessary. Once reality is relational-temporal and emergent (5.8.1), and once knowledge requires complementary analysis and synthesis, methods must preserve what ontology and epistemology require.

Ontology → Epistemology → Methodology

This triad forms the backbone of the dissertation's philosophical architecture.

5.8.3 Ontology→Structure: The Necessity of Open Systems and Boundary Permeability

Introduction

This subsection establishes the third and final interconnection in Section 5.8's bridge sequence: ontological reconstruction entails structural reconstruction. We define open systems as systems with permeable boundaries that maintain themselves through continuous exchange of energy, matter, and information with their environment, in contrast to closed systems, which presuppose impermeable boundaries that isolate internal dynamics from external context. This definition integrates thermodynamic reasoning, complexity science, and relational ontology into a unified framework, clarifying why a CAS-consistent ontology requires open-systems structure.

The core bridge principle is that structure follows ontology. Once we commit to emergence (Section 5.1), process (Section 5.2), and relational constitution (Section 5.3), we cannot coherently model economic systems as closed. Emergent properties depend on environmental coupling; stabilised processes require throughput; relational networks cross formal boundaries. The inherited body of economic theory frequently treats firms as self-contained production functions, markets as isolated domains, and economies as sealed systems – assumptions that are not merely simplifying but structurally incompatible with the ontological commitments established throughout Chapter 5.

This subsection serves three purposes within Chapter 5's architecture. First, it completes the Three-Bridge Sequence, showing that ontological transformation has systematic consequences: for epistemology (Section 5.8.1: process ontology implies fundamental uncertainty), for methodology (Section 5.8.2: emergence necessitates synthesis alongside analysis), and now for structure (CAS ontology entails openness). Second, it provides the ontological warrant for the structural implications developed in Sections 5.6–5.7 – showing that openness is not merely empirically convenient but ontologically necessary. Third, it reconceptualises the firm from a closed transformation mechanism into an open ecosystem node, where value emerges through boundary-crossing interactions rather than internal conversion alone.

Dynamic reinforcement. The ontology–structure relationship is bidirectional. Openness amplifies uncertainty by exposing systems to external perturbations (linking to Section 5.8.1). Relational structure makes network-based methods non-optional (linking to Section 5.8.2). Hierarchy and modularity enable stability under openness via near-decomposability (linking to Section 5.7), allowing adaptation without dissolution.

Before proceeding, one concern must be faced directly. If boundaries are permeable and systems extend indefinitely through relations, where does analysis begin and end? Our framework answers via analytical contingency: boundaries are real and consequential, but their location is partly a function of analytical purpose and scale. Firms have boundaries – yet not as impermeable containers. Boundary permeability varies by degree, and permeability itself becomes an explanandum, not an unexamined premise.

I. The Closed-System Assumption in Mechanistic Economics

Closed-system reasoning pervades the inherited body of economic analysis, typically as an implicit background condition rather than an explicit methodological choice. A closed system may be defined as one "isolated from external influences, where all interactions (the exchange of information, energy and matter), occur within the system's boundaries" (Lifshitz, 1976, p. 8). This assumption appears across multiple levels of economic modelling.

At the firm level, the canonical production function $Q = f(K, L)$ treats the firm as a closed transformation mechanism: capital and labour enter, output emerges, and the relevant causal action is presumed to occur within the bounded interior. The surrounding environment appears only parametrically – as prices that select the optimal input combination. Yet this formulation systematically brackets precisely those features that CAS ontology renders constitutive: supplier networks providing complementary capabilities, customer feedback shaping product development, knowledge spillovers enabling innovation, and institutional contexts structuring feasible strategies. The production function is therefore not merely simplified; it is structurally closed to the relational, processual, and emergent phenomena that generate competitive advantage.

Firm-level translation (boundary permeability). General Electric under Jack Welch (1981–2001) made "boundaryless organisation" a strategic doctrine precisely because performance depended on cross-boundary flows: knowledge circulation across divisions, supplier partnerships, customer co-development, and acquisition integration. GE's valuation expansion (from roughly \$14 billion to over \$400 billion during Welch's tenure) is not intelligible through an internal transformation metaphor alone; it is more coherently interpreted as ecosystem orchestration under varying degrees of boundary permeability. The point is not that the production function is false in all respects, but that it excludes the system-level causal architecture that an open-systems ontology identifies as primary.

At the market level, general equilibrium theory embodies what Lawson (2019, pp. 6–7, 9–15) characterises as an "atomistic ontology" – a tendency to model social reality as a collection of isolated entities operating in closed systems with stable regularities. Markets are treated as self-equilibrating mechanisms in which price adjustments clear supply and demand without requiring network topology, institutional evolution, or technological transformation to enter as constitutive structure. As Paniagua (2023, p. 578) observes, this rests on "ontological presuppositions of closed systems and regular events" – assumptions misaligned with the open-ended, relational character of social reality.

Firm-level translation (systemic risk). The 2008 financial crisis exposed the failure of closure assumptions under interdependence. Risk models treated institutions as separable units whose risks could be assessed independently. Yet Lehman Brothers' collapse (approximately \$639 billion in assets) triggered cascades because financial institutions were relationally constituted – connected through counterparty exposures, shared asset holdings, and confidence contagion. Systemic risk emerged from network topology that closed-system modelling systematically suppressed (Battiston et al., 2012, pp. 4–9).

The conceptual link between closure and substance ontology. The relationship between closed-system structure and substance ontology is not accidental; it is conceptual. If entities

possess intrinsic properties independent of context – if firms are what they are regardless of network position and agents' attributes are invariant to social embedding – then isolation is methodologically defensible. Closure operationalises substance ontology: study the entity in isolation and assume its essential properties remain intact. Mechanistic ontology and closed-system structure thus reinforce one another.

Chapter 5, however, has transformed the ontological basis for this structural choice. Section 5.1 problematised the mechanistic commitments – reductionism, atomism, and decomposability – that make closure appear natural. Sections 5.2–5.4 then established alternatives: emergence (5.1), process ontology (5.2), relational constitution (5.3), and the synthesis that substances are stabilised processes constituted through relations (5.4). If these commitments hold, then closure is not merely a restrictive assumption; it becomes ontologically incoherent. Systems cannot be analysed as isolated when their properties arise from interaction; processes cannot be frozen when stability depends on throughput; context cannot be treated as exogenous when entities are constituted through relations.

For our analytical framework, the failure of closed-system economics is therefore not only empirical inadequacy – though it generates that as well – but a deeper mismatch: closed-system structure applied to open-system reality is categorically inappropriate. The movement from substance to CAS ontology entails a corresponding movement from closure to openness as the structural baseline for economic analysis.

II. Four Bridge Principles: Ontology Requires Openness

This section develops four bridge principles through which Chapter 5's ontological commitments entail open-systems structure. The claim is not that openness is a convenient modelling choice; it is that, given emergence, process, and relational constitution, closure becomes structurally incoherent.

II.1 Emergence Requires Environmental Coupling

The first bridge principle follows from Section 5.1's account of emergent wholeness: emergent properties are context-dependent and therefore require environmental coupling. Closed-system representations suppress the interactions through which emergence is generated and sustained.

Consider firm capabilities. A closed-system view treats capabilities as internal resources – attributes the firm possesses and deploys independently. On the CAS account, many capabilities are emergent from interaction patterns that extend across organisational and legal boundaries. Toyota's lean capability, for example, is not exhaustively located inside the firm. It is sustained through supplier relations: just-in-time delivery requires fine-grained coordination; quality collaboration depends on joint problem-solving; learning depends on shared investments in capability development. Dyer and Nobeoka (2000, pp. 345–367) show how Toyota's knowledge-sharing network generates performance-relevant capabilities that no single firm – including Toyota – possesses in isolation. Closure – whether modelled as boundary insulation or pursued organisationally as self-sufficiency – degrades the relational architecture from which the capability emerges.

The relationship between emergence and openness is therefore reinforcing. Emergence requires openness because emergent properties arise from interaction – including interaction that crosses boundaries. Openness enables emergence by permitting the coupling through which context-sensitive properties arise. The central implication is that many economically relevant properties are partly positional: they depend on where the firm sits within supplier, customer, knowledge, financing, and institutional networks. This positional dependence is invisible in closed-system frameworks but constitutive under CAS ontology.

Firm-level translation (ecosystem capability). Silicon Valley's innovative capacity is not a property any single firm "has." It is an ecosystem-level property sustained by dense inter-firm ties, university linkages, venture capital networks, and labour mobility that enable rapid diffusion and recombination of knowledge. Saxenian (1994, pp. 29–57) demonstrates that isolating firms from these relational structures undermines the very capability the region is known for. Treating firms as closed production functions misses the open-network conditions under which innovative performance becomes systematically reproducible.

Firm-level translation (vaccine development as orchestration). Pfizer's COVID-19 vaccine illustrates emergence via environmental coupling: partnership with BioNTech, public procurement and funding arrangements, clinical trial infrastructures spanning jurisdictions, accelerated regulatory protocols, and manufacturing/logistics partnerships. The operative point is not that Pfizer "lacked capabilities," but that the realised capability was assembled and stabilised through boundary-crossing orchestration. A closed-system reading of "internal resources" cannot capture this emergent configuration.

Implication. Where economically significant properties are interaction-generated and context-dependent, closed-system modelling is not merely incomplete; it is structurally mis-specified because it deletes the coupling that generates the property in question.

II.2 Stabilised Processes Require Throughput

Section 5.6 established that firms are metabolic processors – dissipative structures that maintain their organisation through continuous flows of energy, matter, and information. This thermodynamic insight has a direct structural implication: stabilised processes require throughput, and throughput requires openness.

A completely closed system cannot sustain the flows necessary to maintain organisational complexity. Without exchange, entropy accumulates; without renewal, structure degrades. Closure is therefore not merely limiting – it is incompatible with the persistence of complex organisation.

Firm-level translations. Amazon's logistics network exemplifies throughput dependence. The system processes approximately 1.6 million packages daily through continuous flows of material inputs, energy inputs, and information inputs from over 300 million active customer accounts. The organisational capability that constitutes Amazon's logistics advantage exists only through these flows. Interrupt any one of them and the system rapidly degrades, as supply-chain disruptions during the COVID-19 pandemic made immediately visible.

Tesla's manufacturing system illustrates dissipative structure dynamics. The Fremont factory operates far from equilibrium through continuous flows of materials, energy, and information. During semiconductor shortages (2021–2022), output declined sharply despite unchanged physical infrastructure, demonstrating that organisational capability resides in flow patterns rather than static assets.

Implication. If firms are stabilised processes rather than fixed substances, then their persistence requires continuous throughput – which is definitionally incompatible with closure.

II.3 Relational Constitution Requires Boundary Permeability

Section 5.3 established that entities are constituted through relations rather than possessing intrinsic properties independent of context. If relations are constitutive, then firm identity extends beyond legal and physical boundaries to include network position, alliance structures, and institutional embedding.

This creates a structural requirement: boundaries must be permeable to the relations that constitute the entity. A firm's identity is not contained within the legal perimeter; it includes the developer ecosystem that gives a platform value, the supplier network that enables manufacturing, and the institutional affiliations that confer legitimacy.

Firm-level translation. Apple's identity is not exhausted by its Cupertino headquarters, assembly contracts, and intellectual property. Apple is a platform ecosystem, p. 28 million registered developers, a \$1.1 trillion App Store economy, and strategic dependencies with suppliers like TSMC and Foxconn. Remove the relational structure and Apple ceases to exist in its economically relevant form. The boundary is permeable because identity is relational.

Implication. Relational constitution implies that economically meaningful boundaries are not containers but interfaces – zones of structured permeability through which constitutive relations flow.

II.4 Co-Evolution Requires Openness to Novelty

Firms and their environments mutually transform over time. Technologies create new market possibilities; market feedback shapes technological trajectories; institutional frameworks evolve through firm-state interaction; competitive dynamics reshape industry structure. This co-evolutionary dynamic is structurally incompatible with closure because closed systems cannot receive the novelty necessary for adaptation.

Firm-level translation. The semiconductor industry exemplifies co-evolution. Moore's Law was not a natural phenomenon but a self-fulfilling coordination mechanism: firms invested in capability expecting others to follow, creating the demand that justified the investment. Government policy (CHIPS Act, EU Chips Act), geopolitical tension (US–China technology competition), and firm strategy (TSMC's Arizona fab, Intel's foundry pivot) continuously reshape one another. No firm in this system can be modelled as closed; all are embedded in co-evolutionary dynamics with technology, policy, and competition.

Implication. Co-evolution requires bidirectional exchange between firm and environment. Closed systems cannot evolve with their contexts; open systems can.

III. From Production Functions to Ecosystems

The four bridge principles jointly imply a reconceptualisation of the economic unit of analysis. If emergence requires environmental coupling, stabilised processes require throughput, relational constitution crosses formal boundaries, and system–environment co-evolution is endemic, then the firm cannot be coherently represented as a closed transformation device. The appropriate analytical unit becomes the firm-in-ecosystem: a node in an open network whose capabilities and outcomes are jointly produced through boundary-crossing interaction.

Problematising closed production functions. The canonical production function, $Q = f(K, L)$, is not simply a parsimonious representation of productive activity. It encodes a closed-system structure with three commitments that Chapter 5's ontology has rendered untenable: (1) **Closure** – outputs are fully specified by internal inputs, with the environment entering only through prices and constraints; (2) **Stasis** – the functional mapping is stable, treating production as a timeless relationship rather than a trajectory of becoming; and (3) **Independence** – optimisation is modelled as independent of network position, treating relations as external rather than constitutive.

Under the ontological reconstruction developed in Sections 5.1–5.4, these commitments are not merely restrictive; they are category errors. If many economically salient properties are emergent and context-dependent, then closure deletes the interactions that generate them. If entities are stabilised processes, then stasis suppresses the throughput that maintains persistence. If identity and capability are relationally constituted, then independence eliminates the positional determinants of performance.

Framework consequence. Competitive advantage – typically located by the inherited body of theory in internal, firm-specific resources – must be treated as partly positional: shaped by where the firm sits in supplier, customer, knowledge, financing, and institutional networks. Internal analysis remains necessary but becomes insufficient; it must be complemented by ecosystem analysis.

Firm-level translation (positional advantage in semiconductors). TSMC provides an instructive example. Its performance is not reducible to internal process excellence alone; it is also a function of ecosystem position – ties to equipment frontiers (EUV lithography supply chains), talent ecosystems, long-horizon customer relationships with leading fabless designers, and a supporting institutional environment. Competitors cannot replicate the outcome by copying internal routines alone; they would need to reproduce the network position and boundary architecture through which the capability is assembled and renewed.

Firm-level translation (smartphones as ecosystem competition). The mobile phone industry illustrates the conceptual shift. A closed-firm view treats Apple as producing iPhones from internal inputs. An ecosystem view treats Apple as orchestrating a platform: developers contribute complementary innovation; carriers shape distribution and adoption; suppliers deliver scale and component trajectories; users generate indirect network effects that attract developers and reinforce platform value. Here, value is co-produced through boundary-

crossing relations; capabilities are distributed across the ecosystem; and platform and complements co-evolve. The relevant unit is not "Apple" in isolation but Apple-in-ecosystem.

Toward ecosystem thinking. The direction of travel is clear: from closed production functions to open ecosystems. An ecosystem perspective treats firms as nodes in open networks where value emerges from interaction, not internal transformation alone. Boundaries are not containers but interfaces; capabilities are not concentrated but distributed; and adaptation is governed less by optimisation within fixed constraints than by co-evolution with shifting technological, institutional, and network environments.

Table 138: Table 5.8.3a: From Production Functions to Ecosystems

Dimension	Production Function	Ecosystem
Boundaries	Closed container	Permeable interface
Capabilities	Internal resources	Distributed across network position
Value creation	Internal transformation	Interaction and orchestration
Dynamics	Optimisation within constraints	Co-evolution with environment
Analysis	Firm in isolation	Firm-in-ecosystem
Ontology	Intrinsic properties	Positional, relational properties

IV. The Meso-Level and Boundary Permeability

This implication also reinforces Section 5.5's emphasis on the meso-level as the dissertation's distinctive analytical site. Boundaries are meso-level phenomena: neither micro (single transactions) nor macro (aggregate flows) but intermediate structures where tie density shifts, governance regimes change, and information flows alter character. The boundary is therefore not a line but a zone – a region in which internal coordination mechanisms transition into external interaction patterns.

This has a methodological and structural payoff. Pure micro-to-macro aggregation averages away boundary structure; pure micro analysis treats boundaries as exogenous constraints. By contrast, meso-level analysis – network topology, positional analysis, and institutional configuration – treats boundary permeability as an explanatory variable: how open is the interface, to what kinds of flows, under what governance, and with what consequences for capability formation?

Firm-level translation (Arm's licensing as boundary architecture). Arm's licensing model illustrates boundary permeability as design. Arm creates value not by vertically integrating manufacturing but by structuring an interface: licensing agreements, standards, and royalties coordinate a network of partners across chip design, fabrication, and device markets. The firm's "boundary" is a meso-level governance zone that enables distributed innovation and captures value through orchestrated complementarities. Arm's outcomes are therefore better understood as ecosystem returns to boundary architecture than as internal transformation returns to inputs.

Implication. Open-systems thinking is intrinsically meso-level thinking: openness is a property of boundary structures; boundary structures are meso-level configurations; and

studying openness therefore requires meso-preserving analytical commitments – exactly as established in Section 5.8.2.

V. Process-Relational Synthesis: Irreversibility

Emergence and process together imply a powerful result: economic wholes are historically irreversible. This is not ignorance; it is ontology. Two sources of irreversibility operate simultaneously.

Thermodynamic irreversibility. Dissipative structures exist only through continuous throughput (Prigogine & Stengers, 1984). When flows change, structure changes irreversibly.

Firm-level translation. Detroit could not return to its 1950s form. Skills, suppliers, and institutional knowledge dissipated. Recovery requires building new structures, not restoring old ones.

Historical (path) irreversibility. Learning, routines, standards, and expectations accumulate and constrain future possibilities. Nokia could not become Apple because it had become Nokia. Intel could not pivot to mobile because decades of x86 investment had created irreversible commitments.

Table 139: Four forms of economic irreversibility

Type	Mechanism	Example
Learning	Cumulative experience	Arrow's learning curves
Lock-in	Increasing returns	Windows, QWERTY
Institutions	Precedent	Legal and regulatory systems
Reputation	Trust hysteresis	Wells Fargo, VW

Once accumulated, these structures cannot be undone – only displaced by new trajectories.

Why irreversibility matters for economics. Comparative statics assume reversibility. Optimisation assumes free switching. Equilibrium assumes history can be ignored. CAS reality violates all three. Time has an arrow. The economy ages, learns, ossifies, and transforms. It is not a timeless machine but a historical process.

VI. Counter-Arguments and Responses

Counter-Argument 1: "Closed-System Models Are Useful Simplifications"

Objection. Closed-system assumptions are deliberate simplifications that enable tractable modelling. All models simplify; criticising closure for failing to capture everything misunderstands the nature of modelling.

Response. This objection conflates simplification with ontological distortion. A legitimate simplification preserves the structure that generates the phenomenon while abstracting from

secondary detail. Closed-system modelling does not merely abstract from the environment; it removes the interactions that constitute many economic properties. If firm capabilities, innovation, and competitive advantage emerge from network position, partner co-specialisation, and ecosystem co-evolution, then treating the firm as a closed transformation device does not simplify these phenomena – it eliminates them. As Rodrik (2015, p. 8) notes, "models are never true, but there is truth in models"; that truth requires that the model's structural assumptions match the ontology of what is being modelled. Empirically, the inadequacy of closure is evident in phenomena that are unintelligible without openness: supply-chain contagion, platform competition, ecosystem lock-in, and complementor co-specialisation (Adner, 2017; Kapoor & Lee, 2013). These are not peripheral complications but core drivers of outcomes.

Counter-Argument 2: "Open Systems Are Too Complex to Analyse"

Objection. If systems are open, everything connects to everything. Without closed boundaries, analysis becomes intractable.

Response. This objection rests on a false dichotomy between impermeable closure and boundaryless chaos. CAS ontology implies neither. It implies permeable, structured boundaries whose permeability varies and matters. Complex systems exhibit modularity and near-decomposability (Simon, 1962): dense interaction within subsystems and weaker but consequential coupling between them. This permits tractable analysis without requiring closure. Boundaries become objects of analysis rather than assumptions. Methodologically, this is handled through meso-level tools already developed in Section 5.8.2: network analysis, ecosystem mapping, multi-level modelling, and agent-based simulation. Open-system thinking does not abolish boundaries; it replaces fictional impermeability with empirically grounded permeability.

Counter-Argument 3: "Firms Have Real Legal Boundaries"

Objection. Corporations have owners, employees, and assets defined by law. These boundaries are real and non-negotiable.

Response. Legal boundaries are real but analytically insufficient. The mistake is to treat legal identity as if it coincided with economic, strategic, and capability boundaries. In practice, firms operate across multiple boundary regimes: capability boundaries include suppliers, partners, and knowledge sources; market boundaries include platforms, complementors, and users; strategic boundaries include alliances, standards bodies, and regulatory relationships. These boundaries rarely align with the legal perimeter of the corporation. Treating the firm as a closed container because it has legal identity is a category error: it confuses juridical demarcation with causal structure (Santos & Eisenhardt, 2005; Jacobides & Billinger, 2006). Open-system analysis does not deny legal boundaries; it recognises that economically relevant boundaries are multiple, overlapping, and permeable.

Counter-Argument 4: "Openness Is Empirical, Not Ontological"

Objection. Some firms are more open than others. Openness is a variable, not a structural necessity.

Response. This objection confuses degree with possibility. CAS ontology does not claim that all systems are equally open; it claims that no CAS can be fully closed. If entities are emergent, processual, and relationally constituted, then some degree of environmental coupling, throughput, and boundary permeability is logically required. A completely closed system cannot sustain dissipative structures, generate novelty, or maintain relational identity. Empirical variation exists in how open systems are, not in whether openness exists. Just as all objects above absolute zero possess thermal energy (varying in degree), all CAS possess openness (varying in permeability). Closure is not one end of a continuum; it is ontologically incompatible with CAS (Allen & Holling, 2010; Boisot & Child, 1999; Teece, 2007).

Conclusion: Ontology → Structure

This subsection has completed the third and final interconnection bridge of Section 5.8, demonstrating that CAS ontology entails open-system structure. Where mechanistic economics presupposes closed systems appropriate to substance ontology, CAS commitments to emergence, process, and relational constitution each independently require openness – and together generate co-evolutionary dynamics that closed systems cannot support.

The transformation established here is structural:

FROM closed systems – impermeable boundaries, intrinsic properties, isolated dynamics

TO open systems – permeable interfaces, positional properties, environmental coupling.

This shift is not a modelling preference but the structural expression of the ontological reconstruction carried throughout Chapter 5.

Table 140: Table 5.8.3b: The Ontology → Structure Transformation

Dimension	FROM (Closed Systems)	TO (Open Systems)
Boundaries	Impermeable containers	Permeable interfaces
Properties	Intrinsic to entity	Positional, context-dependent
Dynamics	Isolated, internal	Environmentally coupled
Capabilities	Internal resources	Emergent from networks
Throughput	Optional	Necessary for persistence
Evolution	Within fixed constraints	Co-evolution with environment
Firm model	Production function $Q = f(K, L)$	Ecosystem node in open network

Four bridge principles establish the necessity of this transformation:

- **Emergence requires environmental coupling** because emergent properties are context-dependent: firm capabilities arise from network interactions, not internal resources alone.

- **Process requires throughput** because stabilised entities are dissipative structures: closure means thermodynamic death.
- **Relationality requires boundary permeability** because constitutive relations cross formal boundaries: firm identity includes network position.
- **Co-evolution requires openness to novelty** because firms and environments mutually transform: closed systems cannot evolve.

Together these principles show that closure is ontologically incompatible with CAS. One cannot coherently combine CAS ontology with closed-system structure. If emergence, process, and relationality are accepted, openness follows as a matter of logical entailment.

This completes Section 5.8's three-bridge sequence:

- **5.8.1** showed that ontology determines epistemology (process $\Rightarrow$ fundamental uncertainty).
- **5.8.2** showed that ontology determines methodology (emergence $\Rightarrow$ synthesis, dynamics, networks).
- **5.8.3** shows that ontology determines structure (CAS $\Rightarrow$ open systems).

Ontology, epistemology, methodology, and structure are no longer separable. They form a single coherent architecture.

For firm theory, the implication is decisive. The firm is not a closed production function transforming inputs into outputs. It is an open, network-embedded ecosystem node whose capabilities, value, and identity arise through boundary-crossing interaction. Management therefore becomes boundary management: calibrating permeability, positioning in networks, orchestrating co-evolution.

What the inherited body of economic theory treated as "simplifying assumptions" are now exposed as ontological misclassifications: closed-system methods applied to relational, processual reality.

Openness is not an empirical option but a structural necessity following directly from the ontology established throughout Chapter 5.

5.9 Dynamic Implications – Emergence Enables Novelty Generation

Introduction

This section develops the dynamic implications of Chapter 5's ontological reconstruction. If reality is constituted by process (becoming rather than being) and emergent properties irreducible to their components, then economic systems are capable of genuine novelty – not merely new combinations of existing elements, but the creation of ontologically new properties, organisations, technologies, and markets.

We define generative dynamics as the conception of economic systems as endogenously creative: systems in which novel structures arise through emergent interactions, path-dependent transformation, and irreversible history, rather than through optimisation over a pre-given choice set. This definition operationalises the dynamic consequences of CAS ontology by specifying how emergence, process, and irreversibility together generate open-ended economic evolution.

The central claim is therefore not that economics must "add dynamics" to static models, but that CAS ontology transforms what dynamics mean. Under mechanistic ontology, dynamics describe movement toward equilibrium. Under CAS ontology, dynamics describe the continual creation of new economic possibilities. Economics shifts:

FROM the analysis of equilibrating systems → **TO** the analysis of evolutionary systems that generate novelty

This section serves three functions within the dissertation's architecture. First, it completes the implications sequence (ontology → epistemology → methodology → dynamics). Second, it provides the ontological foundation for non-linearity, innovation, and endogenous technological change. Third, it reconceptualises innovation, competition, and evolution as creative, not mechanical, processes.

The logical dependencies are explicit. Section 5.8.1 established that economic reality is relational-temporal and irreversible. Section 5.8.2 demonstrated that such reality requires expansionist methodology. Section 5.8.3 showed that CAS ontology requires open-system structure. Section 5.9 now shows that these commitments jointly entail generative dynamics – the capacity of economies to create what did not previously exist.

This is not an empirical add-on but a logical implication: *emergence enables novelty; process enables transformation; irreversibility gives history creative force.*

The challenge that shadows any claim of genuine novelty is epistemic: if novel structures do not exist prior to their emergence, they cannot be predicted. But this does not preclude scientific understanding. It means that economics must explain how novelty is generated and amplified, not attempt to forecast which specific novelties will occur. The choice is therefore not between prediction and rigour, but between mechanistic models that exclude novelty by construction and evolutionary models that explain how novelty arises.

5.9.1 The Novelty Problem in Mechanistic Economics

I. The Novelty Problem in Mechanistic Economics

Why the Inherited Body of Theory Cannot Generate Novelty

The inherited body of economic theory faces a fundamental novelty problem. Its core analytical devices – equilibrium, optimisation, and comparative statics – presuppose a pre-given space of possibilities. Dynamics are defined as movement within that space, not transformation of the space itself.

Equilibrium models describe states in which nothing fundamentally new occurs. Growth models describe quantitative expansion – more output, more capital, more labour – but not the emergence of qualitatively new economic forms. Even most innovation models treat novelty as recombination: new products are generated by rearranging existing technologies, routines, or knowledge elements.

This limitation is not accidental; it follows from reductionist ontology. If wholes equal the sum of their parts, then no genuinely new properties can arise. Novelty is restricted to rearrangement. The LEGO metaphor captures this precisely: one can build infinitely many structures, but the bricks remain unchanged. There is no ontological innovation – only reconfiguration.

Schumpeter recognised this constraint and explicitly rejected it. Economic development, he argued, consists of "changes in economic life as are not forced upon it from without, but arise by its own initiative, from within," and are therefore "entirely foreign to what may be observed in the circular flow or in the tendency towards equilibrium" (Schumpeter, 1934, p. 63). Development requires endogenous novelty, not movement toward a pre-defined attractor.

Yet even Schumpeterian economics was largely absorbed into the mechanistic framework by treating innovation either as exogenous shocks (technology drops from outside the economic system) or recombination (new combinations drawn from an existing set of elements).

Nelson and Winter (1982, pp. 3–21) exposed this problem by showing that innovation cannot be modelled as optimisation over a known choice set. Metcalfe (1994, pp. 931–945) further demonstrated that evolutionary economics must explain how variety is generated, not just how it is selected. But even evolutionary models often treat variation as recombinatorial rather than ontologically creative.

The limitation is therefore deeper than equilibrium theory: without emergence, there is no mechanism for the creation of genuinely new properties.

Firm-level translation. The iPhone illustrates the difference. A recombinatorial view treats it as a phone plus a touchscreen plus internet access. But what emerged was a new socio-technical regime: mobile platforms, app ecosystems, digital services, ride-hailing, mobile payments, and new patterns of communication and consumption. These properties did not exist in the components. They emerged at a new organisational and institutional level.

The novelty was therefore not just a new product but a new economic domain. Mechanistic ontology cannot account for this, because it cannot generate new levels of organisation. CAS ontology can – because emergence creates properties that do not exist in the parts.

This is the foundation of generative dynamics.

5.9.2 Five Dynamic Implications

II. First Dynamic Implication: Emergence Creates Ontological Novelty

Recombination versus Ontological Novelty

The relationship between Section 5.1's treatment of emergence and the present account of novelty is foundational. Emergence established that wholes possess properties irreducible to their components. This implies a crucial distinction that mechanistic economics cannot make:

Recombination refers to new arrangements of existing elements in which the elements themselves remain unchanged. Novelty is configurational: the same types of things appear in new patterns.

Ontological novelty refers to the emergence of genuinely new properties that did not exist at the level of the components. Novelty is constitutive: new kinds of things with new causal powers come into being.

Table 141: Table 5.9.1: Recombination vs Ontological Novelty

Dimension	Recombination	Ontological Novelty
Elements	Preserved	Transcended
Novelty	Configurational	Constitutive
Properties	Rearranged	Emergent
Predictability	In principle predictable	Structurally unpredictable
Metaphor	LEGO structures	Life from chemistry
Economic form	Product variants	New markets, organisations, technologies

Recombination preserves ontology; ontological novelty transforms it. Mechanistic economics, grounded in reductionism, can only represent the former. Emergence makes the latter possible.

Economic Forms of Ontological Novelty

1. Technological novelty

Technologies do not merely add new artefacts; they create new economic possibility spaces. Dosi's (1982, pp. 147–162) concept of technological paradigms captures this: innovations reorganise what kinds of problems can be posed and solved.

The steam engine was not merely an improvement in mechanical power. It generated factories, railroads, industrial urbanisation, and wage-labour capitalism – none of which existed in the component technologies. These were emergent socio-technical orders.

Firm-level translation. Tesla's electric-vehicle platform did not merely recombine batteries, motors, and cars. It created over-the-air software-defined vehicles, direct-to-consumer distribution, and integrated energy ecosystems – capabilities that did not exist within the prior automotive ontology.

2. Organisational novelty

The multidivisional (M-form) corporation was not a rearrangement of hierarchical components but the emergence of a new coordination regime: internal capital markets, strategic planning as a function, and portfolio management (Chandler, 1962).

As Dopfer, Foster and Potts (2004, pp. 263–279) show, such meso-level organisational forms carry evolutionary dynamics irreducible to either individual managers or macro aggregates.

Firm-level translation. AWS was not a recombination of IT services but the emergence of cloud computing as a new organisational form – on-demand scalability, pay-per-use infrastructure, and platform-based entrepreneurship. These properties did not exist prior to AWS's emergence.

3. Market novelty

Platform markets are not new venues for exchange but new market ontologies. Reputation systems, two-sided network effects, and algorithmic coordination generate market dynamics absent from classical exchange.

Firm-level translation. Airbnb did not merely connect spare rooms with travellers. It created stranger-to-stranger trust at scale, experience-based hosting, and regulatory arbitrage as competitive advantage – new institutional properties that reconstituted hospitality markets.

In all three cases, novelty is not quantitative (more of the same) but qualitative: new levels of organisation with new causal powers. This is what emergence delivers and what mechanistic economics cannot represent.

III. Second Dynamic Implication: Process Ontology Enables Creative Destruction

Becoming as the Ontology of Transformation

Sections 5.2 and 5.4 jointly establish that economic entities are not fixed substances but stabilised processes. This provides the ontological grounding for creative destruction.

Under substance ontology, entities possess fixed essences; change is external perturbation. Destruction becomes conceptually problematic.

Under process ontology, entities are provisional stabilisations; change is intrinsic. Creation and destruction are natural aspects of transformation.

Schumpeter's (1942, pp. 82–85) "creative destruction" thus becomes ontologically intelligible: not a metaphor but a description of how stabilised economic processes dissolve and re-form.

Arthur (2014, p. 6) makes this explicit: novelty enters from within the system through "new combinations of productive means." These are not disturbances of equilibrium; they are the engine of history itself. As Arthur (2021, p. 137) shows, equilibrium frameworks "make no allowance for the creation of new products or new arrangements... indeed, for history itself." Technologies form by recombination (Arthur, 2021, p. 141), triggering cascades of institutional, organisational, and market restructuring that keep the system permanently far from equilibrium.

The economy, in this view, is not a machine seeking rest but a self-creating process: "alive, ever-changing, organic" (Arthur, 2021, p. 136). This follows directly from irreversible time, endogenous novelty, and feedback-driven dynamics.

Examples of Creative Destruction as Process Transformation

For this framework, creative destruction operates through three structurally distinct but ontologically unified mechanisms.

1. Technological Substitution

Film photography did not merely lose market share to digital photography; the entire socio-technical process of photography was replaced.

Film photography depended on: chemical film, physical development, darkrooms, print distribution, and specialised professional skills.

Digital photography replaced this entire stabilised process with: sensors, software, cloud storage, online sharing, and algorithmic enhancement.

This was not addition but process succession.

Kodak employed 145,000 workers in 1988 to reproduce the film-based photographic regime. By 2012, after the digital regime displaced film, Kodak employed fewer than 8,000. Simultaneously, Instagram – with 13 employees – was sold for \$1 billion.

This is not market reallocation. It is process extinction and process emergence occurring together.

Firm-level translation: Blockbuster versus Netflix. Blockbuster and Netflix competed as trajectories, not strategies.

Blockbuster's dominance in 2004 – 9,000 stores and \$6 billion in revenue – was the product of a trajectory built on retail locations, late-fee economics, and physical distribution. Those same accumulated structures made streaming existentially threatening: adopting Netflix-style delivery would have required destroying its own profit engine. The organisation could not reverse its own history.

Netflix's trajectory accumulated differently: data-driven recommendations, logistics, subscription relationships, and then streaming infrastructure. When the technological regime shifted, Netflix's past became an asset while Blockbuster's past became a trap.

No "essence" explains this divergence. Only process does. Each firm was what it had become. Creative destruction is not an event – it is the unfolding of incompatible trajectories.

Netflix's subsequent evolution illustrates the same logic across multiple transitions:

- DVD-by-mail destroyed video rental stores (Blockbuster filed bankruptcy in 2010)
- Streaming destroyed physical media
- Original content production destroyed licensing dependency

Each stage replaced a stabilised reproduction regime with a new one: Retail rental → postal logistics → digital distribution → platform-based production. These were not incremental improvements. They were successive process regimes.

2. Organisational Form Evolution

Vertically integrated electronics firms were not "outcompeted" by modular networks; they were replaced as organisational process structures.

Vertical integration depended on: internal manufacturing + internal coordination + ownership-based control.

Modular production depends on: global supplier networks + contractual coordination + design-manufacturing separation.

Dell's rise reflected the stabilisation of modular assembly. Dell's later decline reflects its replacement by integrated smartphone ecosystems (Apple–TSMC–Foxconn), which represent a different stabilised coordination regime.

The firm did not "fail." Its organisational process lost viability.

3. Market Structure Transformation

Retail did not "move online." It was reconstituted as a different market process.

Physical retail depended on: geographic discovery, shelf-based search, in-store inventory, and local information asymmetries.

Platform retail depends on: algorithmic discovery, digital reputation, logistics networks, and global information symmetry.

Sears (2018), Toys "R" Us (2017), JCPenney (2020) did not merely lose customers. The retail process they embodied ceased to be reproduced. Amazon and Alibaba stabilised a new retail process.

Ontological Synthesis

In each case, destruction and creation are not sequential events but two sides of the same process transformation. Old stabilised processes dissolve. New stabilised processes emerge.

This is creative destruction grounded in process ontology: not firms failing, not markets reallocating, but processes being replaced by other processes.

That is Schumpeter's insight – now anchored in CAS ontology. Capitalism does not evolve by moving between equilibria but by continuously replacing the structures that constitute it.

IV. Third Dynamic Implication: Non-Ergodicity Creates Path-Dependent Innovation

History Shapes Novelty Generation

The relationship between Section 5.8.1's irreversibility and innovation dynamics is direct and decisive. If emergent wholes accumulate history through learning, lock-in, institutional sedimentation, and network formation, then innovation cannot occur in a fixed, timeless possibility space. Instead, innovation unfolds within a history-dependent, path-dependent landscape.

As Dosi (1988, pp. 1120–1138) showed, technological change proceeds along technological paradigms and trajectories – cumulatively reinforced directions of advance rather than free search across all conceivable designs. Once a trajectory forms, it structures future innovation: resources, skills, expectations, standards, and complementary assets align along particular directions, making some developments easier and others increasingly inaccessible.

This yields a fundamental contrast:

Ergodic Innovation (Mechanistic Assumption) The space of possible innovations is fixed and pre-given. All technologies already "exist" as latent options; innovation merely discovers them. History only affects the timing of discovery, not which technologies will ultimately prevail.

Non-Ergodic Innovation (CAS Reality) The space of possible innovations is created by prior innovations. Each new technology expands some future possibilities while foreclosing others.

The system does not explore a predefined landscape – it constructs the landscape as it moves through it. Sequence matters; different histories generate different futures.

This is why innovation is not merely uncertain but historically contingent.

Firm-level translation: QWERTY. Under an ergodic view, the allegedly superior Dvorak keyboard should eventually replace QWERTY. If the best option always wins, time alone should suffice.

But what actually happened was non-ergodic:

- Early contingent adoption of QWERTY
- → learning investments by typists
- → training standards
- → equipment compatibility
- → network effects
- → lock-in

Once this cumulative process took hold, the system could not "reset" to a different keyboard even if it were technically superior. The path itself became the constraint. History selected the future.

The Adjacent Possible

Stuart Kauffman's concept of the adjacent possible formalises how non-ergodic innovation operates.

The adjacent possible is the set of new configurations that can be formed from what currently exists. It is not the space of all imaginable technologies, but only those that are one step away from the present state.

As Kauffman (1995, p. 24) puts it, order and novelty arise "from the interaction of simple components following simple rules." The adjacent possible is generated endogenously by existing structures.

Three implications follow:

1. Possibility Space Grows Through Innovation

Each innovation expands the space of what can come next.

The transistor created the adjacent possible for integrated circuits. Integrated circuits created the adjacent possible for microprocessors. Microprocessors created the adjacent possible for personal computers.

Innovation does not move through a fixed design space – it creates new dimensions of that space.

Firm-level translation: Smartphones. The iPhone (2007) created the adjacent possible for mobile apps. Apps created the adjacent possible for ride-sharing (Uber, 2009) and mobile payments (Square, 2009). These created the adjacent possible for the gig economy (DoorDash, 2013).

None of these markets existed in 2006 because the adjacent possible had not yet been constructed.

2. Path-Dependent Expansion

Which possibilities are realised determines which new possibilities appear.

If vacuum tubes had remained dominant, the adjacent possible would have evolved differently. Digital computing as we know it would not have existed.

Firm-level translation: Pharmaceuticals. Antibiotics created the adjacent possible of antibiotic families. Resistance created the adjacent possible of resistance-targeting drugs. mRNA research created the adjacent possible for rapid pandemic vaccines.

Moderna and BioNTech did not simply "choose" mRNA in 2020. The adjacent possible that made mRNA vaccines feasible had been built over decades.

3. Unpredictable Futures

Because the adjacent possible itself changes, distant futures are not even definable in advance.

From 1970s packet-switching, email was adjacent. Social media was not – because broadband, browsers, smartphones, and platforms had not yet created the necessary space.

Innovation therefore cannot be predicted as discovery within a known space. It is the historical construction of new spaces.

Table 142: Table 5.9.2: Expansion of the Adjacent Possible

Era	Actual State	Adjacent Possible	Example
1970s	Packet switching	Email, file transfer	ARPANET
1990s	Web & browsers	E-commerce, portals	Amazon, Yahoo
2000s	Platforms & broadband	Social media, cloud	Facebook, AWS
2010s	Smartphones & apps	Sharing economy	Uber, Instagram
2020s	AI platforms	Generative AI, autonomy	OpenAI, Tesla

Ontological Conclusion

Non-ergodicity means that economic evolution is not search but construction. The future is not hidden in the present. It is made by the present.

This is why CAS economics must be evolutionary: not because it is fashionable, but because ontology makes it unavoidable.

V. Fourth and Fifth Dynamic Implications: Feedback Amplifies Novelty, Enabling Open-Ended Evolution

Positive Feedback as the Engine of Structural Transformation

The distinction between negative and positive feedback marks the dividing line between mechanistic and evolutionary economics.

Mechanistic systems rely on negative feedback: deviations are corrected, restoring equilibrium. Complex adaptive systems rely on positive feedback: deviations are amplified, transforming structure.

Table 143: Table 5.9.3: Negative vs Positive Feedback

Feedback Type	Function	Economic Role
Negative	Dampens deviation	Stabilisation, homeostasis, equilibrium
Positive	Amplifies deviation	Lock-in, dominance, path-dependence, innovation cascades

Negative feedback supports equilibrium. Positive feedback creates history.

In CAS, novelty does not diffuse linearly. It snowballs. A small initial deviation – an innovation, a platform, a design – gets amplified through self-reinforcing mechanisms that convert marginal differences into dominant structures.

This is why economic evolution is punctuated rather than smooth: innovations either die or explode. As Foster and Metcalfe (2001, pp. 1–18) show, economic self-organisation is driven not by restoring forces but by amplifying loops that transform micro variations into macro structures. This is how technological paradigms, dominant designs, and industrial regimes arise.

Four Mechanisms of Novelty Amplification

1. Increasing Returns (Arthur, 1989)

Adoption → learning & scale → lower costs → more adoption

Early advantage compounds into dominance.

Firm-level translation. Windows achieved dominance not because it was technically superior but because its early installed base attracted developers, which attracted users, which attracted more developers. By the 1990s, Windows exceeded 90% market share – a path-dependent lock-in, not an equilibrium outcome.

2. Network Effects (Katz & Shapiro, 1985)

Users → value → more users

Value increases endogenously with participation.

Firm-level translation. Facebook's 3 billion users (2023) were not "chosen" by the market. They were produced by self-reinforcing network dynamics: more users → more content → more engagement → more users. Winner-take-all is not a market failure – it is a CAS outcome.

3. Complementarities (Milgrom & Roberts, 1990)

Innovation A increases value of B → adoption of B increases value of A

Technologies co-evolve in clusters.

Firm-level translation. The Wintel ecosystem (Windows + Intel + peripherals + software) dominated computing for two decades not through isolated excellence but through mutual reinforcement among complements. Once established, the ecosystem became self-locking.

4. Learning Curves

Production → learning → cost reduction → more production

Experience becomes advantage.

Firm-level translation. Solar PV costs fell ~89% between 2010 and 2020 because deployment generated learning, which lowered costs, which accelerated deployment. Market adoption created its own competitiveness. The technology did not "compete" into viability – it learned into it.

Table 144: Table 5.9.4: Mechanisms of Novelty Amplification

Mechanism	Feedback Loop	Example	Outcome
Increasing returns	Adoption → learning/scale → adoption	Windows	Dominant design
Network effects	Users → value → users	Facebook	Platform monopoly
Complementarities	A → B → A	PC ecosystem	Co-evolution
Learning curves	Production → learning → cost → production	Solar PV	Cost parity

Positive feedback converts small historical accidents into structural inevitabilities.

Open-Ended Evolution

When ontological novelty (Section II), creative destruction (Section III), non-ergodicity (Section IV), and feedback amplification combine, the economy becomes open-ended.

Not moving toward equilibrium. Not exploring a fixed landscape. But continually creating new landscapes.

Nelson and Winter's (1982) variation-selection-retention framework still applies – but now it operates on ontological novelty, not just recombination. Metcalfe (2005) makes the crucial extension: evolutionary economics must explain not only which variants survive, but how new kinds of variants come into existence.

Four Consequences of Open-Ended Evolution

1. No Equilibrium Endpoint

Dominant designs are only temporary stabilisations.

Firm-level translation. The internal combustion engine ruled for a century. Now EVs are replacing it. EVs will in turn be displaced by autonomous mobility, energy-as-service, or technologies not yet conceivable. "Equilibrium" is just a pause between revolutions.

2. Radical Unpredictability

Future innovations do not yet exist, so they cannot be forecast.

Firm-level translation. In 1990 no one could predict social media because the adjacent possible – broadband, cloud computing, smartphones – did not yet exist. Meta's \$1.1 trillion valuation (2024) was structurally unpredictable from the 1990s.

3. Irreversible Development

Paths create lock-ins that cannot be undone.

Firm-level translation. Toyota's production system reflects 70 years of accumulated routines, supplier coordination, and learning. These cannot be replicated by competitors starting from scratch. History becomes capability.

4. Perpetual Disequilibrium

Capitalism is not a tendency toward rest – it is a far-from-equilibrium process.

As Foster (2005, pp. 873–895) shows, economic systems maintain themselves through continuous structural transformation, not convergence. Stability is exceptional. Disequilibrium is the rule.

Dynamic Conclusion

The economy is not a machine selecting among known options. It is a creative system generating new options.

Emergence creates new kinds of things. Process destroys and rebuilds structures. History channels innovation. Feedback turns novelty into regimes.

This is why economics must be evolutionary: not as a modelling choice, but because ontology demands it.

5.9.3 Counter-Arguments and the Generative Dynamics Framework

VI. Counter-Arguments and Responses

Counter-Argument 1: "Novelty Is Just Recombination at a Different Level"

Objection. What appears as "ontological novelty" is simply recombination at a finer level of description. The iPhone combines existing technologies; platforms combine existing market mechanisms. If analysed deeply enough, nothing genuinely new exists.

Response. The objection confuses composition with reduction. Emergent wholes are indeed composed of parts – but this does not imply that their properties are reducible to those parts.

The criterion is causal powers. If a whole possesses powers that do not exist in its components, then the whole is ontologically novel even if composed of them. Life is composed of chemistry, yet metabolism, reproduction, and evolution do not exist at the chemical level. Organisations are composed of individuals, yet coordination, strategy, and institutional memory do not exist in individuals in isolation.

This is the core insight of emergentist social ontology (Bhaskar, 1975; Archer, 1995; Elder-Vass, 2010): downward causation and system-level powers distinguish emergence from mere aggregation. Recombination rearranges existing causal powers; ontological novelty introduces new ones.

Counter-Argument 2: "Creative Destruction Is Exaggerated – Most Change Is Incremental"

Objection. Most innovation is incremental improvement within existing trajectories. Radical disruption is rare. Creative destruction overstates discontinuity.

Response. The objection is correct at the micro-temporal scale but false at the structural scale.

Most innovation is incremental within paradigms. But paradigms themselves are periodically overturned. Kodak made thousands of incremental improvements to film chemistry – until digital photography destroyed the film paradigm. Blockbuster optimised video retail – until streaming eliminated it. These are not anomalies; they are phase transitions in evolutionary systems.

The process ontology recognises multi-scale dynamics: incremental change dominates within regimes; creative destruction dominates between regimes (Tushman & Anderson, 1986; Utterback, 1994; Christensen, 1997).

Counter-Argument 3: "Path-Dependence Does Not Preclude Optimality"

Objection. Even if history shapes trajectories, selection may still drive systems toward optimal outcomes along each path.

Response. This assumes a path-independent optimum – which non-ergodicity denies.

Lock-in produces stability, not optimality. QWERTY persists not because it is superior but because switching costs dominate performance differences. What becomes "best" is defined by the path taken – which itself depends on early contingencies.

There is no global optimum outside history against which outcomes can be judged (David, 1985; Arthur, 1989; Vergne & Durand, 2010). Economics becomes the study of historically conditioned trajectories, not convergence to timeless ideals.

Counter-Argument 4: "Positive Feedback Is Self-Limiting"

Objection. Dominant designs are often displaced (MySpace → Facebook). Positive feedback does not guarantee monopoly or permanence.

Response. Exactly – and this supports the theory.

Positive feedback does not produce permanence. It produces temporary structural dominance until displaced by new positive-feedback loops triggered by new innovations. History proceeds through successive amplification regimes: VHS → DVD → streaming → platforms.

Creative destruction is feedback-driven succession (Suarez, 2004; Schilling, 2002). Stability exists only between waves.

VII. Integration: The Generative Dynamics Framework

Synthesis

The five dynamic implications now form a single causal architecture:

1. **Emergence** creates **ontological novelty**
2. **Process ontology** enables **creative destruction**
3. **Non-ergodicity** creates **path-dependent innovation**
4. **Positive feedback** amplifies novelty
5. **Open-ended evolution** prevents convergence

This is generative dynamics: economies as systems that create new structures, not merely select among given ones.

Table 145: Failed Dynamic Hybrids

Hybrid	Why Incoherent
CAS ontology + Equilibrium	Emergence prevents convergence
Ontological novelty + Ergodic space	Novelty creates history
Creative destruction + Fixed categories	Destruction transforms categories

Coherent Combination: CAS ontology + generative dynamics

Table 146: Table 5.9.5: Dynamic Coherence

Combination	Reason
CAS + equilibrium	Emergence breaks convergence
Novelty + ergodicity	History is constitutive
Destruction + fixed categories	Categories evolve
CAS + generative dynamics	Dynamics match ontology
Mechanistic + equilibrium	Both assume fixed state space

Table 147: Table 5.9.6: From Equilibrium Mechanics to Generative Dynamics

Dimension	Equilibrium Economics	Generative Dynamics
End-state	Optimum	None
Time	Reversible	Irreversible
Novelty	Exogenous	Endogenous
Structure	Fixed	Evolving
Stability	Normal	Exceptional
Change	Quantitative	Qualitative

Conclusion: The Dynamic Transformation

This section has articulated the dynamic implications of Chapter 5's ontological transformation. The five implications – ontological novelty through emergence, creative destruction through process ontology, path-dependent innovation through non-ergodicity, novelty amplification through positive feedback, and open-ended evolution – converge on a single dynamic stance: generative dynamics.

Table 148: Table 5.9.7: The Dynamic Transformation

Dimension	FROM (Mechanistic)	TO (Generative)
Novelty	Recombination	Ontological emergence
Innovation source	Exogenous shocks	Endogenous creation
Evolution	Search over fixed set	Creation of expanding set
End-state	Equilibrium	Perpetual transformation
History	Irrelevant (ergodic)	Constitutive (non-ergodic)
Feedback	Negative (stabilising)	Positive (amplifying)

Dimension	FROM (Mechanistic)	TO (Generative)
Trajectory	Predictable	Path-dependent

The transformation accomplished here is dynamic: FROM economics as the study of equilibrium states reached through optimisation TO economics as the study of evolutionary processes that generate continuous novelty.

This completes the philosophical reconstruction initiated in Sections 5.8.1–5.8.3. Once reality is relational-temporal (5.8.1), methods must be expansionist (5.8.2), structure must be open (5.8.3), and dynamics must be generative (5.9).

The Four Implications Chain

- **5.8.1 Epistemological** – Knowledge requires analysis–synthesis complementarity under fundamental uncertainty
- **5.8.2 Methodological** – Research must preserve context, relations, and emergence
- **5.8.3 Structural** – Firms exist as open systems within ecosystems
- **5.9 Dynamic** – Economic systems generate genuine novelty through evolution

Together, these form a closed philosophical system: ontology → epistemology → methodology → structure → dynamics. CAS is not merely a different modelling toolkit – it is a different conception of economic reality, knowledge, inquiry, and change.

Firm-Theoretic Implications

Firms exist inside generative dynamics. They:

- create ontological novelty (platforms, ecosystems, organisational forms),
- undergo creative destruction,
- accumulate irreversible capabilities,
- are shaped by positive feedback (learning curves, networks, complementarities),
- and evolve without a predetermined optimal form.

There is no "best" firm – only historically contingent, dynamically stabilised organisational forms embedded in evolving ecosystems.

Arthur captures this precisely: the economy is a system in which "the very explorations agents undertake alter their situation, which requires them to explore and adapt afresh" (Arthur, 2021, p. 137).

The economy does not converge toward a solution – it creates the problem space it must continually solve.

Conclusion

This section concludes Chapter 5's ontological reconstruction. The chapter has established that economic reality is emergent rather than additive (5.1), processual rather than static (5.2), relational rather than atomistic (5.3), trajectory-based rather than substance-based (5.4), meso-structured rather than purely micro-macro (5.5), far-from-equilibrium rather than equilibrium-seeking (5.6), network-organised rather than atomistically distributed (5.7), and open rather than closed (5.8.3).

From this ontology follow necessary implications: epistemologically, knowledge must be complementary rather than purely analytical (5.8.1); methodologically, inquiry must be expansionist rather than reductionist (5.8.2); structurally, systems must be open rather than closed (5.8.3); and dynamically, economic systems must be understood as generative rather than equilibrating (5.9).

These four transformations – ontological, epistemological, methodological, and dynamic – together imply a radical redefinition of what economics studies. The economy is not a system that selects among known alternatives; it is a system that creates new alternatives that did not previously exist. It is not governed by reversible stochastic processes exploring a fixed state space; it is governed by non-ergodic, path-dependent dynamics that continually expand the space of possibilities itself. In such systems, the future is not merely unknown – it is ontologically indeterminate.

The economy does not search a landscape – it builds the landscape.

Emergence creates new kinds of things. Process destroys old stabilisations. History channels innovation. Feedback turns accidents into regimes. And evolution never ends.

Summary of Achievements

Chapter 5 has completed an ontological reconstruction of economic reality around the core commitments of complex adaptive systems. The reconstruction is not a change in modelling taste but a change in the kind of world economic theory claims to describe. Where orthodox foundations presuppose atomistic entities with intrinsic properties, stable functional relations, and equilibrium-seeking dynamics, this chapter has shown that the central phenomena of modern economies – capabilities, systemic fragility, networked competition, platform power, and endogenous technological regimes – require an ontology of emergent wholes, stabilised processes, and relational constitution operating primarily at meso-level scales.

First, the chapter established emergence as ontological (Section 5.1). Organisational capability, market structure, and institutional order are not aggregates of micro attributes but system properties produced by interaction architecture. IDEO's innovation capability does not reside in any designer; Toyota's production system cannot be replicated by hiring its employees; Silicon Valley's dynamism transcends any firm within it. This renders reductionist shortcuts – representative agents, additive aggregation, and closed production-function representations – more than incomplete: they are ontologically mis-specified for domains where interaction-generated properties do the causal work.

Second, the chapter established process ontology as foundational (Section 5.2). Firms and markets persist not because they possess timeless essences, but because they reproduce themselves through throughput, coordination, and boundary-work. Stability is therefore not a starting assumption but an outcome – a continuous achievement in far-from-equilibrium conditions. Nokia and Apple competed as trajectories, not strategies; Blockbuster's 9,000 stores became a trap rather than an asset when the technological regime shifted. Once this is granted, time is no longer a neutral parameter: history becomes constitutive, and irreversibility becomes part of what the object is.

Third, the chapter established relational constitution as foundational (Section 5.3). Relations are not external constraints on pre-formed entities; they are constitutive of identity and capability. Tesla's value derives not from manufacturing assets but from network position – charging infrastructure, regulatory relationships, supplier ecosystems, and platform dynamics. The firm is therefore not adequately conceptualised as a closed transformation device, but as a node embedded in, and partly defined by, the networks and institutions through which it operates.

Fourth, the chapter synthesised these pillars into a unified account of substances as stabilised processes (Section 5.4). Identity is trajectory, not essence. Amazon's transition from bookseller to cloud-computing giant illustrates how firms are what they have become through irreversible historical development. This synthesis resolves the apparent tension between process (flux) and entity (stability) by showing that persistence is work, not given.

Fifth, the chapter established the meso-level as causally privileged (Section 5.5). The micro-level destroys emergent properties; the macro-level averages them away. Only at the meso-level – Silicon Valley's innovation ecosystem, Toyota's supplier network, the pharmaceutical industry's regulatory architecture – are interaction-generated properties preserved and observable. Meso-level analysis is therefore not a methodological convenience but an ontological necessity.

Sixth, the chapter translated ontology into structural implications (Sections 5.6–5.7). CAS are thermodynamically open systems maintaining organisation through continuous throughput. Boundaries are permeable gradients, not impermeable walls. Network architecture – scale-free topology, hub-and-spoke structures, cluster formation – provides the structural substrate through which relational constitution operates. Apple's \$3 trillion valuation reflects not internal assets but orchestration of a trillion-dollar ecosystem; CalPERS' \$440 billion portfolio exercises influence through network position, not ownership alone.

Seventh, the chapter derived a strict implications chain from ontology to epistemology, methodology, and structure (Section 5.8). Epistemologically, emergence and relational constitution entail that analysis alone cannot exhaust understanding: decomposition yields precision but can erase the very structures that generate system behaviour. Hence the requirement for synthesis-guided analysis and disciplined pluralism across model types. Methodologically, ontological openness and meso-level causation entail that inquiry must be expansionist rather than reductionist: methods must preserve interaction, context, trajectory, and intermediate architectures rather than eliminate them via *ceteris paribus* isolation, element-wise decomposition, or micro–macro shortcuts.

Eighth, the chapter completed the implications chain with dynamic implications (Section 5.9). Economies and firms do not merely select among known options; they create new options through endogenous novelty, creative destruction, path-dependent trajectory formation, and positive-feedback amplification. The iPhone created not just a product but a new economic domain – app ecosystems, ride-hailing, mobile payments. Kodak's 145,000 workers in 1988 became fewer than 8,000 by 2012, while Instagram – with 13 employees – sold for \$1 billion. Windows achieved 90% market share not through technical superiority but through self-reinforcing feedback loops. Solar PV costs fell 89% between 2010 and 2020 because deployment generated learning, which lowered costs, which accelerated deployment. The technology did not "compete" into viability – it learned into it.

The Decisive Payoff

The decisive payoff is a redefinition of the economic unit of analysis. Closed production functions and closed-firm representations become limiting cases rather than default foundations. The coherent unit becomes the firm-in-ecosystem: a boundary-permeable, network-embedded process whose capabilities and outcomes are co-produced through interaction, orchestration, and institutional coupling.

The Four Transformations Completed

Table 149: The Four Completed Paradigmatic Transformations: Summary of the Shift from Mechanistic to CAS Foundations Across Ontology, Epistemology, Methodology, and Dynamics

Dimension	FROM (Mechanistic)	TO (CAS)
Ontology	Atomistic substances with intrinsic properties	Emergent wholes constituted through relations
Epistemology	Pure analysis and decomposition	Analysis-synthesis complementarity
Methodology	Reductionism and <i>ceteris paribus</i> isolation	Expansionism preserving interaction and context
Dynamics	Equilibrium-seeking optimisation	Generative novelty through evolution

These transformations are not independent options but form a closed philosophical system. CAS ontology logically entails CAS epistemology, which logically entails CAS methodology, which logically entails CAS dynamics. One cannot coherently adopt CAS ontology while retaining mechanistic epistemology; one cannot coherently adopt expansionist methodology while presupposing equilibrium dynamics. The package is all or nothing.

Bridge to Chapter 6: The Epistemic Consequence

This conclusion creates a decisive break with the inherited body of economic theory. If novelty is endogenous, if structures are emergent, if feedback is amplifying, and if history is irreversible, then the foundational assumptions of equilibrium, optimisation, rational expectations, and probabilistic forecasting cannot be sustained. What economists traditionally treat as "risk" becomes something deeper: fundamental uncertainty. Probabilities presuppose a stable space of possible outcomes; generative systems destroy that space by continually creating new ones.

In 1990, no one could predict social media because the adjacent possible – broadband, cloud computing, smartphones – did not yet exist. Meta's \$1.1 trillion valuation in 2024 was structurally unpredictable from the vantage point of the 1990s. This is not forecasting failure; it is the epistemic signature of non-ergodic, novelty-generating systems.

Chapter 6 therefore does not introduce uncertainty as an add-on to otherwise stable models. It follows logically from Chapter 5's ontology. Once the economy is understood as a non-linear, path-dependent, novelty-generating CAS, uncertainty is not a behavioural imperfection, a market failure, or an information deficit – it is the necessary epistemic consequence of a world that is continuously becoming rather than merely being. Chapter 6 develops this consequence in full, showing how non-linearity, positive feedback, and structural change make Keynesian–Knightian uncertainty not a special case but the normal condition of economic life.

Final Synthesis

In sum, Chapter 5 has delivered an internally consistent ontological platform for the dissertation:

A world in which emergence is real, entities are processual, relations are constitutive, meso-structures are causally primary, and evolution is open-ended.

This platform does not merely accommodate complexity; it makes the inherited assumptions of closure, decomposability, and equilibrium the exceptional case. The economy is not a machine selecting among known options. It is a creative system generating new options.

The economy does not search a landscape – it builds the landscape.

Emergence creates new kinds of things. Process destroys old stabilisations and crystallises new ones. History channels what can be invented. Feedback turns small deviations into structural regimes. Disequilibrium is not failure – it is the normal condition of a creative economy.

And evolution never ends.

Chapter 6 – Conclusions, Contributions, Implications, Limitations, and Future Research Suggested

Introduction

This chapter implements the evaluative stage of problematisation. The dissertation's structure mirrors the three stages articulated in Section 2.3: diagnostic identification of the assumptive architecture generating persistent theoretical failures (Chapter 3); reconstructive transformation of those assumptions across ontological, epistemological, and methodological domains (Chapters 4–5); and evaluative assessment of whether the reconstructed framework generates more adequate theory. Chapter 3 mapped the mechanistic paradigm as a network of fifteen interlocking assumptions stabilised by three reinforcing feedback loops that produce a triple-locked architecture (Sections 3.1–3.2). Chapters 4 and 5 executed the reconstructive cascade – the coordinated sequence of FROM→TO transformations through which ontological reconstruction constrains epistemological possibilities, which in turn constrain methodological options (Section 1.5.1) – yielding a reconceptualisation of the firm as a historically evolving, open, multi-level process-pattern whose properties emerge at the meso-level where neither micro-level reduction nor macro-level aggregation can reach (Section 5.4). That reconstruction is presupposed here.

This chapter does not revisit the arguments developed in the preceding chapters, nor does it introduce new theoretical structures. Its role is evaluative and positioning rather than constructive: it clarifies what the dissertation has achieved, the level at which it intervenes, the criteria by which it should be assessed, and the boundaries it deliberately maintains. It then traces the implications that follow from the reconstruction within those boundaries and identifies the research programmes the reconstruction enables. The separation between reconstruction and evaluation preserves the integrity of the analytical work while allowing examiners and readers to assess the dissertation's contribution without conflating explanation with advocacy. It also reflects the nature of the contribution itself: the dissertation intervenes at the level of foundational assumptions – the paradigmatic commitments that structure what firms are taken to be, what can be known about them, and what forms of inquiry are treated as legitimate – rather than proposing a model or testing hypotheses within an inherited framework.

Paradigmatic contributions of this kind are not assessed primarily through empirical verification or predictive accuracy. Section 2.6.1 established six criteria appropriate to pure theoretical work: internal coherence, problem-solving capacity, theoretical fruitfulness, illumination of mechanisms, interestingness, and comprehensiveness, supplemented by pattern-level intelligibility as the appropriate predictive analogue for complex phenomena. Section 2.6.2 identified four criteria that, although legitimate in other genres of inquiry, constitute category errors when applied to paradigm-level reconstruction: point-predictive accuracy, direct empirical verification, naïve falsifiability, and quantitative precision as a universal standard. Section 2.6.3 developed Lakatosian programme appraisal as the meta-theoretical framework within which both sets of criteria operate: the relevant question is whether the programme resolves anomalies without proliferating ad hoc repairs, generates novel insight, and opens productive inquiry – whether it is progressive rather than degenerating. These are the terms on which this chapter invites the dissertation to be judged.

The chapter proceeds in five sections, ordered to establish the dissertation's contribution before specifying its boundaries, and its boundaries before tracing what the contribution enables. Section 6.1 evaluates the dissertation's contribution – its nature, level, originality, and the

specific achievements of the diagnostic, reconstructive, and integrative phases. Section 6.2 specifies what the thesis does not claim, establishing non-claims that function not as defensive hedges but as conditions of coherence: the thesis is not a unified rival theory, does not provide empirical validation, does not claim ontological emergence of its own ideas, does not reject formal methods, does not constitute a managerial manual, and is not a revolutionary break with existing scholarship. Section 6.3 identifies the conditions under which the dissertation's claims are valid and the vulnerabilities it acknowledges – the cascade condition on which internal coherence depends, the evaluative standards appropriate to the genre, and specific points at which the argument is open to challenge, including the operationalisation gap between philosophical clarification and empirical instrumentarium, the predominantly Western provenance of the theoretical traditions on which it draws, and the institutional conditions under which a CAS-founded programme must operate within a discipline governed by the paradigm it challenges. Section 6.4 traces the implications of the reconceptualisation across six domains – economic theory of the firm, methodology and standards of rigour, management and strategic practice, economic policy and institutional analysis, philosophy and methodology of economics, and the additional urgency created by developments in artificial intelligence – each derived analytically from the dissertation's specific achievements rather than asserted independently. Section 6.5 identifies the research programmes the reconstruction enables – the specific theoretical, empirical, methodological, and comparative investigations that become coherently investigable once the foundational reconstruction is accepted.

The progression from evaluation through delimitation, limitation, implication, and future research reflects the logic of the contribution itself. What the thesis achieves must be specified before its boundaries can be drawn; its boundaries must be drawn before its implications can be traced without overclaiming; and its implications must be established before the research programmes they enable can be identified with precision. Each section presupposes and builds upon its predecessors, and the chapter as a whole completes the evaluative stage that the problematisation methodology requires (Section 2.3).

6.1 Evaluation of the Dissertation's Contribution

6.1.1 Nature and Level of Contribution

This dissertation's contribution must be evaluated in terms of both its nature (what kind of contribution it makes) and its level (where in the architecture of inquiry it intervenes). Misidentifying either leads to inappropriate expectations: if the work is treated as model development, it will be judged by predictive performance; if it is treated as empirical research, it will be judged by data and estimation. Neither standard is appropriate to what the thesis actually attempts.

Nature of the contribution

The dissertation advances a contribution best characterised as paradigm reconstruction rather than model construction, empirical testing, or incremental refinement. It is transformative rather than additive. It does not introduce a new variable, a new optimisation problem, or a new empirical regularity. Nor does it attempt to replace existing theories of the firm with a single rival theory operating on the same plane. Instead, it reconstructs the foundational commitments – ontological and epistemological – that determine (i) what firms are taken to be, (ii) what kinds of knowledge about firms are treated as admissible, and (iii) what standards of rigour are implicitly enforced by those commitments.

Paradigm reconstruction is a distinct scholarly genre. It is not defined by rhetorical ambition, but by the stratum of intervention: the work targets the assumptions that make theories possible and render them intelligible, rather than extending theories within those assumptions. Because it operates at this level, it is not assessed primarily by empirical fit or point prediction. It is assessed by whether it makes the underlying architecture explicit, whether it resolves persistent tensions that remain stable across refinements, and whether it enables disciplined, fruitful inquiry that would otherwise remain blocked.

Level of the contribution

The dissertation is positioned as pure theory and meta-theoretical analysis. Its object is not firm behaviour in a particular industry or period, but the conceptual architecture through which firms are theorised. This positioning explains, and justifies, the absence of an empirical chapter. The claims at stake concern the assumptions that *structure* empirical inquiry—what counts as evidence, what forms of explanation are treated as legitimate, and what the modelling enterprise is implicitly committed to in advance. At this level, empirical testing is not “missing”; it is not yet admissible without circularity, because the thesis is reconstructing the conditions under which testing becomes meaningful.

Root-cause identification as originality

A central contribution of the dissertation is the identification of root generative sources of persistent difficulty in firm theory. Rather than treating failures as local problems – missing variables, incomplete modelling, or imperfect behavioural assumptions – the thesis argues that recurring limitations reflect a deeper mismatch between a mechanistic analytical orientation and the properties firms exhibit under uncertainty, change, and interdependence. Chapter 3 substantiates this argument through systematic reconstruction of the mechanistic paradigm as a network of fifteen interlocking assumptions organised across a Lakatosian hard core and protective belt, stabilised by three reinforcing feedback loops that produce a triple-locked

architecture. This diagnostic apparatus explains why piecemeal reform repeatedly fails: challenging any single assumption while the others remain intact results in the preserved hard core regenerating the challenged derivative or absorbing the modification as a special case (Section 3.2). The originality claim is not that the thesis discovers new empirical phenomena, but that it makes visible – through a specific, traceable analytical mechanism – why familiar failures recur and why successive refinements repeatedly stabilise rather than resolve them.

Epistemological emergence, not invention

The dissertation does not claim to invent an entirely new conceptual vocabulary. Its contribution is epistemologically emergent: it generates new understanding through disciplined reconstruction and integration of insights already present across economics and neighbouring disciplines – critical realism, complexity science, economic methodology, philosophy of science, and organisation theory – while re-specifying the assumptions that govern how those insights can be coherently combined. The novelty is therefore architectural: a reorganisation of the conceptual terrain, rather than an assertion of theoretical creation *ex nihilo* (Section 2.9.1, Part D).

In sum, the dissertation aims to change not what economists and organisational scholars can say about firms within an inherited framework, but the conditions under which such claims are formed, evaluated, and revised. The following subsection specifies what this reconstruction has concretely achieved; Section 6.1.3 then makes explicit the criteria appropriate to evaluating such a contribution.

6.1.2 What the Dissertation Has Achieved

The preceding subsection characterised the nature and level of the contribution. This subsection specifies what the dissertation has concretely accomplished. The Introduction (Section 1.5.5) identified three interrelated contributions; Chapter 2 (Section 2.9.1, Part B) decomposed these into five connected outputs – three answering the subsidiary research questions (SQ1–SQ5) directly, and two addressing procedural and programmatic requirements of the reconstruction – which culminate to the answer of the main research question. With the constructive work now complete, this subsection confirms their delivery and locates each within the dissertation's argumentative structure. The contributions are presented in the order established in Section 2.9.1 for cross-referential consistency.

Contribution 1 (answering SQ1): Identification and mapping of the assumptive architecture of mainstream firm theory.

The dissertation made explicit the background assumptions that are often implicit, dispersed, or treated as methodological conveniences. Chapter 3 reconstructed this architecture as a network of fifteen interlocking assumptions: five ontological commitments forming the Lakatosian hard core (reductionism, determinism, substance ontology, closure, atomism), four epistemological commitments (certainty, universal laws, monism, perfect information), and six derivative assumptions constituting the protective belt (equilibrium, linearity, homogeneity, ergodicity, optimisation, decomposability). The architecture's most consequential structural effect – the systematic elimination of the meso-level at which routines, capabilities, and coordination architectures operate – was shown to be generated by the joint operation of reductionism, atomism, and representative-agent aggregation (Section 3.1).

Contribution 2 (answering SQ2): Diagnosis of why piecemeal reform fails.

Beyond identifying the assumptions, the dissertation demonstrated that they form a mutually reinforcing network – the Triple Lock – that neutralises isolated challenges. Three feedback loops (ontological, epistemic, and structural) interlock at high-centrality nodes, creating redundant stabilisation: challenge to any single assumption activates compensatory defence across the others, so that the paradigm behaves as a closed attractor absorbing perturbations and returning to its stable configuration (Section 3.2). This diagnostic contribution explains why decades of refinement – adding behavioural realism, institutional context, or evolutionary language – have improved local explanatory performance while leaving the overall picture architecturally fragmented. It further explains why the 2008 financial crisis, despite exposing multiple simultaneous failures of the paradigm's foundational commitments, did not trigger paradigm change (Section 3.4).

Contribution 3: A coherent bridge between complexity vocabulary and economic/organisational vocabularies.

Across the literatures engaged, key terms frequently shift meaning – equilibrium, emergence, adaptation, and uncertainty carry different senses in complexity science, mainstream economics, and organisation theory. The dissertation established disciplined translation practices that reduce equivocation and improve conceptual coherence across traditions (Section 2.8.6). This contribution is procedural rather than substantive, but it is a precondition for the integrative work performed in Chapters 4 and 5.

Contribution 4 (answering SQ3–SQ5): Reconstruction of the ontological, epistemological, and methodological foundations required for firms to be coherently understood as complex adaptive systems.

This is the dissertation's central constructive achievement. Chapters 4 and 5 executed a coordinated reconstruction cascade in which ontological commitments constrain epistemological possibilities, which in turn constrain methodological options. Chapter 4 established the negative case: why the inherited ontology – reductionism, substance ontology, and atomism – cannot accommodate the defining properties of firms as complex adaptive systems (emergence, relational constitution, path dependence, adaptive transformation). Chapter 5 delivered the positive reconstruction through systematic FROM→TO transformations across all three domains.

Ontologically: emergent wholeness replaces reductionism (Section 5.1); process ontology replaces substance (Section 5.2); relational primacy replaces atomism (Section 5.4); thermodynamic openness replaces closure (Section 5.6); and the meso-level is restored as the privileged analytical scale (Section 5.5). The result specifies what firms must be assumed to be: historically evolving, open, multi-level process-patterns – relationally constituted configurations whose identity consists in self-reproducing patterns of activity rather than fixed essence, whose boundaries function as selective membranes, whose capabilities emerge from relational configurations, and whose evolution follows path-dependent trajectories (Sections 5.1–5.7).

Epistemologically: fundamental uncertainty is recognised as an ontological condition rather than an information deficit; pattern-based navigation replaces prediction-centred knowledge; and scaffolded rationality replaces unbounded optimisation (Section 5.8.1). These commitments follow from the ontological reconstruction: if firms generate genuine novelty

through interaction under non-ergodic conditions, then prediction-centred epistemology – premised on stable regularities within a knowable distribution – loses its ontological warrant.

Methodologically: structured pluralism replaces monism, governed by admissibility axioms, positioning criteria, and context-specific validation standards that provide rigour without reimposing the closure that the ontological reconstruction rejected (Section 5.8.2). This is not methodological permissiveness but its opposite: each method must satisfy explicit admissibility conditions tied to the level of analysis and the CAS properties at stake.

Contribution 5: Firms reconceptualised as Complex Adaptive Systems

Firms should be reconceptualised as historically evolving, thermodynamically open, multi-level process-patterns – relationally constituted configurations whose identity consists in self-reproducing patterns of coordinated activity rather than fixed essence, whose boundaries function as selective membranes regulating flows of energy, matter, and information, whose capabilities emerge irreducibly from relational configurations, and whose evolution follows path-dependent trajectories through expanding possibility spaces under fundamental uncertainty. They are meso-level phenomena: emergent organisational wholes that exist above individuals but below the economy, possessing causal powers – routines, culture, adaptive capacity, strategic orientation – that belong to the whole and cannot be derived from, decomposed into, or aggregated from the properties of their components. Their stability is not static persistence but dynamic achievement, maintained through continuous self-reproduction far from equilibrium; their knowledge is distributed, partial, and perspectival, enabling pattern-based navigation rather than point prediction; their rationality is procedural and scaffolded by relational and institutional context rather than globally optimising; and their history is constitutive rather than incidental, accumulating as an irreversible trajectory that shapes what futures remain accessible and which are foreclosed.

This reconceptualisation transforms the nature of questions themselves. Strategy becomes navigation through evolving possibility spaces rather than optimisation over known alternatives. Competitive advantage becomes adaptive capacity rather than a defensible position. Organisational design becomes the cultivation of interaction architectures from which desired capabilities can emerge rather than the specification of optimal structures. And what the inherited framework treats as anomalies or residuals — emergence, downward causation, path-dependence, creative destruction, adaptive tension — are revealed as the central causal mechanisms through which firms come into being, persist, transform, and dissolve.

Contribution 6: A basis for subsequent research programmes.

The dissertation identified conceptual openings and methodological implications that can guide subsequent empirical, applied, and theoretical work. It specifies enabling conditions and researchable problem-spaces – empirical investigation of meso-level phenomena, operationalisation of structured pluralism, case-based research examining firms as process-patterns – without claiming implementation (Section 6.5). This final contribution maintains the dissertation's register as foundational reconstruction rather than applied design.

The cascade structure of the argument

These five contributions are not independent outputs. They form a reconstructive cascade whose coherence is structural: just as the mechanistic assumptions mutually require and stabilise one another (the Triple Lock), the CAS-compatible assumptions form a mutually

reinforcing architecture in which emergence requires relational constitution, process ontology necessitates thermodynamic openness, fundamental uncertainty follows from emergence and non-ergodicity, and structured pluralism follows from irreducible multi-scale dynamics. The package stands or falls as a system, as the Chapter 5 conclusion confirms: one cannot coherently adopt CAS ontology while retaining mechanistic epistemology, nor adopt expansionist methodology while presupposing equilibrium dynamics (Sections 1.5.3, 5.8–5.9).

The meso-level as a distinctive analytical contribution

One feature of the reconstruction merits separate emphasis because it distinguishes this dissertation's positioning from existing work in Complexity Economics. The intermediate organisational scale – where routines, capabilities, and coordination architectures constitute the firm's distinctive properties – is systematically eliminated by the joint operation of mechanistic assumptions (the "Missing Middle" diagnosed in Section 3.1). Complexity Economics has reconstructed the micro-to-macro pathway at the level of markets and economies; this dissertation attends to precisely the meso-level that macro-focused complexity work has traversed rather than theorised (Section 1.5.2). The restoration of the meso-level as analytically indispensable – as the scale at which firm-level emergence is empirically observable and causally operative – is among the dissertation's most distinctive contributions (Section 5.5).

A note on unevenness of elaboration

It should be acknowledged that the epistemological and methodological reconstructions developed in Section 5.8, while rigorously derived from the ontological platform, are necessarily less fully elaborated than the ontological reconstruction occupying the bulk of Chapters 4 and 5. This reflects the dissertation's scope priorities and the logical ordering of the cascade: the downstream domains could not be developed in comparable detail without the prior ontological work being complete. Section 6.5 identifies further development of these epistemological and methodological implications as a priority for subsequent research.

6.1.3 Criteria for Evaluation

Given the nature and level of contribution established above, this dissertation should be evaluated against criteria appropriate to meta-theoretical, paradigm-level work. Chapter 2 (Section 2.6.1) articulated six criteria for evaluating pure theory: internal coherence, problem-solving capacity, theoretical fruitfulness, illumination of mechanisms, interestingness, and comprehensiveness. With the constructive work now complete, these criteria remain fully applicable. Two additional criteria – ontological–epistemological alignment and epistemic discipline – proved especially salient as the reconstruction proceeded, reflecting the particular demands of work that intervenes across interdependent philosophical domains. The eight criteria are presented together below.

1. Internal coherence.

Does the reconstructed position hold together as a consistent whole? In particular, do the ontological commitments adopted (what firms are) align with the epistemological standards defended (what can be known about them and with what limits) and the methodological principles advanced (how inquiry should proceed), without contradiction or ad hoc repair? The cascade logic developed in Chapter 5 – in which ontological reconstruction constrains epistemological possibilities, which constrain methodological options – is the thesis's primary coherence mechanism. As the Chapter 5 conclusion states, the package stands or falls as a

system: one cannot coherently adopt CAS ontology while retaining mechanistic epistemology, nor adopt expansionist methodology while presupposing equilibrium dynamics.

2. Problem-solving capacity.

Does the reconstruction address persistent difficulties that remain stable across otherwise diverse theories of the firm – especially those concerning adaptation, innovation, organisational heterogeneity, and breakdown under uncertainty – without merely renaming them? The diagnostic contribution (Chapter 3) identifies these difficulties as structural consequences of mechanistic assumptions; the reconstructive contribution (Chapters 4–5) must demonstrate that they become intelligible rather than residual under the alternative framework. The examiner's question is whether the problems diagnosed as structural become tractable under the reconstruction without ad hoc repairs.

3. Generativity and Theoretical fruitfulness.

Does the reconstructed framework open disciplined lines of inquiry that are difficult to formulate or pursue under an inherited mechanistic orientation? Fruitful theories expand the space of meaningful inquiry rather than merely redescribing what is already known. The examiner's test is whether the framework generates a disciplined agenda of researchable questions – for example, how interaction patterns generate firm-level capabilities irreducible to individual traits, what forms of organisational stability are compatible with continuous adaptation, how historical sequences constrain feasible strategic paths, and which meso-level structures transmit effects between micro actions and macro outcomes – rather than remaining at the level of metaphor.

4. Illumination of mechanisms.

Does the reconstruction render visible generative processes that mechanistic frameworks suppress or distort? Mechanism illumination goes beyond labelling outcomes or reporting correlations; it discloses identifiable processes such as feedback loops, amplification, lock-in, coordination dynamics, and emergence. The CAS framework should illuminate why organisational collapse reflects path-dependent lock-in and capability rigidity rather than merely "bad decisions"; why systemic crises reflect risk amplification through network interconnection rather than isolated regulatory failures; and why persistent firm heterogeneity reflects self-reinforcing, historically contingent capability trajectories rather than static resource differences. The examiner's test is whether the reconstruction allows the articulation of plausible generative accounts – mechanism sketches that are explanatory without pretending to deliver point prediction.

5. Interestingness.

Does the reconstruction challenge taken-for-granted assumptions in a way that is both significant and intelligible? As a meta-theoretical contribution operating through problematisation, the dissertation must interrupt routinised assumptions rather than merely extend them. The examiner's test is whether the work reliably shifts the informed reader from "obvious" to "rethink" – by rendering visible what the mechanistic framework presupposes and obscures – without crossing into implausibility.

6. Comprehensiveness.

Does the reconstruction address the full set of issues that must be addressed for reconceptualisation to be stable and meaningful? Partial reconstruction – changing labels while leaving the foundational architecture intact – produces instability and invites reabsorption into the dominant paradigm. The examiner's test is whether the reconstruction answers the diagnostic task at the appropriate level of generality, across all three domains (ontology, epistemology, methodology), without leaving core issues untreated.

7. Ontological–epistemological alignment.

Are the standards of knowledge and rigour defended in the thesis appropriate to the kind of object being theorised? If firms are treated as open, evolving, and internally organised systems, are prediction-centred, optimisation-first expectations defended, bracketed, or delimited in a principled manner? This criterion is especially important for work that reconstructs across interdependent philosophical domains, because misalignment between ontological commitments and epistemic standards – treating a processual, emergent object as if it were amenable to the epistemic standards designed for closed, decomposable systems – would constitute a form of internal incoherence that criterion 1 alone might not capture.

8. Epistemic discipline and restraint.

Does the dissertation clearly state what it does not claim, the scope conditions under which its claims hold, and the ways in which future empirical work would need to be designed to avoid re-importing the very assumptions the thesis problematises? This criterion reflects the particular risk that paradigm-level work faces: ambitious scope can encourage claims broader than the argument can sustain. As argued in Section 2.9.1 (Part D), calibrated claims are stronger claims – distinguishing what the argument establishes from what it does not is a quality-control device, not a retreat from ambition.

A further evaluative standard, appropriate to the complexity conditions under which this dissertation operates, concerns pattern-level intelligibility. For complex phenomena, the appropriate epistemic aim is not point prediction but the capacity to recognise, classify, and explain recurrent patterns – such as adaptation, lock-in, and emergence – under general conditions (Hayek, 1967; Section 2.6.1). This standard supplements the eight criteria above and informs especially how criteria 3 (theoretical fruitfulness) and 4 (illumination of mechanisms) should be applied.

These criteria do not demand finality. They ask whether the dissertation offers a coherent and productive departure: whether the problems it targets are real, whether the reconstruction is disciplined, and whether it enables a progressive research programme rather than a rhetorical alternative.

6.2 What the Thesis Does Not Claim

Section 6.1 clarified the nature and level of the dissertation's contribution and the criteria appropriate to evaluating it. This section completes that framing by delimiting the object of evaluation: it specifies what the dissertation does not attempt to do, so that it is not judged against expectations it neither adopts nor could satisfy without contradicting its own commitments. These delimitations were established as principled design choices in the

Introduction (Section 1.5.6) and Methodology (Sections 2.9.1, Part C; 2.9.2, Part A); they are restated here so that the evaluation in subsequent sections proceeds on appropriate terms.

Explicit delimitation is particularly important for a dissertation that operates at the level of foundational reconstruction. Without it, the work risks being misread either as an encyclopaedic aggregation of ideas or as an overreaching claim to theoretical unification. The boundaries below are therefore not defensive hedges; they are conditions of coherence for the kind of contribution the thesis makes.

6.2.1 It Does Not Propose a Unified "New Theory of the Firm", due to its paradigmatic nature

The dissertation does not claim to have produced a complete, unified theory of the firm in the sense of a single substantive model that subsumes existing approaches. It does not offer a super-theory designed to replace transaction cost, resource-based, agency, behavioural, evolutionary, or institutional perspectives.

Instead, it operates at a prior level: it reconstructs the ontological, epistemological, and methodological conditions under which theories of the firm can be formulated, compared, and evaluated once firms are treated as complex adaptive systems (Section 1.5.6). The dissertation's aim is therefore architectural rather than totalising: it clarifies why fragmentation persists, identifies the assumptions that generate recurrent limitations, and re-specifies the conceptual commitments required for more adequate theorising.

This does not entail rejection of existing traditions. The dissertation treats them as historically situated achievements containing genuine insight, while arguing that their limits often reflect inherited assumptions whose general applicability to complex organisational phenomena was never systematically examined. The thesis therefore seeks to enable re-positioning and extension, not dismissal. Nor does it treat any single heterodox tradition – critical realism, complexity science, or evolutionary economics – as exhaustive; it draws on multiple streams while maintaining disciplined translation practices between them (Section 2.8.6).

6.2.2 It Does Not Claim Empirical Validation in This Dissertation", due to its paradigmatic nature

The dissertation does not claim empirical validation of its reconstructed framework. It does not introduce new firm-level datasets, estimate models, or test hypotheses intended for immediate econometric adjudication.

This follows from the level at which the dissertation intervenes (Section 2.5). The work targets paradigm-constitutive assumptions – commitments that shape what counts as evidence, what variables are regarded as admissible, and what forms of explanation are treated as legitimate. Such assumptions cannot be "tested" by the evidence they structure without circularity. Where the object of inquiry is itself in question – what the firm is taken to be, what kind of system it is – conceptual reconstruction must precede the design of empirical tests. For the same reason, the framework cannot be "falsified" by isolated counter-instances; paradigm-level propositions are evaluated as research programmes, not as individual hypotheses (Section 2.6.2).

Where empirical and historical material appears in the argument – as in the analysis of the 2008 financial crisis (Section 3.4) or firm-level illustrations in Chapter 4 – its epistemic status is illustrative and diagnostic, not evidential in the population-generalisation sense (Section 2.9.1, Part C). Such material reveals what mechanistic framings systematically suppress; it does not constitute verification of the alternative.

The dissertation therefore provides enabling conditions for future empirical, historical, and organisational research. It clarifies what kinds of empirical questions become meaningful under a complexity-consistent conceptualisation of the firm and what kinds of evidence and methods would be required to investigate emergence, path dependence, and adaptive transformation without re-importing mechanistic defaults (Sections 2.5, 2.9.2). The conditions under which the dissertation's claims are valid – and the standards by which they should be assessed – are addressed in Section 6.4.

6.2.3 It Does Not Reject Formal Modelling, Optimisation, or Prediction Per Se

The dissertation does not argue that formal modelling, optimisation, or predictive methods are illegitimate or should be abandoned. It rejects neither mathematics nor formal reasoning, and it does not treat complexity as an argument for analytical looseness.

Its claim is conditional: the validity of these methods depends on whether their enabling assumptions hold. Within firms there are domains that are complicated rather than complex (Section 2.1.5): decomposable subsystems operating within relatively stable solution spaces, where optimisation and prediction may be appropriate and necessary. The dissertation does not dispute this.

What it disputes is the elevation of optimisation and point prediction into universal standards for the firm as a whole, and for those dimensions of firm behaviour that exhibit complex adaptive properties – non-linearity, feedback, path dependence, endogenous novelty, and fundamental uncertainty. In such contexts, point prediction and global optimisation can become category errors rather than tools (Section 2.6.2). The dissertation therefore repositions formal methods rather than dismissing them: treating them as locally powerful where their conditions hold, and limited where they do not. Where point prediction is inappropriate, pattern recognition, navigation, and historically grounded explanation constitute the appropriate epistemic ideals (Section 1.5.6).

6.2.4 It Does Not Offer a Managerial Manual or Policy Blueprint, due to its meta-theoretical approach

The dissertation is not a managerial handbook and does not present itself as a ready-to-apply policy framework. It does not provide implementation protocols, operational checklists, or prescriptive rules (Section 2.9.1, Part C).

This follows from the nature of the contribution. Managerial tools and policy designs presuppose a conceptual foundation regarding what firms are and how they behave – precisely what the dissertation reconstructs. The thesis provides conceptual resources with practical significance (for diagnosis and orientation), but translation into operational tools, organisational diagnostics, or policy design is positioned as subsequent work requiring further development and empirical grounding.

6.2.5 It Does Not Present Itself as a Revolutionary Rejection of Prior Work

The dissertation does not claim that existing economics or firm theory is "wrong" in a wholesale sense, nor does it position complexity as a replacement doctrine that renders prior scholarship obsolete.

Its stance is transformative, not revolutionary (Section 2.9.1, Part D). Existing theories are treated as analytically valuable, while their recurrent limitations are interpreted as consequences of inherited assumptions whose scope conditions were rarely made explicit. The dissertation therefore aims to strengthen the foundations beneath accumulated insight – so that what has been learned can be retained, reinterpreted, and extended under assumptions better aligned with the properties firms exhibit under uncertainty, change, and interdependence.

Taken together, these five delimitations define what the dissertation asks to be evaluated: not a unified substantive theory, not an empirically verified model, not a claim of conceptual invention, not a rejection of formal methods, not a manual for practitioners, not a polemical dismissal of prior work, and not a reconstruction of economics as a whole – but a meta-theoretical reconstruction, confined to the firm at the meso-level, that seeks to enable more adequate theorising and more disciplined future inquiry.

6.3 Limitations and Conditions of Validity

Section 6.2 specified what the dissertation does not claim. This section addresses a different question: given what the dissertation does claim, under what conditions are those claims valid, what standards of evaluation are appropriate, and what genuine vulnerabilities does the argument acknowledge? The distinction matters. Non-claims delimit scope; conditions of validity specify the assumptions on which the argument's force depends and identify the points at which it is most open to challenge. Following Cilliers (2005b, p. 612), the articulation of limits is not a defensive exercise but a constitutive requirement of inquiry under complexity: "the fact that our knowledge is limited is not a disaster, it is a condition for knowledge. Limits enable knowledge."

The section proceeds in four movements. The first identifies the structural condition on which the framework's internal coherence depends. The second specifies the standards on which the thesis's validity rests, the scope within which those standards apply, and why alternative standards would constitute category errors. The third acknowledges vulnerabilities that an honest assessment of the argument must confront. The fourth explains why these limitations are productive rather than damaging – why they open inquiry rather than undermine it.

6.3.1 The Cascade Condition

The dissertation's central architectural claim is that ontological, epistemological, and methodological reconstruction must proceed as a coordinated cascade: what firms are determines what can be known about them, and both jointly determine how they can legitimately be studied (Section 1.5.1). This cascade logic, developed across Chapters 4 and 5, is not an expository convenience but the argument's governing structure. The framework's coherence depends on the joint acceptance of its principal moves: emergent wholeness replacing reductionism (Section 5.1), process ontology replacing substance (Section 5.2), relational constitution replacing atomism (Section 5.4), thermodynamic openness replacing closure (Section 5.6), the meso-level restored as the privileged analytical scale (Section 5.5), and the epistemological and methodological consequences that follow necessarily from these ontological commitments (Sections 5.8–5.9).

This non-decomposability is both the framework's strength and its principal condition of validity. It is its strength because the mutual reinforcement of ontological, epistemological, and methodological commitments is what produces explanatory power that no single replacement could achieve in isolation – just as the mechanistic paradigm's power derives from the mutual reinforcement of its own assumptions across a triple-locked architecture (Section 3.2). It is its condition of validity because rejecting any principal move while accepting others would produce an internally inconsistent hybrid. The framework, like the paradigm it diagnoses, stands or falls as a system.

This raises a question an examiner may reasonably pose: does the framework degrade gracefully, or does it collapse entirely if one element is disputed? The answer requires distinguishing between the cascade's load-bearing entailments and its reinforcing connections. Some links are logically necessary. Section 5.8 establishes three such bridges: process ontology entails fundamental uncertainty, because a world of genuine becoming cannot be fully pre-determined (Section 5.8.1); emergence requires synthetic rather than purely analytical methods, because decomposition destroys the interaction-dependent properties it purports to explain (Section 5.8.2); and CAS ontology requires open-system structure, because closure

becomes logically incompatible with relational, processual reality (Section 5.8.3). At the level of the Lakatosian hard core, the dissertation identifies further necessary connections: emergence requires relational constitution; process ontology necessitates thermodynamic openness; fundamental uncertainty follows from emergence and non-ergodicity; structured pluralism follows from irreducible multi-scale dynamics (Section 1.5.5; Abstract). These are not elective associations. Accepting emergence while retaining reductionist methodology would be incoherent, because the properties emergence generates are precisely those that decomposition eliminates.

Other connections, while strongly motivated, admit degrees of acceptance – not at the level of principle but at the level of specific formulation. One could accept that CAS ontology requires open-system structure (the load-bearing principle established in Section 5.8.3) while remaining agnostic about the specific characterisation of openness as selective-membrane dynamics, or accept that capabilities are relationally constituted while questioning particular claims about how boundary permeability operates in specific institutional contexts. What admits variation is not whether the firm is open or relational – the bridges establish that it must be – but how that openness or relationality is best elaborated in particular analytical settings. The reconstruction is therefore more resilient than a simple chain: it is a network of mutually supporting commitments whose load-bearing connections are logically necessary and whose reinforcing elaborations admit refinement without structural collapse. But it is less resilient than a modular system in which components can be independently evaluated. The conditionality is not a weakness but a structural feature of non-decomposable argumentation – it reflects the same ontological insight, that properties of wholes cannot be recovered from properties of parts, that the dissertation advances as its central claim about firms.

The dissertation makes this dependency explicit rather than concealing it behind claims of neutrality. The assumptions under which its conclusions hold are stated, not smuggled. This is what distinguishes principled conditionality from fragility.

6.3.2 The Standards on Which Validity Rests

Because the dissertation is positioned as pure theory operating at the level of paradigm rather than model (Section 2.5), the standards by which its claims should be evaluated differ categorically from those appropriate to empirical hypothesis testing. Section 2.6 established this evaluative framework in detail; the present subsection restates its implications as conditions of validity.

The thesis's claims are valid insofar as the reconstructed framework satisfies six criteria articulated in Section 2.6.1: internal coherence across ontological, epistemological, and methodological commitments; problem-solving capacity relative to the mechanistic paradigm's documented limitations; theoretical fruitfulness in generating research questions that were previously incoherent; illumination of generative mechanisms that mechanistic assumptions suppress; interestingness in the sense of rendering visible what was previously treated as background; and comprehensiveness in addressing the full scope of the diagnostic task. Section 6.1 evaluated the dissertation's performance against each of these criteria, supplemented by two additional standards – ontological–epistemological alignment and epistemic discipline – that emerged as especially salient through the reconstruction itself.

Conversely, four standards that are legitimate in other genres of inquiry constitute category errors when applied to this work (Section 2.6.2): predictive accuracy in the point-prediction sense, replaced by pattern-level intelligibility as the appropriate epistemic ideal under non-

ergodicity; direct empirical verification of foundational commitments, replaced by evaluation through explanatory reach and generative capacity; naïve falsifiability through isolated counter-instances, replaced by Lakatosian programme appraisal; and quantitative precision as a universal badge of scientific adequacy, replaced by conceptual precision – clarity of constructs, disciplined distinctions, and explicit scope conditions. These replacements are not weaker standards. They are different standards, appropriate to the level at which the dissertation intervenes.

The circularity condition explains why. Because the thesis interrogates pre-empirical commitments – the ontological, epistemological, and methodological assumptions that determine what counts as evidence, what variables are admissible, and what forms of explanation are treated as legitimate – empirical testing conducted within inherited frameworks would presuppose the very assumptions under reconstruction (Section 2.5). The circularity is categorical, not practical: it would require validating pre-empirical commitments through methods whose admissibility those commitments themselves determine (Section 2.9.2, Part E). Empirical inquiry is enabled by the reconstruction; it cannot logically precede it (Section 2.9.2, Part C). The validity of the thesis therefore rests on logical coherence, architectural consistency, explanatory adequacy, and generative capacity – not on statistical fit, predictive performance, or the decisiveness of any single illustration.

These standards apply within a defined domain. The reconceptualisation applies to organisational forms that exhibit the defining properties of complex adaptive systems: non-decomposability, adaptive behaviour, historical path dependence, relational constitution, and operation under fundamental uncertainty (Section 1.5.6). Where these properties do not hold – in purely mechanical subsystems, short-horizon optimisation problems with stable solution spaces, or contexts in which interaction effects and organisational learning are negligible – mechanistic or optimisation-based models may retain local and conditional validity (Section 2.1.5). Where prediction is possible and meaningful within such complicated rather than complex subsystems, the framework does not prohibit it; it denies only the elevation of point prediction into a universal criterion of scientific adequacy (Section 2.6.2). Similarly, while the framework's implications extend toward markets, institutions, and policy (Section 6.4), the claims are strongest at the firm level and become enabling rather than definitive as one moves outward to higher levels of aggregation.

6.3.3 Acknowledged Vulnerabilities

Beyond the cascade condition and evaluative framework, five vulnerabilities require honest acknowledgment. The first four were identified in the Methodology's self-assessment (Section 2.9.2, Part B); the fifth concerns the framework's relationship to its own analytical discipline. They are restated here because the Conclusions chapter is where a dissertation's settled understanding of its own limitations should reside.

The evidential status of firm-level illustrations. The dissertation uses firm-level material throughout Chapters 3–5 – Toyota's production system, Nokia's trajectory, IDEO's emergent capabilities, Silicon Valley's ecosystem dynamics, the 2008 financial crisis – to render its ontological claims intelligible and to demonstrate what the mechanistic paradigm systematically obscures. The epistemic status of this material is illustrative and diagnostic, not evidential in the population-generalisation sense (Section 2.9.2, B.1; Section 2.9.1, Part C). A potential vulnerability lies in interpretive contestation: alternative readings may argue that mechanistic explanations could be extended to account for the same phenomena through auxiliary adjustments – protective-belt modifications in Lakatosian terms – rather than

paradigmatic reconstruction. The dissertation mitigates this risk through careful linguistic discipline (consistently employing "illustrates" and "renders visible" rather than "proves" or "confirms"), through contrastive framing that identifies what each case makes visible under CAS assumptions that remains invisible under mechanistic ones, and through the selection of structurally revealing cases rather than convenient anecdotes. But the possibility of competing interpretations cannot be eliminated within a pure theory project, nor should it be. What would constitute a decisive challenge is not an alternative reading of any single illustration, but a demonstration that the mechanistic paradigm can accommodate the full range of phenomena the thesis identifies – emergence, path dependence, relational constitution, adaptive tension – without the ad hoc proliferation that signals degenerating rather than progressive problem-shifts (Section 2.6.3).

The pluralism–eclecticism boundary. Structured pluralism is a central methodological consequence of the reconstruction (Section 5.8.2). It is disciplined by admissibility axioms, positioning criteria, transformation protocols, and context-specific validation standards that distinguish it from undisciplined eclecticism. However, the dissertation establishes these disciplining mechanisms at the level of principle; it does not fully resolve how adjudication proceeds in cases of deep theoretical incompatibility – cases where two methods, both admissible under CAS assumptions, generate conflicting accounts of the same phenomenon and no higher-order criterion determines precedence. This is not a flaw unique to the present work; it is a standing challenge for all pluralist research programmes (Section 2.9.2, B.2). The dissertation identifies the architectural constraints that prevent pluralism from collapsing into permissiveness; the operational refinement of adjudication criteria in applied and empirical contexts is positioned as subsequent work (Section 6.5.3).

The operationalisation gap. The dissertation clarifies the conceptual roles of CAS-relevant constructs – emergence, relational constitution, path dependence, adaptive tension, meso-level causation – and establishes the ontological and epistemological conditions under which they become analytically coherent. It does not provide a complete empirical instrumentarium for measuring these constructs (Section 2.9.2, B.3). Translating philosophical clarity into empirical operationalisation is itself a substantial research programme: some constructs – network topology, path dependence – are more tractable than others – emergence, relational constitution – and measurement strategies will necessarily be context-dependent. The gap is real, but it is a consequence of the thesis's positioning as foundational reconstruction rather than applied methodology. A thesis that attempted both would risk doing neither adequately. The specific empirical programmes this gap enables are identified in Section 6.5.4.

Institutional and cultural boundaries. The argument draws predominantly on Western institutional contexts and theoretical traditions – mainstream Anglo-American economics, European critical realism, Santa Fe Institute complexity science – and its firm-level illustrations are drawn largely from North American, European, and East Asian corporate settings (Section 2.9.2, B.4). Two distinct questions arise, and they should not be conflated. The first is diagnostic: does the fifteen-assumption architecture (Section 3.1) and the Triple Lock mechanism (Section 3.2) apply with equal force to economic traditions that developed under different intellectual genealogies? The mechanistic paradigm as the thesis characterises it has a specific historical trajectory rooted in the Newtonian inheritance (Section 3.1); economic traditions shaped by different philosophical foundations – developmental-state frameworks, post-Soviet transition economics, or traditions grounded in non-Western ontologies – may exhibit different assumptive architectures requiring separate diagnostic work. This is properly a scope condition: the diagnosis targets the mechanistic paradigm specifically, and its

applicability to other paradigmatic traditions is an open question rather than a settled result. The second question is reconstructive: does the CAS-compatible ontology developed in Chapters 4–5 require modification when applied to firms embedded in fundamentally different institutional orders? Because the reconstruction's ontological claims – emergence, process, relationality, openness – are advanced as general properties of complex adaptive systems rather than as culturally specific descriptions, there is reason to expect broad applicability. But whether specific formulations – the selective-membrane model of boundaries, the particular characterisation of meso-level structures – require adaptation across institutional contexts remains to be investigated. The dissertation achieves depth by accepting limits on breadth. Cross-cultural and comparative-institutional extension, addressing both the diagnostic and reconstructive dimensions, is a genuine research frontier (Section 6.5.6).

The analytical discipline condition. The framework's validity depends on maintaining the rigour with which emergence, synthesis, and complexity are treated throughout the thesis. The invocation of complexity language alone does not constitute explanation (Section 2.1.5). This is not merely a warning against careless extension by others; it is a condition that the thesis itself must satisfy. The reconceptualisation claims analytical force because it specifies what becomes visible under the CAS frame, what becomes less central, and how this reframing reshapes the questions that firm theory can coherently ask (Section 2.1.5). If complexity were treated as a rhetorical gesture rather than an analytical commitment – if "emergence" functioned as a label rather than a mechanism, or "path dependence" as a narrative convenience rather than an ontological claim – the framework would lose its explanatory power. The validity of the thesis rests on the substantive content of its ontological, epistemological, and methodological reconstructions, not on the vocabulary in which they are expressed.

6.3.4 Why These Limitations Are Productive & Necessary Given the Paradigmatic Nature of the Thesis

The limitations identified above do not weaken the dissertation's contribution; they specify the conditions under which it can be extended without category error (Section 2.9.2, Part E).

The cascade condition identifies the minimum commitment required for the framework to function – and in doing so, it provides a clear test: a reader who accepts the ontological reconstruction but disputes the epistemological derivation knows precisely where the disagreement lies, which of the three bridges (Section 5.8) is contested, and what would need to be shown to resolve it. The evaluative framework specifies the standards by which the thesis invites assessment, preventing category-error objections while maintaining full accountability under criteria appropriate to its genre. The circularity condition explains why the thesis cannot be its own empirical test and thereby identifies what the next stage of the research programme must be: not empirical testing of foundational assumptions, but empirical inquiry conducted from within the reconstructed framework, investigating whether CAS-consistent methods reveal patterns that mechanistic methods systematically miss.

The five acknowledged vulnerabilities convert similarly. The illustrative-evidence boundary identifies the precise form subsequent empirical work should take: not replication of the thesis's illustrations, but systematic investigation using the methods the reconstruction licenses – longitudinal process tracing, network analysis, comparative institutional studies – applied to the phenomena the reconstruction renders visible. The pluralism–eclecticism boundary identifies the most pressing methodological challenge: developing operational adjudication criteria within the architectural constraints the thesis provides. The operationalisation gap

identifies the bridge that must be built between philosophical reconstruction and empirical programme: translating conceptual constructs into context-sensitive measurement strategies. The institutional boundary identifies the geographical and intellectual frontier across which the framework's generality requires testing. The analytical discipline condition identifies the standing requirement that all subsequent work employing the framework must meet: substantive specification rather than rhetorical invocation.

Each limitation, in other words, converts into a specific research challenge rather than an unresolved deficiency. This is not coincidental. It follows from the thesis's own epistemological commitments. If knowledge of complex systems is necessarily perspectival, provisional, and bounded (Section 2.4), then a framework that claimed to resolve all questions it raises would contradict its own epistemic stance – just as a firm theory that claimed to eliminate uncertainty would contradict the ontology of emergence and non-ergodicity on which it rests. The productive character of limitations is therefore not a concession but a consequence: a CAS-compatible framework that acknowledged no open boundaries would, by its own criteria, be less adequate than one that did.

This returns the argument to the cascade condition with which this section began. The framework's non-decomposability – the very feature that makes it vulnerable to coordinated rejection – is also what gives its limitations their productive character: because the commitments are mutually reinforcing, each boundary opened by a limitation is constrained by the architecture as a whole, ensuring that subsequent work extends the framework rather than dissolving it. The contribution should be evaluated not by whether it achieves closure – no paradigm-level reconstruction could, and the thesis's own epistemology explains why – but by whether it raises the right questions, specifies the conditions under which they can be coherently pursued, and provides sufficient architectural discipline to ensure that subsequent work does not simply reproduce the difficulties the reconstruction was designed to overcome.

As Cilliers (2005b, p. 613) observes, the aim is not defensive limitation but epistemic humility: recognising that the absence of certain claims is a condition of credibility rather than a deficiency. Together with Section 6.2's specification of what the dissertation does not claim, this section completes the thesis's account of its own boundaries – non-claims delimiting scope from one direction, conditions of validity delimiting it from another. The framework is neither totalising nor final. It is precise about where it applies, explicit about the standards by which it should be judged, honest about what could go wrong, and clear about the commitments on which it rests. Within these conditions, the next section traces what the reconceptualisation makes visible – the implications it enables across six domains of inquiry.

6.4 Implications Enabled by the Reconceptualisation of firms as CAS

Section 6.3 established the conditions under which the dissertation's claims hold, the standards by which they should be assessed, and the vulnerabilities the argument acknowledges. Within those conditions, the implications traced in this section follow analytically from the dissertation's specific achievements: the diagnostic exposure of the mechanistic paradigm's assumptive architecture and triple-locked persistence (Chapter 3), and the coordinated reconstruction of the firm's ontological, epistemological, and methodological foundations through the FROM→TO transformation framework (Chapters 4–5).

The question therefore shifts from "what does the framework assert?" to "what becomes visible and intelligible once firms are treated as complex adaptive systems?" The point is not to replace one set of results with another, but to re-specify the conditions under which results can be meaningfully produced and interpreted – within economics and across the organisational and managerial disciplines that implicitly depend on a theory of the firm.

This section proceeds in three movements. The first establishes the general structure of what the reconceptualisation makes visible – the architecture of the visibility shift that runs through every domain-specific consequence. The second traces the implications of that shift across six domains: economic theory of the firm, methodology and standards of rigour, management and strategic practice, economic policy and institutional analysis, the philosophy and methodology of economics, and the additional urgency created by developments in artificial intelligence. Each domain-specific subsection is an application of the general argument, not an independent treatment. The third returns to what the reconceptualisation opens when the discipline's self-imposed constraints are made visible and can be consciously relaxed.

The General Argument: What Becomes Visible

The deepest implication of the reconceptualisation is not any specific ontological replacement. It is the act of making the invisible medium visible.

The mechanistic assumptions that the dissertation surfaces – reductionism, substance ontology, atomism, closure, determinism, and their epistemological and methodological derivatives – do not function as explicit commitments that economists consciously adopt and could consciously revise. They function as the medium through which economic reasoning about firms is conducted: the background conditions that determine what questions can be formulated, what counts as evidence, what forms of explanation are treated as legitimate, and what phenomena register as real. One does not typically argue for substance ontology; one thinks through it without noticing. One does not defend reductionism as a philosophical position; one builds models whose mathematical structure presupposes it. The assumptions are the water the fish does not see.

Once that medium is surfaced, two consequences follow simultaneously. First, what the paradigm has made possible becomes intelligible as achievement rather than necessity – the specific explanatory power of equilibrium analysis, optimisation, and formal modelling is preserved, but its scope conditions are made explicit. Second, what the paradigm has structurally excluded becomes visible for the first time as exclusion rather than absence. What could not be seen was not absent from reality; it was absent from the apparatus. The reconceptualisation does not discover new phenomena. It reveals that familiar phenomena – adaptation, innovation, organisational heterogeneity, crisis, capability formation and decay – were structurally invisible to the inherited framework, not because they were unimportant, but

because the framework's foundational assumptions defined them out of existence before analysis began.

The term "visibility shift" captures this combined consequence of the diagnostic and reconstructive work accomplished in Chapters 3–5. It does not introduce a new theoretical claim; it names the structure of what the thesis's specific achievements, taken together, make possible. Six dimensions of this visibility shift pervade the dissertation and generate the domain-specific implications that follow.

The first concerns the causal locus that was erased. The meso-level – where routines, capabilities, coordination architectures, and network structures constitute the firm's distinctive properties – is the scale at which the firm's most consequential dynamics are produced. The joint operation of reductionism, atomism, and representative-agent aggregation systematically eliminates this scale (Section 3.1). Reductionism collapses it downward to individual choice; aggregation collapses it upward to macro outcomes. The meso-level vanishes – not because it was investigated and found irrelevant, but because the analytical apparatus cannot represent it. The consequence is that the dominant explanatory strategies in economics explain from the wrong level: they locate causation in individuals (leadership, talent, entrepreneurial vision) or in macro forces (market selection, competitive pressure, technological change) while the causal work is being done by the interaction architectures that mediate between them.

The second concerns the ontological status of stability. Under substance ontology, the firm is a thing that persists and occasionally undergoes change. Change is deviation from a stable state. The replacement of substance with process ontology (Section 5.2) reverses the explanatory burden: the firm is a pattern of activity that must be continuously reproduced, and stability is achievement rather than default. What this makes visible is that the daily work of sustaining coordination, maintaining relational architectures, and reproducing routines – work that appears as mere "operations" under substance ontology – is the constitutive activity through which the firm exists at all. When that reproductive process degrades, the firm does not "break" like a thing; the pattern dissolves. Failure becomes intelligible not as the breakdown of an otherwise stable object but as the cessation of a generative process.

The third concerns where properties reside. Under atomism, firms are independent entities that enter into relations; properties are intrinsic. The replacement of atomism with relational constitution (Section 5.4) reveals that what a firm is – its capabilities, identity, competitive position – is partly produced through relational configurations that extend beyond its formal boundaries. This makes visible why a whole class of analytical and managerial practices systematically fail: when capabilities are treated as possessed assets rather than emergent achievements of specific relational architectures, interventions that disrupt those architectures – restructuring, acquisition, "best-practice" transplantation – destroy the very properties they seek to capture or replicate.

The fourth concerns the space of possibilities. Under ergodicity, where you have been does not determine where you can go; the possibility space is given. Under non-ergodicity and path dependence (Sections 5.2, 5.9), history is constitutive: the firm's accumulated trajectory produces the possibility space within which it operates. Two firms in ostensibly the same industry may face radically different possibility spaces because their trajectories differ. This makes visible why "strategic choice" under fundamental uncertainty is not selection among pre-existing options but navigation through a space that is partly created by the act of navigating – and why "best practices" are not transferable algorithms but outcomes of specific, non-reproducible paths.

The fifth concerns a hidden trade-off. Equilibrium analysis cannot represent endogenous fragility because it treats disturbances as exogenous and responses as convergent. The network architecture and feedback analyses (Sections 5.7, 5.9) make visible that the features which produce efficiency under stable conditions – tight coupling, eliminated redundancy, optimised resource allocation – are precisely the features that produce fragility under stress. Standard performance metrics reward the systematic production of invisible fragility. This is not a risk that can be captured by adding a risk variable, because the fragility is architectural: it inheres in the system's topology and becomes apparent only when conditions change – which is precisely the possibility that equilibrium assumptions exclude.

The sixth concerns the relationship between explanation and prediction. The mechanistic paradigm assumes symmetry: if you can explain, you can predict, and vice versa. Under CAS assumptions, this symmetry breaks (Section 5.8.1). You can explain why organisational collapse followed a particular trajectory – through identifiable generative mechanisms – without being able to have predicted that it would occur, because the system generates genuine novelty under non-ergodic conditions. This means that the phenomena that matter most for understanding firms – innovation, transformation, crisis, capability emergence – become amenable to rigorous explanatory analysis freed from the requirement that explanation entail prediction. Pattern recognition, mechanism identification, and historically grounded accounts become the appropriate epistemic ideals, not impoverished substitutes for a prediction that was never available (Section 2.6.1).

These six dimensions are not independent. They form a mutually reinforcing visibility shift, just as the assumptions they surface form a mutually reinforcing architecture. Erasing the meso-level, treating stability as default, locating properties in entities, assuming a given possibility space, ignoring endogenous fragility, and demanding prediction-explanation symmetry – these are not separate analytical choices. They are facets of a single ontological commitment, inherited from mechanism, that pre-determines what the apparatus can see. The reconceptualisation makes that commitment visible as a commitment – contingent, historically inherited, and revisable – rather than as the natural order of inquiry.

The remainder of this section traces what this visibility shift means across six domains. The selection is illustrative rather than exhaustive, in both the domains addressed and the implications drawn within each – a demonstration of the kind of inquiry the reconceptualisation opens, not a delimitation of it.

6.4.1 Implications for Economic Theory of the Firm

What the theory of the firm seeks to explain

The theory of the firm addresses a persistent cluster of foundational questions: why firms exist rather than all production being coordinated through markets; what determines firm boundaries; how firms coordinate, decide, learn, and adapt; and why durable differences in performance, structure, and innovative capacity persist across firms.

Since Coase, these questions have generated major research programmes – including Williamson's transaction cost economics, agency approaches, the resource-based tradition associated with Penrose and Barney, evolutionary approaches associated with Nelson and Winter, and institutional traditions in multiple variants. Yet the field remains characterised by a familiar difficulty: multiple powerful partial lenses without an agreed architecture that relates them. Firms appear alternately as production functions, contract nexuses, incentive systems,

capability bundles, or routine-based evolving organisations. Each tradition illuminates something real – yet their relationships remain under-specified, and "integration" is often sought through add-ons rather than through a shared conceptual architecture. Recent strategic-management scholarship has recognised this integration problem explicitly, including the synthesis calls advanced by Leiblein and Reuer, and the fragmentation diagnosis highlighted by Drnevich, Mahoney, and Schendel.

The visibility shift established above transforms how this landscape is understood.

Two structural implications

First, the unit of theorising shifts from "the firm as a thing" to "the firm as an organised process." The replacement of substance ontology with process ontology (Section 5.2) implies that the firm cannot be adequately represented as a timeless mapping, a static container of resources, or a passive nexus of contracts. Under a complexity-consistent ontology, the firm is better conceived as an historically evolving, organised pattern of coordination that must be continually reproduced through routines, interaction, and learning – what the dissertation characterises as a "historically evolving, open, multi-level process-pattern" (Section 5.4). This shift does not "invalidate" production, contracts, incentives, or resources. It repositions them as partial descriptions of a system whose defining properties arise from interaction, feedback, and temporally extended adaptation.

Second, integration becomes an architectural question. The standard response to fragmentation is to "combine" theories – often by adding variables, appending institutional qualifiers, or stacking mechanisms. The dissertation's diagnostic argument implies something stricter: theories become comparable and connectable only once their background assumptions and scope conditions are made explicit (Section 3.1). The problem is not merely that theories are specialised; it is that they often rely on different implicit ontologies of what a firm is and different epistemic ideals of what counts as explanation. The fifteen-assumption architecture identified in Section 3.1 provides the diagnostic instrument for making these differences visible, while the reconstructive cascade (Sections 5.1–5.8) specifies what a shared alternative architecture requires.

How core questions are reframed

With these two implications in place, several classic firm-theory questions are reframed – not by adding complexity as a qualifier, but because specific FROM→TO transformations alter the terms in which the questions are posed.

(a) Why firms exist: beyond economising on transaction costs. Transaction cost reasoning retains its insight: asset specificity and dependency create governance problems, and governance structures respond to transaction characteristics. But under fundamental uncertainty and endogenous novelty (Section 5.9), the question "why firms exist" cannot be exhausted by economising on contracting costs within already-defined possibility spaces. The replacement of closure with thermodynamic openness (Section 5.6), and of atomism with relational constitution (Section 5.4), implies that firms also function as organised systems for generating, maintaining, and recombining capabilities through time – activities that require sustained coordination, learning, and the stabilisation of routines under uncertainty. More precisely: because the possibility space within which transactions occur is itself endogenously generated through organisational processes, the question bifurcates. Governance and capability

generation become co-equal explananda, and a theory that addresses only the former necessarily underdetermines the phenomenon.

(b) What firm boundaries are: not only a legal line, but a coordination achievement. Boundary questions ("make or buy") are typically treated as categorical decisions explained by governance costs, incentive alignment, or resource ownership. The replacement of closure with thermodynamic openness, and the reconceptualisation of boundaries as selective membranes (Section 5.6), implies a more processual interpretation: boundaries are stabilised patterns of coupling and decoupling – how tightly the firm integrates certain relations, how it governs interfaces, and how it maintains identity while remaining open to exchanges that enable learning and co-evolution. The question ceases to be "where should the boundary be?" and becomes "what relational architecture produces the capabilities the firm requires, and how should different zones of coupling intensity be managed as a continuous adaptive challenge?"

(c) Why imitation and capability transfer fail: not a mystery, but a structural consequence. The persistent difficulty of copying "successful firms" becomes more intelligible once emergence is recognised as ontological (Section 5.1) and relational constitution is accepted as foundational (Section 5.4). If capabilities are not possessed objects but emergent organisational achievements dependent on specific relational and temporal configurations, then copying visible components will often fail. Inimitability is not a mystical property of "resources"; it is a consequence of configuration dependence and history dependence – precisely the kind of mechanism that substance ontology and decomposability assumptions render invisible. Similarly, value destruction in mergers becomes intelligible: what is "acquired" is the legal entity and its assets; the capability that made the entity valuable may dissolve when the relational architecture that produced it is disrupted.

(d) The "missing middle" becomes central rather than residual. The dissertation's diagnostic claim that firm-level phenomena are systematically squeezed between micro-individualism and macro-aggregation (Section 3.1.4), together with the restoration of the meso-level as the privileged analytical scale (Section 5.5), implies that firm theory requires explicit attention to routines, capabilities, organisational architecture, networks, and institutional scaffolding as causally efficacious structures. Once emergence and path dependence are treated as constitutive, the meso-level is not a convenient compromise. It is where the firm exists as a theoretical object.

Heterogeneity, innovation, and temporality

The reconceptualisation also reframes three phenomena that the mechanistic paradigm persistently struggles with – each a direct consequence of the visibility shift established above.

Heterogeneity as an expected outcome, not an anomaly. Under mechanistic assumptions – particularly homogeneity and equilibrium (Section 3.1.3) – persistent heterogeneity is difficult to justify. If firms face similar optimisation problems and competitive selection pushes them toward equilibrium, then durable differences appear as residuals. The visibility shift reverses this inference. If history is constitutive, if learning is path-dependent, if interaction effects are non-linear, and if novelty is endogenously generated (Section 5.9), heterogeneity becomes the expected outcome. Firms do not traverse identical decision landscapes; they co-create and reshape their own possibility spaces through cumulative choices and relational entanglements. Variation across firms ceases to be treated as noise around a representative benchmark and becomes a primary object of explanation.

Innovation as constitutive, not exogenous. In many orthodox representations, innovation enters as a disturbance: an external shock shifting parameters and initiating adjustment toward a new steady state. Section 5.9 established that novelty generation is endogenous – arising from emergence, creative destruction through process succession, and path-dependent recombination. A firm's innovative capacity is inseparable from its history and organisational architecture. Innovation is therefore not merely "something that happens to firms," but something that firms do – through being organised in particular ways. The paradigm's inability to accommodate genuine novelty – the emergence of possibilities not contained in the prior state – is perhaps the most consequential blindness the visibility shift corrects.

Temporality as constitutive, not parametric. The replacement of substance ontology with process ontology (Section 5.2) entails that time becomes constitutive: firms carry histories, learn through irreversibilities, and encounter futures that are not statistical shadows of the past. Firms operate across multiple temporal horizons simultaneously – operational execution, strategic positioning, evolutionary transformation, and institutional-historical embedding – and different mechanisms become determinative at different horizons. A theory that is strong at one horizon can become misleading when applied to another – not because it is "wrong," but because it is being asked to explain patterns governed by different temporal mechanisms.

From rivalry among theories to scope-aware repositioning

The dissertation's diagnostic and reconstructive work together imply that much of the fragmentation of firm theory does not stem from intellectual failure, but from the fact that different traditions highlight different slices of a multi-level, historically evolving object. Transaction cost reasoning illuminates governance and hold-up problems; the resource-based tradition highlights capability-based performance differences; evolutionary approaches illuminate routines, search, and selection; institutional perspectives illuminate how organisational forms are enabled, constrained, and stabilised by field-level structures.

The implication is not that one lens must defeat the others, but that their explanatory reach is conditional: each framework is informative under certain ontological and epistemic presuppositions and becomes strained when extended beyond them. Treating the firm as a complex adaptive system therefore encourages scope-aware repositioning of these traditions: preserving their genuine mechanisms, while making explicit their background assumptions and the boundaries of their claims (Section 2.8.6). The diagnostic apparatus identifies where each tradition operates within relatively decomposable, stable sub-problems – where formal optimisation may remain useful – and where the object is inherently processual, historically contingent, and non-decomposable, requiring different forms of explanation.

Transaction cost economics, for example, retains its explanatory power for governance problems within bounded contracting environments, but its scope conditions – specifiable alternatives, calculable costs, given possibility spaces – become visible as conditions rather than as the universal character of the problem. The resource-based view's insight that configuration matters is preserved, but "resources" are reinterpreted as stabilised process-patterns rather than possessed assets (Section 5.4), which explains both the tradition's insight and its persistent operationalisation difficulty: the object is processual, not static, and methods designed to capture static properties will systematically miss it. Integration, if it is to be more than eclecticism, requires this kind of scope discipline.

6.4.2 Implications for Methodology and Standards of Rigour

The reconceptualisation advanced in Chapters 2–5 carries methodological implications that operate at two levels: what the thesis demonstrates through its own inferential practice, and what follows for the discipline once the reconstructive cascade is accepted. The central claim is not that rigour should be relaxed, but that rigour must be made ontologically adequate (Section 2.6.2): standards appropriate to complicated, closed, decomposable systems cannot serve as universal benchmarks for open, emergent, historically constituted systems.

The visibility shift has a direct methodological consequence: it reveals that the current hierarchy of methods in economics is not a neutral quality-control mechanism but a paradigm-reproducing institution. The mechanistic paradigm generates the methods; the methods generate the data; the data confirm the paradigm's assumptions – because the methods are designed to capture only what the assumptions say exists. Meso-level dynamics, relational constitution, path dependence, and emergence leave traces in the data that mechanistic methods produce, but these traces appear as residuals, unexplained variance, or "qualitative" material positioned as inferior. The visibility shift reveals that this hierarchy of evidence is itself an artefact of ontological commitments disguised as methodological standards.

Structured pluralism as methodological consequence

If firms are ontologically emergent and epistemically characterised by fundamental uncertainty, then methodological monism becomes incoherent. This is not a preference for pluralism as diplomatic concession; it is a necessity imposed by the object of inquiry. The dissertation develops structured pluralism – disciplined by admissibility axioms, positioning criteria, and transformation protocols (Section 5.8.2) – as the methodological architecture that preserves rigour without reimposing monism. Ashby's Law of Requisite Variety formalises the principle: explanatory variety must match phenomenal variety. Because emergence generates irreducible properties at each level of analysis, no single model can span micro behaviour, meso interaction patterns, and macro outcomes simultaneously.

The implication for the field is substantial. The current evaluation of firm-theoretical research overwhelmingly privileges methodological uniformity – econometric estimation, optimisation-derived prediction, or formal model-building – as markers of rigour. The reconceptualisation implies that rigour is better assessed by whether the method is appropriate to the object at the level of analysis under investigation. Formal modelling and quantitative precision retain their force where the complicated/complex distinction (Section 2.1.5) indicates that enabling conditions approximately hold. Where they do not – where emergence, feedback, path dependence, and fundamental uncertainty are constitutive – insistence on these methods becomes a category error rather than a standard of excellence.

The relationship between qualitative and quantitative methods is accordingly respecified. Under mechanistic assumptions, qualitative work is often treated as a preliminary to quantitative work – useful for hypothesis generation but insufficient for "rigorous" testing. Under CAS assumptions, this hierarchy dissolves: case-based, processual, and network methods are ontologically necessary for the properties they study. They are not second-best alternatives to econometric estimation but the appropriate instruments for emergence, relational constitution, and path dependence.

Prediction reconceived

Once the complicated/complex distinction is taken seriously, prediction-centred evaluation loses its status as a universal benchmark. In complicated systems, precise prediction can be demanded because the enabling conditions – closure, stable parameters, decomposability – approximately hold. In complex adaptive systems, prediction is limited in principle: emergence, non-linearity, and historical contingency restrict inquiry to pattern understanding, conditional expectation, and domain-bounded anticipation rather than point forecasting (Section 2.6.2). Pattern recognition, navigation, and historically grounded explanation constitute the appropriate epistemic ideals (Section 1.5.6).

The implication is not that theory is excused from accountability. It is that accountability must be expressed through standards appropriate to the object: whether the theory captures the relevant causal tendencies, whether it identifies mechanisms plausibly capable of generating observed patterns, and whether it specifies the boundary conditions under which its claims apply. The symmetry between explanation and prediction that the mechanistic paradigm assumes – if you can explain, you can predict – breaks under CAS conditions. Explanation becomes richer than prediction: understanding the generative mechanisms that produce a class of outcomes is more analytically powerful than forecasting specific instances, because it reveals why things happen rather than merely that they happen. This is not a retreat from rigour but its reconstitution for objects that generate genuine novelty.

What the thesis demonstrates through its own practice

The dissertation does not merely argue for methodological reconstruction; it demonstrates reconstructed practice across seven dimensions. First, it performs problematisation rather than gap-spotting (Section 2.3), showing that persistent limitations require structural rather than local critique. Second, it performs systematic FROM→TO reconstruction (Section 2.3.6), replacing inherited assumptions with coherent alternatives rather than merely relaxing constraints. Third, it performs integrative synthesis under constraint (Section 2.8.6), drawing from multiple traditions while maintaining compatibility at the level of foundational commitments. Fourth, it adopts retroductive inference, reasoning from observed explanatory patterns toward underlying generative structures. Fifth, it restores the meso-level as analytically indispensable (Section 5.5). Sixth, it treats axiological transparency as a component of rigour (Section 2.2.5). Seventh, it treats inquiry as reflexive rather than externally neutral (Section 2.2.2), rendering pre-understandings available to critique rather than embedding them silently.

These are not idiosyncratic choices. Each follows from the reconstructive cascade: once ontology is processual and emergent, epistemology must accommodate fundamental uncertainty, which means methodology must be pluralist and reflexive. The implication for the field is that any inquiry into firms as complex adaptive systems will face the same constraints and require the same methodological reorientation – not because the thesis prescribes it, but because the object demands it. The persistent failures of firm theory – capability non-transferability, merger value destruction, organisational rigidity under turbulence – arise not from missing variables or inadequate optimisation techniques, but from a coherent yet restrictive background orientation that treats firms as bounded entities, locates explanatory power at the node level, and compresses time into synchronic states (Section 3.2). Partial reforms fail because the preserved hard core regenerates challenged derivatives or absorbs modifications as special cases. This is why the dissertation pursues coordinated reconstruction rather than incremental extension.

6.4.3 Implications for Management and Strategic Practice

The foundational shift

If the firm is reconceptualised as a complex adaptive system – as a process-pattern reproduced through ongoing activity, embedded in relational configurations, and operating under fundamental uncertainty – then the implicit image of management embedded in mechanistic firm theory must be revised. The point is not to derive a "new management doctrine," but to state what becomes intelligible about managerial action once the visibility shift is applied to organisational life. Three consequences follow directly from specific transformations accomplished in Chapters 4–5.

From controlling to enabling. Under substance ontology, management is naturally framed as control: resources are deployed, capabilities are directed, outcomes are produced. The replacement of substance with process ontology (Section 5.2) implies that what appears as "capability" is continuously reproduced through patterned activity and interaction. Managerial action is therefore better conceptualised as the maintenance (and sometimes redesign) of the conditions under which those patterns can be sustained and adapted. Where control presupposes stable properties that persist independently of their reproduction, enabling recognises that stability is an achievement – maintained rather than possessed.

From optimisation to navigation. Under genuine uncertainty – where the possibility space is not fully specifiable, and where decisions help constitute the terrain through which later decisions must move – optimisation loses its defining conditions (Section 5.8.1). This does not imply irrationality. It implies a different rationality appropriate to open systems: action that preserves adaptive capacity, learns through feedback, and commits selectively as uncertainty resolves. Calculation becomes local and conditional: appropriate where closure and stability approximate, unreliable where novelty, feedback, and open-ended change dominate.

From possession-language to reproduction-language. The replacement of substance ontology with process ontology reinterprets many managerial descriptions. Statements such as "the firm has a capability," "the firm owns a competence," or "the firm possesses culture" risk reifying what is in fact reproduced. The reconceptualisation redirects attention toward the activity-patterns and interaction-structures through which those "possessions" persist, decay, or transform.

These shifts are interdependent: they follow jointly from the reconstructive cascade. If capabilities are reproduced, management is enabling rather than controlling. If the future is not fully specifiable, decision-making is navigational rather than optimising. If the firm is open and relationally constituted, managerial action is necessarily attentive to interfaces and dependencies rather than confined to legal boundaries.

Why certain classes of intervention systematically fail

The visibility shift reveals that a whole class of managerial and organisational interventions fail not because of poor execution but because of ontological misidentification. The general mechanism is consistent: interventions designed for decomposable systems are applied to non-decomposable ones.

Mergers that treat capabilities as transferable assets destroy the relational configurations that produce those capabilities – the capability does not survive the disruption of its generative

architecture. Culture-change programmes that target individual attitudes miss the interaction patterns that generate organisational culture – culture is an emergent meso-level property, not an aggregate of individual beliefs. Restructurings that decompose for efficiency eliminate the redundancy that enables adaptive response – what appears as inefficiency under mechanistic metrics may be the slack through which the system maintains its capacity to reconfigure under stress. Best-practice transplantation that copies visible procedures without reproducing the trajectories that produced them fails because the practices are path-dependent outcomes, not context-free algorithms.

These are not implementation failures. They are predictable consequences of treating emergent, relational, path-dependent properties as if they were decomposable, possessable, and transferable. The visibility shift converts them from puzzling exceptions into a structurally intelligible class.

Resilience, fragility, and multi-scale causation

Chapters 4–5 imply that firm phenomena are irreducibly meso-level (Section 5.5): they cannot be explained as mere aggregation of individuals, nor as macro outcomes without a meso account of organisational configuration and emergence. This has a direct consequence for managerial reasoning: firm-level outcomes cannot be understood (or influenced) by attending exclusively to any single scale. A single-scale view systematically misattributes causation. Micro-only narratives overstate individual agency; macro-only narratives treat firms as passive recipients of environmental forces; meso-only narratives risk ignoring the institutional and competitive conditions that constitute the firm's possibility space.

The hidden trade-off between efficiency and fragility – the fifth dimension of the visibility shift – has particular force at the managerial level. If capabilities are emergent and path-dependent, then organisational resilience is not a fixed property but a dynamic capacity dependent on the firm's relational architecture, its redundancy, and its capacity for reconfiguration (Section 5.7). Fragility, conversely, is a structural condition – tight coupling, path-dependent lock-in, and reduced variety create brittleness that may be invisible under normal conditions and catastrophic under stress. Standard performance metrics, calibrated for stable equilibrium conditions, reward the systematic elimination of precisely the features – redundancy, loose coupling, exploratory slack, diversity of approaches – that provide resilience. The efficiency-fragility trade-off is invisible within the mechanistic paradigm because the paradigm's evaluation criteria are designed for conditions under which fragility does not manifest.

A further consequence concerns time. The reconceptualisation implies that firms operate across multiple temporal horizons simultaneously: routine reproduction, strategic development, and longer-run transformation under shifting conditions. Short-horizon performance signals can be misleading if read as indicators of long-horizon viability. In complex systems, delays, feedback, and path dependence can produce systematic gaps between present indicators and future consequences. The efficiency-fragility trade-off operates precisely through this temporal disjunction: what optimises short-horizon metrics may accumulate invisible fragility at longer horizons.

What this subsection does and does not claim

These implications are not offered as managerial prescriptions, nor as an applied "toolkit." They are consequences that follow from treating firms as complex adaptive systems: they specify what becomes intelligible about management once the mechanistic conception is abandoned. The practical relevance is real, but the thesis's contribution at this stage is conceptual: it clarifies the kinds of managerial questions that become coherent (and those that become category errors) under the reconceptualisation. Translation into operational guidance is positioned as subsequent work (Section 6.2.5).

6.4.4 Implications for Economic Policy and Institutional Analysis

The reconceptualisation has direct implications for how policy problems are framed, even where it does not generate policy blueprints. Under mechanistic assumptions, policy is naturally imagined as control: authorities identify levers (tax rates, interest rates, subsidies, rules), adjust parameters, and expect the system to converge toward intended outcomes. This image presupposes a relatively stable mapping from intervention to effect, and it encourages evaluation of policy by whether it delivers predicted results.

The visibility shift transforms this image. If the systems that policy addresses are open, non-decomposable, and path-dependent, then the mapping from intervention to effect is neither stable nor context-independent. The central shift is from outcome engineering toward institutional and architectural reasoning: policy is understood less as selecting targets and tuning parameters, and more as shaping the conditions under which adaptation, coordination, learning, and resilience occur. Five implications follow.

First, interventions become architecture-dependent. In open, non-decomposable systems, causal effects propagate through networks of interdependence and often return through feedback loops (Section 5.7). Outcomes therefore depend not only on the magnitude of an intervention but on the system's architecture: the structure of incentives, constraints, informational channels, and institutional complementarities through which responses are formed. A policy that appears "well-designed" in isolation can generate unintended consequences when it couples with existing feedback structures. This is not merely a practical warning; it follows analytically from the ontological commitments of openness, nonlinearity, and feedback established in Chapters 4–5.

Second, stability is reconceived as adaptive capacity rather than equilibrium enforcement. Under the mechanistic ideal, stability is associated with convergence: reducing deviation from target states, dampening fluctuations, and restoring a steady configuration. Under CAS conditions, stability is better interpreted as the capacity of institutions to maintain function under disturbance – absorbing shocks, reorganising without collapse, and learning in response to novelty (Section 5.6). The relevant policy question shifts from "how do we restore equilibrium?" to "what institutional forms enable resilience, learning, and coordination when equilibrium is neither stable nor reliably definable?"

Third, institutions become constitutive scaffolds rather than external constraints. The mechanistic image treats institutions as constraints on agents who exist prior to them. The reconstruction developed in Chapters 4–5 supports a different interpretation: institutions shape what problems are conceivable, what actions are feasible, and what forms of coordination

become possible (Section 5.4). They are not merely restrictions on behaviour but enabling structures – scaffolds that stabilise expectations, reduce uncertainty where it is reducible, and create the conditions for collective adaptation where optimisation is infeasible.

Fourth, competition policy is reframed by the abandonment of the representative firm. If firms are heterogeneous by construction – path-dependent, relationally constituted, and capability-differentiated (Sections 5.4, 5.9) – then competition cannot be modelled as interaction among identical agents converging toward equilibrium. The representative-firm assumption, which the dissertation identifies as a mechanistic derivative (Section 3.1), eliminates by construction the heterogeneity, interaction effects, and relational configurations that constitute competitive dynamics in actual markets. Competition is better understood as a co-evolutionary process among differentiated organisations embedded in relational architectures – with implications for how market power, entry barriers, and competitive dynamics are conceptualised.

Fifth, crisis preparedness becomes an architectural question. The dissertation's analysis of the 2008 financial crisis (Section 3.4) reveals that crises in interconnected systems reflect network-level dynamics – feedback amplification, contagion through coupling, cascading failure across interdependent nodes – rather than isolated failures of individual institutions or regulators. The policy implication is topological: crisis preparedness concerns the structure of interdependence – concentration of connectivity, tightness of coupling, presence or absence of buffers – not merely the capitalisation or regulation of individual nodes. This follows directly from the network architecture analysis (Section 5.7) and the feedback dynamics established in Section 5.9.

A further implication concerns temporality and irreversibility. If path dependence is constitutive, then policy interventions do not merely shift parameters within a stable system; they alter trajectories. The path not taken is not costlessly available later. Policy evaluation that assumes reversibility – that failed interventions can be corrected by returning to prior settings – presupposes ergodicity. Under non-ergodicity, policy choices create irreversibilities, and evaluation must include what futures are foreclosed, not just what targets are hit in the short term.

These implications do not yield prescriptive policy programmes within the scope of the dissertation (Section 6.2.5). Their contribution is to legitimise and clarify a different logic of policy reasoning: one grounded in system properties – openness, feedback, nonlinearity, path dependence, fundamental uncertainty – rather than in the control imagery inherited from mechanistic foundations. Policy evaluation correspondingly shifts from point-prediction and target-hitting toward criteria such as robustness across scenarios, institutional fit to context, resilience under disturbance, and the capacity to learn and adapt over time.

6.4.5 Implications for the Philosophy and Methodology of Economics

The visibility shift has implications not only for how firms are theorised but for how economics understands itself as a discipline – what kind of science it is, what standards of evaluation are appropriate, and how it relates to other forms of inquiry into complex social systems. The implications move from what economics has been assuming, through what its methods have been importing, to what standards and disciplinary relationships follow from making these commitments explicit.

The ontological turn

The dissertation's most fundamental implication for economic methodology is that ontological questions cannot be avoided. The mechanistic paradigm's assumptions about what firms are – stable, decomposable, equilibrium-seeking entities – are not empirically derived descriptions but pre-empirical commitments that determine what counts as evidence, explanation, and rigour (Section 3.1). These commitments do much of the theoretical work before formal analysis begins, yet they are rarely made explicit within normal scientific practice. The thesis demonstrates that when they are surfaced and subjected to systematic examination, many of the discipline's persistent difficulties become intelligible as structural consequences of specific ontological choices rather than as unsolved empirical puzzles.

The implication is not that every economist must become a philosopher. It is that a discipline whose foundational commitments remain invisible to itself cannot diagnose the source of its own recurring limitations. The diagnostic apparatus developed in Chapter 3 – the fifteen-assumption network, the three reinforcing feedback loops, the Triple Lock mechanism – provides a concrete instrument for this kind of self-examination: a way to identify which assumptions generate which limitations, and why addressing them individually fails while the network remains intact.

The hidden ontology of formal methods

A related implication concerns the ontological content of mathematical formalism itself. The visibility shift reveals that formal methods in economics – equilibrium analysis, optimisation, econometric estimation – are not ontologically neutral tools. Each carries implicit commitments: equilibrium analysis presupposes closure and convergence; optimisation presupposes a knowable objective function within a specifiable possibility space; standard econometric estimation presupposes stable relationships between identifiable variables. These are not technical requirements; they are ontological commitments embedded in the mathematical apparatus. The choice of formal method pre-determines what can be found. If you model with equilibrium, you will find equilibrium or deviation from it. If you optimise, you will find optima or failure to reach them. What appears as methodological rigour is often ontological smuggling: importing assumptions about the nature of reality through the formalism while appearing to be merely "doing science."

The paradigm as a self-reproducing system

The thesis's diagnostic framework carries a reflexive implication. The Triple Lock mechanism (Section 3.2) explains why the mechanistic paradigm absorbs challenges rather than evolving: the assumptions form a mutually reinforcing network stabilised by semantic, methodological, and institutional feedback loops. But economics-as-institution is itself a complex adaptive system. It exhibits the very properties the thesis identifies as constitutive of CAS: network

effects (journal standards, hiring criteria, training curricula reinforce shared assumptions), positive feedback (researchers trained in the paradigm train the next generation in the same paradigm), and path dependence (institutional investment in existing methods creates switching costs that increase with each generation of accumulated work). The thesis's own framework therefore explains why the paradigm it critiques is so resilient – not because its content is correct, but because its architecture is self-stabilising.

This is not a sociological aside. It is a genuinely grounded implication: the thesis establishes that CAS exhibit lock-in, path dependence, and resistance to piecemeal reform; economics-as-institution is such a system; therefore paradigm change in economics cannot be achieved through individual papers or incremental adjustments – it requires architectural intervention of the kind the thesis attempts through coordinated reconstruction across ontology, epistemology, and methodology simultaneously.

From theoretical monism to structured pluralism

The Abstract states the implication directly: "In economic methodology, theoretical monism dissolves into structured pluralism under explicit scope conditions." This is not a concession to methodological permissiveness; it is a strict consequence of the ontological reconstruction. If firms generate irreducible properties at the meso-level (Section 5.5), if emergence produces qualitatively novel phenomena at each level of organisation (Section 5.1), and if these phenomena require different forms of explanation at different scales, then no single method or model can span the full range of firm-level phenomena without distortion.

Structured pluralism, as developed in Section 5.8.2, differs from eclecticism precisely because it is disciplined by admissibility conditions tied to the level of analysis and the CAS properties at stake. Different methods are not simply "added" according to researcher preference; they are positioned within an architecture that specifies their domain of applicability, the ontological conditions under which they are appropriate, and the translation protocols required when moving across levels.

Programme appraisal as evaluation standard

The thesis's Lakatosian self-appraisal (Section 2.6.3) carries a broader implication: it models how paradigm-level theoretical work can be evaluated without collapsing into either naïve falsificationism or relativistic paradigm sociology. The relevant question for evaluating foundational frameworks is not whether they survive isolated empirical tests (which are paradigm-dependent), but whether they exhibit progressive character – resolving anomalies without proliferating ad hoc repairs, generating novel insight, and opening productive inquiry. The four criteria identified as inappropriate for paradigm-level work – point-predictive accuracy, direct empirical verification, naïve falsifiability, and quantitative precision as a universal standard (Section 2.6.2) – are not rejected as evaluation criteria in general; they are identified as category errors when applied to foundational reconstruction.

Problematisation as methodological contribution

The dissertation's use of problematisation (Section 2.3) – generating theoretical contributions through systematic assumption-challenging rather than gap-spotting – has implications beyond the present work. Gap-spotting extends literatures while leaving foundational commitments intact; it is the normal mode of theoretical development within a stable paradigm. But when limitations arise from the assumptive configuration rather than from isolated premises, gap-spotting reinforces the very framework that generates the difficulties. The implication is that economics needs problematisation as a recognised methodological register – not as a replacement for empirical work, but as a complement that addresses the foundational level where empirical work cannot reach.

Cross-disciplinary reunification

A final implication concerns the relationship between economics and adjacent disciplines. The mechanistic paradigm maintains disciplinary boundaries that separate economics from organisation theory, sociology, network science, and complexity science – boundaries maintained partly by incompatible assumptions. Once the foundational assumptions are reconstructed on CAS-compatible grounds, these disciplines discover that they have been studying different aspects of the same kind of system with different vocabularies. The dissertation's translation practices (Contribution 3, Section 2.8.6) and structured pluralism (Section 5.8.2) provide the methodological architecture for disciplined cross-disciplinary synthesis: not the diplomatic "combination" of theories that currently passes for integration, but genuine synthesis governed by shared ontological commitments and explicit scope conditions.

6.4.6 Implications Amplified by Artificial Intelligence

The Introduction (Section 1.4) argued that developments in artificial intelligence amplify both the relevance and the urgency of the reconceptualisation. That argument carries forward as an implication with two dimensions, both grounded in the visibility shift.

First, AI enables complexity-informed approaches in ways that were previously computationally prohibitive. Pattern recognition under uncertainty, navigation through evolving possibility spaces, multi-scale analysis, and scenario-based reasoning – the epistemic ideals that the reconceptualisation establishes as appropriate for CAS inquiry – are precisely the capabilities that AI systems increasingly provide. For the first time, the insights of complex systems science can be made actionable for real-world decision-making about firms and organisations. The thesis's reconceptualisation therefore does not merely redefine what should be studied; it connects with emerging tools that make the redefined inquiry practically feasible.

Second, the widespread deployment of AI is accelerating the very complexity that renders the mechanistic framework inadequate. Organisational boundaries are becoming more fluid, coordination architectures more distributed, adaptation cycles shorter, and emergent dynamics more consequential. AI systems are themselves adaptive, non-linear, and generative of genuine novelty – they exhibit the properties that the mechanistic paradigm was designed to suppress. More fundamentally, AI systems increasingly constitute part of the relational architectures through which organisational capabilities emerge (Section 5.4). They participate in the interaction patterns through which meso-level properties are produced. If firm capabilities are emergent properties of relational configurations, then the introduction of AI agents into those configurations does not merely automate tasks; it changes the interaction architecture and therefore the emergent properties. Understanding what firms become when their constitutive processes are partly conducted by adaptive, non-linear systems requires the full CAS apparatus this dissertation develops.

The implication is that the reconceptualisation developed in this dissertation is not merely theoretically preferable but practically necessary: it offers a coherent framework for understanding firms as entities that must navigate, incorporate, and be transformed by technologies that are themselves emergent, adaptive, and fundamentally unpredictable. This does not constitute an additional argument for the thesis's validity – that must rest on the internal coherence and explanatory adequacy established in the preceding chapters – but it does establish its contemporary significance and the cost of continued reliance on a conceptual architecture that was never designed for the conditions it now faces.

What the Reconceptualisation Opens

The implications traced above are consequences of what the thesis has established. But there is a further implication that concerns not what the reconstruction achieves but what it releases: the liberation of economics from constraints it imposed on itself in order to preserve the coherence of a paradigm that was never designed for the phenomena it now confronts.

The mechanistic paradigm was built on the aspiration to achieve for economic systems what Newtonian mechanics achieved for physical systems – universal laws, precise prediction, mathematical formalisation, and deductive explanation. This aspiration was fruitful: it generated powerful analytical tools, disciplined formal reasoning, and explanatory achievements that this dissertation acknowledges and preserves within their scope conditions. But the aspiration also imposed constraints – on what could count as a question, what could

count as evidence, what could count as explanation, and what could count as rigour – that progressively narrowed the discipline's capacity to engage with the properties that economic systems most distinctively exhibit: genuine novelty, endogenous transformation, relational constitution, fundamental uncertainty, and historical irreversibility.

The thesis makes these constraints visible as constraints – contingent, historically inherited, and revisable – rather than as the natural order of inquiry. Once they are visible, the discipline is freed to become the science of a different kind of system: open, creative, historical, relational, emergent. Not a lesser science, constrained to imprecise claims about messy phenomena. A science whose object is arguably more demanding than the mechanical systems that inspired the paradigm – because it must grapple with systems that generate the genuinely new, that constitute themselves through their own history, and that are irreducible to any description simpler than themselves.

What opens is not a single alternative programme but a space of disciplined inquiry: an economics of novelty that can study where new possibilities come from, rather than allocating within given possibility spaces; a processual science of organisation that can study how coordination is sustained, how it degrades, and how it reconfigures, rather than treating organisation as a static input; a relational economics that can study how network configurations generate the properties conventionally attributed to individual firms, rather than analysing nodes in isolation; a multi-temporal analysis that can study how mechanisms operating at different speeds interact, rather than collapsing all dynamics onto a single equilibrium horizon; and an economics of fundamental uncertainty that can study how firms, institutions, and economies actually navigate genuinely uncertain environments – connecting the discipline to the traditions of Knight, Keynes, Hayek, and Shackle that the mechanistic paradigm marginalised, while grounding their insights in the ontological and methodological foundations this dissertation provides.

These are not speculative aspirations. Each corresponds to a specific ontological move accomplished in the thesis – the restoration of novelty, the adoption of process ontology, the acceptance of relational constitution, the recognition of multiple temporalities, the acknowledgement of fundamental uncertainty as ontological rather than epistemic – and each identifies a domain of inquiry that was structurally inaccessible under the mechanistic paradigm's foundational commitments. The thesis does not conduct this inquiry; it establishes the conditions under which it can be conducted without category error.

The reconceptualisation therefore does not close a research programme. It opens one. It specifies the foundations on which a more complex understanding of firms – and of the economic systems within which firms operate – can be progressively built. The test of the reconstruction is not whether it provides final answers, but whether it enables better questions – questions that the inherited framework could not formulate, that the phenomena demand, and that the discipline is now, for the first time, equipped to ask. The next section identifies the specific research programmes through which these questions can begin to be pursued.

6.5 Future Research Programmes

The preceding sections have established what the dissertation achieves (Section 6.1), what it does not claim (Section 6.2), the conditions under which its claims hold (Section 6.3), and the implications those claims carry across six domains (Section 6.4). This final section turns to the research programmes that the reconceptualisation enables – not as an inventory of possible investigations, but as a characterisation of what disciplined inquiry into firms looks like once the foundational reconstruction is taken seriously.

The distinction matters. A foundational reconstruction does not generate future research in the way that a specific theoretical model generates testable hypotheses. It transforms the conditions under which research is conducted: what counts as evidence, what methods are admissible, what questions are coherent, what forms of explanation are legitimate. The difference is between extending a programme and inhabiting a different one. The programmes identified below are not topics that the mechanistic paradigm overlooked and that a CAS framework can now address. They are investigations that become articulable for the first time once the ontological, epistemological, and methodological reconstruction accomplished in Chapters 4–5 is in place – investigations whose very formulation requires the apparatus the thesis provides.

Three features distinguish the research programme that follows from its mechanistic counterpart and structure the presentation of this section.

The first is cascade logic. The dissertation's central architectural claim – that ontological, epistemological, and methodological commitments must be reconstructed as a coordinated system (Section 1.5.1) – applies equally to the programme the reconstruction enables. Theoretical extensions, empirical investigations, and methodological innovations are not independent activities that can be pursued in isolation and assembled afterward. Each constrains and enables the others. The empirical programme requires methods that do not yet fully exist; the methodological programme requires theoretical specificity about what is to be investigated; theoretical extensions require empirical feedback to determine where the framework's generative power is greatest. The six subsections that follow are presented sequentially, but they operate as a network in which progress on any front conditions progress on others.

The second is generative tension. Every foundational reconstruction creates tensions that cannot be dissolved but that drive inquiry forward precisely because they resist final resolution. The reconstruction developed in this thesis produces several such tensions: between foundational principles and empirical purchase, between non-decomposability and analytical tractability, between universal ontological commitments and radical context-dependence, and between the framework's aspiration to transform inquiry and its operation within institutions governed by the paradigm it challenges (Section 3.2). These tensions are not problems awaiting solution. They are the permanent engines of the research programme – the structural features that ensure it remains generative rather than merely programmatic.

The third is reflexive survival. The dissertation's diagnosis of the Triple Lock (Section 3.2) – the mutually reinforcing semantic, methodological, and institutional mechanisms that reproduce the mechanistic paradigm – applies to any research programme that operates within the discipline the Triple Lock governs. A programme built on CAS foundations must navigate institutional conditions shaped by mechanistic commitments: publication standards that privilege prediction, hiring criteria that reward methodological uniformity, evaluation systems calibrated to the very paradigm the programme challenges. The thesis's own analysis implies

that the institutional conditions under which the programme can survive and develop are themselves a legitimate object of inquiry – not as a sociological aside, but as a direct consequence of treating the discipline as a system exhibiting the network effects, positive feedback, and path dependence the thesis identifies as constitutive of CAS (Sections 3.2, 6.4.5). This reflexive dimension is addressed not as a standalone programme but as a condition that pervades the others and is returned to in the section's closing.

6.5.1 Theoretical Extensions

Three theoretical frontiers follow directly from the reconstructive cascade and require development before the framework's full implications can be pursued empirically.

The first concerns the micro–meso–macro linkage. The dissertation restores the meso-level as the privileged analytical scale (Section 5.5), arguing that the firm's distinctive properties – capabilities, adaptive capacity, coordination architectures – emerge at the interaction level rather than residing in individuals or being imposed by macro forces. But the relationship between scales remains incompletely specified. How do individual-level actions aggregate into meso-level patterns without the aggregation destroying the interaction-dependent properties that make the meso-level causally distinctive? How do meso-level configurations condition macro-level outcomes – industry dynamics, market structure, institutional evolution – without the conditioning relationship being reducible to composition effects? This is the complexity-consistent version of the micro-foundations question, and its answer must be non-reductive: it must preserve the causal autonomy of each level while specifying the mechanisms through which levels interact. The thesis provides the ontological warrant for this inquiry (emergent wholeness, Section 5.1; multi-level causation, Section 5.5); what is required is the theoretical articulation of specific cross-level mechanisms in the context of firm organisation.

The second concerns firm identity as process-pattern. If the firm is an historically evolving, open, multi-level process-pattern (Section 5.4) rather than a substance that persists through change, then the question of what makes a firm the same firm over time requires reformulation. Under substance ontology, identity is given by the persistence of an underlying entity. Under process ontology, identity must be constituted through the continuity of reproductive processes – the ongoing patterned activity through which the firm maintains its distinctive configuration. This theoretical question bifurcates into two empirical programmes: first, specifying the minimal reproduction conditions – which processes, relationships, and coordination mechanisms must be sustained for the pattern to persist, and at what point does degradation constitute dissolution rather than transformation? Second, investigating what happens when those conditions are disrupted – mergers, acquisitions, radical restructuring – and whether the characteristic failures of such interventions (Section 6.4.3) are intelligible as disruptions to reproductive processes rather than misallocations of assets.

The third concerns multi-temporal analysis. The replacement of synchronic equilibrium with process ontology (Section 5.2) and the recognition of path dependence and irreversibility (Section 5.9) imply that mechanisms operating at different speeds interact to produce the firm's trajectory. Routines reproduce on daily cycles; capabilities evolve over years; industry architectures shift over decades; institutional frameworks transform over generations. The thesis establishes that time is constitutive rather than parametric – that temporal sequence matters ontologically, not merely as an index (Section 5.2). But the specific mechanisms through which processes at different temporal scales interact – how short-term efficiency gains produce long-term fragility (Section 5.7), how path-dependent commitments foreclose future

possibility spaces (Section 5.9), how phase transitions and tipping points emerge from the interaction of slow and fast dynamics – constitute a domain of inquiry that the mechanistic paradigm could not formulate because its equilibrium orientation collapses all dynamics onto a single horizon. This is one of the five domains the implications section identifies as structurally inaccessible under the inherited framework's foundational commitments (Section 6.4, closing), and it connects directly to the thesis's reconceptualisation of prediction as pattern recognition under non-ergodic conditions (Section 2.6.2).

The fourth concerns the economics of fundamental uncertainty. The replacement of risk with ontological uncertainty (Section 5.8.1) reconnects the theory of the firm to traditions – Knight, Keynes, Hayek, Shackle – that the mechanistic paradigm marginalised because their central insights were incompatible with its foundational commitments. The thesis provides what these traditions lacked: a systematic ontological and methodological architecture within which fundamental uncertainty is not a residual category but a constitutive feature of the economic landscape. What is required is theoretical development of how firms navigate environments characterised by genuine novelty, endogenous possibility spaces, and irreversible commitments – what the dissertation terms "scaffolded rationality," in which organisational routines, relational architectures, and institutional structures provide the cognitive scaffolding that enables bounded agents to act effectively under conditions that formal optimisation cannot address (Sections 5.5, 5.9). This programme connects the thesis's ontological commitments to practical questions about how firms actually make consequential decisions under conditions of irreducible uncertainty.

These four extensions are not independent. Multi-temporal analysis requires the micro–meso–macro linkage to specify how processes at different scales interact. Firm identity as process-pattern requires multi-temporal analysis to explain how reproductive continuity is maintained across different temporal horizons. The economics of fundamental uncertainty requires all three to specify how firms navigate possibility spaces that are constituted through their own path-dependent activity at multiple scales. The cascade logic that structures the thesis also structures its research programme.

6.5.2 The Empirical Programme

The thesis is a foundational reconstruction, not an empirical investigation. Its force is established through the internal coherence of the reconstructive cascade and its capacity to render visible phenomena that the mechanistic paradigm structurally excludes (Section 6.3). But its value is ultimately demonstrated through the empirical inquiry it enables. This subsection identifies five investigative streams, each derived from a specific achievement of the reconceptualisation and each specifying what the first steps of investigation look like from within the reconstructed framework.

The first stream concerns meso-level process studies. If capabilities are emergent properties of relational architectures (Section 5.4) rather than assets possessed by firms, then the dynamics through which capabilities are produced, sustained, and degraded become primary objects of inquiry. Longitudinal, process-tracing studies of specific firms – tracking how coordination patterns form, how routines reproduce, how interaction architectures evolve – would investigate whether the meso-level dynamics the thesis identifies as causally fundamental are empirically accessible and whether they explain outcomes that entity-level analysis systematically misses. A specific testable expectation follows: capability decay should follow relational disruption even when the individual competences that the mechanistic framework

treats as the capability's constituents remain intact. The asymmetry between entity-level continuity and process-level disruption is a signature prediction of the CAS framework that the mechanistic paradigm cannot generate.

The second stream concerns relational architecture mapping. The replacement of atomism with relational constitution (Section 5.4) implies that what a firm is – its distinctive capabilities, competitive position, and adaptive capacity – is partly constituted by network configurations that extend beyond its formal boundaries. Network-analytic studies that map the relational architectures through which capabilities emerge would test this claim directly. The specific question is whether network topology – density, centrality distribution, bridging patterns, redundancy – explains capability differences between firms better than firm-level attributes alone. This stream provides the thesis's strongest bridge between ontological commitment and empirical method: network analysis is the natural toolkit for investigating relational constitution, and the thesis provides the ontological warrant for treating network configurations as constitutive rather than merely contextual.

The third stream concerns the hidden efficiency–fragility trade-off. The network architecture and feedback analyses (Sections 5.7, 5.9) identify a structural trade-off invisible to equilibrium analysis: features that produce efficiency under stable conditions – tight coupling, eliminated redundancy, optimised allocation – produce fragility under stress. Standard performance metrics reward the systematic generation of invisible fragility. Crisis episodes constitute natural experiments: they reveal whether firms whose pre-crisis configurations exhibit the architectural features the thesis identifies as fragility-producing are disproportionately vulnerable, and whether resilience correlates with the redundancy and loose coupling that efficiency metrics penalise. This investigative stream converts a theoretical insight into an empirical object – the trade-off becomes measurable once the analytical apparatus registers it.

The fourth stream concerns intervention failures. Mergers, acquisitions, restructurings, and best-practice transplantations frequently destroy the capabilities they seek to capture. The thesis offers a structural explanation: these interventions fail because they treat capabilities as assets that can be transferred between organisational containers, when capabilities are in fact emergent properties of specific relational architectures that cannot survive the architectural disruption the intervention inflicts (Section 6.4.3). This constitutes ontological misidentification – intervention design premised on an incorrect ontology of what capabilities are. Systematic investigation of whether intervention outcomes correlate with the degree of relational-architectural disruption rather than with the financial or strategic logic of the intervention would test this explanation against alternatives. The practical relevance is immediate: if intervention failures are structurally intelligible under CAS assumptions in ways that the mechanistic paradigm cannot explain, the framework demonstrates its empirical value precisely where existing theory is weakest.

The fifth stream – contrastive paradigm studies – constitutes the programme's most direct evidential strategy. The thesis's perspectival epistemology (Section 2.4) implies that different foundational commitments render different aspects of the same phenomenon visible. The cleanest way to demonstrate this is to study the same organisational phenomenon under both frameworks simultaneously, documenting what each makes visible and what each systematically misses. Such studies would not "test" the CAS framework in the hypothetico-deductive sense that the mechanistic paradigm's epistemology demands; they would investigate the framework's generative capacity – whether it reveals patterns, mechanisms, and dynamics that the mechanistic framework cannot register. This is the form of empirical accountability the thesis's own epistemology licenses (Section 2.6.1): not prediction against a known

benchmark, but systematic demonstration that the framework illuminates phenomena the inherited apparatus renders invisible.

These five streams share a common methodological character. Each requires longitudinal, process-sensitive methods – case-based research, network analysis, comparative process tracing – applied to phenomena that the mechanistic paradigm structurally excludes. Each generates expectations that are specific enough to be investigated and distinctive enough that they could not have been generated from within the framework being challenged. Together, they constitute the empirical programme through which the reconceptualisation's explanatory power can be progressively demonstrated. The Abstract's three named programmes – empirical investigation of meso-level phenomena, operationalisation of structured pluralism, and case-based research examining firms as process-patterns – converge here, because the thesis's own ontological commitments imply that meso-level phenomena and process-patterns are not separate research objects but different descriptions of the same ontological reality (Section 5.4).

6.5.3 Methodological Development

The empirical programme described above requires methods that are not yet fully developed. The operationalisation gap identified in Section 6.3 – between the thesis's philosophical clarification of CAS-relevant constructs and the empirical instrumentarium needed to investigate them – is real rather than rhetorical. Closing that gap is itself a substantial research programme, and this subsection identifies four priorities.

The first concerns composition-sensitive measurement. The central methodological challenge the thesis creates is the measurement problem for emergent properties: how does one measure a property that exists only at the system level and is destroyed by decomposition? Emergence (Section 5.1) means that the relevant properties – adaptive capacity, organisational resilience, coordination quality – cannot be recovered from measurements of components. The required innovation is what might be termed composition-sensitive methods: measurement strategies that decompose for analytical purposes while preserving awareness that the decomposition eliminates the interaction-dependent properties under investigation. Network-analytic measures, information-theoretic indicators of organisational complexity, and multi-level modelling techniques that explicitly represent cross-level interactions are candidate approaches; what is needed is their systematic development and validation for specific CAS constructs in organisational contexts.

The second concerns adjudication criteria for structured pluralism. The thesis establishes structured pluralism as a methodological consequence of CAS ontology (Section 5.8.2), disciplined by admissibility axioms, positioning criteria, and transformation protocols. But Section 6.3 acknowledges that the dissertation does not fully resolve how adjudication proceeds in cases of deep theoretical incompatibility – cases where two methods, both admissible under CAS assumptions, generate conflicting accounts of the same phenomenon and no higher-order criterion determines precedence. Developing operational adjudication protocols for such cases is the most pressing methodological challenge the reconstruction creates. The protocols cannot be purely formal – the thesis's own epistemology implies that adjudication will be context-sensitive and historically informed – but they must be principled enough to prevent pluralism from collapsing into the undisciplined eclecticism the thesis diagnoses as the current field's characteristic failure.

The third concerns pattern prediction under non-ergodicity. The thesis reconceives prediction as pattern recognition, conditional expectation, and domain-bounded anticipation rather than

point forecasting (Section 2.6.2). But what do these epistemic ideals look like in practice? How does one specify what counts as a pattern, what makes an expectation conditional in the right way, and how domain boundaries are drawn? The explanation–prediction asymmetry that the thesis identifies as the sixth dimension of the visibility shift (Section 6.4) – the insight that rigorous explanation need not entail prediction under CAS conditions – creates an epistemological gap that subsequent methodological work must fill. Agent-based modelling, historical-comparative analysis, and simulation under non-ergodic conditions are candidate approaches, but their epistemological status within the reconstructed framework requires explicit articulation.

The fourth concerns the hidden ontological commitments of formal methods. Section 6.4.5 argues that formal methods – equilibrium analysis, optimisation, econometrics – carry implicit ontological commitments that constrain what they can represent before any empirical application begins. A systematic programme of ontological analysis applied to the discipline's formal toolkit – identifying what each method presupposes about the nature of its objects, what it can and cannot capture, and where its scope conditions hold – would be immediately productive. This is not a critique of formalism as such; it is an extension of the thesis's diagnostic method (Section 3.1) to the analytical instruments through which the discipline conducts inquiry. Such analysis would clarify where formal methods retain their full explanatory force, where they require supplementation, and where their ontological presuppositions render them actively misleading – contributing to the scope-aware repositioning of the discipline's existing toolkit that the next subsection addresses.

6.5.4 From Reconstruction to Empirical Inquiry

The gap between foundational reconstruction and empirical inquiry is not merely a matter of developing new methods. It requires the systematic repositioning of the discipline's existing theoretical and analytical resources within the architecture the thesis provides. This subsection identifies three dimensions of that bridging work.

The first concerns the scope-aware repositioning of existing theories. Section 6.4.1 argues that the major research programmes in firm theory – Williamson's transaction cost economics, the resource-based view, evolutionary approaches, institutional theories – are not invalidated by the reconceptualisation but repositioned: each captures something real about firms, and the diagnostic architecture (Section 3.1) makes explicit the foundational assumptions that determine each programme's scope and limits. What has not been done – and what constitutes perhaps the most immediately productive theoretical programme the thesis enables – is the systematic working-through of what each tradition explains under what conditions, what it captures and what it misses, and where its scope conditions hold and where they fail. This is not a taxonomic exercise. It requires treating the FROM→TO transformation framework (Section 2.3.6) as an analytical instrument: for each existing theory, identifying which mechanistic assumptions it retains, what those assumptions exclude, and what becomes accessible when the relevant assumptions are replaced with their CAS-compatible alternatives. The output would be a scope map of the discipline's existing theoretical resources – a specification of how each tradition relates to the others within a shared architectural framework, governed by explicit scope conditions rather than implicit ontological incompatibilities.

The second concerns the transformation of the FROM→TO framework itself into empirically tractable propositions. The fifteen-assumption architecture (Section 3.1) and the coordinated

replacement strategy (Sections 5.1–5.8) are constructed as philosophical arguments about the structure of economic reasoning. But each FROM→TO transformation carries empirical implications that can be investigated: does the shift from substance to process ontology generate different empirical expectations about capability dynamics? Does the shift from atomism to relational constitution predict different patterns of inter-firm performance variation? Does the shift from closure to openness generate different expectations about crisis propagation? Systematically converting these foundational transformations into testable empirical expectations would constitute a direct bridge between the thesis's philosophical reconstruction and the empirical programme described in Section 6.5.2. The contrastive paradigm studies identified there provide the investigative framework; the FROM→TO conversions provide the specific expectations to be investigated.

The third concerns cross-disciplinary synthesis. Section 6.4.5 argues that the mechanistic paradigm maintains disciplinary boundaries separating economics from organisation theory, sociology, network science, and complexity science – boundaries sustained partly by incompatible foundational assumptions. Once those assumptions are reconstructed on CAS-compatible grounds, these disciplines discover shared ontological commitments that enable genuine synthesis rather than the diplomatic combination of incommensurable perspectives. The thesis provides the architectural framework for such synthesis – its translation practices (Section 2.8.6) and structured pluralism (Section 5.8.2) specify the methodological conditions under which cross-disciplinary integration can proceed without collapsing into eclecticism. What is required is the practical demonstration that researchers working from CAS foundations across disciplinary boundaries can produce integrated accounts of organisational phenomena that no single discipline achieves in isolation. This programme is not merely additive. It tests whether the thesis's claim to provide a shared foundational architecture is practically realisable – whether the ontological common ground the reconstruction establishes is sufficient to support the cross-disciplinary synthesis it promises.

6.5.5 Artificial Intelligence, Computation, and the Transformation of Firms

Section 6.4.6 establishes that developments in artificial intelligence both validate and complicate the reconceptualisation. This subsection identifies three dimensions of a research programme that connects the thesis's foundational work to the technological transformation currently reshaping the phenomena it analyses.

The first dimension concerns AI as enabler of CAS inquiry. The epistemic ideals the thesis establishes as appropriate for complex adaptive systems – pattern recognition under uncertainty, navigation through evolving possibility spaces, multi-scale analysis, scenario-based reasoning (Section 2.6.2) – are precisely the capabilities that AI systems increasingly provide. For the first time, the analytical demands of complexity-informed research can be met with practical tools: agent-based models can simulate non-ergodic dynamics at scale; machine learning can identify patterns in process data that human analysis cannot detect; network algorithms can map relational architectures across thousands of nodes. The research programme must develop these computational tools explicitly for CAS-consistent inquiry – not importing AI as a black-box technique but designing computational methods whose epistemological assumptions align with the framework's ontological commitments.

The second dimension concerns AI as relational constituent. If firm capabilities are emergent properties of relational configurations (Section 5.4), then the introduction of AI agents into those configurations does not merely automate tasks; it changes the interaction architecture and

therefore changes what the firm is – not merely what it can do. AI systems are themselves adaptive, non-linear, and generative of genuine novelty; they exhibit precisely the properties that the mechanistic paradigm was designed to suppress. Understanding what firms become when their constitutive processes are partly conducted by adaptive, non-linear computational agents requires the full CAS apparatus the thesis develops. This is not an application of the framework to a new topic; it is a domain in which the framework's adequacy is most urgently tested, because the phenomena are evolving faster than the discipline's inherited conceptual resources can track.

The third dimension concerns governance under complexity. If AI accelerates the very complexity that renders the mechanistic framework inadequate – making boundaries more fluid, coordination more distributed, adaptation cycles shorter, and emergent dynamics more consequential (Section 6.4.6) – then governance frameworks designed for mechanistically understood firms face systematic inadequacy. The thesis implies that regulatory, competitive, and institutional evaluation must contend with entities whose properties are emergent, whose boundaries are relational, and whose trajectories are path-dependent. Developing governance approaches adequate to this reality connects the thesis's foundational work to pressing policy questions – including how meso-level configurations condition policy effects and how policy evaluation must accommodate path dependence and contextual sensitivity (Section 2.9.2, D.4) – while demonstrating that the reconceptualisation has practical consequences beyond theoretical reorientation.

6.5.6 Comparative, Cross-Cultural, and Meta-Theoretical Frontiers

Section 6.3 acknowledges that the dissertation draws predominantly on Western institutional contexts and theoretical traditions. Whether the framework's diagnostic and reconstructive achievements generalise beyond these contexts is an open question that defines the programme's comparative frontier. Two distinct questions must be separated.

The first is whether the diagnostic architecture generalises. The fifteen-assumption network (Section 3.1) and the Triple Lock (Section 3.2) characterise the mechanistic paradigm as it developed within a specific intellectual tradition – the Newtonian inheritance that shaped Anglo-American and Continental European economics. Economic traditions that developed under different philosophical genealogies – developmental-state frameworks, post-Soviet institutional economics, traditions grounded in non-Western ontologies – may exhibit different assumptive architectures. A comparative programme that applies the thesis's diagnostic method to these traditions would investigate whether the mechanistic paradigm's structure is culturally specific or whether analogous lock-in mechanisms operate across different intellectual environments. The method itself – systematic identification and assessment of foundational assumptions – is transferable even if the specific assumptions differ.

The second is whether the reconstructive cascade generalises. The thesis argues that process ontology, relational constitution, emergence, and structured pluralism are ontologically appropriate for complex adaptive systems as a class. If firms embedded in fundamentally different institutional orders – different legal frameworks, different governance traditions, different cultural conceptions of organisational identity – exhibit different specific configurations while sharing CAS properties, then the reconstructive framework should accommodate this variation through its context-sensitivity rather than requiring separate foundations for each context. Cross-cultural and comparative-institutional studies that investigate the same CAS dynamics in different institutional settings would test the

framework's generality at the level that matters: not whether the specific relational architectures are identical across contexts (the framework predicts they will not be), but whether the CAS properties – emergence, path dependence, relational constitution, multi-temporal dynamics – are empirically accessible across institutional orders using the framework's analytical apparatus.

A broader frontier extends beyond the firm to meta-theoretical territory. The thesis develops a method – systematic identification of foundational assumptions, diagnostic assessment of their consequences, coordinated reconstruction through a FROM→TO transformation framework – that is not inherently limited to the theory of the firm. The problematisation methodology the thesis employs (Section 2.3) – treating persistent limitations as symptoms of structural rather than local deficiencies – constitutes a methodological register that economics currently lacks as a standard complement to gap-spotting. The diagnostic architecture (Section 3.1) could in principle be applied to any discipline whose inquiry is structured by implicit foundational assumptions: development economics, health policy, urban planning, ecological science, and any field where inherited conceptual commitments may constrain what can be investigated. What the thesis terms "paradigm science" – the systematic study of how foundational assumptions structure, enable, and constrain inquiry across disciplines – is a potential contribution that the dissertation may underappreciate. It would extend the thesis's methodological innovations beyond its specific domain, testing whether the diagnostic-reconstructive sequence the thesis demonstrates is a transferable research methodology or a one-time achievement.

What Firm Theory Becomes

The research programme outlined above is itself a complex adaptive system. Its components interact: theoretical extensions will be developed, empirical findings will modify them, methodological innovations will enable new forms of empirical investigation, cross-disciplinary synthesis will reshape the questions that theoretical work addresses. Its trajectory is path-dependent: early theoretical developments and their accompanying empirical results will constrain and enable later investigations, and the order in which methodological challenges are addressed will shape the programme's evolution. Its outcomes will exhibit emergence: the understanding of firms that the programme progressively builds will have properties that no individual investigative stream could have generated alone. A framework that remained unchanged by the work it enabled would, by its own criteria, be less adequate than one that evolved – because responsiveness to feedback, openness to transformation, and the generation of novelty through interaction are precisely the properties the framework attributes to the systems it studies.

This self-referential quality is not rhetorical. The thesis's own analysis implies that the institutional conditions under which the programme operates matter substantively to its prospects. Economics-as-institution is itself a complex adaptive system exhibiting the network effects, positive feedback, and path dependence that the thesis identifies as constitutive of CAS (Section 3.2). The Triple Lock's semantic, methodological, and institutional mechanisms do not passively resist alternative frameworks. The research programme must therefore be self-conscious about its own institutional ecology: about how its methods will be evaluated, how its results will be received, how its practitioners will adopt it, and how its relationship to the dominant paradigm will be navigated. Reflexive survival – the capacity of the programme to maintain its coherence while operating within institutions shaped by the commitments it challenges – is not a sociological afterthought but a condition of the programme's viability that the thesis's own analysis makes explicit.

The dissertation began with a question: what kind of entity must the firm be assumed to be if economic theory is to explain how firms emerge, adapt, and transform under conditions of genuine uncertainty and continuous change (Section 1.5.1)? The preceding chapters have provided a systematic answer: the firm is a complex adaptive system – an historically evolving, open, multi-level process-pattern, constituted through relational configurations, operating under fundamental uncertainty, and exhibiting emergent properties irreducible to its components. The research programmes identified in this section specify what becomes investigable once that answer is taken seriously – not as a metaphor to be acknowledged and then set aside, but as a foundational commitment to be inhabited, tested, and progressively refined. The foundations the thesis provides are not the last word; they are the first disciplined move in a conversation that the discipline has been unable to begin under the constraints it inherited. What is built upon them belongs to the discipline and to the generations of scholars who will inhabit, test, extend, and ultimately transform the architecture this thesis provides.

References List

A

Abbott, A. (1988). *The System of Professions: An Essay on the Division of Expert Labor*. Chicago: University of Chicago Press.

Abbott, A. (2004). *Methods of Discovery: Heuristics for the Social Sciences*. New York: W.W. Norton.

Acemoglu, D. and Robinson, J.A. (2012). *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*. New York: Crown Business.

Ackoff, R.L. (1971). Towards a system of systems concepts. *Management Science*, 17(11), 661-671.

Ackoff, R.L. (1981). *Creating the Corporate Future: Plan or be Planned For*. New York: Wiley.

Ackoff, R.L. (1993). Idealized design: Creative corporate visioning. *Omega*, 21(4), 401-410.

Ackoff, R.L. (1999). *Re-Creating the Corporation: A Design of Organizations for the 21st Century*. Oxford University Press.

Ackoff, R.L. and Gharajedaghi, J. (2003). On the mismatch between systems and their models. Ackoff Center Weblog. [Also revision of: Ackoff, R.L. and Gharajedaghi, J. (1996). Reflections on systems and their models. *Systems Research*, 13(1), 13-23.].

Ackoff, R.L. and Gharajedaghi, J. (2003). On the mismatch between systems and their models. *Systems Research and Behavioral Science*, 13(5), 367-371.

Adner, R. (2017). Ecosystem as Structure: An Actionable Construct for Strategy. *Journal of Management*, 43(1), 39–58.

Adner, R. and Helfat, C.E. (2003). Corporate effects and dynamic managerial capabilities. *Strategic Management Journal*, 24(10), 1011-1025.

Ahuja, G. (2000). Collaboration networks, structural holes, and innovation: A longitudinal study. *Administrative Science Quarterly*, 45(3), 425-455.

Akerlof, G.A. (1970). The market for lemons: Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3), 488-500.

Akerlof, G.A. (2020). Sins of omission and the practice of economics. *Journal of Economic Literature*, 58(2), 405-418.

Aldrich, H.E. (1979). *Organizations and Environments*. Englewood Cliffs, NJ: Prentice-Hall.

Aldrich, H.E. and Ruef, M. (2006). *Organizations Evolving* (2nd ed.). London: Sage.

Alhadeff-Jones, M. (2013). Complexity, methodology and method: Crafting a critical process of research. *Complicity: An International Journal of Complexity and Education*, 10(1/2), 19-44.

- Allen, C.R. and Holling, C.S. (2010). Novelty, adaptive capacity, and resilience. *Ecology and Society*, 15(3), 24.
- Allen, P.M. (2015). Evolution: Complexity, Uncertainty and Innovation. *Journal of Evolutionary Economics*, 25(1), 1–28.
- Allin, B.W. (1949). Is economics a science? *Canadian Journal of Economics and Political Science*, 15(1), 1-11.
- Alston, L.J. (2023). Institutions and economic development. In C. Ménard and M.M. Shirley (Eds.), *Handbook of New Institutional Economics* (pp. 3-28). Berlin: Springer.
- Alvesson, M. and Sandberg, J. (2011). Generating research questions through problematization. *Academy of Management Review*, 36(2), 247-271.
- Alvesson, M. and Sandberg, J. (2013). *Constructing Research Questions: Doing Interesting Research*. London: Sage.
- Alvesson, M. and Sandberg, J. (2020). The problematizing review: A counterpoint to Elsbach and Van Knippenberg's argument for integrative reviews. *Journal of Management Studies*, 57(6), 1290-1304.
- Alvesson, M. and Sandberg, J. (2022). Re-imagining the research process: Conventional and alternative metaphors. *SAGE Open*, 12(1), 1-13.
- Alvesson, M. and Sandberg, J. (2024). Pre-understandings and their revision in theorization. *Academy of Management Review*. [Forthcoming].
- Anderson, P. (1999). Complexity theory and organization science. *Organization Science*, 10(3), 216-232.
- Anderson, P. and Tushman, M.L. (1990). Technological discontinuities and dominant designs: A cyclical model of technological change. *Administrative Science Quarterly*, 35(4), 604-633.
- Anderson, P.W. (1972). More Is Different. *Science*, 177(4047), 393–396.
- Apel, K.-O. (2011). *Paradigms of First Philosophy: Metaphysical, Transcendental-Pragmatic, and Hermeneutic Perspectives*. Hamburg: Humanities Online.
- Archer, M.S. (1995). *Realist Social Theory: The Morphogenetic Approach*. Cambridge University Press.
- Arrow, K.J. (1962). Economic welfare and the allocation of resources for invention. In *The Rate and Direction of Inventive Activity*. Princeton: Princeton University Press.
- Arrow, K.J. (1994). Methodological Individualism and Social Knowledge. *American Economic Review*, 84(2), 1–9.
- Arrow, K.J. and Debreu, G. (1954). Existence of an equilibrium for a competitive economy. *Econometrica*, 22(3), 265-290.

- Arthur, W.B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *Economic Journal*, 99(394), 116-131.
- Arthur, W.B. (1994). *Increasing Returns and Path Dependence in the Economy*. University of Michigan Press.
- Arthur, W.B. (1999). Complexity and the economy. *Science*, 284(5411), 107–109.
- Arthur, W.B. (2013). *Complexity Economics: A Different Framework for Economic Thought*. Santa Fe Institute Working Paper.
- Arthur, W.B. (2014). *Complexity and the Economy*. Oxford: Oxford University Press. [First published as SFI working paper].
- Arthur, W.B. (2015). *Complexity and the Economy*. Oxford University Press.
- Arthur, W.B. (2021). Foundations of Complexity Economics. *Nature Reviews Physics*, 3, 136–145.
- Ashby, W.R. (1956). *An Introduction to Cybernetics*. London: Chapman and Hall.
- Astley, W.G. (1985). Administrative science as socially constructed truth. *Administrative Science Quarterly*, 30, 497-513.
- Astley, W.G. and Van de Ven, A.H. (1983). Central perspectives and debates in organization theory. *Administrative Science Quarterly*, 28(2), 245-273.
- Axelrod, R. (1997). *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*. Princeton: Princeton University Press.
- Axelrod, R. & Cohen, M.D. (2000). *Harnessing Complexity: Organizational Implications of a Scientific Frontier*. Basic Books.
- Axelrod, R. and Cohen, M.D. (1999). *Harnessing Complexity: Organizational Implications of a Scientific Frontier*. New York: Free Press.
- Axtell, R.L. (2001). Zipf distribution of US firm sizes. *Science*, 293(5536), 1818-1820.
- Ayres, R.U. and Warr, B. (2009). *The Economic Growth Engine: How Energy and Work Drive Material Prosperity*. Cheltenham: Edward Elgar.

B

- Bacharach, S.B. (1989). Organizational theories: Some criteria for evaluation. *Academy of Management Review*, 14(4), 496-515.
- Bachelard, G. (1938). *La Formation de l'Esprit Scientifique: Contribution à une Psychanalyse de la Connaissance Objective*. Paris: Vrin. [English trans.: *The Formation of the Scientific Mind*. Manchester: Clinamen Press, 2002].
- Bak, P. (1996). *How Nature Works: The Science of Self-Organized Criticality*. New York: Copernicus/Springer.
- Baker, G., Gibbons, R. & Murphy, K.J. (2002). Relational Contracts and the Theory of the Firm. *Quarterly Journal of Economics*, 117(1), 39–84.
- Bankes, S.C. (1993). Exploratory modeling for policy analysis. *Operations Research*, 41(3), 435-449.
- Barabási, A.-L. and Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439), 509-512.
- Barney, J.B. (1991). Firm resources and sustained competitive advantage. *Journal of Management*, 17(1), 99-120.
- Bartunek, J.M., Rynes, S.L. and Ireland, R.D. (2006). What makes management research interesting? *Academy of Management Journal*, 49(1), 9-15.
- Batterman, R.W. (2002). *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*. Oxford: Oxford University Press.
- Battiston, S., Farmer, J.D., Flache, A., Garlaschelli, D., Haldane, A.G., Heesterbeek, H., Hommes, C., Jaeger, C., May, R. & Scheffer, M. (2016). Complexity Theory and Financial Regulation. *Science*, 351(6275), 818–819.
- Battiston, S., Puliga, M., Kaushik, R., Tasca, P. and Caldarelli, G. (2012). DebtRank: Too central to fail? Financial networks, the FED and systemic risk. *Scientific Reports*, 2, 541.
- Bedau, M.A. (1997). Weak emergence. *Philosophical Perspectives*, 11, 375-399.
- Bedau, M.A. (2008). Downward Causation and Autonomy in Weak Emergence. In *Emergence: Contemporary Readings* (pp. 155–172). MIT Press.
- Bennett, A. and Elman, C. (2006). Qualitative research: Recent developments in case study methods. *Annual Review of Political Science*, 9, 455-476.
- Bergmann, G. (2012). *The Metaphysics of Logical Positivism* (2nd ed.). Whitefish, MT: Literary Licensing.
- Bergson, H. (1907). *L'Évolution Créatrice*. Paris: Alcan. [English trans.: *Creative Evolution*. New York: Henry Holt, 1911].

- Bernanke, B. (2010). Causes of the Recent Financial and Economic Crisis. Testimony before the Financial Crisis Inquiry Commission.
- Bertalanffy, L. von (1968). General System Theory: Foundations, Development, Applications. New York: George Braziller.
- Bhaskar, R. (1975). A Realist Theory of Science. Leeds: Leeds Books. [2nd ed. 1978, Brighton: Harvester Press].
- Bhaskar, R. (1978). A Realist Theory of Science (2nd ed.). Brighton: Harvester.
- Bhaskar, R. (1979). The Possibility of Naturalism. Brighton: Harvester.
- Bhaskar, R. (1986). Scientific Realism and Human Emancipation. London: Verso.
- Biddle, J. (1997). Towards a history of American economics. *History of Political Economy*, 29(5), 1-8.
- Bizzarri, M. (2021). The concept of nature between Heraclitus and Prigogine. *Organisms: Journal of Biological Sciences*, 5(1), 67-71. doi:10.13133/2532-5876/17540
- Bizzarri, M. (Ed.). (2020). Complexity in Biological Organization: A Challenge for Reductionism. Special issue of *Progress in Biophysics and Molecular Biology*.
- Boisot, M.H. (1998). Knowledge Assets: Securing Competitive Advantage in the Information Economy. Oxford: Oxford University Press.
- Borgatti, S.P. & Foster, P.C. (2003). The Network Paradigm in Organizational Research. *Journal of Management*, 29(6), 991–1013.
- Borgatti, S.P. and Halgin, D.S. (2011). On network theory. *Organization Science*, 22(5), 1168-1181.
- Borgatti, S.P., Mehra, A., Brass, D.J. and Labianca, G. (2009). Network analysis in the social sciences. *Science*, 323(5916), 892-895.
- Boulding, K.E. (1956). General systems theory – The skeleton of science. *Management Science*, 2(3), 197-208.
- Boulding, K.E. (1978). *Ecodynamics: A New Theory of Societal Evolution*. Sage.
- Bourdieu, P. (1986). The forms of capital. In J. Richardson (Ed.), *Handbook of Theory and Research for the Sociology of Education* (pp. 241-258). Westport, CT: Greenwood.
- Bourdieu, P. (1991). Language and Symbolic Power. Edited by J.B. Thompson. Translated by G. Raymond and M. Adamson. Cambridge: Polity Press.
- Bourdieu, P. (2004). *Science of Science and Reflexivity*. Chicago: University of Chicago Press.
- Braithwaite, R.B. (1953). *Scientific Explanation: A Study of the Function of Theory, Probability and Law in Science*. Cambridge: Cambridge University Press.

Bronk, R. (2011). Uncertainty, Modelling Monocultures and the Financial Crisis. *Business Economist*, 42(2), 5–18.

Bunge, M. (1973). *Method, Model and Matter*. Dordrecht: D. Reidel.

Bunge, M. (1997). Mechanism and explanation. *Philosophy of the Social Sciences*, 27(4), 410-465.

Bunge, M. (2003). *Emergence and Convergence: Qualitative Novelty and the Unity of Knowledge*. University of Toronto Press.

Burt, R.S. (1992). *Structural Holes: The Social Structure of Competition*. Cambridge, MA: Harvard University Press.

Burt, R.S. (2004). Structural Holes and Good Ideas. *American Journal of Sociology*, 110(2), 349–399.

Bylund, P.L. (2021). Understanding the limits of pure theory in economics: Knight and Mises. In L. Fiorito, S. Scheall, and C.E. Suprinyak (Eds.), *Research in the History of Economic Thought and Methodology* (Vol. 39, pp. 3-18). Emerald Publishing.

C

- Caballero, R.J. (2010). Macroeconomics after the crisis: Time to deal with the pretense-of-knowledge syndrome. *Journal of Economic Perspectives*, 24(4), 85-102.
- Campbell, D.T. (1974). 'Downward causation' in hierarchically organised biological systems. In F.J. Ayala and T. Dobzhansky (Eds.), *Studies in the Philosophy of Biology* (pp. 179-186). London: Macmillan.
- Capra, F. and Luisi, P.L. (2014). *The Systems View of Life: A Unifying Vision*. Cambridge: Cambridge University Press.
- Cartwright, N. (1999). *The Dappled World: A Study of the Boundaries of Science*. Cambridge University Press.
- Catmull, E. & Wallace, A. (2014). *Creativity, Inc.* Random House.
- Caves, R.E. and Porter, M.E. (1977). From entry barriers to mobility barriers: Conjectural decisions and contrived deterrence to new competition. *Quarterly Journal of Economics*, 91(2), 241-261.
- Chalmers, D.J. (2006). Strong and weak emergence. In P. Clayton and P. Davies (Eds.), *The Re-Emergence of Emergence* (pp. 244-254). Oxford: Oxford University Press.
- Chandler, A.D. (1962). *Strategy and Structure: Chapters in the History of the Industrial Enterprise*. Cambridge, MA: MIT Press.
- Chesbrough, H.W. (2003). *Open Innovation: The New Imperative for Creating and Profiting from Technology*. Boston: Harvard Business School Press.
- Child, J. (1972). Organizational structure, environment and performance: The role of strategic choice. *Sociology*, 6(1), 1-22.
- Choi, T.Y. & Krause, D.R. (2006). The Supply Base and Its Complexity. *Journal of Operations Management*, 24(5), 637-652.
- Christensen, C.M. (1997). *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Boston: Harvard Business School Press.
- Cilliers, P. (1998). *Complexity and Postmodernism: Understanding Complex Systems*. Routledge.
- Cilliers, P. (2005). Complexity, Deconstruction and Relativism. *Theory, Culture & Society*, 22(5), 255-267.
- Cilliers, P. (2005). Knowledge, limits and boundaries. *Futures*, 37(7), 605-613.
- Clayton, P. and Davies, P. (Eds.). (2004). *The Re-Emergence of Emergence: The Emergentist Hypothesis from Science to Religion*. Oxford: Oxford University Press.
- Cleveland, H. (1982). Information as a resource. *The Futurist*, 16(6), 34-39.

- Coase, R.H. (1937). The nature of the firm. *Economica*, 4(16), 386-405.
- Cohen, I.B. (1985). *Revolution in Science*. Harvard University Press.
- Cohen, M.D. and Bacdayan, P. (1994). Organizational routines are stored as procedural memory: Evidence from a laboratory study. *Organization Science*, 5(4), 554-568.
- Cohen, W.M. and Levinthal, D.A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1), 128-152.
- Colander, D. (2000). The Death of Neoclassical Economics. *Journal of the History of Economic Thought*, 22(2), 127–143.
- Colander, D., Föllmer, H., Haas, A., Goldberg, M.D., Juselius, K., Kirman, A., Lux, T. and Sloth, B. (2009). The financial crisis and the systemic failure of academic economics. *Critical Review*, 21(2-3), 249-267.
- Conant, J.B. (1947). *On Understanding Science: An Historical Approach*. New Haven: Yale University Press.
- Craver, C.F. & Bechtel, W. (2007). Top-down causation without top-down causes. *Biology and Philosophy*, 22(4), 547–563.
- Cusumano, M.A., Gawer, A. and Yoffie, D.B. (2019). *The Business of Platforms: Strategy in the Age of Digital Competition, Innovation, and Power*. New York: Harper Business.
- Cyert, R.M. and March, J.G. (1963). *A Behavioral Theory of the Firm*. Englewood Cliffs, NJ: Prentice-Hall.

D

Daft, R.L. and Lewin, A.Y. (1990). Can organization studies begin to break out of the normal science straitjacket? An editorial essay. *Organization Science*, 1(1), 1-9.

Danermark, B., Ekström, M., Jakobsen, L. and Karlsson, J.C. (2002). *Explaining Society: Critical Realism in the Social Sciences*. London: Routledge.

David, P.A. (1985). Clio and the Economics of QWERTY. *American Economic Review*, 75(2), 332–337.

Davidson, P. (1991). Is probability theory relevant for uncertainty? A Post Keynesian perspective. *Journal of Economic Perspectives*, 5(1), 129-143.

Davidson, P. (2012). Is economics a science? Should economics be rigorous? *Real-World Economics Review*, 59, 86-96.

Davis, L.E. and North, D.C. (1971). *Institutional Change and American Economic Growth*. Cambridge: Cambridge University Press.

Davis, M.S. (1971). That's interesting! Towards a phenomenology of sociology and a sociology of phenomenology. *Philosophy of the Social Sciences*, 1(4), 309-344.

De Rond, M. and Miller, A.N. (2005). Publish or perish: Bane or boon of academic life? *Journal of Management Inquiry*, 14(4), 321-329.

Deacon, T.W. (2000). *Incomplete Nature: How Mind Emerged from Matter*. New York: W.W. Norton.

Deacon, T.W. (2011). *Incomplete Nature: How Mind Emerged from Matter*. New York: W.W. Norton.

Denzin, N.K. (1978). *The Research Act: A Theoretical Introduction to Sociological Methods* (2nd ed.). New York: McGraw-Hill.

DiMaggio, P.J. and Powell, W.W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*, 48(2), 147-160.

Dopfer, K. and Potts, J. (2008). *The General Theory of Economic Evolution*. London: Routledge.

Dopfer, K., Foster, J. & Potts, J. (2004). Micro-Meso-Macro. *Journal of Evolutionary Economics*, 14(3), 263–279.

Dosi, G. (1982). Technological paradigms and technological trajectories: A suggested interpretation of the determinants and directions of technical change. *Research Policy*, 11(3), 147-162.

Dosi, G. (1988). Sources, procedures, and microeconomic effects of innovation. *Journal of Economic Literature*, 26(3), 1120-1171.

Dosi, G. and Marengo, L. (1994). Some elements of an evolutionary theory of organizational competences. In R.W. England (Ed.), *Evolutionary Concepts in Contemporary Economics* (pp. 157-178). Ann Arbor: University of Michigan Press.

Dosi, G. and Nelson, R.R. (1994). An introduction to evolutionary theories in economics. *Journal of Evolutionary Economics*, 4(3), 153-172.

Dosi, G., Freeman, C., Nelson, R., Silverberg, G. and Soete, L. (eds.) (1988). *Technical Change and Economic Theory*. London: Pinter.

Dow, S.C. (1997). Methodological pluralism and pluralism of method. In A. Salanti and E. Screpanti (Eds.), *Pluralism in Economics* (pp. 89-99). Cheltenham: Edward Elgar.

Dow, S.C. (2004). Structured pluralism. *Journal of Economic Methodology*, 11(3), 275–290.

Duhigg, C. (2016). *Smarter Faster Better: The Secrets of Being Productive in Life and Business*. New York: Random House.

Durlauf, S.N. (2001). A framework for the study of individual behavior and social interactions. *Sociological Methodology*, 31(1), 47-87.

Dyer, J.H. and Nobeoka, K. (2000). Creating and managing a high-performance knowledge-sharing network: The Toyota case. *Strategic Management Journal*, 21(3), 345-367.

Dyer, J.H. and Singh, H. (1998). The relational view: Cooperative strategy and sources of interorganizational competitive advantage. *Academy of Management Review*, 23(4), 660-679.

E

Eisenhardt, K.M. (1989). Building theories from case study research. *Academy of Management Review*, 14(4), 532-550.

Eisenhardt, K.M. and Martin, J.A. (2000). Dynamic capabilities: What are they? *Strategic Management Journal*, 21(10-11), 1105-1121.

Eisenhardt, K.M. and Sull, D.N. (2001). Strategy as simple rules. *Harvard Business Review*, 79(1), 106-116.

Elder-Vass, D. (2010). *The Causal Power of Social Structures: Emergence, Structure and Agency*. Cambridge University Press.

Emirbayer, M. (1997). Manifesto for a relational sociology. *American Journal of Sociology*, 103(2), 281-317.

Epstein, J.M. (2006). *Generative Social Science: Studies in Agent-Based Computational Modeling*. Princeton: Princeton University Press.

Epstein, J.M. (2008). Why model? *Journal of Artificial Societies and Social Simulation*, 11(4), 12.

Epstein, J.M. & Axtell, R. (1996). *Growing Artificial Societies: Social Science from the Bottom Up*. MIT Press.

Epstein, J.M. and Axtell, R. (1996). *Growing Artificial Societies: Social Science from the Bottom Up*. Washington, DC: Brookings Institution Press.

Estrada, F. (2023). Complexity and Economics: A Research Agenda. *Journal of Economic Methodology*, 30(2), 1155–1170.

F

- Fang, F.C. & Casadevall, A. (2011). Reductionist and Holistic Science. *Infection and Immunity*, 79(4), 1401–1404.
- Farmer, J.D. and Foley, D. (2009). The economy needs agent-based modelling. *Nature*, 460(7256), 685-686.
- Farmer, J.D. and Geanakoplos, J. (2009). The virtues and vices of equilibrium and the future of financial economics. *Complexity*, 14(3), 11-38.
- Feldman, M.S. and Pentland, B.T. (2003). Reconceptualizing organizational routines as a source of flexibility and change. *Administrative Science Quarterly*, 48(1), 94-118.
- Felin, T. and Foss, N.J. (2005). Strategic organization: A field in search of micro-foundations. *Strategic Organization*, 3(4), 441-455.
- Felin, T., Foss, N.J., Heimeriks, K.H. and Madsen, T.L. (2012). Microfoundations of routines and capabilities: Individuals, processes, and structure. *Journal of Management Studies*, 49(8), 1351-1374.
- Feyerabend, P. (1975). *Against Method: Outline of an Anarchistic Theory of Knowledge*. London: New Left Books.
- Feyerabend, P. (1978). *Science in a Free Society*. London: New Left Books.
- Fiss, P.C. (2007). A set-theoretic approach to organizational configurations. *Academy of Management Review*, 32(4), 1180-1198.
- Fiss, P.C. (2011). Building better causal theories: A fuzzy set approach to typologies in organization research. *Academy of Management Journal*, 54(2), 393-420.
- Flyvbjerg, B. (2001). *Making Social Science Matter: Why Social Inquiry Fails and How It Can Succeed Again*. Cambridge: Cambridge University Press.
- Flyvbjerg, B. (2006). Five misunderstandings about case-study research. *Qualitative Inquiry*, 12(2), 219-245.
- Fontana, M. (2008). *The Complexity Approach to Economics*. CESMEP Working Paper 01/2008.
- Fontana, M. (2010). Can neoclassical economics handle complexity? The fallacy of the oil spot dynamic. *Journal of Economic Behavior & Organization*, 76(3), 584–596.
- Foster, J. (2005). From simplistic to complex systems in economics. *Cambridge Journal of Economics*, 29(6), 873-892.
- Foster, J. & Metcalfe, J.S. (2012). Economic Emergence: An Evolutionary Economic Perspective. *Journal of Economic Behavior & Organization*, 82(2-3), 420–432.
- Foster, J. and Metcalfe, J.S. (eds.) (2001). *Frontiers of Evolutionary Economics: Competition, Self-Organization and Innovation Policy*. Cheltenham: Edward Elgar.

- Foucault, M. (1984). Polemics, politics, and problematizations. In P. Rabinow (Ed.), *The Foucault Reader* (pp. 381-390). New York: Pantheon Books.
- Foucault, M. (1985). *The Use of Pleasure: Volume 2 of The History of Sexuality*. Translated by R. Hurley. New York: Pantheon.
- Freeman, C. (1982). *The Economics of Industrial Innovation* (2nd ed.). London: Frances Pinter.
- Freeman, C. (1987). *Technology Policy and Economic Performance: Lessons from Japan*. London: Pinter Publishers.
- French, S. (2014). *The Structure of the World: Metaphysics and Representation*. Oxford University Press.
- Friedman, M. (1953). The methodology of positive economics. In *Essays in Positive Economics*. Chicago: University of Chicago Press.
- Fromm, J. (2005). Types and forms of emergence. arXiv preprint nlin/0506028.
- Frydman, R. and Goldberg, M.D. (2007). *Imperfect Knowledge Economics: Exchange Rates and Risk*. Princeton: Princeton University Press.
- Fuhse, J.A. (2022). How can theories represent social phenomena? *Sociological Theory*, 40(2), 99-123.

G

Gadamer, H.-G. (1960). *Truth and Method*. [Wahrheit und Methode]. Translated by J. Weinsheimer and D.G. Marshall (1989). New York: Crossroad.

Gadamer, H.-G. (1975). *Truth and Method*. Translated by J. Weinsheimer and D.G. Marshall. New York: Seabury Press.

Gallegati, M. & Kirman, A. (2025). Complexity and Economic Policy. *Journal of Economic Dynamics and Control*.

Gavetti, G. and Levinthal, D. (2000). Looking forward and looking backward: Cognitive and experiential search. *Administrative Science Quarterly*, 45(1), 113-137.

Gavetti, G., Levinthal, D.A. and Rivkin, J.W. (2005). Strategy making in novel and complex worlds: The power of analogy. *Strategic Management Journal*, 26(8), 691-712.

Gawer, A. (2014). Bridging differing perspectives on technological platforms: Toward an integrative framework. *Research Policy*, 43(7), 1239-1249.

Geanakoplos, J. (2009). The leverage cycle. *NBER Macroeconomics Annual*, 24(1), 1-65.

George, A.L. and Bennett, A. (2005). *Case Studies and Theory Development in the Social Sciences*. Cambridge, MA: MIT Press.

Georgescu-Roegen, N. (1971). *The Entropy Law and the Economic Process*. Harvard University Press.

Gershenson, C. and Heylighen, F. (2006). When can we call a system self-organizing? In W. Banzhaf et al. (Eds.), *Advances in Artificial Life* (pp. 606-614). Berlin: Springer.

Gershenson, C., Trianni, V., Werfel, J. & Sayama, H. (2011). When Can we Call a System Self-organizing? *Advances in Artificial Life*.

Gharajedaghi, J. (2011). *Systems Thinking: Managing Chaos and Complexity* (3rd ed.). Burlington, MA: Morgan Kaufmann.

Giddens, A. (1984). *The Constitution of Society: Outline of the Theory of Structuration*. Polity Press.

Giere, R.N. (2006). *Scientific Perspectivism*. Chicago: University of Chicago Press.

Gigerenzer, G. (2008). *Rationality for Mortals: How People Cope with Uncertainty*. Oxford University Press.

Gigerenzer, G. and Todd, P.M. (1999). *Simple Heuristics That Make Us Smart*. New York: Oxford University Press.

Gilbert, N. (2008). *Agent-Based Models*. Thousand Oaks, CA: Sage.

Giorgi, A. (2005). The phenomenological movement and research in the human sciences. *Nursing Science Quarterly*, 18(1), 75-82.

- Girvan, M. and Newman, M.E.J. (2002). Community structure in social and biological networks. *Proceedings of the National Academy of Sciences*, 99(12), 7821-7826.
- Goodhart, C.A.E. (1975). Monetary relationships: A view from Threadneedle Street. *Papers in Monetary Economics*, Reserve Bank of Australia.
- Gould, S.J. & Eldredge, N. (1977). Punctuated Equilibria: The Tempo and Mode of Evolution Reconsidered. *Paleobiology*, 3(2), 115–151.
- Gouldner, A.W. (1970). *The Coming Crisis of Western Sociology*. New York: Basic Books.
- Goyal, S. (2007). *Connections: An Introduction to the Economics of Networks*. Princeton University Press.
- Grandori, A. (1997). An organizational assessment of interfirm coordination modes. *Organization Studies*, 18(6), 897-925.
- Granovetter, M. (1985). Economic action and social structure: The problem of embeddedness. *American Journal of Sociology*, 91(3), 481-510.
- Granovetter, M.S. (1973). The strength of weak ties. *American Journal of Sociology*, 78(6), 1360-1380.
- Grant, R.M. (1996). Toward a knowledge-based theory of the firm. *Strategic Management Journal*, 17(Winter Special Issue), 109-122.
- Greenstein, S. (2012). Nurturing the accumulation of innovations: Lessons from the Internet. NBER Working Paper No. 17858.
- Greenstein, S. and Zhu, F. (2012). Is Wikipedia biased? *American Economic Review*, 102(3), 343-348.
- Gulati, R. (1998). Alliances and networks. *Strategic Management Journal*, 19(4), 293-317.
- Gulati, R. and Gargiulo, M. (1999). Where do interorganizational networks come from? *American Journal of Sociology*, 104(5), 1439-1493.

H

Hacking, I. (1983). *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science*. Cambridge: Cambridge University Press.

Hacking, I. (1999). *The Social Construction of What?* Cambridge, MA: Harvard University Press.

Haldane, A.G. and Madouros, V. (2012). The dog and the frisbee. Speech at Federal Reserve Bank of Kansas City's 36th Economic Policy Symposium, Jackson Hole, Wyoming.

Haldane, A.G. and May, R.M. (2011). Systemic risk in banking ecosystems. *Nature*, 469(7330), 351-355.

Hall, P.A. and Soskice, D. (Eds.). (2001). *Varieties of Capitalism: The Institutional Foundations of Comparative Advantage*. Oxford: Oxford University Press.

Hannan, M.T. and Freeman, J. (1984). Structural inertia and organizational change. *American Sociological Review*, 49(2), 149-164.

Hannan, M.T. and Freeman, J. (1989). *Organizational Ecology*. Cambridge, MA: Harvard University Press.

Hausman, D.M. (1992). *The Inexact and Separate Science of Economics*. Cambridge: Cambridge University Press.

Hayek, F.A. (1945). The use of knowledge in society. *American Economic Review*, 35(4), 519-530.

Hayek, F.A. (1967). *Studies in Philosophy, Politics and Economics*. London: Routledge & Kegan Paul.

Hayek, F.A. (1974). *The Pretence of Knowledge*. Nobel Prize Lecture.

Hayek, F.A. (1989). *The Pretence of Knowledge*. *American Economic Review*, 79(6), 3–7.

Hayek, F.A. (2014). Competition as a Discovery Procedure. In *The Market and Other Orders: The Collected Works of F.A. Hayek, Volume XV*. University of Chicago Press.

Hegel, G.W.F. (1807). *Phenomenology of Spirit*. [Phänomenologie des Geistes]. Translated by A.V. Miller (1977). Oxford: Oxford University Press.

Heidegger, M. (1927). *Being and Time*. [Sein und Zeit]. Translated by J. Macquarrie and E. Robinson (1962). New York: Harper & Row.

Helfat, C.E. and Peteraf, M.A. (2003). The dynamic resource-based view: Capability lifecycles. *Strategic Management Journal*, 24(10), 997-1010.

Helfat, C.E., Finkelstein, S., Mitchell, W., Peteraf, M.A., Singh, H., Teece, D.J. and Winter, S.G. (2007). *Dynamic Capabilities: Understanding Strategic Change in Organizations*. Malden, MA: Blackwell.

- Henderson, R.M. and Clark, K.B. (1990). Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms. *Administrative Science Quarterly*, 35(1), 9-30.
- Hendry, D.F. and Mizon, G.E. (2014). Unpredictability in economic analysis, econometric modeling and forecasting. *Journal of Econometrics*, 182(1), 176-195.
- Hendry, J. (2014). *Science and Contingency: A History of Science and Policy in the Modern Financial Industry*. New York: Polity.
- Hochberg, Y.V., Ljungqvist, A. and Lu, Y. (2007). Whom you know matters: Venture capital networks and investment performance. *Journal of Finance*, 62(1), 251-301.
- Hodgson, G.M. (1998). The approach of institutional economics. *Journal of Economic Literature*, 36(1), 166–192.
- Hodgson, G.M. (2001). *How Economics Forgot History: The Problem of Historical Specificity in Social Science*. London: Routledge.
- Hodgson, G.M. (2004). *The Evolution of Institutional Economics: Agency, Structure and Darwinism in American Institutionalism*. London: Routledge.
- Hodgson, G.M. (2006). What are institutions? *Journal of Economic Issues*, 40(1), 1–25.
- Hodgson, G.M. (2007). Meanings of methodological individualism. *Journal of Economic Methodology*, 14(2), 211-226.
- Hodgson, G.M. (2010). Choice, habit and evolution. *Journal of Evolutionary Economics*, 20(1), 1-18.
- Hodgson, G.M. and Knudsen, T. (2010). *Darwin's Conjecture: The Search for General Principles of Social and Economic Evolution*. Chicago: University of Chicago Press.
- Holland, J.H. (1992). *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence* (2nd ed.). Cambridge, MA: MIT Press.
- Holland, J.H. (1995). *Hidden Order: How Adaptation Builds Complexity*. Reading, MA: Addison-Wesley.
- Holland, J.H. (1998). *Emergence: From Chaos to Order*. Addison-Wesley.
- Holland, J.H. (2002). Complex adaptive systems. *Daedalus*, 121(1), 17-30.
- Holland, J.H. (2014). *Complexity: A Very Short Introduction*. Oxford University Press.
- House, P.J. (2010). Financial crisis: Lessons and policy responses. In A.N. Apergis (Ed.), *Financial Crises: Causes, Context, and Consequences*. New York: Nova Science.
- Hox, J.J. (2010). *Multilevel Analysis: Techniques and Applications* (2nd ed.). New York: Routledge.

Humphreys, P. (2016). *Emergence: A Philosophical Account*. Oxford: Oxford University Press.

Hynes, W., Lees, M. and Müller, J.M. (Eds.). (2021). *Systemic Thinking for Policy Making: The Potential of Systems Analysis for Addressing Global Policy Challenges in the 21st Century*. Paris: OECD Publishing.

Ireland, R.D., Hitt, M.A. and Sirmon, D.G. (2006). A model of strategic entrepreneurship: The construct and its dimensions. *Journal of Management*, 32(4), 587-602.

I

Isaacson, W. (2011). *Steve Jobs*. New York: Simon & Schuster.

J

Jackson, M.O. (2008). *Social and Economic Networks*. Princeton: Princeton University Press.

Jacobides, M.G., Cennamo, C. and Gawer, A. (2018). Towards a theory of ecosystems. *Strategic Management Journal*, 39(8), 2255-2276.

K

- Kahneman, D. (2011). *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.
- Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263-291.
- Kapoor, R. and Lee, J.M. (2013). Coordinating and competing in ecosystems: How organizational forms shape new technology investments. *Strategic Management Journal*, 34(3), 274-296.
- Katz, D. and Kahn, R.L. (1966). *The Social Psychology of Organizations*. New York: Wiley.
- Katz, M.L. and Shapiro, C. (1985). Network externalities, competition, and compatibility. *American Economic Review*, 75(3), 424-440.
- Kauffman, S.A. (1993). *The Origins of Order: Self-Organization and Selection in Evolution*. New York: Oxford University Press.
- Kauffman, S.A. (1995). *At Home in the Universe: The Search for the Laws of Self-Organization and Complexity*. Oxford University Press.
- Kaufmann, F. (1936). *Methodology of the Social Sciences*. Oxford: Oxford University Press.
- Kaufmann, F. (1942). On the subject-matter of economic science. *Economica*, 9(36), 381-401.
- Kay, J. and King, M. (2020). *Radical Uncertainty: Decision-Making Beyond the Numbers*. New York: W.W. Norton.
- Keynes, J.M. (1936). *The General Theory of Employment, Interest and Money*. London: Macmillan.
- Keynes, J.M. (1937). The general theory of employment. *Quarterly Journal of Economics*, 51(2), 209–223.
- Kim, J. (1999). Making sense of emergence. *Philosophical Studies*, 95(1-2), 3-36.
- Kirman, A. (1992). Whom or what does the representative individual represent? *Journal of Economic Perspectives*, 6(2), 117–136.
- Kirman, A. (2010). The economic crisis is a crisis for economic theory. *CESifo Economic Studies*, 56(4), 498-535.
- Kirman, A. (2011). *Complex Economics: Individual and Collective Rationality*. London: Routledge.
- Kirman, A. (2016). Complexity and economic policy: A paradigm shift or a change in perspective? *Journal of Economic Literature*, 54(2), 534-572.
- Kirman, A. (2017). The economy as a complex system. In B. Arthur, S. Durlauf and D. Lane (Eds.), *The Economy as an Evolving Complex System II* (pp. 491-531). New York: Routledge. [Reprint/update].

- Kitto, K. (2014). A contextualised general systems theory. *Systems*, 2(4), 541-565.
- Klein, P.G. (2010). *The Capitalist and the Entrepreneur: Essays on Organizations and Markets*. Auburn, AL: Mises Institute.
- Knight, F.H. (1921). *Risk, Uncertainty and Profit*. Houghton Mifflin.
- Knight, F.H. (1940). Preface to the re-issue. In *Risk, Uncertainty and Profit* (pp. xi-xxxvi). London: London School of Economics.
- Kogut, B. and Zander, U. (1992). Knowledge of the firm, combinative capabilities, and the replication of technology. *Organization Science*, 3(3), 383-397.
- Krugman, P. (2009). How did economists get it so wrong? *New York Times Magazine*, September 2.
- Kuhn, T.S. (1962). *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
- Kuhn, T.S. (1970). *The Structure of Scientific Revolutions* (2nd ed.). Chicago: University of Chicago Press.
- Kuiper, E. (2023). *A Different History of Economics: History of Economic Thought in Feminist Critical Perspective*. Cambridge: Cambridge University Press.
- Kurz, H.D. (Ed.). (2008). *The Theory of Value and Distribution in Economics: Discussions Between Pierangelo Garegnani and Paul Samuelson*. London: Routledge.

L

Ladyman, J. & Ross, D. (2007). *Every Thing Must Go: Metaphysics Naturalized*. Oxford University Press.

Ladyman, J., Lambert, J. & Wiesner, K. (2013). What is a complex system? *European Journal for Philosophy of Science*, 3(1), 33–67.

Lakatos, I. (1970). Falsification and the methodology of scientific research programmes. In I. Lakatos & A. Musgrave (Eds.), *Criticism and the Growth of Knowledge* (pp. 91–196). Cambridge University Press.

Lakatos, I. (1974). Falsification and the methodology of scientific research programmes. In I. Lakatos and A. Musgrave (Eds.), *Criticism and the Growth of Knowledge* (pp. 91-196). Cambridge: Cambridge University Press. [Revised edition].

Lakatos, I. (1978). *The Methodology of Scientific Research Programmes: Philosophical Papers Volume 1*. Cambridge: Cambridge University Press.

Lambert, R.A. (2001). Contracting theory and accounting. *Journal of Accounting and Economics*, 32(1-3), 3-87.

Landau, L.D. and Lifshitz, E.M. (1976). *Mechanics* (3rd ed.). Oxford: Pergamon Press.

Landau, L.D. and Lifshitz, E.M. (1980). *Statistical Physics, Part 1* (3rd ed.). Oxford: Pergamon Press.

Landini, S., Ferraris, G. and Gallegati, M. (2025). Non-linearity and complexity in economic systems. *Journal of Economic Dynamics and Control*. [Forthcoming].

Lane, D. and Maxfield, R. (2005). Ontological uncertainty and innovation. *Journal of Evolutionary Economics*, 15(1), 3-50.

Lane, D.A. (2005). Hierarchy, complexity, society. In D. Pumain (Ed.), *Hierarchy in Natural and Social Sciences* (pp. 81-119). Dordrecht: Springer.

Langley, A. (1999). Strategies for theorizing from process data. *Academy of Management Review*, 24(4), 691-710.

Langley, A., Smallman, C., Tsoukas, H. and Van de Ven, A.H. (2013). Process studies of change in organization and management: Unveiling temporality, activity, and flow. *Academy of Management Journal*, 56(1), 1-13.

Laplace, P.-S. (1814). *A Philosophical Essay on Probabilities*. Translated by F.W. Truscott and F.L. Emory (1902). New York: John Wiley & Sons.

Laplace, P.-S. (1902). *A Philosophical Essay on Probabilities*. Translated by F.W. Truscott and F.L. Emory. New York: John Wiley & Sons. [Original work published 1814].

Larkin, J.H. and Simon, H.A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, 11(1), 65-100.

- Laudan, L. (1977). *Progress and Its Problems: Towards a Theory of Scientific Growth*. University of California Press.
- Laughlin, R.B. (2005). *A Different Universe: Reinventing Physics from the Bottom Down*. New York: Basic Books.
- Lawson, T. (1988). Probability and uncertainty in economic analysis. *Journal of Post Keynesian Economics*, 11(1), 38-65.
- Lawson, T. (1997). *Economics and Reality*. London: Routledge.
- Lawson, T. (2003). *Reorienting Economics*. London: Routledge.
- Lawson, T. (2019). *The Nature of Social Reality: Issues in Social Ontology*. London: Routledge.
- Levinthal, D.A. (1997). Adaptation on rugged landscapes. *Management Science*, 43(7), 934-950.
- Levinthal, D.A. and March, J.G. (1993). The myopia of learning. *Strategic Management Journal*, 14(S2), 95-112.
- Lewin, K. (1951). *Field Theory in Social Science: Selected Theoretical Papers*. Edited by D. Cartwright. New York: Harper.
- Lewin, R. (1992). *Complexity: Life at the Edge of Chaos*. New York: Macmillan.
- Little, D. (1991). *Varieties of Social Explanation: An Introduction to the Philosophy of Social Science*. Boulder, CO: Westview Press.
- Little, D. (1993). On the scope and limits of generalizations in the social sciences. *Synthese*, 97(2), 183-207.
- Locke, K. (1996). Rewriting the discovery of grounded theory after 25 years? *Journal of Management Inquiry*, 5(3), 239-245.
- Locke, K. and Golden-Biddle, K. (1997). Constructing opportunities for contribution: Structuring intertextual coherence and 'problematizing' in organizational studies. *Academy of Management Journal*, 40(5), 1023-1062.
- Lucas, R.E. (1976). Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy*, 1, 19-46.

M

Machlup, F. (1955). The problem of verification in economics. *Southern Economic Journal*, 22(1), 1-21.

Machlup, F. (1967). Theories of the firm: Marginalist, behavioral, managerial. *American Economic Review*, 57(1), 1-33.

Mahoney, J. and Thelen, K. (Eds.). (2010). *Explaining Institutional Change: Ambiguity, Agency, and Power*. Cambridge: Cambridge University Press.

Mahoney, J.T. and Thelen, K. (2010). *Explaining Institutional Change: Ambiguity, Agency, and Power*. Cambridge: Cambridge University Press.

Mäki, U. (1992). On the method of isolation in economics. *Poznan Studies in the Philosophy of the Sciences and the Humanities*, 26, 319-354.

Malerba, F. (2002). Sectoral systems of innovation and production. *Research Policy*, 31(2), 247-264.

Malerba, F. and Orsenigo, L. (1997). Technological regimes and sectoral patterns of innovative activities. *Industrial and Corporate Change*, 6(1), 83-117.

Mandelbrot, B. and Hudson, R.L. (2004). *The (Mis)Behavior of Markets: A Fractal View of Financial Turbulence*. New York: Basic Books.

Manski, C.F. (1993). Identification of endogenous social effects: The reflection problem. *Review of Economic Studies*, 60(3), 531-542.

March, J.G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2(1), 71-87.

March, J.G. and Olsen, J.P. (1989). *Rediscovering Institutions: The Organizational Basis of Politics*. New York: Free Press.

March, J.G. and Simon, H.A. (1958). *Organizations*. New York: Wiley.

Max-Neef, M.A. (2005). Foundations of transdisciplinarity. *Ecological Economics*, 53(1), 5-16.

Maxfield, R.R. (2005). General equilibrium and the theory of directed graphs. *Journal of Mathematical Economics*, 8(1), 27-51.

May, R.M. and Arinaminpathy, N. (2011). Systemic risk: The dynamics of model banking systems. *Journal of the Royal Society Interface*, 7(46), 823-838.

Meadows, D.H. (2008). *Thinking in Systems: A Primer*. White River Junction, VT: Chelsea Green Publishing.

Metcalf, J.S. (2005). Evolutionary concepts in relation to evolutionary economics. In K. Dopfer (Ed.), *The Evolutionary Foundations of Economics* (pp. 391-430). Cambridge: Cambridge University Press.

- Metcalfe, J.S. (2008). Accounting for economic evolution: Fitness and the population method. *Journal of Bioeconomics*, 10(1), 23-49.
- Michaels, E., Handfield-Jones, H. and Axelrod, B. (2001). *The War for Talent*. Boston: Harvard Business School Press.
- Milgrom, P. and Roberts, J. (1990). The economics of modern manufacturing: Technology, strategy, and organization. *American Economic Review*, 80(3), 511-528.
- Milgrom, P. and Roberts, J. (1992). *Economics, Organization and Management*. Englewood Cliffs, NJ: Prentice-Hall.
- Mill, J.S. (1843). *A System of Logic, Ratiocinative and Inductive*. London: John W. Parker.
- Miller, J.G. (1978). *Living Systems*. New York: McGraw-Hill.
- Miller, J.H. and Page, S.E. (2007). *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton: Princeton University Press.
- Minsky, H.P. (1986). *Stabilizing an Unstable Economy*. New Haven: Yale University Press.
- Mintzberg, H. (1978). Patterns in strategy formation. *Management Science*, 24(9), 934-948.
- Mintzberg, H. (1979). *The Structuring of Organizations*. Englewood Cliffs, NJ: Prentice-Hall.
- Mintzberg, H. (1994). *The Rise and Fall of Strategic Planning*. New York: Free Press.
- Mintzberg, H. and Waters, J.A. (1985). Of strategies, deliberate and emergent. *Strategic Management Journal*, 6(3), 257-272.
- Mirowski, P. (1989). *More Heat Than Light: Economics as Social Physics, Physics as Nature's Economics*. Cambridge University Press.
- Mirowski, P. (2002). *Machine Dreams: Economics Becomes a Cyborg Science*. Cambridge: Cambridge University Press.
- Mitchell, M. (2009). *Complexity: A Guided Tour*. Oxford University Press.
- Mitchell, S.D. (2002). Integrative pluralism. *Biology and Philosophy*, 17(1), 55-70.
- Mitchell, S.D. (2003). *Biological Complexity and Integrative Pluralism*. Cambridge University Press.
- Mizon, G.E. (2014). Econometric methods and the evaluation of economic forecasts. In G. Elliott and A. Timmermann (Eds.), *Handbook of Economic Forecasting* (Vol. 2, pp. 693-712). Amsterdam: Elsevier.
- Moon, J.D. (1999). Practical discourse and communicative ethics. In S.K. White (Ed.), *The Cambridge Companion to Habermas* (pp. 143-164). Cambridge: Cambridge University Press.
- Morin, E. (2008). *On Complexity*. Cresskill, NJ: Hampton Press.

Murphy, N.C. and Ellis, G.F.R. (2002). *On the Moral Nature of the Universe: Theology, Cosmology, and Ethics*. Minneapolis: Fortress Press.

Musgrave, A. (1981). 'Unreal assumptions' in economic theory: The F-twist untwisted. *Kyklos*, 34(3), 377-387.

Myers, M.D. (1993). Hermeneutics in information systems research. In J. Mingers and L. Willcocks (Eds.), *Social Theory and Philosophy for Information Systems* (pp. 103-128). Chichester: Wiley.

N

- Nason, R.S. and Wiklund, J. (2017). An assessment of resource-based theorizing on firm growth and suggestions for the future. *Journal of Management*, 44(1), 32-60.
- Nelson, R.R. (ed.) (1993). *National Innovation Systems: A Comparative Analysis*. New York: Oxford University Press.
- Nelson, R.R. and Winter, S.G. (1982). *An Evolutionary Theory of Economic Change*. Cambridge, MA: Harvard University Press.
- Newell, W.H. (2010). Undergraduate general education. In R. Frodeman, J.T. Klein, and C. Mitcham (Eds.), *The Oxford Handbook of Interdisciplinarity* (pp. 360-371). Oxford: Oxford University Press.
- Newman, M.E.J. (2002). Assortative mixing in networks. *Physical Review Letters*, 89(20), 208701.
- Newman, M.E.J. (2006). Modularity and community structure in networks. *Proceedings of the National Academy of Sciences*, 103(23), 8577-8582.
- Newman, M.E.J. (2018). *Networks* (2nd ed.). Oxford: Oxford University Press.
- Newman, M.E.J. (2019). *Networks* (2nd ed., corrected printing). Oxford: Oxford University Press.
- Nickles, T. (1997). A multi-pass conception of discovery. In I. Kuby (Ed.), *Scientific Discovery, Logic, and Rationality* (pp. 105-135). Dordrecht: Springer.
- Nicolescu, B. (2002). *Manifesto of Transdisciplinarity*. Albany, NY: SUNY Press.
- Nicolis, G. and Prigogine, I. (1977). *Self-Organization in Nonequilibrium Systems: From Dissipative Structures to Order Through Fluctuations*. New York: Wiley.
- Nohria, N. and Eccles, R.G. (eds.) (1992). *Networks and Organizations: Structure, Form, and Action*. Boston: Harvard Business School Press.
- North, D.C. (1990). *Institutions, Institutional Change and Economic Performance*. Cambridge: Cambridge University Press.
- North, D.C. (1999). Dealing with a non-ergodic world: Institutional economics, property rights, and the global environment. *Duke Environmental Law and Policy Forum*, 10(1), 1-12.
- North, D.C. (2005). *Understanding the Process of Economic Change*. Princeton University Press.

O

Oliver, C. (1991). Strategic responses to institutional processes. *Academy of Management Review*, 16(1), 145-179.

P

Padgett, J.F. and Powell, W.W. (2012). *The Emergence of Organizations and Markets*. Princeton: Princeton University Press.

Page, S.E. (2007). *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton: Princeton University Press.

Paniagua, P. (2023). The economic calculation debate: New lessons for the profession. *Journal of Economic Methodology*, 30(2), 93-109.

Patton, M.Q. (1999). Enhancing the quality and credibility of qualitative analysis. *Health Services Research*, 34(5), 1189-1208.

Pearl, J. (2000). *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press.

Penrose, E.T. (1959). *The Theory of the Growth of the Firm*. Oxford: Basil Blackwell.

Penrose, E.T. (2009). *The Theory of the Growth of the Firm* (4th ed., 50th anniversary edition). Oxford: Oxford University Press.

Pentland, B.T. and Feldman, M.S. (2005). Organizational routines as a unit of analysis. *Industrial and Corporate Change*, 14(5), 793-815.

Pisano, G.P. (2006). *Science Business: The Promise, the Reality, and the Future of Biotech*. Boston: Harvard Business School Press.

Polanyi, K. (1944). *The Great Transformation: The Political and Economic Origins of Our Time*. New York: Farrar & Rinehart.

Polanyi, M. (1958). *Personal Knowledge: Towards a Post-Critical Philosophy*. Chicago: University of Chicago Press.

Polanyi, M. (1966). *The Tacit Dimension*. Garden City, NY: Doubleday.

Poli, R. (2010). The many aspects of anticipation. *Foresight*, 12(3), 7–17.

Poli, R. (2013). A note on the difference between complicated and complex social systems. *Cadmus*, 2(1), 142-147.

Popper, K.R. (1963). *Conjectures and Refutations: The Growth of Scientific Knowledge*. London: Routledge & Kegan Paul.

Porter, M.E. (1980). *Competitive Strategy: Techniques for Analyzing Industries and Competitors*. New York: Free Press.

Porter, M.E. (1985). *Competitive Advantage: Creating and Sustaining Superior Performance*. New York: Free Press.

Potts, J. (2000). *The New Evolutionary Microeconomics: Complexity, Competence and Adaptive Behaviour*. Cheltenham: Edward Elgar.

Potts, J. (2004). Liberty bubbles. *Policy*, 20(3), 15-21.

Powell, W.W. (1990). Neither market nor hierarchy: Network forms of organization. *Research in Organizational Behavior*, 12, 295-336.

Powell, W.W. and DiMaggio, P.J. (eds.) (1991). *The New Institutionalism in Organizational Analysis*. Chicago: University of Chicago Press.

Powell, W.W., Koput, K.W. and Smith-Doerr, L. (1996). Interorganizational collaboration and the locus of innovation: Networks of learning in biotechnology. *Administrative Science Quarterly*, 41(1), 116-145.

Prigogine, I. (1980). *From Being to Becoming: Time and Complexity in the Physical Sciences*. W.H. Freeman.

Prigogine, I. (1996). *The End of Certainty: Time, Chaos, and the New Laws of Nature*. New York: Free Press. [Also published 1997].

Prigogine, I. (1997). *The End of Certainty: Time, Chaos and the New Laws of Nature*. Free Press.

Prigogine, I. (2003). *Is Future Given?* Singapore: World Scientific.

Prigogine, I. and Stengers, I. (1984). *Order Out of Chaos: Man's New Dialogue with Nature*. New York: Bantam Books.

Prigogine, I. and Stengers, I. (1984). *Order Out of Chaos: Mans New Dialogue with Nature*. New York: Bantam Books.

Q

Quine, W.V.O. (1951). Two dogmas of empiricism. *Philosophical Review*, 60(1), 20-43.

R

Ragin, C.C. (2008). *Redesigning Social Inquiry: Fuzzy Sets and Beyond*. Chicago: University of Chicago Press.

Reagans, R. and McEvily, B. (2003). Network structure and knowledge transfer: The effects of cohesion and range. *Administrative Science Quarterly*, 48(2), 240-267.

Rescher, N. (1993). *Pluralism: Against the Demand for Consensus*. Oxford University Press.

Rescher, N. (1996). *Process Metaphysics: An Introduction to Process Philosophy*. SUNY Press.

Rescher, N. (2000). *Process Philosophy: A Survey of Basic Issues*. Pittsburgh: University of Pittsburgh Press.

Richardson, G.P. (1991). *Feedback Thought in Social Science and Systems Theory*. University of Pennsylvania Press.

Richardson, K.A. and Cilliers, P. (2001). What is complexity science? A view from different directions. *Emergence*, 3(1), 5-23.

Ricoeur, P. (1981). *Hermeneutics and the Human Sciences*. Edited and translated by J.B. Thompson. Cambridge: Cambridge University Press.

Ritzer, G. (1991). *Metatheorizing in Sociology*. Lexington, MA: Lexington Books.

Ritzer, G. (1993). *The McDonaldization of Society*. Thousand Oaks, CA: Pine Forge Press.

Robinson, H. and Weir, R. (2024). Substance. In E.N. Zalta and U. Nodelman (Eds.), *Stanford Encyclopedia of Philosophy*. Stanford: Metaphysics Research Lab.

Robinson, J.A. (2024). Power, politics, and state capacity. In *The Handbook of Economic Development* (pp. forthcoming). Cambridge: Cambridge University Press.

Rodrik, D. (2015). *Economics Rules: The Rights and Wrongs of the Dismal Science*. New York: W.W. Norton.

Romer, P.M. (1990). Endogenous technological change. *Journal of Political Economy*, 98(5), S71-S102.

Runde, J. (1998). Clarifying Frank Knight's discussion of the meaning of risk and uncertainty. *Cambridge Journal of Economics*, 22(5), 539-546.

S

Sandberg, J. (2022). Problematization and the construction of theoretical categories. *Organization Studies*, 43(2), 269-290.

Sandberg, J. (2024). Theory development through problematization. *Academy of Management Review*. [Forthcoming or recent].

Sandberg, J. and Alvesson, M. (2020). Meanings of theory: Clarifying theory through typification. *Journal of Management Studies*, 58(2), 487-516.

Sandberg, J. and Tsoukas, H. (2011). Grasping the logic of practice: Theorizing through practical rationality. *Academy of Management Review*, 36(2), 338-360.

Santos, F.M. and Eisenhardt, K.M. (2005). Organizational boundaries and theories of organization. *Organization Science*, 16(5), 491-508.

Sawyer, R.K. (2005). *Social Emergence: Societies as Complex Systems*. Cambridge: Cambridge University Press.

Saxenian, A. (1994). *Regional Advantage: Culture and Competition in Silicon Valley and Route 128*. Cambridge, MA: Harvard University Press.

Sayer, A. (1992). *Method in Social Science: A Realist Approach* (2nd ed.). London: Routledge.

Sayer, A. (2000). *Realism and Social Science*. London: Sage.

Scheffer, M. (2009). *Critical Transitions in Nature and Society*. Princeton University Press.

Schelling, T.C. (1978). *Micromotives and Macrobehavior*. New York: Norton.

Schickore, J. (2015). Scientific discovery. *Stanford Encyclopedia of Philosophy*.

Schilling, M.A. (2002). Technology success and failure in winner-take-all markets: The impact of learning orientation, timing, and network externalities. *Academy of Management Journal*, 45(2), 387-398.

Schoemaker, P.J.H. (1995). Scenario planning: A tool for strategic thinking. *Sloan Management Review*, 36(2), 25-40.

Schofield, J.W. (1983). Increasing the generalizability of qualitative research. In E.W. Eisner and A. Peshkin (Eds.), *Qualitative Inquiry in Education* (pp. 201-232). New York: Teachers College Press.

Schumpeter, J.A. (1934). *The Theory of Economic Development: An Inquiry into Profits, Capital, Credit, Interest, and the Business Cycle*. Cambridge, MA: Harvard University Press.

Schumpeter, J.A. (1942). *Capitalism, Socialism and Democracy*. New York: Harper & Brothers.

Seibt, J. (2023). Process philosophy. *Stanford Encyclopedia of Philosophy*.

- Senge, P.M. (1990). *The Fifth Discipline: The Art and Practice of the Learning Organization*. New York: Doubleday/Currency.
- Senge, P.M. (2006). *The Fifth Discipline: The Art and Practice of the Learning Organization* (Rev. ed.). New York: Doubleday.
- Shackle, G.L.S. (1972). *Epistemics and Economics: A Critique of Economic Doctrines*. Cambridge: Cambridge University Press.
- Shapiro, C. and Varian, H.R. (1999). *Information Rules: A Strategic Guide to the Network Economy*. Boston: Harvard Business School Press.
- Siggelkow, N. (2002). Evolution toward fit. *Administrative Science Quarterly*, 47(1), 125-159.
- Siggelkow, N. (2007). Persuasion with case studies. *Academy of Management Journal*, 50(1), 20-24.
- Siggelkow, N. and Levinthal, D.A. (2003). Temporarily divide to conquer: Centralized, decentralized, and reintegrated organizational approaches to exploration and adaptation. *Organization Science*, 14(6), 650-669.
- Siggelkow, N. and Rivkin, J.W. (2005). Speed and search: Designing organizations for turbulence and complexity. *Organization Science*, 16(2), 101-122.
- Simon, H.A. (1947). *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organization*. New York: Macmillan.
- Simon, H.A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69(1), 99-118.
- Simon, H.A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society*, 106(6), 467-482.
- Simon, H.A. (1979). Rational Decision Making in Business Organizations. *American Economic Review*, 69(4), 493-513.
- Simon, H.A. (1983). *Reason in Human Affairs*. Stanford University Press.
- Simon, H.A. (1997). *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organizations* (4th ed.). Free Press.
- Singh, J.V., House, R.J. and Tucker, D.J. (1986). Organizational change and organizational mortality. *Administrative Science Quarterly*, 31(4), 587-611.
- Skinner, Q. (1969). Meaning and understanding in the history of ideas. *History and Theory*, 8(1), 3-53.
- Slife, B.D. and Williams, R.N. (1995). *What's Behind the Research? Discovering Hidden Assumptions in the Behavioral Sciences*. Thousand Oaks, CA: Sage.
- Slovik, P. and Cournède, B. (2011). Macroeconomic impact of Basel III. OECD Economics Department Working Papers No. 844.

Small, M.L. (2011). How to conduct a mixed methods study: Recent trends in a rapidly growing literature. *Annual Review of Sociology*, 37, 57-86.

Solow, R.M. (2010). Building a science of economics for the real world. Prepared statement to the House Committee on Science and Technology, July 20.

Soros, G. (2009). The Crash of 2008 and What It Means: The New Paradigm for Financial Markets. PublicAffairs.

Soros, G. (2013). Fallibility, reflexivity, and the human uncertainty principle. *Journal of Economic Methodology*, 20(4), 309-329.

Spear, S. and Bowen, H.K. (1999). Decoding the DNA of the Toyota production system. *Harvard Business Review*, 77(5), 96-106.

Sperry, R.W. (1986). Macro- versus micro-determinism. *Philosophy of Science*, 53(2), 265-270.

Stacey, R.D. (1995). The science of complexity: An alternative perspective for strategic change processes. *Strategic Management Journal*, 16(6), 477-495.

Star, S.L. and Griesemer, J.R. (1989). Institutional ecology, 'translations' and boundary objects: Amateurs and professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39. *Social Studies of Science*, 19(3), 387-420.

Staw, B.M. and Epstein, L.D. (2000). What bandwagons bring: Effects of popular management techniques on corporate performance, reputation, and CEO pay. *Administrative Science Quarterly*, 45(3), 523-556.

Sterman, J.D. (2000). *Business Dynamics: Systems Thinking and Modeling for a Complex World*. Boston: Irwin McGraw-Hill.

Stiglitz, J.E. (2018). Where modern macroeconomics went wrong. *Oxford Review of Economic Policy*, 34(1-2), 70-106.

Stiglitz, J.E. (2019). *People, Power, and Profits: Progressive Capitalism for an Age of Discontent*. New York: W.W. Norton.

Stone, B. (2013). *The Everything Store: Jeff Bezos and the Age of Amazon*. New York: Little, Brown.

Storm, S. (2017). The new normal: Demand, secular stagnation, and the vanishing middle class. *International Journal of Political Economy*, 46(4), 169-210.

Storm, S. (2021). Cordon of conformity: Why DSGE models are not the future of macroeconomics. *International Journal of Political Economy*, 50(2), 77-98.

Strathern, M. (1997). 'Improving ratings': Audit in the British university system. *European Review*, 5(3), 305-321.

Stuart, T.E. (2000). Interorganizational alliances and the performance of firms: A study of growth and innovation rates in a high-technology industry. *Strategic Management Journal*, 21(8), 791-811.

Suarez, F.F. and Lanzolla, G. (2005). The half-truth of first-mover advantage. *Harvard Business Review*, 83(4), 121-127.

Suarez, F.F. and Utterback, J.M. (1995). Dominant designs and the survival of firms. *Strategic Management Journal*, 16(6), 415-430.

Sutton, R.I. and Staw, B.M. (1995). What theory is not. *Administrative Science Quarterly*, 40(3), 371-384.

Sydow, J., Schreyögg, G. and Koch, J. (2009). Organizational path dependence: Opening the black box. *Academy of Management Review*, 34(4), 689-709.

T

- Taleb, N.N. (2007). *The Black Swan: The Impact of the Highly Improbable*. Random House.
- Taleb, N.N. (2012). *Antifragile: Things That Gain from Disorder*. Random House.
- Taylor, F.W. (1911). *The Principles of Scientific Management*. New York: Harper & Brothers.
- Teece, D.J. (2007). Explicating dynamic capabilities: The nature and microfoundations of (sustainable) enterprise performance. *Strategic Management Journal*, 28(13), 1319–1350.
- Teece, D.J., Pisano, G. and Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, 18(7), 509-533.
- Tesfatsion, L. and Judd, K.L. (Eds.). (2006). *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*. Amsterdam: Elsevier.
- Thompson, J.D. (1967). *Organizations in Action: Social Science Bases of Administrative Theory*. New York: McGraw-Hill.
- Toulmin, S. (1958). *The Uses of Argument*. Cambridge: Cambridge University Press.
- Tripsas, M. and Gavetti, G. (2000). Capabilities, cognition, and inertia: Evidence from digital imaging. *Strategic Management Journal*, 21(10-11), 1147-1161.
- Tsai, W. (2001). Knowledge transfer in intraorganizational networks: Effects of network position and absorptive capacity on business unit innovation and performance. *Academy of Management Journal*, 44(5), 996-1004.
- Tsoukas, H. (2005). *Complex Knowledge: Studies in Organizational Epistemology*. Oxford: Oxford University Press.
- Tsoukas, H. and Chia, R. (2002). On organizational becoming: Rethinking organizational change. *Organization Science*, 13(5), 567-582.
- Tuma, N.B. and Hannan, M.T. (1984). *Social Dynamics: Models and Methods*. Orlando: Academic Press.
- Tushman, M.L. and Anderson, P. (1986). Technological discontinuities and organizational environments. *Administrative Science Quarterly*, 31(3), 439-465.
- Tversky, A. and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124-1131.

U

Utterback, J.M. (1994). *Mastering the Dynamics of Innovation: How Companies Can Seize Opportunities in the Face of Technological Change*. Boston: Harvard Business School Press.

Utterback, J.M. and Abernathy, W.J. (1975). A dynamic model of process and product innovation. *Omega*, 3(6), 639-656.

Uzzi, B. (1997). Social structure and competition in interfirm networks: The paradox of embeddedness. *Administrative Science Quarterly*, 42(1), 35-67.

V

Van de Ven, A.H. (2007). *Engaged Scholarship: A Guide for Organizational and Social Research*. Oxford: Oxford University Press.

Van de Ven, A.H. and Poole, M.S. (2016). Process research methods. In A. Langley and H. Tsoukas (Eds.), *The SAGE Handbook of Process Organization Studies* (pp. 477-490). London: Sage.

Van Maanen, J. (2011). *Tales of the Field: On Writing Ethnography* (2nd ed.). Chicago: University of Chicago Press.

Veblen, T. (1898). Why is economics not an evolutionary science? *Quarterly Journal of Economics*, 12(4), 373-397.

Veblen, T. (1899). *The Theory of the Leisure Class: An Economic Study of Institutions*. New York: Macmillan.

Vergne, J.-P. and Durand, R. (2010). The missing link between the theory and empirics of path dependence: Conceptual clarification, testability issue, and methodological implications. *Journal of Management Studies*, 47(4), 736-759.

von Bertalanffy, L. (1968). *General System Theory: Foundations, Development, Applications*. New York: George Braziller.

Von Hippel, E. (1988). *The Sources of Innovation*. New York: Oxford University Press.

von Mises, L. (1949). *Human Action: A Treatise on Economics*. New Haven: Yale University Press.

W

- Wacker, J.G. (1998). A definition of theory: Research guidelines for different theory-building research methods in operations management. *Journal of Operations Management*, 16(4), 361-385.
- Wasserman, S. and Faust, K. (1994). *Social Network Analysis: Methods and Applications*. Cambridge: Cambridge University Press.
- Watts, D.J. (1999). *Small Worlds: The Dynamics of Networks between Order and Randomness*. Princeton University Press.
- Watts, D.J. (2004). *Six Degrees: The Science of a Connected Age*. New York: W.W. Norton.
- Watts, D.J. and Strogatz, S.H. (1998). Collective dynamics of 'small-world' networks. *Nature*, 393(6684), 440-442.
- Weick, K.E. (1979). *The Social Psychology of Organizing* (2nd ed.). Reading, MA: Addison-Wesley.
- Weick, K.E. (1995). *Sensemaking in Organizations*. Thousand Oaks, CA: Sage Publications.
- Weidlich, W. (2008). *Sociodynamics: A Systematic Approach to Mathematical Modelling in the Social Sciences*. Mineola, NY: Dover.
- Welch, J.F. (1981). Growing fast in a slow-growth economy. Speech at Hotel Pierre, New York.
- West, J. and Gallagher, S. (2006). Challenges of open innovation: The paradox of firm investment in open-source software. *R&D Management*, 36(3), 319-331.
- Whetten, D.A. (1989). What constitutes a theoretical contribution? *Academy of Management Review*, 14(4), 490-495.
- White, H. (2018). *Markets from Networks: Socioeconomic Models of Production*. Princeton: Princeton University Press.
- Whitehead, A.N. (1929). *Process and Reality: An Essay in Cosmology*. Macmillan.
- Williamson, O.E. (1975). *Markets and Hierarchies: Analysis and Antitrust Implications*. New York: Free Press.
- Williamson, O.E. (1985). *The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting*. New York: Free Press.
- Williamson, O.E. (1991). Comparative economic organization: The analysis of discrete structural alternatives. *Administrative Science Quarterly*, 36(2), 269-296.
- Williamson, O.E. (2000). The new institutional economics: Taking stock, looking ahead. *Journal of Economic Literature*, 38(3), 595-613.

Winter, S.G. (1982). An essay on the theory of production. In S.H. Hymans (Ed.), *Economics and the World Around It* (pp. 55-91). Ann Arbor: University of Michigan Press.

Winter, S.G. (2003). Understanding dynamic capabilities. *Strategic Management Journal*, 24(10), 991-995.

Witt, U. (2001). Learning to consume – A theory of wants and the growth of demand. *Journal of Evolutionary Economics*, 11(1), 23–36.

Witt, U. (2008). What is specific about evolutionary economics? *Journal of Evolutionary Economics*, 18(5), 547-575.

Wolfram, S. (2002). *A New Kind of Science*. Wolfram Media.

Woodward, J. (1965). *Industrial Organization: Theory and Practice*. London: Oxford University Press.

Woodward, J. (2003). *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.

Y

Yin, R.K. (2018). *Case Study Research and Applications: Design and Methods* (6th ed.). Thousand Oaks, CA: Sage.

Z

Zaheer, A. and Bell, G.G. (2005). Benefiting from network position: Firm capabilities, structural holes, and performance. *Strategic Management Journal*, 26(9), 809-825.

Zenger, T.R. and Hesterly, W.S. (1997). The disaggregation of corporations: Selective intervention, high-powered incentives, and molecular units. *Organization Science*, 8(3), 209-222.

Zhu, F. and Iansiti, M. (2012). Entry into platform-based markets. *Strategic Management Journal*, 33(1), 88-106.