UNIVERSITY OF PIRAEUS
DEPARTMENT OF BANKING & FINANCIAL MANAGEMENT

'Trading spreads in implied volatility indices'

Msc Project

Rigkou L. Eleni

Supervisor: George Skiadopoulos
Assistant Professor

Athens
June 2009

Table of Contents

Acknowledgements

I would like to take the opportunity to thank my supervisor, Professor George Skiadopoulos, for the suggestions, the encouragement and all the support he has given me during this project.

I would also like to thank Eirini Konstantinidi, Ph.D student of the department of Banking and Financial Management, for all the information she supplied and the help she was willing to give.

I would also like to thank my mother Maria and all my friends for the moral support and encouragement they have given me throughout the year.

1. Introduction

Implied volatility is a theoretical measure which comes from option pricing models such as Black-Scholes, rather than using historical data and it is affected by the option’s strike price, the riskless rate of return, the time to maturity and the price of the option. Options are wonderfully powerful tools for managing risk. When used in a prudent, risk-controlled fashion, they can reduce portfolio volatility. So, this measure is important as it allows funds to develop trading and hedging strategies based upon it. Many studies have tried to model implied volatility indices in order to produce useful forecasts for the traders.

In this paper, in an effort to differentiate, we are going to explore the dynamics of the implied volatility *spread*, meaning the difference between two implied volatility indices, by studying two US implied volatility indices, the CBOE VIX and VXD. We will see the predictable power of two models- an economic variables model and an AR(1) model-, form point and interval forecasts and counting on these forecasts of the models we will employ trading strategies based on VIX and VXD volatility futures. The trading rule will concern the simultaneously taking position in VIX and VXD volatility futures.

In the next section, in retrospect, we put forward different studies based upon implied volatility and its dynamics. In section 3, we describe the data set. In section 4, we present the models to be used for forecasting. Section 5 shows the in-sample evidence. In sections 5 and 6, we assess the out-of-sample performance of the models and examine whether the trading strategies generate economic profits, respectively. Finally, in section 8, we present our results.

2. Literature review

Last years, we see an attempt to model the implied volatility, by using different models and subsequently a variety of trading strategies, mainly for profitable reasons. In the following, we cite some interesting studies.

Ahoniemi (2006) models the implied volatility of the S&P 500 index in order to forecast the future direction of implied volatility, useful for option traders. According to the study, VIX index is high persistent, seems to be skewed to the right and exhibits excess kurtosis. ARIMA models are augmented with GARCH errors and exogenous regressors so as to model VIX. Among the regressors are the returns in the S&P 500 index, the historical volatility of the S&P500, the trading volume of the S&P 500 index, the MSCI EAFE index, the three-month USD LIBOR interest rate, the 10-year U.S. government bond yield and the price of crude oil from the next expiring futures contract. The study proves that negative shocks have a greater impact on VIX than the positive ones. Apart from ARIMA models, binary probit models are estimated but the first outperformed the latter both in in-sample performance and out-of-sample. In oreder to assess the economic significance of the VIX forecasts, option trading strategies are formed. More specifically, volatility spreads, such as straddles are used (buy or sell an equal amount of call and put options – care only for the volatility of the underlying asset price and not for its direction) and yield positive profits, without, though, taking into account transaction costs. There is a certain degree of predictability in the direction of change of the VIX index that can be exploited profitably by option traders.

Konstantinidi et al. (2008) examine whether the evolution of implied volatility can be forecasted by using a number of European and US implied volatility indices, such as VIX, VXO, VXN, VXD (American implied volatility indices) and VDAX-New, VCAC and VSTOXX (European implied volatility indices). They use different models in order to evaluate the in-sample, out-of-sample performance and economic significance. Among the models are the economic variables model, univariate autoregressive and VAR models, the principal components model and ARIMA/ARFIMA models. The VAR and the PCA models perform best. The changes in the implied volatility index do not exhibit long memory. In the out-of-sample performance, they form both point and interval forecasts. At point forecasts, in 28% of the cases one of the models performs better than the random walk, indicating a statistically predictable pattern (stronger for European indices). At interval forecasts, there is also observed a predictable pattern (especially for the US indices) but no single model yields accurate forecasts. The economic significance is tested by forming trading strategies based on VIX and VXD futures. However, the generated point and interval forecasts are not economically significant.

Goncalves and Guidolin (2006) study whether the resulting predictability in the implied volatility surface[1] can be exploited. They model jointly the cross-sectional features and the dynamics of the implied volatility surface of S&P 500 index options. Firstly, they fit daily cross-sectional models that describe implied volatilities as a function of moneyness and time to maturity and secondly, they fit time-series models of a VAR type to capture the presence of time variation in the first step estimated coefficients. Comparing their model with other benchmarks (such as random walk and NGARCH(1,1) of Heston and Nandi (2000)), they find that the forecast accuracy of the model is good, the forecasts ability to support portfolio decisions display satisfactory performance and that simulated delta-hedged trading strategies generate positive out-of-sample returns when low transaction costs are taken into account. Moreover, their strategy is more profitable in out-of-the-money and short to medium term contracts. However, when transaction costs are raised to higher levels, all profits disappear.

Poon and Pope (2000) examine whether there are common components in the volatilities on S&P 100 and S&P 500. They support that if two assets share common factors, the respective volatilities of the two assets will be related. Among their findings is that volatility on those indices exhibits strong persistence, long memory. They use options written on the S&P 100 (OEX) and S&P 500 (SPX), whom find the implied standard deviation (isd), in order to create the volatility spread, which in this case defines as the quotient of isd_{OEX} to isd_{SPX}. The volatility spread has thin tails, pointing out that volatilities of the two indices share common factors. In an effort to test the joint efficiency of the two markets, they adopt a trading rule which includes vega-delta-neutral[2] hedge positions. When the volatility spread (as defined above) is too high, the investor sell OEX calls and buys SPX calls and conversely. The strategy

[1] Implied volatility surface is a plot which depicts the relation of implied volatility with time to maturity and the strike price. In index options, its is observed the volatility smile, meaning that implied volatility is lower for at-the-money options and higher for in- or out-of the money options.

[2] Vega neutrality protects against variations in volatility σ. Delta neutrality provides protection against relatively small stock price moves between hedge rebalancing.

based upon the volatility spread seems to generate profits even if transaction costs are taken into account. However, the trading rule concludes that OEX and SPX markets are not jointly efficient.

Aboura (1999) examine whether there are interactions between returns and implied volatility and between volatilities belonging to different markets. The implied volatility indices under consideration are the US VIX, the German VDAX and the French VX1. Causality tests are used in order to investigate if the markets are integrated and if there is a lead-lag relation among the markets. The German market influence the French market and US market for 1 day lag and the US market can predict the German and French market for 1 day lag. They also show that implied volatility and future realized volatility do not always move in the same way and that volatility is more sensitive to negative shocks than positive ones. Another test they employ is that use model to assess the reaction of the implied volatility indices around an extreme event. The French VX1 is more sensitive to shock and the influence of events which raise volatility on this index is less persistent than VIX and VDAX. Finally, the study confirms that VIX is dominant upon the other two European indices.

Fernandes, Medeiros and Scharth (2007) exploit the time series properties of VIX resulting in the long range dependence in the VIX time series. They use ARMA, ARFIMA[3], several HAR-type processes and smooth transition autoregressive trees for modeling and forecasting purposes. They found a strong, negative, asymmetric relation between VIX and contemporaneous index returns. In addition, VIX time series seem to be skewed to the right, leptokurtic and far from Gaussian. As far as the ou-of-sample performance is concerned, ARMA and ARFIMA models are better in the short run, rather than nonlinear START model which outperform for their predictive ability in the long run.

Franks and Schwartz examine the time series properties of volatility, various economic variables and if they are significant with expected volatility and explore if there is any relationship between capital structure and changes in volatility. They find that leverage is a significant explanatory variable, as inflation too.

Harvey and Whaley (1991) study the time variation in implied volatility of S&P 100 option prices (OEX). They test the hypothesis that market volatility is

---

[3] Models that are capable of dealing with long memory time-series such as VIX index.

unpredictable; however, according their results, it is statistically predictable. Based on this, by using the out-of-sample forecasts, they employ trading strategies, which do not generate abnormal profits when transaction costs are taken into account. The trading rule concern delta-hedged strategies (i.e. in order to become delta neutral we buy an amount of call options equal to the reciprocal of the delta[4] value and sell a unit of the underlying index). The non-existence of abnormal profits indicates market efficiency.

3. The data set

In this study we use daily data on the spread of two implied volatility indices, the CBOE volatility futures (settlement prices) and a set of economic variables. The first volatility index is the CBOE volatility index (VIX) which is based on real-time prices of options on the S&P500 index. The second is CBOE DJIA volatility index (VXD) which is based on real-time prices of options on the Dow Jones Industrial Average. Both volatility indices reflect market sentiment of future expected stock market volatility over the next 30-days. The sample period will be January 4, 2000 to 21 November, 2008. The subset January 4, 2000 to January 4, 2006 will be used for the in-sample evaluation and the remaining data will be used for the out-of-sample one.

3.1 Description of implied volatility by B-S model

The implied volatility of an option contract is the volatility implied by the market price of the option based on an option pricing model. An ordinary pricing model is the one of Black-Scholes, which uses a variety of inputs to derive a theoretical value of an option. The Black-Scholes implied volatility is the annualized volatility that equates the Black-Scholes formula value to the options market quote.

The following formula shows the B-S model for a call option:

---

[4] Delta is defined as the rate of change of the option price (c) with respect to the price of the underlying asset (S): $\Delta = \frac{\partial c}{\partial S}$

$$C_t = S_t N(d_1) - Xe^{-r(T-t)} N(d_2) \tag{1}$$

where

$$\begin{aligned} d_1 &= \frac{\ln(S_t / X) + (r + v^2/2)(T-t)}{v\sqrt{T-t}} \\ d_2 &= d_1 - v\sqrt{T-t} \end{aligned} \tag{2}$$

C_t denotes the price of a European call option, S_t is the spot price of the underlying index, X is the strike price of the option, r is the risk-free interest rate, (T-t) is the time to maturity of the option, N is the cumulative normal distribution function and v is the volatility in the returns of the underlying index during the life of the option.

There are several assumptions underlying the Black-Scholes model. The most significant is that volatility, a measure of how much the price of an underlying index can be expected to move in the near-term, is a constant over time. However, volatility does not remain constant; creating problem to the proper valuation of options. The Black-Scholes model also assumes that the logarithmic returns of the underlying index are normally distributed, in comparison to what is observed in financial markets where returns exhibit skewness and kurtosis.

3.2 Description of VIX & VXD

Both CBOE Volatility Index (VIX) and DJIA Volatility Index (VXD) are designed to reflect investors' consensus view of expected volatility over the next 30 days in the S&P 500 Index and Dow Jones Industrial Average, respectively, and as such, can be used as benchmarks of investor sentiment. VIX was introduced in 1993 and was originally designed to measure the market’s expectation of 30-day volatility implied by at-the-money S&P 100 Index (OEX) option prices. Since 2003, VIX is based on the S&P 500 Index (SPX), the core index for U.S. equities, and estimates expected volatility by averaging the weighted prices of SPX puts and calls over a wide range of strike prices. By supplying a script for replicating volatility exposure with a portfolio of SPX options, this new methodology transformed VIX from an abstract concept into a practical standard for trading and hedging volatility. CBOE has listed derivatives on VIX since 2004, and with the trading of these new products there has been increasing demand for other VIX-based calculations.

VIX is a volatility index comprised of options rather than stocks, with the price of each option reflecting the market's expectation of future volatility. The generalized VIX formula[5] has been modified as follows:

$$\sigma^2 = \frac{2}{T} \sum_i \frac{\Delta K_i}{K_i^2} e^{RT} Q(K_i) - \frac{1}{T} \left[\frac{F}{K_0} - 1 \right]^2 \tag{3}$$

In the above formula, σ is the value of the VIX divided by 100, T is time to expiration of the option contract, F is the forward index level derived from index option prices, K_0 is the first strike below the forward level F, K_i is the strike price of ith out-of-the-money option (call if $K_i > F$ and put if $K_i < F$), ΔKi is the interval between strike prices $\left(K_i = \frac{K_{i+1} - K_{i-1}}{2} \right)$, R is the risk-free interest rate to expiration and $Q(Ki)$ is the midpoint of the bid-ask spread for each option with strike Ki. The contribution of a single option to the VIX value is proportional to K and the price of that option, and inversely proportional to the square of the option's strike price.

The components of VIX are near- and next-term put and call options, usually in the first and second SPX contract months. "Near-term" options must have at least one week to expiration; a requirement intended to minimize pricing anomalies that might occur close to expiration. When the near-term options have less than a week to expiration, VIX "rolls" to the second and third SPX contract months.

As volatility rises and falls, the strike price range of options with nonzero bids tends to expand and contract. As a result, the number of options used in the VIX calculation may vary from month-to-month, day-to-day and possibly, even minute-to-minute.

[5] *www.cboe.com/micro/vix/vixwhite.pdf*

For the CBOE DJIA volatility index (VXD), the selection of component options and calculation are identical to the method applied for the VIX index.

VIX can be considered as variance swap (Carr, 2006). The latter is actually a contract where the buyer receives the difference between the realised variance of the returns and pays a fixed variance rate, which is the variance swap rate. So, VIX^2 approximates the 30-day variance swap rate. This concept stands for VXD volatility index, too.

Figure1 shows the daily level of VIX and VXD for the whole sample period. Clear spikes in the value of both indices coincide with the internet bubble burst in 2000, the 9/11/2001 terrorist attack, the corporate scandals in 2002 and of course with the contemporaneous crisis at the mid of 2008. This comment is consistent with the view that VIX index reflects investor's sentiment, fear and risk aversion.

Figure 1: VIX & VXD indices 3/1/2000-21/11/2008

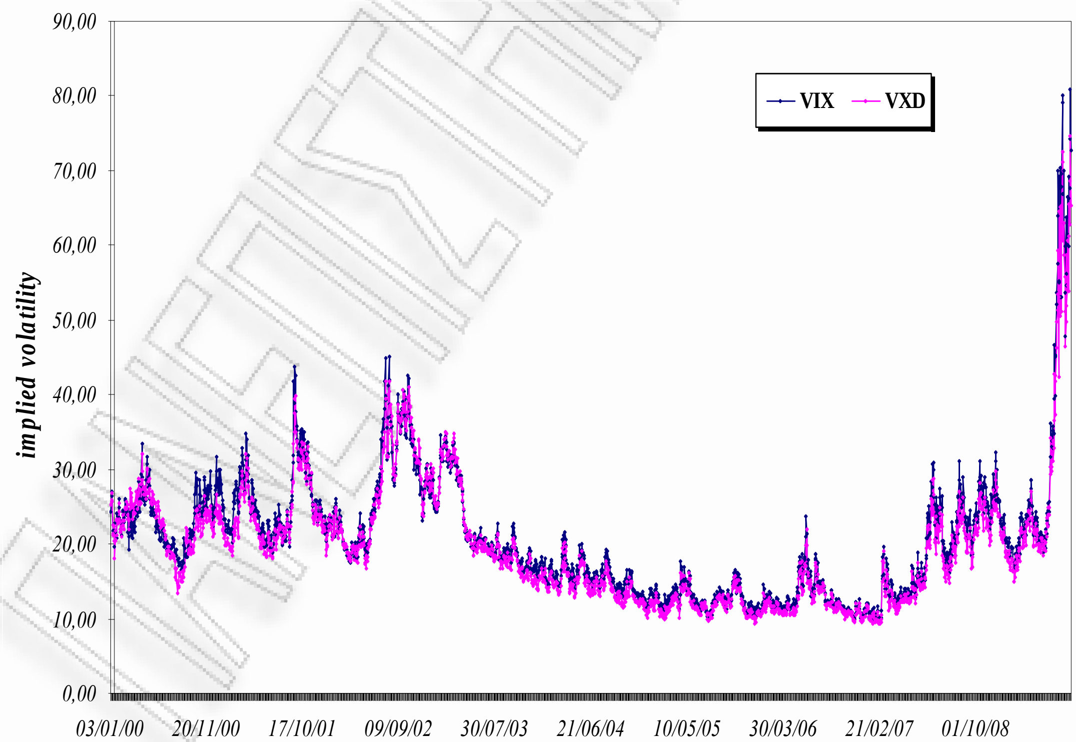

Similar to figure 1, at figure 2, we observe fluctuations of the implied volatility spread, especially the last period and the so called mean reversion. Taking spread first differences, we see points of heteroskedasticity.

Figure 2: Implied volatility spread (VIX-VXD) and first differences

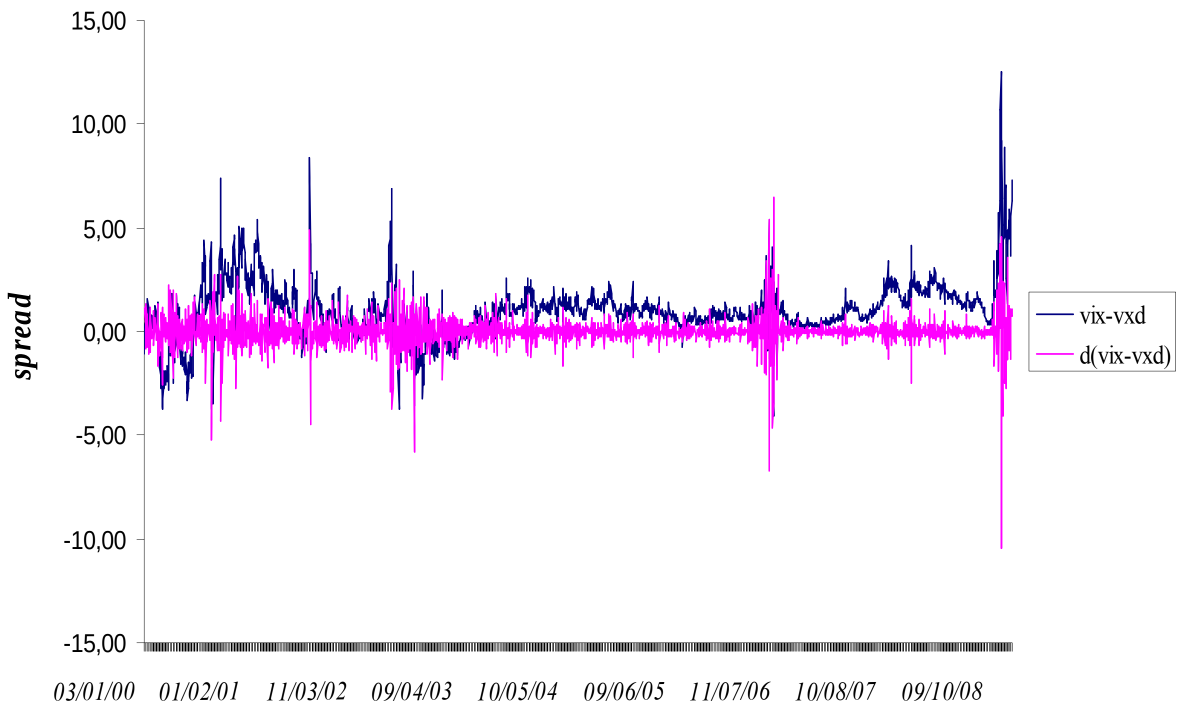

3.3 Other data

The set of the economic variables consists of the returns of the S&P 500 and the DJIA, which are the stock indices of the underlying asset to the options that are used to construct VIX and VXD respectively, (the S&P 500 is a value weighted index published since 1957 of the prices of 500 large-cap common stocks actively traded in the United States, one-month US interbank interest rate for the US market and the USD monthly LIBOR, the Euro/USD exchange rate, the WTI (brent crude oil) price for the American market, the volume of the futures contract of the S&P 500 index (high trading volume may be considered as a signal of e.g. panic selling, linked to rising implied volatility), the slope of the yield curve calculated as the difference between the yield of the ten year government bond and the one-month interbank interest rate. The data of the economic variables were downloaded from Datastream and the time series of implied volatility indices and futures from the CBOE site.

On March 26, 2004, the CBOE launched a new exchange, the Chicago Futures Exchange, and started trading futures on the new volatility index. In less than five years, the combined trading activity in VIX options and futures has grown to more than 100,000 contracts per day. Among the other data, we will use the CBOE VIX and VXD volatility futures which were listed in March 26, 2004 and April 25, 2005, respectively. The contract size of the volatility futures is $1000 times the volatility index. The final settlement day is the Wednesday that is thirty days prior to the third Friday of the calendar month immediately following the month in which the contract expires. In line with, Konstantinidi et al. (2008) we construct three time series of future prices by ranking the data according to their expiry date: the shortest, the second shortest and third shortest series. In case where the shortest contract has less than five days to maturity, we roll to the second shortest series. Moreover, we discard prices that their trading volume is less than five contracts.

Prior to March 26, 2007, the underlying asset of the VIX (VXD) futures contract was an ‘‘Increased-Value index” termed VBI (DVB) that was 10 times the value of VIX (VXD) at any point in time. The contract size of the volatility futures was $100 times the value of the underlying index, so our series have been rescaled accordingly.

Table 1 shows the summary statistics of the implied volatility spread in full sample, in-sample and out-of-sample respectively, in levels (Panel A) and in daily differences (Panel B). It also exhibits volatility futures, VIX and VXD in levels and daily differences (Panel C & D respectively) for the period January 5, 2006 to November 21, 2008. We can see the volatility spread in levels exhibits strong positive autocorrelation, in contrast to first difference which is weak. According to the augmented Dickey-Fuller test for unit roots the implied volatility spread is not stationary in the levels, but stationary in daily differences.

The economic variables, we use in forecasting models, are examined whether they are stationary or not. We can observe that all of them have a unit root in levels, whereas in daily differences, they are stationary.

The implied volatility spread exhibits excess kurtosis both in levels and first differences, meaning that the distribution is leptokurtic, with fatter tails, and extreme values happen with higher probability. The volatility spread is skewed to the right in levels and in daily differences it is left-skewed, the left tail is longer.

As far as the volatility futures are concerned, both VIX and VXD show positive autocorrelation in levels, in contrast to first differences where they display negative autocorrelation. In all the three series, VIX and VXD are skewed to the right. In first differences they exhibit excess kurtosis, whereas in levels the kurtosis is lower for the 3rd shortest series.

Table 1

Summary statistics

VIX-VXD

Panel A: Summary statistics for implied volatility spread (levels)

	January 4, 2000 to November 21, 2008	January 4, 2000 to January 4, 2006	January 5, 2006 to November 21, 2008
Mean	1.02	0.84	1.40
Median	0.96	0.87	1.13
Maximum	12.54	8.38	12.54
Minimum	-4.02	-3.75	-4.02
Std. Dev.	1.34	1.29	1.34
Skewness	1.16	0.21	3.08
Kurtosis	11.35	5.85	20.02
$\rho 1$	0.84	0.86	0.78
Jarque-Bera	6991.47	521.56	9925.43
Probability	0.00	0.00	0.00
Observations	2237	1510	727

Panel B: Summary statistics for implied volatility spread (daily differences)

	January 4, 2000 to November 21, 2008	January 4, 2000 to January 4, 2006	January 5, 2006 to November 21, 2008
Mean	0.00	0.00	0.01
Median	0.00	0.00	-0.01
Maximum	6.45	4.90	6.45
Minimum	-10.48	-5.76	-10.48
Std. Dev.	0.75	0.69	0.87
Skewness	-1.43	-0.81	-2.00
Kurtosis	34.02	14.10	46.68
$\rho 1$	-0.24	-0.24	-0.25
Jarque-Bera	90387.84	7908.93	58281.55
Probability	0.00	0.00	0.00
Observations	*2236*	1509	727

Panel C: Summary statistics for VIX futures: January 5, 2006 to November 21, 2008

	LEVELS			DAILY DIFFERENCES		
	Shortest	2nd shortest	3rd shortest	Shortest	2nd shortest	3rd shortest
Mean	192.69	185.97	188.10	0.00	0.00	0.00
Median	163.20	160.40	167.20	0.00	0.00	0.00
Maximum	637.10	427.00	275.60	0.26	0.14	0.11
Minimum	105.40	118.50	128.10	-0.31	-0.13	-0.12
Std. Dev.	85.59	54.32	46.83	0.05	0.03	0.02
Skewness	2.24	0.65	0.23	0.14	0.72	0.29
Kurtosis	9.87	2.85	1.40	10.09	5.92	5.92
ρ1	0.96	0.96	0.95	-0.11	-0.09	-0.13
Jarque-Bera	2011.50	49.43	69.22	1496.68	301.49	214.18
Probability	0.00	0.00	0.00	0.00	0.00	0.00
Observations	716	689	602			

Panel D: Summary statistics for VXD futures: January 5, 2006 to November 21, 2008

	LEVELS			DAILY DIFFERENCES		
	Shortest	2nd shortest	3rd shortest	Shortest	2nd shortest	3rd shortest
Mean	183.70	186.70	188.00	0.00	0.00	0.00
Median	163.50	186.80	196.90	0.00	0.00	0.00
Maximum	571.20	474.90	327.50	0.30	0.18	0.12
Minimum	103.90	113.20	121.60	-0.30	-0.15	-0.13
Std. Dev.	78.32	58.11	44.29	0.05	0.04	0.03
Skewness	2.22	1.17	0.08	0.26	0.53	0.30
Kurtosis	10.13	5.90	1.84	8.85	5.67	5.33
ρ1	0.80	0.73	0.71	-0.11	-0.09	-0.10
Jarque-Bera	1895.00	320.44	25.53	846.35	154.12	83.46
Probability	0.00	0.00	0.00	0.00	0.00	0.00
Observations	643	550	446			

Table 1 reports the summary statistics of the volatility spread and VIX and VXD futures in levels and first differences, in full sample, in-sample and out-of-sample period. The Jarque-Bera test values and the first order autocorrelation ρ_1 are also presented. The null hypothesis for the Jarque-Bera test is that the series are normally distributed.

4. The forecasting models

4.1 The economic variables model

We use a set of lagged economic variables to forecast the evolution of the implied volatility spread (VIX-VXD) (see Konstantinidi et al. 2008 and Ahoniemi, 2006). We estimate the following regression:

$$\begin{aligned}\Delta(VIX-VXD)_t &= c_1 + a_1 i_{t-1} + b_1 fx_{t-1} + d_1^+ R^+_{t-1(DJIA)} + d_1^- R^-_{t-1(DJIA)} + f_1^+ R^+_{t-1(S\&P500)} + \\ &+ f_1^- R^-_{t-1(S\&P500)} + g_1 \Delta ys_{t-1} + h_1 oil_{t-1} + j_1 \Delta HV_{(DJIA)_{t-1}} + k_1 \Delta HV_{(S\&P500)_{t-1}} + l_1 vol_{t-1} \\ &+ m_1 \Delta(VIX-VXD)_{t-1} + \varepsilon_t\end{aligned} \tag{4}$$

where $\Delta(VIX\text{-}VXD)_t$ denotes the daily changes of the implied volatility spread, c_1 is a constant, i_t denotes the one-month US interbank interest rate for the US market, fx_t the Euro/USD exchange rate, oil_t the WTI (brent crude oil) price for the American market, ΔHV_t the historical volatility of DJIA and S&P 500 respectively and vol_t the volume of the futures contract of the S&P 500 index; all these six variables are measured in log-differences. R^+, R^- denote the underlying stock index positive and negative log-returns (e.g., R^+ is filled with positive returns and zeroes elsewhere). Δys_t denotes the changes of the slope of the yield curve as the difference between the yield of the ten year government bond and the one-month interbank interest rate. The historical volatility is calculated as a 30-days moving average of equally weighted past squared returns (Figlewski, 1994). Positive and negative returns of both underlying indices are separated as negative shocks have the tendency of raising volatility more than positive shocks (Ahoniemi, 2006). Returns on S&P 500 and DJIA are lagged as traders open their position at the start of the day, without knowing how the stock market will develop during the day, so the best available information come from the previous day.

4.2 Univariate autoregressive

Univariate autoregressive model is a tool to explore whether past volatility values can be used for predictive purposes. In this study, we employ an AR(1) model

$$\Delta(VIX - VXD)_t = c_1 + a_1 \Delta(VIX - VXD)_{t-1} \tag{5}$$

and the choice of one lag is based on the BIC criterion (minimum value within a range of ten lags).

5. In-sample evidence

Table 2 and 3 presents the in-sample evaluation of the economic variables model and AR(1) model, respectively. In both tables are reported the estimated coefficients, the t-statistics in parentheses and the adjusted R^2. For the economic variables model, the regressors that are found to have statistically significant coefficients are the returns on brent crude oil, the returns on the historical volatility for both underlying indices, S&P 500 and Dow Jones Industrial Average, and the lagged term of implied volatility spread. The latter is also statistically significant for the AR(1) model (see Table 3). The adjusted R^2 is almost 7% for both models. As in Franks (1991) the volume of the futures contract of the S&P 500 index cannot be considered as a significant explanatory variable. Moreover, despite the past literature, none of the returns on underlying indices (S&P 500 and DJIA) are statistically significant.

Figure 3 and 4 display the residuals of the two models under consideration. Points of heteroskedasticity are observed, and this fact is corroborated by the tests for heteroskedasticity, which show rejection of the null hypothesis of homoskedasticity of the residuals.

Table 2

Forecasting with the economic variables model

	Dependent variable: Δ(VIX-VXD)t
Included observations:	1395
	Coeff.*(t-statistic)*
C_1	0.040
	(-1.404)
I_{t-1}	2.549
	(1.249)
Oil_{t-1}	1.442*
	(2.041)
$R_{(DJIA)t-1}^{+}$	-5.071
	(-0.685)
$R_{(DJIA)t-1}^{-}$	2.680
	(0.344)
$R_{(S\&P500)t-1}^{+}$	4.116
	(0.565)
$R_{(S\&P500)t-1}^{-}$	7.132
	(0.929)
Δys_{t-1}	0.419
	(1.431)
Fx_{t-1}	-2.342
	(-0.817)
$\Delta HV_{(DJIA)t-1}$	-1.329*
	(-2.365)
$\Delta HV_{(S\&P500)t-1}$	1.525*
	(2.675)
Vol_{t-1}	-0.095
	(-1.533)
$\Delta(VIX\text{-}VXD)_{t-1}$	-0.193*
	(-7.111)
Adj.R2	0.069

Table 2 reports the results from the regression of the volatility spread. The estimated regression is $\Delta(VIX\text{-}VXD)_t = c_1 + a_1 i_{t-1} + b_1 fx_{t-1} + d^+_1 R^+_{t-1}(DJIA) + d^-_1 R^-_{t-1}(DJIA) + f^+_1 R^+_{t-1}(S\&P500) + f^-_1 R^-_{t-1}(S\&P500) + g_1 \Delta ys_{t-1} + h_1 oil_{t-1} + j_1 \Delta HV(DJIA)_{t-1} + k_1 \Delta HV(S\&P500)_{t-1} + l_1 vol_{t-1} + m_1 \Delta(VIX\text{-}VXD)_{t-1} + \varepsilon_t$ where $\Delta(VIX\text{-}VXD)$: the changes of the implied volatility spread, i: the one-month US interbank interest rate for the US market, fx: the Euro/USD exchange rate, $R^+(.)$: the positive stock return of DJIA and S&P 500 respectively, $R^-(.)$: the negative stock return of the underlying indices, Δys: the changes of the slope of the yield curve as the difference between the yield of the ten year government bond and the one-month interbank interest rate, oil: the WTI (brent crude oil) price for the American market, $\Delta HV(.)$: the historical volatility of DJIA and S&P 500 and *vol:* the volume of the futures contract of the S&P 500 index. The estimated coefficients, t-statistics in parentheses, and the adjusted R^2 are presented. One asterisk denotes rejection of the null hypothesis of a zero coefficient at 5% level. The model has been estimated for the period January 4, 2000 to January 4, 2006.

Figure 3: Residuals for economic variables model

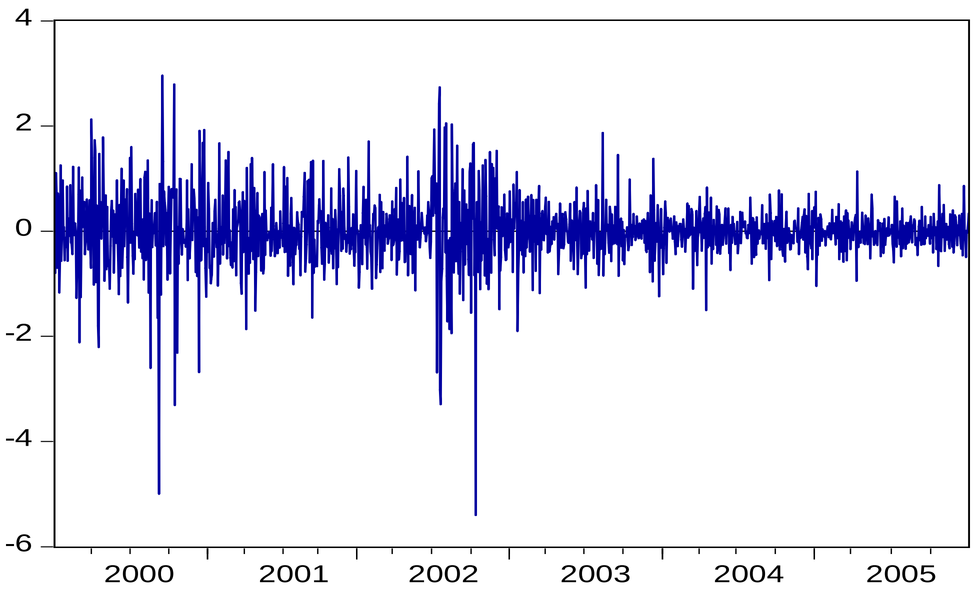

Table 3

Forecasting with the univariate autoregressive	
	AR(1) model
	Dependent variable: Δ(VIX-VXD)t
Included observations:	1453
	Coeff.*(t-statistic)*
C	-0.009
	(-0.663)
$\Delta(VIX\text{-}VXD)_{t-1}$	-0.252^{*}
	(-0.269)
Adj.R^2	0.066

Table 3 reports the results from the univariate autoregressive model. The estimated specification is $\Delta(VIX\text{-}VXD)_t = c + \Delta(VIX\text{-}VXD)_{t-1} + \varepsilon_t$. The estimated coefficients, t-statistics in parentheses, and the adjusted R^2 are presented. One asterisk denotes rejection of the null hypothesis of a zero coefficient at 5% level. The model has been estimated for the period January 4, 2000 to January 4, 2006.

Figure 3: Residuals for AR(1)

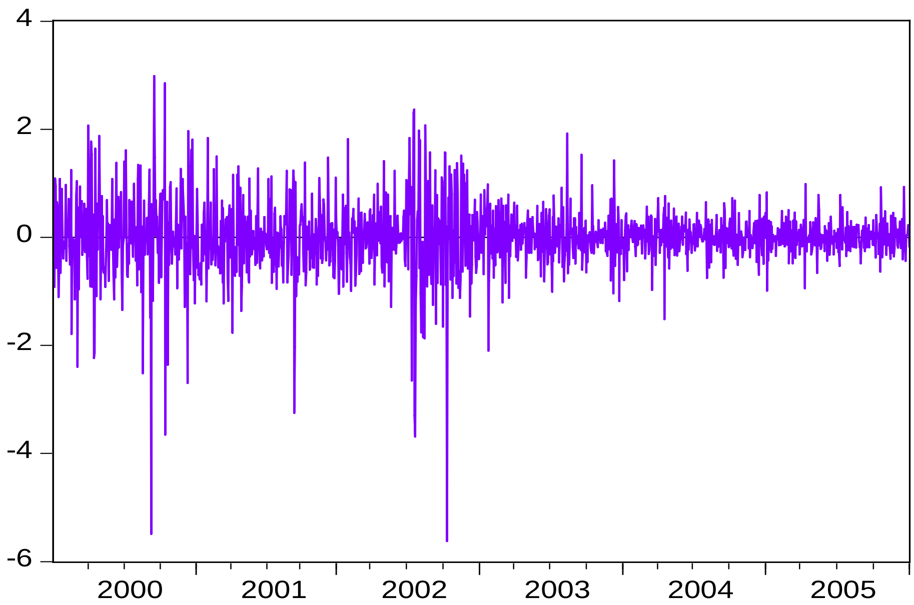

6. Out-of-sample performance

In order to assess the out-of-sample performance of the two models we have already analyzed, we form point and interval forecasts for each model. As mentioned above, the out-of-sample period is January 5, 2006 to November 21, 2008. So, to construct the point forecasts, we use an expanding window, in other words, we re-estimate the models by adding each observation, without dropping any, to the in-sample data. As far as the construction of interval forecasts is concerned, for every day (out-of-sample period), we generate 10.000 simulation runs.

6.1 Point forecasts

In this study, in line with Konstantinidi et al.(2008) and Goncalves and Guidolin (2006) we use three measures to assess the out-of-sample performance of the fitted models. The root mean squared prediction error (RMSE), the mean absolute prediction error (MAE) and the mean correct prediction (MCP). The RMSE is calculated as the square root of the average squared deviations of actual prices (spread) from the models' forecast prices. The MAE is computed as the absolute differences between the realized spread of the two IV indices and the models' forecasts spread. The MCP is the average frequency for which the change in the forecasted spread by the models has the same sign as the actual change in spread.

To this point, we have to say that random walk model is going to be used as the benchmark model. The forecasts for the random walk model are based on that today's implied volatility is the best forecast of tomorrow's implied volatility. So, the forecasts of each model will be compared to those from the random walk model. To perform these comparisons, we use the modified Diebold-Mariano test (see Harvey et al., 1997) for the RMSE/MAE metrics and a ratio test for the MCP metric. The null hypothesis for both tests is that the models under consideration perform equally the random walk model. These test are used because sometimes one set of forecasts appears more successful than the other, however, we do not know if this outcome is random.

For the modified Diebold-Mariano test, the null hypothesis is $H_o : E\left[g\left(e^i{}_t\right) - g\left(e^{rw}{}_t\right)\right] = 0$, where $g\left(e_t^i\right)$ and $g\left(e_t^{rw}\right)$ are the loss functions and $\left\{e_t^i\right\}_{t=1}^n$ and $\left\{e_t^{rw}\right\}_{t=1}^n$ are the forecasts error for the ith model specification and the random walk model, respectively. We also define, the loss differential $d_t^i = g\left(e_t^i\right) - g\left(e_t^{rw}\right), t = 1, \ldots, n$. The alternative hypothesis we test is that the random walk outperforms each of the models $\left(H_1 : E\left(d_t^i\right) > 0\right)$ or that each of the models we examine outperform the random walk $\left(H_1 : E\left(d_t^i\right) < 0\right)$. Here, we suppose one-step ahead forecasts, so the modified Diebold-Mariano (S_1) test statistic is the following:

$$S_1 = \frac{\overline{d}_t^i}{\sqrt{\mathrm{var}\left(\overline{d}_t^i\right)}} \tag{6}$$

where $\overline{d}_t^i = \frac{\sum_{t=1}^{n} d_t^i}{n}$ and $\mathrm{var}\left(\overline{d}_t^i\right)$ the variance of $\overline{d}_t^i$. The modification of the Diebold-Mariano test statistic is compared to critical values from the Student's t distribution with (n-1) degrees of freedom. At 5% significance level, we reject the null hypothesis if $S_1 > S_T$ (=1,645) (where S_T is Student's t t-statistic) and then we accept that the models under consideration perform better than the random walk.

As we have already mentioned, we use a ratio test to compare the MCP metric of the respective models and we consider random walk model as a naïve model with MCP=50%. So, the null hypothesis is that the models under consideration perform equally well with the naïve rule ($H_o : MCP = 50\%$). The alternative one is that the models perform better than the naïve rule ($H_1 : MCP > 50\%$). The test statistic is:

$$T_R = \frac{\frac{X}{n} - 0.5}{\sqrt{\frac{0.5 \times 0.5}{n}}} \tag{7}$$

where X is the number of times that the forecasting model predicts correctly the sign of the price change, n is the number of the observations (the number of forecasts) and $\frac{X}{n}$ is the MCP. The test statistic follows approximately a normal distribution $N(0,1)$, so at 5% significance level, we reject H_o in the cases where $T_R > T_N$ (=1,645) (where T_N is t-statistic of standardized normal distribution).

Table 4 presents the results of the out-of-sample performance of the random walk model, the economic variables model and AR(1) model for the implied volatility spread. None of the combinations of implied volatility spread and predictability metrics show that any of the models outperform the random walk model. So, at 5% significance level, we cannot observe any statistically predictable pattern in the implied volatility spread.

Table 4

Out-of-sample performance of the model specifications

	VIX-VXD
Panel A: Random Walk	
RMSE	1,37
MAE	0,62
Panel B: Economic variables model	
RMSE	0,81
MAE	0,37
MCP	54,67%
Panel C: AR(1) Model	
RMSE	0,839
MAE	0,366
MCP	61,95%

Table 4 reports the root mean squared prediction error (RMSE), the mean absolute prediction error (MAE) and the mean correct prediction (MCP) for the random walk model (Panel A), the economic variables model(Panel B) and AR(1) model (Panel C). The null hypothesis is that the random walk and the model under consideration perform equally well, against the alternative that the model under consideration performs better, have been tested via the Modified Diebold-Mariano test (for RMSE and MAE) and the ratio test (for MCP). The models have been estimated for the period January 5, 2006 to November 21, 2008.

6.2 Interval forecasts

Apart from point forecasts, we form interval forecast, as recently financial market participants have shown increasing interest in interval forecasts as measures of uncertainty. Christoffersen's (1998) likelihood ratio test of unconditional coverage is used to assess the constructed interval forecasts. The procedure is the following: observe a sample path, $\{y_t\}_{t=1}^{T}$, of the time series of the implied volatility spread and a corresponding sequence of out-of-sample constructed interval forecasts, $\{(L_{t/t-1}(a), U_{t/t-1}(a))\}_{t=1}^{T}$, at a% significance level. $L_{t/t-1}(a)$ and $U_{t/t-1}(a)$ are the lower and upper limits of the interval forecasts for time t made at time t-1.

We set an indicator variable, I_t, for a given interval forecast, $(L_{t/t-1}(a), U_{t/t-1}(a))$ as,

$$I_t = \begin{cases} 1, & y_t \in [(L_{t/t-1}(a), U_{t/t-1}(a))] \\ 0, & y_t \notin [(L_{t/t-1}(a), U_{t/t-1}(a))] \end{cases} \tag{8}$$

To test the unconditional coverage, the hypothesis that $E[I_t] = a$ should be tested against the alternative $E[I_t] \neq a$. The likelihood under the null hypothesis is

$$L(a; I_1, I_2, ..., I_T) = (1-a)^{n_o} a^{n_1}, \tag{9}$$

and under the alternative

$$L(\pi; I_1, I_2, ..., I_T) = (1-\pi)^{n_o} \pi^{n_1}. \tag{10}$$

The likelihood ratio unconditional test (LR_{unc}) can be formulated as:

$$LR_{unc} = -2\ln[L(a; I_1, I_2, ..., I_T) / L(\hat{\pi}; I_1, I_2, ..., I_T)] \approx \chi_1^2$$
$$LR_{unc} = -2\ln[(1-a)^{n_o} a^{n_1} / (1-\hat{\pi})^{n_o} \hat{\pi}^{n_1}] \approx \chi_1^2$$

Where $\hat{\pi} = n_1/(n_o + n_1)$ is the maximum likelihood estimate of π and χ_1^2 is the distribution with one degree of freedom. Furthermore, n_1 and n_o are the numbers of 1s and 0s in the sample, respectively, so $\hat{\pi}$ is the observed fraction of violations in the sequence, which we report in table 5.

In table 5 we also provide the p-value associated with the test statistic which is calculated as

$$p - value \equiv 1 - F_{x_1^2}(LR_{unc}) \tag{11}$$

If the P-value is below the desired significance level, then we reject the null hypothesis. As we observe, p-value is less than 5% significance level, so the null hypothesis that the percentage of times that the actually realized index value falls outside the constructed 5%-intervals is 5% is rejected.

According to table 5, as in the case of point forecasts, none of the models yield accurate forecasts. So, there cannot be observed any predictable pattern in the implied volatility spread based on interval forecasts formed by the economic variables model and the AR(1) model.

Table 5

Statistical accuracy of the interval forecasts	
VIX-VXD	
Panel A: Regression model based on economic variables - interval forecasts	
#Violations	5.35%
p-value(LRunc)	0.00*
Panel B: AR(1) Model- interval forecasts	
#Violations	3.86%
p-value(LRunc)	0.00*

Table 5 presents the percentage of observations that fall outside the constructed intervals and the p-value of Christoffersen's likelihood ratio test of unconditional coverage for the implied volatility spread. The null hypothesis is that the percentage of times that the actually realized index value falls outside the constructed $a\%$-intervals is $a\%$.The results are reported for daily 5% - interval forecasts over the period January 5, 2006 to November 11, 2008.

7. Economic significance

In this section, we evaluate the economic significance of the point and interval forecasts formed already for the implied volatility spread and in contrast to other studies, we construct trading strategies based on intercommodity spread (same time-to-maturity, different underlying). In other words, we take simultaneously position in VIX and VXD futures for each maturity series. To employ the strategy, three futures series have been constructed according to their expiry date. Even if the two models under consideration do not generate statistically significant forecasts, we apply the trading rule to both of them separately.

To assess the economic significance of the point and interval forecasts, we calculate two measures: Sharpe ratio (SR) and Leland's (1999) alpha (A_p). Sharpe ratio shows the excess return per unit of risk in a trading strategy. Sharpe takes into account only the first two moments of the distribution, supposing that the mean and standard deviation of the distribution of returns are sufficient statistics for the assessment of the prospects of an investment portfolio. The SR is defining as:

$$SR \equiv \frac{\overline{D}}{\sigma_D} \tag{12}$$

Where $D_t = R_{It} - R_{ft}$ is the differential return in period t, R_{It} is the return on the index in period t and R_{ft} the return on the risk-free interest rate. $\overline{D} \equiv \frac{1}{T}\sum_{t=1}^{T} D_t$ is the average value of D_t over the period from t=1 through T and $\sigma_D \equiv \sqrt{\frac{\sum_{t=1}^{T}(D_t - \overline{D})^2}{T-1}}$ is the standard deviation of the differential return.

Instead of using Jensen's alpha, A_p is preferred since it takes into account the non-normality of the distribution of the returns. Leland's alpha is given by the following equation:

$$A_p = E(r_p / M) - E(r_p) = E(r_p / M) - B_p\left[E(r_m) - r_f\right] - r_f \tag{13}$$

Where $B_p = \frac{\operatorname{cov}(r_p, -(1+r_m)^{-b})}{\operatorname{cov}(r_m, -(1+r_m)^{-b})}$ and $b = \frac{\ln[E(1+r_m)] - \ln(1+r_f)}{\operatorname{var}[\ln(1+r_m)]}$. In order to obtain A_p, we regress the following specification:

$$\left[r_{pt} - r_{ft}\right] - B_p\left[r_{mt} - r_{ft}\right] = \widehat{A}_p + \varepsilon_t \tag{14}$$

r_p are the returns from the trading strategy, r_f are the returns on the risk-free interest rate and r_m are the returns of the market portfolio. As risk-free interest rate we use one-month Libor interest rate and as the market returns, we consider returns on S&P 500 and DJIA, separately. To assess the economic significance of Sharpe ratio and Leland's alpha, for point and interval forecasts, 95% confidence interval have been bootstrapped and are reported within parentheses.

7.1 Trading strategy based on point forecasts

To evaluate the economic significance of the point forecasts, we employ the following strategy. The investor goes long in VIX and short in VXD volatility futures when the forecasted value of the implied volatility spread is greater than its current value. Correspondingly, the investor goes short in VIX and long in VXD volatility futures in the case where the forecasted value of the implied volatility spread is smaller than its current value. Table 6 and 8 present the annualized Sharpe ratio and Leland's (1999) alpha obtained for each of the three series (shortest, 2nd shortest, 3rd shortest) for point forecasts formed by economic variables model (Panel A) and the AR(1) model (Panel B). At first, we set the trading rule taking into account the transaction costs (see Table 8), however Sharpe ratio and Leland's alpha are statistically insignificant for both models and these models cannot generate abnormal profits. Subsequently, under a theoretical viewpoint, we employ the same trading strategy, ignoring the transaction costs, in order to see whether the models under consideration, although they offer no statistically significant forecasts, can generate profits. The results coincide with the first case (table 6), as we accept the null hypothesis of a zero Sharpe ratio and Leland's alpha, at 5% level of significance. The results for Leland's alpha remain the same, whether we use S&P 500 or DJIA as market returns.

7.2 Trading strategy based on interval forecasts

To assess the economic significance of the constructed interval forecasts, the following trading strategy is employed:

If $IVS_{t-1} < \frac{U_{t/t-1}(a) + L_{t/t-1}(a)}{2}$, then go long in VIX and short in VXD

If $IVS_{t-1} > \frac{U_{t/t-1}(a) + L_{t/t-1}(a)}{2}$, then go short in VIX and long in VXD

If $IVS_{t-1} = \frac{U_{t/t-1}(a) + L_{t/t-1}(a)}{2}$, then do nothing

Table 7 and 9 report the annualized Sharpe ratio and Leland's (1999) alpha obtained for each of the three series (shortest, 2^{nd} shortest, 3^{rd} shortest) for interval forecasts formed by economic variables model (Panel A) and the AR(1) model (Panel B). At first, as in the case of point forecasts, we employ the trading strategy by taking into account the transaction costs (see Table 9), however Sharpe ratio and Leland's alpha are statistically insignificant for both models and these models cannot generate abnormal profits. Subsequently, we employ the same trading strategy, ignoring the transaction costs, in order to see whether the models under consideration, although they offer no statistically significant forecasts, can generate profits. The results (Table 7) are at the same direction with the first case, as the models do not generate economic significant profits. The results for Leland's alpha remain the same, whether we use S&P 500 or DJIA as market returns.

Table 6

*Trading strategies with VIX-VXD spread based on point forecasts from January, 5, 2006 to November 21, 2008 (**without** transaction costs)*

	shortest	2nd shortest	3rd shortest
Panel A: Economic variables model point forecasts			
Sharpe Ratio	0.01	0.01	-0.02
95% CI	(-0.02, 0.03)	(-0.02, 0.05)	(-0.07, 0.03)
$A_{p(S\&P\ 500)}$	0.03	0.03	-0.06
95% CI	(-7.37, 6.32)	(-7.08, 7.57)	(-8.05, 7.86)
$A_{p(DJIA)}$	0.03	0.03	-0.06
95% CI	(-7.43, 5.80)	(-8.71, 7.69)	(-10.34, 7.32)
Panel B: AR(1) point forecasts			
Sharpe Ratio	0.01	0.00	-0.00
95% CI	(-0.02, 0.04)	(-0.04, 0.05)	(-0.03, 0.04)
$A_{p(S\&P\ 500)}$	0.03	0.00	-0.01
95% CI	(-5.34, 6.02)	(-6.67, 8.06)	(-10.23, 8.23)
$A_{p(DJIA)}$	0.03	0.00	-0.01
95% CI	(-5.47, 6.15)	(-6.82, 8.24)	(-10.47, 8.42)

Table 6 reports the annualized Sharpe ratio and Leland's alpha and their respective bootstrapped 95% confidence intervals (CI). The trading strategy is based on point forecasts obtained from the economic variables model (Panel A) and AR(1) model (PanelB), without taking into account transaction costs. The Sharpe ratio for the S&P 500 and the Dow Jones Industrial Average is 0.0536 [95% CI=(0.05, 0.15)] and 0.0693 [95% CI=(0.06, 0.16)].

Table 7

*Trading strategies with VIX-VXD spread based on interval forecasts from January, 5, 2006 to November 21, 2008 (**without** transaction costs)*

	VIX-VXD		
	shortest	2nd shortest	3rd shortest
Panel A: Economic variables model point forecasts			
Sharpe Ratio	-0.03	-0.00	-0.07
95% CI	(-0.08, 0.03)	(-0.07, 0.06)	(-0.11, -0.02)
$A_{p(S\&P\ 500)}$	-0.11	-0.00	-0.18
95% CI	(-5.58, 6.47)	(-7.62, 7.36)	(-9.22, 7.72)
$A_{p(DJIA)}$	-0.11	-0.00	-0.18
95% CI	(-6.11, 7.08)	(-8.33, 8.06)	(-10.10, 8.46)
Panel B: AR(1) point forecasts			
Sharpe Ratio	-0.00	-0.02	-0.01
95% CI	(-0.08, 0.08)	(-0.09, 0.04)	(-0.07, 0.04)
$A_{p(S\&P\ 500)}$	-0.02	-0.04	-0.04
95% CI	(-7.04, 6.55)	(-8.02, 7.40)	(-10.26, 7.76)
$A_{p(DJIA)}$	-0.02	-0.04	-0.04
95% CI	(-7.70, 7.17)	(-8.78, 8.11)	(-11.24, 8.49)

Table 7 reports the annualized Sharpe ratio and Leland's alpha and their respective bootstrapped 95% confidence intervals (CI). The trading strategy is based on interval forecasts obtained from the economic variables model (Panel A) and AR(1) model (PanelB), without taking into account transaction costs. The Sharpe ratio for the S&P 500 and the Dow Jones Industrial Average is 0.0536 [95% CI= (0.05, 0.15)] and 0.0693 [95% CI= (0.06, 0.16)].

Table 8

*Trading strategies with VIX-VXD spread based on point forecasts from January, 5, 2006 to November 21, 2008 (**with** transaction costs)*			
	shortest	2nd shortest	3rd shortest
Panel A: Economic variables model point forecasts			
Sharpe Ratio	-0.01	-0.01	-0.05
95% CI	(-0.05, 0.00)	(-0.07, 0.01)	(-0.11, 0.01)
$A_{p(S\&P\ 500)}$	-0.03	-0.03	-0.14
95% CI	(-8.45, 7.26)	(-9.39, 8.07)	(-8.64, 8.98)
$A_{p(DJIA)}$	-0.03	-0.03	-0.14
95% CI	(-9.25, 7.95)	(-10.26, 8.84)	(-9.45, 9.83)
Panel B: AR(1) point forecasts			
Sharpe Ratio	0.00	-0.02	-0.02
95% CI	(-0.04, 0.03)	(-0.07, 0.02)	(-0.06, 0.02)
A_p	0.00	-0.05	-0.07
95% CI	(-6.78, 7.79)	(-7.54, 8.99)	(-8.63, 8.49)
$A_{p(DJIA)}$	0.00	-0.05	-0.07
95% CI	(-7.38, 7.78)	(-9.06, 9.15)	(-11.64, 12.11)

Table 8 reports the annualized Sharpe ratio and Leland's alpha and their respective bootstrapped 95% confidence intervals (CI). The trading strategy is based on point forecasts obtained from the economic variables model (Panel A) and AR(1) model (PanelB), by taking into account transaction costs. The Sharpe ratio for the S&P 500 and the Dow Jones Industrial Average is 0.0536 [95% CI=(0.05, 0.15)] and 0.0693 [95% CI=(0.06, 0.16)].

Table 9

*Trading strategies with VIX-VXD spread based on interval forecasts from January, 5, 2006 to November 21, 2008 (**with** transaction costs)*			
	shortest	2nd shortest	3rd shortest
Panel A: Economic variables model point forecasts			
Sharpe Ratio	-0.06	-0.05	-0.11
95% CI	(-0.12, 0.00)	(-0.12, -0.01)	(-0.16, -0.07)
$A_{p(S\&P\ 500)}$	-0.21	-0.12	-0.31
95% CI	(-6.35, 6.18)	(-8.70, 7.64)	(-10.14, 9.22)
$A_{p(DJIA)}$	-0.21	-0.12	-0.31
95% CI	(-6.95, 6.77)	(-9.51, 8.38)	(-11.09, 10.10)
Panel B: AR(1) point forecasts			
Sharpe Ratio	-0.02	-0.04	-0.05
95% CI	(-0.08, 0.05)	(-0.12, 0.02)	(-0.10, 0.00)
A_p	-0.07	-0.10	-0.13
95% CI	(-7.62, 7.33)	(-8.89, 8.33)	(-11.82, 10.50)
$A_{p(DJIA)}$	-0.07	-0.10	-0.13
95% CI	(-8.15, 7.86)	(-9.72, 9.14)	(-12.93, 11.52)

Table 9 reports the annualized Sharpe ratio and Leland's alpha and their respective bootstrapped 95% confidence intervals (CI). The trading strategy is based on interval forecasts obtained from the economic variables model (Panel A) and AR(1) model (Panel B), by taking into account transaction costs The Sharpe ratio for the S&P 500 and the Dow Jones Industrial Average is 0.0536 [95% CI=(0.05, 0.15)] and 0.0693 [95% CI=(0.06, 0.16)]

8. Conclusion

In this paper, we have examined the dynamics of the implied volatility spread. We have investigated the predictability of two models, an economic variables model and an AR(1) model, by forming point and interval forecasts. The out-of-sample performance of the models has been compared with the random walk model. Based on forecasts, we have developed trading strategies which concern the simultaneous taking position on two different volatility futures, VIX and VXD, with the same time to maturity.

None of the models display statistically significant forecasts, as the naïve rule of random walk model outperforms in all out-of-sample performance metrics. Moreover, there seems to be no predictable pattern in the implied volatility spread, according to the two models under consideration.

Furthermore, as far as the economic significance is concerned, point and interval forecasts, for both models, do not generate abnormal profits, even if transaction costs are not taken into account. This signifies that CBOE volatility futures market are informational efficient, as prices in this markets embody any new information and investors have no margins to make abnormal profits.

References

Ahoniemi, K., 2006. Modeling and forecasting implied volatility: An econometric analysis of the VIX index. Working paper, Helsinki School of Economics.

Aboura, S., 2003. International transmission of volatility: A study on the volatility indexes VX1, VDAX, VIX. Working paper, ESSEC Business School.

Carr, P., Wu, L., 2006. A tale of two indices. Journal of Derivatives 13, 13–29.

Christoffersen, P.F., 1998. Evaluating interval forecasts. International Economic Review 39, 841–862.

Fernandes, M., Medeiros, M.C., Scharth, M., 2007. Modeling and predicting the CBOE market volatility index. Working paper, Queen Mary University.

Figlewski, S., 1994. Forecasting volatility using historical data. Working paper, New York University.

Franks, J.R., Schwartz, E.S, 1991. The stochastic behaviour of market variance implied in the prices of index options. The Economic Journal, Vol. 101, No 409, 1460-1475.

Goncalves, S., Guidolin, M., 2006. Predictable dynamics in the S&P 500 index options implied volatility surface. Journal of Business 79, 1591–1635.

Harvey, C.R., Whaley, R.E., 1992. Market volatility prediction and the efficiency of the S&P 100 index option market. Journal of Financial Economics 31, 43–73.

Harvey, D.I., Leybourne, S.J., Newbold, P., 1997. Testing the equality of prediction mean squared errors. International Journal of Forecasting 13, 281–291.

Konstantinidi, E., Skiadopoulos, G., Tzagkaraki, E., 2008. Can the evoloution of implied volatility be forecasted? Evidence from European and US implied volatility indices. Journal of Banking & Finance 32, 2401-2411

Konstantinidi, E., Skiadopoulos, G., 2009. Are VIX future prices predictable? An empirical investigation.

Poon, S.H., Pope, P.F., 2000. Trading volatility spreads: A test of index option market efficiency. European Financial Management Journal 6, 235–260.

Skiadopoulos, G., Hodges, S.D., Clewlow, L., 1999. The dynamics of the S&P 500 implied volatility surface. Review of Derivatives Research 3, 263–282.