A Simplified Pricing Model for Volatility Futures

BRICE DUPOYET
ROBERT T. DAIGLER*
ZHIYAO CHEN

We develop a general model to price VIX futures contracts. The model is adapted to test both the constant elasticity of variance (CEV) and the Cox–Ingersoll–Ross formulations, with and without jumps. Empirical tests on VIX futures prices provide out-of-sample estimates within 2% of the actual futures price for almost all futures maturities. We show that although jumps are present in the data, the models with jumps do not typically outperform the others; in particular, we demonstrate the important benefits of the CEV feature in pricing futures contracts. We conclude by examining errors in the model relative to the VIX characteristics. © 2010 Wiley Periodicals, Inc. Jrl Fut Mark 31:307–339, 2011

*Correspondence author, Chapman Graduate School of Business, Department of Finance RB206, Miami, Florida 33199. Tel: 305-348-3325, Fax: 305-348-4245, e-mail: daiglerr@fiu.edu

- *Brice Dupoyet is an Associate Professor of Finance, Florida International University, Miami, Florida.*

- *Robert T. Daigler is Knight Ridder Research Professor of Finance, Florida International University, Miami, Florida.*

- *Zhiyao Chen is a Doctoral Candidate in Finance, The University of Washington, Seattle, Washington and Florida International University, Florida.*

INTRODUCTION

Volatility and Financial Markets

Volatility remains an important, but opaque, topic in financial economics after four decades of study. Researchers have attempted to forecast volatility by using more sophisticated quantitative models and by employing GARCH and implied volatility models; however, the explanatory power of these forecasts is relatively low. Poon and Granger (2003) review the literature on forecasting volatility and related issues, concluding that much needs to be done to understand the process and characteristics of volatility.

The introduction of exchange-traded volatility futures and over-the-counter volatility swaps allows for speculation and hedging of future volatility. These products have created increased interest in forecasting the direction and size of changes in volatility. Previous models to price volatility futures encountered several problems in their quest to develop a precise and accurate method to value these contracts, not the least of which is the complexity of the time series behavior of volatility (such as successfully incorporating jumps into the process). Here we propose and test a simple yet flexible constant elasticity of variance (CEV) model with and without jumps to price volatility futures, and to compare the model against the often used Cox–Ingersoll–Ross formulation.

Volatility Futures

Implied volatility has long been a popular measure of the "fear" and "exuberance" in the stock market (see Hibbert, Daigler, & Dupoyet, 2008; Low, 2004; Whaley, 2000), as well as a measure of the market's forecast of future volatility. The current spot VIX (as opposed to the futures VIX price) measures expected future volatility by determining the constant 30-day implied volatility of priced out-of-the-money S&P 500 option strikes as a weighted average of the nearby and first deferred option expirations.[1]

Volatility futures are a relatively new instrument that allows users to trade volatility directly rather than indirectly attempting to manage volatility through option sensitivities. The volume of the VIX futures fluctuates based on the volatility in the market and the time to futures expiration, with changes in the bid–ask spread occurring frequently throughout the day. VIX futures prices are affected by the time to expiration of the futures, the mean reversion characteristics of the expected implied volatility, and potentially the market's forecast of possible jumps in the VIX.

[1]When the S&P 500 options approach expiration then the nearby maturity is dropped from the calculation when five days are left until expiration; at that time and the first and second deferred maturities are employed in the calculation.

Testing the Pricing of Volatility Futures

Futures pricing models for stock index futures (and for similar contracts such as currency futures) are arbitrage cost-of-carry models that relate the underlying asset price to the futures price via holding the asset. VIX futures possess different characteristics as it is a non-traded asset, since holding the "underlying asset" for the VIX futures (a portfolio of S&P 500 options for almost all priced out-of-the-money calls and puts) is infeasible due to cost and liquidity issues. Moreover, the characteristics of the spot VIX include mean-reversion and jumps in the level, as well as changing volatility levels, making it difficult to keep a zero-risk hedged portfolio. Several volatility futures models have emerged in the literature; however, these models generally do not include all the above factors and therefore typically are too basic to price VIX futures adequately.

Our approach differs from previous research that theoretically models the VIX in that it includes both the CEV and Poisson jumps (previous models are discussed shortly). Therefore, compared to previous attempts to model the VIX for the purpose of pricing derivatives, our model includes more of the components that could drive the implied volatility of the S&P 500, while still managing a relatively simple structure. The model involves a two-step process. First, we posit various dynamics of the VIX and then estimate the parameters that determine the evolution of this VIX process over time. Second, we derive the VIX futures values based on our various VIX model specifications, and then test the models' pricing abilities using actual market values.

This study provides comprehensive empirical testing of fair VIX futures values against actual VIX futures prices.[2] The out-of-sample empirical results from our CEV model track the VIX futures well, especially for the 60 days before expiration, such that the average Mean Signed Percentage Error (MSPE) is typically below 1.2% (13 basis points). The Mean Absolute Percentage Error (MAPE) for these time intervals is always less than 4.5% (67 basis points). Importantly, the model outperforms other alternatives, including the Grunbichler and Longstaff (1996) and Zhang and Zhu (2006) formulations, with these alternatives based on the Cox–Ingersoll–Ross (1985) approach, a nested model of our general formulation. The size of the average errors increases slightly for a longer time to futures maturity. However, the empirical results also show that models with jumps are *not* optimal for the best out-of-sample results. In the final section of the study we examine potential reasons for the difference between actual and model futures prices.

[2]Zhang and Zhu (2006) test their model against futures prices, but they only examine the closing price from *one* day. Dotsis et al. (2007) perform empirical comparisons of models but these are limited to a geometric Brownian motion and a square root process. Lin's (2007) empirical results are not the focus of her study and are limited in scope, as discussed later.

CHARACTERISTICS OF THE VIX AND VIX FUTURES

Calculation of the Spot VIX

In October 2003, the calculation of the spot VIX changed in order to make it compatible for derivative instruments to be based on the VIX, to reflect the institutional interest in the S&P 500 index instead of the previously used S&P 100 index, and to include the effect of the smile (skew) that occurs with S&P 500 index option prices. The new VIX uses a new formula that is not adversely affected by the assumptions of the Black–Scholes option pricing model, especially the resultant implication of an equal volatility across all strike prices. In addition, the options on the S&P 500 are European in nature, whereas the previously used OEX 100 options are American options.[3]

The spot VIX represents a constant 30-day implied volatility based on the S&P 500 index options and is calculated as follows:

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

where T is the time until the option's expiration expressed in years, F is the forward index level derived using the S&P 500 index option prices, K_i is the strike price of the ith out-of-the-money option relative to F, ΔK_i equals the interval between strike prices and is defined as $\Delta K_i = (K_{i+1} - K_i)/2$, K_o is the first strike below the forward index level F, R is the annualized risk-free interest rate to expiration, and $Q(K_i)$ equals the midpoint of the bid–ask spread for the option with strike price K_i. The VIX keeps a constant 30-day maturity by using a weighted average of the nearby and first deferred option maturities.[4] The formula is a discretization of the mathematical derivation of return variance, as shown by Carr and Wu (2006).[5]

This calculation employs the S&P 500 call and put index options that are both out-of the money relative to the forward index F and have a non-zero bid price.[6] Therefore, essentially the entire tradable range of the implied volatility

[3]The previous VIX value (now referred to as the VXO) is calculated using the Black–Scholes implied volatilities on only the four call and four put S&P 100 (OEX) index options closest to the at-the-money strike. While both measures are provided to the marketplace, the new VIX has become the standard for measuring implied volatility in the market.

[4]When the S&P 500 options approach expiration, the nearby maturity is dropped from the calculation when five days are left until expiration; at that time and the first and second deferred maturities are employed in the calculation.

[5]Biases in the discrete VIX formulation are discussed later.

[6]Technically one chooses all out-of-the-money calls and puts with non-zero bid prices as one moves away from the forward index value until one encounters two consecutive options with a zero bid-price. No other away-from-the-money options are employed in the calculation after this point.

smile (smirk) is included in the VIX calculation. Since the calculation is based on bid and ask prices, the implied volatilities do not change as option transactions prices change from the bid to the ask price. Theoretically, this new calculation procedure for the VIX yields an asset that could be hedged using a weighted combination of options.[7]

VIX Futures and the Data

The VIX futures started trading on the Chicago Futures Exchange in March 2004. Here we examine contracts that expire from May 2004 (the first expiration) through September 2006. In order to estimate our VIX model, we employ the spot VIX index closings for the two years preceding the day where the VIX futures theoretical price needs to be calculated and then we compare the model prices against the observed VIX futures market price. Thus, each day the parameters for the VIX model are updated using the most recent two years of historical data in order to determine the long-term mean-reverting value for the VIX. We take advantage of an interesting result found in Henderson, Hobson, Howison, and Kluge (2004), Daouk and Guo (2004), Psychoyios and Skiadopoulos (2006), and Dotsis, Psychoyious, and Skiadopoulous (2007) who show that model rankings do not depend on risk premium values as model prices are a monotonic function of the price of risk of volatility; for model's comparison purposes we therefore assume the price of risk to be zero. Next, the most recent spot market data is used to obtain the VIX futures estimates for the four models developed here. Finally, we compare the actual VIX futures prices to the estimates of each of the four models to determine which model is most compatible with actual VIX daily futures prices. As with all such models, any errors can be the result of an imperfect model or incorrectly valued actual futures prices.

Figure 1 illustrates the historical behavior of the VIX during the time period of this study, with the VIX ranging from 10.23 to 23.81%. Several spikes in volatility are apparent, as is the mean-reverting nature of the VIX after both spikes in volatility and longer-term trends in the VIX.

---

[7]Note that the spot VIX squared can be replicated with a static combination of S&P500 index options, but that VIX futures require a dynamic hedge. Moreover, the formula for the spot VIX is model-free, but the one for VIX futures is model dependent. Therefore, the volatility model employed for the VIX futures is critical. The dynamic hedge involves a strip of options. As the relative value of the S&P 500 options change, the weights of the options in the dynamic portfolio change, complicating the objective of reducing the basis risk by hedging. The futures settle to the 30-day S&P 500 option VIX value due to the settlement procedure. Moreover, the settlement procedure uses the *actual open prices* of options traded on the settlement day, rather than bid–ask spreads; this procedure has caused biased VIX futures settlement prices compared to the equivalent spot VIX (see Pavlova and Daigler, 2008).
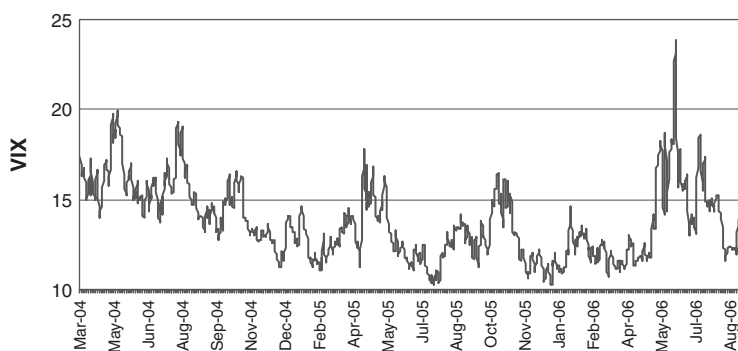
FIGURE 1
Values of the spot VIX for the forecast period.

Current VIX Derivatives Pricing Models

Researchers believe that a number of factors describe the time series behavior of volatility, and therefore the appropriate price of volatility futures. These factors include the mean-reversion of volatility, jumps in the S&P 500 index and in volatility, and the stochastic volatility of the VIX and possibly of the S&P 500 index. However, some factors are difficult to model, such as both the frequency and the size of jumps. Other factors, such as a risk premium, are difficult to estimate accurately with market data. Thus, most volatility futures models ignore or simplify many of these characteristics in order to provide a tractable model that is usable by traders. Consequently, a model to price futures could be biased, depending on the procedure employed to estimate the parameters of the factors included in the model. Therefore, the form of the appropriate model to price volatility futures essentially becomes an empirical question. This study examines the empirical aspect of pricing volatility futures by developing four models with differing characteristics and then tests these models empirically for several years of data.

Grunbichler and Longstaff (1996), Carr and Wu (2006) Zhang and Zhu (2006), Dotsis et al. (2007) and Lin (2007) provide VIX derivatives pricing models. Grunbichler and Longstaff employ a square root process for volatility, which is less general than the CEV approach, and does not include the possibility of jumps in the VIX; conversely, Duan and Yeh (2007) show that incorporating a jump risk factor is critically important for the VIX index, and they conclude that both jumps and volatility risks are priced. Zhang and Zhu model VIX futures based on the variance swap rate and a maximum likelihood estimation function by specifying a process for the square of the VIX instead of the

VIX itself.[8] Thus, Zhang and Zhu choose the same stochastic differential equation as Grunbichler and Longstaff (which is essentially a Cox, Ingersoll, and Ross (1985) model), with the volatility component being limited to a square root process with no jump components. In addition, their model includes a mean-reverting component and the variance of the S&P 500 index returns is tied to the S&P 500 index through a constant correlation coefficient. However, using the variance of the S&P 500 creates some estimation difficulties and the resultant Zhang and Zhu futures model is not a closed-form one. Moreover, in limiting valuation to (only) one specific day, their estimation procedure does not provide a reliable value for the volatility risk premium, so they consequently ignore the output of their own model and employ the estimate of the risk premium from previous research. None of these studies test their model on actual VIX futures data for more than one day.

Carr and Wu (2006) develop lower and upper bounds for the VIX futures price based on the forward volatility and forward variance swap rates. However, the arbitrage bounds that they develop are very wide, averaging $1\frac{1}{2}$ points above and below the fair price. Such a range is quite large given the typical level of the VIX futures and is significantly larger than the typical VIX futures bid–ask spread. Therefore, compared to the previous attempts to model VIX derivatives discussed above, our model includes more of the components that could drive the implied volatility of the market and is more straightforward in its approach.

Dotsis et al. (2007) explore the ability of several diffusion and jump-diffusion models to capture the dynamics of implied volatility indices. They also examine the pricing performance of the corresponding volatility futures pricing models. However, although they allow for jumps, the models are limited to either a mean-reverting square root process or a geometric Brownian motion. Our CEV formulation provides additional flexibility not found in such specifications.

Lin (2007) develops the most complete VIX futures price model. Her model includes jumps and stochastic volatility of the S&P 500 and risk-premium variables. However, as stated by the author (Lin, 2007, p. 1187) "This study primarily focuses on parameter estimation" rather than testing the accuracy of the model. More importantly, the relative completeness of the model is offset by its complexity. Consequently, it is unlikely that practitioners would adopt the Lin model, since the interaction of the terms is not easy to understand and the model would be difficult to program for normal trading purposes. Moreover, the likely sensitivity of fitting the model's parameters to changes in the sample

---

[8]Zhang and Zhu (2006, p. 522) specify that "The VIX squared is equal to the 30-calendar-day variance swap rate."

period (unreported in her study) raises the question of the accuracy and stability of the model.[9] In particular, the risk premium parameter fits the data to minimize the error. Consequently, a change in the behavior of volatility can cause a major change in the risk premium value. In fact, Garcia, Lewis, Pastorello, and Renault (2006) show that the estimated volatility risk premium for mark-dollar futures options differ across time periods for their data. Finally, the Lin and the Dotsis et al. (2007) articles are the only ones to empirically test their model, and the size and consistency of the errors reported in the Lin study are not clearly identified.[10]

EXAMINING VOLATILITY DERIVATIVE FACTORS

Factors Affecting the VIX and VIX Futures

There are a number of factors affecting the VIX and thus the VIX futures. It is well known that volatility is stochastic and mean-reverting over time. Jumps can exist in volatility, although the frequency and size of the jumps is not well known. The next two sections discuss the mean-reversion and jump characteristics of the VIX and its effect on VIX futures. We also discuss the bias in the method to calculate the VIX and issues relating to the settlement procedure of the VIX futures.

Figure 1, shown earlier, illustrates the mean-reverting characteristic of the VIX. Figure 2 provides an average term structure of the VIX futures over the entire sample, which shows that the futures flatten out for longer-term expirations, showing the longer-term effect of the mean-reverting characteristic on futures prices.

Figure 3 shows the slope coefficient obtained by regressing the change in the spot VIX on the change in the VIX futures. Thus, the futures only partially respond when there are more than 10 trading days until expiration, with the response coefficient increasing to near 1.0 when the contract is near expiration. The muted futures responsiveness for longer times to expiration supports the market's interpretation of the mean-reversion characteristic of the VIX.

Another factor affecting the new VIX is the existence of biases associated with the VIX calculation procedure compared to a model-free implied variance.

---

[9]An example of the potential sensitivity of Lin's parameter estimation is hinted at in her results; in particular, stock index price jumps are associated with better pricing for short-dated contracts, whereas volatility jumps are associated with better results for long-dated contracts. These differing results by time to expiration imply parameter instability. Moreover, the parameter estimates are only updated once per month.

[10]Interpreting the results in Lin is difficult. In Table I, a 15% volatility shows as 0.15. In Tables II and IV, the mean absolute pricing errors (MAE) range from 2.33 to 5.90. In addition, the results apparently are averages of monthly results, with no information on the variability of the results. Such broad averages can be misleading concerning the accuracy of a model.
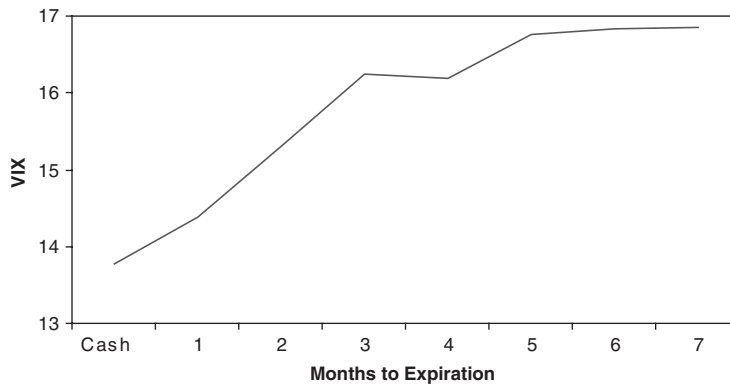
FIGURE 2
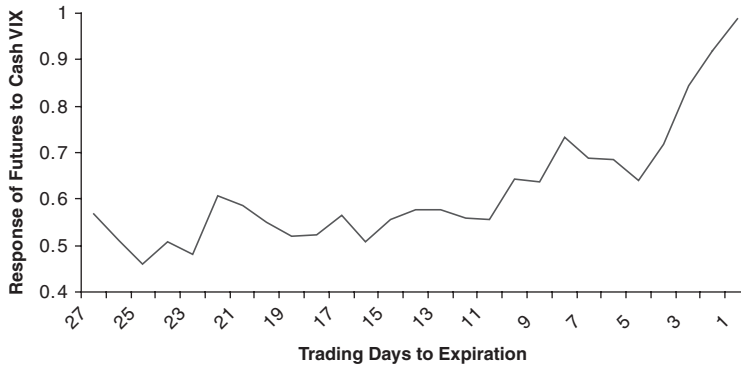The term structure of VIX futures prices.

FIGURE 3
Mean reversion of VIX futures: response of futures to the spot VIX.

Jiang and Tian (2007) conclude that the CBOE procedures underestimate (overestimate) true volatility when the underlying spot VIX is high (low). When the true volatility is 20%, then the error averages less than three basis points. For low volatility levels (10%) they find simulated differences of 19–25 basis points. Most of the bias occurs due to the larger strike price intervals that exist between options for farther out-of-the-money strikes. However, to the extent that the bias is systematically present then the bias should not significantly affect modeling the VIX process or the rankings of the models. More importantly, the Jiang and Tian simulations show that the biases are generally small for typical characteristics of the underlying S&P 500 options.

Finally, the settlement bias due to the procedure used to determine the settle price for the VIX futures is variable and can be large. Although the spot

VIX is determined based on the average of the bid and ask quotes of the relevant out-of-the-money S&P 500 options, the settlement procedure employs the actual *trade* prices of the nearby S&P 500 options at the opening on the futures settlement day (for those options trading at the open). This variation from the typical spot VIX calculation procedure is employed in order to reduce the risk for dealers/arbitrageurs of VIX futures so that better pricing will occur in the futures market during the life of the contract. Using opening option prices causes the settlement VIX value to tend toward either the average bid or the average ask price as VIX futures dealers unwind their strip of option positions. However, the size and direction of the difference from the actual spot VIX is uncertain and on average differs from the spot VIX by 0.26 index points ($260 per contract). The consequence of this bias is that the spot VIX and the VIX futures do not converge at expiration.[11] The market will therefore adjust for this characteristic of the VIX futures and price this factor into the futures, especially into the nearby expiration. Pavlova and Daigler (2008) discuss this issue in detail.

Jumps in the VIX

A key issue when modeling the spot VIX index is whether the time-series behavior of the VIX contains jumps, or whether the VIX is a continuous process.[12] Eraker, Johannes, and Polson (2003) provide evidence regarding the presence of jumps for individual implied volatility strikes for the S&P 500 and NASDAQ 100 index options by estimating jump-related parameters from Markov Chain Monte Carlo methods; they show that jumps in underlying asset volatility do impact option prices. Since jumps in volatility affect individual option prices, they would directly affect the value of the VIX. Related work by Bakshi and Cao (2004) shows that a model allowing for jumps in both the returns and volatility of the underlying equity provides the best fit for the data, as well as improving the equity option pricing performance relative to models constraining volatility to a purely continuous mean-reverting diffusion process. Similarly, Dotsis et al. (2007) show that jumps are important in capturing the dynamics of the implied volatility index over time.

Empirically determining whether a given series is continuous or whether it contains discrete jumps is a challenging endeavor. The distinction between a continuous series and a series with jumps is a subtle one, since observed data are essentially discrete and therefore are made up of a series of small jumps

---

[11]In addition, the reported spot VIX is based on a weighted average of two near-to-expiration S&P 500 options, whereas the VIX settlement employs only the next S&P 500 option expiration.
[12]In fact, the interest in jumps for financial markets has increased substantially in recent years, e.g. see Duffie, Pan, and Singleton (2000), Pan (2002), Eraker et al. (2003), Eraker (2004), and Dotsis et al. (2007).

(even when collected at high frequencies). Work in this area includes Wang (1995), Ait-Sahalia (2002, 2004), Ait-Sahalia and Jacod (2009), Carr and Wu (2003), Johannes (2004), and Johannes, Polson, and Stroud (2009). However, the difficulty with the various methods and jump criteria derived in these articles is that they either require tests on the process using the options prices of the underlying asset, or they need a pre-specification of the process itself in the form of a stochastic differential equation. These requirements create difficulties in our case, since only a short history of options on the VIX exists, and we do not want to restrict the VIX process by defining a stochastic equation before testing for jumps.

In order to disentangle jumps from a pre-specified diffusion process using only spot VIX data, we implement a new statistical non-parametric model-free test derived by Jiang and Oomen (2008). The methodology builds on another model-free technique, the bi-power variation test of Barndorff-Nielsen and Shephard (2006). We choose the Jiang and Oomen approach because their methodology possesses a faster rate of convergence to its asymptotic distribution and is more powerful in detecting jumps relative to the Barndorff–Nielsen and Shephard test. Moreover, the Jiang and Oomen test distinguishes between jumps and continuous movements without the need for option prices or a pre-specified stochastic process for the VIX itself.

The null hypothesis of the Jiang and Oomen test is that the sample path of the process is continuous. The alternative hypothesis is that the path does contain jumps. Identifying the correct hypothesis involves the computation of the "Swap Variance Jump Ratio Test" statistic, defined by:

$$V_{(0,T)}N\left(1 - \frac{RV_M(T)}{S\omega V_M(T)}\right) \to N(0, \sqrt{\Omega_{S\omega V}})$$ (2)

with T = the number of days, M = the number of intervals within a day, N = MT the total number of observations, and the remaining elements of Equation (2) defined or estimated robustly by:

$$V_{(0,T)} = \frac{1}{\mu_1^2}\frac{N}{N-1}\sum_{i=1}^{N-1}|r_{\delta,i+1}||r_{\delta,i}| \quad \text{where} \quad \mu_p = \frac{2^{p/2}\Gamma\left(\dfrac{p+1}{2}\right)}{\sqrt{\pi}}$$

$$\text{and} \quad r_{\delta,i} = \ln(VIX_{i\delta}) - \ln(VIX_{(i-1)\delta})$$

$$RV_M(T) = \sum_{i=1}^{N}r_{\delta,i}^2, \ S\omega V_M(T) = 2\sum_{j=1}^{N}\frac{VIX_{j\delta} - VIX_{(j-1)\delta}}{VIX_{(j-1)\delta}} - 2\ln(S_T/S_0)$$

$$\Omega_{S\omega V} = \frac{\mu_6}{9}\frac{N^3\mu_{6/6}^{-6}}{N-6+1}\sum_{j=0}^{N-6}\prod_{k=1}^{6}|r_{\delta,j+k}|^{6/6}$$ (3)

The null hypothesis of a continuous sample path is rejected if the test statistic is significantly negative or positive, since the swap variance jump test is two-sided. To calculate the test statistic we employ 5-minute VIX data ranging from May 2002 to September 2006, setting T equal to 1 month (21 trading days on average), and testing each month separately for the presence of jumps. The number of intervals within a day, M, therefore is equal to 80 and the total number of intervals N averages 21×80 or 1,680 per month. The time period covered is chosen to match the time period used both for the estimation of the parameters and for the pricing tests conducted later in the study. The jump test statistic results in Table I show that the null hypothesis of continuous sample paths (free of volatility jumps) is rejected for almost 90% of the cases. This overwhelming presence of jumps leads us to conclude that the inclusion of a jump component in the VIX model is warranted. Moreover, and relevant to changes in volatility, the *possibility* of a large downside jump (crash) in stock prices creates "fear" in traders' mindsets. This fear of a negative jump in returns is sufficient to increase option prices (especially in the put options skew) and therefore the VIX value, although the historical experience of downside jumps in returns is less than the model's prediction of the actual probability of a jump in the VIX. This fear has existed since the 1987 crash, increasing the implied volatility smile for out-of-the-money puts. Since the new VIX includes all tradable out-of-the-money options, changes in fear levels would directly affect the level of the VIX. Therefore, the *perception* of jumps can be an important factor for option prices and the VIX, supported by larger implied volatilities than realized volatilities. This potential for an abrupt VIX increase, combined with our swap variance jump ratio test results, warrants the inclusion of jumps as a feature of the VIX futures pricing model developed in the next section.

THE SPOT VIX PROCESS

A Model for the VIX Process

Developing a model to describe the evolution of the spot VIX over time is challenging, given the lack of understanding of how market participants fully "value" the various factors affecting volatility (e.g. since the VIX is determined from option values on the S&P500 index, the aggregate market's perception of future volatility levels affects the value of the VIX as well as how different S&P 500 strikes react to these perceptions). Therefore, we allow flexibility in the evolution of the spot VIX process by incorporating the following features: mean-reversion, CEV, and the potential for discontinuous jumps.

We propose the following class of processes to model the evolution of changes in the spot VIX over time:

TABLE I
Swap Variance Jump Ratio Test for the Spot VIX

Month	Critical Value (+/−)	Test Statistic	Continuous Path Null Hypothesis	Month	Critical Value (+/−)	Test Statistic	Continuous Path Null Hypothesis
May-02	0.1024	0.0898	"Cannot reject"	Aug-04	0.0193	0.0242	"Reject"
Jun-02	0.4002	1.1602	"Reject"	Sep-04	0.0431	0.1416	"Reject"
Jul-02	0.1722	0.1179	"Cannot reject"	Oct-04	0.0286	0.0925	"Reject"
Aug-02	0.1304	0.4289	"Reject"	Nov-04	0.0124	−0.1555	"Reject"
Sep-02	0.0706	0.6943	"Reject"	Dec-04	0.0207	0.1770	"Reject"
Oct-02	0.0557	0.0182	"Cannot reject"	Jan-05	0.0109	0.0694	"Reject"
Nov-02	0.0765	0.0220	"Cannot reject"	Feb-05	0.0117	0.2361	"Reject"
Dec-02	0.0530	0.1586	"Reject"	Mar-05	0.0179	0.0752	"Reject"
Jan-03	0.0762	0.1982	"Reject"	Apr-05	0.1283	−0.1482	"Cannot reject"
Feb-03	0.0365	0.1645	"Reject"	May-05	0.0425	−0.0314	"Reject"
Mar-03	0.0498	0.3546	"Reject"	Jun-05	0.0399	−0.0402	"Reject"
Apr-03	0.0385	0.0210	"Cannot reject"	Jul-05	0.0227	0.7103	"Reject"
May-03	0.0909	0.2010	"Reject"	Aug-05	0.0602	0.0488	"Cannot reject"
Jun-03	0.0738	0.1757	"Reject"	Sep-05	0.0252	0.1856	"Reject"
Jul-03	0.0535	0.1474	"Reject"	Oct-05	0.0491	−0.0336	"Cannot reject"
Aug-03	0.0808	0.1771	"Reject"	Nov-05	0.0302	−0.0026	"Cannot reject"
Sep-03	0.0408	0.1853	"Reject"	Dec-05	0.0201	−0.0692	"Reject"
Oct-03	0.0161	−0.0234	"Reject"	Jan-06	0.0165	0.1756	"Reject"
Nov-03	0.0102	0.1476	"Reject"	Feb-06	0.0197	0.0442	"Reject"
Dec-03	0.0175	0.0252	"Reject"	Mar-06	0.0104	−0.0113	"Reject"
Jan-04	0.0123	0.1585	"Reject"	Apr-06	0.0112	0.1297	"Reject"
Feb-04	0.0057	0.0470	"Reject"	May-06	0.1423	0.2091	"Reject"
Mar-04	0.0354	0.1918	"Reject"	Jun-06	0.0937	−0.2009	"Reject"
Apr-04	0.0595	0.0968	"Reject"	Jul-06	0.0890	0.1117	"Reject"
May-04	0.0259	0.2095	"Reject"	Aug-06	0.0208	0.1248	"Reject"
Jun-04	0.0138	−0.0286	"Reject"	Sep-06	0.0110	0.1648	"Reject"
Jul-04	0.1315	0.2347	"Reject"				

This table reports the Swap Variance Jump Ratio Test statistics using Equation (2), with the corresponding two-sided critical values at the 95% level. The jump ratio statistic tests for the presence of jumps in the 5-min VIX process for each month between May 2002 and September 2006. The null hypothesis is that a continuous sample path for the VIX exists, while the alternative hypothesis is that a jump is present in the process.

$$dV_t = (\alpha - \beta V_t)dt + \sigma V_t^\gamma dW_t + y_t dq_t \tag{4}$$

where V_t represents the level of the VIX at time t, W_t is the standard Brownian motion, and y_t is the jump size that would shift the VIX from V_t to $V_t + y_t$ should a jump occur. The frequency of jumps is driven by a Poisson process represented by q_t, where the probability of a jump taking place between time t and time $t + dt$ is given by λdt. Therefore,

$$\text{Prob}[dq_t = 1] = \lambda dt$$
$$\text{Prob}[dq_t = 0] = 1 - \lambda dt \tag{5}$$

Furthermore, we assume that the jump size y_t is exponentially, identically, and independently distributed as:

$$y_t \sim \text{Exponential}(\mu_y)$$

where μ_y is the mean size of the VIX jump distribution. Duffie et al. (2000), in their analytical treatment of general affine jump-diffusion processes, note that in order to preserve the positive nature of the volatility (or VIX), one must preclude two-sided volatility jumps. The exponential distribution is well suited to this task as well as constituting a fairly robust class of one-parameter distributions, despite its relative simplicity. In particular, actual market behavior shows that market volatility easily jumps to higher levels during sudden market turmoil, whereas it typically declines more gradually when the market returns to a normal state. The use of one-sided jumps is consistent with this stylized feature of the market.

One line of research in the volatility literature focuses on finding and/or incorporating factors that explain changes in volatility over time. An example of an early work in this area is Franks and Schwartz (1991), who attempt to explain changes in implied volatility with various economic variables. A more recent article that specifically examines VIX modeling is Lu and Zhu (2010), who examine the term structure of VIX futures contracts. They show that a log-linear model, along with three underlying state variables, performs best for futures pricing purposes.[13] Another line of research focuses on the volatility process directly, without necessarily investigating which underlying factors might be relevant and/or without linking the volatility process to a market index.[14]

---

[13]Other research examining factors affecting volatility (and examined in Lu and Zhu) includes Bollerslev and Zhou (2002), Alizadeh, Brandt, and Diebold (2002), Chernov, Ghysels, Gallant, and Tauchen (2003), Chacko and Viceira (2003), Zhang and Zhu (2006), Lin (2007), Christoffersen, Jacobs, Ornthanalai, and Wang (2008), Li and Zhang (2008), and Egloff, Leippold, and Wu (2009).

[14]A non-exhaustive list of research opting for this approach includes Grunbichler and Longstaff (1996), Moraux, Navatte, and Villa (1999), Detemple and Osakwe (2000), Whaley (2000), Daouk and Guo (2004), Brenner, Ou, and Zhang (2006), Dotsis et al., (2007), Huskaj and Nossman (2009), Chourdakis and Dotsis (2009), and Mencia and Sentana (2009).

We employ the latter approach. Consequently, we model the volatility index as a "reduced-form" statistic in the sense that additional hidden variables that could be attributable to its evolution over time are neither searched for nor specified; this approach is described in Dotsis et al. (2007). Whereas such a modeling approach at first can seem incomplete relative to including potential underlying volatility-driving factors, Bakshi, Cao, and Chen (1997), Dumas, Fleming, and Whaley (1998), and Psychoyios and Skiadopoulos (2006), among others, demonstrate that relatively basic models are able to deliver hedging capabilities comparable to significantly more complex models. Additionally, whereas identifying latent, underlying volatility-driving factors can have many useful purposes, it can also make hedging more challenging. In particular, statistically identified factors are not tradable because they do not correspond to any actual tradable asset; they can therefore make hedging a complex endeavor. For all these reasons we choose relative simplicity over complexity and model the VIX dynamics directly.

We name the model in Equation (4) the mean-reverting constant elasticity of variance with jumps model (CEVJ). This general model essentially allows for the modeling of mean reversion of the VIX, the presence of jumps in market volatility, and a flexible diffusion component. This general model nests several others. When $\gamma = 1/2$ and $\lambda = 0$ the model collapses to the well-known Cox et al. (1985) term structure model (CIR) initially developed to model interest rates and adapted by Grunbichler and Longstaff (1996) to model the VIX process. The typical CIR model allows for mean reversion but not for jumps in the process, and constrains the diffusion component to a square root process. When $\lambda = 0$ the model collapses to the mean-reverting CEV model, first proposed by Chan, Karolyi, Longstaff, and Anthony (1992) to study the dynamics of short-term interest rates. It is similar to the CIR model but offers added flexibility in the diffusion component. Finally, when $\gamma = 1/2$ the model becomes the CIR model with Jumps (CIRJ).

Estimating the Model

As a first step in our analysis we employ spot VIX index closings from April 2002 to September 2006 to estimate and compare the four potential models. In a subsequent section we employ the complete dataset to examine the pricing accuracy of the futures model. We estimate the parameters of the continuous-time VIX process using Hansen's (1982) Generalized Method of Moments (GMM) as applied to the following discrete-time econometric specification:

$$V_{t+1} - V_t = (\alpha - \beta V_t)\Delta t + \sigma V_t^\gamma \sqrt{\Delta t}\xi_t + \gamma_t \Delta Q_t \qquad (6)$$

where ξ_t is i.i.d. $N(0, 1)$ and is uncorrelated with the discrete Poisson process ΔQ_t. We define θ to be the parameter vector containing the six elements α, β, σ^2, γ, μ_y and λ. Given $\varepsilon_{t+1} = V_{t+1} - V_t - (\alpha - \beta V_t)\Delta t$, we compute six moments of ε_{t+1} and condition on the volatility level, yielding the following vector of moment conditions:

$$
f_t(\theta) = \begin{bmatrix}
\varepsilon_{t+1} - \mu_y \lambda \Delta t \\
\varepsilon_{t+1}^2 - \sigma^2 V_t^{2\gamma}\Delta t - 2\mu_y^2 \lambda \Delta t \\
\varepsilon_{t+1}^3 - 6\mu_y^3 \lambda \Delta t \\
\dfrac{\pi}{2}|\varepsilon_{t+1}\varepsilon_t| - (V_{t+1}V_t)^\gamma \sigma^2 \Delta t \\
|\varepsilon_{t+1}\varepsilon_t\varepsilon_{t-1}|^{4/3} - (V_{t+1}V_tV_{t-1})^{4\gamma/3}\sigma^4(\Delta t)^2\left(2^{2/3}\dfrac{\Gamma(7/6)}{\Gamma(1/2)}\right)^3 \\
|\varepsilon_{t+1}\varepsilon_t\varepsilon_{t-1}\varepsilon_{t-2}| - (V_{t+1}V_tV_{t-1}V_{t-2})^\gamma \sigma^4\left(\dfrac{2}{\pi}\Delta t\right)^2
\end{bmatrix} \otimes \begin{bmatrix} 1 \\ V_t \end{bmatrix}
$$

$$(7)$$

with $E[f_t(\theta)] = 0$. The first three moments simply are developed from evaluating the expressions for $E(\varepsilon)$, $E(\varepsilon^2)$, and $E(\varepsilon^3)$. The last three moment conditions are specified to disentangle the conditional variance and jumps in the spirit of Ait-Sahlia (2004). They are called bi-power variation, tri-power variation, and quadratic variation, respectively in Barndorff–Nielsen and Shephard (2004, 2006). They are derived from theorems 1 and 2 of Barndorff–Nielsen and Shephard (2004) and Equation (7) of Tauchen and Zhou (2010), whereby:

$$
\mu_k = 2^{k/2}\Gamma\left(\frac{k+1}{2}\right)\Big/\Gamma\left(\frac{1}{2}\right)
$$

Therefore, we have $\mu_1 = 2^{1/2}\Gamma(1)/\Gamma(\frac{1}{2}) = (2/\pi)^{1/2}$ and $\mu_{4/3} = 2^{2/3}\Gamma(7/6)/\Gamma(\frac{1}{2})$, allowing us to compute the following moments:

$$
|\varepsilon_{t+1}\varepsilon_t| = (\mu_1)^2(V_{t+1}V_t)^\gamma\sigma^2\Delta t
$$

$$
|\varepsilon_{t+1}\varepsilon_t\varepsilon_{t-1}|^{4/3} = (\mu_{4/3})^3(V_{t+1}V_tV_{t-1})^{4\gamma/3}\sigma^4(\Delta t)^2
$$

$$
|\varepsilon_{t+1}\varepsilon_t\varepsilon_{t-1}\varepsilon_{t-2}| = (\mu_1)^4(V_{t+1}V_tV_{t-1}V_{t-2})^\gamma\sigma^4
$$

The theoretical expectation's sample counterpart is $g_T(\theta) = 1/T \sum_{t=1}^{T} f_t(\theta)$, with this expression theoretically converging to zero for an infinite sample size. The parameter vector is chosen to minimize the quadratic form or J-statistic

$$
J_T(\theta) = g'_T(\theta)W_T(\theta)g_T(\theta) \tag{8}
$$

where $W_T(\theta)$ is a positive-definite symmetric weighting matrix. The optimal weighting matrix is given by $W_T(\theta) = S^{-1}(\theta)$, where $S(\theta) = E[f_t(\theta)f_t'(\theta)]$, as shown by Hansen (1982). To account for possible autocorrelation and heteroskedasticity, we use the Newey–West (1987) estimator with lags of up to one-third of the length of the data. The minimized value of $J_T(\theta)$ asymptotically follows a χ^2 distribution with degrees of freedom equal to the number of moments, net of the number of the parameters estimated. The χ^2 statistic provides a goodness-of-fit test for the model estimated.

We perform criterion-difference D-tests to evaluate the performance of the various nested models. The D-statistic is given by $D = T[J_T^R(\theta) - J_T^u(\theta)]$ and is analogous to a likelihood ratio test. It asymptotically follows a χ^2 distribution with degrees of freedom equal to the difference in the number of parameters between the unrestricted and restricted models. $J_T^u(\theta)$ represents the J-statistic of the unrestricted general CEVJ model, and $J_T^R(\theta)$ represents the J-statistic of the restricted alternative, both of which are estimated from the minimized quadratic form. However, for purposes of conducting D-tests, the restricted nested models must use the weighting matrix estimated from the unrestricted model when computing the objective functions for both the unrestricted and restricted cases. The purpose of the D-test is to determine whether fewer parameters in the model are justified, which is particularly important for the jump component. In our tests the CEVJ model (the most complete model) is employed as the base case reference model, since CEVJ is used as the "unrestricted" model, has the most parameters, and includes all parameters found in the other models.

Results for the Spot VIX Model

Panel A of Table II reports the GMM estimates of the parameters and their associated t-statistics for all of the discussed unrestricted and restricted models for the dynamics of the VIX. The table also provides the relevant J-statistics and their associated p-values to test the individual models. The J-statistics are χ^2 distributed, with the degrees of freedom equal to the number of moment conditions net the number of parameters. In the GMM setting, a model is successful if the overall model estimation yields a large p-value, since not rejecting the null hypothesis that the moment conditions are equal to zero (on average) means that the model is correctly specified. The CEVJ model clearly displays the highest level of performance, with a J statistic of 4.68 and a p-value of 0.5855, followed by the CEV and CIR-J models, which yield J statistics of 14.16 and 13.99 and p-values of 0.0777 and 0.0513, respectively. The simplest CIR model is rejected at a significance level of less than 1%.

TABLE II

Estimates of Alternative Models on the Spot VIX Process

	Unrestricted				Restricted			
	CEV-J		CIR-J		CEV		CIR	
	Values	t-stats	Values	t-stats	Values	t-stats	Values	t-stats
Panel A. Estimates for Alternative Models on the Spot VIX Process								
α	3.1750	4.8414	3.2472	5.0462	1.9626	3.4811	1.3129	2.4470
β	23.536	4.7660	23.967	4.9364	14.6770	3.3914	9.315	2.2576
α^2	8.4259	0.91316	0.097613	9.3807	8.1598	0.97903	0.48035	37.324
γ	1.5937	5.7028			1.6066	6.3871		
μ_y	0.20991	0.99107	0.46576	8.7238				
λ	0.02546	0.61387	0.018684	2.7549				
J-stats	4.680		13.9910		14.157		58.148	
d.f.	6		7		8		9	
p-values	0.58546		0.051345		0.077753		0.000000	
Panel B. D-test of the Alternative Models								
D-stats			14.0840		4.8169		17.7598	
d.f.			1		2		3	
p-values			0.00017		0.0900		0.00014	

This table reports the estimates of the parameters for the spot VIX process and their associated *t*-statistics using four alternative models to describe spot VIX prices. The data range from April 2002 to September 2006. The Generalized Methods of Moments (GMM) procedure is employed for the estimation. The VIX data have been scaled by 1/100. The *J*-statistics are reported with the associated degree of freedom (d.f.) and *P*-values. The parameters are estimated from the discrete-form equation:

$$V_{t+1} - V_t = (\alpha - \beta V_t)\Delta t + \sigma V_t^\gamma \sqrt{\Delta t}\xi_t + y_t \Delta Q_t$$

$$E[y_t] = \mu_y, \ E[\Delta Q_t] = \lambda \Delta t$$

All estimates of the α and β parameters are statistically significant. The significance of the β parameter across all models is a strong indication of the presence of mean-reversion in the spot VIX index. The conditional variance of the spot VIX in our model is sensitive to the interaction between the CEV γ and the σ^2 values, since the conditional variance at time t depends on both γ and σ^2. Moreover, the σ^2 parameter is significant in the CIRJ and CIR models (with γ constrained to 0.5), since the constant σ^2 must include most of the effect of the conditional variance in order to compensate for the restriction on the γ parameter. Conversely, the σ^2 parameter in the CEVJ and CEV models (with unconstrained CEV γ values) is insignificant, since the unconstrained CEV γ can capture significant percentage of the dynamics of the conditional variance.

Our estimates of the γ parameter in the CEVJ and CEV models are almost identical, showing that the estimates of the CEV's γ are independent from whether or not jumps are included. This proves that we successfully disentangled the conditional variance from the jumps by incorporating the bivariate power variation through the last three GMM moment conditions. The estimate of jump frequency λ is not statistically significant, a finding consistent with the volatility-jump literature and potentially due to the fact that the CEV already allows for fairly rich dynamics on its own. However, the average jump size μ_y and the jump intensity λ become statistically significant in the CIRJ estimation when the CEV's γ is restricted to 0.5. Intuitively, since the constrained CEV's γ of no more than 0.5 is much less than the non-constrained estimated CEV's value of 1.6, the misspecification in the CIRJ model forces part of the significant conditional variation to be absorbed by the jump parameters.

Panel B of Table II computes GMM D-statistics in order to provide additional justification for the ranking of our models. In this table all D-statistics are computed in relation to the CEVJ unrestricted model, our most comprehensive model. When performing a D-test in GMM, the null hypothesis that the restrictions on the parameters are "correct" is tested by calculating the difference in J-values between the restricted and unrestricted models, while using the weighting matrix that comes from the unrestricted (CEVJ) model in both cases. The D-statistic is distributed as a χ^2 value with degrees of freedom equal to the number of restrictions that the restricted model imposes on the unrestricted model. The p-values obtained for each restricted model show whether the corresponding restrictions are appropriate (if a p-value is greater than the critical tail mass of 5%) or inappropriate (if a p-value is less than the critical tail mass of 5%). The D-statistics reveal that the CEV model is a close second to the CEVJ model, as the p-value of 0.09 suggests that the null hypothesis of no difference in performance between the CEV and CEVJ models cannot be rejected at the 5% significance level, although the null hypothesis can be rejected at the 10% significance level. On the other hand, the CIR and CIRJ

models are rejected at the 5 and 1% significance levels, showing that the restrictions associated with CIR modeling are strongly rejected.

In conclusion, Table II tells us that the unrestricted CEVJ model empirically outperforms the restricted models, in particular outperforming the CIR model proposed by Grunbichler and Longstaff (1996). Additionally, models with an unconstrained exponent γ (the CEVJ and CEV models) outperform models with γ constrained to 0.5 (the CIRJ and CIR). The results show that the spot VIX process is mean-reverting, sensitive to CEV, and potentially affected by jump risk, depending on the model. Our general unrestricted specification encompasses all of these unique features and is the model that performs best for the in–sample data. We now examine the pricing of the VIX futures contract and testing the "out-of-sample" performance of the four models for pricing actual VIX futures contracts.

PRICING THE VIX FUTURES CONTRACT

We show in the previous section how alternative stochastic models can perform differently in their attempt to capture the dynamics of the spot VIX index. We now compare the pricing performance of these models using VIX futures data. First, we derive the expected VIX futures price for the general class of CEVJ models employed here. Second, we discuss the properties of these related theoretical prices. Finally, we empirically examine the performance of these alternative models in the pricing of VIX futures. Of course, any such test jointly examines the pricing ability of the model in conjunction with any mispricing of the VIX futures by market participants.

The VIX Futures Model

Our model to price VIX futures includes the mean-reversion characteristic of the VIX, jumps (for two of the four models), Brownian motion associated with random changes in the VIX, and the time to (futures) expiration. Therefore, the model is more complete than the current models in the literature attempting to price volatility futures (with the possible exception of Lin, 2007). In order to develop the VIX futures model let $F(V_t, t, T)$ denote the price of a futures contract expiring at time T on the underlying spot VIX price V_t at time t. Following Grunbichler and Longstaff (1996), we use the futures price as the expected value of the spot VIX price V_t at time T, under the risk-neutral measure Q:

$$F(V_t, t, T) = E^Q[V_T] \tag{9}$$

Based on our specification of the spot VIX process, we show in the appendix that for our most general CEVJ model Equation (9) implies a VIX futures price of:

$$F(V_t, t, T) = V_t e^{-\beta(T-t)} + \frac{\alpha + \mu\lambda}{\beta}(1 - e^{-\beta(T-t)}). \tag{10}$$

We take advantage of a result established in Henderson et al. (2004), Daouk and Guo (2004), Psychoyios and Skiadopoulos (2006), and Dotsis et al. (2007) to deal with the risk premium issue. These authors show that model rankings do not depend on the magnitude of the volatility risk premium, since model prices are a monotonic function of the price of risk. In particular, Henderson et al. (2004) show that the prices of European call and put options are a monotonically decreasing function of the price of volatility risk when a fairly general stochastic volatility model setup is employed. More specifically, Daouk and Guo (2004) compare volatility options pricing performance using various types of GARCH models, as well as providing a comparison with the Grunbichler and Longstaff (1996) model, while assuming that volatility risk is not priced. In the process, they also test whether the presence of the volatility risk premium (or lack thereof) affects their results. They find that their model comparison is not affected by their choice of the volatility risk premium, even for a wide range of risk premium values. Psychoyios and Skiadopoulos (2006) study whether volatility options are superior to standard options when trying to hedge volatility risk, as well as investigate the pricing and hedging performance of several volatility option pricing models. Similar to Daouk and Guo (2004), they also assume away the price of volatility risk, as well as test whether their results would differ due to the "injection" of volatility risk into the volatility parameter. They conclude that the order of the performance of their compared models is not perturbed by the presence of a volatility risk premium. Dotsis et al. (2007) estimate and compare the pricing performance of various diffusion processes aimed at capturing the dynamics of volatility indices over time. In the vein of Daouk and Guo (2004), as well as Psychoyios and Skiadopoulos (2006), Dotsis et al. (2007) set the volatility risk premium equal to zero in their calibration exercise. They confirm the robustness of their results to the choice of different levels of the volatility risk premium by examining the pricing errors calculated for a wide range of the volatility risk premia, concluding that the rankings of the processes do not depend on the chosen risk premium level. Finally, in our model the volatility risk premium would "enter" via the β parameter. From Equation (10) it is fairly straightforward to conclude, everything else held constant, that different risk premium levels would simply monotonically affect the predicted futures price of the model. Therefore, for purposes of ranking models and their respective features, we set the volatility price of risk equal to zero.

As with the spot models, when $\gamma = 1/2$ and $\lambda = 0$, the above model becomes the CIR model; when $\lambda = 0$ the model is the mean-reverting CEV model, and when $\gamma = 1/2$ the model becomes the CIR model with Jumps

(CIRJ).[15] Note that the γ parameter disappears as it is eliminated when the expectation is taken in the derivation of the formula.

The formula for the VIX futures is similar in *structure* to the models described in Grunbichler and Longstaff (1996) and in Dotsis et al. (2007), but our model includes both a jump and a CEV feature. Some of the properties of the theoretical futures price stemming from these models are similar. However, the estimation of the spot VIX process is sensitive to the unconstrained CEV γ and the presence of a jump. For instance, the estimated values of α and the speed-of-mean-reversion parameter β decline significantly when moving from the CEVJ to the CEV model. Consequently, there are important fundamental differences among these various specifications.

Results for VIX Futures Pricing

We estimate the conditional parameters for each of our four models using the generalized method of moments procedure described previously by separately employing the last two years of spot VIX data prior to the date of each day's estimation in order to determine the model parameters. Thus, the model is re-estimated each day using a new set of input values (the oldest one day of data is replaced with the most recent day's data). The resultant estimated parameters using data prior to day $t - 1$ are then employed in combination with the current value of the VIX to determine the fair VIX futures price. Actual futures prices are compared to the estimated fair futures price from the model for the period March 29, 2004 to September 12, 2006.[16] We then calculate the signed percentage error (SPE) and the absolute percentage error (APE) between the actual futures price and the modeled fair futures price for each daily futures forecast and for each futures expiration. The APE is calculated as follows:

$$APE = \frac{|f_t - F(V_t, t, T)|}{F(V_t, t, T)}. \tag{11}$$

The SPE is equivalent to the APE, except that one removes the absolute value restriction. Errors are determined in basis points of the VIX (one basis point is worth \$10 per futures contract). Finally, we compare the alternative models' performance in terms of the MSPE and the MAPE and their basis point equivalents.

Panel A of Table III provides the MSPE and MAPE results for each of the four models for each futures expiration from May 2004 to September 2006,

---

[15]The first stage model of the evolution of the VIX over time does not provide an estimate of the current value of the VIX. Therefore, the current value of the spot VIX is used as the current expected value of the VIX.
[16]Both the spot VIX and VIX futures "trade" until 4:15, eliminating any timing issues.

TABLE III

Average Percentage and Basis Point Errors from VIX Models

Days until Expir.	No. Exp. Mo.	Avg. future volume	Avg. level spot VIX	MSPE				MAPE				St Dev Signed Error				St Dev Absolute Error			
				CEVJ	CEV	CIR	CIRJ	CEVJ	CEV	CIR	CIRJ	CEVJ	CEV	CIR	CIRJ	CEVJ	CEV	CIR	CIRJ
Panel A: Mean Percentage Error Close Results																			
1–15	25	365.71	13.89	2.95	2.12	2.78	3.10	3.55	3.03	3.42	3.59	2.54	2.34	2.30	2.36	2.11	1.85	1.93	2.00
16–30	25	193.83	13.78	3.31	1.08	2.56	3.45	4.26	3.24	3.73	4.09	2.72	2.24	2.06	2.14	2.35	1.86	1.78	1.82
31–45	25	244.75	13.93	4.80	1.15	3.07	4.11	6.01	4.41	4.90	5.07	3.07	2.36	2.14	2.29	2.52	2.10	1.87	1.98
46–60	25	134.02	13.82	5.59	0.70	2.91	4.17	6.45	4.13	4.89	5.29	2.08	1.59	1.46	1.56	1.98	1.42	1.37	1.40
61–75	13	214.61	13.60	8.02	1.75	4.35	5.63	8.33	4.32	5.83	6.54	2.73	1.91	1.83	2.10	2.48	1.78	1.55	1.88
75–90	12	143.83	13.64	9.96	2.56	5.24	7.04	10.01	4.98	6.60	7.38	2.33	1.76	1.71	1.99	2.26	1.52	1.58	1.87
91–120	10	254.83	13.85	10.87	1.88	4.58	6.22	10.91	4.72	6.27	6.94	2.31	1.76	1.89	2.13	2.23	1.54	1.85	2.09
121–150	9	126.52	13.41	11.10	0.78	3.73	5.57	11.12	3.51	4.65	6.02	2.33	1.60	1.65	1.97	2.26	1.51	1.49	1.72
151–230	8	66.45	13.61	12.19	−0.89	1.70	3.59	12.21	5.12	5.44	5.89	3.85	3.44	3.62	4.07	3.82	3.08	2.98	3.32
Panel B: Mean Basis Point Error Results																			
1–15	25	365.71	13.89	41.50	28.97	37.82	42.62	49.81	42.00	47.11	49.69	35.42	32.69	32.06	32.91	29.56	25.85	26.79	27.79
16–30	25	193.83	13.78	46.99	12.21	32.92	47.10	61.48	46.93	52.55	57.78	40.04	33.04	30.41	31.53	34.17	27.54	26.09	26.48
31–45	25	244.75	13.93	67.98	9.92	38.11	55.64	88.24	66.42	71.41	73.43	46.51	35.95	32.76	34.89	37.80	32.20	28.42	29.94
46–60	25	134.02	13.82	81.56	3.46	37.25	58.86	96.65	64.56	73.28	79.00	31.89	24.38	22.48	23.86	30.21	21.82	21.07	21.19
61–75	13	214.61	13.60	121.78	21.45	61.16	82.84	126.27	66.12	87.56	98.41	42.52	29.78	28.88	32.89	38.70	27.57	24.23	29.44
75–90	12	143.83	13.64	149.36	32.24	72.78	102.5	150.25	77.13	98.43	108.84	35.94	27.31	26.47	30.67	34.66	23.60	24.44	28.66
91–120	10	254.83	13.85	167.03	20.59	63.41	92.26	167.65	77.65	98.32	107.13	36.15	27.81	29.77	33.15	34.82	24.48	28.93	32.21
121–150	9	126.52	13.41	181.09	8.86	57.64	90.08	181.43	58.49	74.76	97.81	39.30	27.65	28.46	33.16	38.24	25.86	25.53	28.76
151–230	8	66.45	13.61	202.50	−21.72	22.54	57.02	202.72	88.95	91.69	96.90	64.11	57.65	60.58	67.66	63.66	51.68	49.89	54.41

This table provides the MSPE (mean signed percentage error), MAPE (mean absolute percentage error), MSBPE (mean signed basis point error), and MABPE (mean average basis point error) for all four models for the difference between the estimated fair (model) futures price and the actual futures price. The standard deviations of the errors are also provided.

segregated by the number of days until futures expiration. The table also shows the standard deviations for the MSPE and MAPE. The results are dominated by the CEV model (without jumps), which has the lowest MSPE and MAPE values across the differing ranges of days until expiration categories. Panel A shows that the CEV model's average signed error is typically below 2%, with errors hovering around 1% for 16–60 days until futures expiration.[17] The CIR model is consistently second, with the jump models providing larger errors.

The absolute errors in Table III are larger, but the CEV model still dominates. The larger increase in the size of the percentage errors from MSPE to MAPE for the CEV model shows that the CEV model's errors are more consistently distributed around zero, while the other models tend to be consistently positive. The CEV model's distribution is consistent with a model employed by market participants. The standard deviations of the errors are also smaller for the CEV model compared to the other models.[18] Therefore, these results show that more complicated models under-perform relative to simpler models for out-of-sample data, in spite of better in-sample accuracy. Specifically, our results demonstrate the superiority of the CEV model over the CEVJ model, even though the *J*- and *D*-test results from the previous section show the CEVJ model to be superior to or statistically equivalent to the in-sample results. Thus, these out-of-sample results show the somewhat surprising lack of importance or sensitivity of the model to the jump parameter when actual futures pricing is employed. However, the *J* and *D* tests do perform a type of "best-fit" bias in its evaluation, whereas the market results provide an "out-of-sample test" of the model. These results contradict Lin (2007), where jump models provide the best results and the best model varies by the time to futures expiration.

Panel B of Table III shows the equivalent errors from Panel A in basis points (1 basis point is 0.01 of a VIX index point). These results give dollar equivalent errors not scaled by the size of the VIX. The CEV signed basis point errors are 3.46–12.21 for 16–60 days until futures expiration, a very small error, which would be similar to the typical bid–ask spread for the futures contract. The signed basis point errors for the other models are much larger. However, the absolute basis point errors are much larger, i.e. over 42 basis points for all time ranges until futures expiration.

Table III provides only average errors over all of the futures expirations. Figure 4 shows the size of the MSPEs for each separate expiration for all four

---

[17]The larger average percentage error for the 1–15 days until futures expiration for the CEV model is most likely due to the futures settlement procedure and its resultant bias, as previously discussed.
[18]In order to examine the effect of liquidity, the analysis was repeated using the settlement futures prices, which circumvents adverse effects due to stale trading. The results were comparable to the closing value model results presented here.
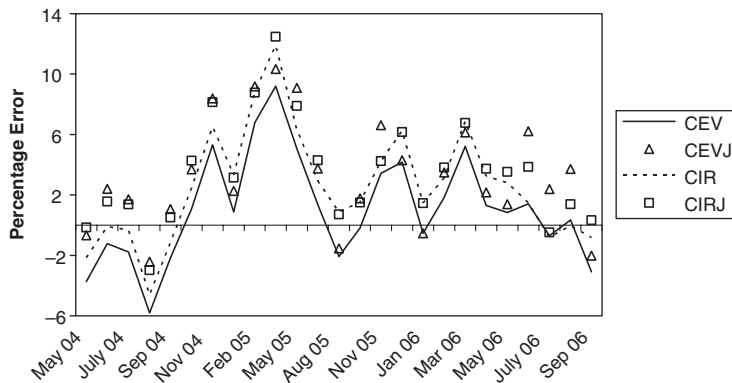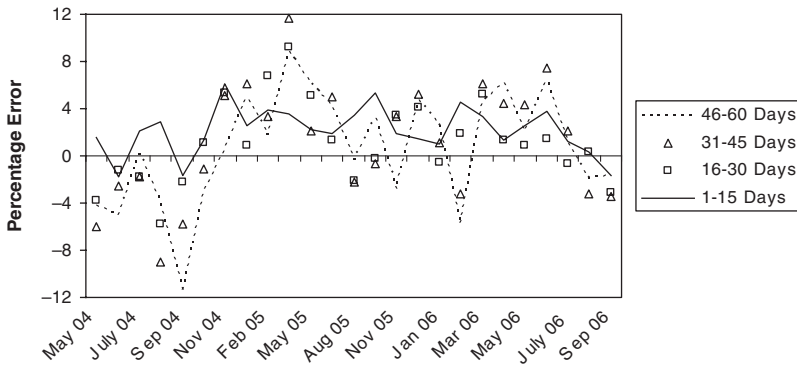
FIGURE 4
Comparison of VIX futures models.

models, using the 16–30 days until futures expiration category. Figure 4 shows that the CEV model consistently has the lowest percentage errors when the errors are positive, but often has larger negative errors when they occur. From month-to-month the tendency is for the models to cluster near a similar level of percentage error, suggesting that the errors may be associated more with how futures are priced than with differences between the models.

The futures pricing performance of our parsimonious model compares favorably with studies that showcase more complex specifications. For example, Lu and Zhu (2010) show that a log-linear model with three underlying state variables performs best for futures pricing purposes. Their *in-sample* percentage pricing errors (shown on a graph) vary between approximately −3% and +4% for their best model. Our CEV model's *out-of-sample* mean APEs vary from approximately +3% and +5%.[19] Lin's (2007) complex VIX futures pricing formula displays out-of-sample percentage pricing errors ranging from nearly −7% to +2%, depending on the model used, which typically indicates an over-pricing fit for the VIX futures. Zhang and Zhu (2006) report out-of-sample percentage pricing errors ranging from +2% to +44%, whereas Dotsis et al. (2007) report mean squared percentage pricing errors ranging from 0.01 to 0.29, translating roughly into mean APEs from +7% to +53%.

Figure 5 shows the percentage errors for the superior CEV model for differing times to futures expiration for each trading month. While variability does exist in the errors shown in Figure 5, there is a general tendency for the errors to cluster by the trading month, showing that unique factors relating to pricing for given expiration months influence the size of the errors. Moreover, there is a tendency for the errors to be similar for adjacent months. Therefore, the level

[19]Lu and Zhu's out-of-sample performance only encompasses several days of data and is not in percentage form, making direct comparison impossible.

FIGURE 5
Comparing the CEV model for differing times to expiration.
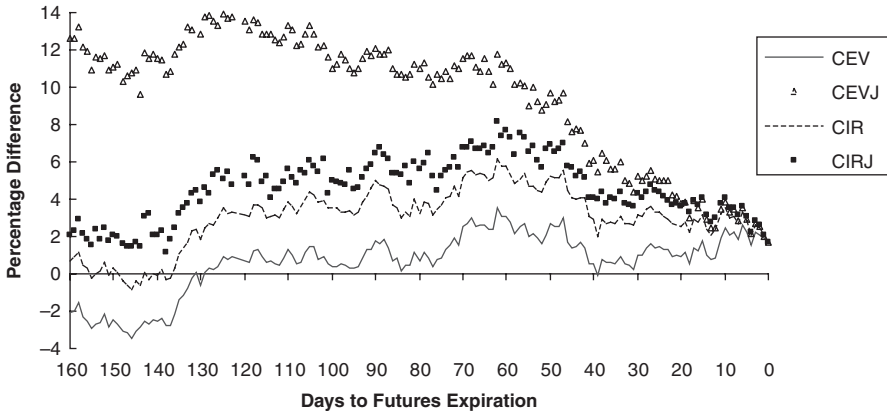
FIGURE 6
Comparing the VIX models by day to futures expiration.

of the VIX could be a factor affecting the size of the errors. Consequently, the results in Figures 4 and 5 suggest that all stochastic volatility models could be missing an important factor, or the models possess a bias in estimating the model's parameters.

Figure 6 shows the percentage errors for each model by day until futures expiration. Again, the CEV model is consistently the best model across time to expiration, with the CEVJ having much larger errors than the other models for most of the days until expiration.

Overall, the results show that the CEV model is consistently superior to the other models in pricing actual VIX futures, whether examined by mean signed or APEs, basis point errors, time to futures expiration, or expiration

month. In Dotsis et al. (2007) the geometric Brownian motion specification outperforms the square root specification, showing that a model with a γ parameter set to 1 is superior to a model with a γ parameter set to 1/2. As seen earlier, our CEV model actually allows γ to be unrestricted; when estimated γ is near 1.6, which implies that the VIX process "needs" the less restrictive CEV feature for the futures model to perform the best.

Examining the Model vs. Market Differences

In this section we examine the factors that affect the differences between the model values and the actual VIX futures market values. The regression of the differences between model and market prices on these factors is:

$$M(t) - P_F(t) = a_1 + a_2\sqrt{T} + a_3 log(P_F(t)) + a_4 100(P_F(t) - P_C(t))/P_C(t)$$

$$+ a_5 log(V_t) + a_6 100(P_C(t) - P_C(t-1))/P_C(t-1) \qquad (12)$$

where $M(t)$ = the model's value at time t, $P_F(t)$ = the price of the VIX futures at time t, $P_C(t)$ = the spot VIX value at time t, T = the time to futures expiration in terms of a fraction of a year, and V_t = volume of the futures contract at time t. The reasoning for including each of these variables is as follows:

- Time to futures expiration T: examines whether the model incorporates the time to expiration in the same way as does the market.
- Level of the futures price $P_F(t)$: a higher VIX futures value reflects more volatility and uncertainty in the market (the volatility of volatility), which could be valued differently than a low level of the VIX futures.
- Difference between the futures and spot VIX values $(P_F(t) - P_C(t))/P_C(t)$: shows the extent that the spot VIX differs from the mean-reverting estimate given by the VIX futures value.
- Volume V_t: the size of the volume provides a measure of the liquidity in the market.
- Change in the spot VIX $(P_C(t) - P_C(t))/P_C(t)$: the percentage change in the spot VIX is a measure of jumps in the VIX.

Table IV provides the results of the regression explained above. These results could be employed to adjust the estimation of the model, *if* the goal is to bring the model's prices closer to the market. The R^2 in the table are significant and large for all four models, with the CEVJ model possessing the largest R^2 value. Time to futures expiration and the level of the VIX futures price are not

Factors Affecting the Difference Between the Model and Market Values

	CEVJ	CEV	CIR	CIRJ
Intercept	62.335	95.724	110.870	82.122
	17.29	**29.16**	**29.89**	**19.86**
Time	0.021	0.031	0.039	0.013
	1.19	1.91	2.16	0.63
Level VIX Futures	0.326	0.177	0.148	0.283
	2.12	1.27	0.94	1.60
Difference Futures minus Spot VIX	−0.462	−0.291	−0.223	−0.292
	−32.38	**−22.38**	**−15.22**	**−17.85**
Level of Volume	−13.209	−19.275	−22.163	−16.360
	−18.11	**−29.01**	**−29.52**	**−19.55**
Difference VIX(t) − VIX(t−1)	1.991	0.644	0.737	0.900
	55.51	**19.72**	**19.97**	**21.87**
R^2 (%)	61.6	45.1	38.3	29.7

This table shows the results of regressions of Model vs. Market differences on time to expiration, the level of the VIX futures price, the percentage difference between the futures and spot VIX values, the level of volume (liquidity), and the percentage difference between the time t and time $t − 1$ values in the spot VIX (jumps). The equation estimated on the model vs. market differences for each of the four models examined in this study is $M(t) − P_F(t) = a_1 + a_2\sqrt{t} + a_3 \log(P_F(t)) + a_4 100(P_F(t) − P_C(t))/P_C(t) + a_5 \log(V_t) + a_6 100(P_C(t) − P_C(t − 1))/P_C(t − 1)$, where t = time to futures expiration in a fraction of a year. The first line for each variable is the coefficient value and the second line is the t-value; significant t-values at the 5% level are shown in bold.

significant for any of the models. The relevance of the remaining three variables is described as follows:

- The difference between the futures and spot VIX values shows that the mean-reverting characteristic of the VIX is valued differently in the market compared to the models given here. The models employ historical means of the VIX to determine the mean-reverting value, while the market would naturally use its estimate of the appropriate current average value. During most of the time period of this study the spot VIX is below the long-run average for the VIX. Therefore, when there is a long time to futures expiration, the futures contract is much higher than the spot VIX, which is associated with larger negative values in the model vs. market values. Consequently, the negative coefficient is consistent with the model underpricing the longer-term effect of the mean-reverting component relative to how the market prices this factor.

- The significance of the volume variable is associated with the liquidity of the VIX futures contract, especially for the more deferred contracts with a longer time to expiration. A larger volume variable is associated with a model value that is farther below the market value of the futures. Conversely, deferred contracts with lower volume generally fluctuate less in price, negating the effect of a low volume.

- The larger the percentage change in the spot VIX (with large changes considered to be a proxy for jumps), the greater the model value is above the futures price. This variable possesses the largest *t*-value for the CEVJ (jump) model. Thus, the model values jump to be more important for the VIX futures price than does the market. This result could be associated with the lower spot VIX value for most of this period, i.e. less concern in the stock market for a "crash event."

SUMMARY AND CONCLUSIONS

Volatility derivatives are a new class of assets that have received considerable interest in the industry, but only limited interest in academic research. In fact, the volume of exchange-traded and over-the-counter volatility derivatives has continued to grow over time. Volatility derivatives provide the opportunity to speculate and hedge risks that are difficult to duplicate in existing instruments, even options. One problem with research into volatility derivatives is the lack of comprehensive and accurate models to value these derivatives.

This study develops and tests a comprehensive model to price VIX futures contracts. The model is separated into four sub-models, the CEV model and the Cox–Ingersoll–Ross (CIR) model, both with and without jump components. The CEV model is the model that best prices actual futures values, regardless of time to futures expiration or contract month. Examination of the differences between the model and market prices find that variables representing mean reversion, liquidity, and jumps all affect these differences. Further research in this area should consider these factors, as well as whether the model needs improvement or whether the market is "wrong" in its valuation procedures.

The results from this study leave some intriguing questions. For example, is the VIX futures model efficient in its pricing across trades? Can one earn profits from using the model to decide when to buy and sell VIX futures? Are the differences between the model and market prices due to an incomplete model, model calibration problems, or mispriced futures contracts? Further research into volatility derivatives could provide further insights to these questions.

APPENDIX: PROOF OF THE VIX FUTURES PRICING FORMULA

Let the risk-neutral process for the volatility be defined as:

$$dV_t = (\alpha - \beta V_t)dt + \sigma V_t^\gamma dW_t + y_t dq_t \tag{A1}$$

The integral form of Equation (A1) is

$$V_t = V_0 + \int_0^t (\alpha - \beta V_u)du + \sigma \int_0^t V_u^\gamma dW_u + \int_0^t y_u dq_u \tag{A2}$$

Taking expectations and using the properties of the jump distribution, we obtain

$$E(V_t) = V_0 + \int_0^t [\alpha - \beta E(V_u)]du + \int_0^t \mu\lambda du \tag{A3}$$

Differentiating (A3) with respect to time provides

$$\frac{dE(V_t)}{dt} = \alpha - \beta E(V_t) + \mu\lambda \tag{A4}$$

Noticing that

$$\frac{d}{dt}[e^{\beta t}E(V_t)] = e^{\beta t}\left[\beta E(V_t) + \frac{dE(V_t)}{dt}\right] \tag{A5}$$

We can see that

$$\frac{d}{dt}[e^{\beta t}E(V_t)] = e^{\beta t}(\alpha + \mu\lambda) \tag{A6}$$

Finally, integrating and solving for $E(V_t)$ produces the desired result:

$$E(V_t) = V_0 e^{-\beta t} + \frac{(\alpha + \mu\lambda)}{\beta}(1 - e^{-\beta t}) \qquad\qquad \square$$

BIBLIOGRAPHY

Ait-Sahalia, Y. (2002). Telling from discrete data whether the underlying continuous-time model is a diffusion. The Journal of Finance, 57, 2075–2112.

Ait-Sahalia, Y. (2004). Disentangling diffusion from jumps. Journal of Financial Economics, 74, 487–528.

Ait-Sahalia, Y., & Jacod, J. (2009). Testing for jumps in a discretely observed process. Annals of Statistics, 37, 184–222.

Alizadeh, S., Brandt, M., & Diebold, F. (2002). Range-based estimation of stochastic volatility models. Journal of Finance, 52, 1047–1092.

Bakshi, G., & Cao, C. (2004). Risk-neutral kurtosis, jumps, and option pricing: Evidence from most actively traded firms. University of Maryland/Pennsylvania State Working Paper.

Bakshi, G., Cao, C., & Chen, Z. (1997). Empirical performance of alternative option pricing models. Journal of Finance, 52, 2003–2049.

Barndorff-Nielsen, O., & Shephard, N. (2004). Power and bipower variation with stochastic volatility and jumps. Journal of Financial Econometrics, 2, 1–37.

Barndorff-Nielsen, O., & Shephard, N. (2006). Econometrics of testing for jumps in financial economics using bipower variation. Journal of Financial Econometrics, 4, 1–30.

Bollerslev, T., & Zhou, H. (2002). Estimating stochastic volatility diffusion using conditional moments of integrated volatility. Journal of Econometrics, 109, 33–65.

Brenner, M., Ou, E., & Zhang, J. (2006). Hedging volatility risk. Journal of Banking and Finance, 30, 811–821.

Carr, P., & Wu, L. (2003). What type of process underlies options? A simple robust test. The Journal of Finance, 58, 2581–2610.

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

Chacko, G., & Viceira, L. (2003). Spectral GMM estimation of continuous-time processes. Journal of Econometrics, 116, 259–292.

Chan, K. C., Karolyi, A., Longstaff, F., & Anthony, S. (1992). An empirical comparison of alternative models of the short-term interest rate. The Journal of Finance, 47, 1209–1227.

Chernov, M., Ghysels, E., Gallant, A. R., & Tauchen, G. (2003). Alternative models for stock price dynamics. Journal of Econometrics, 116, 225–257.

Chourdakis, K., & Dotsis, G. (2009). Maximum likelihood estimation and dynamic asset allocation with non-affine volatility processes. University of Essex, working paper.

Christoffersen, P., Jacobs, K., Ornthanalai, C., & Wang, Y. (2008). Option valuation with long-run and short-run volatility. Journal of Financial Economics, 90, 272–297.

Cox, J., Ingersoll, J., & Ross, S. (1985). A theory of the term structure of interest rates. Econometrica, 53, 363–384.

Daouk, H., & Guo, J. (2004). Switching asymmetric GARCH and options on a volatility index. Journal of Futures Markets, 24, 251–283.

Detemple, J., & Osakwe, C. (2000). The valuation of volatility options. European Finance Review, 4, 21–50.

Dotsis, G., Psychoyios, D., & Skiadopoulous, G. (2007). An empirical comparison of continuous-time models of implied volatility indices. Journal of Banking and Finance, 31, 3584–3603.

Duan, J.-C., & Yeh, C.-Y. (2007). Jumps and volatility risk premiums implied by VIX. University of Toronto and National Taiwan University Working Paper.

Duffie, D., Pan, J., & Singleton, K. (2000). Transform analysis and asset pricing for affine jump-diffusions. Econometrica, 68, 1343–1376.

Dumas, B., Fleming, J., & Whaley, R. E. (1998). Implied volatility functions: Empirical tests. Journal of Finance, 53, 2059–2106.

Egloff, D., Leippold, M., & Wu, L. (2010). The term structure of variance swap rates and optimal variance swap investments. Journal of Financial and Quantitative Analysis, forthcoming.

Eraker, B. (2004). Do stock prices and volatility jump? Reconciling evidence from spot and option prices. Journal of Finance, 59, 1367–1403.

Eraker, B., Johannes, M., & Polson, N. (2003). The impact of jumps in volatility and returns. The Journal of Finance, 58, 1269–1300.

Franks, J., & Schwartz, E. (1991). The stochastic behaviour of market variance implied in the prices of index options. The Economic Journal, 101, 1460–1475.

Garcia, R., Lewis, M. A., Pastorello, S., & Renault, E. (2006). Estimation of objective and risk-neutral distributions based on moments of integrated volatility. Working paper, CIRANO and CIREQ, HEC Montreal, Montreal, Canada.

Grunbichler, A., & Longstaff, F. (1996). Valuing futures and options on volatility. Journal of Banking and Finance, 20, 985–1001.

Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. Econometrica, 50, 1029–1054.

Henderson, V., Hobson, D., Howison, S., & Kluge, T. (2004). A comparison of option prices under different pricing measures in a stochastic volatility model with correlation (working paper). Princeton, NJ: Princeton University.

Hibbert, A.-M., Daigler, R., & Dupoyet, B. (2008). A behavioral explanation for the negative asymmetric return-volatility relation. Journal of Banking and Finance, 32, 2254–2256.

Huskaj, B., & Nossman, M. (2009). A market model for the VIX futures market. Lund University, working paper.

Jiang, G. J., & Oomen, R. C. (2008). Testing for jumps when asset prices are observed with noise—A swap variance approach. Journal of Econometrics, 144, 352–370.

Jiang, G. J., & Tian, Y. S. (2007). Extracting model-free volatility from option prices: An examination of the VIX index. The Journal of Derivatives, 32, 35–60.

Johannes, M. (2004). The statistical and economic role of jumps in continuous-time interest rate models. Journal of Finance, 59, 227–260.

Johannes, M., Polson, N., & Stroud, J. (2009). Optimal filtering of jump diffusions: Extracting latent states from asset prices. Review of Financial Studies, 22, 2759–2799.

Li, G., & Zhang, C. (2008). On the number and dynamic features of state variables in option pricing (working paper). The Hong Kong University of Science & Technology.

Lin, Y.-N. (2007). Pricing VIX futures: Evidence from integrated physical and risk-neutral probability measures. The Journal of Futures Markets, 27, 1175–1217.

Low, C. (2004). The fear and exuberance from implied volatility of S&P 100 index options. Journal of Business, 77, 527–546.

Lu, Z., & Zhu, Y. (2010). Volatility components: The term structure dynamics of VIX futures. The Journal of Futures Markets, 30, 230–256.

Mencia, J., & Sentana, E. (2009). Valuation of VIX derivatives. CEMFI and Bank of Spain working paper.

Moraux, F., Navatte, P., & Villa, C. (1999). The predictive power of the French market volatility index: A multi horizons study. European Finance Review, 2, 303–320.

Pan, J. (2002). The jump-risk premia implicit in options: Evidence from an integrated time-series study. Journal of Financial Economics, 63, 3–50.

Pavlova, I., & Daigler, R. (2008). The non-convergence of the VIX futures and expiration. Review of Futures Markets, 17, 201–223.

Poon, S., & Granger, C. W. (2003). Forecasting volatility in financial markets: A review. Journal of Economic Literature, 41, 478–539.

Psychoyios, D., & Skiadopoulous, G. (2006). Volatility options: hedging effectiveness, pricing, and model error. The Journal of Futures Markets, 26, 1–31.

Tauchen, G., & Zhou, H. (2010). Realized jumps on financial markets and predicting credit spreads. Journal of Econometrics, forthcoming.

Wang, Y. (1995). Jump and sharp cusp detection by wavelets. Biometrika, 82, 385–397.

Whaley, R. E. (2000). The investor fear gauge. Journal of Portfolio Management, 26, 12–17.

Zhang, J. E., & Zhu, Y. (2006). VIX futures. Journal of Futures Markets, 26, 521–532.