DISCUSSION PAPER SERIES

No. 7619

VALUATION OF VIX DERIVATIVES

Javier Mencía and Enrique Sentana

FINANCIAL ECONOMICS

Centre for Economic Policy Research

www.cepr.org

Available online at: www.cepr.org/pubs/dps/DP7619.asp

ISSN 0265-8003

VALUATION OF VIX DERIVATIVES

Javier Mencía, Banco de España
Enrique Sentana, Centre for Monetary and Financial Studies (CEMFI) and CEPR

Discussion Paper No. 7619
January 2010

Centre for Economic Policy Research
53–56 Gt Sutton St, London EC1V 0DG, UK
Tel: (44 20) 7183 8801, Fax: (44 20) 7183 8820
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 7619

January 2010

ABSTRACT

Valuation of VIX Derivatives*

We conduct an extensive empirical analysis of VIX derivative valuation models over the 2004-2007 bull market and the subsequent financial crisis. We show that existing models yield large distortions during the crisis because of their restrictive volatility mean reverting assumptions. We propose generalisations with a time varying central tendency, jumps and stochastic volatility, analyse their pricing performance, and their implications for the term structures of VIX futures and options, and the option volatility "skews". We find that a model combining central tendency and stochastic volatility is required to reliably price VIX futures and options, respectively, across bull and bear markets.

JEL Classification: G13
Keywords: central tendency, jumps, stochastic volatility, term structure and volatility skews

Javier Mencía
Banco de España
Alcalá 48
28014 Madrid
SPAIN
Email: javier.mencia@bde.es

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=163095

Enrique Sentana
CEMFI
Casado del Alisal 5
E-28014 Madrid
SPAIN
Email: sentana@cemfi.es

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=111330

* We are grateful to Dante Amengual, Torben Andersen, Max Bruche, Peter Carr, Rob Engle, Javier Gil-Bazo, Francisco Peñaranda, Eduardo Schwartz and Neil Shephard as well as seminar participants at Universidad de Murcia, Universitat Pompeu Fabra and the SoFiE 2009 European Conference (Geneva) for helpful comments, discussions and suggestions. Of course, the usual caveat applies. Financial support from the Spanish Ministry of Science and Innovation through grant ECO 2008-00280 (Sentana) is gratefully acknowledged. The views expressed in this paper are those of the authors, and do not reflect those of the Bank of Spain.

Submitted 16 December 2009

1 Introduction

It is now widely accepted that the volatility of financial assets changes stochastically over time, with fairly calmed phases being followed by more turbulent periods of uncertain length. For financial market participants, it is of the utmost importance to understand the nature of those variations because volatility is a crucial determinant of their investment decisions. Although many model-based and model-free volatility measures have been proposed in the academic literature (see Andersen, Bollerslev, and Diebold, 2009, for a recent survey), the Chicago Board Options Exchange (CBOE) volatility index, widely known by its ticker symbol VIX, has effectively become the standard measure of volatility risk for investors in the US stock market. The goal of the VIX index is to capture the volatility (i.e. standard deviation) of the S&P 500 over the next month implicit in stock index option prices. Formally, it is the square root of the risk neutral variance of the S&P 500 over the next 30 calendar days, reported on an annualised basis. Despite this rather technical definition, both financial market participants and the media pay a lot of attention to its movements. To some extent, its popularity is due to the fact that VIX changes are negatively correlated to changes in stock prices. The most popular explanation is that investors trade options on the S&P 500 to buy protection in periods of market turmoil, which increases the value of the VIX. For that reason, some commentators refer to it as the market's fear gauge, even though a high value does not necessarily imply negative future returns.

But apart from its role as a risk indicator, nowadays it is possible to directly invest in volatility as an asset class by means of VIX derivatives. Specifically, on March 26, 2004, trading in futures on the VIX began on the CBOE Futures Exchange (CFE). They are standard futures contracts on forward 30-day implied vols that cash settle to a special opening quotation (VRO) of the VIX on the Wednesday that is 30 days prior to the 3rd Friday of the calendar month immediately following the expiring month. Further, on February 24, 2006, European-style options on the VIX index were also launched on the CBOE. Like VIX futures, they are cash settled according to the difference between the value of the VIX at expiration and their strike price. They can also be interpreted as options on VIX futures, which one can exploit to simplify their valuation. VIX options and VIX futures are among the most actively traded contracts at CBOE and

CFE, averaging close to 140,000 contracts combined per day. The high interest in these products is mainly due to their ability to hedge the risks of investments in the S&P 500 index. In this sense, Szado (2009) finds that such strategies do indeed provide significant protection, especially in downturns.

Although these new assets certainly offer additional investment and hedging opportunities, their correct use requires reliable valuation models that adequately capture the features of the underlying volatility index. Somewhat surprisingly, several theoretical approaches to price VIX derivatives appeared in the academic literature long before they could be traded. Specifically, Whaley (1993) priced volatility futures assuming that volatility follows a Geometric Brownian Motion (GMB). As a result, his model does not allow for mean reversion in volatility, which, as we shall see, is at odds with the empirical evidence. The two most prominent mean reverting models proposed so far have been the square root process (SQR) process considered by Grünbichler and Longstaff (1996), and the log-normal Ornstein-Uhlenbeck (log-OU) process analysed by Detemple and Osakwe (2000).

Several authors have recently assessed the empirical performance of these volatility derivative pricing models. In particular, Zhang and Zhu (2006) study the empirical validity of the SQR model by first estimating its parameters from VIX historical data, and then assessing the pricing errors of VIX futures implied by those estimates. Following a similar estimation strategy, Dotsis, Psychoyios, and Skiadopoulos (2007) also use VIX futures data to evaluate the gains of adding jumps to the SQR diffusion. In addition, they estimate a GMB process. Surprisingly, this model yields reasonably good results, but the time span of their sample is perhaps too short for the mean reverting features of the VIX to play any crucial role.

In turn, Wang and Daigler (2009) compare the empirical fit of the SQR and GMB models using option data. They also find evidence supporting the GMB assumption. However, one could alternatively interpret their findings as evidence in favour of the log-OU process, which also yields the Black (1976) formula if the underlying instrument is a VIX futures contract, but at the same time is consistent with mean reversion. Despite this empirical evidence, both the SQR and the log-OU processes show some glaring deficiencies in capturing the strong persistence of the VIX, which produces large and lasting deviations of this index from its long run mean. In contrast, the implicit assumption in those models is that volatility mean reverts at a simple, non-negative exponential rate. Such a limitation becomes particularly apparent during bearish stock markets, in which volatility typically experiences large increases and remains at high levels for long periods. Unfortunately, the sample periods considered in the existing studies cover mostly the relatively long and quiet bull market that ended in the summer of 2007.

In this context, the initial objective of our paper is to study the empirical ability of existing mean-reverting models to price derivatives on the VIX over a longer sample span that also includes data from the unprecedented 2007-2009 financial crisis, which provides a unique testing ground for our study. To do so, we use an extensive database that comprises the entire history of futures and European options on the VIX since their inception until April 2009. As a result, we can study whether the SQR and log-OU models are able to yield reliable in- and out-of-sample prices in a variety of market circumstances.

Our findings indicate that although the log-OU process provides a better fit than the SQR model, especially for options, its performance deteriorates during the market turmoil of the last part of our sample. For that reason, we generalise the log-OU process by considering several empirically relevant extensions: a time-varying central tendency in the mean, jumps, and stochastic volatility. The presence of jumps in volatility has already been considered in the literature (see e.g. Eraker, Johannes, and Polson, 2003). A central tendency, which was first introduced by Jegadeesh and Pennacchi (1996) and Balduzzi, Das, and Foresi (1998) in the context of term structure models, allows the "average" volatility level to be time-varying, while stochastic volatility permits a changing dispersion for the (log) volatility index.

Since we combine futures and options data, we can also assess which of those additional features is more relevant to price futures and which one is more important for options. As typically done in the literature, we calibrate the parameters of the models by minimising the discrepancies between the actual derivative prices and the theoretical prices, which we often derive by inverting the conditional characteristic function using Fourier methods. As forcefully argued by Christoffersen and Jacobs (2004), such an approach ensures the internal consistency of our analysis, since we use the same loss function to estimate and validate the models. But we also go beyond pricing errors, and analyse the implications of the aforementioned extensions for the term structures of VIX futures and options, and the option volatility "skews", all of which are of considerable independent interest.

Given that the VIX index is computed from options data on the S&P 500, our analysis also has important implications for modelling this broad stock market index. In particular, our results imply that stochastic volatility models for the S&P 500 should allow for a slower mean reversion than usually considered, as well as for time-varying volatility of volatility.

The rest of the paper is organised as follows. We describe the data in Section 2, and explain our estimation strategy in Section 3. Then, we assess the empirical performance of existing models in Section 4 and consider our proposed extensions in Section 5. Finally, we conclude in Section 6. Auxiliary results are gathered in an Appendix.

2 Preliminary data analysis

2.1 The CBOE Volatility Index

VIX was originally introduced in 1993 to track the Black-Scholes implied volatilities of options on the S&P 100 with near-the-money strikes (see Whaley, 1993). The CBOE redefined the index in 2003, renamed the original index as VXO, and released a time series of daily closing prices starting in January 1990 (see Carr and Wu, 2006). Nowadays, VIX is computed in real time using as inputs the mid bid-ask market prices for most calls and puts on the S&P 500 index for the front month and the second month expirations with at least 8 days left (see CBOE, 2009).[1] Figure 1a displays the entire historical evolution of the VIX. Between January 1990 and April 2009 its average closing value was 19.9. As other volatility measures, though, it is characterised by swings from low to high levels, with a temporal pattern that shows mean reversion over the long run but displays strongly persistent deviations from the mean during extended periods. The lowest closing price (9.31) corresponds to December 22, 1993. Figure 1b, which focuses on the sample period in our derivatives database, shows that volatility was also remarkably low between March 2004 and July 2007, with values well below 20. During this period,

[1]Currently, CBOE applies the VIX methodology to 3-month options on the S&P 500 (VXV), as well as 1-month options on the most important US stock market indices: DJIA (VXD), S&P100 (VXO), Nasdaq-100 (VXN) and Russell 2000 (RVX). They also construct analogous short term volatility indices for Crude Oil (OVX), Gold (GVZ) and the US $/€ exchange rate (EVZ).

the lowest value was 9.89 on January 24, 2007, in what some have called "the calm before the storm". Over the following year, though, VIX increased to values between 20 and 35. Finally, in the autumn of 2008 it reached unprecedented levels. In particular, the largest historical closing price (80.86) took place on November 20, 2008, although on October 24 the VIX reached an intraday value of 89.53. These three markedly different regimes offer a very interesting testing ground to analyse the out-of-sample performance of valuation models for volatility derivatives. In particular, we can assess if the models calibrated with pre-crisis data perform well under the extreme conditions of the 2008-2009 financial crisis.

In order to characterise the time series dynamics of the VIX, we have estimated several ARMA models using the whole VIX historical observations from 1990 to 2009. Figure 2a compares the sample autocorrelations of the log VIX with those implied by the estimated models. There is clear evidence of high persistence, with a first-order autocorrelation above 0.98 and a slow rate of decay for higher orders. Consequently, an AR(1) model seems unable to capture the shape of the sample correlogram. An alternative illustration of the failure of this model is provided by the presence of positive partial autocorrelations of orders higher than one, as Figure 2b confirms. Therefore, it is necessary to introduce a moving average component to take into account this feature. An ARMA(1,1) model, though, only offers a slight improvement. In contrast, an ARMA(2,1) model turns out to yield autocorrelations and partial autocorrelations that are much closer to the sample values. As we shall see below, our preferred continuous time models have ARMA(2,1) representations in discrete time.

2.2 VIX derivatives

Our sample contains daily closing bid-ask mid prices of futures and European put and call options on the VIX, which we downloaded from Bloomberg. We consider the whole history of these series until April 2009. In practice, this means that our sample period starts in March 2004 for futures, and in February 2006 in the case of options. In terms of maturity, we have data on all the contracts with expirations between May 2004 and November 2009 for futures, and March 2006 and August 2009 for options. CFE may list futures for up to 9 near-term serial months, as well as 5 months on the February quarterly cycle associated to the March quarterly cycle for options on the S&P 500. In turn, CBOE initially lists some in-, at- and out-of-the-money strike prices, and then adds new strikes as the VIX index moves up or down. Generally, the options expiration dates are up to 3 near-term months plus up to 3 additional months on the February quarterly cycle.

Following Dumas, Fleming, and Whaley (1998), we exclude derivatives with fewer than six days to expiration due to their illiquidity. In addition, we only consider those prices for which open interests and volumes are available. All in all, we have 7,158 and 45,280 prices of futures and options, respectively. Of those option prices, 29,240 correspond to calls and 16,040 to puts. We proxy for the riskless interest rate by using the daily Eurodollar rates at 1-week, 1, 3 and 6 months and 1 year, which we interpolate to match the maturities of the futures and option contracts that we observe.

Figure 3a shows the number of futures prices per day in our database. We have between 2 and 4 daily prices until 2006; afterwards, around 8 prices become available on average. Figure 3b indicates that the number of option prices per day is also smaller at the beginning of the sample, but it tends to stabilise in 2007 at around 20 for puts and 40 for calls.

Figure 4a shows the evolution of the VIX term structure implicit in VIX futures. This figure clearly confirms that there was a substantial level shift in the last quarter of 2008. The negative slope during that period, though, indicates that the market did not expect the VIX to remain at such high values forever. Figure 4b compares futures prices with the VIX. As can be seen, most future prices were above the volatility index until mid 2007, which suggests that market participants perceived that the VIX was too low during those years. For instance, on July 14, 2004, the VIX was 13.76, while the price of VIX futures expiring on November 20, 2004, was 20.14. The spot price reflected the expected volatility for the period of July 14 to August 13, while the futures price reflected the expected volatility for the period of November 20 to December 20, 2004. In contrast, most prices were below the VIX in the last months of our sample, which confirms that volatility was then perceived to be above its long run mean.

Figure 4c focuses on the VIX term structure at four particularly informative days. Futures prices were extremely low on July 16, 2007, just before the Dow Jones Industrial Average closed above 14,000 for the first time in its history, and Bear Stearns disclosed that its two subprime hedge funds had lost nearly all of their value. Future prices had already risen significantly by August 15, 2008, one month before the collapse of Lehman Brothers. Nevertheless, prices were much higher on November 20, 2008, which is the day in which the VIX reached its maximum historical closing value at 80.86. The increase is particularly remarkable at the short end of the curve. Since then, though, VIX futures prices have significantly come down. Still, on April 7, 2009, they remained well above their 2007 levels.

3 Pricing and estimation strategy

We assume that there is a risk free asset with instantaneous rate r. Let $V(t)$ be the VIX value at time t. We define $F(t,T)$ as the actual strike price of a futures contract on $V(t)$ that matures at $T > t$. Similarly, we will denote the prices at t of call and put options maturing at T with strike price K by $c(t,T,K)$ and $p(t,T,K)$, respectively. Importantly, since the VIX index is a risk neutral volatility forecast, not a directly traded asset, there is no cost of carry relationship between the price of the futures and the VIX (see Grünbichler and Longstaff, 1996, for more details). There is no convenience yield either, as in the case of futures on commodities. Therefore, absent any other market information, VIX derivatives must be priced according to some model for the risk neutral evolution of the VIX. This situation is similar, but not identical, to term structure models.

Let M index the asset pricing models that we consider. Then, the theoretical futures price implied by model M will be:

$$F_M(t,T,V(t),\boldsymbol{\phi}) = E_M^{\mathbb{Q}}[V(T)|I(t),\boldsymbol{\phi}], \tag{1}$$

where $\mathbb{Q}$ indicates that the expectation is evaluated at the risk neutral measure, $\boldsymbol{\phi}$ is the vector of free parameters of model M, and $I(t)$ denotes the information available at time t, which includes $V(t)$ and its past values.

We can analogously express the theoretical value of a European call option with strike K and maturity at T under this model as

$$c_M[t,T,K,F(t,T),\boldsymbol{\phi}] = \exp(-r\tau)E_M^{\mathbb{Q}}[\max(V(T)-K,0)|I(t),\boldsymbol{\phi}], \tag{2}$$

where $\tau = T-t$. Nevertheless, we can exploit the fact that $V(T) = F(T,T)$ to price calls using the futures contract that expires on the same date as the underlying instrument, instead of the actual volatility index. In this way, we make sure that the pricing errors of options are not caused by distortions in our futures valuation formulas. Similarly, European put prices $p(t, T, K)$ under M can be easily obtained from the put-call-forward parity relationship

$$p_M[t, T, K, F(t, T), \boldsymbol{\phi}] = c_M[t, T, K, F(t, T), \boldsymbol{\phi}] - \exp(-r\tau)[F(t, T) - K].$$

We estimate the parameters of all the models that we consider by minimising the sum of the square differences between the actual and model-based prices of futures and options. In this way, we use VIX futures and option prices to choose the parameters of the risk neutral stochastic process for the VIX volatility index that provide the best empirical fit. We weight those square differences by the corresponding volume of the contracts at each day to ensure that the more liquid quotes have a higher weight. We consider three estimation samples that correspond to the distinct volatility phases described in Section 2: (i) until 15-July-2007; (ii) until 15-Aug-2008 and (iii) full sample. Finally, we evaluate the in- and out-of-sample empirical fit of the models over those periods. Thus, we can assess model performance following two major volatility increases in global stock markets.

4 Existing one-factor models

4.1 Model specification

We first compare the two mean-reverting volatility models that have been used so far in the literature: the square root and the log Ornstein-Uhlenbeck processes. As we mentioned before, Grünbichler and Longstaff (1996) proposed the square root process (SQR) to model a standard deviation index. This model, which was used by Cox, Ingersoll, and Ross (1985) for interest rates and Heston (1993) for the instantaneous variance of stock prices, satisfies the diffusion

$$dV(t) = \kappa[\bar{\theta} - V(t)]dt + \sigma\sqrt{V(t)}dW(t),$$

where $W(t)$ is a Brownian motion. As is well known, the distribution of $2cV(T)$ given $V(t)$ is a non-central chi-square with $\nu = 4\kappa\bar{\theta}/\sigma^2$ degrees of freedom and non-centrality parameter $\psi = 2cV(t)\exp(-\kappa\tau)$, where

$$c = \frac{2\kappa}{\sigma^2(1 - \exp(-\kappa\tau))}.$$

As a result, the price of the futures contract (1) can be expressed in this case as

$$F_{SQR}(t, T, V(t), \kappa, \bar{\theta}, \sigma) = \bar{\theta} + \exp(-\kappa\tau)[V(t) - \bar{\theta}]. \tag{3}$$

We can interpret $\bar{\theta}$ as the long-run mean of $V(t)$, since the conditional expected value of the volatility index converges to $\bar{\theta}$ as τ goes to infinity. In addition, κ is usually interpreted as a mean-reversion parameter because the higher it is, the more quickly the process reverts to its long run mean.[2]

The call price formula (2) for this model becomes

$$\begin{aligned} c_{SQR}(t, T, K, \kappa, \bar{\theta}, \sigma) &= V(t)\exp(-(\kappa + r)\tau)[1 - F_{NC2}(2cK; \nu + 4, \psi)] \\ &\quad + \bar{\theta}[1 - \exp(-\kappa\tau)]\exp(-r\tau)[1 - F_{NC2}(2cK; \nu + 2, \psi)] \\ &\quad - K\exp(-r\tau)[1 - F_{NC2}(2cK; \nu, \psi)], \end{aligned} \tag{4}$$

where $F_{NC2}(\cdot; \nu, \psi)$ is the cumulative distribution function (cdf) of a non-central chi-square distribution with ν degrees of freedom and non-centrality parameter ψ. Hence, it is straightforward to express the call price as a function of $F(t, T)$ by exploiting the relationship between futures prices and $V(t)$ in (3).

More recently, Detemple and Osakwe (2000) considered the log-normal Ornstein-Uhlenbeck (log-OU) diffusion:

$$d\log V(t) = \kappa[\bar{\theta} - \log V(t)]dt + \sigma dW(t).$$

As is well known, this model implies that $\log V(t)$ would follow a Gaussian AR(1) process if sampled at equally spaced discrete intervals. More generally, the conditional distribution of $\log V(T)$ given $V(t)$ would be Gaussian with mean

$$\mu(t, \tau) = \bar{\theta} + \exp(-\kappa\tau)[\log V(t) - \bar{\theta}]$$

and variance

$$\varphi^2(\tau) = \frac{\sigma^2}{2\kappa}[1 - \exp(-2\kappa\tau)]. \tag{5}$$

As in the SQR process, $\bar{\theta}$ and κ can be interpreted as the long-run mean and mean-reversion parameters, respectively, but now it is the log of $V(t)$ that mean reverts to $\bar{\theta}$. In this context, it is straightforward to show that the futures price is

$$F_{Log-OU}(t, T, V(t), \kappa, \bar{\theta}, \sigma) = \exp(\mu(t, \tau) + 0.5\varphi^2(\tau)),$$

[2] As remarked by Hansen and Scheinkman (2009), the density of this distribution will be positive at $V(T) = 0$ if the Feller condition $2\kappa\bar{\theta} \geq \sigma^2$ is violated.

while the call price can be expressed as

$$
\begin{aligned}
c_{Log-OU}(t,T,K,\kappa,\bar{\theta},\sigma) &= \exp(-r\tau)F(t,T)\Phi\left[\frac{\log(F(t,T)/K)+\frac{1}{2}\varphi^2(\tau)}{\varphi(\tau)}\right] \\
&\quad -K\exp(-r\tau)\Phi\left[\frac{\log(F(t,T)/K)-\frac{1}{2}\varphi^2(\tau)}{\varphi(\tau)}\right] \qquad (6)
\end{aligned}
$$

if we take the futures contract as the underlying instrument, where $\Phi(\cdot)$ denotes the standard normal cdf. This is the well known Black (1976) formula, although in this case the implied volatility $\varphi(\tau)$ in (5) is not constant across maturities. In this sense, the pricing formula proposed by Whaley (1993) based on a geometric Brownian motion can also be expressed as (6) if $\varphi(\tau)$ is taken as a constant irrespective of τ.

4.2 Empirical performance

We have estimated the parameters of these two models over the three sample periods described at the end of section (3). Table 1 reports the in- and out-of-sample root mean square pricing errors (RMSE). Both models yield a similar fit to futures prices but their performance deteriorates substantially in the August 08 - April 09 period, especially when we use the parameters estimated with data prior to August 2008. In contrast, the log-OU model displays smaller RMSE's in the case of options, both in- and out-of-sample. This is confirmed in Figure 5, which compares the daily RMSE's of the two models. As can be seen, both models provide similar daily RMSE's for futures, but the log-OU process is clearly superior in the case of options. Nevertheless, the RMSE's obtained with the log-OU model also tend to deteriorate at the end of the sample.

Table 2 reports the parameter estimates that we obtain. Although the values are not directly comparable because volatility is expressed in levels in the SQR model and in logs in the log-OU model, in both cases the mean reversion parameter κ is quite sensitive to the sample period used for estimation purposes. In contrast, the volatility parameter σ is more stable for the log-OU process, probably because the log transformation is more appropriate to capture the distortions produced by the large movements of the VIX that took place at the end of our sample.

An alternative, more illustrative way to assess the validity of the log-OU model is by computing the implied volatilities with the Black (1976) formula (6). As we mentioned before, the implied volatility of the log-OU process is constant for different degrees of moneyness, but not across different maturities because (5) depends on τ. Hence, if this model were correct then we should obtain constant implied volatilities for a given maturity regardless of moneyness. Figure 6 shows that this is not at all the case. In fact, we generally observe a positive slope. In addition, implied volatilities are much higher in November 2008 than in the other three dates for a similar range of maturities. These empirical results suggest that the volatility of the VIX might not be constant over time, as the SQR and log-OU models assume. In this sense, Figure 6 clearly shows that the log-OU process is not only unable to generate the observed volatility skews, but it also fails to capture the average level of implied volatilities.

5 Extensions

5.1 Model specification

The results of the previous section indicate that a log-OU process offers a better empirical fit than a SQR process, especially for options. Unfortunately, its performance tends to deteriorate during the recent financial crisis. For that reason, in this section we explore several ways to extend the log-OU model. Specifically, we introduce three empirically relevant features: a time varying mean, jumps and stochastic volatility. We consider these extensions first in isolation, and then in combination. All in all, we compare the following cases:

- Central tendency (CT):

$$d \log V(t) = \kappa[\theta(t) - \log V(t)]dt + \sigma dW_v(t),$$

with

$$d\theta(t) = \bar{\kappa}[\bar{\theta} - \theta(t)]dt + \bar{\sigma} dW_\theta(t), \tag{7}$$

where $W_v(t)$ and $W_\theta(t)$ are independent Brownian motions. As we mentioned before, an important deficiency of previously existing models is that they assume that volatility mean reverts at a simple, non-negative exponential rate, which is in fact 0 in the GMB model proposed by Whaley (1993). However, the long, persistent swings in the VIX in Figure 1a suggest that we need to allow for more complex dynamics. In this sense, we show in Appendix A.2 that the exact discretisation of $\log V(t)$ in the above model is a Gaussian ARMA(2,1) process, which is consistent with the evidence reported in Section 2 (see Figure 2). As discussed in Jegadeesh and Pennacchi (1996) and Balduzzi, Das, and Foresi (1998), (7) allows volatility to revert towards a central tendency, which in turn, fluctuates stochastically over time around a long run mean $\bar{\theta}$. As a consequence, the conditional mean of $\log V(T)$ is a function of the distance between $\log V(t)$ and the central tendency $\theta(t)$, as well as the distance between $\theta(t)$ and $\bar{\theta}$. More explicitly,

$$\begin{aligned} E[\log V(T)|I(t)] &= \bar{\theta} + \delta(\tau)[\theta(t) - \bar{\theta}] \\ &\quad + \exp(-\kappa\tau)[\log V(t) - \theta(t)], \end{aligned} \tag{8}$$

where

$$\delta(\tau) = \frac{\kappa}{\kappa - \bar{\kappa}} \exp(-\bar{\kappa}\tau) - \frac{\bar{\kappa}}{\kappa - \bar{\kappa}} \exp(-\kappa\tau).$$

Importantly, this model can also be expressed as the "superposition" (i.e. sum) of two log-OU processes. Therefore, the nested structure in (7) is not restrictive.

- Jumps (J):

$$d\log V(t) = \kappa[\bar{\theta} - \log V(t)]dt + \sigma dW(t) + dZ(t) - \frac{\lambda}{\kappa\delta}dt, \tag{9}$$

where $Z(t)$ is a pure jump process independent of $W(t)$, with intensity λ, and whose jump amplitudes are exponentially distributed with mean $1/\delta$, or $Exp(\delta)$ for short. Note that the last term in (9) simply introduces a constant shift in the distribution of $\log V(t)$ which ensures that $\bar{\theta}$ remains the long-run mean of $\log V(t)$. Jumps in volatility have been previously considered by Duffie, Pan, and Singleton (2000) and Eraker, Johannes, and Polson (2003), among others. Unlike pure diffusions, this model allows for sudden movements in volatility, which nevertheless have lasting effects due to the fact that the mean-reversion parameter κ is bounded.

- Stochastic volatility (SV):

$$d\log V(t) = \kappa[\bar{\theta} - \log V(t)]dt + \gamma\sqrt{\omega(t)}dW(t)$$

where $\omega(t)$ follows an OU-Γ process, which belongs to the class of Lévy OU processes considered by Barndorff-Nielsen and Shephard (2001). Specifically,

$$d\omega(t) = -\bar{\lambda}\omega(t)dt + d\bar{Z}(t) \tag{10}$$

where $\bar{Z}(t)$ is a pure jump process with intensity $\bar{\lambda}$ and $Exp(\bar{\delta})$ jump amplitude, while $W(t)$ is an independent Brownian motion. We use this extension to assess to what extent the price distortions in the previous models are due to the assumption of constant volatility over time. Importantly, the model that we adopt is consistent with the presence of mean reversion in $\omega(t)$, since

$$E\left[\omega(T))|\omega(t)\right] = \bar{\delta}^{-1} + \exp(-\bar{\lambda}\tau)\left[\omega(t) - \bar{\delta}^{-1}\right]. \qquad (11)$$

Hence, $\bar{\delta}^{-1}$ can be interpreted as the long run mean of the instantaneous volatility of the log VIX, while $\bar{\lambda}$ will be the corresponding mean reversion parameter.

- Central tendency and jumps (CTJ):

$$d\log V(t) = \kappa[\theta(t) - \log V(t)]dt + \sigma dW_v(t) + dZ(t) - \frac{\lambda}{\kappa\delta}dt \qquad (12)$$

where $\theta(t)$ follows the diffusion (7) and $Z(t)$ is a pure jump process with intensity λ and $Exp(\delta)$ jump amplitude, while $W_v(t)$ and $W_\theta(t)$ are independent Brownian motions. We again introduce a constant shift in (12) to ensure that $\bar{\theta}$ is the long run mean of both $\theta(t)$ and $\log V(t)$.

- Central tendency and stochastic volatility (CTSV):

$$d\log V(t) = \kappa[\theta(t) - \log V(t)]dt + \gamma\sqrt{\omega(t)}dW_V(t)$$

where $\theta(t)$ follows the diffusion (7) while $\omega(t)$ is defined in (10). As in previous cases, the jump variable $\bar{Z}(t)$ and the Brownian motions are mutually independent.

Except for the CT model, it is not generally possible to price derivative contracts for these extensions in closed form. However, it is possible to obtain the required prices by Fourier inversion of the conditional characteristic function. In particular, we use formula (5) in Carr and Madan (1999) to invert the relevant characteristic functions, which we derive in the Appendix.

Some of the previous extensions introduce as additional factors a time varying tendency (CT, CTJ), a time varying volatility (SV) or both (CTSV),which we need to filter out for estimation purposes. Following standard practice, our filtering strategy in this context yields the values of $\theta(t)$ and $\omega(t)$ that minimise the (weighted) root mean square pricing errors for any choice of parameters. And although the jump variable $Z(t)$ in (9) can also be interpreted as an additional factor, its impact cannot be separately identified from the impact of the diffusion shocks $dW(t)$ because both variables share the same mean reversion coefficient κ.

5.2 Empirical performance

Table 3 reports the in- and out-of-sample RMSE's of the extensions introduced in the previous section. As expected, the CT model is able to yield much smaller RMSE's for futures than the standard log-OU model, both in- and out-of-sample. Therefore, the inclusion of a time-varying mean seems to be crucial for adequately capturing the behaviour of futures prices. At the same time, given that we show in Appendix A.2 that option prices do not depend on $\theta(t)$ once we condition on the current futures price, it is not surprising that the difference in parameter estimates hardly affects the RMSE's of options.

In contrast, the introduction of jumps in the baseline log-OU process slightly improves the in-sample pricing of options, but it does not change much the out-of-sample results. As for futures prices, jumps yield almost no change, which is again to be expected because we can also show that jumps generate identical futures prices as a log-OU model up to first order.[3] Similarly, adding jumps to the CT model does not introduce any remarkable improvements. In this sense, it is important to emphasise that we do not assess the importance of jumps on the historical dynamics of the VIX, only their pricing implications, as we estimate the different models using derivative prices.

But when we consider stochastic volatility, we observe the opposite effect: the improvement over the standard log-OU model is greater for options than for futures. Hence, stochastic volatility turns out to be more important to value options, while central tendency is crucial to price futures. Intuitively, the reason is that the futures formula (1) only involves the conditional mean of $V(T)$, whereas (2) is more sensitive to higher order moments, such as the volatility of volatility. Not surprisingly, we obtain the best results when we combine central tendency and stochastic volatility. In particular, the pricing performance of the CTSV model during the crisis (August 08 - April 09) is remarkably good compared to the other models, even if we only focus on out-of-sample statistics.

Figure 7 compares the daily RMSE's of the CTSV model with the analogous series for the standard log-OU model. Apart from the superior fit, we can clearly see that the performance of the CTSV model is relatively stable over the entire sample. As an alternative comparison, we also compute the empirical cumulative distribution functions

[3] Formally, when we consider a Taylor expansion of the futures price formula of model J around λ/δ for $\lambda/\delta > 0$ but small, we only observe deviations from the standard log-OU expression for the second and higher terms.

of the daily square pricing errors that the different models generate. The smaller the errors, the more shifted to the left those distributions will be. In this respect, Figure 8 shows that the CTSV model dominates all the other models in the first order stochastic sense. The ordering of the remaining models, though, depends on the type of derivative asset considered. In the case of futures, the CTSV is closely followed by CTJ and CT, which provide an almost identical fit. In turn, they are followed by the SV, standard log-OU and SQR models. But for options, the second best model is SV, which is followed by the CTJ, CT, J, standard log-OU and SQR models. Thus, we observe again that a central tendency is relatively more important for futures, while stochastic volatility offers greater gains on options. At the same time, a central tendency does not harm the option pricing performance of the CTSV model, while stochastic volatility does not cause any distortions to futures prices. As for jumps, Figure 8 confirms that they do indeed help in pricing options but they hardly provide any improvement for futures. In any case, compared to stochastic volatility, jumps seem to yield a minor improvement even for options. Lastly, the SQR model undoubtedly offers the worst results, especially for options.

Table 4 reports the parameter estimates that we obtain in the different extensions of the log-OU model introduced in section 5. In the models with central tendency we observe fast mean reversion of $\log V(t)$ to $\theta(t)$, which in turn mean-reverts rather more slowly to its long rung mean $\bar{\theta}$ (i.e. $\kappa \gg \bar{\kappa}$). Jump intensities can vary substantially depending on the sample period considered for estimation. When we include the last part of the sample, we observe the highest values of λ, which are consistent with more than 100 jumps on average per year. On the other hand, the estimates of the mean reversion parameter $\bar{\lambda}$ in the stochastic volatility models tend to be larger when central tendency is not simultaneously included. Since this parameter is also responsible for jump intensity in the OU-Γ model, this result has two interesting implications. First, a smaller value of $\bar{\lambda}$ tends to reduce jump activity in $\bar{Z}(t)$. Specifically, the expected number of jumps per year decreases from 12 in the SV model to 2 in the CTSV extension. Second, the deviations of $\omega(t)$ from its long run mean are more persistent in the CTSV case because mean reversion is slower the smaller $\bar{\lambda}$ is, as (11) indicates. These features are also important from a time series perspective. As mentioned before, central tendency is consistent with ARMA(2,1) dynamics in discrete time, while stochastic volatility introduces GARCH-type persistent variances for the log VIX.

5.3 Evolution of the factors

Figure 9a shows that the filtered values of the central tendency factor $\theta(t)$ are rather insensitive to the particular specification that we use. Similarly, Figure 9b indicates that our preferred CTSV model also generates almost indistinguishable filtered values for $\theta(t)$ across the different estimation subsamples that we consider. Not surprisingly, these figures also confirm that the log of the VIX oscillates around $\theta(t)$, which in turn changes over time rather more slowly. In fact, the main reason for central tendency models to work so well during the recent financial crisis is because they allow for large temporal deviations of $\theta(t)$ from its long term value $\bar{\theta}$, thereby reconciling the large increases of the VIX observed during that period with mean reversion over the long term.

To assess the extent to which the filtered values of $\theta(t)$ are realistic, Figure 9c compares the "actual" futures prices for a constant 30-day maturity, which we obtain by interpolation of the adjacent contracts, with the daily estimates of 30-day futures prices generated by the log-OU, CT and CTSV models. By focusing on a constant maturity, we can not only compare the absolute magnitude of the pricing errors, as in Figure 7a, but also their sign and persistence. As can be seen, the pricing errors of the log-OU process are not only substantially larger than those of the two CT extensions, but they also display much stronger persistence. For instance, the log-OU model systematically underprices futures from October 2008 until the end of the sample. In addition, the pricing errors of the CT and CTSV models are almost identical. This feature shows once again that central tendency is the most relevant extension for pricing futures.

Figure 10a compares the filtered instantaneous volatilities of the SV and CTSV models. Not surprisingly, the scale of $\omega(t)$ is much larger in the SV model because part of the apparent variability in the volatility of the log VIX is captured in the CTSV model by the variability in the central tendency. Nevertheless, the estimated volatility associated to the CTSV model still varies substantially over time. Therefore, the inclusion of a central tendency reduces the need for stochastic volatility but does not eliminate it completely. As it was the case for the central tendency $\theta(t)$ in Figure 9b, Figure 10b confirms that the filtered values of $\omega(t)$ obtained from the CTSV model are also quite insensitive to the sample period used to estimate the model parameters.

Finally, to assess the extent to which the filtered values of $\omega(t)$ make sense, we compute the 30-day ahead standard deviations of $\log V(t)$ implied by the CTSV model. In Figure 10c we compare those standard deviations with the Black (1976) implied volatilities of at-the-money options that are exactly 30-day from expiration, which we again obtain by interpolation. The high correlation between the two series shows that the filtered values of stochastic volatility are indeed related to the changing perceptions of the market about the standard deviation of the VIX.

5.4 Term structures of derivatives and implied volatility skews

So far, we have focused on the overall empirical performance of the different models. In principle, though, our results could change for different time horizons or different degrees of moneyness. For that reason, Table 5 shows the RMSE's of futures contracts for different ranges of maturity. We observe that all models tend to yield larger distortions for longer maturities. Nevertheless, the worst RMSE of the CTSV model, which corresponds to maturities higher than six months, is still much smaller than the best RMSE of either the SQR or the standard log-OU models, which correspond to contracts with less than three months to maturity.

To gain some additional insight, in Figure 11 we look at the term structure of futures prices for the four days considered in Figure 4c. This figure confirms that central tendency is crucial to reproduce the changes in the level and slope of the actual term structure of VIX futures prices. We can also observe that the J model and the standard log-OU process yields prices which are almost identical, while the CTJ model can be barely distinguished from the CT extension.

Tables 6 and 7 provide the RMSE's of calls and puts for different ranges of maturity and moneyness. In this case, we observe the highest price distortions for at-the-money call and put options. For maturities smaller than 3 months, the SQR also seems to suffer large price distortions for deep out-of-the-money calls and deep in-the-money puts. A similar effect is observed for the standard log-OU model, but this pattern is less clear for its generalisations. In particular, stochastic volatility models provide the best fit uniformly across all moneyness and maturity ranges.

In Figure 12 we consider the term structure of call prices for a given strike for the same four days as in Figures 4c and 11. All models manage to capture reasonably well the shape of the term structure of option prices, which was upward sloping on 16-July-2007 and 15-August-2008, downward sloping on 20-November-2008, and hump shaped on 7-April-2009. However, only those models with stochastic volatility are able to reproduce the level of call options market prices. The remaining models either underprice (20-November-2008) or overprice (remaining dates) calls for all maturities.

Finally, in Figures 13 and 14 we assess whether jumps or stochastic volatility are better at capturing the actual implied volatilities skews depicted in Figure 6. For the sake of completeness, we have also plotted the implied volatilities generated by the standard log-OU model and its CT extension. These figures show that while the introduction of jumps can reproduce the positive slope of the actual implied vols, only stochastic volatility seems to correctly capture both their level and slope.

6 Conclusions

We carry out an extensive empirical analysis of VIX derivatives valuation models. We consider daily prices of futures and European options over a sample period that covers the turbulences that took place between August 2007 and August 2008, the worst months of the recent financial crisis (September 2008 to April 2009), as well as the previous bull market years from 2004 to 2007. These markedly different periods provide a very useful testing ground to assess the empirical performance of the different pricing models. We initially focus on the two existing mean-reversion models: the square root (SQR) and the log-normal Ornstein-Uhlenbeck processes (log-OU). Although SQR is more popular in the empirical literature, we find that the log-OU model yields a similar fit for future prices and a better fit for options. However, both models yield large price distortions during the crisis. In addition, they do not seem to capture either the level or the slope of the term structure of futures and option prices, or indeed the volatility skew. Part of the problem is that these models implicitly assume that volatility mean reverts at a simple exponential rate, which cannot accommodate the long and persistent swings of the VIX observed at the end of our sample. In this sense, we show that the simple AR(1) structure that they imply in discrete time is not consistent with the empirical evidence of ARMA dynamics for the VIX.

We investigate the potential sources of mis-pricing by considering several empirically relevant generalisations of the standard log-OU model. In particular, we introduce a time varying central tendency, jumps and stochastic volatility. Our parameter estimates indicate that the VIX rapidly mean-reverts to a central tendency, which in turn reverts more slowly to a long run constant mean. This flexible structure can reconcile the large variations of the VIX over our sample period with mean-reversion to a long run constant value. Except for the central tendency model, though, it is not generally possible to price derivatives in closed form for the extensions that we consider. For that reason, we obtain the required prices by Fourier inversion of the conditional characteristic function.

Interestingly, our results indicate that a time varying central tendency is crucial to price futures. At the same time, it is not detrimental for the valuation of options. We also find evidence of time varying volatility in the VIX. But stochastic volatility plays a much more important role for options while leaving futures prices almost unaffected. In contrast, jumps do not change futures prices (up to first order), and only provide a minor improvement for options. Importantly, our results remain valid when we focus exclusively on the out-of-sample performance. In view of these findings, we conclude that a generalised log-OU model that combines a time varying central tendency with stochastic volatility is needed to obtain a good pricing performance during bull and bear markets, as well as to capture the term structures of VIX futures and options, and the implied volatilities of options.

Given the definition of the VIX index, our analysis also has important implications for models of the S&P 500. Specifically, it implies that stochastic volatility models for this broad stock index should allow for slow mean reversion, and a time-varying volatility of volatility. Amengual (2009) considers such a model for volatility swaps.

We could extend our empirical exercise to other volatility derivatives such as binary options. Another interesting avenue for future research would be to explore the relationship between the actual and risk-neutral VIX measures to obtain the prices of risk. A rather more ambitious generalisation would be to integrate the valuation of VIX derivatives with the valuation of the underlying options on the S&P 500 that are used to compute this volatility index, as proposed by Lin and Chang (2009). Thus, one could formally explore the ability of VIX derivatives to successfully hedge against falls in the value of the S&P 500. Once again, the main lesson of our paper is that it would be necessary to take into account the presence of a time varying central tendency and stochastic volatility in the volatility of the S&P 500.

References

Amengual, D. (2009). The term structure of variance risk premia. mimeo Princeton University.

Andersen, T. G., T. Bollerslev, and F. X. Diebold (2009). Parametric and nonparametric volatility measurement. In Y. Aït-Sahalia and L. P. Hansen (Eds.), *Handbook of Financial Econometrics*, Volume 1, pp. 67–138. North-Holland Press.

Balduzzi, P., S. R. Das, and S. Foresi (1998). The central tendency: a second factor in bond yields. *Review of Economics and Statistics 80*, 62–72.

Barndorff-Nielsen, O. and N. Shephard (2001). Non-Gaussian Ornstein-Uhlenbeck-based models and some of their uses in financial economics. *Journal of the Royal Statistical Society. Series B (Statistical Methodology) 63*, 167–241.

Black, F. (1976). The pricing of commodity contracts. *Journal of Financial Economics 3*, 167–179.

Carr, P. and D. B. Madan (1999). Option valuation using the fast Fourier transform. *Journal of Computational Finance 2*, 61–73.

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

CBOE (2009). The CBOE volatility index - VIX. White Paper (available at www.cboe.com/micro/vix).

Christoffersen, P. and K. Jacobs (2004). The importance of the loss function in option valuation. *Journal of Financial Economics 72*, 291–318.

Cox, J. C., J. E. Ingersoll, and S. A. Ross (1985). A theory of the term structure of interest rates. *Econometrica 53*, 385–407.

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

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

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

Dumas, B., J. Fleming, and E. Whaley (1998). Implied volatility functions: empirical tests. *Journal of Finance 53*, 2059–2106.

Eraker, B., M. Johannes, and N. Polson (2003). The impact of jumps in volatility and returns. *Journal of Finance 53*, 1269–1300.

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

Hansen, L. P. and J. A. Scheinkman (2009). Long-term risk: an operator approach. *Econometrica 77*, 177–234.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies 6*, 327–343.

Jegadeesh, N. and G. G. Pennacchi (1996). The behavior of interest rates implied by the term structure of eurodollar futures. *Journal of Money, Credit and Banking 28*, 426–446.

León, A. and E. Sentana (1997). Pricing options on assets with predictable white noise returns. mimeo CEMFI.

Lin, Y. N. and C. H. Chang (2009). VIX option pricing. *Journal of Futures Markets 29*, 523–543.

Szado, E. (2009). VIX futures and options - A case study of portfolio diversification during the 2008 financial crisis. mimeo Isenberg School of Management.

Trolle, A. B. and E. S. Schwartz (2009). Unspanned stochastic volatility and the pricing of commodity derivatives. *Review of Financial Studies* (Forthcoming).

Wang, Z. and R. T. Daigler (2009). The performance of VIX option pricing models: empirical evidence beyond simulation. mimeo Florida International University.

Whaley, R. E. (1993). Derivatives on market volatility: hedging tools long overdue. *Journal of Derivatives 1*, 71–84.

Zhang, J. E. and Y. Zhu (2006). Vix futures. *Journal of Futures Markets 26*, 521–531.

A Extensions of log-OU processes

A.1 General case

The extensions that we consider belong to the class of affine jump-diffusion state processes analysed by Duffie, Pan, and Singleton (2000). In particular, consider an N-dimensional vector $\mathbf{Y}(t)$ that satisfies the diffusion

$$d\mathbf{Y}(t) = \mathbf{K}(\mathbf{\Theta} - \mathbf{Y}(t))dt + \sqrt{\mathbf{S}(t)}d\mathbf{W}(t) + d\mathbf{Z}(t),$$

where $\mathbf{W}(t)$ is an N-dimensional vector of independent standard Brownian motions, $\mathbf{K}$ is an $N \times N$ matrix, $\mathbf{\Theta}$ is a vector of dimension N, $\mathbf{S}(t)$ is a diagonal matrix of dimension N whose i^{th} diagonal element is $c_{i0} + \mathbf{c}'_{i1}\mathbf{Y}(t)$, and finally, $\mathbf{Z}(t)$ is a multivariate pure jump process with intensity λ whose jump amplitudes have joint density $f_J(\cdot)$.

Duffie, Pan, and Singleton (2000) show that the conditional characteristic function of $\mathbf{Y}(T)$ can be expressed as

$$\begin{aligned} \phi_Y(t,T,\mathbf{u}) &= E[\exp(i\mathbf{u}'\mathbf{Y}(T))|I(t)] \\ &= \exp(\varphi_0(\tau) + \boldsymbol{\varphi}'_Y(\tau)\mathbf{Y}(t)), \end{aligned}$$

where $\varphi_0(\tau)$ and $\boldsymbol{\varphi}_Y(\tau)$ satisfy the following system of differential equations:

$$\dot{\boldsymbol{\varphi}}_Y(\tau) = -\mathbf{K}'\boldsymbol{\varphi}_Y(\tau) + \frac{1}{2}\boldsymbol{\varsigma}_Y(\tau),$$

$$\dot{\varphi}_0(\tau) = \mathbf{\Theta}'\mathbf{K}\boldsymbol{\varphi}_Y(\tau) + \frac{1}{2}\boldsymbol{\varphi}'_Y(\tau)diag(\mathbf{c}_0)\boldsymbol{\varphi}_Y(\tau) + \lambda[J(\boldsymbol{\varphi}_Y(\tau)) - 1],$$

where

$$J(\mathbf{u}) = \int \exp(\mathbf{u}'\mathbf{x})f_J(\mathbf{x})d\mathbf{x},$$

and $\boldsymbol{\varsigma}_Y(\tau)$ is an N-dimensional vector whose k^{th} element is

$$\varsigma_{Y,k}(\tau) = \boldsymbol{\varphi}'_Y(\tau)diag(\mathbf{c}_{k1})\boldsymbol{\varphi}_Y(\tau).$$

If we place $\log V(t)$ as the first element in of $\mathbf{Y}(t)$, then the conditional characteristic function of $\log V(t)$ will be

$$\phi(t,T,u) = \phi_Y(t,T,\mathbf{u}_0),$$

where $\mathbf{u}_0 = (u, 0, \cdots, 0)'$. Following Carr and Madan (1999), the price of a call option with strike K can then be expressed as

$$c(t,T,K) = \frac{\exp(-\alpha\log(K))}{\pi}\int_0^\infty \exp(-iu\log(K))\psi(u)du, \qquad \text{(A1)}$$

where

$$\psi(u) = \frac{\exp(-r\tau)\phi(t,T,u-(1+\alpha)i)}{\alpha^2+\alpha-u^2+i(1+2\alpha)u}$$

and α is a smoothing parameter. We evaluate (A1) by numerical integration. In our experience, $\alpha = 1.1$ yields good results.

Given that the estimation algorithm requires the evaluation of the objective function at many different parameter values, we linearise option prices with respect to $\omega(t)$ in the models with stochastic volatility to speed up the calculations. Our procedure is analogous to the treatment of other stochastic volatility models, which are sometimes linearised to employ the Kalman filter (see e.g. Trolle and Schwartz, 2009). Specifically, we linearise call prices for day t around the volatility of the previous day as follows:

$$\begin{aligned} c_{SV}(t,T,K,\omega(t)) \;\approx\; & c_{SV}\left(t,T,K,\omega\left(t-\frac{1}{360}\right)\right) \\ & + \left.\frac{\partial c_{SV}(t,T,K,x)}{\partial x}\right|_{x=\omega\left(t-\frac{1}{360}\right)} \left[\omega\left(t\right)-\omega\left(t-\frac{1}{360}\right)\right]. \end{aligned}$$

Due to the high persistence of $\omega(t)$, its previous day value turns out to be a very good predictor, which reduces the approximation error of the above expansion. In fact, the linearisation error, expressed in terms of the RMSE's of options, is very small: around 0.57% when we use data up to July 2007, and below 0.25% in the remaining samples. In any case, we calculate the exact pricing errors once we have obtained the final parameter estimates.

A.2 Central tendency

Following León and Sentana (1997), it can be shown that the conditional distribution of $\log V(T)$ given information up to time t is Gaussian with mean $\mu_{CT}(t,\tau)$ given in (8) and variance

$$\begin{aligned} \varphi^2_{CT}(\tau) \;=\; & \frac{\sigma^2}{2\kappa}[1-\exp(-2\kappa\tau)] \\ & +\bar{\sigma}^2\left(\frac{\kappa}{\kappa-\bar{\kappa}}\right)^2\left[\begin{array}{c}\frac{1-\exp(-2\bar{\kappa}\tau)}{2\bar{\kappa}}+\frac{1-\exp(-2\kappa\tau)}{2\kappa} \\ -2\frac{1-\exp(-(\kappa+\bar{\kappa})\tau)}{\kappa+\bar{\kappa}}\end{array}\right]. \end{aligned}$$

By exploiting log-normality, we can write futures prices as

$$F_{CT}(t,T,V(t),\kappa,\theta,\sigma) = \exp(\mu_{CT}(t,\tau)+0.5\varphi^2_{CT}(\tau)),$$

while call prices follow the usual Black (1976) formula with volatility $\varphi_{CT}(\tau)$.

In terms of time series dynamics, it can be shown that $\theta(t)$ and $\log V(t)$ jointly follow a Gaussian VAR(1) if sampled at equally spaced intervals. Specifically,

$$\begin{pmatrix} \theta(T) \\ \log V(T) \end{pmatrix} = \mathbf{g}_\tau + \mathbf{F}_\tau \begin{pmatrix} \theta(t) \\ \log V(t) \end{pmatrix} + \boldsymbol{\varepsilon}_\tau,$$

where

$$\mathbf{g}_\tau = \begin{bmatrix} 1-\exp(-\bar{\kappa}\tau) \\ 1-\exp(-\kappa\tau) - \frac{\kappa}{\kappa-\bar{\kappa}}\left(\exp(-\bar{\kappa}\tau)-\exp(-\kappa\tau)\right) \end{bmatrix} \bar{\theta},$$

$$\mathbf{F}_\tau = \begin{bmatrix} \exp(-\bar{\kappa}\tau) & 0 \\ \frac{\kappa}{\kappa-\bar{\kappa}}[\exp(-\bar{\kappa}\tau)-\exp(-\kappa\tau)] & \exp(-\kappa\tau) \end{bmatrix}.$$

and $\boldsymbol{\varepsilon}_\tau \sim iid\ N(\mathbf{0}, \boldsymbol{\Sigma}_\tau)$, where $\boldsymbol{\Sigma}_\tau$ is a symmetric 2×2 matrix with elements

$$\boldsymbol{\Sigma}_\tau(1,1) = \frac{\bar{\sigma}^2}{2\bar{\kappa}}[1-\exp(-2\bar{\kappa}\tau)]$$

$$\boldsymbol{\Sigma}_\tau(1,2) = \frac{\kappa\bar{\sigma}^2}{\kappa-\bar{\kappa}}\left[\frac{1-\exp(-2\bar{\kappa}\tau)}{2\bar{\kappa}} - \frac{1-\exp(-(\kappa+\bar{\kappa})\tau)}{\kappa+\bar{\kappa}}\right]$$

and $\boldsymbol{\Sigma}_\tau(2,2) = \varphi^2_{CT}(\tau)$. From here, it is straightforward to obtain the marginal process followed by $\log V(t)$, which corresponds to the following ARMA(2,1) model:

$$\log V(t) = h_0(\tau) + h_1(\tau)\log V(t-\tau) + h_2(\tau)\log V(t-2\tau) + u(t) + g(\tau)u(t-\tau)$$

where $u(t), u(t-\tau), \cdots \sim iid\ N(0, p^2(\tau))$ and

$$\begin{aligned}
h_0(\tau) &= \bar{\theta}\left(1-h_1(\tau)-h_2(\tau)\right), \\
h_1(\tau) &= \mathbf{F}_\tau(1,1)+\mathbf{F}_\tau(2,2), \\
h_2(\tau) &= -\mathbf{F}_\tau(1,1)\mathbf{F}_\tau(2,2), \\
g(\tau)p^2(\tau) &= \begin{pmatrix} \mathbf{F}_\tau(2,1) & -\mathbf{F}_\tau(1,1) \end{pmatrix} \boldsymbol{\Sigma}_\tau \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \\
(1+g^2(\tau))p^2(\tau) &= \begin{pmatrix} \mathbf{F}_\tau(2,1) & -\mathbf{F}_\tau(1,1) \end{pmatrix} \boldsymbol{\Sigma}_\tau \begin{pmatrix} \mathbf{F}_\tau(2,1) \\ -\mathbf{F}_\tau(1,1) \end{pmatrix} + \boldsymbol{\Sigma}_\tau(2,2).
\end{aligned}$$

A.3 Jumps

The conditional characteristic function reduces to

$$\phi_J(t,T,u) = \exp\left[\varphi_{J,0}(\tau) + \varphi_{J,V}(\tau)\log V(t)\right],$$

where

$$\begin{aligned}
\varphi_{J,0}(\tau) &= iu\left(\bar{\theta} - \frac{\lambda}{\kappa\delta}\right)[1-\exp(-\kappa\tau)] - \frac{\sigma^2 u^2}{4\kappa}[1-\exp(-2\kappa\tau)] \\
&\quad + \frac{\lambda}{\kappa}\log\left[\frac{\delta - iu\exp(-\kappa\tau)}{\delta - iu}\right]
\end{aligned}$$

and

$$\varphi_{J,V}(\tau) = iu\exp(-\kappa\tau).$$

A.4 Stochastic volatility

The conditional characteristic function simplifies to

$$\phi_{SV}(t,T,u) = \exp\left[\varphi_{SV,0}(\tau) + \varphi_{SV,V}(\tau)\log V(t) + \varphi_{SV,\omega}(\tau)\omega(t)\right],$$

where

$$\begin{aligned}\varphi_{SV,V}(\tau) &= iu\exp(-\kappa\tau),\\ \varphi_{SV,\omega}(\tau) &= \frac{\gamma^2u^2}{2(2\kappa-\lambda)}\left[\exp(-2\kappa\tau)-\exp(-\lambda\tau)\right],\end{aligned}$$

and

$$\varphi_{SV,0}(\tau) = i\bar{\theta}u[1-\exp(-\kappa\tau)] + \lambda\left[\varkappa(\tau,u)-\tau\right],$$

with

$$\varkappa(\tau,u) = \int_0^\tau \frac{\delta}{\delta - \frac{\gamma^2u^2}{2(2\kappa-\lambda)}\left[\exp(-2\kappa x)-\exp(-\lambda x)\right]}dx. \tag{A2}$$

A.5 Central tendency and jumps

The conditional characteristic function becomes

$$\phi_{CTJ}(t,T,u) = \exp\left[\varphi_{CTJ,0}(\tau) + \varphi_{CTJ,\theta}(\tau)\theta(t) + \varphi_{CTJ,V}(\tau)\log V(t)\right],$$

where

$$\begin{aligned}\varphi_{CTJ,0}(\tau) &= iu\frac{\kappa\bar{\kappa}}{\kappa-\bar{\kappa}}\bar{\theta}\left[\frac{1-\exp(-\bar{\kappa}\tau)}{\bar{\kappa}} - \frac{1-\exp(-\kappa\tau)}{\kappa}\right]\\ &\quad -\frac{\sigma^2u^2}{4\kappa}[1-\exp(-2\kappa\tau)]\\ &\quad -\frac{\bar{\sigma}^2u^2}{2}\left(\frac{\kappa}{\kappa-\bar{\kappa}}\right)^2\left[\begin{array}{c}\frac{1-\exp(-2\bar{\kappa}\tau)}{2\bar{\kappa}} + \frac{1-\exp(-2\kappa\tau)}{2\kappa}\\ -2\frac{1-\exp(-(\kappa+\bar{\kappa})\tau)}{\kappa+\bar{\kappa}}\end{array}\right]\\ &\quad +\frac{\lambda}{\kappa}\log\left[\frac{\delta-iu\exp(-\kappa\tau)}{\delta-iu}\right]\\ &\quad -iu\frac{\lambda}{(\kappa-\bar{\kappa})\delta}\left[\exp(-\bar{\kappa}\tau)-\exp(-\kappa\tau)\right],\end{aligned}$$

$$\varphi_{CTJ,\theta}(\tau) = iu\frac{\kappa}{\kappa-\bar{\kappa}}\left[\exp(-\bar{\kappa}\tau)-\exp(-\kappa\tau)\right],$$

and

$$\varphi_{CTJ,V}(\tau) = iu\exp(-\kappa\tau).$$

A.6 Central tendency and stochastic volatility

The conditional characteristic function is

$$\phi_{CTSV}(t,T,u) = \exp\left[\begin{array}{c} \varphi_{CTSV,0}(\tau) + \varphi_{CTSV,V}(\tau)\log V(t) \\ +\varphi_{CTSV,\omega}(\tau)\omega(t) + \varphi_{CTSV,\theta}(\tau)\theta(t) \end{array}\right],$$

where

$$\begin{aligned}
\varphi_{CTSV,V}(\tau) &= iu\exp(-\kappa\tau), \\
\varphi_{CTSV,\omega}(\tau) &= \frac{\gamma^2 u^2}{2(2\kappa-\lambda)}\left[\exp(-2\kappa\tau) - \exp(-\lambda\tau)\right], \\
\varphi_{CTSV,\theta}(\tau) &= iu\frac{\kappa}{\kappa-\bar{\kappa}}\left[\exp(-\bar{\kappa}\tau) - \exp(-\kappa\tau)\right],
\end{aligned}$$

and

$$\begin{aligned}
\varphi_{CTSV,0}(\tau) &= iu\bar{\kappa}\bar{\theta}\frac{\kappa}{\kappa-\bar{\kappa}}\left[\frac{1-\exp(-\bar{\kappa}\tau)}{\bar{\kappa}} - \frac{1-\exp(-\kappa\tau)}{\kappa}\right] \\
&\quad -\frac{1}{2}\bar{\sigma}^2 u^2\left(\frac{\kappa}{\kappa-\bar{\kappa}}\right)^2\left[\begin{array}{c} \frac{1-\exp(-2\bar{\kappa}\tau)}{2\bar{\kappa}} + \frac{1-\exp(-2\kappa\tau)}{2\kappa} \\ -2\frac{1-\exp(-(\kappa+\bar{\kappa})\tau)}{\kappa+\bar{\kappa}} \end{array}\right] \\
&\quad +\lambda\left[\varkappa(\tau,u) - \tau\right],
\end{aligned}$$

with $\varkappa(\tau,u)$ defined in (A2).

Table 1

Root mean square pricing errors in standard one-factor models

(a) Futures

Model	Estimation	Validation			
		May 04-Jul 07	Jul 07-Aug 08	Aug 08-Apr 09	Full sample
SQR	Full sample	1.985	2.635	8.603	4.144
	Until Aug 08	1.789	1.890	*6.337*	3.132
	Until Jul 07	1.357	*4.869*	*11.938*	5.758
Log-OU	Full sample	2.167	1.877	8.461	4.013
	Until Aug 08	1.681	1.847	*6.958*	3.320
	Until Jul 07	1.358	*5.191*	*14.780*	6.891

(a) Options

Model	Estimation	Validation			
		May 04-Jul 07	Jul 07-Aug 08	Aug 08-Apr 09	Full sample
SQR	Full sample	0.755	0.572	0.992	0.782
	Until Aug 08	0.435	0.433	*1.225*	0.785
	Until Jul 07	0.294	*0.395*	*1.371*	0.836
Log-OU	Full sample	0.263	0.440	0.634	0.475
	Until Aug 08	0.261	0.428	*0.672*	0.488
	Until Jul 07	0.252	*0.369*	*0.630*	0.447

Notes: The root mean square errors that correspond to out-of-sample periods are reported in italics. "SQR" denotes square root model while "log-OU" refers to the log-normal Ornstein-Uhlenbeck process.

Table 2
Parameters estimates of the standard one-factor models

		Drift		Volatility
Model	Estimation	κ	θ	σ
SQR	Full sample	2.676	19.503	4.920
	Until Aug 08	1.590	20.688	3.921
	Until Jul 07	3.354	16.186	3.879
Log-OU	Full sample	2.306	2.917	0.953
	Until Aug 08	1.685	2.879	0.871
	Until Jul 07	3.540	2.693	1.009

Notes: "SQR" denotes square root model while "log-OU" refers to the log-normal Ornstein-Uhlenbeck process.

Table 3
Root mean square pricing errors in the extended log-OU models

(a) Futures

Extension	Estimation	Validation			
		May 04-Jul 07	Jul 07-Aug 08	Aug 08-Apr 09	Full sample
CT	Full sample	0.534	0.444	1.278	0.704
	Until Aug 08	0.489	0.429	*1.266*	0.680
	Until Jul 07	0.447	*0.663*	*2.344*	1.097
J	Full sample	2.164	1.877	8.449	4.008
	Until Aug 08	1.678	1.841	*6.916*	3.302
	Until Jul 07	1.355	*5.146*	*14.645*	6.830
SV	Full sample	1.921	2.032	8.500	3.985
	Until Aug 08	1.246	1.873	*6.135*	2.913
	Until Jul 07	0.672	*6.412*	*18.328*	8.466
CTJ	Full sample	0.518	0.437	1.260	0.691
	Until Aug 08	0.469	0.430	*1.303*	0.685
	Until Jul 07	0.450	*1.046*	*3.712*	1.691
CTSV	Full sample	0.445	0.419	1.295	0.672
	Until Aug 08	0.424	0.413	*1.265*	0.654
	Until Jul 07	0.404	*0.451*	*1.353*	0.685

(b) Options

Extension	Estimation	Validation			
		May 04-Jul 07	Jul 07-Aug 08	Aug 08-Apr 09	Full sample
CT	Full sample	0.235	0.351	0.621	0.435
	Until Aug 08	0.227	0.334	*0.654*	0.445
	Until Jul 07	0.226	*0.342*	*0.650*	0.444
J	Full sample	0.227	0.401	0.617	0.448
	Until Aug 08	0.215	0.386	*0.664*	0.463
	Until Jul 07	0.204	*0.333*	*0.669*	0.448
SV	Full sample	0.225	0.274	0.377	0.300
	Until Aug 08	0.226	0.257	*0.436*	0.319
	Until Jul 07	0.203	*0.479*	*0.612*	0.472
CTJ	Full sample	0.211	0.322	0.601	0.413
	Until Aug 08	0.186	0.298	*0.664*	0.434
	Until Jul 07	0.164	*0.345*	*0.773*	0.498
CTSV	Full sample	0.166	0.224	0.317	0.245
	Until Aug 08	0.168	0.221	*0.324*	0.247
	Until Jul 07	0.165	*0.270*	*0.347*	0.274

Notes: The root mean square errors that correspond to out-of-sample periods are reported in italics. “CT” indicates time varying Central Tendency. “J” refers to jumps. “SV” denotes Stochastic Volatility. “CTJ” combines CT and J. “CTSV” combines CT and SV.

Table 4

Parameters estimates of the extended log-OU models

		Drift			Jumps		Volatility				
Extension	Estimation	κ	$\bar{\kappa}$	$\bar{\theta}$	λ	δ	σ	$\bar{\sigma}$	γ	$\bar{\lambda}$	$\bar{\delta}$
CT	Full sample	12.748	0.671	2.913	-	-	1.409	0.494	-	-	-
	Until Aug 08	11.813	0.588	2.886	-	-	1.316	0.498	-	-	-
	Until Jul 07	7.744	0.845	2.801	-	-	1.140	0.561	-	-	-
J	Full sample	2.301	-	2.922	123.140	17.245	0.042	-	-	-	-
	Until Aug 08	1.671	-	2.888	68.647	14.242	0.000	-	-	-	-
	Until Jul 07	3.480	-	2.698	6.956	4.918	0.537	-	-	-	-
SV	Full sample	2.272	-	2.922	-	-	-	-	0.703	11.759	0.692
	Until Aug 08	1.256	-	2.880	-	-	-	-	0.840	19.553	1.499
	Until Jul 07	5.985	-	2.640	-	-	-	-	1.961	0.465	0.700
CTJ	Full sample	11.908	0.647	2.920	102.326	11.939	0.542	0.478	-	-	-
	Until Aug 08	10.134	0.575	2.870	55.674	9.984	0.638	0.511	-	-	-
	Until Jul 07	6.194	1.186	2.733	12.465	5.772	0.620	0.698	-	-	-
CTSV	Full sample	12.007	0.667	2.862	-	-	-	0.204	2.039	2.110	0.832
	Until Aug 08	12.533	0.603	2.846	-	-	-	0.172	1.816	2.082	0.627
	Until Jul 07	12.480	0.657	2.838	-	-	-	0.039	1.993	3.565	0.807

Notes: "CT" indicates Central Tendency, a log-OU process where the long run mean follows another log-OU diffusion. "J" introduces jumps in the standard log-OU model, whose size follows an exponential distribution. "SV" denotes Stochastic Volatility, a log-OU model where the volatility is a Gamma OU Lévy process. "CTJ" combines central tendency and jumps. "CTSV" combines central tendency and stochastic volatility.

Table 5

Root mean square pricing errors of futures prices by maturity

			Log-OU models					
Maturity	N	SQR	Standard	CT	J	SV	CTJ	CTSV
Less than 1 month	985	1.844	1.742	0.541	1.741	1.736	0.540	0.491
From 1 to 3 months	1996	3.281	3.186	0.466	3.184	3.114	0.465	0.448
From 3 to 6 months	2278	4.450	4.320	0.703	4.313	4.292	0.712	0.676
More than 6 months	1899	5.281	5.107	0.946	5.101	5.100	0.901	0.904
Total	7158	4.144	4.013	0.704	4.008	3.985	0.691	0.672

Notes: "SQR" denotes square root model and "log-OU" refers to a log-normal Ornstein-Uhlenbeck process. "CT" indicates Central Tendency, a log-OU process where the long run mean follows another log-OU diffusion. "J" introduces jumps in the standard log-OU model, whose size follows an exponential distribution. "SV" denotes Stochastic Volatility, a log-OU model where the volatility is a Gamma OU Lévy process. "CTJ" combines central tendency and jumps. "CTSV" combines central tendency and stochastic volatility. The column labeled N gives the number of prices per category. Results are based on full sample estimates.

Table 6

Root mean square pricing errors of call prices by moneyness and maturities

(a) $\tau < 1$ month

			Log-OU models					
Moneyness	N	SQR	Standard	CT	J	SV	CTJ	CTSV
< -0.3	1560	0.183	0.148	0.144	0.169	0.128	0.147	0.131
$[-0.3, -0.1)$	1302	0.610	0.395	0.399	0.475	0.234	0.382	0.208
$[-0.1, 0.1)$	1165	1.021	0.620	0.569	0.630	0.325	0.549	0.196
$[0.1, 0.3)$	1192	0.827	0.567	0.423	0.461	0.302	0.398	0.155
≥ 0.3	1255	0.418	0.354	0.288	0.267	0.200	0.218	0.105

(b) 1 month$< \tau < 3$ months

			Log-OU models					
Moneyness	N	SQR	Standard	CT	J	SV	CTJ	CTSV
< -0.3	2227	0.375	0.232	0.202	0.201	0.211	0.188	0.171
$[-0.3, -0.1)$	2457	0.785	0.517	0.448	0.458	0.337	0.421	0.224
$[-0.1, 0.1)$	2429	0.913	0.565	0.523	0.554	0.321	0.523	0.217
$[0.1, 0.3)$	2390	0.831	0.480	0.499	0.476	0.287	0.467	0.208
≥ 0.3	2782	0.565	0.359	0.404	0.313	0.203	0.330	0.214

(c) $\tau > 3$ months

			Log-OU models					
Moneyness	N	SQR	Standard	CT	J	SV	CTJ	CTSV
< -0.3	1901	0.947	0.338	0.246	0.262	0.220	0.224	0.233
$[-0.3, -0.1)$	2161	1.056	0.582	0.404	0.498	0.335	0.374	0.265
$[-0.1, 0.1)$	2372	0.926	0.512	0.399	0.486	0.332	0.401	0.307
$[0.1, 0.3)$	1973	0.743	0.475	0.482	0.493	0.384	0.468	0.335
≥ 0.3	2074	0.597	0.343	0.466	0.337	0.301	0.411	0.368

(d) All maturities

			Log-OU models					
Moneyness	N	SQR	Standard	CT	J	SV	CTJ	CTSV
< -0.3	5688	0.603	0.255	0.205	0.216	0.195	0.191	0.185
$[-0.3, -0.1)$	5920	0.863	0.518	0.422	0.477	0.317	0.396	0.237
$[-0.1, 0.1)$	5966	0.940	0.556	0.488	0.544	0.326	0.484	0.254
$[0.1, 0.3)$	5555	0.800	0.498	0.478	0.479	0.328	0.453	0.253
≥ 0.3	6111	0.550	0.352	0.406	0.313	0.240	0.342	0.263

Notes: Moneyness is defined as $\log(K/F(t,T))$, where K and $F(t,T)$ are the strike and futures prices, respectively. "SQR" denotes square root model and "log-OU" refers to a log-normal Ornstein-Uhlenbeck process. "CT" indicates Central Tendency, a log-OU process where the long run mean follows another log-OU diffusion. "J" introduces jumps in the standard log-OU model, whose size follows an exponential distribution. "SV" denotes Stochastic Volatility, a log-OU model where the volatility is a Gamma OU Lévy process. "CTJ" combines central tendency and jumps. "CTSV" combines central tendency and stochastic volatility. The column labeled N gives the number of prices per category. τ denotes time to maturity. Results are based on full sample estimates.

Table 7

Root mean square pricing errors of put prices by moneyness and maturities

(a) $\tau < 1$ month

			Log-OU models					
Moneyness	N	SQR	Standard	CT	J	SV	CTJ	CTSV
< -0.3	348	0.290	0.229	0.194	0.296	0.156	0.228	0.158
$[-0.3, -0.1)$	1030	0.660	0.436	0.438	0.530	0.257	0.423	0.221
$[-0.1, 0.1)$	1143	1.024	0.620	0.575	0.628	0.332	0.552	0.209
$[0.1, 0.3)$	837	0.951	0.650	0.494	0.537	0.340	0.466	0.175
≥ 0.3	602	0.440	0.368	0.305	0.288	0.222	0.253	0.141

(b) 1 month$< \tau < 3$ months

			Log-OU models					
Moneyness	N	SQR	Standard	CT	J	SV	CTJ	CTSV
< -0.3	1125	0.387	0.281	0.254	0.280	0.250	0.263	0.175
$[-0.3, -0.1)$	2415	0.780	0.539	0.469	0.481	0.366	0.448	0.230
$[-0.1, 0.1)$	1908	0.944	0.597	0.564	0.585	0.352	0.563	0.223
$[0.1, 0.3)$	1079	0.963	0.568	0.598	0.551	0.340	0.564	0.226
≥ 0.3	784	0.582	0.378	0.420	0.325	0.245	0.349	0.207

(c) $\tau > 3$ months

			Log-OU models					
Moneyness	N	SQR	Standard	CT	J	SV	CTJ	CTSV
< -0.3	974	0.705	0.346	0.295	0.300	0.241	0.298	0.242
$[-0.3, -0.1)$	1536	0.968	0.607	0.449	0.523	0.350	0.435	0.305
$[-0.1, 0.1)$	1153	0.887	0.568	0.502	0.541	0.353	0.512	0.369
$[0.1, 0.3)$	606	0.836	0.526	0.617	0.536	0.431	0.600	0.390
≥ 0.3	500	0.482	0.325	0.408	0.342	0.288	0.366	0.423

(d) All maturities

			Log-OU models					
Moneyness	N	SQR	Standard	CT	J	SV	CTJ	CTSV
< -0.3	2447	0.528	0.302	0.264	0.290	0.235	0.273	0.202
$[-0.3, -0.1)$	4981	0.821	0.542	0.456	0.505	0.341	0.439	0.254
$[-0.1, 0.1)$	4204	0.951	0.595	0.550	0.585	0.347	0.547	0.268
$[0.1, 0.3)$	2522	0.930	0.587	0.571	0.543	0.364	0.543	0.262
≥ 0.3	1886	0.514	0.361	0.383	0.319	0.250	0.326	0.268

Notes: Moneyness is defined as $\log(K/F(t,T))$, where K and $F(t,T)$ are the strike and futures prices, respectively. "SQR" denotes square root model and "log-OU" refers to a log-normal Ornstein-Uhlenbeck process. "CT" indicates Central Tendency, a log-OU process where the long run mean follows another log-OU diffusion. "J" introduces jumps in the standard log-OU model, whose size follows an exponential distribution. "SV" denotes Stochastic Volatility, a log-OU model where the volatility is a Gamma OU Lévy process. "CTJ" combines central tendency and jumps. "CTSV" combines central tendency and stochastic volatility. The column labeled N gives the number of prices per category. τ denotes time to maturity. Results are based on full sample estimates.

Figure 1: Historical evolution of the VIX index

(a) 1990-2009

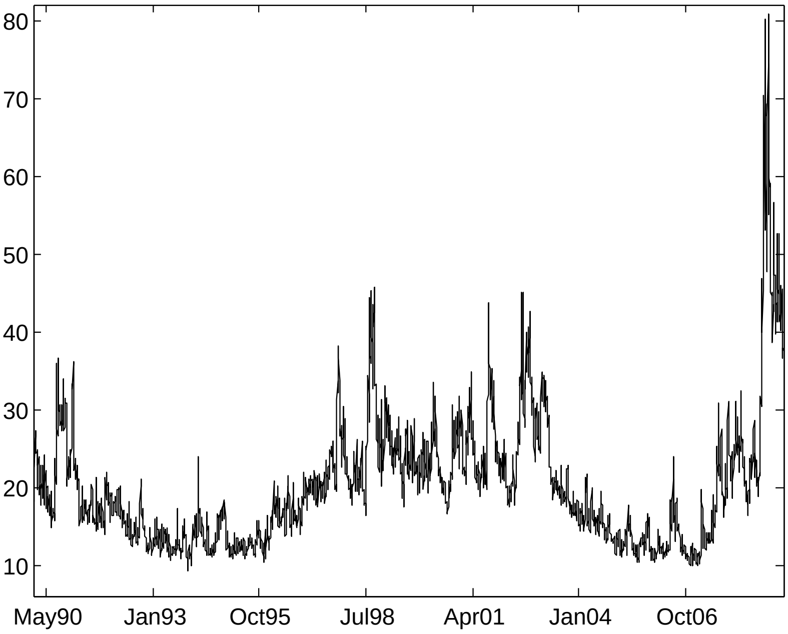

(b) 2004-2009

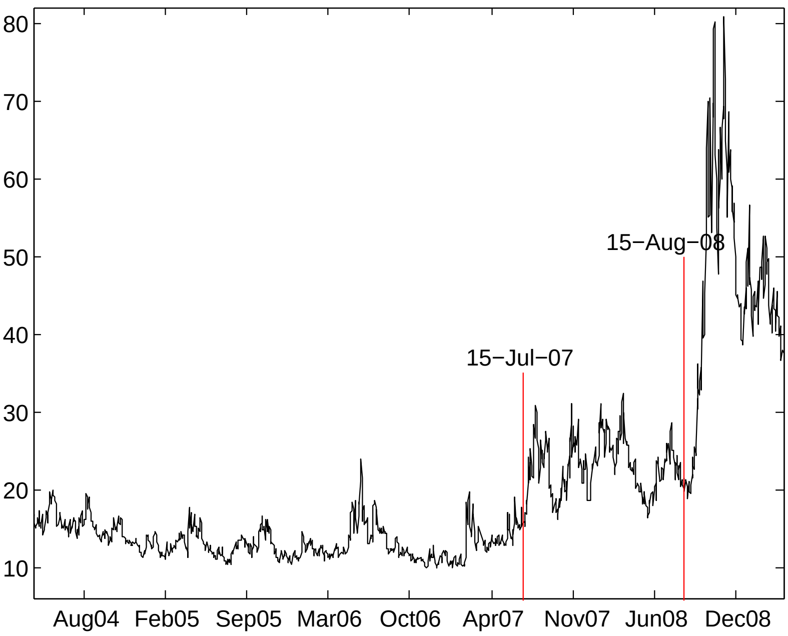

Figure 2: Time series autocorrelations of the log-VIX and estimated ARMA models

(a) Autocorrelations

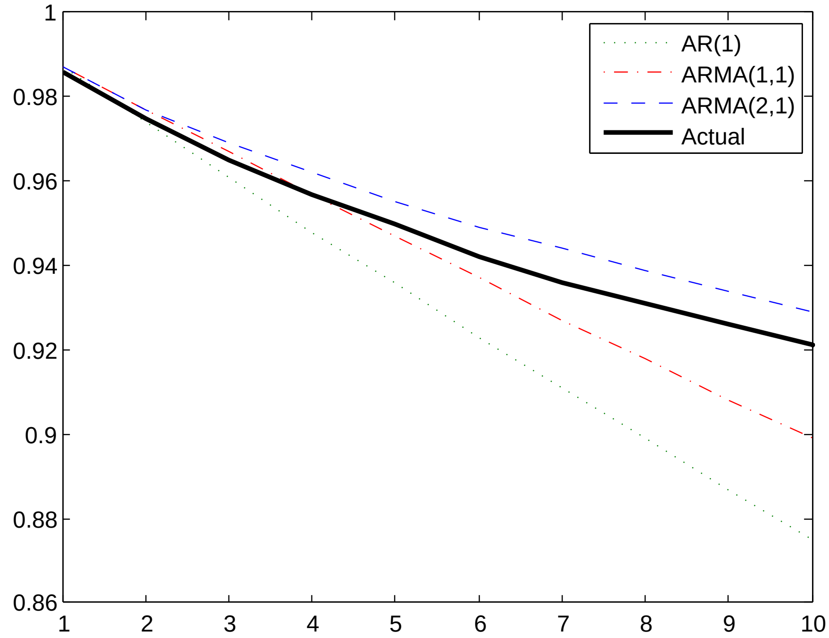

(b) Partial autocorrelations

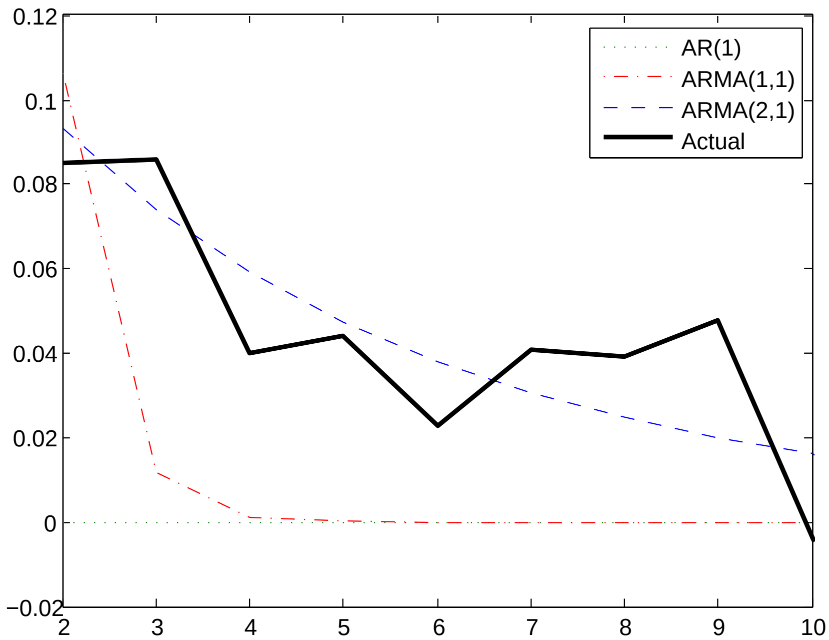

Note: Results are based on the 1990-2009 sample (4859 daily observations).

Figure 3a: Number of futures prices per day

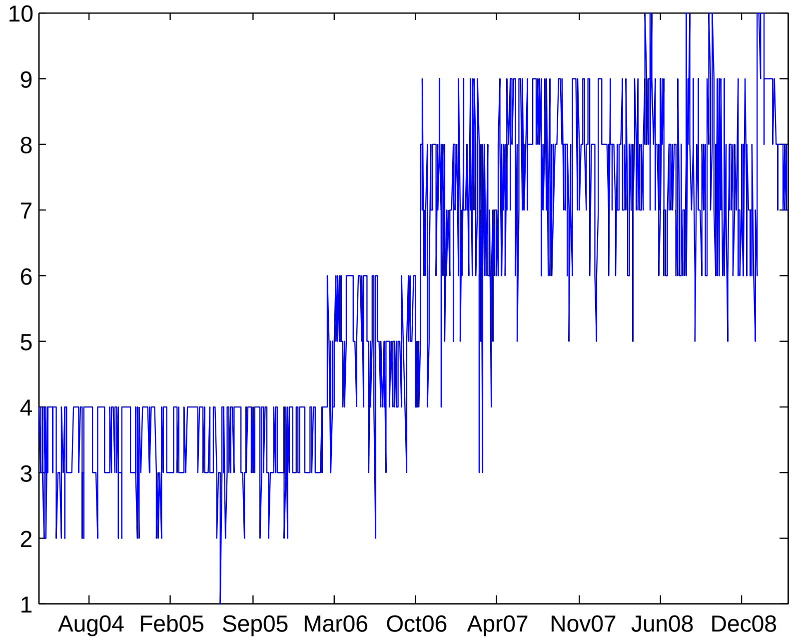

Figure 3b: Number of option prices per day

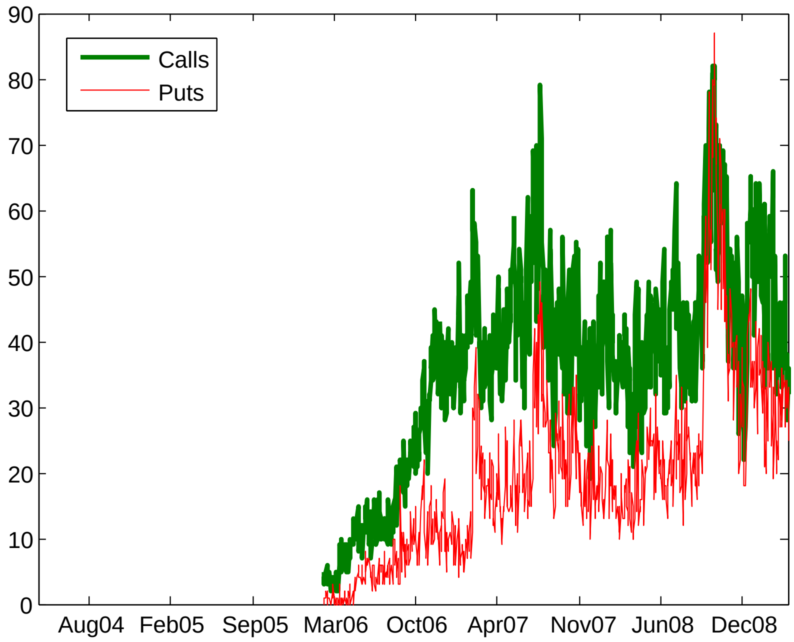

Figure 4a: Term structure of VIX futures

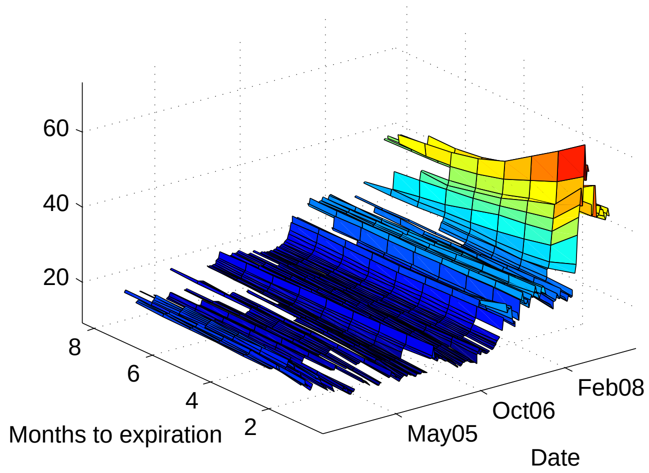

Figure 4b: VIX futures prices vs. VIX index

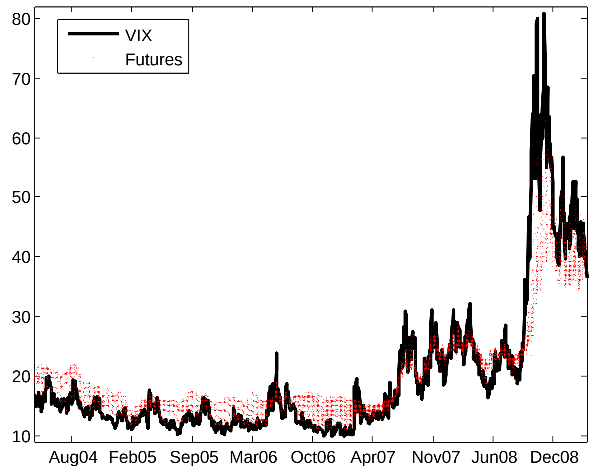

Figure 4c: Term structure of VIX futures at four particular days

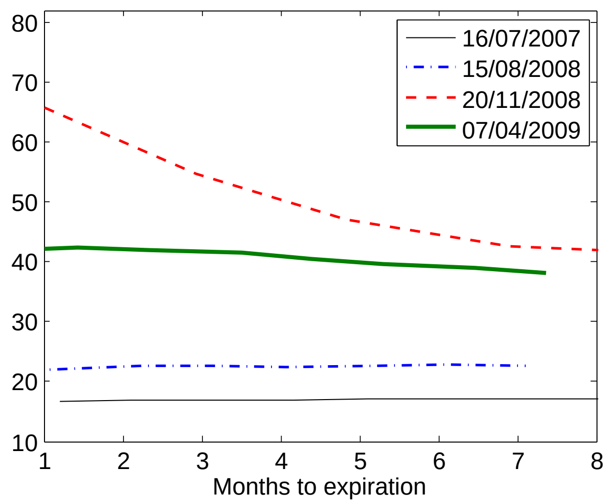

Figure 5: Daily root mean square errors of standard one-factor pricing models

(a) Futures

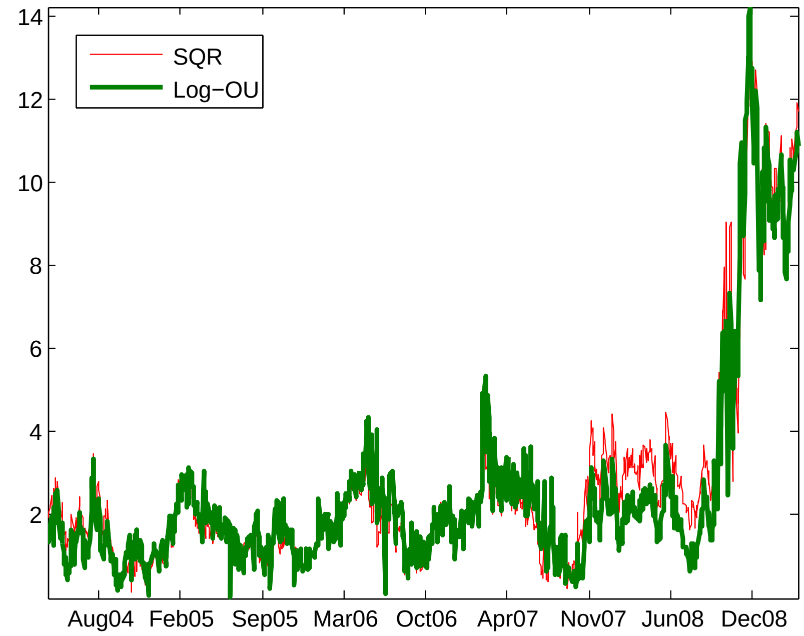

(b) Options

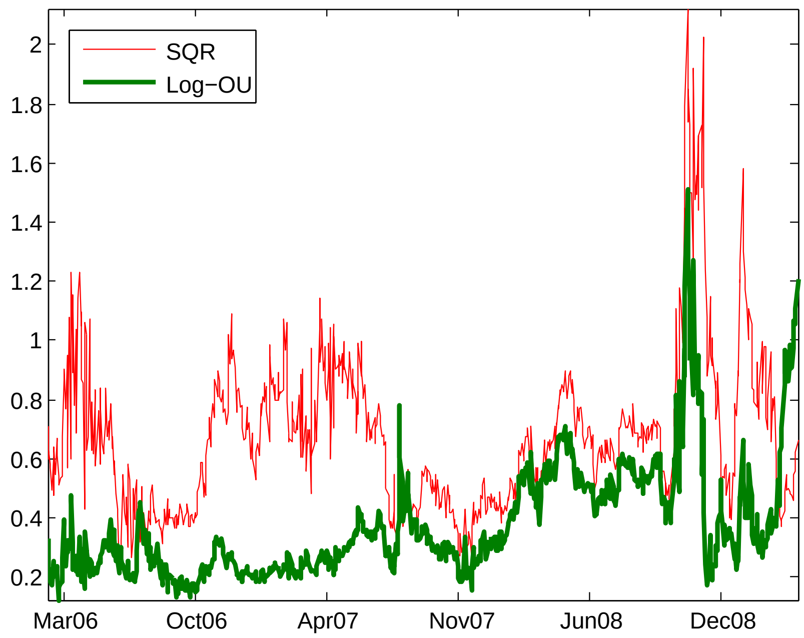

Note: "SQR" denotes square root model and "log-OU" refers to a log-normal Ornstein-Uhlenbeck process. Results are based on full sample estimates.

Figure 6: Implied volatilities of call prices

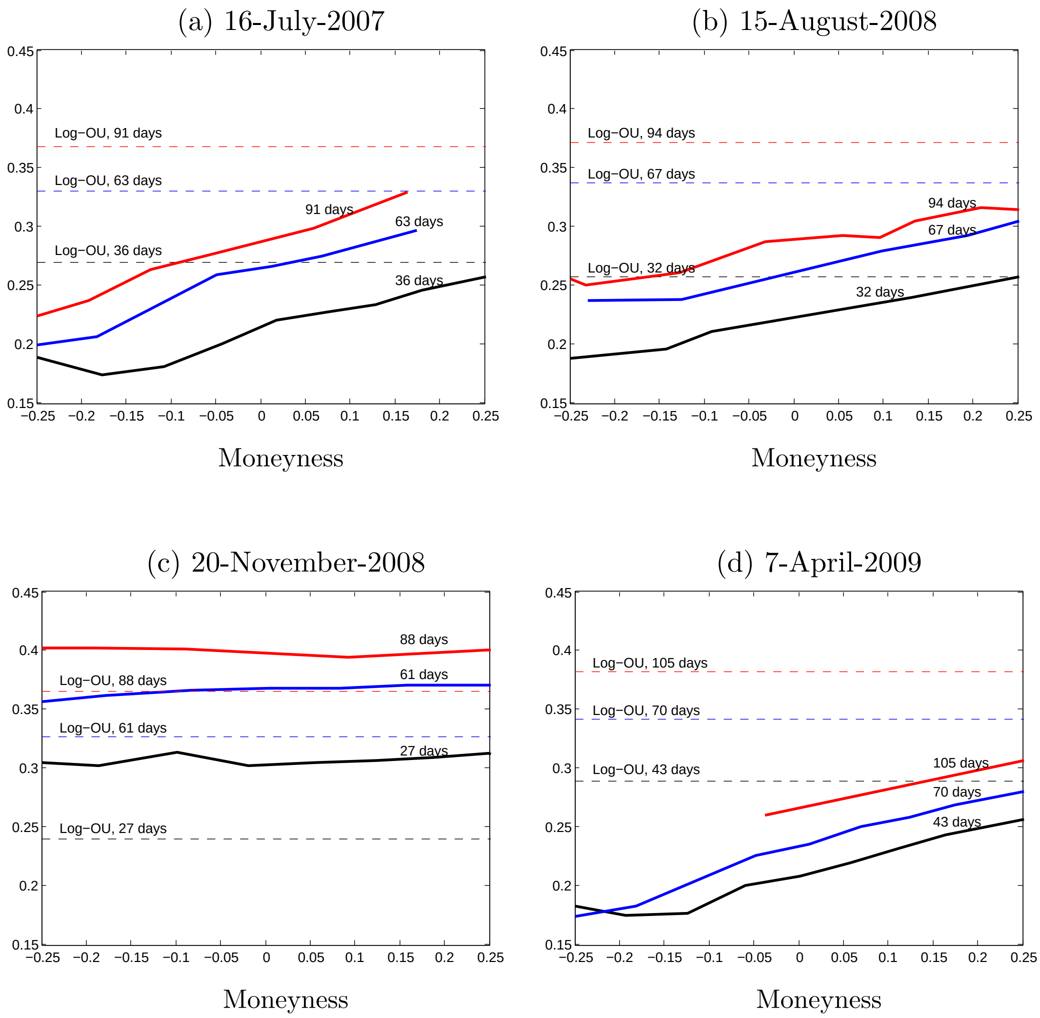

Note: Implied volatilities have been obtained by inverting the Black (1976) call price formula (5). Moneyness is defined as $\log(K/F(t,T))$.

Figure 7: Daily root mean square errors of standard log-OU and central tendency log-OU with stochastic volatility

(a) Futures

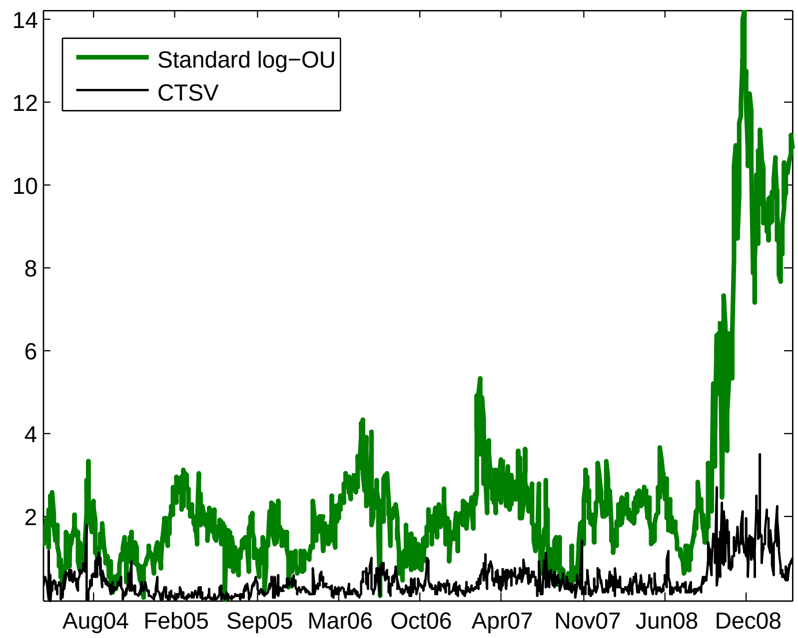

(b) Options

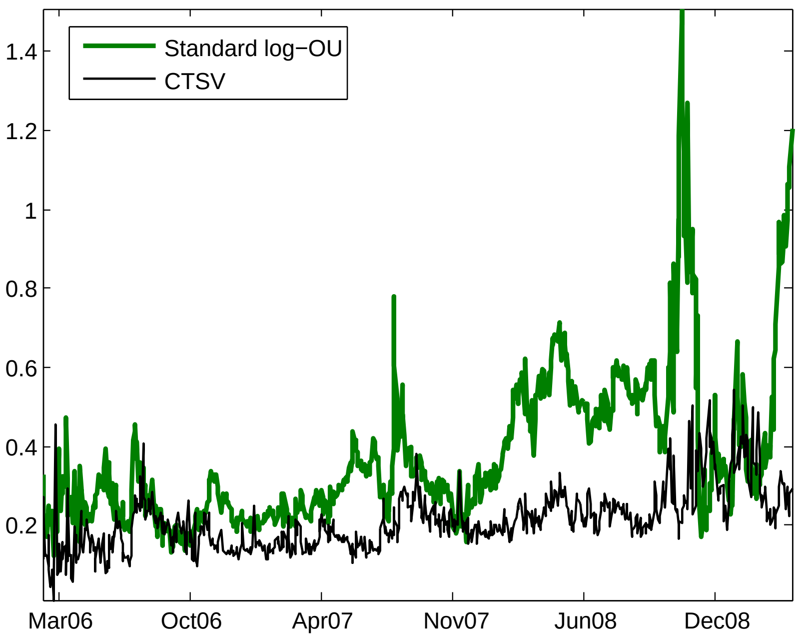

Notes: "Standard log-OU" refers to a log-normal Ornstein-Uhlenbeck process. "CTSV" generalises the standard log-OU by introducing a time varying central tendency and stochastic volatility. Results are based on full sample estimates.

Figure 8: Empirical cumulative distribution function of the square pricing errors

(a) Futures

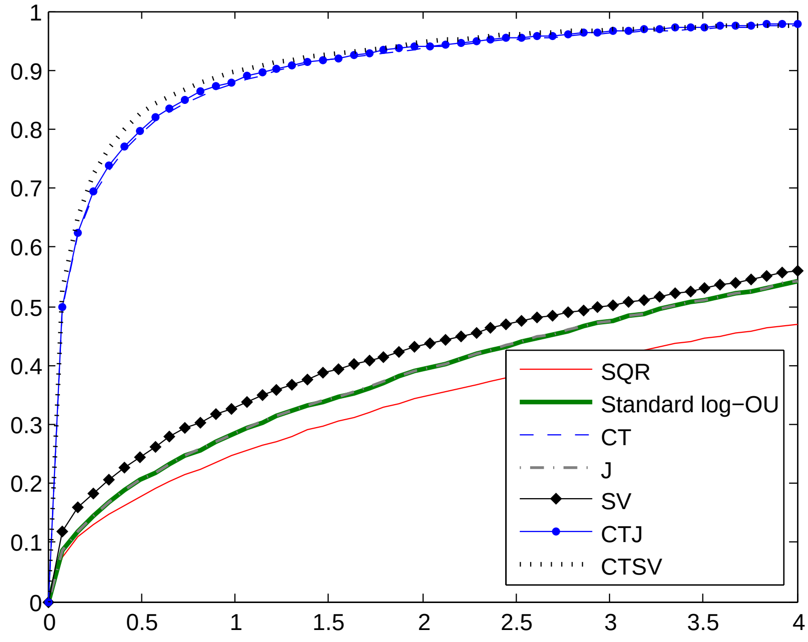

(b) Options

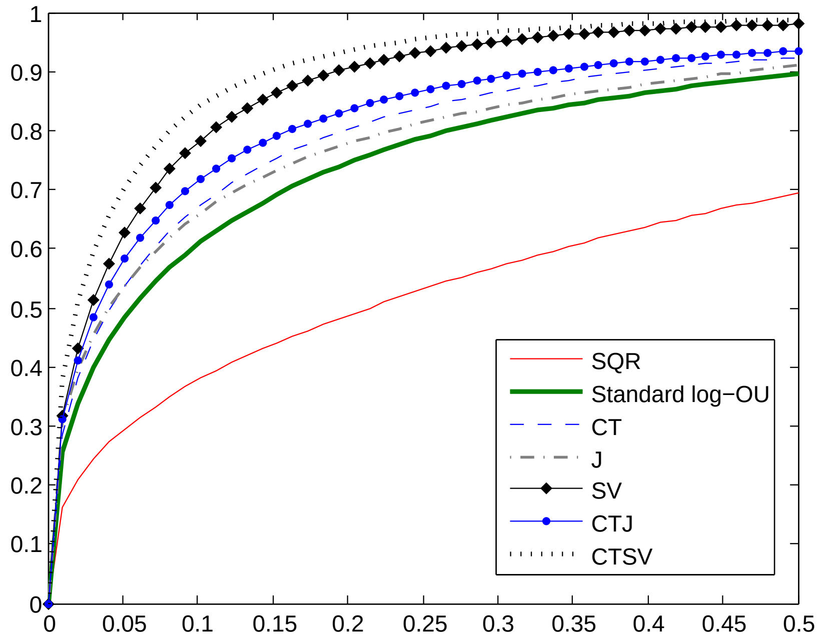

Notes: "SQR" denotes square root model and "log-OU" refers to a log-normal Ornstein-Uhlenbeck process. "J" introduces jumps in the standard log-OU model, whose size follows an exponential distribution. "CT" indicates Central Tendency, a log-OU process where the long run mean follows another log-OU diffusion. "SV" denotes Stochastic Volatility, a log-OU model where the volatility is a Gamma OU Lévy process. "CTJ" combines central tendency and jumps. "CTSV" combines central tendency and stochastic volatility. Results are based on full sample estimates.

Figure 9a: Filtered $\theta(t)$ for different log-OU models with central tendency

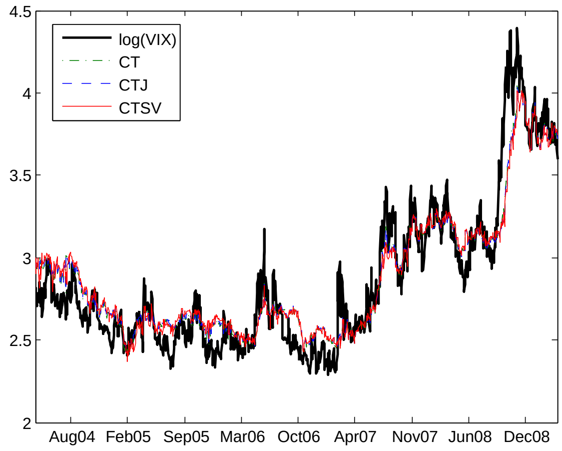

Figure 9b: Filtered $\theta(t)$ for different estimation samples in the log-OU model with central tendency and stochastic volatility

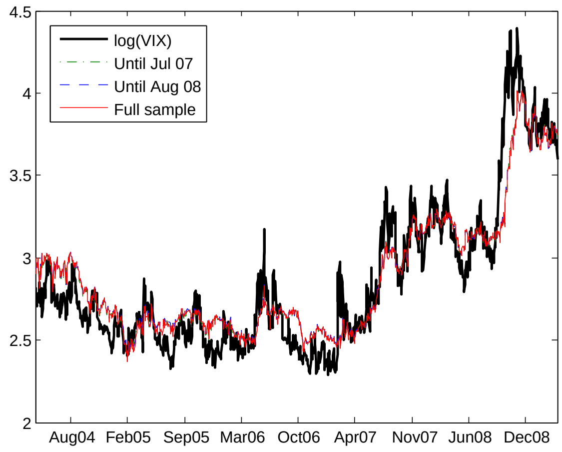

Figure 9c: Differences between model-based and actual one-month futures prices

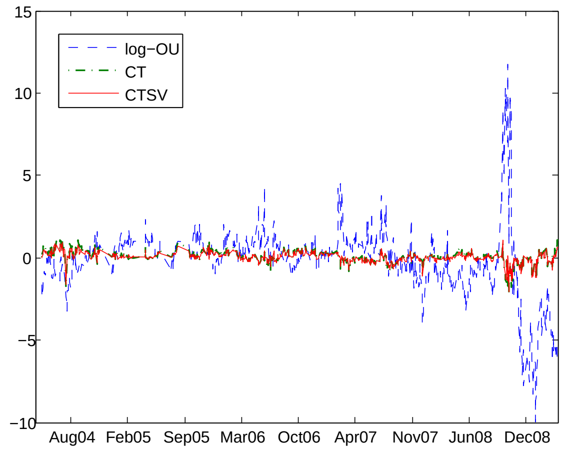

Notes: $\theta(t)$ denotes the time varying central tendency around which the VIX mean-reverts in central tendency models. "CT" indicates Central Tendency, a log-OU process where the long run mean follows another log-OU diffusion. "CTJ" combines central tendency and jumps with exponential sizes. "CTSV" combines central tendency and stochastic volatility. Results in Figures 9a and 9c are based on full sample estimates. One month actual futures prices in Figure 9c have been obtained by interpolation of the prices at the adjacent maturities.

Figure 10a: Filtered $\omega(t)$ for different log-OU models with stochastic volatility

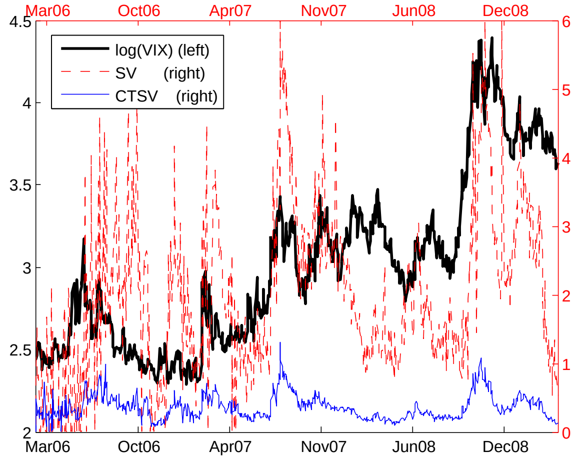

Figure 10b: Filtered $\omega(t)$ for different estimation samples in the log-OU model with central tendency and stochastic volatility

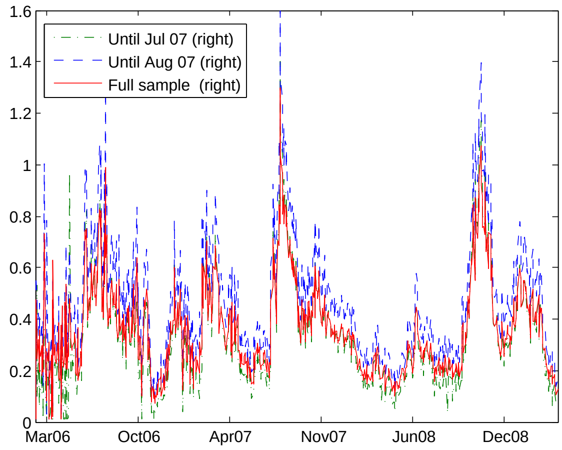

Figure 10c: One month volatilities of the VIX

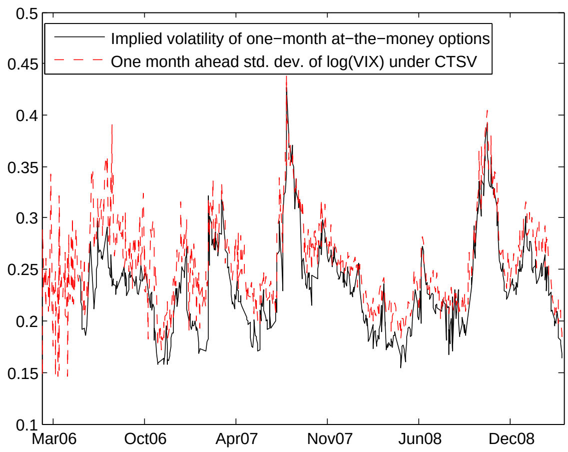

Notes: Different vertical scales are used for the VIX and the volatilities in Figure 9a. "SV" denotes a log-OU model with stochastic volatility, while "CTSV" refers to a log-OU model with central tendency and stochastic volatility. One-month implied vols in Figure 9c have been obtained by interpolation of the implied vols of options with moneyness $|\log(K/F(t,T))| < .1$, which in turn result from inverting the Black (1976) call price formula. Results in Figures 10a and 10c are based on full sample estimates.

Figure 11: Fit of the term structure of futures prices

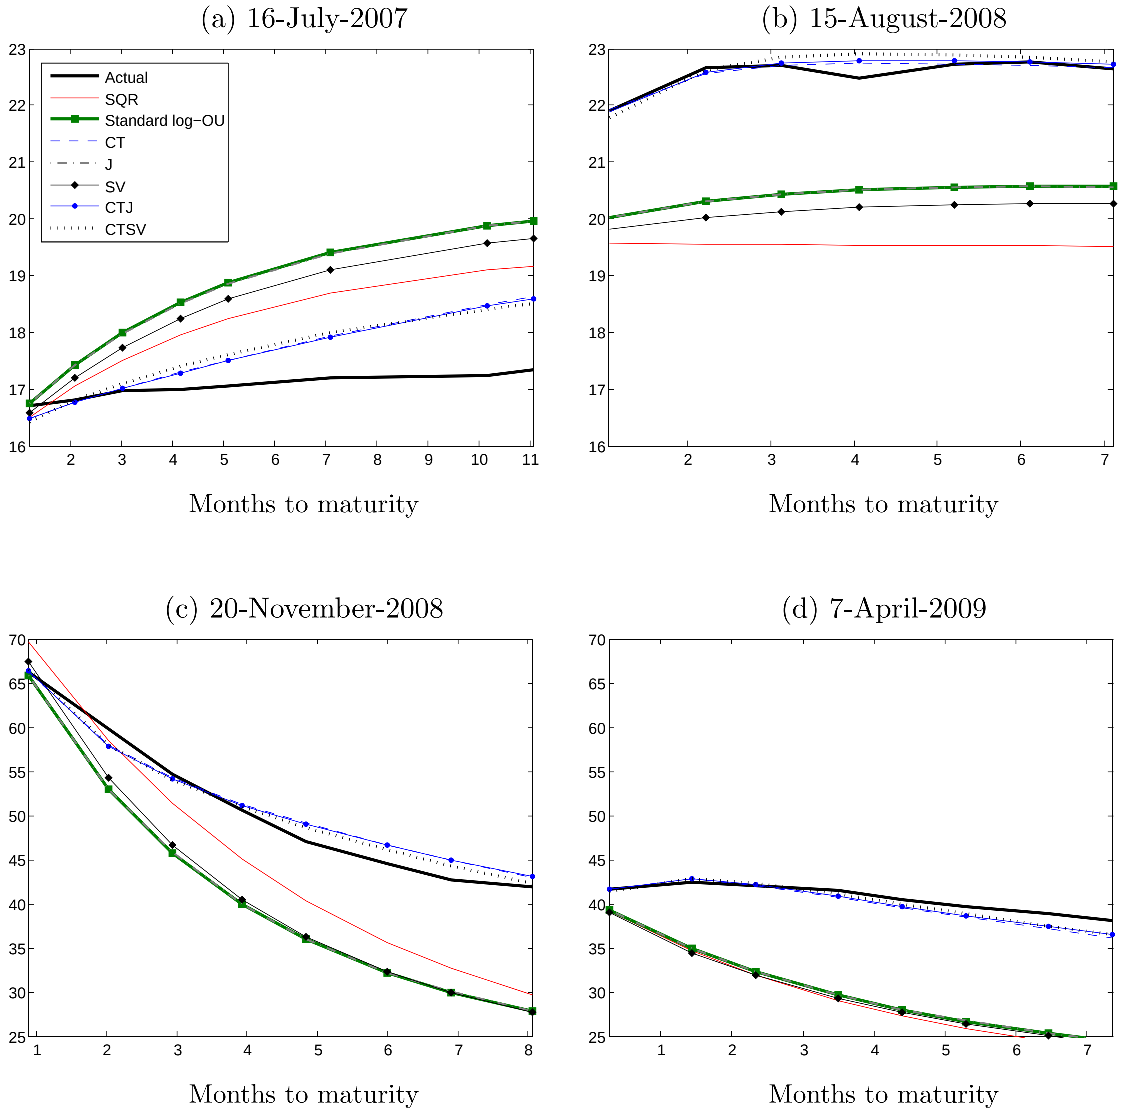

Notes: "SQR" denotes square root model and "log-OU" refers to a log-normal Ornstein-Uhlenbeck process. "J" introduces jumps in the standard log-OU model, whose size follows an exponential distribution. "CT" indicates Central Tendency, a log-OU process where the long run mean follows another log-OU diffusion. "SV" denotes Stochastic Volatility, a log-OU model where the volatility is a Gamma OU Lévy process. "CTJ" combines central tendency and jumps. "CTSV" combines central tendency and stochastic volatility. Results are based on full sample estimates.

Figure 12: Fit of the term structure of call prices for fixed strikes

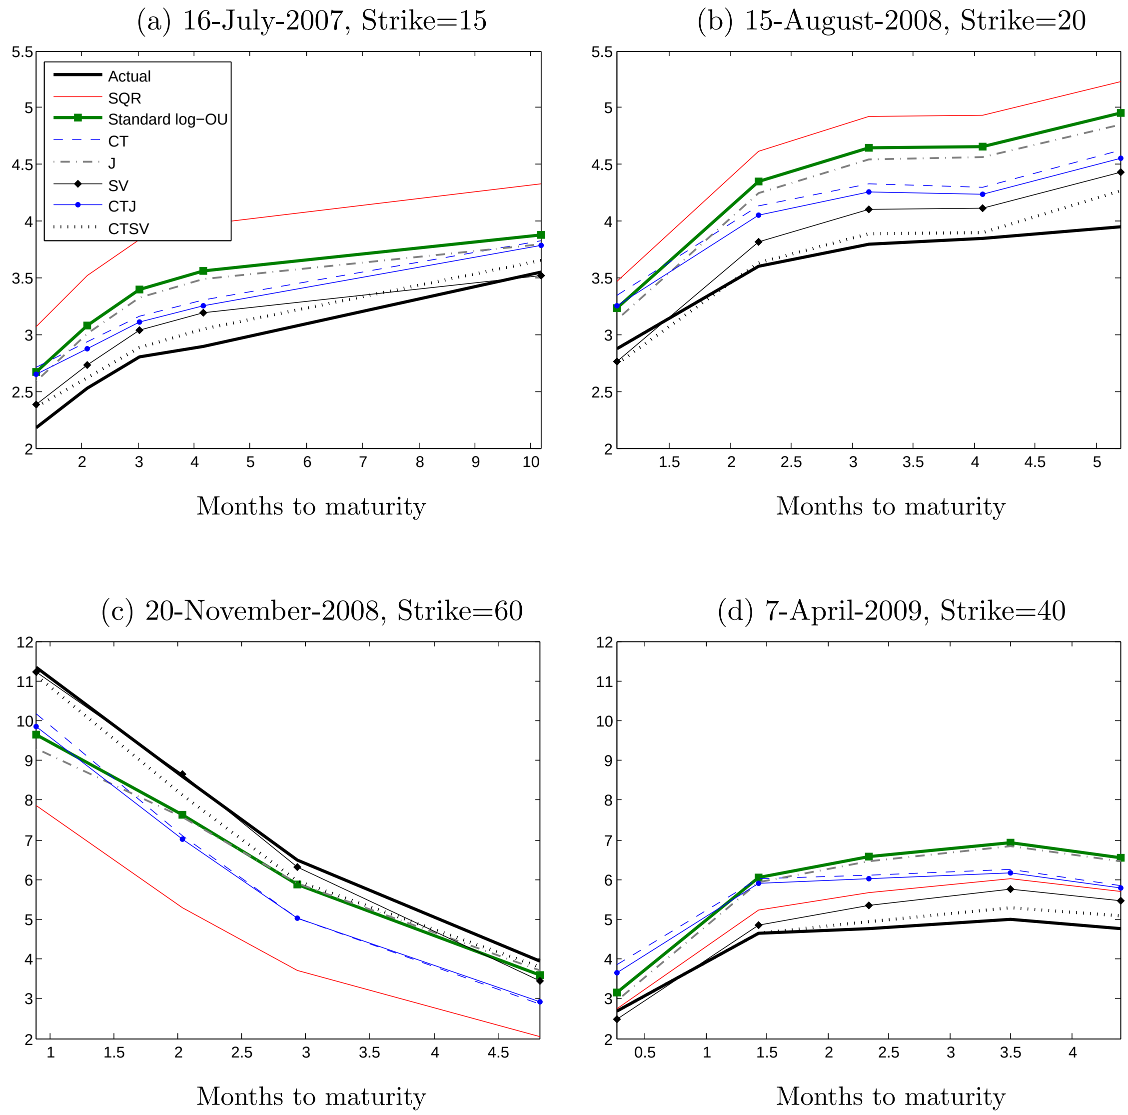

Notes: “SQR” denotes square root model and “log-OU” refers to a log-normal Ornstein-Uhlenbeck process. “J” introduces jumps in the standard log-OU model, whose size follows an exponential distribution. “CT” indicates Central Tendency, a log-OU process where the long run mean follows another log-OU diffusion. “SV” denotes Stochastic Volatility, a log-OU model where the volatility is a Gamma OU Lévy process. “CTJ” combines central tendency and jumps. “CTSV” combines central tendency and stochastic volatility. Results are based on full sample estimates.

Figure 13: Fit of the implied volatilities of call prices (i)

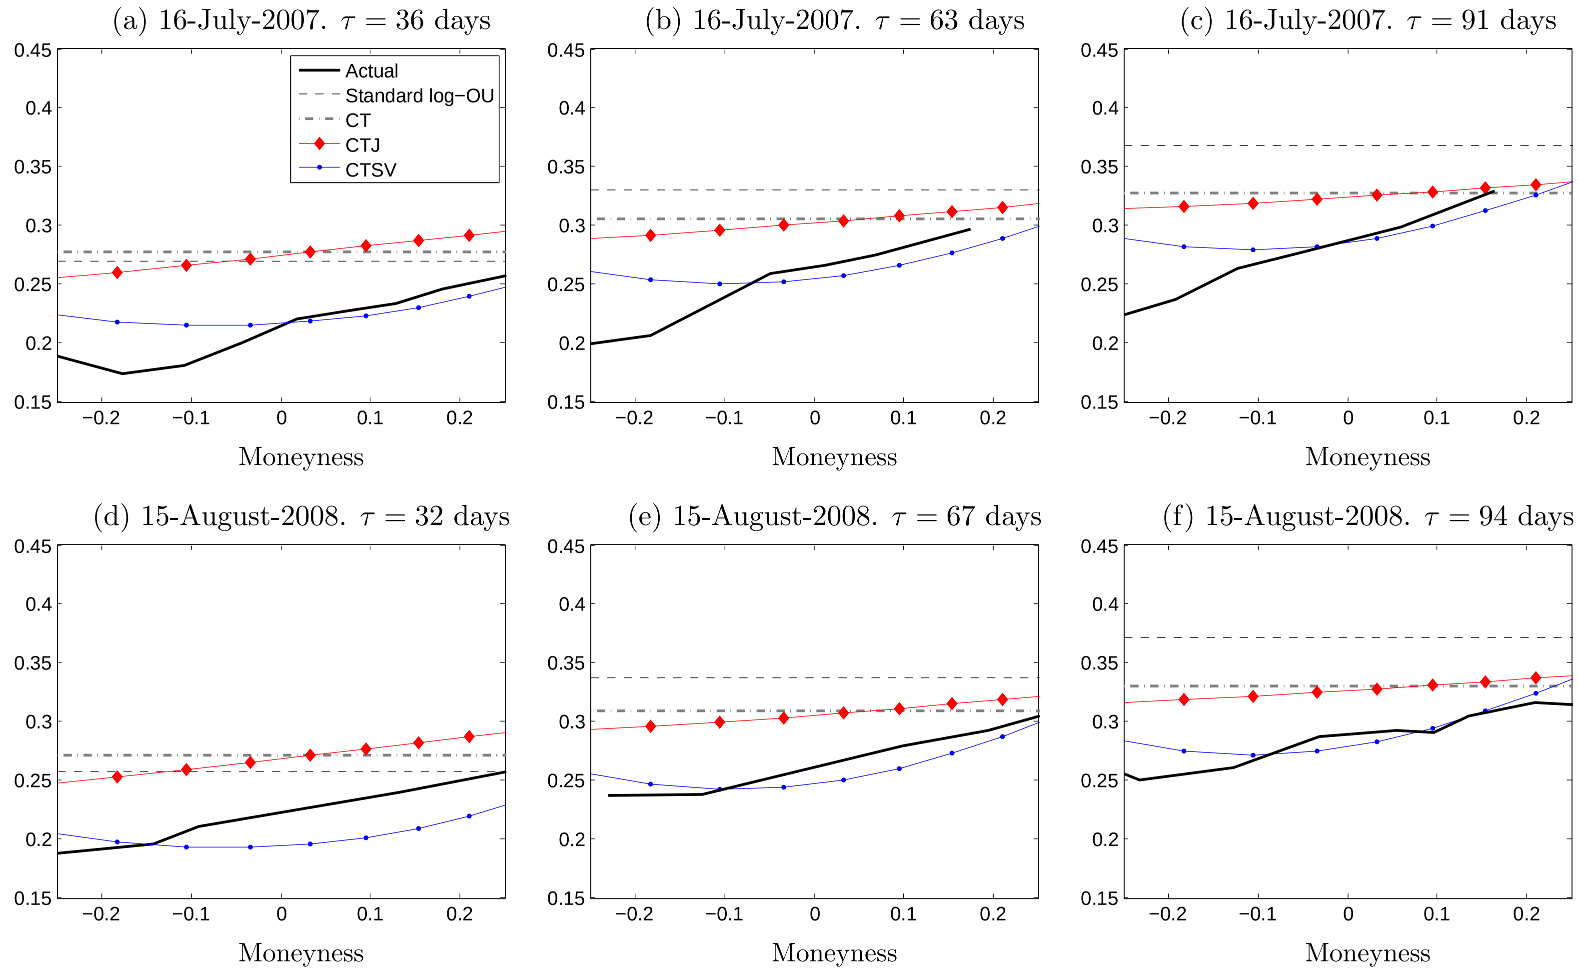

Note: Implied volatilities have been obtained by inverting the Black (1976) call price formula. Moneyness is defined as $\log(K/F(t,T))$. "Log-OU" refers to a log-normal Ornstein-Uhlenbeck process. "CTJ" combines central tendency and jumps. "CTSV" combines central tendency and stochastic volatility. Results are based on full sample estimates.

Figure 14: Fit of the implied volatilities of call prices (ii)

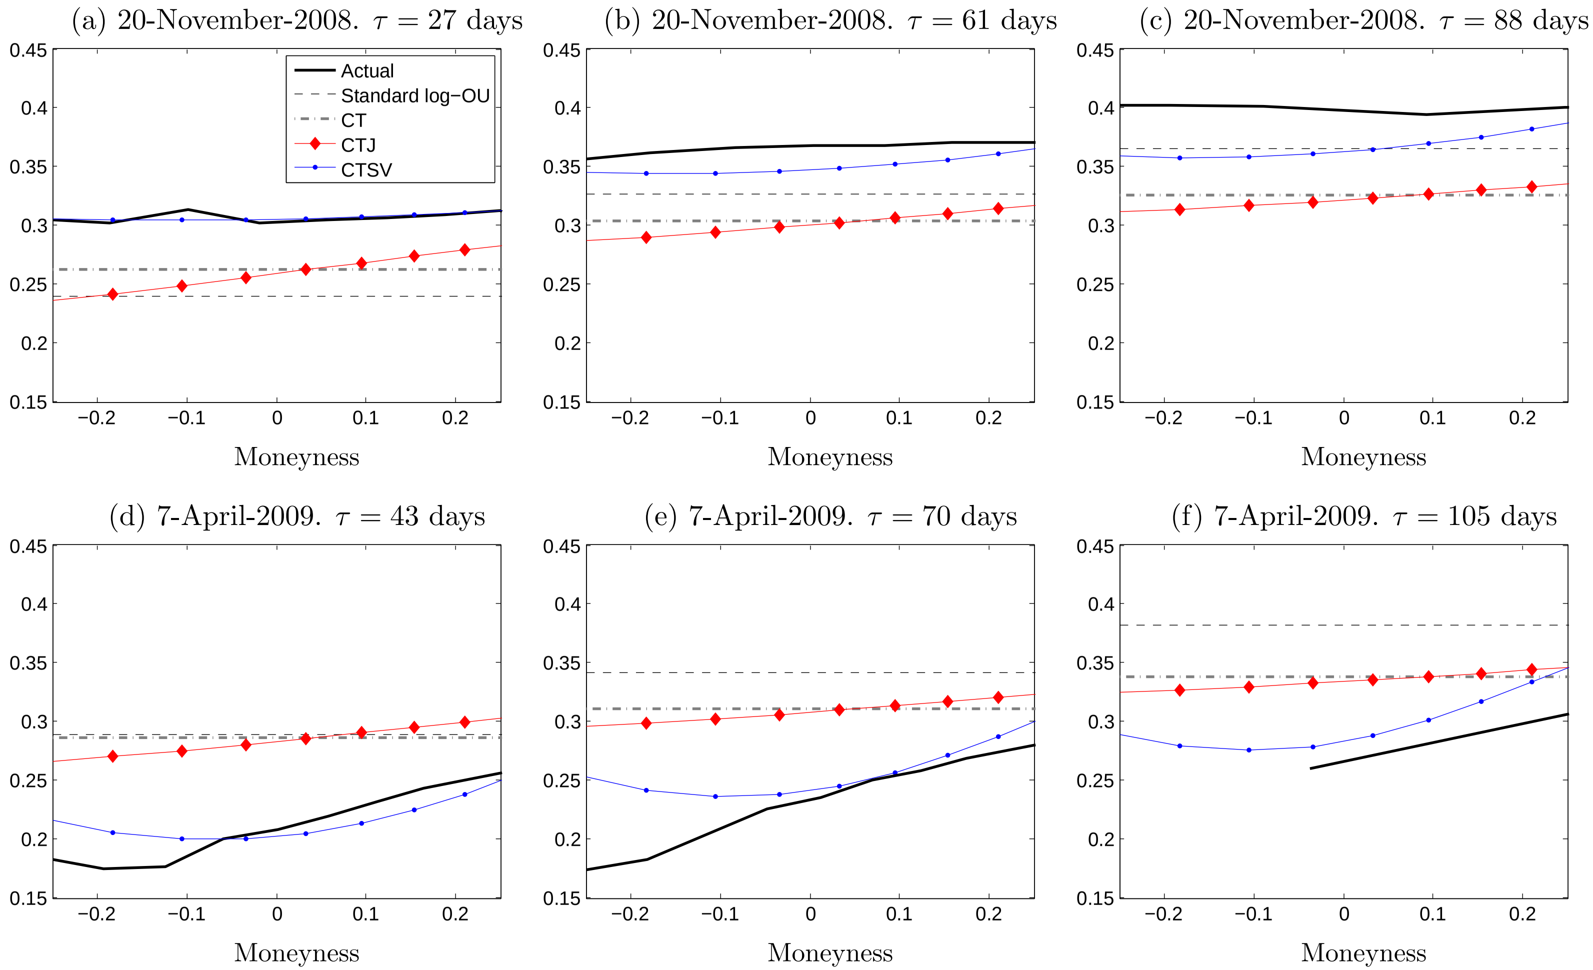

Note: Implied volatilities have been obtained by inverting the Black (1976) call price formula. Moneyness is defined as $\log(K/F(t,T))$. "Log-OU" refers to a log-normal Ornstein-Uhlenbeck process. "CTJ" combines central tendency and jumps. "CTSV" combines central tendency and stochastic volatility. Results are based on full sample estimates.