Valuation of VIX derivatives[☆]

Javier Mencía[a,*], Enrique Sentana[b]

[a]*Banco de España, Alcalá, 48, E-28014 Madrid, Spain*
[b]*CEMFI, Casado del Alisal, 5, E-28014 Madrid, Spain*

Abstract

We conduct an extensive empirical analysis of VIX derivative valuation models before, during, and after the 2008–2009 financial crisis. Since the restrictive mean-reversion and heteroskedasticity features of existing models yield large distortions during the crisis, we propose generalisations with a time-varying central tendency, jumps, and stochastic volatility, analyse their pricing performance, and implications for term structures of VIX futures and volatility "skews." We find that a process for the log of the observed VIX combining central tendency and stochastic volatility reliably prices VIX derivatives. We also uncover a significant risk premium that shifts the long-run volatility level.

Keywords: Central Tendency, Stochastic Volatility, Jumps, Term Structure, Volatility Skews
JEL: G13

[☆]We are grateful to Francis Longstaff (the referee), Dante Amengual, Torben Andersen, Max Bruche, Peter Carr, Rob Engle, Javier Gil-Bazo, Ioannis Paraskevopoulos, Francisco Peñaranda, Eduardo Schwartz, Neil Shephard and Dacheng Xiu, as well as seminar participants at Oxford MAN Institute, Universidad de Murcia, Universitat Pompeu Fabra, the SoFiE European Conference (Geneva, June 2009), the Finance Forum (Elche, November 2010) and the Symposium of the Spanish Economic Association (Madrid, December 2010) for helpful comments, discussions and suggestions. Of course, the usual caveat applies. Financial support from the Spanish Ministry of Science and Innovation through grants ECO 2008-00280 and 2011-26342 (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.

[*]Corresponding author. Tel: +34 91 338 5414, fax: +34 91 338 6102.

Email addresses: `javier.mencia@bde.es` (Javier Mencía), `sentana@cemfi.es` (Enrique Sentana)

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 Standard and Poor's 500 (S&P500) over the next month implicit in stock index option prices. Formally, it is the square root of the risk-neutral expectation of the integrated variance of the S&P500 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&P500 to buy protection in periods of market turmoil, which increases the value of the VIX. In fact, as Andersen and Bondarenko (2007) and many others show, the VIX almost uniformly exceeds realised volatility because investors are, on average, willing to pay a sizeable premium to acquire a positive exposure to future equity-index volatility. 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. Importantly, they can also be interpreted as options on VIX futures, which one can exploit to simplify their valuation and avoid compounding errors made in valuing futures. VIX options and futures are among the most actively traded contracts at CBOE and CFE, averaging close to 260,000 contracts combined per day in 2010, with much larger volumes on certain days.[1] One of the main reasons for the high interest in these products is that VIX derivative positions can be used to hedge the risks of investments in the S&P500 index. Szado (2009) finds that such strategies do indeed provide significant protection, especially in downturns. Moreover, by holding VIX derivatives investors can achieve exposure to S&P500 volatility without having to delta hedge their S&P500 option positions with the stock index. As a result, it is often cheaper to be long in out-of-the-money VIX call options than to buy out-of-the-money puts on the S&P500. Due to this possibility, VIX options are the only asset in which open interest is highest for out-of-the-money call strikes (Rhoads, 2011).

Although these new assets certainly offer additional investment and hedging opportunities for financial market participants, their correct use requires reliable valuation models that adequately capture the features of the underlying volatility index. In turn, the empirical performance of those valuation models can shed some light on the stochastic process for the VIX, and stock market volatility more generally, which is of interest to academic researchers. Moreover, the prices of volatility derivatives contain valuable information about the views financial market participants hold about the future, and we need reliable valuation models to make the correct inferences in a world with risk-averse agents.

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 the observed volatility index on which they are written follows a Geometric Brownian Motion (GBM). As a result, his model does not allow for mean-reversion in the VIX, which, as we shall see, seems to be at odds with the recent empir-

[1] For comparison purposes, this value represents about 37% of the average daily volume of S&P500 options, which was 695,000 in 2010. However, the volume of VIX derivatives is growing at a much faster rate. Specifically, the volume of VIX options already represented 52% of the volume of S&P500 options during the first ten months of 2011. In this sense, Rhoads (2011) points out that the volume traded is executed over a smaller time frame (7:20 am - 3:15 pm Central Time) than for other derivatives markets.

ical evidence.[2] The two most prominent mean-reverting models proposed so far for volatility indexes have been the square root process (SQR) considered by Grünbichler and Longstaff (1996), and the log-normal Ornstein-Uhlenbeck (LOU) process analysed by Detemple and Osakwe (2000).

Several authors have previously looked at the empirical performance of these pricing models for VIX derivatives. 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 an SQR diffusion. In addition, they estimate a GBM 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 (2011) compare the empirical fit of the SQR and GBM models using data on options written on the VIX. They also find evidence supporting the GBM assumption. However, one could alternatively interpret their findings as evidence in favour of the LOU process, which also yields the Black (1976) option 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 LOU 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 this volatility index mean-reverts at a simple, non-negative exponential rate. Such a limitation becomes particularly apparent during bearish stock markets, in which volatility measures such as the VIX typically experience large increases and remain at high levels for long periods. Arguably, the apparent success of those models is to a large extent due to the fact that the sample periods considered in the existing studies only cover 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 for volatility indexes to price deriva-

[2]The stationarity of volatility seems to depend on the historical period considered. Schwert (1990) and Pagan and Schwert (1990) find strong evidence for a unit root in stock volatility if the data span the 1930s. In contrast, Schwert (2011) convincingly argues that during the 2008–2009 crisis, volatility exhibits more mean-reversion than in the past.

tives on the VIX over a longer sample span that includes data before, during, and after the unprecedented recent financial crisis, which provides a unique testing ground for our study. To do so, we use an extensive database that comprises all futures and European option prices on the VIX from February 2006, when options were introduced, until December 2010. As a result, we can study whether the SQR and LOU models are able to yield reliable in- and out-of-sample prices in a variety of market circumstances.

Our findings indicate that although a LOU process for the VIX provides a better fit than an SQR model, especially for VIX options, its performance deteriorates during the market turmoil of the second part of our sample. For that reason, we consider an extension of the SQR model advocated by Bates (2012), among others, which allows for a time-varying central tendency in the mean as well as stochastic volatility that is unspanned by the VIX. Similarly, we generalise the LOU process by proposing several novel but empirically relevant extensions: a time-varying central tendency in the mean, jumps, and stochastic volatility. 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, and together with jumps, introduces non-normality in its conditional distribution. Importantly, we study the role that risk premiums play in reconciling the dynamics of the VIX with the prices of VIX derivatives in order to determine the existence of economically important systematic risks.

We estimate the SQR and LOU models by maximum likelihood. However, a closed-form expression for the likelihood is not available for many of the extensions. As a consequence, following Trolle and Schwartz (2009), among others, we specify our extensions in state space form and calibrate their parameters by pseudo maximum likelihood. We use both derivative prices and historical observations on the VIX itself since we want to estimate both real and risk-neutral measures. In order to compute the theoretical derivatives prices, we often need to invert the conditional characteristic function using Fourier methods (see Carr and Madan, 1999; and Amengual and Xiu, 2012).

To validate the models, we analyse the discrepancies between actual and theoretical derivatives prices. But we also go beyond pricing errors, and analyse the implications of the aforementioned extensions for the term structures of VIX futures, and the option volatility "skews," all of which are of considerable independent interest. Since we combine futures and options data, we can also assess which of those additional features is more relevant for pricing futures, and which one is more important for options.

Importantly, our approach differs from Sepp (2008), who prices VIX derivatives by using data on the S&P500 only. Although the relevant distribution of future values of the VIX conditional on current information might arguably depend not only on the current level of the VIX but also on the value of the S&P500, it seems odd to completely neglect the information content of contemporaneously observed VIX values. This is particularly relevant in view of the fact that one cannot reproduce the spot VIX index by using the model proposed by Sepp (2008) for the purposes of obtaining the required S&P500 option values that feed in the closed-form CBOE formula for this volatility index.[3] In order to avoid compounding valuation errors, therefore, we prefer to treat the current VIX level as our "sufficient statistic,"and to treat VIX futures likewise for the purposes of valuing options, instead of assigning that role to the S&P500. The same assumption has recently been made by Song and Xiu (2012). Ideally, though, one would like to use both, but we leave this for future research. In this sense, it is important to emphasise that the univariate stochastic processes that we consider for the VIX are not necessarily incompatible with models for S&P500 options capable of reproducing this observed volatility index.

Our analysis can also shed some light on the long-lasting debate surrounding the modeling and stationarity of the (instantaneous) volatility of this broad stock market index. In particular, given the relationship between the observed VIX and the unobserved integrated volatility of the S&P500, our results imply that the sophisticated stochastic volatility models that many researchers and practitioners use for valuing S&P options should probably 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 several Appendices.

[3]A relevant analogy would be the use of a model for the dividends and risk premia of the 500 individual stocks that constitute the index for the purposes of valuing S&P500 derivatives discarding the information in the index itself, even though any such multivariate model would very likely fail to reproduce the value of the index.

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&P100 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&P500 index for the front month and the second month expirations with at least eight days left (see Chicago Board Options Exchange, 2009).[4] Since VIX is expressed in annualised terms, investors typically divide it by $16(=\sqrt{256})$ to gauge the expected size of the daily movements in the stock market implied by this index (see Rhoads, 2011).

Fig. 1a displays the entire historical evolution of the VIX. Between January 1990 and December 2010, its average closing value was 20.4. 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. Fig. 1b, which focuses on the sample period in our derivatives database, shows that volatility was also remarkably low between February 2006 and July 2007, with values well below 20. During this period, the lowest value was 9.89 on January 24, 2007, in what some have called "the calm before the storm."Indeed, although the Dow Jones Industrial Average closed above 14,000 for the first time in history in the summer of 2007, some warning signals were observed around this period. In particular, on June 22, 2007, Bear Stearns pledged up to $3.2 billion in loans to bail out one of its hedge funds, which was collapsing due to bad bets on subprime mortgages. Moreover, on July 18, 2007, this investment bank disclosed that its two subprime hedge funds had lost all of their value, and one day later Fed Chairman Ben Bernanke warned the US Senate Banking Committee that subprime losses

[4] Currently, CBOE applies the VIX methodology to three-month options on the S&P500 (VXV), as well as one-month options on the most important US stock market indexes: DJIA (VXD), S&P100 (VXO), Nasdaq-100 (VXN), and Russell 2000 (RVX). They also construct analogous short-term volatility indexes for Crude Oil (OVX), Gold (GVZ), Soybean (SIV), Corn (CIV), and the US $/€exchange rate (EVZ).

could top $100 billion. Perhaps not surprisingly, over the following year 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. After this peak, VIX followed a decreasing trend over the following months until the beginning of April 2010, when the Greek debt crisis started worsening. These markedly different regimes offer a very interesting testing ground to analyse the 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 autoregressive-moving-average (ARMA) models using the entirety of VIX historical observations from 1990 to 2010; see French, Schwert, and Stambaugh (1987) for a related analysis. Fig. 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 Fig. 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. Importantly, this result is not due to the behaviour of the VIX during the financial crisis, since an ARMA(2,1) is also necessary to capture the autocorrelation profile when we only consider data from 1990 until the summer of 2007.

We have also analysed the presence of time-varying features in the volatility of the log VIX. In particular, we have tested for generalised autoregressive conditional heteroskedasticity (GARCH) effects in the residuals of the ARMA(2,1) estimation.[5] Interestingly, we can easily reject the conditional

[5]Following Demos and Sentana (1998), we carry out a one-sided Lagrange multiplier (LM) test of conditional homoskedasticity against GARCH as $T\tilde{R}^2$ from the regression of

homoskedasticity of the log VIX at all conventional levels. As we shall in Section 5.4, the volatility of our preferred continuous time specification closely matches the volatility of an ARMA(2,1)-GARCH(1,1) model.

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 apply several filters to ensure the reliability of the data that we finally use. Following Dumas, Fleming, and Whaley (1998), we exclude derivatives with fewer than six days to expiration due to their illiquidity. In addition, we only retain those prices for which open interests and volumes are available. We consider the entire history of these series since options were introduced in February 2006 until December 2010. In terms of maturity, we have data on all the contracts with expirations between March 2006 and May (July) 2011 for options (futures). CFE may list futures for up to nine near-term serial months, as well as five months on the February quarterly cycle associated to the March quarterly cycle for options on the S&P500. 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 three near-term months plus up to three additional months on the February quarterly cycle.[6]

All in all, we have 8,665 and 87,870 prices of futures and options, respectively. Of those option prices, 58,099 correspond to calls and 29,771 to puts. We have between four and six daily futures prices during 2006; afterwards, eight prices per day become available, on average. 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. The number of call prices per day increases again after mid-2009 and is greater than 60 at the end of the sample. We proxy for the riskless interest rate by using the daily Eurodollar rates at one-week, one, three and six months, and one year, which we interpolate to match the maturities of the futures and option contracts

---

the squared ARMA(2,1) residuals on a constant and the RiskMetrics volatility estimate (see RiskMetrics Group, 1996).

[6]VIX futures contracts have a multiplier of 1,000, while VIX option contracts have a multiplier of 100. This means that the value of the VIX futures contract is determined by multiplying $1,000 times the quoted futures price. On March 2, 2009, the CBOE introduced mini-VIX futures, which have a multiplier of 100 (see Rhoads, 2011).

that we observe.[7]

Fig. 3a shows the evolution of the VIX term structure implicit in VIX futures. Futures prices remained low albeit 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 October 12, 2006, the VIX was 11.09, while the price of VIX futures expiring on February 14, 2007, was 15.72. The spot price reflected the expected volatility for the period of October 12 to November 11, while the futures price reflected the expected volatility for the period of February 14 to March 16, 2007. Fig. 3a clearly confirms that there was an initial increase in the summer of 2007 and a substantial level shift in the last quarter of 2008. The negative slope during the latter period, though, indicates that the market did not expect the VIX to remain at such high values forever, a fact already highlighted by Schwert (2011). Eventually, this prediction turned out to be on the winning side and futures prices significantly came down until April 2010, right before the beginning of the European sovereign debt crisis.

Figs. 3b and 3c show the average implied volatility skews implicit in VIX option prices for different times to maturity. We consider low volatility days on Fig. 3b and high volatility days on Fig. 3c. We define low (high) volatility days as those in which the implied volatility of one-month at-the-money options remains below (above) their average value over the sample. Interestingly, we observe a consistent pattern of implied volatility smirks with positive slope regardless of the time to maturity and volatility of the VIX. There only seems to be a small change in slope for deep in-the-money calls (or out-of-the-money puts). Therefore, any pricing model must adequately capture these empirical features. Finally, although Figs. 3b and 3c show that option prices were higher for longer maturities on average, this is not necessarily the case on all the days in the sample. In particular, option prices decreased with maturity at the peak of the financial crisis (November 2008) because the market expected a fall in the volatility of the VIX.

[7]Given that Eurodollar and London interbank offered rates (LIBOR) are very similar, the results are unlikely to be sensitive to this choice. As usual in the option pricing literature, we assume constant interest rates in our models. Although interest rates are of course not constant in practice, we have checked that their changes are not significantly correlated with changes in the VIX.

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 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 $\mathcal{M}$ index the asset pricing models that we consider. Then, the theoretical futures price implied by model $\mathcal{M}$ will be:

$$F_{\mathcal{M}}(t,T,V(t),\boldsymbol{\phi}) = E^{Q}_{\mathcal{M}}[V(T)|I(t),\boldsymbol{\phi}], \tag{1}$$

where Q indicates that the expectation is evaluated at the risk-neutral measure, $\boldsymbol{\phi}$ is the vector of free parameters of model $\mathcal{M}$, 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

$$c_{\mathcal{M}}[t,T,K,F(t,T),\boldsymbol{\phi}] = \exp(-r\tau)E^{Q}_{\mathcal{M}}[\max(V(T)-K,0)|I(t),\boldsymbol{\phi}], \tag{2}$$

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 $\mathcal{M}$ can be easily obtained from the put-call-forward parity relationship

$$p_{\mathcal{M}}[t,T,K,F(t,T),\boldsymbol{\phi}] = c_{\mathcal{M}}[t,T,K,F(t,T),\boldsymbol{\phi}] - \exp(-r\tau)[F(t,T)-K].$$

We use futures and options prices, as well as data on the VIX index, which allows us to estimate the risk premia implicit in the market of VIX derivatives. We estimate the parameters by maximum likelihood for the models whose likelihood is known in closed form, and by pseudo maximum likelihood for the remaining ones. In both cases, we assume the existence of pricing errors such that

$$F(t,T) = F_{\mathcal{M}}(t,T,V(t),\boldsymbol{\phi}) + \xi_{ft} + \varepsilon_{ft,T},$$

$$c(t,T,K) = c_{\mathcal{M}}[t,T,K,F(t,T),\boldsymbol{\phi}] + \varpi(M(t,T))\left[(1-\eta_o)\xi_{ot} + \eta_o\varepsilon_{ct,T}(K)\right],$$

$$p(t,T,K) = p_{\mathcal{M}}[t,T,K,F(t,T),\boldsymbol{\phi}] + \varpi(M(t,T))\left[(1-\eta_o)\xi_{ot} + \eta_o\varepsilon_{pt,T}(K)\right],$$

where $\varepsilon_{ft,T} \sim N(0,\sigma^2_{f\varepsilon})$, $\varepsilon_{ct,T}(K) \sim N(0,1)$, and $\varepsilon_{pt,T}(K) \sim N(0,1)$ are orthogonal and independent and identically distributed (*iid*) over time and across strikes and maturities, while $\xi_{ft} \sim N(0,\sigma^2_{f\xi})$ and $\xi_{ot} \sim N(0,1)$ are also orthogonal and *iid* over time but common for all futures and options, respectively. Thus, we decompose pricing errors into common and idiosyncratic terms. Thanks to this parametrisation, we allow for contemporaneous correlation between the pricing errors. As noted by Bates (2000), ignoring this feature would affect the relative weighting that the options and futures would receive in the estimation. In contrast, the inclusion of common factors avoids overestimating the amount of truly independent information in the data.[8] In addition, we parametrise the variances of option pricing errors by means of the quadratic function

$$\varpi^2(M(t,T)) = \sigma_a + \sigma_b(M(t,T) - \sigma_c)^2, \tag{3}$$

where $M(t,T) = \log(K/F(t,T))$ measures the moneyness of the option. In this way, we avoid giving undue weight to higher priced or in-the-money options whose pricing errors are higher simply because of their larger scale. Importantly, we also employ different parameters for the pricing error variances of futures and options, thereby implicitly adjusting their relative weight in the likelihood.

We can write the log likelihood at a particular time t as

$$\begin{aligned} l[\mathbf{o}_t,\mathbf{f}_t,V(t)|I(t-1)] &= l_o(\mathbf{o}_t|\mathbf{f}_t,V(t),I(t-1)) + l_f(\mathbf{f}_t|V(t),I(t-1)) \\ &\quad + l_v[V(t)|I(t-1)], \end{aligned}$$

where $\mathbf{o}_t$ and $\mathbf{f}_t$ are, respectively, the set of option and futures prices traded at t, $l_o(\mathbf{o}_t|\mathbf{f}_t,V(t),I(t-1))$ is the log density of the option prices conditional

[8]Nevertheless, empirical results without the common factors ξ_{ft} and ξ_{ot}, which are available from the authors upon request, yield qualitatively similar results.

on futures prices and VIX, $l_f(\mathbf{f}_t|V(t), I(t-1))$ is the log likelihood of the futures prices conditional on the volatility index, and $l_v[V(t)|I(t-1)]$ is the log density of the volatility index.[9] Thus, the log likelihood reflects that we price futures conditional on the value of the VIX, in the same way as we price options conditional on the value of the futures. In addition, we take into account the contribution of the VIX to the log likelihood in order to obtain the parameters under the real measure.

A nontrivial advantage of our estimation method over traditional calibration procedures is that we do not have to treat the model parameters as deterministic functions of time. In addition, we can also obtain standard errors for our parameter estimators, which can in turn be used to conduct formal hypothesis tests. Furthermore, we can exploit the relationship between actual and risk-neutral measures to estimate prices of risk (see Appendices Appendix B and Appendix C for details).

We consider two estimation samples that correspond to two distinct volatility phases described in Section 2: until 15-Aug-2008 and full sample. We use the first sample to evaluate the out-of-sample empirical fit of the models.[10] Thus, we can assess model performance following a major volatility increase 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

[9] The usefulness of this additive decomposition is limited, though, because most model parameters affect more than one component.

[10] We analyse out-of-sample performance using the parameter estimates obtained in-sample. We also use the current value of the VIX to compute futures prices and the current value of futures prices to compute option prices, just as a financial market participant would do in real time. In addition, in the models with more than one factor, we follow the standard procedure of filtering the additional factor using the past evolution of the VIX and its derivatives with the in-sample estimates.

prices, satisfies the diffusion

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

where $W^P(t)$ is a Brownian motion under the real measure. Following Dai and Singleton (2000), we specify a price of risk that is proportional to the instantaneous volatility. Specifically, we assume that $\Lambda_v(t) = \varsigma\sqrt{V(t)}$, so that $dW^Q(t) = dW^P(t) + \varsigma\sqrt{V(t)}dt$. Then, it is straightforward to show that $V(t)$ satisfies the square root diffusion

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

under the equivalent risk-neutral measure, where $\kappa = \kappa^P + \sigma\varsigma$ and $\bar{\theta} = \kappa^P\bar{\theta}^P/\kappa$.

As is well known, the risk-neutral 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

$$F_{SQR}(t, T, V(t), \kappa, \bar{\theta}, \sigma) = \bar{\theta} + \exp(-\kappa\tau)[V(t) - \bar{\theta}]. \tag{4}$$

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.[11] Interestingly, the SQR model introduces stochastic volatility because the conditional variance of the VIX implied by this model is an affine function of $V(t)$ (see Appendix Appendix D).

The call price formula (2) for this model becomes

$$\begin{aligned} c_{SQR}(t, T, K, \kappa, \bar{\theta}, \sigma) = {} & V(t)\exp(-(\kappa + r)\tau)[1 - F_{NC2}(2cK; \nu + 4, \psi)] \\ & +\bar{\theta}[1 - \exp(-\kappa\tau)]\exp(-r\tau)[1 - F_{NC2}(2cK; \nu + 2, \psi)] \\ & -K\exp(-r\tau)[1 - F_{NC2}(2cK; \nu, \psi)], \end{aligned} \tag{5}$$

[11] 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.

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 (4).

Subsequently, Detemple and Osakwe (2000) considered the log-normal Ornstein-Uhlenbeck (LOU) diffusion:

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

If once again we follow Dai and Singleton (2000) in specifying the price of risk as $\Lambda_v(t) = \varsigma$, then $V(t)$ will satisfy the LOU diffusion

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

under the risk-neutral measure, where $\bar{\theta} = \bar{\theta}^P - \sigma\varsigma/\kappa$.

As is well known, this model implies that $\log V(t)$ would follow a conditionally homoskedastic Gaussian AR(1) process if sampled at equally spaced discrete intervals. More generally, the conditional risk-neutral 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

$$\varphi^2(\tau) = \frac{\sigma^2}{2\kappa}[1 - \exp(-2\kappa\tau)]. \tag{6}$$

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 $\bar{\theta}$. And although (6) shows that the LOU model is homoskedastic for the log VIX, it can be shown that the process followed by the VIX is heteroskedastic due to Jensen's inequality (see Appendix Appendix D for more details). In this context, it is straightforward to show that the futures price is

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

while the call price can be expressed as

$$\begin{aligned} c_{LOU}(t,T,K,\kappa,\bar{\theta},\sigma) &= \exp(-r\tau)F(t,T)\Phi\left[\frac{\log(F(t,T)/K) + \frac{1}{2}\varphi^2(\tau)}{\varphi(\tau)}\right] \\ &\quad -K\exp(-r\tau)\Phi\left[\frac{\log(F(t,T)/K) - \frac{1}{2}\varphi^2(\tau)}{\varphi(\tau)}\right] \end{aligned} \tag{7}$$

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 (6) 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 (7) 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 two sample periods described at the end of Section 3. Table 1 reports the in- and out-of-sample root mean square pricing errors (RMSE). We observe that the SQR model yields larger in-sample price distortions than the LOU model. Those distortions remain in the out-of-sample period. We have re-estimated the two models with VIX and futures data only to check that the use of option data in the estimation of the parameters is not driving the results. We have found that the LOU model still yields a better performance in that case.

Table 2 reports the parameter estimates that we obtain under the real measure, as well as the parameters of the pricing error variances. Although the values are not directly comparable because volatility is expressed in levels in the SQR model and in logs in the LOU model, in both cases the mean-reversion parameter κ is quite sensitive to the sample period used for estimation purposes (see footnote 2 and references therein). In contrast, the volatility parameter σ is more stable for the LOU 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. Table 2 also shows that the parameters of the quadratic function (3) are highly significant. Therefore, allowing for different variances of pricing errors across strikes seems to be relevant in ensuring that higher priced options do not receive undue weight.

An alternative, more illustrative way to assess the validity of the LOU model is by considering the implied volatilities obtained with the Black (1976) formula (7). As we mentioned before, the implied volatility of the LOU process is constant for different degrees of moneyness, but not across different maturities because (6) depends on τ. Hence, if this model were correct, then we should obtain constant implied volatilities for a given maturity regardless of moneyness. As we have seen in Figs. 3b and 3c, this is not at all the case. In fact, we generally observe that implied vols have a positive moneyness slope. In addition, the comparison of these figures suggests that the volatility of the log VIX is not constant over time, as the LOU model assumes. In this sense, the LOU process is not only unable to generate the observed volatility skews, but it also fails to capture the average level of implied volatilities. As we will see in Section 5.5, though, the SQR model generates volatility skews, but they have the opposite slope to the actual skews in Figs. 3b and 3c.

5. Extensions

5.1. Model specification

The results of the previous section indicate that a LOU process offers a better empirical fit than a SQR process. Unfortunately, their performance tends to deteriorate during the recent financial crisis. For that reason, in this section we explore several extensions to those models. For the sake of brevity, though, we focus on the risk-neutral measure and the prices of risk. Further details can be found in Appendices Appendix B and Appendix C.

Specifically, we extend the SQR model by considering the concatenated SQR (CSQR) proposed by Bates (2012), among others:

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

with

$$d\theta(t) = \bar{\kappa}[\bar{\theta} - \theta(t)]dt + \bar{\sigma}\sqrt{\theta(t)}dW_\theta^Q(t).$$

We specify the prices of risk that link the Wiener processes under P and Q as

$$\begin{aligned} dW_v^Q(t) &= dW_v^P(t) + \varsigma\sqrt{V(t)}dt, && (8) \\ dW_\theta^Q(t) &= dW_\theta^P(t) + \bar{\varsigma}\sqrt{\theta(t)}dt, && (9) \end{aligned}$$

where $W_v^Q(t)$ and $W_\theta^Q(t)$ are independent Brownian motions. These specifications ensure that the VIX follows the same process under the real measure. As we mentioned before, an important deficiency of previously existing models is that they assume that the volatility index mean-reverts at a simple, non-negative exponential rate, which is in fact zero in the GBM model proposed by Whaley (1993). However, the long, persistent swings in the VIX in Fig. 1a suggest that we need to allow for more complex dynamics. In this sense, the CSQR allows the VIX 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 $V(T)$ is a function of the distance between $V(t)$ and the central tendency $\theta(t)$, as well as the distance between $\theta(t)$ and $\bar{\theta}$.[12] More explicitly,

$$E^Q[V(T)|I(t)] = \bar{\theta} + \delta(\tau)[\theta(t) - \bar{\theta}] + \exp(-\kappa\tau)[V(t) - \theta(t)], \tag{10}$$

where

$$\delta(\tau) = \frac{\kappa}{\kappa - \bar{\kappa}} \exp(-\bar{\kappa}\tau) - \frac{\bar{\kappa}}{\kappa - \bar{\kappa}} \exp(-\kappa\tau). \tag{11}$$

Importantly, the CSQR model introduces stochastic volatility that is not spanned by the VIX, since the conditional variance under this model is an affine function of $V(t)$ and $\theta(t)$ (see Appendix Appendix B).

In turn, we add three empirically relevant features to the LOU model: a time-varying mean, jumps, and stochastic volatility. We consider these extensions first in isolation, and then in combination. As a general rule, we model the price of risk of the continuous part of the diffusions as in the one-factor models, while we assume that jump risk is not priced. All in all, we compare the following cases:

- Central tendency (CTOU):

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

  with

  $$d\theta(t) = \bar{\kappa}[\bar{\theta} - \theta(t)]dt + \bar{\sigma} dW_\theta^Q(t), \tag{12}$$

  where $W_v^Q(t)$ and $W_\theta^Q(t)$ are independent Brownian motions. We specify the prices of risk that link the Wiener processes under P and Q as

  $$\begin{aligned} dW_v^Q(t) &= dW_v^P(t) + \varsigma dt, & (13)\\ dW_\theta^Q(t) &= dW_\theta^P(t) + \bar{\varsigma} dt. & (14) \end{aligned}$$

  These specifications ensure that the process followed by the log VIX is also affine under the real measure.

  We show in Appendix Appendix C that the exact discretisation of $\log V(t)$ in the above model is a Gaussian ARMA(2,1) process, which

[12] Unlike in the SQR model, the convergence of (10) to the long-run mean is not necessarily a monotonic function of τ.

is consistent with the evidence reported in Section 2 (see Fig. 2). Therefore, its likelihood function can be computed in closed form. As discussed in Jegadeesh and Pennacchi (1996) and Balduzzi, Das, and Foresi (1998), (12) allows the volatility index to revert towards a time-varying central tendency whose long-run mean is $\bar{\theta}$. 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 $\bar{\theta}$.[13] More explicitly,

$$\begin{aligned} E^Q[\log V(T)|I(t)] &= \bar{\theta} + \delta(\tau)[\theta(t) - \bar{\theta}] \\ &\quad + \exp(-\kappa\tau)[\log V(t) - \theta(t)], \end{aligned} \qquad (15)$$

where $\delta(\tau)$ is given by (11). Notice that, although (10) and (15) have a similar structure, the two formulas are not equivalent because the VIX appears in levels in (10). In addition, the CTOU model can equivalently be expressed as the "superposition" (i.e., sum) of two LOU factors. Therefore, the structure in (12) is not restrictive for the CTOU.[14]

- Jumps (LOUJ):

$$d\log V(t) = \kappa[\bar{\theta} - \log V(t)]dt + \sigma dW_v^Q(t) + dZ(t) - \frac{\lambda}{\kappa\delta}dt, \qquad (16)$$

where $Z(t)$ is a pure jump process independent of $W_v^Q(t)$, 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 (16) 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 models for instantaneous volatility have been previously considered by Duffie, Pan, and Singleton (2000) and Eraker, Johannes, and Polson (2003), among others, while Todorov and Tauchen (2011) also consider jumps in modeling the VIX. Unlike pure diffusions, this model allows for sudden movements in volatility indexes, which nevertheless have lasting effects

---

[13] As in the CSQR model, the convergence of (15) to the long-run mean is not necessarily a monotonic function of τ.

[14] The "component" GARCH model proposed by Ding and Granger (1996) to capture "long memory" features of volatility (see also Engle and Lee, 1999) can be regarded as a discrete time analogue to (12).

due to the fact that the mean-reversion parameter κ is bounded. Again, we assume (13) for the price of diffusion risk.

- Stochastic volatility (LOUSV):

$$d \log V(t) = \kappa[\bar{\theta} - \log V(t)]dt + \sqrt{\omega(t)} dW_v^Q(t),$$

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,

$$d\omega(t) = -\bar{\lambda}\omega(t)dt + d\bar{Z}(t), \tag{17}$$

where $\bar{Z}(t)$ is a pure jump process with intensity $\bar{\lambda}$ and $Exp(\bar{\delta})$ jump amplitude, while $W_v^Q(t)$ 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

$$E^Q\left[\omega(T))|\omega(t)\right] = \bar{\delta}^{-1} + \exp(-\bar{\lambda}\tau)\left[\omega(t) - \bar{\delta}^{-1}\right]. \tag{18}$$

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. Another nontrivial advantage of this model over other alternatives such as a square root process for $\omega(t)$ is that it allows the valuation of derivatives by inverting the conditional characteristic function (see Appendix Appendix C for details). Once again, we consider a price of diffusion risk proportional to instantaneous volatility to ensure that the log VIX under the real measure remains an affine process (see Dai and Singleton, 2000). Specifically, we assume

$$dW_v^Q(t) = dW_v^P(t) + \varsigma_\omega \sqrt{\omega(t)} dt. \tag{19}$$

Under this specification, the price of risk has a stronger impact when the volatility of the VIX is larger.

- Central tendency and jumps (CTOUJ):

$$d \log V(t) = \kappa[\theta(t) - \log V(t)]dt + \sigma dW_v^Q(t) + dZ(t) - \frac{\lambda}{\kappa\delta}dt, \tag{20}$$

where $\theta(t)$ follows the diffusion (12) and $Z(t)$ is a pure jump process with intensity λ and $Exp(\delta)$ jump amplitude, while $W_v^Q(t)$ and $W_\theta^Q(t)$ are independent Brownian motions. We again introduce a constant shift in (20) to ensure that $\bar{\theta}$ is the long-run mean of both $\theta(t)$ and $\log V(t)$. As in the CTOU model, we assume that the prices of risk are given by (13) and (14).

- Central tendency and stochastic volatility (CTOUSV):

$$d \log V(t) = \kappa[\theta(t) - \log V(t)]dt + \sqrt{\omega(t)} dW_v^Q(t),$$

where $\theta(t)$ follows the diffusion (12) while $\omega(t)$ is defined in (17). As in previous cases, the jump variable $\bar{Z}(t)$ and the Brownian motions are mutually independent. Similarly, we consider the prices of risk specifications (14) and (19).

Despite the apparent differences between (13), (14), and (19), they are all special cases of a generic specification. In particular, following Cheridito, Filipović, and Kimmel (2007), we can write the risk-neutral probability measure in terms of the real measure for all the extensions that we consider as

$$q = \exp\left[-\int_t^T \left(\Lambda_v(s)dW_v(s) + \Lambda_\theta(s)dW_\theta(s)\right) - \frac{1}{2}\int_t^T \left(\Lambda_v^2(s) + \Lambda_\theta^2(s)\right) ds\right],$$

where $\Lambda_v(t) = \varsigma_\omega \sqrt{\omega(t)}$ and $\Lambda_\theta(t) = \bar{\varsigma}$ in our case.[15]

Except for the CTOU 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 of the extensions to the LOU model (see Appendix Appendix C). However, this approach is not valid for the CSQR model, where we need to follow recent results by Amengual and Xiu (2012) to invert the characteristic function (see Appendix Appendix B).

Some of the previous extensions introduce as additional factors a time-varying tendency, a time-varying volatility, or both, which we need to filter

[15] Note that $\omega(t) = \sigma^2$ in the extensions without stochastic volatility. Accordingly, we define $\varsigma = \varsigma_\omega \sigma$ in those cases.

out for estimation purposes. In this regard, we use the standard Kalman filter to "estimate" $\theta(t)$. And although the jump variable $Z(t)$ in (16) can also be interpreted as an additional factor, its impact cannot be separately identified from the impact of the diffusion shocks $dW_v^Q(t)$ because both variables share the same mean-reversion coefficient κ. Finally, following Trolle and Schwartz (2009) and others, we employ the extended Kalman filter to deal with $\omega(t)$ (see once again Appendix Appendix C for further details).

5.2. Empirical performance

Table 3 reports the in- and out-of-sample RMSEs of the extensions introduced in the previous section. As expected, the CSQR process is able to yield much smaller RMSEs than the SQR and LOU models, both in- and out-of-sample. Nevertheless, the CTOU model achieves a slightly superior fit, especially out-of-sample. In any case, a persistent time-varying mean seems to capture a crucial feature of the data. On the other hand, the introduction of jumps in the LOU model does not introduce improvements in the aggregate RMSEs. Similarly, adding jumps to the CTOU model does not substantially improve the fit either. However, it is important to emphasise that Table 3 does not assess the importance of jumps on the historical dynamics of the VIX, only their pricing implications.

In contrast, we find significant improvements when we consider stochastic volatility. The LOUSV introduces important reductions in the RMSEs, but the combination of central tendency and stochastic volatility provided by the CTOUSV model yields the overall best fit. Importantly, this model is able to describe the out-of-sample behaviour of the VIX during the more extreme periods of the financial crisis.

Fig. 4 compares the empirical CDFs of the square pricing errors of futures and option prices separately. This figure shows that the CTOUSV model dominates in the first-order stochastic sense all the other models for both futures and options. The ordering of the remaining models, though, depends on the type of derivative asset considered. In the case of futures, the CTOUSV is closely followed by the CSQR and then by the other central tendency specifications (CTOU and CTOUJ), which display almost identical results. In turn, they are followed by the LOUSV, LOUJ, LOU, and SQR models. But for options, the second-best model is LOUSV, which is followed by the CTOUJ, LOUJ, CTOU, LOU, SQR, and CSQR models. Thus, we observe that the 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 CTOUSV model, while stochastic volatility does not cause any distortions to futures prices. This is not surprising, given that we show in Appendix Appendix C that option prices for the extensions of the LOU model do not depend on $\theta(t)$ once we condition on the current futures price. As for jumps, Fig. 4 shows that they do indeed help in pricing options but they hardly provide any improvement for futures. The small impact of jumps on futures is again to be expected because we can also show that jumps in the LOUJ model generate identical futures prices as a LOU model up to first order.[16] And, compared to stochastic volatility, jumps seem to yield a minor improvement even for options. In this sense, Bakshi, Cao, and Chen (1997) find that, once stochastic volatility is modeled, adding jumps only leads to second-order pricing improvements. Lastly, the SQR model offers the worst overall results, while the CSQR model also yields a poor fit for option prices.

Fig. 5a 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 SQR, LOU, and CTOU models. By focusing on a constant maturity, we can not only compare the absolute magnitude of the pricing errors, but also their sign and persistence. As can be seen, the pricing errors of the SQR process are larger than those of the LOU process, especially after mid-2007. In turn, those of the LOU model are not only substantially larger than those of the CTOU, but they also display much stronger persistence. For instance, the SQR and LOU models systematically underprice futures from October 2008 until the end of the sample. The slower mean-reverting properties of the VIX are probably responsible for these persistent biases in the one-factor models. We can also observe in Fig. 5b that the pricing errors of the CSQR, CTOU, and CTOUSV models are almost identical, albeit with slightly smaller oscillations for the CTOUSV model. This feature confirms once again that central tendency is the most relevant extension for pricing futures.

Table 4 reports the parameter estimates that we obtain under the real

[16] Formally, when we consider a Taylor expansion of the futures price formula of model LOUJ around λ/δ for $\lambda/\delta > 0$ but small, we only observe deviations from the LOU expression for the second and higher terms. This probably reflects the fact that the conditional mean of the VIX is essentially unaffected by the presence of jumps although they alter skewness and kurtosis.

measure in the extensions to the SQR and LOU models we are considering. In the models with central tendency, we observe fast mean-reversion of the VIX to $\theta(t)$, which in turn mean-reverts rather more slowly to its long-run mean $\bar{\theta}$ (i.e., $\kappa \gg \bar{\kappa}$). Importantly, the estimates of these parameters are very stable in models with central tendency. In contrast, jump intensities vary substantially depending on the sample period considered for estimation. Specifically, we obtain smaller values of λ when we include the crisis period. For the full sample, we estimate around five jumps per year in the LOUJ specification, and almost seven jumps per year in the CTOUJ model. 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 15 in the LOUSV model to two in the CTOUSV extension. Second, the deviations of $\omega(t)$ from its long-run mean are more persistent in the CTOUSV case because mean-reversion is slower the smaller $\bar{\lambda}$ is, as (18) 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. Risk pricing

Table 5 shows the point estimates of the price of risk parameters for our preferred model. We also conduct Wald tests to assess the statistical significance of these estimates. The price of risk related to the VIX Eq. (19), ς_ω, is highly significant in the two samples. In contrast, $\bar{\varsigma}$, which is the price of risk related to the innovations in (12), is insignificant in both cases.

The negative sign of ς_ω implies a more adverse distribution in the Q measure, because the VIX mean-reverts towards a higher long-run level than in the P measure. This result, which is consistent with the evidence found by Egloff, Leippold, and Wu (2010), should be taken into account in inferring the market expectation about future values of the VIX from its derivatives.

The difference between the properties of the VIX under the real and risk-neutral measures implies, within the framework of our model, that an economically important systematic risk is priced in the market of VIX derivatives. In addition, the statistical significance of the estimate of ς_ω implies that the impact of the price of risk is larger when the VIX is more volatile (see Eq. 19).

5.4. Evolution of the factors

Fig. 6a shows that the filtered values of the central tendency factor $\theta(t)$ are rather insensitive to the particular specification that we use. We have also found that our preferred CTOUSV model generates almost indistinguishable filtered values for $\theta(t)$ for the two estimation samples that we consider. Not surprisingly, Fig. 6a also confirms 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.

Fig. 6b compares the filtered instantaneous volatilities of the LOUSV and CTOUSV models. Interestingly, both series display an almost identical pattern, with substantial persistent oscillations over time. This figure also shows that large increases in the value of the VIX are associated with volatile periods, as measured by $\omega(t)$. This is particularly visible in August 2007, October 2008, and May 2010. As with central tendency, we have found that the filtered values of $\omega(t)$ obtained from the CTOUSV model are also quite insensitive to the sample period used to estimate the model parameters.

We compute the 30-day-ahead standard deviations of $\log V(t)$ implied by the CTOUSV model to assess the extent to which the filtered values of $\omega(t)$ make sense. In Fig. 6c we compare those standard deviations with the Black (1976) implied volatilities of at-the-money options that are exactly 30 days 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.

Finally, Fig. 6d compares the one-day-ahead standard deviations under the real measure implied by the models with the standard deviations obtained from a discrete ARMA(2,1)-GARCH(1,1) model for the log VIX. We only consider the two most relevant extensions to the LOU model to avoid cluttering the picture. As we have already mentioned in Section 4, both the SQR and LOU models yield time-varying variances for the VIX. However, the variance implied by the SQR process seems to be too high until mid-2007 and too low afterwards. The LOU model performs better over tranquil periods, but it underestimates actual volatility levels in the most severe phases of the financial crisis. In contrast, the CTOUJ and especially the CTOUSV do a much better job, while the CSQR significantly overestimates the standard deviations during low volatility periods.

5.5. 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 6 shows the RMSEs of futures contracts for different ranges of maturity. We observe that the models without central tendency tend to yield larger distortions for longer maturities. In contrast, the models with central tendency are relatively worse at pricing the shortest maturity. Nevertheless, even the worst RMSE of the CTOUSV model is still much smaller than the best RMSE of either the SQR or the LOU models without central tendency.

To gain some additional insight, in Fig. 7 we look at the term structure of futures prices for four particularly relevant days. Futures prices were extremely low on June 21, 2007, even though the first warning signals about the impending crisis were starting to appear. Prices had already risen significantly by August 15, 2008, one month before the collapse of Lehman Brothers. Nevertheless, they 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 significantly came down until the beginning of April 2010, right before the European sovereign debt crisis. The figure confirms that a central tendency is crucial for the purposes of reproducing the changes in the level and slope of the actual term structure of VIX futures prices. We can also observe that the LOU and LOUJ processes yield futures prices which are almost identical, while the CTOUJ model can be barely distinguished from the CTOU extension.

Tables 7 and 8 provide the RMSEs of calls and puts for different ranges of maturity and moneyness. In this case, we generally observe the highest price distortions for at-the-money call and put options. The SQR, CSQR, LOU, LOUJ, CTOU, and CTOUJ models seem to do a better job for moneyness smaller than -0.3. However, once again we can confirm that stochastic volatility models provide the best fit uniformly across all moneyness and maturity ranges.

In Fig. 8 we assess the ability of the different models to fit the average implied volatility skews in Figs. 3b and 3c. Once again, we only plot the most relevant extensions to the LOU model to avoid cluttering the pictures. In addition to the average implied skews of actual prices, we also consider 5% and 95% percentiles as a measure of dispersion. The figures confirm that the LOU model yields a constant implied volatility for all strikes. In contrast, the SQR and CSQR yield volatility skews but with a negative slope, which is inconsistent with the positive slope in the data. In contrast, both the CTOUJ and CTOUSV models are able to reproduce this positive slope, with the latter generally providing the best fit. Specifically, the CTOUJ model does not capture the shifts of the implied skews due to rises in volatility and it also performs poorly for low strikes.

6. Conclusions

We carry out an extensive empirical analysis of VIX derivatives valuation models. We consider daily prices of futures and European options from February 2006 until December 2010. Therefore, we not only cover an unusually tranquil period, but also the early turbulences that took place between August 2007 and August 2008, the worst months of the recent financial crisis (autumn 2008), as well as the months in 2010 in which the European sovereign debt crisis unfolded. These markedly different periods provide a very useful testing ground to assess the empirical performance of the different pricing models. We estimate the models using not only futures and options data, but also historical data on the VIX itself, which allows us to look at the relationship between real and risk-neutral measures. We initially focus on the two existing mean-reversion models: the square root (SQR) and the log-normal Ornstein-Uhlenbeck processes (LOU). Although SQR is more popular in the empirical literature, we find that the LOU model yields a better fit, especially during the crisis. 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 prices, or indeed the volatility skews. Part of the problem is that these models implicitly assume that volatility either mean-reverts at a simple exponential rate or does not mean-revert at all, which cannot accommodate the long and persistent swings of the VIX observed in 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. In addition, these models are also inconsistent with the strong presence of GARCH-type heteroskedasticity that we find in this volatility index.

We investigate the potential sources of mispricing by considering several empirically relevant generalisations. In particular, we consider the concatenated SQR process (CSQR), which substantially extends the SQR model by introducing a time-varying central tendency and allowing for unspanned stochastic volatility. We also extend the LOU model by introducing 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 CTOU 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 for pricing futures, regardless of whether the model is expressed in levels or logs. We also find evidence of time-varying volatility in the VIX. As expected, stochastic volatility plays a much more important role for options while leaving futures prices almost unaffected. The CSQR model, though, seems unable to generate the positive slope of the option-implied volatility skews. It is also worth mentioning that jumps only provide a minor improvement for options and do not change futures prices (up to first order). Nevertheless, we would like to emphasise that jumps seem to be a relevant feature to describe the historical dynamics of the VIX, even though they only yield second-order gains for pricing VIX derivatives.

Importantly, our results remain valid when we focus exclusively on the out-of-sample performance with parameters estimated using data prior to the autumn of 2008. In view of these findings, we conclude that a generalised LOU 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 positive slope of the implied volatility skews of options. Interestingly, the price of risk of this specification is highly significant and implies that the VIX mean-reverts towards a higher long-run mean under the risk-neutral measure than under the real measure. The difference between the properties of the VIX under the real and risk-neutral measures implies that an economically important systematic risk is priced in the market of

VIX derivatives.

Given the relationship between the observed VIX index and the unobserved integrated volatility of the S&P500, our analysis also has important implications for the models and stationarity of this broad stock index commonly used by participants in markets for stock index options. Specifically, our results imply that stochastic volatility models for the S&P500 should allow for slow mean-reversion by including two volatility factors, and a time-varying volatility of volatility. Amengual (2009) considers such a model for volatility swaps.

We could extend our empirical exercise to other recently introduced volatility derivatives such as binary options or American options on VIX futures, or even the futures and options on the CBOE Gold ETF Volatility Index. Sophisticated filtering procedures for the volatility of VIX might also be worth exploring, as well as tractable ways of modeling jump and stochastic volatility risk for Lévy processes. It would also be interesting to investigate the incremental information content of the contemporaneous observations of the S&P500 over and above the spot VIX for pricing VIX futures, and above and beyond VIX futures for pricing VIX options. This question is also relevant for the purposes of integrating the valuation of VIX derivatives with the valuation of the underlying options on the S&P500 that are used to compute this volatility index, as suggested by Lin and Chang (2009). We plan to address these points in subsequent research.

Appendix A. Affine model conditional characteristic function

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

$$d\mathbf{Y}(t) = \mathbf{K}(\mathbf{\Theta} - \mathbf{Y}(t))dt + \sqrt{\mathbf{S}(t)}d\mathbf{W}(t) + d\mathbf{Z}(t), \tag{A.1}$$

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

Appendix B. The CSQR model

Given the prices of risk (8) and (9), we can write the CSQR process under the real measure as

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

with

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

where $\kappa = \kappa^P + \sigma\varsigma$, $\bar{\kappa} = \bar{\kappa}^P + \bar{\sigma}\bar{\varsigma}$,

$$\begin{aligned}\bar{\theta} &= \frac{\kappa^P \bar{\kappa}^P}{(\kappa^P + \sigma\varsigma)(\bar{\kappa}^P + \bar{\sigma}\bar{\varsigma})}\bar{\theta}^P, \\ \bar{\sigma} &= \bar{\sigma}^P \sqrt{\frac{\kappa^P}{\kappa^P + \sigma\varsigma}},\end{aligned}$$

and

$$\theta(t) = \frac{\kappa^P}{\kappa^P + \sigma\varsigma}\theta^P(t).$$

Let $\mathbf{X}(t) = [\theta(t), V(t)]'$. Following Fackler (2000), it can be shown that the conditional mean and variance of $X(T)$ given information known at time t is the affine function of $\mathbf{X}(t)$:

$$\begin{bmatrix} E[X(T)|I(t)] \\ vec[V[X(T)|I(t)]] \end{bmatrix} = \begin{bmatrix} \mathbf{m}_0 \\ \mathbf{v}_0 \end{bmatrix} + \begin{bmatrix} \mathbf{M}_1 \\ \mathbf{V}_1 \end{bmatrix} \mathbf{X}(t),$$

where $\mathbf{m}_0$, $\mathbf{v}_0$, $\mathbf{M}_1$ and $\mathbf{V}_1$are 2×1, 4×1, 2×2 and 4×2 vectors and matrices, respectively, such that

$$\begin{aligned}\begin{bmatrix} \mathbf{m}_0 \\ \mathbf{v}_0 \end{bmatrix} &= \left[\exp(\tau\mathbf{A}) - \mathbf{I}_6\right]\mathbf{A}^{-1}\mathbf{a}, \\ \begin{bmatrix} \mathbf{M}_1 \\ \mathbf{V}_1 \end{bmatrix} &= \exp(\tau\mathbf{A}) \begin{bmatrix} \mathbf{I}_2 \\ \mathbf{0}_{4\times2} \end{bmatrix},\end{aligned}$$

$\mathbf{a}$ is a 6×1 vector whose first element is $\bar{\kappa}\bar{\theta}$ and all the other elements are zero; $\mathbf{I}_6$ is the identity matrix of order 6; $\mathbf{0}_{4\times2}$ is a 4×2 matrix of zeros; and

$$\mathbf{A} = \begin{bmatrix} \mathbf{R} & \mathbf{0}_{2\times4} \\ \mathbf{\Sigma} & (\mathbf{R}\otimes\mathbf{I}_2) + (\mathbf{I}_2\otimes\mathbf{R}) \end{bmatrix},$$

with

$$\mathbf{R} = \begin{bmatrix} -\bar{\kappa} & 0 \\ \kappa & -\kappa \end{bmatrix}, \quad \mathbf{\Sigma} = \begin{bmatrix} \bar{\sigma}^2 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & \sigma^2 \end{bmatrix}.$$

It can be tediously shown that $\mathbf{m}_0 + \mathbf{M}_1\mathbf{X}(t)$ yields (10).

Using the results form Appendix Appendix A, we can express the conditional characteristic function of this model as

$$\begin{aligned}\phi_{CSQR}(\tau, u) &= E\left[\exp((\alpha + iu)V(T))|I(t)\right] \\ &= \exp\left[\varphi_{CSQR,0}(\tau) + \varphi_{CSQR,\theta}(\tau)\theta(t) + \varphi_{CSQR,V}(\tau)V(t)\right],\end{aligned}$$

where

$$\begin{aligned}
\dot{\varphi}_{CSQR,\theta}(\tau) &= -\bar{\kappa}\varphi_{CSQR,\theta}(\tau) + \kappa\varphi_{CSQR,V}(\tau) + \frac{1}{2}\bar{\sigma}^2\varphi^2_{CSQR,\theta}(\tau), \\
\dot{\varphi}_{CSQR,V}(\tau) &= -\kappa\varphi_{CSQR,V}(\tau) + \frac{1}{2}\sigma^2\varphi^2_{CSQR,V}(\tau), \\
\dot{\varphi}_{CSQR,0}(\tau) &= \bar{\kappa}\bar{\theta}\varphi_{CSQR,\theta}(\tau),
\end{aligned}$$

with the conditions $\varphi_{CSQR,\theta}(0) = 0$, $\varphi_{CSQR,V}(0) = \alpha + iu$, and $\varphi_{CSQR,0}(0) = 0$. Furthermore, it can be shown that

$$\varphi_{CSQR,V}(t) = \frac{(\alpha + iu)\exp(-\kappa\tau)}{1 - (\alpha + iu)\dfrac{\sigma^2(1 - \exp(-\kappa\tau))}{2\kappa}}.$$

Then, we can follow Amengual and Xiu (2012) in showing that the price of a European call option with strike K can be expressed as

$$c(t, T, K) = \frac{\exp(-r\tau)}{\pi}\int_0^\infty \mathrm{Re}\left[\phi_{CSQR}(\tau, u)\frac{\exp[-K(\alpha + iu)]}{(\alpha + iu)^2}\right] du,$$

where the smoothing parameter α must be such that

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

Appendix C. Extensions of LOU processes

Appendix C.1. General case

If we place $\log V(t)$ as the first element of $\mathbf{Y}(t)$ in (A.1), 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

$$c(t, T, K) = \frac{\exp(-\alpha\log(K))}{\pi}\int_0^\infty \exp(-iu\log(K))\psi(u)du, \qquad \text{(C.1)}$$

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 (C.1) 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:

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

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 RMSEs of options, is very small (below 0.25%). In any case, we calculate the exact pricing errors once we have obtained the final parameter estimates.

Appendix C.2. Central tendency

Given the prices of risk (13) and (14), the diffusions under the real measure can be expressed as

$$\begin{aligned} d\log V(t) &= \kappa[\theta^P(t)-\log V(t)]dt+\sigma dW_v^P(t), \\ d\theta^P(t) &= \bar{\kappa}[\bar{\theta}^P-\theta^P(t)]dt+\bar{\sigma}dW_\theta^P(t), \end{aligned}$$

where

$$\bar{\theta}^P=\bar{\theta}+\frac{\sigma\varsigma}{\kappa}+\frac{\bar{\sigma}\bar{\varsigma}}{\bar{\kappa}},$$

and

$$\theta^P(t)=\theta(t)+\frac{\sigma\varsigma}{\kappa}.$$

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 $\mu_{CTOU}(t,\tau)$ given in (15) and variance

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

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

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

while call prices follow the Black (1976) formula with volatility $\varphi_{CTOU}(\tau)$. This confirms that the prices of options do not depend on $\theta(t)$ once we condition on the futures price.

In terms of time series dynamics, it can be shown that $\theta(t)$ and $\log V(t)$ jointly follow a Gaussian first-order vector autoregressive process (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 $\boldsymbol{\Sigma}_\tau(2,2) = \varphi^2_{CTOU}(\tau)$. 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

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

Appendix C.3. Jumps

Once again, we assume that $dW_v^Q(t) = dW_v^P(t) + \varsigma dt$ holds and that jump risk is not priced. Then, it can be shown that $\log V(t)$ satisfies the diffusion

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

under the real measure, where $\bar{\theta}^P = \bar{\theta} + \sigma\varsigma/\kappa$.

The conditional characteristic function reduces to

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

where

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

and

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

Appendix C.4. Stochastic volatility

We consider a price of risk such that $dW_v^Q(t) = dW_v^P(t) + \varsigma_\omega\sqrt{\omega(t)}dt$, but we again assume that jump risk related to $Z(t)$ is not priced. Then, the process under the real measure can be expressed as

$$\begin{aligned} d\log V(t) &= \kappa\left[\bar{\theta} + \frac{\varsigma_\omega}{\kappa}\omega(t) - \log V(t)\right]dt + \sqrt{\omega(t)}dW_v^P(t), \\ d\omega(t) &= -\bar{\lambda}\omega(t)dt + d\bar{Z}(t). \end{aligned}$$

Since in this case the processes under the real and risk-neutral measures are different, we will describe them separately.

Appendix C.4.1. Risk-neutral measure

The conditional characteristic function simplifies to

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

where

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

and

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

with

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

Appendix C.4.2. Real measure

We can write the characteristic function under the real measure as

$$\begin{gathered}\phi^P_{LOUSV}(t,T,u) = E^P\left[\exp(iu\log V(T)+iv\omega(T))|V(t),\omega(t)\right] \\ = \exp\left[\varphi^P_{LOUSV,0}(\tau)+\varphi^P_{LOUSV,V}(\tau)\log V(t)+\varphi^P_{LOUSV,\omega}(\tau)\omega(t)\right],\end{gathered}$$

where

$$\begin{gathered}\varphi^P_{LOUSV,0}(\tau) = iu\bar{\theta}[1-\exp(-\kappa\tau)]+\bar{\lambda}[\varkappa(\tau)-\tau], \\ \varphi^P_{LOUSV,V}(\tau) = iu\exp(-\kappa\tau),\end{gathered}$$

$$\begin{aligned}\varphi^P_{LOUSV,\omega}(\tau) &= iv\exp(-\bar{\lambda}\tau) - iu\frac{\varsigma_\omega}{\kappa-\bar{\lambda}}\left[\exp(-\kappa\tau)-\exp(-\bar{\lambda}\tau)\right] \\ &\quad +\frac{u^2}{2(2\kappa-\bar{\lambda})}\left[\exp(-2\kappa\tau)-\exp(-\bar{\lambda}\tau)\right]\end{aligned}$$

and

$$\varkappa(\tau) = \int_0^\tau \frac{\bar{\delta}}{\left[\begin{array}{c}\bar{\delta}-iv\exp(-\bar{\lambda}x)+iu\frac{\varsigma_\omega\left[\exp(-\kappa x)-\exp(-\bar{\lambda}x)\right]}{\kappa-\bar{\lambda}} \\ -\frac{u^2\left[\exp(-2\kappa x)-\exp(-\bar{\lambda}x)\right]}{2(2\kappa-\bar{\lambda})}\end{array}\right]}dx.$$

Based on the characteristic function, it is possible to show that

$$
\begin{aligned}
E\left[\log V(T)|V(t),\omega(t)\right] &= \bar{\theta}[1-\exp(-\kappa\tau)] \\
&\quad -\frac{\bar{\lambda}\varsigma_\omega}{\bar{\delta}(\kappa-\bar{\lambda})}\left[\frac{1-\exp(-\kappa\tau)}{\kappa}-\frac{1-\exp(-\bar{\lambda}\tau)}{\bar{\lambda}}\right] \\
&\quad +\exp(-\kappa\tau)\log V(t) \\
&\quad -\frac{\varsigma_\omega}{\kappa-\bar{\lambda}}\left[\exp(-\kappa\tau)-\exp(-\bar{\lambda}\tau)\right]\omega(t) \\
E\left[\omega(T)|V(t),\omega(t)\right] &= \frac{1}{\bar{\delta}}[1-\exp(-\bar{\lambda}\tau)]+\exp(-\bar{\lambda}\tau)\omega(t),
\end{aligned}
$$

$$
V\left[\log V(T)|V(t),\omega(t)\right] = \frac{\bar{\lambda}}{(2\kappa-\bar{\lambda})\bar{\delta}}\left[\frac{1-\exp(-\bar{\lambda}\tau)}{\bar{\lambda}}-\frac{1-\exp(-2\kappa\tau)}{2\kappa}\right]
$$

$$
+\frac{2\varsigma_\omega^2\bar{\lambda}}{\bar{\delta}^2(\kappa-\bar{\lambda})^2}\left[\frac{1-\exp(-2\kappa\tau)}{2\kappa}+\frac{1-\exp(-2\bar{\lambda}\tau)}{2\bar{\lambda}}-2\frac{1-\exp(-(\kappa+\bar{\lambda})\tau)}{\kappa+\bar{\lambda}}\right]
$$

$$
+\frac{1}{2\kappa-\bar{\lambda}}\left[\exp(-\bar{\lambda}\tau)-\exp(-2\kappa\tau)\right]\omega(t),
$$

$$
V^P\left[\omega(T)|V(t),\omega(t)\right] = \frac{1-\exp(-2\bar{\lambda}\tau)}{\bar{\delta}^2},
$$

and

$$
cov\left[\log V(T),\omega(T)|V(t),\omega(t)\right] = -\frac{2\varsigma_\omega\bar{\lambda}}{\bar{\delta}^2(\kappa-\bar{\lambda})}\left[\frac{1-\exp(-(\kappa+\bar{\lambda})\tau)}{\kappa+\bar{\lambda}}-\frac{1-\exp(-2\bar{\lambda}\tau)}{2\bar{\lambda}}\right].
$$

Appendix C.5. Central tendency and jumps

For this case, we use the same prices of risk as in the CTOU model and assume that jump risk is not priced. Then, we obtain

$$
\begin{aligned}
d\log V(t) &= \kappa[\theta^P(t)-\log V(t)]dt+\sigma dW_v^P(t)+dZ(t)-\frac{\lambda}{\kappa\delta}dt, \\
d\theta^P(t) &= \bar{\kappa}[\bar{\theta}^P-\theta^P(t)]dt+\bar{\sigma}dW_\theta^P(t),
\end{aligned}
$$

under the real measure, where

$$
\bar{\theta}^P = \bar{\theta}+\frac{\sigma\varsigma}{\kappa}+\frac{\bar{\sigma}\bar{\varsigma}}{\bar{\kappa}},
$$

and

$$
\theta^P(t) = \theta(t)+\frac{\sigma\varsigma}{\kappa}.
$$

The conditional characteristic function becomes

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

where

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

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

and

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

Using the characteristic function, we can show that

$$E\left[\log V(T)|V(t),\theta(t)\right] = \bar{\theta}\left[\begin{array}{c}1-\exp(-\kappa\tau) \\ -\frac{\kappa}{\kappa-\bar{\kappa}}\left(\exp(-\bar{\kappa}\tau)-\exp(-\kappa\tau)\right)\end{array}\right]$$
$$+\frac{\kappa}{\kappa-\bar{\kappa}}\left[\exp(-\bar{\kappa}\tau)-\exp(-\kappa\tau)\right]\theta(t)+\exp(-\kappa\tau)\log V(t),$$

$$V\left[\log V(T)|V(t),\theta(t)\right] = \left(\frac{\sigma^2}{2\kappa}+\frac{\lambda}{\kappa\delta^2}\right)\left[1-\exp(-2\kappa\tau)\right]$$
$$+\bar{\sigma}^2\left(\frac{\kappa}{\kappa-\bar{\kappa}}\right)^2\left[\begin{array}{c}\frac{1-\exp(-2\bar{\kappa}\tau)}{2\bar{\kappa}}+\frac{1-\exp(-2\kappa\tau)}{2\kappa} \\ -2\frac{1-\exp(-(\kappa+\bar{\kappa})\tau)}{\kappa+\bar{\kappa}}\end{array}\right],$$

$$V\left[\log V(T),\theta(T)|V(t),\theta(t)\right]=\frac{\kappa\bar{\sigma}^2}{\kappa-\bar{\kappa}}\left[\begin{array}{c}\frac{1-\exp(-2\bar{\kappa}\tau)}{2\bar{\kappa}}\\-\frac{1-\exp(-(\kappa+\bar{\kappa})\tau)}{\kappa+\bar{\kappa}}\end{array}\right].$$

Appendix C.6. Central tendency and stochastic volatility

The price of risk is such that $dW_v^Q(t)=dW_v^P(t)+\varsigma_\omega\sqrt{\omega(t)}dt$ and $dW_\theta^Q(t)=dW_\theta^P(t)+\bar{\varsigma}dt$, which yields

$$\begin{aligned} d\log V(t) &= \kappa\left[\theta(t)+\frac{\varsigma_\omega}{\kappa}\omega(t)-\log V(t)\right]dt+\sqrt{\omega(t)}dW_v^P(t),\\ d\theta(t) &= \bar{\kappa}[\bar{\theta}^P-\theta(t)]dt+\bar{\sigma}dW_\theta^P(t),\\ d\omega(t) &= -\bar{\lambda}\omega(t)dt+d\bar{Z}(t), \end{aligned}$$

under the real measure, where

$$\bar{\theta}^P=\bar{\theta}+\frac{\bar{\varsigma}\bar{\sigma}}{\bar{\kappa}}.$$

Appendix C.6.1. Risk-neutral measure

The conditional characteristic function is

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

where

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

and

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

with $\varkappa(\tau,u)$ defined in (C.2).

Appendix C.6.2. Real measure

We can write the characteristic function under the real measure as

$$\phi^P(t,T,u_1,u_2,u_3) = E\left[\exp\left[\begin{array}{c} iu_1 \log V(T) + iu_2\theta(T) \\ +iu_3\omega(T) \end{array}\right]\middle| V(t),\theta(t),\omega(t)\right]$$
$$= \exp\left[\begin{array}{c} \varphi^P_{CTOUSV,0}(\tau) + \varphi^P_{CTOUSV,V}(\tau)\log V(t) \\ +\varphi^P_{CTOUSV,\theta}(\tau)\theta(t) + \varphi^P_{CTOUSV,\omega}(\tau)\omega(t) \end{array}\right],$$

where

$$\begin{gathered}
\varphi^P_{CTOUSV,0}(\tau) = iu_2\bar{\theta}(1-\exp(-\bar{\kappa}\tau)) \\
+iu_1\bar{\theta}\frac{\kappa\bar{\kappa}}{\kappa-\bar{\kappa}}\left[\frac{1-\exp(-\bar{\kappa}\tau)}{\bar{\kappa}} - \frac{1-\exp(-\kappa\tau)}{\kappa}\right] \\
-u_2^2\frac{\bar{\sigma}^2}{2}\frac{1-\exp(-2\bar{\kappa}\tau)}{2\bar{\kappa}} \\
-u_1^2\frac{\bar{\sigma}^2}{2}\left(\frac{\kappa}{\kappa-\bar{\kappa}}\right)^2\left[\begin{array}{c} \frac{1-\exp(-2\bar{\kappa}\tau)}{2\bar{\kappa}} + \frac{1-\exp(-2\kappa\tau)}{2\kappa} \\ -2\frac{1-\exp(-(\kappa+\bar{\kappa})\tau)}{\kappa+\bar{\kappa}} \end{array}\right] \\
-u_1u_2\bar{\sigma}^2\frac{\kappa}{\kappa-\bar{\kappa}}\left[\frac{1-\exp(-2\bar{\kappa}\tau)}{2\bar{\kappa}} - \frac{1-\exp(-(\kappa+\bar{\kappa})\tau)}{\kappa+\bar{\kappa}}\right] \\
+\bar{\lambda}[\varkappa(\tau)-\tau],
\end{gathered}$$

$$\varphi^P_{CTOUSV,V}(\tau) = iu_1\exp(-\kappa\tau),$$

$$\varphi^P_{CTOUSV,\theta}(\tau) = iu_2\exp(-\bar{\kappa}\tau) + iu_1\frac{\kappa}{\kappa-\bar{\kappa}}\left[\exp(-\bar{\kappa}\tau)-\exp(-\kappa\tau)\right]$$

$$\begin{aligned}
\varphi^P_{CTOUSV,\omega}(\tau) &= iu_3\exp(-\bar{\lambda}\tau) - iu_1\frac{\varsigma_\omega}{\kappa-\bar{\lambda}}\left[\exp(-\kappa\tau)-\exp(-\bar{\lambda}\tau)\right] \\
&\quad +\frac{u_1^2}{2(2\kappa-\bar{\lambda})}\left[\exp(-2\kappa\tau)-\exp(-\bar{\lambda}\tau)\right]
\end{aligned}$$

and

$$\varkappa(\tau) = \int_0^\tau \frac{\delta}{\left[\begin{array}{c} \delta - iu_3\exp(-\bar{\lambda}x) + iu_1\frac{\varsigma_\omega\left[\exp(-\kappa x)-\exp(-\bar{\lambda}x)\right]}{\kappa-\bar{\lambda}} \\ -\frac{u_1^2\left[\exp(-2\kappa x)-\exp(-\bar{\lambda}x)\right]}{2(2\kappa-\bar{\lambda})} \end{array}\right]}dx.$$

Based on the characteristic function, it is possible to show that

$$
\begin{aligned}
E\left[\log V(T)|V(t),\theta(t),\omega(t)\right] = \bar{\theta}^P \frac{\kappa\bar{\kappa}}{\kappa-\bar{\kappa}}\left[\frac{1-\exp(-\bar{\kappa}\tau)}{\bar{\kappa}} - \frac{1-\exp(-\kappa\tau)}{\kappa}\right] \\
-\frac{\bar{\lambda}\varsigma_\omega}{\bar{\delta}(\kappa-\bar{\lambda})}\left[\frac{1-\exp(-\kappa\tau)}{\kappa} - \frac{1-\exp(-\bar{\lambda}\tau)}{\bar{\lambda}}\right] \\
+\exp(-\kappa\tau)\log V(t) \\
+\frac{\kappa}{\kappa-\bar{\kappa}}\left[\exp(-\bar{\kappa}\tau)-\exp(-\kappa\tau)\right]\theta(t) \\
-\frac{\varsigma_\omega}{\kappa-\bar{\lambda}}\left[\exp(-\kappa\tau)-\exp(-\bar{\lambda}\tau)\right]\omega(t) \\
E\left[\theta(T)|V(t),\theta(t),\omega(t)\right] = \bar{\theta}^P\left[1-\exp(-\bar{\kappa}\tau)\right]+\exp(-\bar{\kappa}\tau)\theta(t) \\
E\left[\omega(T)|V(t),\theta(t),\omega(t)\right] = \frac{1}{\bar{\delta}}[1-\exp(-\bar{\lambda}\tau)]+\exp(-\bar{\lambda}\tau)\omega(t),
\end{aligned}
$$

$$
\begin{aligned}
V\left[\log V(T)|V(t),\theta(t),\omega(t)\right] = \frac{\bar{\lambda}}{(2\kappa-\bar{\lambda})\bar{\delta}}\left[\frac{1-\exp(-\bar{\lambda}\tau)}{\bar{\lambda}} - \frac{1-\exp(-2\kappa\tau)}{2\kappa}\right] \\
+\bar{\sigma}^2\left(\frac{\kappa}{\kappa-\bar{\kappa}}\right)^2\left[\frac{1-\exp(-2\bar{\kappa}\tau)}{2\bar{\kappa}} + \frac{1-\exp(-2\kappa\tau)}{2\kappa} - 2\frac{1-\exp(-(\kappa+\bar{\kappa})\tau)}{\kappa+\bar{\kappa}}\right] \\
+\frac{2\varsigma_\omega^2\bar{\lambda}}{\bar{\delta}^2(\kappa-\bar{\lambda})^2}\left[\frac{1-\exp(-2\kappa\tau)}{2\kappa} + \frac{1-\exp(-2\bar{\lambda}\tau)}{2\bar{\lambda}} - 2\frac{1-\exp(-(\kappa+\bar{\lambda})\tau)}{\kappa+\bar{\lambda}}\right] \\
+\frac{1}{2\kappa-\bar{\lambda}}\left[\exp(-\bar{\lambda}\tau)-\exp(-2\kappa\tau)\right]\omega(t),
\end{aligned}
$$

$$
\begin{aligned}
V\left[\theta(T)|V(t),\theta(t),\omega(t)\right] = \bar{\sigma}^2\frac{1-\exp(-2\bar{\kappa}\tau)}{2\bar{\kappa}} \\
V\left[\omega(T)|V(t),\theta(t),\omega(t)\right] = \frac{1-\exp(-2\bar{\lambda}\tau)}{\bar{\delta}^2}, \\
cov\left[\log V(T),\theta(T)|V(t),\theta(t),\omega(t)\right] = \frac{\bar{\sigma}^2\kappa}{\kappa-\bar{\kappa}}\left[\begin{array}{c}\frac{1-\exp(-2\bar{\kappa}\tau)}{2\bar{\kappa}} \\ -\frac{1-\exp(-(\kappa+\bar{\kappa})\tau)}{\kappa+\bar{\kappa}}\end{array}\right]
\end{aligned}
$$

and

$$
cov\left[\log V(T),\omega(T)|V(t),\theta(t),\omega(t)\right] = -\frac{2\varsigma_\omega\bar{\lambda}}{\bar{\delta}^2(\kappa-\bar{\lambda})}\left[\begin{array}{c}\frac{1-\exp(-(\kappa+\bar{\lambda})\tau)}{\kappa+\bar{\lambda}} \\ -\frac{1-\exp(-2\bar{\lambda}\tau)}{2\bar{\lambda}}\end{array}\right]
$$

Appendix D. One-factor model variances and autocorrelations

Let $V(t)$ follow the SQR process. Then, it can be shown that

$$V[V(T)|V(t)] = \frac{\sigma^2}{2\kappa}(1-\exp(-\kappa\tau))^2 + \frac{\sigma^2}{\kappa}\exp(-\kappa\tau)(1-\exp(-\kappa\tau))V(t),$$

and $corr[V(T), V(t)] = \exp(-\kappa\tau)$.

Alternatively, if $V(t)$ follows the LOU process, it holds that

$$\begin{aligned} V[V(T)|V(t)] &= \exp\left[2\bar{\theta}(1-\exp(-\kappa\tau)) + \frac{\sigma^2}{2\kappa}(1-\exp(-2\kappa\tau))\right] \\ &\quad \times \left[\exp\left[\frac{\sigma^2}{2\kappa}(1-\exp(-2\kappa\tau))\right] - 1\right][V(t)]^{2\exp(-\kappa\tau)}, \end{aligned}$$

and

$$corr[V(T), V(t)] = \frac{\exp\left[\frac{\sigma^2}{2\kappa}\exp(-\kappa\tau)\right] - 1}{\exp\left[\frac{\sigma^2}{2\kappa}\right] - 1}.$$

References

Amengual, D., 2009. The term structure of variance risk premia. Unpublished working paper, CEMFI.

Amengual, D., Xiu, D., 2012. Delving into risk premia: reconciling evidence from the S&P 500 and VIX derivatives. Unpublished working paper, CEMFI and University of Chicago.

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

Andersen, T. G., Bondarenko, O., 2007. Construction and interpretation of model-free implied volatility, NBER Working Paper No. W13449.

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

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

Barndorff-Nielsen, O., Shephard, N., 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.

Bates, D. S., 2000. Post-'87 crash fears in the S&P 500 futures option market. Journal of Econometrics 94, 181–238.

Bates, D. S., 2012. U.S. stock market crash risk, 1926–2010. Journal of Financial Economics 105, 229–259.

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

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

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

Cheridito, P., Filipović, D., Kimmel, R. L., 2007. Market price of risk specifications for affine models: theory and evidence. Journal of Financial Economics 83, 123–170.

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

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

Dai, Q., Singleton, K. J., 2000. Specification analysis of affine term structure models. Journal of Finance 55, 1943–1978.

Demos, A., Sentana, E., 1998. Testing for GARCH effects: a one-sided approach. Journal of Econometrics 86, 97–127.

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

Ding, Z., Granger, C. W. J., 1996. Modeling volatility persistence of speculative returns: a new approach. Journal of Econometrics 73, 185–215.

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

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, 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 45, 1279–1310.

Engle, R. F., Lee, G. J., 1999. A permanent and transitory component model of stock return volatility. In: Engle, R. F., White, H. (Eds.), Cointegration, Causality, and Forecasting: A Festschrift in Honor of Clive W.J. Granger. Oxford University Press, New York, pp. 475–497.

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

Fackler, P. L., 2000. Moments of affine diffusions. Unpublished working paper, North Carolina State University.

French, K. R., Schwert, G. W., Stambaugh, R. F., 1987. Expected stock returns and volatility. Journal of Financial Economics 19, 3–29.

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

Hansen, L. P., Scheinkman, J. A., 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., Pennacchi, G. G., 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., Sentana, E., 1997. Pricing options on assets with predictable white noise returns, LSE FMG Discussion Papers 267.

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

Pagan, A. R., Schwert, G. W., 1990. Alternative models for conditional stock volatility. Journal of Econometrics 45, 267–290.

Rhoads, R., 2011. Trading VIX Derivatives: Trading and Hedging Strategies Using VIX Futures, Options and Exchange Traded Notes. John Wiley and Sons, New York.

RiskMetrics Group, 1996. RiskMetrics Technical Document.

Schwert, G. W., 1990. Stock returns and real activity: a century of evidence. Journal of Finance 45, 1237–1257.

Schwert, G. W., 2011. Stock volatility during the recent financial crisis. European Financial Management 17, 789–805.

Sepp, A., 2008. VIX option pricing in a jump-diffusion model. Risk Magazine, 84–89 April.

Song, Z., Xiu, D., 2012. A tale of two option markets: state-price densities implied from S&P500 and VIX option prices. Unpublished working paper, Federal Reserve Board and University of Chicago.

Szado, E., 2009. VIX futures and options: A case study of portfolio diversification during the 2008 financial crisis. Journal of Alternative Investments 12, 68–85.

Todorov, V., Tauchen, G., 2011. Volatility jumps. Journal of Business and Economic Statistics 29, 356–371.

Trolle, A. B., Schwartz, E. S., 2009. Unspanned stochastic volatility and the pricing of commodity derivatives. Review of Financial Studies 22, 4423–4461.

Wang, Z., Daigler, R. T., 2011. The performance of VIX option pricing models: empirical evidence beyond simulation. Journal of Futures Markets 31, 251–281.

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

Zhang, J. E., Zhu, Y., 2006. VIX futures. Journal of Futures Markets 26, 521–531.

Table 1
Root mean square pricing errors in existing one-factor models.
The pricing errors of futures and options on the VIX have been considered. Estimation of the parameters by maximum likelihood using data on the VIX index and its options and futures. "SQR" denotes square root model while "LOU" refers to the log-normal Ornstein-Uhlenbeck process.

	Aug08–Mar10 estimates		Full sample estimates
	Feb06–Aug08	Aug08–Dec10	Feb06–Dec10
	In-sample	*Out-of-sample*	
SQR	0.807	2.037	2.454
LOU	0.779	1.982	1.533

Table 2

Parameter estimates of the existing one-factor models.

Estimation by maximum likelihood using data on the VIX index and its options and futures from Feb 06 until Dec 10. "SQR" denotes square root model while "LOU" refers to the log-normal Ornstein-Uhlenbeck process. Standard errors, displayed in parentheses, have been obtained by using the outer-product of the score to estimate the information matrix. $\sigma_{f\varepsilon}$ and $\sigma_{f\xi}$ are the parameters of the futures pricing errors. η_o, σ_a, σ_b, and σ_c are the parameters of the option pricing errors, whose variance is the quadratic function of moneyness (3).

		Real measure parameters			Pricing error parameters					
Model	Estimation	κ	$\bar{\theta}$	σ	$\sigma_{f\varepsilon}$	$\sigma_{f\xi}$	σ_a	η_o	σ_b	σ_c
SQR	Full sample	2.555	24.350	4.504	0.122	0.081	0.006	0.497	0.437	-0.901
		(0.071)	(0.661)	(0.005)	(0.001)	(0.009)	($2.8\ 10^{-4}$)	(0.019)	(0.017)	(0.002)
	Until Aug 08	1.495	19.666	2.916	0.040	0.008	0.006	0.369	0.275	-0.727
		(0.015)	(0.150)	(0.006)	($4.1\ 10^{-4}$)	(0.001)	($2.6\ 10^{-4}$)	(0.024)	(0.009)	(0.003)
LOU	Full sample	2.898	3.099	0.884	0.075	0.030	0.006	0.455	0.207	-0.901
		(0.004)	(0.112)	($4.4\ 10^{-4}$)	(0.001)	(0.003)	($2.4\ 10^{-4}$)	(0.017)	(0.007)	(0.002)
	Until Aug 08	1.732	2.927	0.710	0.042	0.008	0.005	0.412	0.174	-0.773
		(0.009)	(0.171)	(0.001)	($4.2\ 10^{-4}$)	(0.001)	($2.2\ 10^{-4}$)	(0.021)	(0.006)	(0.005)

Table 3
Root mean square pricing errors in the extended models.
The pricing errors of futures and options on the VIX have been considered. Estimation of the parameters by maximum likelihood for the CTOU model, and by pseudo maximum likelihood in the remaining ones. The data employed in the estimation includes the VIX index and its options and futures. "LOUJ" introduces jumps in the log-normal Ornstein-Uhlenbeck process (LOU), whose size follows an exponential distribution. "CTOU" adds central tendency to the LOU process. "LOUSV" denotes a LOU model with stochastic volatility modeled with a Gamma OU Lévy process. "CTOUJ" adds central tendency and jumps to the LOU model, while "CTOUSV" introduces central tendency and stochastic volatility. "CSQR" is the concatenated square root process.

	Aug08-Mar10 estimates		Full sample estimates
	Feb06–Aug08	Aug08–Dec10	Feb06–Dec10
	In-sample	*Out-of-sample*	
CTOU	0.348	0.713	0.655
LOUJ	0.837	2.160	1.634
LOUSV	0.664	1.617	1.303
CTOUJ	0.354	0.691	0.631
CTOUSV	0.232	0.344	0.306
CSQR	0.424	1.156	0.691

Table 4
Real measure parameter estimates of the extended models.

Estimation by maximum likelihood for the CTOU model, and by pseudo maximum likelihood in the remaining ones. The data employed includes the VIX index and its options and futures from Feb 06 until Dec 10. "LOUJ" introduces jumps in the log-normal Ornstein-Uhlenbeck process (LOU), whose size follows an exponential distribution. "CTOU" adds central tendency to the LOU process. "LOUSV" denotes a LOU model with stochastic volatility modeled with a Gamma OU Lévy process. "CTOUJ" adds central tendency and jumps to the LOU model, while "CTOUSV" introduces central tendency and stochastic volatility. "CSQR" is the concatenated square root model. Standard errors, displayed in parentheses, have been obtained by using the outer-product of the score to estimate the information matrix.

Extension	Estimation	κ	$\bar{\kappa}$	$\bar{\theta}$	λ	δ	σ	$\bar{\sigma}$	$\bar{\lambda}$	$\bar{\delta}$
CTOU	Full sample	5.827	0.300	3.019			1.037	0.446		
		(0.017)	(0.003)	(0.710)			(0.001)	(0.001)		
	Until Aug08	6.946	0.422	2.851			1.012	0.447		
		(0.041)	(0.006)	(0.867)			(0.002)	(0.001)		
LOUJ	Full sample	3.632		3.093	5.370	4.870	0.728			
		(0.004)		(0.112)	(0.109)	(0.037)	(0.002)			
	Until Aug08	2.289		2.897	14.621	7.802	0.382			
		(0.008)		(0.159)	(0.721)	(0.149)	(0.010)			
LOUSV	Full sample	0.679		2.978					14.937	3.026
		(0.004)		(0.002)					(0.038)	(0.007)
	Until Aug08	0.573		2.842					15.631	3.474
		(0.005)		(0.002)					(0.126)	(0.014)
CTOUJ	Full sample	5.452	0.285	3.021	6.671	5.196	0.836	0.486		
		(0.011)	(0.003)	(0.873)	(0.071)	(0.013)	(0.002)	(0.001)		
	Until Aug08	6.094	0.360	2.861	11.382	5.962	0.640	0.450		
		(0.029)	(0.006)	(1.055)	(0.222)	(0.030)	(0.004)	(0.001)		
CTOUSV	Full sample	8.903	0.357	2.934				0.222	2.140	0.271
		(0.009)	(0.002)	(0.163)				(0.001)	(0.005)	(0.001)
	Until Aug08	10.029	0.477	2.826				0.227	0.953	0.170
		(0.038)	(0.006)	(0.209)				(0.003)	(0.013)	(0.001)
CSQR	Full sample	3.076	0.673	22.476			6.424	2.145		
		(0.009)	(0.021)	(0.645)			(0.005)	(0.007)		
	Until Aug08	2.575	0.446	19.795			3.732	1.360		
		(0.126)	(0.032)	(1.565)			(0.006)	(0.013)		

Table 5
Prices of risk in the CTOUSV model.
The estimates have been obtained by pseudo maximum likelihood using data on the VIX index and its options and futures from Feb 06 until Aug 10. Standard errors are displayed in parentheses below the estimates, and p-values are reported below the Wald tests. "CTOUSV" introduces central tendency and stochastic volatility in a log-normal Ornstein-Uhlenbeck process.

	Estimates		Wald tests		
	ς_ω	$\bar{\varsigma}$	$\varsigma_\omega = 0$	$\bar{\varsigma} = 0$	Joint
Full sample	-1.27	0.15	76.35	0.32	76.99
	(0.15)	(0.26)	(0.00)	(0.57)	(0.00)
Until Aug 08	-1.19	0.14	45.63	0.11	46.67
	(0.18)	(0.44)	(0.00)	(0.74)	(0.00)

Table 6

Root mean square pricing errors of futures prices by maturity.

Estimation of the parameters by maximum likelihood for the SQR, LOU and CTOU models, and by pseudo maximum likelihood in the remaining ones. The data employed in the estimation includes the VIX index and its options and futures from Feb 06 until Dec 10. "SQR" denotes square root model and "LOU" refers to a log-normal Ornstein-Uhlenbeck process. "LOUJ" introduces jumps in the LOU model, whose size follows an exponential distribution. "CTOU" adds central tendency to the LOU process. "LOUSV" denotes a LOU model with stochastic volatility modeled with a Gamma OU Lévy process. "CTOUJ" adds central tendency and jumps to the LOU model, while "CTOUSV" introduces central tendency and stochastic volatility. "CSQR" is the concatenated SQR model. The column labeled N gives the number of prices per category.

Maturity	N	SQR	LOU	CTOU	LOUJ	LOUSV	CTOUJ	CTOUSV	CSQR
Less than 1 month	1036	2.607	1.957	1.606	2.068	2.552	1.652	0.982	1.109
From 1 to 3 months	2290	6.128	3.953	2.046	4.334	4.332	2.115	0.844	1.029
From 3 to 6 months	3093	9.039	5.250	1.561	5.755	4.797	1.626	0.494	0.696
More than 6 months	2246	9.370	6.054	1.278	6.489	4.031	1.341	0.567	0.743
Total	8665	7.916	4.892	1.645	5.311	4.262	1.708	0.688	0.862

Table 7
Root mean square pricing errors of call prices by moneyness and maturities.
Estimation of the parameters by maximum likelihood for the SQR, LOU and CTOU models, and by pseudo maximum likelihood in the remaining ones. The data employed in the estimation includes the VIX index and its options and futures from Feb 06 until Dec 10. 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 "LOU" refers to a log-normal Ornstein-Uhlenbeck process. "LOUJ" introduces jumps in the LOU model, whose size follows an exponential distribution. "CTOU" adds central tendency to the LOU process. "LOUSV" denotes a LOU model with stochastic volatility modeled with a Gamma OU Lévy process. "CTOUJ" adds central tendency and jumps to the LOU model, while "CTOUSV" introduces central tendency and stochastic volatility. "CSQR" is a concatenated SQR model. The column labeled N gives the number of prices per category. τ denotes time to maturity.

Panel A: $\tau < 1$ month

Moneyness	N	SQR	LOU	CTOU	LOUJ	LOUSV	CTOUJ	CTOUSV	CSQR
< -0.3	2199	0.155	0.138	0.133	0.142	0.110	0.204	0.142	0.214
$[-0.3, -0.1)$	2080	0.516	0.373	0.352	0.381	0.140	0.375	0.237	0.601
$[-0.1, 0.1)$	2260	0.867	0.597	0.527	0.531	0.245	0.509	0.227	0.800
$[0.1, 0.3)$	2254	0.771	0.615	0.535	0.414	0.344	0.409	0.187	0.640
≥ 0.3	3085	0.389	0.361	0.338	0.214	0.253	0.213	0.160	0.349

Panel B: 1 month$< \tau < 3$ months

Moneyness	N	SQR	LOU	CTOU	LOUJ	LOUSV	CTOUJ	CTOUSV	CSQR
< -0.3	3553	0.250	0.176	0.177	0.174	0.151	0.253	0.214	0.390
$[-0.3, -0.1)$	4217	0.560	0.381	0.384	0.394	0.166	0.372	0.270	0.768
$[-0.1, 0.1)$	4701	0.787	0.494	0.486	0.495	0.199	0.476	0.173	0.792
$[0.1, 0.3)$	4638	0.864	0.573	0.557	0.452	0.368	0.434	0.192	0.704
≥ 0.3	8089	0.573	0.448	0.443	0.272	0.328	0.261	0.178	0.482

Panel C: $\tau > 3$ months

Moneyness	N	SQR	LOU	CTOU	LOUJ	LOUSV	CTOUJ	CTOUSV	CSQR
< -0.3	3229	0.321	0.198	0.215	0.187	0.195	0.251	0.250	0.606
$[-0.3, -0.1)$	3719	0.486	0.360	0.356	0.370	0.222	0.351	0.280	0.806
$[-0.1, 0.1)$	4491	0.614	0.451	0.426	0.430	0.254	0.407	0.256	0.741
$[0.1, 0.3)$	3971	0.860	0.608	0.583	0.455	0.401	0.432	0.297	0.734
≥ 0.3	5613	0.758	0.564	0.555	0.312	0.395	0.319	0.289	0.648

Panel D: All maturities

Moneyness	N	SQR	LOU	CTOU	LOUJ	LOUSV	CTOUJ	CTOUSV	CSQR
< -0.3	8981	0.260	0.176	0.183	0.171	0.160	0.241	0.213	0.451
$[-0.3, -0.1)$	10016	0.524	0.372	0.367	0.382	0.185	0.365	0.267	0.751
$[-0.1, 0.1)$	11452	0.742	0.500	0.472	0.479	0.231	0.457	0.219	0.774
$[0.1, 0.3)$	10863	0.844	0.595	0.562	0.446	0.376	0.428	0.235	0.702
≥ 0.3	16787	0.615	0.476	0.467	0.277	0.340	0.274	0.219	0.524

Table 8

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

Estimation of the parameters by maximum likelihood for the SQR, LOU and CTOU models, and by pseudo maximum likelihood in the remaining ones. The data employed in the estimation includes the VIX index and its options and futures from Feb 06 until Dec 10. 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 "LOU" refers to a log-normal Ornstein-Uhlenbeck process. "LOUJ" introduces jumps in the LOU model, whose size follows an exponential distribution. "CTOU" adds central tendency to the LOU process. "LOUSV" denotes a LOU model with stochastic volatility modeled with a Gamma OU Lévy process. "CTOUJ" adds central tendency and jumps to the LOU model, while "CTOUSV" introduces central tendency and stochastic volatility. "CSQR" is a concatenated SQR model. The column labeled N gives the number of prices per category. τ denotes time to maturity.

Panel A: $\tau < 1$ month

Moneyness	N	SQR	LOU	CTOU	LOUJ	LOUSV	CTOUJ	CTOUSV	CSQR
< -0.3	457	0.268	0.232	0.212	0.248	0.101	0.345	0.199	0.268
$[-0.3, -0.1)$	1657	0.568	0.415	0.391	0.425	0.141	0.420	0.259	0.636
$[-0.1, 0.1)$	2239	0.870	0.598	0.533	0.536	0.242	0.517	0.240	0.813
$[0.1, 0.3)$	1787	0.833	0.660	0.579	0.464	0.350	0.459	0.192	0.691
≥ 0.3	1385	0.442	0.409	0.386	0.283	0.275	0.281	0.188	0.400

Panel B: 1 month$< \tau < 3$ months

Moneyness	N	SQR	LOU	CTOU	LOUJ	LOUSV	CTOUJ	CTOUSV	CSQR
< -0.3	1970	0.304	0.224	0.222	0.231	0.159	0.347	0.244	0.380
$[-0.3, -0.1)$	4274	0.581	0.399	0.400	0.410	0.160	0.394	0.268	0.746
$[-0.1, 0.1)$	3966	0.814	0.515	0.507	0.513	0.199	0.496	0.178	0.792
$[0.1, 0.3)$	2272	0.968	0.663	0.640	0.510	0.373	0.505	0.202	0.765
≥ 0.3	1505	0.651	0.512	0.498	0.333	0.346	0.321	0.218	0.538

Panel C: $\tau > 3$ months

Moneyness	N	SQR	LOU	CTOU	LOUJ	LOUSV	CTOUJ	CTOUSV	CSQR
< -0.3	1698	0.366	0.269	0.276	0.279	0.205	0.394	0.260	0.687
$[-0.3, -0.1)$	2764	0.547	0.411	0.410	0.418	0.240	0.410	0.299	0.810
$[-0.1, 0.1)$	2174	0.745	0.523	0.515	0.489	0.271	0.476	0.288	0.805
$[0.1, 0.3)$	990	1.068	0.746	0.730	0.556	0.404	0.542	0.344	0.920
≥ 0.3	633	0.747	0.566	0.548	0.376	0.337	0.368	0.404	0.628

Panel D: All maturities

Moneyness	N	SQR	LOU	CTOU	LOUJ	LOUSV	CTOUJ	CTOUSV	CSQR
< -0.3	4125	0.327	0.245	0.245	0.253	0.175	0.367	0.246	0.521
$[-0.3, -0.1)$	8695	0.568	0.406	0.402	0.415	0.186	0.404	0.277	0.748
$[-0.1, 0.1)$	8379	0.812	0.540	0.516	0.513	0.231	0.497	0.228	0.801
$[0.1, 0.3)$	5049	0.944	0.679	0.638	0.504	0.371	0.497	0.234	0.773
≥ 0.3	3523	0.598	0.485	0.468	0.323	0.318	0.315	0.252	0.507

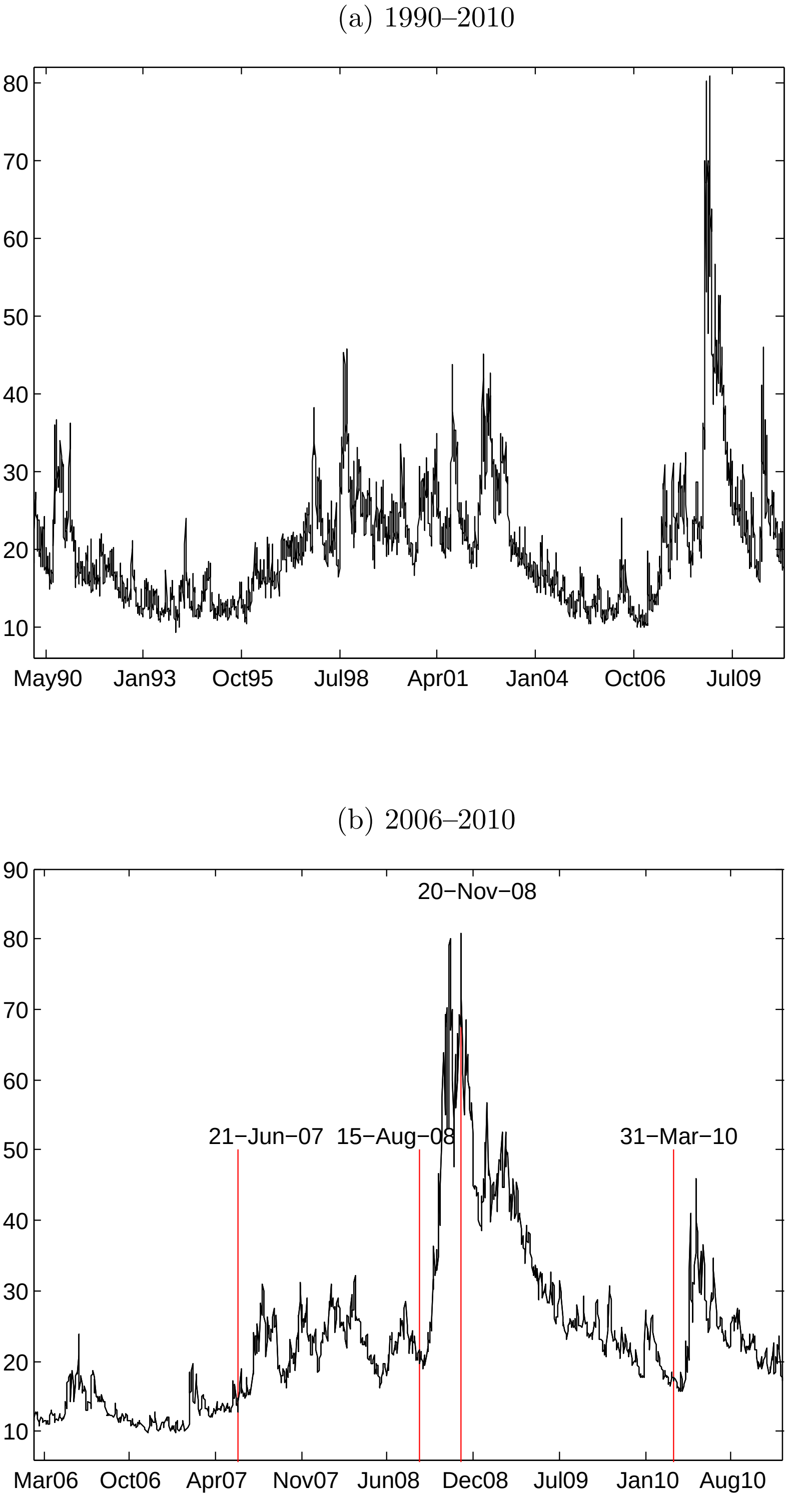

Fig. 1. Historical evolution of the VIX index.

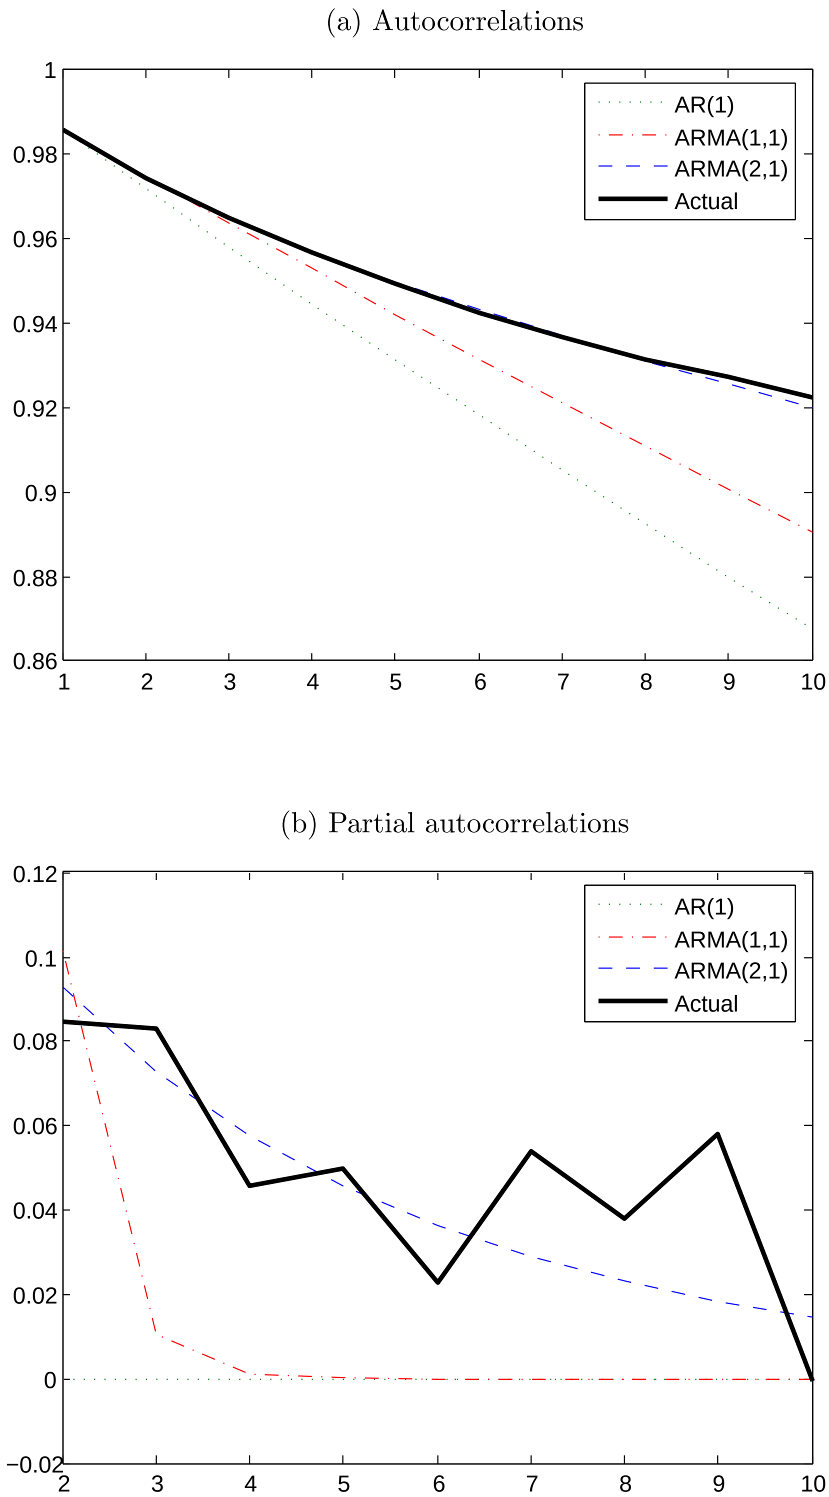

Fig. 2. Time series autocorrelations of the log-VIX and estimated ARMA models. Results are based on the 1990-2010 sample (5280 daily observations).

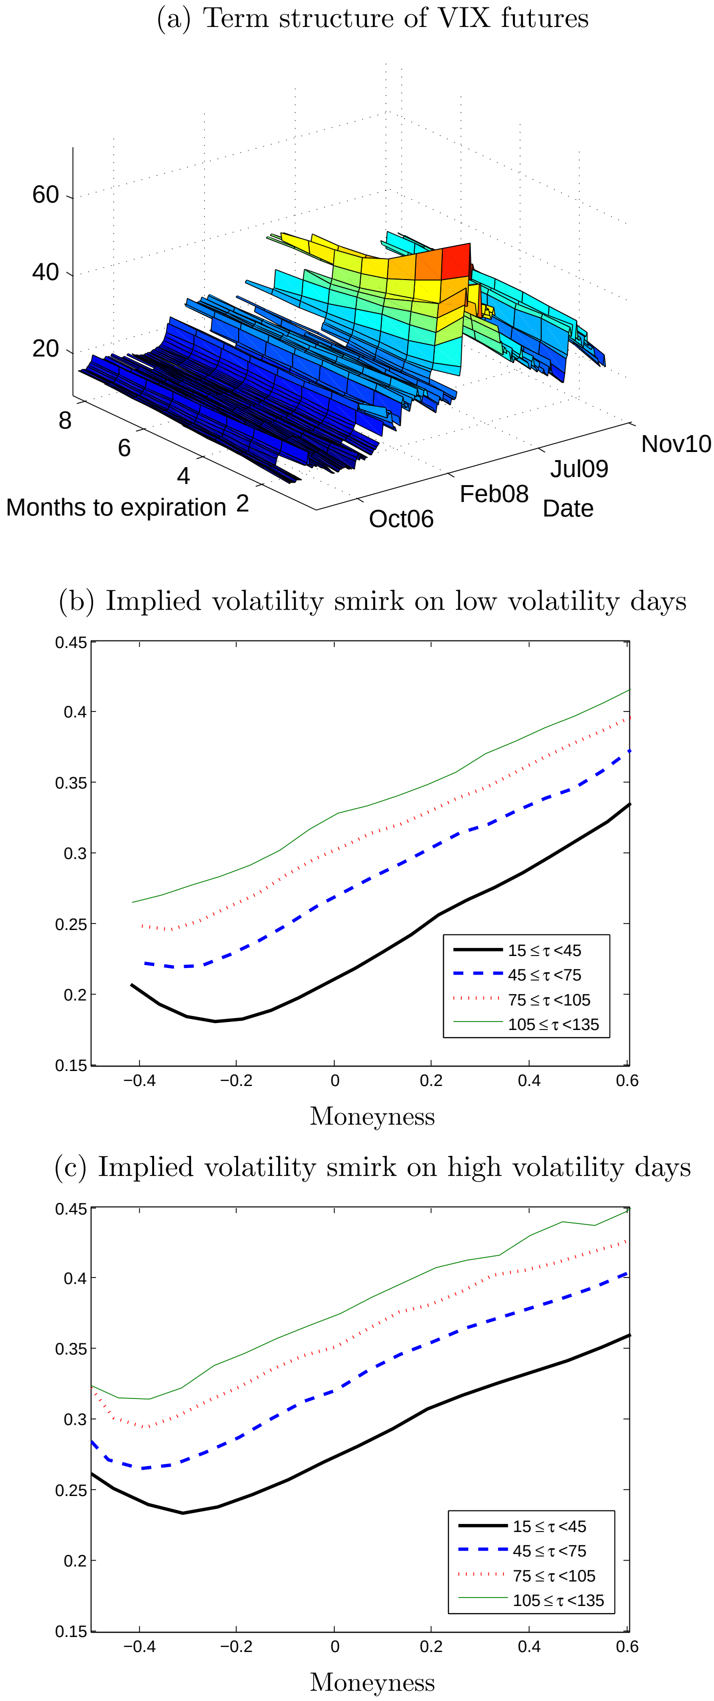

Fig. 3. Term structure of VIX derivatives. The lines in Panels B and C show the average implied volatility for a certain moneyness and time to maturity (denoted by τ) within a given interval. Implied volatilities have been obtained by inverting the Black (1976) call price formula. Moneyness is defined as $\log(K/F(t,T))$. Low (high) volatility days are those in which one-month at-the-money implied volatilities are below (above) their average value over the sample.

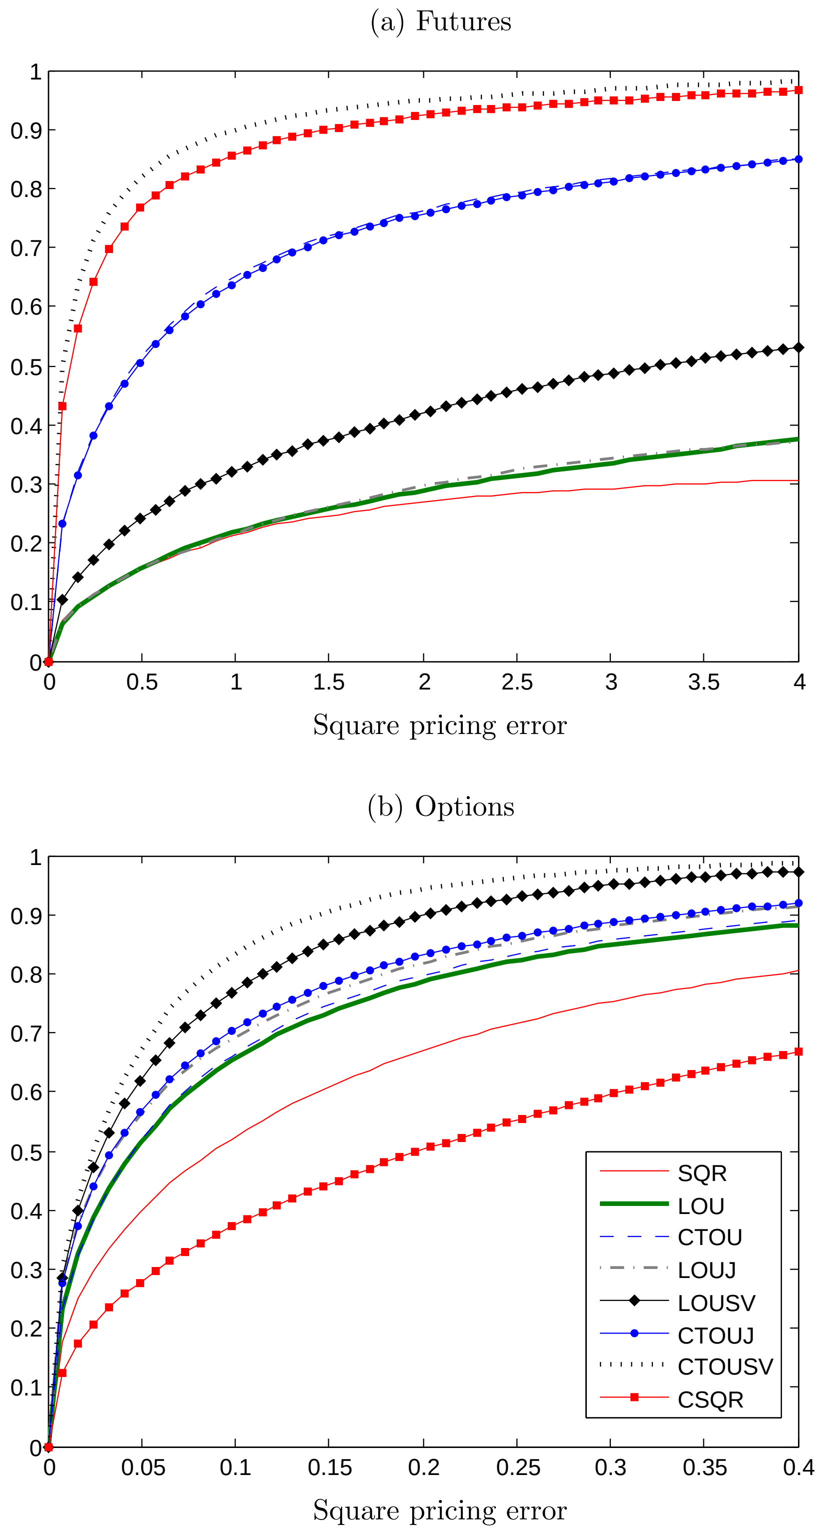

Fig. 4. Empirical cumulative distribution function of the square pricing errors. "SQR" denotes square root model and "LOU" refers to a log-normal Ornstein-Uhlenbeck process. "LOUJ" introduces jumps in the LOU model, whose size follows an exponential distribution. "CTOU" adds central tendency to the LOU process. "LOUSV" denotes a LOU model with stochastic volatility modeled with a Gamma OU Lévy process. "CTOUJ" adds central tendency and jumps to the LOU model, while "CTOUSV" introduces central tendency and stochastic volatility. "CSQR" is the concatenated SQR model. Estimation by maximum likelihood for the SQR, LOU and CTOU models, and by pseudo maximum likelihood in the remaining ones. The data employed in the estimation includes the VIX index and its options and futures from Feb 06 until Dec 10.

(a) Differences between model-based and actual one-month futures prices

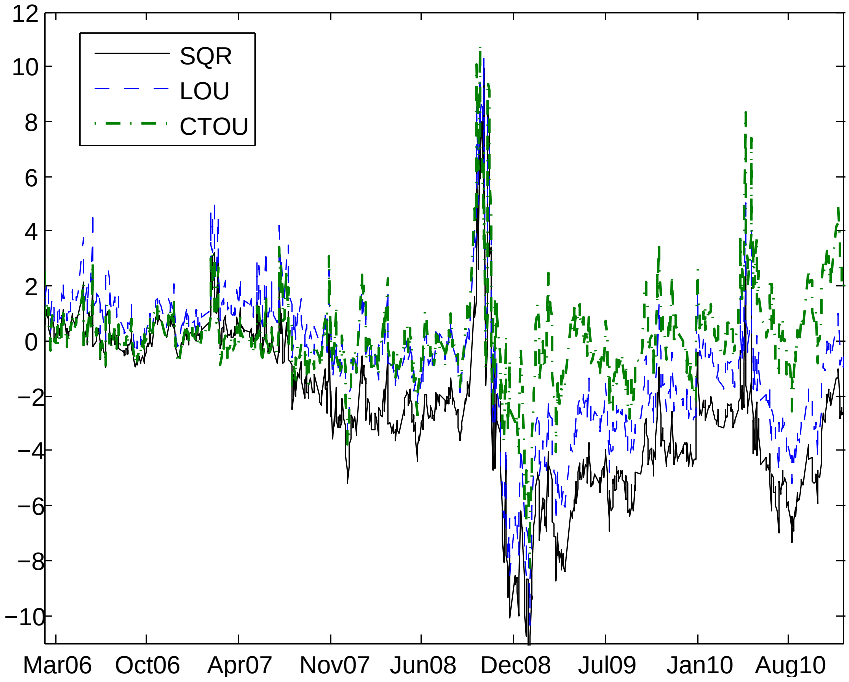

(b) Differences between model-based and actual one-month futures prices for the two best performing models

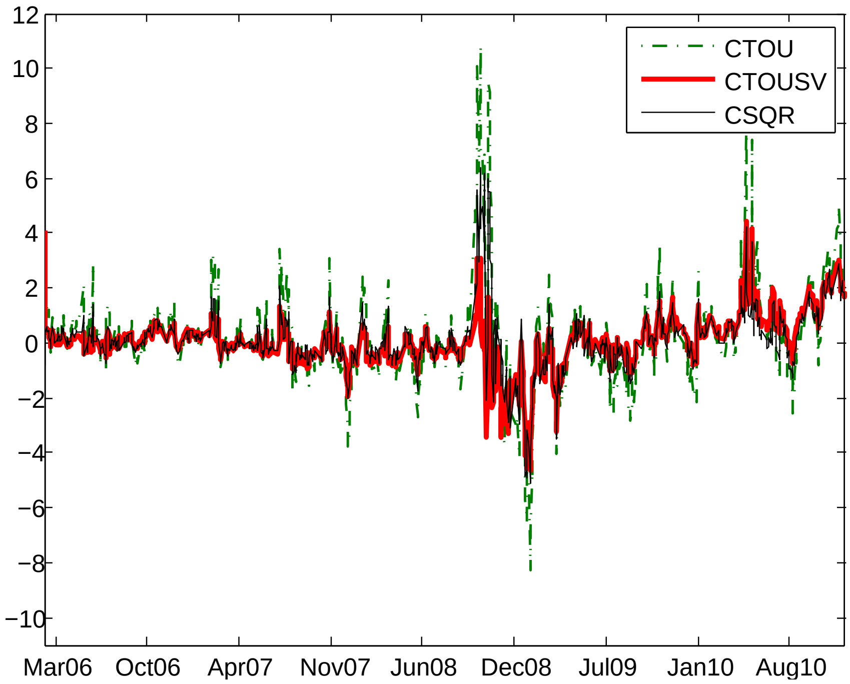

Fig. 5. Historical evolution of futures pricing errors. "SQR" denotes square root model and "LOU" refers to a log-normal Ornstein-Uhlenbeck process. "CTOU" adds central tendency to the LOU process. "CTOUSV" introduces central tendency and stochastic volatility in the LOU model. "CSQR" is the concatenated SQR model. One-month actual futures prices have been obtained by interpolation of the prices of the adjacent maturities. Estimation by maximum likelihood for the SQR, LOU and CTOU models, and by pseudo maximum likelihood in the remaining ones. The data employed in the estimation includes the VIX index and its options and futures from Feb 06 until Dec 10.

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

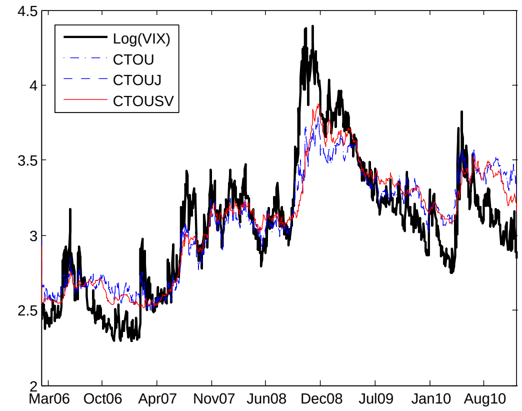

(b) Filtered $\omega(t)$ for different log-OU models with stochastic volatility

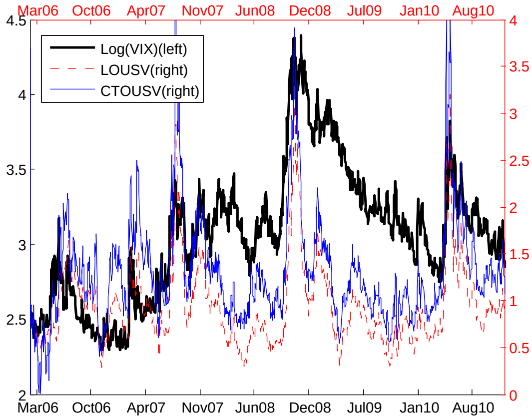

(c) One-month volatilities of the VIX. Risk-neutral measure

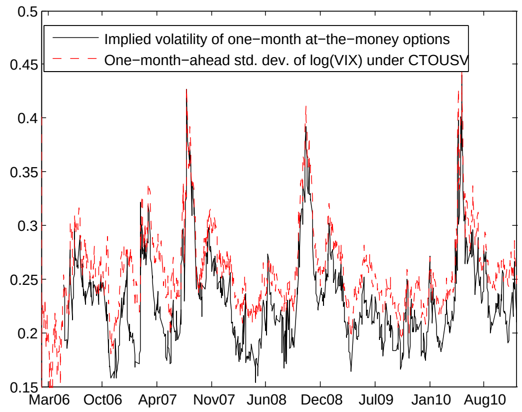

(d) Logs of one-day-ahead standard deviations of the VIX. Real measure

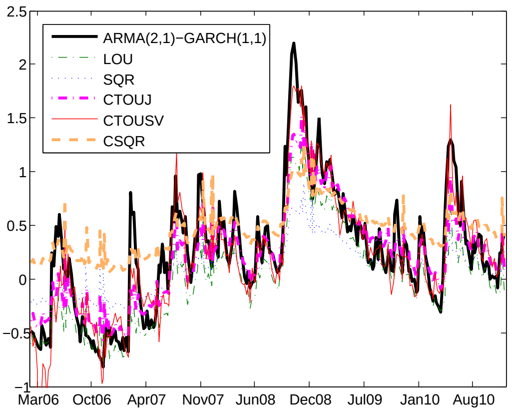

Fig. 6. Historical evolution of central tendency and stochastic volatility. $\theta(t)$ denotes the time-varying central tendency around which the VIX mean-reverts in central tendency models. $\omega(t)$ denotes the instantaneous volatility of stochastic volatility models. Different vertical scales are used for the VIX and the volatilities in Panel B. "SQR" denotes square root model and "LOU" refers to a log-normal Ornstein-Uhlenbeck process. "LOUSV" denotes a LOU model with stochastic volatility modeled with a Gamma OU Lévy process. "CTOUJ" adds central tendency and jumps to the LOU model, while "CTOUSV" introduces central tendency and stochastic volatility. "CSQR" is the concatenated SQR model. One-month implied vols in Panel C have been obtained by interpolation of the implied vols of options with moneyness $|\log(K/F(t,T))| < 0.1$, which in turn result from inverting the Black (1976) call price formula. The black line in Panel D is the conditional standard deviation of the VIX, obtained from an ARMA(2,1)-GARCH(1,1) model estimated for the daily data of the log VIX. Only Thursdays are plotted on Panel D to avoid cluttering the picture. Estimation by maximum likelihood for the SQR, LOU and CTOU models, and by pseudo maximum likelihood in the remaining ones. The data employed in the estimation includes the VIX index and its options and futures from Feb 06 until Dec 10.

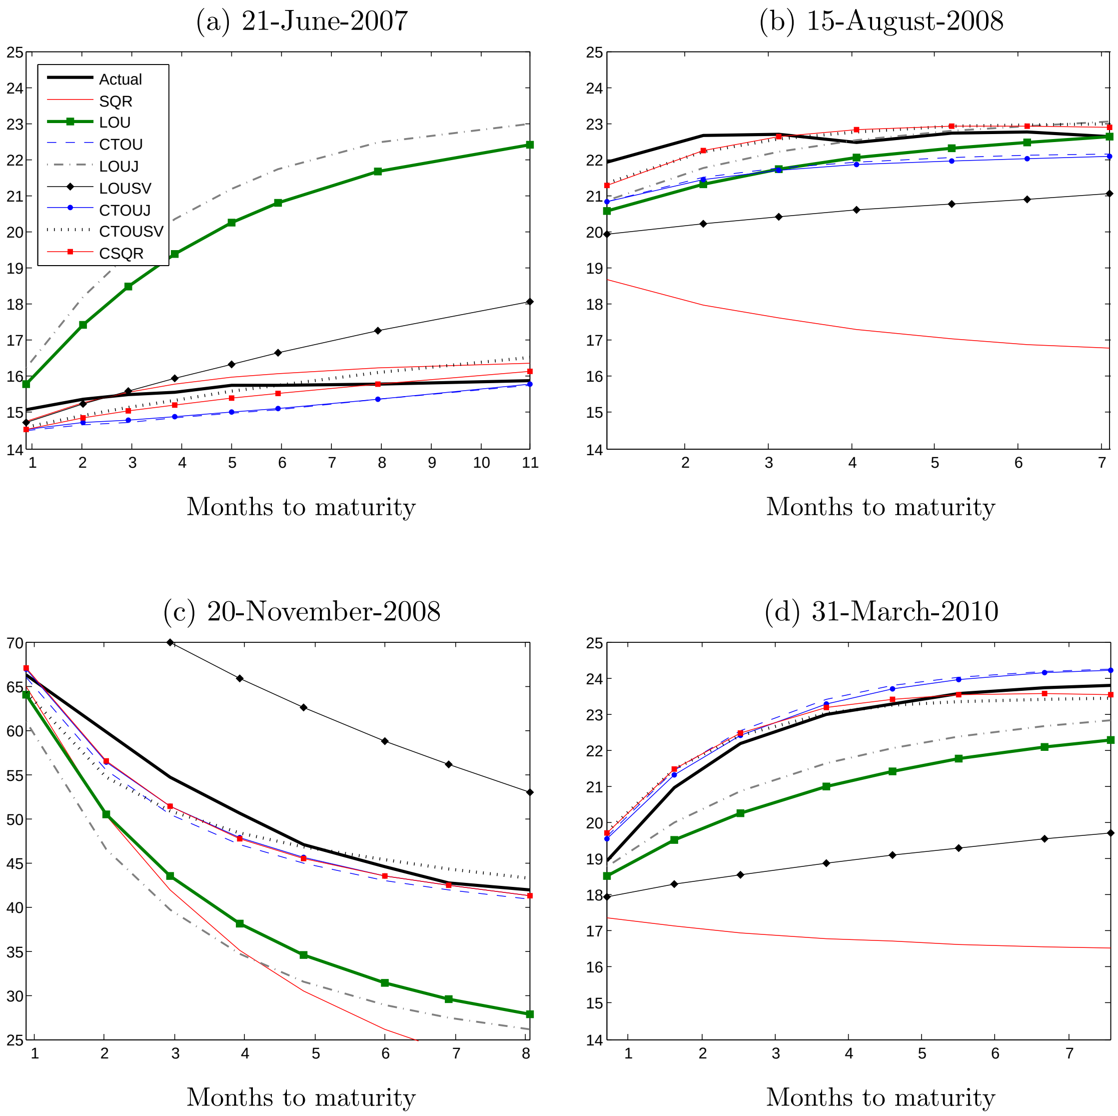

Fig. 7. Fit of the term structure of futures prices. "SQR" denotes square root model and "LOU" refers to a log-normal Ornstein-Uhlenbeck process. "LOUJ" introduces jumps in the LOU model, whose size follows an exponential distribution. "CTOU" adds central tendency to the LOU process. "LOUSV" denotes a LOU model with stochastic volatility modeled with a Gamma OU Lévy process. "CTOUJ" adds central tendency and jumps to the LOU model, while "CTOUSV" introduces central tendency and stochastic volatility. "CSQR" is the concatenated SQR model. Estimation by maximum likelihood for the SQR, LOU and CTOU models, and by pseudo maximum likelihood in the remaining ones. The data employed in the estimation includes the VIX index and its options and futures from Feb 06 until Dec 10.

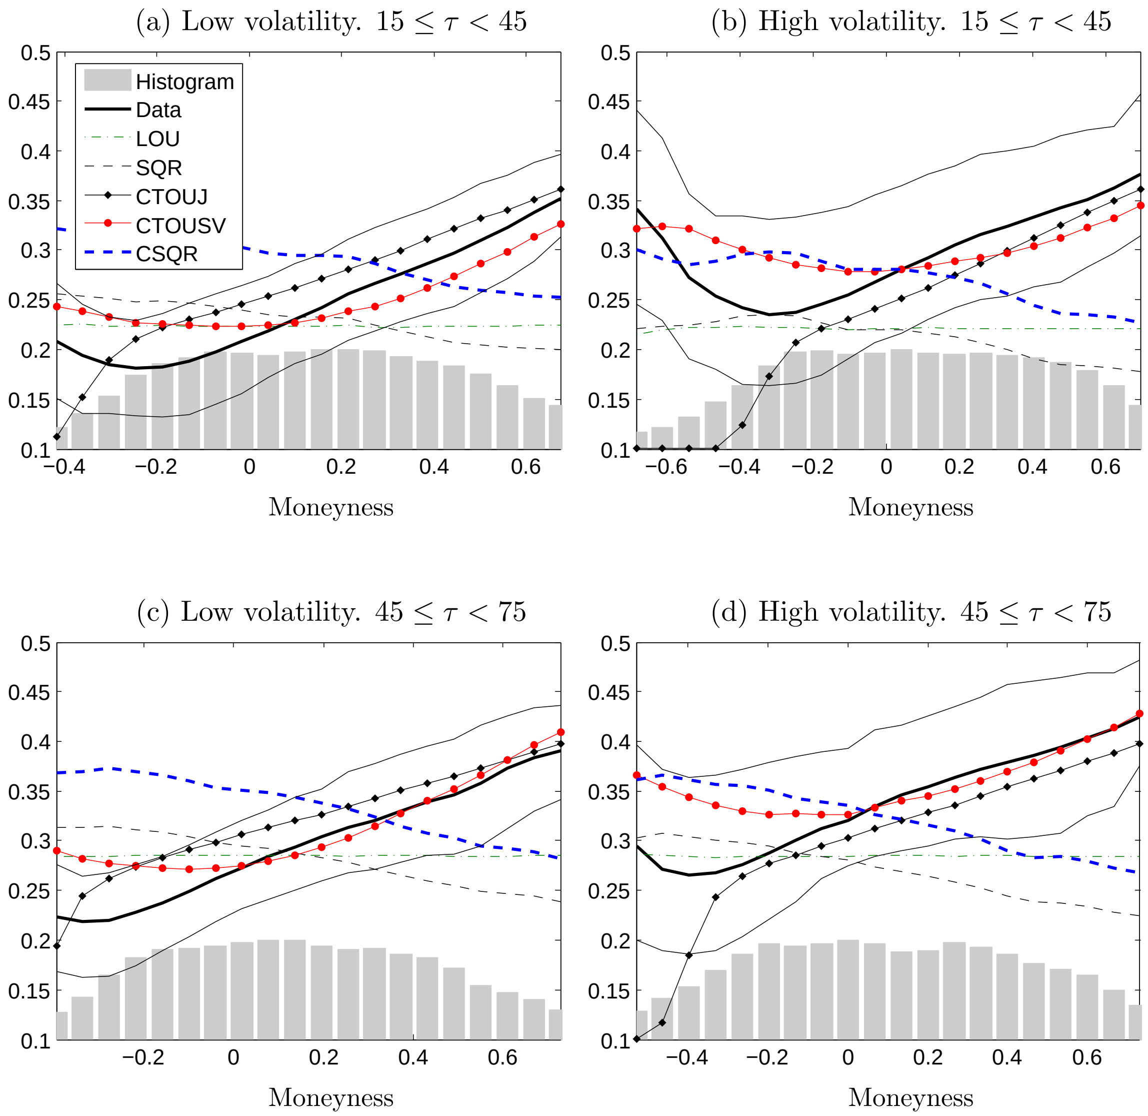

Fig. 8.1. Fit of the implied volatility smirks. Short maturities. The lines show the average implied volatility for a certain moneyness and time to maturity (denoted by τ) within a given interval. The thick black lines correspond to the average implied volatilities of the actual call prices, while the thin black lines show the 5% and 95% percentiles. Grey bars at the bottom of the plots show the histogram of call prices across maturities. The interval of moneyness in each panel has been chosen to cover the central 90% section of the data. "SQR" denotes square root model and "LOU" refers to a log-normal Ornstein-Uhlenbeck process. "CTOUJ" adds central tendency and jumps to the LOU model, while "CTOUSV" introduces central tendency and stochastic volatility. "CSQR" is the concatenated SQR model. Moneyness is defined as $\log(K/F(t,T)$. Estimation by maximum likelihood for the SQR, and LOU models, and by pseudo maximum likelihood in the remaining ones. The data employed in the estimation includes the VIX index and its options and futures from Feb 06 until Dec 10.

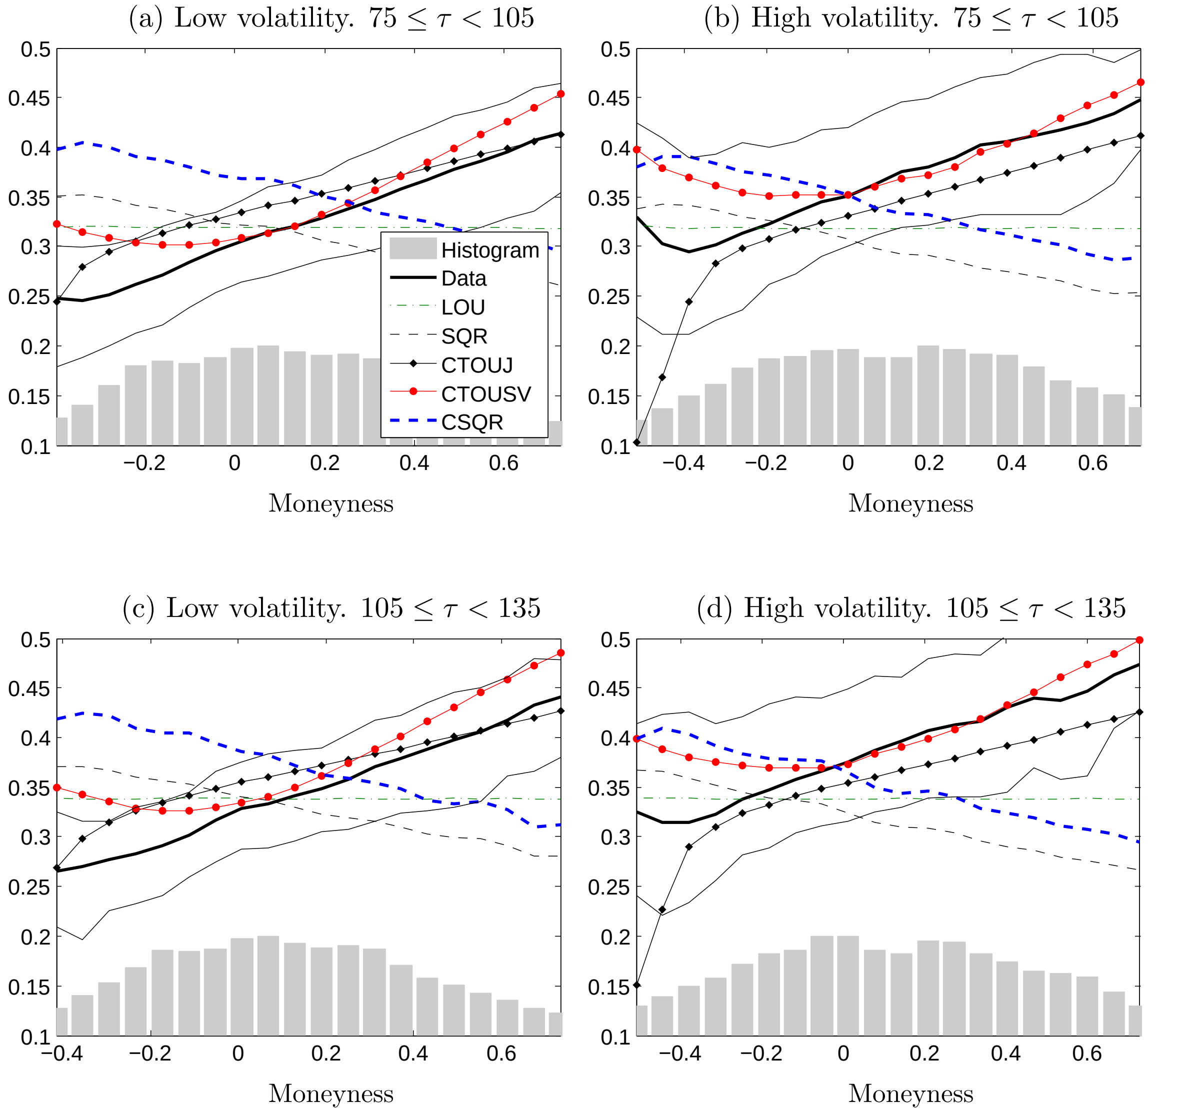

Fig. 8.2. Fit of the implied volatility smirks. Long maturities. The lines show the average implied volatility for a certain moneyness and time to maturity (denoted by τ) within a given interval. The thick black lines correspond to the average implied volatilities of the actual call prices, while the thin black lines show the 5% and 95% percentiles. Grey bars at the bottom of the plots show the histogram of call prices across maturities. The interval of moneyness in each panel has been chosen to cover the central 90% section of the data. "SQR" denotes square root model and "LOU" refers to a log-normal Ornstein-Uhlenbeck process. "CTOUJ" adds central tendency and jumps to the LOU model, while "CTOUSV" introduces central tendency and stochastic volatility. "CSQR" is the concatenated SQR model. Moneyness is defined as $\log(K/F(t,T)$. Estimation by maximum likelihood for the SQR, and LOU models, and by pseudo maximum likelihood in the remaining ones. The data employed in the estimation includes the VIX index and its options and futures from Feb 06 until Dec 10.