

Unspanned Stochastic Volatility and the Pricing of Commodity Derivatives

Anders B. Trolle

Copenhagen Business School and Ecole Polytechnique Fédérale de Lausanne

Eduardo S. Schwartz

UCLA Anderson School of Management and NBER

Commodity derivatives are becoming an increasingly important part of the global derivatives market. Here we develop a tractable stochastic volatility model for pricing commodity derivatives. The model features unspanned stochastic volatility, quasi-analytical prices of options on futures contracts, and dynamics of the futures curve in terms of a low-dimensional affine state vector. We estimate the model on NYMEX crude oil derivatives using an extensive panel data set of 45,517 futures prices and 233,104 option prices, spanning 4082 business days. We find strong evidence for two predominantly unspanned volatility factors. (*JEL G13*)

1. Introduction

The market for commodity derivatives has exhibited phenomenal growth over the past few years. For exchange-traded commodity derivatives, the Bank for International Settlements (BIS) estimates that the number of outstanding contracts has more than doubled from 12.4 million in June 2003 to 32.1 million in June 2006; see BIS (2007). For over-the-counter (OTC) commodity derivatives, the growth has been even stronger with the BIS estimating that, over the same period, the notional value of outstanding contracts increased fivefold from USD 1.04 trillion to USD 6.39 trillion.¹

We thank Michael Brennan, Peter Christoffersen, Pierre Collin-Dufresne, Alexander Eydeland, Hélyette Geman, Bjarne Astrup Jensen, Dev Joneja, Vassilis Koulovassilopoulos, David Lando, Claus Munk, Carsten Sørensen, and seminar participants at Bocconi University, Copenhagen Business School, University of Miami, Lehman Brothers, Natixis, the Commodities 2007 conference in London, the 2008 Nippon Finance Association meeting, and the 2008 European Finance Association meeting for comments. We are especially grateful for suggestions by Raman Uppal (the editor) and two anonymous referees that have improved the paper significantly. Anders B. Trolle thanks the Danish Social Science Research Council for financial support. Send correspondence to Anders B. Trolle, Copenhagen Business School, Solbjerg Plads 3, A5, DK-2000 Frederiksberg, Denmark; telephone: +45-3815-3058. E-mail: abt.fi@cbs.dk.

¹ Furthermore, during this period, commodity derivatives' share of the total OTC derivatives market increased from 0.61% to 1.73% in terms of notional values, and from 1.26% to 7.13% in terms of gross market value. Hence, commodity derivatives are becoming an increasingly important segment of the total OTC derivatives market.

Importantly, a large and increasing fraction of the commodity derivatives are options (as opposed to futures, forwards, and swaps). According to BIS statistics, options now constitute over one-third of the number of outstanding exchange-traded contracts and almost two-thirds of the notional value of outstanding OTC contracts.

In order to price, hedge, and risk-manage commodity options, it is critical to understand the dynamics of volatility in commodity markets. While volatility is clearly stochastic, it is not clear to what extent volatility risk can be hedged by trading in the commodities themselves or, more generally, their associated futures, forward, or swap contracts; in other words, the extent to which volatility is spanned. If, for a given commodity, volatility contains important unspanned components, options are not redundant securities and cannot be fully hedged and risk-managed, using only the underlying instruments.

Existing equilibrium models typically imply that volatility in commodity markets is largely spanned by the futures contracts. For instance, models that emphasize the embedded timing option in inventories (see, e.g., Deaton and Laroque 1992, 1996; Chambers and Bailey 1996; Routledge, Seppi, and Spatt 2000) have the implication that the convenience yield and, therefore, the degree of backwardation of the futures curve is increasing in volatility. The model in Litzenberger and Rabinowitz (1995), which incorporates the embedded option in reserves of extractable resource commodities, has similar implications.² The empirical analyses by Litzenberger and Rabinowitz (1995) and Routledge, Seppi, and Spatt (2000) for crude oil and Ng and Pirrong (1994) for metals show that the degree of backwardation is indeed positively related to volatility, implying that volatility does contain a component that is spanned by the futures contracts. However, whether volatility also contains important unspanned components remains an open question.

In this paper, we conduct a comprehensive analysis of unspanned stochastic volatility in commodity markets and make both theoretical and empirical contributions to the literature. Our theoretical analysis is applicable to most commodities. However, in the empirical analysis, we limit our attention to the crude oil market, which is by far the largest and most liquid commodity derivatives market.

The first main contribution of the paper is to develop a tractable framework for pricing commodity derivatives in the presence of unspanned stochastic volatility. The model is specified directly under the risk-neutral probability measure and is based on the Heath, Jarrow, and Morton (1992) (HJM) framework. In its most general form, futures prices are driven by three factors (one factor being the spot price of the commodity and two factors affecting the forward cost of carry curve), and option prices are driven by two additional volatility factors.

² Other papers that emphasize production/extraction and investment decisions for the formation of futures prices include Casassus, Collin-Dufresne, and Routledge (2003); Kogan, Livdan, and Yaron (2005); and Carlson, Khoker, and Titman (2007). In these models, the relationship between volatility and the slope of the futures curve is highly nonlinear and possibly nonmonotone.

Both volatility factors may contain a spanned and an unspanned component and both factors may affect the instantaneous volatility of the spot price and the forward cost of carry. The model features quasi-analytical prices of European options on futures contracts based on transform techniques. By a suitable parameterization of the shocks to the forward cost of carry curve, the dynamics of the futures curve can be described in terms of a low-dimensional affine state vector, which makes the model suited for pricing complex commodity derivatives by simulation.

The second main contribution of the paper is to conduct an extensive empirical analysis of the model. We use a very large panel data set of the New York Mercantile Exchange (NYMEX) crude oil futures and options spanning 4082 business days from January 1990 to May 2006. It consists of a total of 45,517 futures prices and 233,104 option prices. This is, to our knowledge, by far the most extensive data set that has been used in empirical studies of commodity derivatives.³ Estimation is facilitated by parameterizing the market prices of risk, such that the state vector is also described by an affine diffusion under the actual probability measure. The estimation procedure is quasi-maximum likelihood in conjunction with the extended Kalman filter.

In the empirical analysis, in addition to the general model, we also consider two nested, more parsimonious, specifications. In one specification, futures prices are driven by two factors, and option prices are driven by two additional volatility factors, the second of which is completely unspanned by the futures contracts and does not affect the instantaneous volatility of the spot price or the forward cost of carry. In another specification, futures prices are driven by two factors, and option prices are driven by one additional volatility factor that may contain a spanned and an unspanned component.

We find that two volatility factors are necessary to match options on futures contracts. While the specification with one volatility factor captures the overall time variation in implied volatilities, the specifications with two volatility factors perform much better at capturing the variation in implied volatilities across option maturity and moneyness. In the general two-factor specification, both volatility factors are predominantly unspanned by the futures contracts, and the first volatility factor drives virtually all of the instantaneous volatility of the spot price and the front end of the forward cost of carry curve (with the second volatility factor being more important for the instantaneous volatility of longer-term forward cost of carry rates). Therefore, the more parsimonious two-factor specification performs almost as well as the general specification in terms of pricing short-term and medium-term options. This holds true both in-sample and out-of-sample.

³ Doran and Ronn (2006) use futures and options on crude oil (as well as natural gas and heating oil) to study the market price of volatility risk in energy commodity markets. However, they use only at-the-money (ATM) options, whereas we use options with a wide range of strike prices. Richter and Sørensen (2002) use futures and options to investigate volatility and seasonality dynamics in the soybean derivatives market. However, they also use a much more limited option data set than we do.

We also assess the importance of unspanned stochastic volatility when hedging an option portfolio. Hedging the option portfolio solely with futures contracts causes only a small reduction in the variation in the portfolio profit–loss. However, consistent with the finding that volatility is largely unspanned by the futures contracts, adding one or two options to the set of hedge instruments significantly reduces the variation in the portfolio profit–loss. Again, these results hold true both in-sample and out-of-sample as well as for different weighting schemes of the individual options in the portfolio.

The model proposed in this paper appears to be the first stochastic volatility HJM-type model for pricing commodity derivatives.⁴ Previous HJM-type commodity models such as Cortazar and Schwartz (1994), Amin, Ng, and Pirrong (1995), Miltersen and Schwartz (1998), Clewlow and Strickland (1999), and Miltersen (2003) all assume deterministic volatilities.⁵ The advantage of working in an HJM setting is that unspanned stochastic volatility arises naturally.

An alternative approach to pricing commodity derivatives relies on specifying the (typically affine) dynamics of a limited set of state variables and deriving futures prices endogenously. Examples of this approach include Gibson and Schwartz (1990); Brennan (1991); Schwartz (1997); Hilliard and Reis (1998); Schwartz and Smith (2000); Richter and Sørensen (2002); Nielsen and Schwartz (2004); and Casassus and Collin-Dufresne (2005). Of these papers, only Richter and Sørensen (2002) and Nielsen and Schwartz (2004) explicitly allow for stochastic volatility. The main drawback of this modeling approach is that volatility is almost invariably completely spanned by the futures contracts.⁶

The paper is organized as follows. Section 2 lays out the model for pricing commodity derivatives in the presence of unspanned stochastic volatility. Section 3 describes the crude oil derivatives data and the estimation procedure. Section 4 discusses the empirical results. Section 5 concludes. Three appendices contain proofs and additional information.

2. A Model for Commodity Derivatives Featuring Unspanned Stochastic Volatility

Here, we develop a tractable framework for pricing commodity derivatives in the presence of unspanned stochastic volatility. In Sections 2.1–2.4, we present a general model that nests some interesting and more parsimonious specifications that are described in Section 2.5.

⁴ See Casassus, Collin-Dufresne, and Goldstein (2005) and Trolle and Schwartz (2009) for HJM-type models for pricing interest rate derivatives.

⁵ Eydeland and Geman (1998) propose a stochastic volatility Heston (1993) model for pricing energy derivatives. However, they model only the evolution of the spot price, not the evolution of the entire futures curve.

⁶ Indeed, this is the case for the models in Richter and Sørensen (2002) and Nielsen and Schwartz (2004). Collin-Dufresne and Goldstein (2002) derive the parameter restrictions necessary for volatility to be unspanned in affine term structure models. Similar conditions can be derived for affine commodity models.

2.1 The model under the risk-neutral measure

Let $S(t)$ denote the time- t spot price of the commodity with an instantaneous spot cost of carry given by $\delta(t)$. Furthermore, let $y(t, T)$ denote the time- t instantaneous forward cost of carry at time T , with $y(t, t) = \delta(t)$. Standard models of commodity derivatives, such as Gibson and Schwartz (1990), Schwartz (1997), and Casassus and Collin-Dufresne (2005), among others, typically specify a process for $S(t)$ and $\delta(t)$. Here instead, we follow Cortazar and Schwartz (1994) and Miltersen and Schwartz (1998), among others, and specify a process for $S(t)$ and $y(t, T)$; that is, we model the evolution of the entire forward cost of carry curve.

The main theoretical contribution of this paper is to extend the framework to accommodate unspanned stochastic volatility. Specifically, we assume that the volatility of both $S(t)$ and $y(t, T)$ may depend on two volatility factors, $v_1(t)$ and $v_2(t)$, and postulate the following very general process for $S(t)$, $y(t, T)$, $v_1(t)$, and $v_2(t)$ under the risk-neutral measure:

$$\frac{dS(t)}{S(t)} = \delta(t)dt + \sigma_{S1}\sqrt{v_1(t)}dW_1^Q(t) + \sigma_{S2}\sqrt{v_2(t)}dW_2^Q(t), \quad (1)$$

$$dy(t, T) = \mu_y(t, T)dt + \sigma_{y1}(t, T)\sqrt{v_1(t)}dW_3^Q(t) + \sigma_{y2}(t, T)\sqrt{v_2(t)}dW_4^Q(t), \quad (2)$$

$$dv_1(t) = (\eta_1 - \kappa_1 v_1(t) - \kappa_{12} v_2(t))dt + \sigma_{v1}\sqrt{v_1(t)}dW_5^Q(t), \quad (3)$$

$$dv_2(t) = (\eta_2 - \kappa_{21} v_1(t) - \kappa_2 v_2(t))dt + \sigma_{v2}\sqrt{v_2(t)}dW_6^Q(t), \quad (4)$$

where $W_i^Q(t)$, $i = 1, \dots, 6$, denote Wiener processes under the risk-neutral measure.⁷ We allow $W_1^Q(t)$, $W_3^Q(t)$, and $W_5^Q(t)$ to be correlated, with ρ_{13} , ρ_{15} , and ρ_{35} denoting pairwise correlations, and we also allow $W_2^Q(t)$, $W_4^Q(t)$, and $W_6^Q(t)$ to be correlated, with ρ_{24} , ρ_{26} , and ρ_{46} denoting pairwise correlations. This is the most general correlation structure that preserves the tractability of the model.

The forward cost of carry is given by the forward interest rate minus the forward convenience yield, and the model could be extended with separate processes for the forward interest rate and the forward convenience yield.⁸ However, for pricing most commodity futures, this extension is of minor importance; see, e.g., the discussion in Schwartz (1997). Furthermore, for pricing

⁷ Existence of this process requires $\eta_1 \geq 0$, $\eta_2 \geq 0$, $\kappa_{12} \leq 0$, and $\kappa_{21} \leq 0$; see, e.g., the discussion of affine processes in Dai and Singleton (2000).

⁸ When we allow for stochastic interest rates, we should, strictly speaking, distinguish between forward and future convenience yield and forward and future cost of carry; see Miltersen and Schwartz (1998) for a further discussion on this issue.

short-term or medium-term options on most commodity futures, the pricing error that arises from not explicitly modeling stochastic interest rates is negligible, since the volatility of interest rates is typically orders of magnitudes smaller than the volatility of futures returns, and the correlation between interest rates and futures returns tends to be very low.⁹

Let $F(t, T)$ denote the time- t price of a futures contract maturing at time T . By definition, we have

$$F(t, T) \equiv S(t) \exp \left\{ \int_t^T y(t, u) du \right\}. \tag{5}$$

In the absence of arbitrage opportunities, the futures price process must be a martingale under the risk-neutral measure; see, e.g., Duffie (2001). Applying Ito’s lemma to Equation (5) and setting the drift to zero, it follows that the dynamics of $F(t, T)$ are given by

$$\begin{aligned} \frac{dF(t, T)}{F(t, T)} = & \sqrt{v_1(t)} \left(\sigma_{S1} dW_1^Q(t) + \int_t^T \sigma_{y1}(t, u) du dW_3^Q(t) \right) \\ & + \sqrt{v_2(t)} \left(\sigma_{S2} dW_2^Q(t) + \int_t^T \sigma_{y2}(t, u) du dW_4^Q(t) \right). \end{aligned} \tag{6}$$

Volatility of futures prices depends on $v_1(t)$ and $v_2(t)$. These are driven by $W_5^Q(t)$ and $W_6^Q(t)$, which do not appear in Equation (6). Therefore, volatility risk and, by implication, options on futures contracts cannot be completely hedged by trading in futures contracts alone. Hence, the model features unspanned stochastic volatility. To the extent that $W_5^Q(t)$ is correlated with $W_1^Q(t)$ and $W_3^Q(t)$, and $W_6^Q(t)$ is correlated with $W_2^Q(t)$ and $W_4^Q(t)$, both volatility factors contain a spanned component, and volatility risk is partly hedgeable. If these correlations are all zero, volatility risk is completely unhedgeable.

From the requirement that the drift of the futures price process is zero, we obtain the following condition on the drift of the forward cost of carry process.

Proposition 1. *Absence of arbitrage implies that the drift term in Equation (2) is given by*

$$\begin{aligned} \mu_y(t, T) = & - \left(v_1(t) \sigma_{y1}(t, T) \left(\rho_{13} \sigma_{S1} + \int_t^T \sigma_{y1}(t, u) du \right) \right) \\ & + v_2(t) \sigma_{y2}(t, T) \left(\rho_{24} \sigma_{S2} + \int_t^T \sigma_{y2}(t, u) du \right). \end{aligned} \tag{7}$$

⁹ For instance, for the sample period 2 January 1990 to 25 August 2003, Casassus and Collin-Dufresne (2005) estimate the crude oil spot return volatility to 0.397 and the spot interest rate volatility to 0.009. Furthermore, the instantaneous correlation between the spot return and the spot interest rate is estimated to be 0.051 and is insignificant. Accounting explicitly for stochastic interest rates may become more important when pricing long-term options on futures contracts.

Proof. See Appendix A. ■

This condition is analogous to the Heath, Jarrow, and Morton (1992) drift condition in forward rate term structure models.

2.2 An affine model for the dynamics of the futures curve

So far we have left $\sigma_{y_1}(t, T)$ and $\sigma_{y_2}(t, T)$ unspecified. Intuitively, long-term forward cost of carry rates should be less volatile than short-term forward cost of carry rates. This requirement is satisfied with the following time-homogeneous specification:

$$\sigma_{y_i}(t, T) = \alpha_i e^{-\gamma_i(T-t)}, \tag{8}$$

$i = 1, 2$.¹⁰ In the following proposition, we show that $y(t, T)$ is then an affine function of four state variables, $x_1(t)$, $x_2(t)$, $\phi_1(t)$, and $\phi_2(t)$, where $x_1(t)$ and $x_2(t)$ are stochastic, while $\phi_1(t)$ and $\phi_2(t)$ are “auxiliary,” locally deterministic, state variables.¹¹

Proposition 2. *The time- t instantaneous forward cost of carry at time T , $y(t, T)$, is given by*

$$y(t, T) = y(0, T) + \sum_{i=1}^2 (\alpha_i e^{-\gamma_i(T-t)} x_i(t) + \alpha_i e^{-2\gamma_i(T-t)} \phi_i(t)), \tag{9}$$

where $x_1(t)$, $x_2(t)$, $\phi_1(t)$, and $\phi_2(t)$ evolve according to

$$dx_1(t) = \left(-\gamma_1 x_1(t) - \left(\frac{\alpha_1}{\gamma_1} + \rho_{13} \sigma_{S1} \right) v_1(t) \right) dt + \sqrt{v_1(t)} dW_3^Q(t), \tag{10}$$

$$dx_2(t) = \left(-\gamma_2 x_2(t) - \left(\frac{\alpha_2}{\gamma_2} + \rho_{24} \sigma_{S2} \right) v_2(t) \right) dt + \sqrt{v_2(t)} dW_4^Q(t), \tag{11}$$

$$d\phi_i(t) = \left(-2\gamma_i \phi_i(t) + \frac{\alpha_i}{\gamma_i} v_i(t) \right) dt, \quad i = 1, 2, \tag{12}$$

subject to $x_1(0) = x_2(0) = \phi_1(0) = \phi_2(0) = 0$.

Proof. See Appendix A. ■

¹⁰ With this specification, σ_{S_i} , α_i , η_i , and σ_{v_i} , $i = 1, 2$ are not simultaneously identified; see, e.g., the discussion on invariant affine transformations in Dai and Singleton (2000). In our empirical analysis, we normalize η_i to one to achieve identification.

¹¹ This result draws on a branch of the term structure literature that has investigated under which conditions HJM models are Markovian with respect to a finite number of state variables; see, e.g., Ritchken and Sankarasubramaniam (1995) and Bhar and Chiarella (1997) for two early papers.

It follows that the instantaneous spot cost of carry is given by

$$\delta(t) = y(0, t) + \sum_{i=1}^2 \alpha_i(x_i(t) + \phi_i(t)), \tag{13}$$

and from Equation (5) we have that $F(t, T)$ is given by

$$F(t, T) = S(t) \frac{F(0, T)}{F(0, t)} \exp \left\{ \sum_{i=1}^2 \left(B_{x_i}(T - t)x_i(t) + B_{\phi_i}(T - t)\phi_i(t) \right) \right\}, \tag{14}$$

where

$$B_{x_i}(\tau) = \frac{\alpha_i}{\gamma_i} (1 - e^{-\gamma_i \tau}), \tag{15}$$

$$B_{\phi_i}(\tau) = \frac{\alpha_i}{2\gamma_i} (1 - e^{-2\gamma_i \tau}), \tag{16}$$

$i = 1, 2$.

It is convenient to use $s(t) \equiv \log(S(t))$ instead of $S(t)$ as a state variable. In this case, the log futures price is an affine function of $s(t)$, $x_1(t)$, $x_2(t)$, $\phi_1(t)$, and $\phi_2(t)$,

$$\begin{aligned} \log F(t, T) &= \log F(0, T) - \log F(0, t) + s(t) \\ &\quad + \sum_{i=1}^2 (B_{x_i}(T - t)x_i(t) + B_{\phi_i}(T - t)\phi_i(t)), \end{aligned} \tag{17}$$

and the dynamics of $s(t)$ are given by

$$\begin{aligned} ds(t) &= \left(y(0, t) + \sum_{i=1}^2 \alpha_i(x_i(t) + \phi_i(t)) - \frac{1}{2} (\sigma_{S1}^2 v_1(t) + \sigma_{S2}^2 v_2(t)) \right) dt \\ &\quad + \sigma_{S1} \sqrt{v_1(t)} dW_1^Q(t) + \sigma_{S2} \sqrt{v_2(t)} dW_2^Q(t). \end{aligned} \tag{18}$$

2.3 Pricing options on futures contracts

To price options on futures contracts, we introduce the transform

$$\psi(u, t, T_0, T_1) = E_t^Q [e^{u \log(F(T_0, T_1))}], \tag{19}$$

which has an exponentially affine solution as demonstrated in the following proposition:

Proposition 3. *The transform in Equation (19) is given by*

$$\psi(u, t, T_0, T_1) = e^{M(T_0-t) + N_1(T_0-t)v_1(t) + N_2(T_0-t)v_2(t) + u \log(F(t, T_1))}, \tag{20}$$

where $M(\tau)$, $N_1(\tau)$, and $N_2(\tau)$ solve the following system of ODEs

$$\frac{dM(\tau)}{d\tau} = N_1(\tau)\eta_1 + N_2(\tau)\eta_2, \tag{21}$$

$$\begin{aligned} \frac{dN_1(\tau)}{d\tau} = & -N_2(\tau)\kappa_{21} + N_1(\tau)(-\kappa_1 + u\sigma_{v1}(\rho_{15}\sigma_{S1} + \rho_{35}B_{x1}(T_1 - t))) \\ & + \frac{1}{2}N_1(\tau)^2\sigma_{v1}^2 + \frac{1}{2}(u^2 - u)(\sigma_{S1}^2 + B_{x1}(T_1 - t)^2 \\ & + 2\rho_{13}\sigma_{S1}B_{x1}(T_1 - t)), \end{aligned} \tag{22}$$

$$\begin{aligned} \frac{dN_2(\tau)}{d\tau} = & -N_1(\tau)\kappa_{12} + N_2(\tau)(-\kappa_2 + u\sigma_{v2}(\rho_{26}\sigma_{S2} + \rho_{46}B_{x2}(T_1 - t))) \\ & + \frac{1}{2}N_2(\tau)^2\sigma_{v2}^2 + \frac{1}{2}(u^2 - u)(\sigma_{S2}^2 + B_{x2}(T_1 - t)^2 \\ & + 2\rho_{24}\sigma_{S2}B_{x2}(T_1 - t)), \end{aligned} \tag{23}$$

subject to the boundary conditions $M(0) = 0$, $N_1(0) = 0$, and $N_2(0) = 0$.

Proof. See Appendix A. ■

Following Bakshi and Madan (2000), Duffie, Pan, and Singleton (2000), Collin-Dufresne and Goldstein (2003), and others, we can now price European options on futures contracts by applying the Fourier inversion theorem.¹²

Proposition 4. *The time- t price of a European put option expiring at time T_0 with strike K on a futures contract expiring at time T_1 , $\mathcal{P}(t, T_0, T_1, K)$, is given by*

$$\begin{aligned} \mathcal{P}(t, T_0, T_1, K) &= E_t^Q \left[e^{-\int_t^{T_0} r(s)ds} (K - F(T_0, T_1)) \mathbf{1}_{F(T_0, T_1) < K} \right] \\ &= P(t, T_0)(KG_{0,1}(\log(K)) - G_{1,1}(\log(K))), \end{aligned} \tag{24}$$

where $P(t, T_0)$ denotes the time- t price of a zero-coupon bond maturing at time T_0 , and $G_{a,b}(y)$ is defined as

$$G_{a,b}(y) = \frac{\Psi(a, t, T_0, T_1)}{2} - \frac{1}{\pi} \int_0^\infty \frac{\text{Im}[\Psi(a + iub, t, T_0, T_1)e^{-iuy}]}{u} du, \tag{25}$$

where $i = \sqrt{-1}$.

Proof. See Appendix A. ■

¹² We could also follow the approach of Carr and Madan (1999) to obtain an alternative pricing formula that permits the use of the computationally efficient fast Fourier transform algorithm. This might be preferable for applications requiring real-time calibration of the model. However, for our purpose, the formula derived here is sufficiently fast and numerically stable.

This formula is exact when interest rates are uncorrelated with futures prices under Q . To the extent that the correlation is low and the volatility of interest rates is significantly lower than the volatility of futures prices, the formula gives a very accurate approximation of the true price of short-term or medium-term options.¹³

2.4 Market price of risk specification

For the purpose of estimation, we also need the dynamics of the state vector under the actual probability measure P . This is achieved by specifying the market prices of risk, Λ_i , that link the Wiener processes under Q and P through

$$dW_i^P(t) = dW_i^Q(t) - \Lambda_i(t)dt, \quad i = 1, \dots, 6. \quad (26)$$

We apply the following “completely affine” specification, which has been widely used in the literature:

$$\Lambda_i(t) = \lambda_i \sqrt{v_1(t)}, \quad i = 1, 3, 5, \quad (27)$$

$$\Lambda_i(t) = \lambda_i \sqrt{v_2(t)}, \quad i = 2, 4, 6. \quad (28)$$

This specification preserves the affine structure of the state vector under the change of measure.¹⁴

2.5 Nested specifications

The model presented above is quite general. Futures prices are driven by three factors, $S(t)$, $x_1(t)$, and $x_2(t)$, and option prices are driven by two additional volatility factors, $v_1(t)$ and $v_2(t)$. Both volatility factors may contain a spanned and an unspanned component, and both factors may affect the instantaneous volatility of $S(t)$ and $y(t, T)$. This specification will be denoted by SV2gen in the following. We have also experimented with a wide range of nested, more parsimonious alternatives. To streamline the exposition, we will focus on the two most interesting specifications, which will be denoted by SV2 and SV1.

¹³ To see this, note that in general, we have

$$\begin{aligned} \mathcal{P}(t, T_0, T_1, K) &= P(t, T_0)E_t^Q[(K - F(T_0, T_1))\mathbf{1}_{F(T_0, T_1) < K}] \\ &\quad + \text{Cov}_t^Q \left[\exp \left\{ - \int_t^{T_0} r(s)ds \right\}, (K - F(T_0, T_1))\mathbf{1}_{F(T_0, T_1) < K} \right]. \end{aligned}$$

The covariance term is zero if $r(t)$ and $F(t, T_1)$ are uncorrelated, in which case the formula is exact. The covariance term is insignificant relative to the option price, provided that the correlation between $r(t)$ and $F(t, T_1)$ is low and the volatility of $r(t)$ is low relative to the volatility of $F(t, T_1)$.

¹⁴ In their Gaussian model, Casassus and Collin-Dufresne (2005) apply the more flexible “essentially affine” specification, which allows the market prices of risk to depend on all the state variables. However, market price of risk parameters are notoriously difficult to estimate and, indeed, they find that most of the additional parameters in the “essentially affine” specification are insignificant. For this reason, and because we are concerned mainly with the dynamics of volatility, for which the market prices of risk are the same in the “completely affine” and “essentially affine” specifications, we choose the more parsimonious specification.

In the SV2 specification, we set $\sigma_{S2} = 0$ and $\alpha_2 = 0$. This has several implications; futures prices are driven by only two factors, $S(t)$ and $x_1(t)$, and the second volatility factor is completely unspanned by the futures contracts and does not affect the instantaneous volatility of $S(t)$ and $y(t, T)$ (it will still affect expected future volatility and, therefore, option prices). We also impose $\eta_1 = 0$, $\kappa_{12} = -\kappa_1$, and $\kappa_{21} = 0$, so that the volatility process simplifies to

$$dv_1(t) = \kappa_1(v_2(t) - v_1(t))dt + \sigma_{v1}\sqrt{v_1(t)}dW_5^Q(t), \quad (29)$$

$$dv_2(t) = (\eta_2 - \kappa_2v_2(t))dt + \sigma_{v2}\sqrt{v_2(t)}dW_6^Q(t), \quad (30)$$

and $v_2(t)$ obtains the intuitive interpretation as the stochastic long-run mean of $v_1(t)$ (under the risk-neutral measure). This volatility process was suggested by Duffie, Pan, and Singleton (2000) in the context of pricing equity derivatives.¹⁵

In the SV1 specification, we set $\sigma_{S2} = 0$, $\alpha_2 = 0$, $\sigma_{v2} = 0$, and $\kappa_{12} = 0$, which implies that the volatility process simplifies to

$$dv_1(t) = (\eta_1 - \kappa_1v_1(t))dt + \sigma_{v1}\sqrt{v_1(t)}dW_5^Q(t). \quad (31)$$

There is only one volatility factor that may be partially spanned by the futures contracts.

3. Data and Estimation Procedure

3.1 Data

We estimate the three specifications, SV1, SV2, and SV2gen, on an extensive panel data set of crude oil futures and options trading on the NYMEX. The NYMEX crude oil derivatives market is the world's largest and most liquid commodity derivatives market. The range of maturities covered by futures and options and the range of strike prices on the options are also greater than for other commodities. This makes it an ideal market for studying commodity derivatives pricing. The raw data set consists of daily data from 2 January 1990 until 18 May 2006, on settlement prices, open interest, and daily volume for all available futures and options.^{16,17}

¹⁵ Several specifications in between SV2gen and SV2 are possible. For instance, with $\gamma_1 = \gamma_2 = \gamma$ one can show that futures prices are driven by only two factors, but both volatility factors may still contain a spanned and an unspanned component, and both factors may still affect the instantaneous volatility of $S(t)$ and $y(t, T)$. As another example, $\rho_{26} = \rho_{46} = 0$ implies that $v_2(t)$ is completely unspanned by the futures contracts but may still affect the instantaneous volatility of $S(t)$ and $y(t, T)$.

¹⁶ The NYMEX light, sweet crude oil futures contract trades in units of 1000 barrels. Prices are quoted as U.S. dollars and cents per barrel.

¹⁷ Note that all computations in the paper are based on settlement prices. Settlement prices for all contracts are determined by a "Settlement Price Committee" at the end of regular trading hours (currently 2:30 p.m. EST) and represent a very accurate measure of the true market prices at the time of close. Settlement prices are widely scrutinized by all market participants since they are used for marking to market all account balances.

The number of futures and options and their maximum maturities have increased significantly through the sample period. From the first trading day in the sample to the last trading day, the number of futures with positive open interest increased from 17 to 45, and the number of options with positive open interest increased from 77 to 1435. The maximum maturity among futures with positive open interest increased from 499 to 2372 days, while the maximum maturity among options with positive open interest increased from 164 to 2008 days.¹⁸

Liquidity has also increased significantly through the sample period. From the first to the last trading day, open interest (daily volume) for the first futures contract with more than 14 days to expiration increased from 66,925 (45,177) to 273,746 (86,622) contracts. The combined open interest (daily volume) for all options on that futures contract increased from 92,083 (16,427) to 376,694 (32,820) contracts.

We make two observations regarding liquidity. First, open interest for futures contracts tends to peak when expiration is a couple of weeks away, after which open interest declines sharply. Second, among futures and options with more than a couple of weeks to expiration, the first six monthly contracts tend to be very liquid. Beyond approximately six months, liquidity is concentrated in the contracts expiring in March, June, September, and December. Beyond approximately one year, liquidity is concentrated in the contracts expiring in December. Due to these liquidity patterns, we screen the available futures and options according to the following procedure: we discard all futures contracts with 14 or fewer days to expiration. Among the remaining, we retain the first six monthly contracts. Beyond these, we choose the first two contracts with expiration in either March, June, September, or December. Beyond these, we choose the next four contracts with expiration in December. This procedure leaves us with twelve generic futures contracts, which we label M1, M2, M3, M4, M5, M6, Q1, Q2, Y1, Y2, Y3, and Y4.

We use options on the first eight futures contracts, M1–Q2. We do not consider options on the remaining futures contracts, Y1–Y4, for two reasons. First, the quasi-analytical expression for option prices that we develop in Section 2.3 does not take stochastic interest rates into account. While it appears that ignoring stochastic interest rates results in negligible pricing errors for short-term and medium-term options, it may result in nonnegligible pricing errors for long-term options. Second, and more important, the options are American, whereas our pricing formula is for European options.¹⁹ For computational reasons, estimation is feasible only with European options, necessitating a conversion

¹⁸ Futures contracts expire on the third business day prior to the 25th calendar day of the month preceding the delivery month. If the 25th calendar day of the month is a non-business day, expiration is on the third business day prior to the business day preceding the 25th calendar day. Options expire three business days prior to the expiration of the underlying futures contract.

¹⁹ Recently, NYMEX has introduced European crude oil options. However, the trading history is much shorter and liquidity is much lower than for the American options.

of American prices to European prices. This requires an approximation of the early exercise premium, described in detail in Appendix B, and since the size of the early exercise premium as a fraction of the total option price increases with option maturity, any error in the early exercise approximation becomes more serious for longer-term options. Considering options on the first eight futures contracts appears to strike a reasonable balance between including information from the maturity dimension of option prices, while limiting our exposure to the approximation errors associated with the American-to-European conversions and not explicitly modeling stochastic interest rates.

For each option maturity, we consider eleven moneyness intervals, 0.78–0.82, 0.82–0.86, 0.86–0.90, 0.90–0.94, 0.94–0.98, 0.98–1.02, 1.02–1.06, 1.06–1.10, 1.10–1.14, 1.14–1.18, and 1.18–1.22, where moneyness is defined as option strike divided by the price of the underlying futures contract. Among the options within a given moneyness interval, we select the one that is closest to the mean of the interval. To minimize the effect of any errors in the early exercise approximation, we use only out-of-the-money (OTM) and at-the-money (ATM) options (i.e., puts with moneyness less than one and calls with moneyness greater than one). An added benefit of this choice is that OTM options tend to be more liquid than in-the-money (ITM) options.²⁰

After sorting the data, we are left with 45,517 futures contracts and 233,104 option contracts over a period of 4082 business days. On a given trading day, the number of futures is between 8 and 12, while the number of options is between 23 and 87.²¹

Figure 1 displays the futures data. The run-up in crude oil prices since 2002 is striking. Using the M1 futures contract as a proxy for the spot price, the Q2 futures contract is backwarddated 82.6% of the time and strongly backwarddated 66.3% of the time.²²

Figure 2 displays the ATM lognormal implied volatilities for options on the first eight of the futures contracts. Volatility is evidently stochastic. The question is the extent to which volatility is unspanned.

3.2 Estimation procedure

The estimation approach is quasi-maximum likelihood (QML) in conjunction with the extended Kalman filter. To apply the Kalman filter we cast the model in state space form, which consists of a measurement equation and a transition

²⁰ In addition, we consider only options that have open interest in excess of 100 contracts and options with prices larger than 0.10 dollars. The reason for the latter is that prices are reported with a precision of 0.01 dollars in the data set.

²¹ The discount function $P(t, T)$ in Equation (24) is obtained by fitting a Nelson and Siegel (1987) curve each trading day to a LIBOR/swap curve consisting of the 1-month, 3-month, 6-month, 9-month, and 12-month LIBOR rates and the 2-year swap rate.

²² Let $S(t)$ denote the time- t spot price and $F(t, T)$ [$P(t, T)$] the time- t price of a futures contract (zero-coupon bond) with maturity $T - t$. The futures contract is backwarddated if $S(t) - P(t, T)F(t, T) > 0$ and strongly backwarddated if $S(t) - F(t, T) > 0$. The numbers reported here are slightly lower than those reported by Litzenberger and Rabinowitz (1995) for an earlier sample.




Figure 1
Prices of futures contracts

Prices of M1, M2, M3, M4, M5, M6, Q1, Q2, Y1, Y2, Y3, and Y4 futures contracts. Our data set spans 4082 trading dates from 2 January 1990 to 18 May 2006. To avoid cluttering the figure, we display only the futures term structures on Wednesdays.

equation. The measurement equation describes the relationship between the state variables and the prices of futures and options, while the transition equation describes the discrete-time dynamics of the state variables.

The state vector, X_t , is given by

$$X_t = (s_t, x_{1,t}, x_{2,t}, \phi_{1,t}, \phi_{2,t}, v_{1,t}, v_{2,t})'. \tag{32}$$

The continuous-time dynamics of X_t under the actual measure are obtained from the risk-neutral dynamics in Equations (3), (4), (10), (11), (12), and (18) along with the market price of risk specification described in Section 2.4. The general form of the transition equation is given by

$$\begin{aligned} X_{t+1} &= \Phi_0 + \Phi_X X_t + w_{t+1}, \quad w_{t+1} \text{ iid,} \\ E[w_{t+1}] &= 0, \\ \text{Var}[w_{t+1}] &= Q_0 + Q_{v1} v_{1,t} + Q_{v2} v_{2,t}, \end{aligned} \tag{33}$$

where Φ_0 , Φ_X , Q_0 , Q_{v1} , and Q_{v2} can be computed in closed form.




Figure 2

ATM lognormal implied volatility of futures options

ATM lognormal implied volatility of options on the M1, M2, M3, M4, M5, M6, Q1, and Q2 futures contracts. Implied volatilities are computed from option prices by inverting the Barone-Adesi and Whaley (1987) formula. Our data set spans 4082 trading dates from 2 January 1990 to 18 May 2006. To avoid cluttering the figure, we display only the volatility term structures on Wednesdays.

The measurement equation is given by

$$z_t = h(X_t) + u_t, \quad u_t \sim \text{iid } N(0, \Omega), \tag{34}$$

where z_t is the data vector, h is the pricing function, and u_t is a vector of iid. Gaussian measurement errors with covariance matrix Ω . The measurement equation is constructed from Equation (17), which relates log futures prices to $s(t)$, $x_1(t)$, $x_2(t)$, $\phi_1(t)$, and $\phi_2(t)$ through an affine expression, and Equations (20), (24), and (25), which relate option prices to $v_1(t)$ and $v_2(t)$ through nonlinear expressions.²³

Ideally, we would like to fit the model directly to implied Black (1976) (i.e., lognormal) volatilities, which are more stable than prices along the moneyness, maturity, and time-series dimensions. This is not practical, however,

²³ Note that option prices do not depend directly on $s(t)$, $x_1(t)$, $x_2(t)$, $\phi_1(t)$, and $\phi_2(t)$ since we price options based on the actual, rather than fitted, futures prices. This implies that an imperfect fit to the futures contracts does not get reflected in option prices, which in turn should provide us with a cleaner estimate of the volatility process.

since computing implied volatilities from prices requires a numerical inversion for each option, which would add an extra layer of complexity to the estimation procedure. Instead, we fit the model to option prices scaled by their Black (1976) vegas (i.e., the sensitivities of the option prices to variations in lognormal volatilities). This essentially converts option pricing errors in terms of prices into option pricing errors in terms of implied volatilities through a linear approximation.²⁴

More formally, suppose at time t we observe m futures prices, $F_{t,1}, \dots, F_{t,m}$, and n option prices, $O_{t,1}, \dots, O_{t,n}$, with associated Black (1976) vegas, $\mathcal{V}_{t,1}, \dots, \mathcal{V}_{t,n}$. Then z_t is given by

$$z_t = (\log F_{t,1}, \dots, \log F_{t,m}, O_{t,1}/\mathcal{V}_{t,1}, \dots, O_{t,n}/\mathcal{V}_{t,n}). \tag{35}$$

The dimension of the z_t -vector varies over time, but the Kalman filter easily handles missing observations.

We make two assumptions to reduce the number of parameters in Ω . First, we assume that the measurement errors are cross-sectionally uncorrelated (that is, Ω is diagonal). Second, we assume that one variance applies to all measurement errors for the log of futures prices, and that another variance applies to all measurement errors for scaled option prices.²⁵

The model derived in Section 2 is time-inhomogeneous and fits the initial futures curve by construction. For the purpose of estimation, it is more convenient to work with the model's time-homogeneous counterpart. We, therefore, assume that the initial forward cost of carry curve, $y(0, t)$, is flat and equal to a constant φ . This amounts to setting $y(0, t) = \varphi$ in Equation (18) and setting $\log F(0, T) - \log F(0, t) = \varphi(T - t)$ in Equation (17). φ is estimated as part of the estimation procedure and has the interpretation as the infinite-maturity forward cost of carry rate; see Equation (9).

Application of the extended Kalman filter involves linearizing the h -function in Equation (34) and making the assumption that the disturbance term w_{t+1} in Equation (33) is Gaussian. From the Kalman filter recursions, we can compute the likelihood function.²⁶

²⁴ To see this, let $\tilde{O}_{t,i}$ and $\hat{O}_{t,i}$ denote the fitted and actual option prices, respectively, for option i at time t , and let $\tilde{\sigma}_{t,i}$ and $\hat{\sigma}_{t,i}$ denote the corresponding lognormal implied volatilities. Furthermore, let $\hat{\mathcal{V}}_{t,i} \equiv \frac{\partial O}{\partial \sigma} |_{\sigma=\hat{\sigma}_{t,i}}$ denote the Black (1976) vega computed at the true option price. Then

$$\tilde{O}_{t,i} \approx \hat{O}_{t,i} + \hat{\mathcal{V}}_{t,i} (\tilde{\sigma}_{t,i} - \hat{\sigma}_{t,i}) \Leftrightarrow \frac{\tilde{O}_{t,i} - \hat{O}_{t,i}}{\hat{\mathcal{V}}_{t,i}} \approx \tilde{\sigma}_{t,i} - \hat{\sigma}_{t,i}.$$

A similar approach is used by Carr and Wu (2007); Bakshi, Carr, and Wu (2008); and Trolle and Schwartz (2009), among others.

²⁵ Note that the assumption of normally distributed additive measurement errors on the log of futures prices implies lognormally distributed multiplicative measurement errors on futures prices.

²⁶ For completeness, the extended Kalman filter recursions and the likelihood function are stated in Appendix C; Harvey (1989) and Hamilton (1994) are classic references. In the Appendix, we also comment on the consistency of the estimation procedure and discuss various numerical issues related to the implementation of the model and the optimization of the likelihood function.

4. Results

4.1 Parameter estimates

Tables 1 and 2 display parameter estimates of the SV1, SV2, and SV2gen specifications, both for the entire sample 1990–2006 and for two subsamples 1990–1997 and 1998–2006.²⁷ Estimating the specifications on subsamples shows the stability of the parameter estimates and also allows us to determine the out-of-sample pricing and hedging performance of the model. In our discussion of the parameter estimates, we will mainly consider the results for the full-sample estimation. However, we will also briefly discuss the stability of the parameter estimates across subsamples.

The volatility process is moderately persistent, under the risk-neutral measure, in the SV1 specification ($\kappa_1 = 1.0125$). In the SV2 specification, $v_1(t)$ is very highly mean-reverting toward $v_2(t)$ ($\kappa_1 = 7.9035$), while $v_2(t)$ is moderately persistent ($\kappa_2 = 0.9240$). Hence, $v_1(t)$ captures transitory shocks to volatility, while $v_2(t)$ captures more persistent shocks to volatility, which in turn implies that $v_1(t)$ primarily affects prices of very short-term options, while $v_2(t)$ affects prices of all options. In the SV2gen specification, the estimate of κ_{21} is close to zero, while the estimate of κ_{12} is close to $-\kappa_1$, which suggests that imposing $\kappa_{21} = 0$ and $\kappa_{12} = -\kappa_1$ in the SV2 specification will not significantly affect the pricing of options. In both the SV2gen and SV2 specifications, σ_{v_1} is significantly higher than σ_{v_2} , and since $v_1(t)$ and $v_2(t)$ are of similar magnitudes,²⁸ this implies that shocks to the transitory component of volatility are much more variable than shocks to the more persistent component of volatility.

The SV2gen specification allows both volatility factors to contain a spanned component. However, the correlations between innovations to the volatility factors and innovations to the spot price and the forward cost of carry curve are very low, with $\rho_{15} = -0.039$, $\rho_{35} = -0.103$, $\rho_{26} = -0.131$, and $\rho_{46} = -0.001$, implying that volatility is predominantly unspanned by the futures contracts. This is consistent with model-free regression-based results, reported in the NBER working paper version of the article, that strongly suggest the presence of unspanned stochastic volatility in the crude oil market.²⁹ The SV2

²⁷ The asymptotic covariance matrix of the estimated parameters is computed from the outer-product of the first derivatives of the likelihood function. Theoretically, it would be more appropriate to compute the asymptotic covariance matrix from both the first and second derivatives of the likelihood function. In reality, however, the second derivatives of the likelihood function are somewhat numerically unstable. In general, parameters identified under Q are strongly identified with very low standard errors. In contrast, the market price of risk parameters is quite imprecisely estimated, and many of them are statistically insignificant at conventional levels.

²⁸ The average values of $v_1(t)$ and $v_2(t)$ in the SV2 (SV2gen) specification are 3.367 (2.900) and 3.151 (4.274), respectively.

²⁹ For brevity, we excluded this analysis from the final version of the article. Instead, we briefly summarize it here. For different option maturities, we regress returns on ATM option straddles and changes in implied volatilities—both reasonable proxies for changes in the true but unobservable volatility—on futures returns and find low R^2 s, indicating that most volatility risk cannot be hedged by trading in the futures contracts. Furthermore, there is large common variation in regression residuals across option maturities, strongly indicating the presence of a few unspanned volatility factors. This holds true regardless of the number of options included in the analysis or the length of the sample. These findings are model-free in the sense that no pricing model is applied.

Table 1
Parameter estimates

	1990–2006			1990–1997			1998–2006		
	SV1	SV2	SV2gen	SV1	SV2	SV2gen	SV1	SV2	SV2gen
κ_1	1.0125 (0.0123)	7.9035 (0.0127)	6.7965 (0.0251)	1.1599 (0.0216)	9.7846 (0.0366)	8.2222 (0.0497)	1.9404 (0.0142)	4.6820 (0.0123)	4.0333 (0.0227)
κ_{12}	–	–7.9035 (–0.0127)	–5.4153 (–1.7547)	–	–9.7846 (–0.0366)	–10.2910 (–4.1150)	–	–4.6820 (–0.0123)	–3.1057 (–0.0941)
σ_{v1}	2.8051 (0.0211)	7.3088 (0.0447)	8.7350 (1.2901)	2.8635 (0.0325)	3.1844 (0.0134)	7.4003 (0.2779)	1.7368 (0.0106)	3.4402 (0.0155)	5.3689 (0.0575)
σ_{S1}	0.2289 (0.0017)	0.2099 (0.0013)	0.2085 (0.0307)	0.2268 (0.0025)	0.5038 (0.0015)	0.2094 (0.0079)	0.5024 (0.0021)	0.3827 (0.0016)	0.3090 (0.0035)
α_1	0.1373 (0.0024)	0.1542 (0.0017)	0.3156 (0.0464)	0.1927 (0.0043)	0.5511 (0.0069)	0.3860 (0.0151)	0.3594 (0.0049)	0.3380 (0.0031)	0.4317 (0.0058)
γ_1	0.7796 (0.0011)	0.8098 (0.0011)	1.6446 (0.0030)	1.2961 (0.0033)	1.3917 (0.0031)	1.9080 (0.0087)	0.8125 (0.0012)	0.8378 (0.0013)	1.3916 (0.0037)
ρ_{13}	–0.8797 (0.0040)	–0.8982 (0.0023)	–0.8914 (0.0028)	–0.9267 (0.0039)	–0.8892 (0.0042)	–0.8720 (0.0031)	–0.8851 (0.0033)	–0.9256 (0.0014)	–0.8992 (0.0024)
ρ_{15}	–0.0912 (0.0018)	–0.0522 (0.0007)	–0.0390 (0.0009)	–0.0082 (0.0020)	–0.0113 (0.0013)	–0.0153 (0.0015)	–0.0928 (0.0021)	–0.0691 (0.0007)	–0.0311 (0.0010)
ρ_{35}	–0.1128 (0.0116)	–0.2049 (0.0043)	–0.1031 (0.0044)	–0.0443 (0.0109)	–0.0374 (0.0081)	–0.0018 (0.0066)	–0.2391 (0.0102)	–0.2004 (0.0041)	–0.0300 (0.0044)
λ_1	0.3423 (0.1900)	0.2939 (0.1742)	–0.0209 (0.1884)	0.2234 (0.2893)	0.5577 (0.3312)	0.4964 (0.1444)	0.9827 (0.5455)	0.7125 (0.4726)	–0.6043 (0.3817)
λ_3	–0.2391 (0.1992)	–0.1768 (0.1936)	–0.1293 (0.1982)	–0.3765 (0.2913)	–0.7950 (0.5083)	–0.3370 (0.2144)	–0.5273 (0.6393)	–0.3213 (0.5357)	–0.0656 (0.4187)
λ_5	0.2344 (0.1395)	0.0709 (0.0881)	–0.1752 (0.0905)	0.2452 (0.1604)	–2.5174 (0.2284)	–1.2888 (0.1051)	0.3167 (0.6151)	0.0317 (0.2608)	–0.1461 (0.2393)
φ	0.0054 (0.0001)	0.0059 (0.0001)	0.0087 (0.0001)	0.0181 (0.0001)	0.0170 (0.0001)	0.0051 (0.0002)	–0.0021 (0.0001)	–0.0009 (0.0001)	0.0063 (0.0002)

Maximum-likelihood estimates of the SV1, SV2, and SV2gen specifications for the full-sample period and two subsamples. The three periods are 2 January 1990 to 18 May 2006 (4082 daily observations), 2 January 1990 to 31 December 1997 (2008 daily observations), and 2 January 1998 to 18 May 2006 (2074 daily observations). Outer-product standard errors are in parentheses. σ_{futures} denotes the standard deviation of log futures price measurement errors and σ_{options} denotes the standard deviation of scaled option price measurement errors. For the specifications to be identified, we set $\eta_1 = \eta_2 = 1$ in SV2gen, $\eta_1 = 0$ and $\eta_2 = 1$ in SV2, and $\eta_1 = 1$ in SV1. This table is continued in Table 2.

Table 2
Parameter estimates (Table 1 cont.)

	1990–2006			1990–1997			1998–2006		
	SV1	SV2	SV2gen	SV1	SV2	SV2gen	SV1	SV2	SV2gen
κ_2	–	0.9240 (0.0066)	0.8960 (0.0069)	–	4.8837 (0.0334)	4.3744 (0.0332)	–	1.0343 (0.0057)	1.4450 (0.0102)
κ_{21}	–	–	–0.0522 (–0.0167)	–	–	–0.0328 (–0.0063)	–	–	–0.0294 (–0.0143)
σ_{v2}	–	1.6109 (0.0127)	1.8223 (0.0300)	–	1.0623 (0.0102)	1.0207 (0.0075)	–	0.9349 (0.0126)	1.8441 (0.0127)
σ_{s2}	–	–	0.0704 (0.0011)	–	–	0.0006 (0.0055)	–	–	0.1703 (0.0009)
α_2	–	–	0.0518 (0.0010)	–	–	0.1503 (0.0036)	–	–	0.0749 (0.0016)
γ_2	–	–	0.5473 (0.0019)	–	–	0.7504 (0.0054)	–	–	0.5365 (0.0028)
ρ_{24}	–	–	–0.1270 (0.0088)	–	–	–0.0908 (0.0122)	–	–	–0.0633 (0.0120)
ρ_{26}	–	–	–0.1312 (0.0045)	–	–	–0.0093 (0.0097)	–	–	–0.3078 (0.0044)
ρ_{46}	–	–	–0.0009 (0.0185)	–	–	0.0577 (0.0218)	–	–	–0.0010 (0.0275)
λ_2	–	–	0.6877 (0.1345)	–	–	–2.3536 (0.9580)	–	–	1.8619 (0.2148)
λ_4	–	–	–0.0984 (0.1391)	–	–	0.2686 (0.4155)	–	–	–0.3755 (0.3256)
λ_6	–	0.3901 (0.0817)	0.3780 (0.0728)	–	3.1460 (0.2197)	2.9110 (0.1925)	–	0.2929 (0.3542)	0.5169 (0.4183)
σ_{futures}	0.0150 (0.0000)	0.0152 (0.0000)	0.0051 (0.0000)	0.0106 (0.0000)	0.0107 (0.0000)	0.0051 (0.0000)	0.0151 (0.0000)	0.0155 (0.0000)	0.0052 (0.0000)
σ_{options}	0.0265 (0.0000)	0.0164 (0.0000)	0.0163 (0.0000)	0.0286 (0.0000)	0.0170 (0.0000)	0.0169 (0.0000)	0.0235 (0.0000)	0.0131 (0.0000)	0.0130 (0.0000)
$\log L$	–652255.2	–546317.8	–503544.4	–246108.8	–208012.7	–191497.0	–389080.3	–306585.4	–282874.0

This table is the continuation of Table 1.

specification imposes that the second volatility factor is completely unspanned by the futures contracts, which is not an unreasonable restriction given that ρ_{26} and ρ_{46} are both close to zero in the SV2gen specification. Since $\rho_{15} = -0.052$ and $\rho_{35} = -0.205$ in the SV2 specification, and $\rho_{15} = -0.091$ and $\rho_{35} = -0.113$ in the SV1 specification, the finding that volatility is mostly unspanned holds true regardless of the precise model specification.³⁰

Both volatility factors may affect the instantaneous volatility of $S(t)$ and $y(t, T)$ in the SV2gen specification. However, since σ_{S1} and α_1 are significantly larger than σ_{S2} and α_2 , $v_1(t)$ accounts for most of the instantaneous volatility of the spot price and the spot cost of carry. At the same time, since $\gamma_1 > \gamma_2$, the proportion of the instantaneous volatility of the forward cost of carry that $v_1(t)$ accounts for decreases with maturity. Straightforward computations show that $v_1(t)$, on average, accounts for 85.6% of the instantaneous variance of the spot price, 96.2% of the instantaneous variance of the spot cost of carry, and 76.8% of the instantaneous variance of the one-year forward cost of carry. Furthermore, since $v_1(t)$ is much more volatile than $v_2(t)$, $v_1(t)$ accounts for virtually all of the variation in the instantaneous volatility of the spot price and the front end of the forward cost of carry curve. As a result, the instantaneous volatility of the spot price and the front end of the forward cost of carry curve is very highly correlated. This suggests that the SV2 specification, where $v_1(t)$ is the sole driver of the instantaneous volatility of the spot price and the forward cost of carry curve, may be reasonable in the context of pricing short-term or medium-term options. However, for pricing long-term options, there may be more of a difference between the SV2gen and SV2 specifications.³¹

The 1990–1991 Gulf Crisis is by far the largest shock to the crude oil derivatives market in our data set. As can be seen from Figure 2, implied volatility increased dramatically during this period, reaching more than 100% for short-dated options. Not surprisingly, this event affects parameter estimates. Note, in particular, that the mean-reversion coefficients in the SV2 and SV2gen specifications are significantly higher in the first subsample than in the second subsample.

Judging from the loglikelihood values, the parameter restrictions imposed in the SV2 and SV1 specifications are rejected at extreme levels of statistical

³⁰ The negative sign on ρ_{35} (and ρ_{26} in the case of the SV2gen specification) implies that volatility is (weakly) positively correlated with the degree of backwardation. This is consistent with the findings of Litzenberger and Rabinowitz (1995) and Routledge, Seppi, and Spatt (2000).

³¹ Since the estimation relies on both time-series and cross-sectional information, it will not necessarily be the case that the model-implied conditional volatilities of the spot price and cost of carry will be consistent with time-series estimates of the conditional volatilities. To investigate this issue, we back out time series of $S(t)$ and $\delta(t)$ from the Kalman filtered state variables, estimate a GARCH(1,1) model for each of them, and compare the fitted GARCH volatilities with the corresponding model-implied conditional volatilities. For the SV2gen specification, we find correlations of 0.847 and 0.769 for $S(t)$ and $\delta(t)$, respectively, confirming that this specification is indeed internally consistent (we would not expect to find a perfect correlation, since the GARCH volatilities are themselves model-dependent benchmarks). Not surprisingly, the correlations are lower for the nested specifications. Furthermore, for the SV2gen specification, we find a correlation of 0.924 between the fitted GARCH volatilities of $S(t)$ and $\delta(t)$, supporting the conclusion that the volatility of the spot price and the front end of the forward cost of carry curve is very highly correlated.

significance, which is not surprising, given the very large number of contracts used for estimation. More interesting is whether there are economically significant differences between the models. Table 2 shows that σ_{futures} , the standard deviation of the log futures pricing errors, is virtually the same for the SV1 and SV2 specifications, but significantly lower for the SV2gen specification, whereas σ_{options} , the standard deviation of scaled option pricing errors, is virtually the same for the SV2 and SV2gen specifications, but significantly higher for the SV1 specification. In other words, the parameter restrictions imposed in the SV2 specification affect the fit to futures prices but have negligible impact on the fit to option prices, while the additional parameter restrictions imposed in the SV1 specification affect the fit to option prices but not the fit to futures prices. In the next section, we compare the pricing performance of the different specifications in more detail.

4.2 Pricing performance

We compute the fitted prices of futures and options based on the filtered state variables. For the futures contracts, the pricing errors are given by the differences between the fitted and actual prices divided by the actual prices. For the option contracts, the pricing errors are given by the differences between fitted and actual lognormal implied volatilities. On each day in the sample, we compute the root mean squared pricing errors (RMSEs) of the futures and options trading on that day. This way we construct time series of RMSEs of futures and options for each specification and each set of parameter estimates.

To compare the pricing performance of the different specifications, we use the approach of Diebold and Mariano (1995). Suppose, say, the SV1 and SV2 specifications generate time series of option RMSEs, $RMSE_{\text{options}}^{SV1}(t)$ and $RMSE_{\text{options}}^{SV2}(t)$. We then compute the mean of the difference $RMSE_{\text{options}}^{SV2}(t) - RMSE_{\text{options}}^{SV1}(t)$ and the associated t -statistics. A significantly negative mean implies that the SV2 specification has a significantly better fit to options than the SV1 specification (according to the RMSE criterion).

When parameter estimates are obtained from the full sample, we obtain an in-sample RMSE time series. When parameter estimates are obtained from the first (second) subsample, we compute an in-sample RMSE time series for the first (second) subsample and an out-of-sample RMSE time series for the second (first) subsample.³² Therefore, for each specification (SV1, SV2, and SV2gen) and contract type (futures and options), we have five RMSE time series—three in-sample and two out-of-sample. Table 3 displays the means of these RMSE time series and compares the pricing performance of the different

³² An alternative way of computing out-of-sample RMSEs would be to use all information up to time t when computing the RMSE at time $t + 1$. This would require reestimating the model at every observation date, which would be a very demanding task. In some sense, the out-of-sample RMSEs that we report provide an upper bound on the model's true out-of-sample RMSEs, since we do not update the parameters of the model.

Table 3
Overall comparison of SV1, SV2, and SV2gen specifications

	1990–2006 <i>IS</i>	1990–1997 <i>IS</i>	1998–2006 <i>IS</i>	1990–1997 <i>OOS</i>	1998–2006 <i>OOS</i>
Panel A: Futures contracts					
SV1	1.227	0.854	1.286	1.255	2.599
SV2	1.247	0.861	1.316	1.276	2.789
SV2gen	0.391	0.369	0.431	0.400	0.508
SV2-SV1	0.021 (1.384)	0.007 (0.356)	0.030 (1.070)	0.020 (0.859)	0.191*** (4.224)
SV2gen-SV2	-0.856*** (-15.007)	-0.492*** (-9.629)	-0.885*** (-11.242)	-0.876*** (-9.336)	-2.281*** (-12.357)
Panel B: Options on futures contracts					
SV1	2.103	1.891	2.098	3.031	2.550
SV2	1.316	1.222	1.204	1.574	2.508
SV2gen	1.307	1.216	1.197	1.558	2.459
SV2-SV1	-0.787*** (-10.511)	-0.668*** (-5.265)	-0.894*** (-7.761)	-1.457*** (-10.473)	-0.042 (-0.328)
SV2gen-SV2	-0.010 (-1.512)	-0.006 (-1.251)	-0.007 (-0.453)	-0.015 (-0.570)	-0.050*** (-3.260)

Means of root mean squared pricing errors (RMSEs) for futures (Panel A) and options (Panel B). The three periods are 2 January 1990 to 18 May 2006 (4082 daily observations), 2 January 1990 to 31 December 1997 (2008 daily observations), and 2 January 1998 to 18 May 2006 (2074 daily observations). *IS* refers to in-sample, i.e., the fit during the estimation period. *OOS* refers to out-of-sample, i.e., the fit during the first (second) subsample when the model is estimated on the second (first) subsample. The futures pricing errors are defined as the differences between the fitted and actual prices divided by the actual prices and reported in percentages. The option pricing errors are defined as the differences between fitted and actual lognormal implied volatilities and reported in percentages. *T*-statistics, corrected for serial correlation up to 50 lags, are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

specifications.³³ For the futures contracts, there are no statistically significant differences in mean RMSEs between SV1 and SV2 (except in the second out-of-sample period), which is not surprising, given that the two specifications differ only in terms of the volatility dynamics. In contrast, the SV2gen specification has significantly lower mean RMSEs, due to the extra factor driving futures prices. For instance, the mean in-sample RMSE for the full sample is around 1.25% for the SV1 and SV2 specifications but drops to 0.39% for the SV2gen specification.³⁴ The mean out-of-sample RMSEs, while still fairly small, are larger than the corresponding mean in-sample RMSEs, reflecting that the parameters do change across the two samples.

For the option contracts, the mean RMSEs are significantly lower for the SV2 specification, compared with the SV1 specification. This holds true for the in-sample RMSEs and for the out-of-sample RMSE in the first subsample (for the out-of-sample RMSE in the second subsample the improvement is not significant). At the same time, the SV2gen specification does not significantly reduce the mean RMSEs (except for the out-of-sample RMSE in the second

³³ The *t*-statistics reported in this table, as well as Tables 5 and 6, are computed from Newey and West (1987) standard errors with 50 lags to correct for heteroscedasticity and autocorrelation. The results are robust to variations in the lag length.

³⁴ For comparison, Schwartz (1997) reports RMSEs of around 1% for his two- and three-factor models fitted to crude oil futures.

subsample), suggesting that the greater complexity of this specification is not justified when it comes to pricing the options in this data set, which is also what our analysis of the parameter estimates indicated. For instance, the mean in-sample RMSE for the full sample drops from 2.10% for the SV1 specification to 1.32% for the SV2 specification and 1.31% for the SV2gen specification. Again, the mean out-of-sample RMSEs are larger than the corresponding mean in-sample RMSEs, but not dramatically so. We conclude that while a second volatility factor is important, it appears sufficient to model this factor in a fairly parsimonious fashion in order to price short-term and medium-term options.³⁵

While these results show that a model with two or three factors driving futures prices needs at least two additional, largely unspanned, volatility factors to match option prices, they do not rule out that a five-factor model, where all the factors are spanned by the futures contracts, would perform even better. Given that the model-free analysis, described in the NBER working paper version of the article and summarized in footnote 29, corroborates the existence of unspanned stochastic volatility, we doubt that this is the case. Nevertheless, to check that the volatility factors are only weakly related to the futures term structure, we consider the SV2gen specification, estimated on the entire period, and regress innovations to the volatility factors on the returns to all M1–Q2 futures contracts.³⁶ The R^2 s are 0.040 and 0.012 for the first and the second volatility factor, respectively, confirming that these are indeed largely unspanned by the futures term structure.

Figure 3 displays the time series of the RMSEs of futures and options when the specifications are estimated on the entire data set. It also highlights the dates of four major shocks to the crude oil market: the Iraqi invasion of Kuwait on 2 August 1990, the beginning of the U.S.-led liberation of Kuwait (Operation Desert Storm) on 17 January 1991, the 11 September 2001 terrorist attacks, and the U.S.-led invasion of Iraq on 20 March 2003. Since the futures RMSEs for the SV1 and SV2 specifications and the option RMSEs for the SV2 and SV2gen specifications are visually indistinguishable, we display the RMSEs only for the SV1 and SV2gen specifications. The futures RMSEs for the SV1 (and SV2) specifications mostly fluctuate between 0.5% and 2.5%, but do reach almost 5% at the beginning of Operation Desert Storm as well as in early 1996. In contrast, the futures RMSEs for the SV2gen specification mostly fluctuate between 0% and 1%.

The option RMSEs for the SV1 specification spike around all four events that we have highlighted plus a few others. Otherwise, they lie mostly in a

³⁵ The finding of multiple volatility factors is consistent with studies on other asset markets. For instance, Bates (2000) and Trolle and Schwartz (2009) develop and test models for pricing equity and interest rate derivatives, respectively, with multi-factor specifications for volatility. Both studies find evidence for multiple volatility factors with important unspanned components. Various papers, applying a diverse set of time-series techniques, also find that volatility is driven by multiple factors; see, e.g., Andersen et al. (2001b); Bollerslev and Zhou (2002); and Alizadeh, Brandt, and Diebold (2002) for studies on foreign exchange rate volatility, and Engle and Lee (1999); Andersen et al. (2001a); and Chernov et al. (2003) for studies on equity volatility.

³⁶ We use the M1–Q2 futures contracts because they are available throughout the entire sample period.




Figure 3
Time series of RMSEs of futures and options

Panel A shows time series of root mean squared errors (RMSEs) of the percentage differences between fitted and actual futures prices. Panel B shows time series of RMSEs of the differences between fitted and actual lognormal implied option volatilities. The grey lines refer to the SV1 specification and the black lines refer to the SV2gen specification. Both specifications are estimated on the entire data set. The vertical dotted lines mark the Iraqi invasion of Kuwait on 2 August 1990, the beginning of the U.S.-led liberation of Kuwait (Operation Desert Storm) on 17 January 1991, the 11 September 2001 terrorist attacks, and the U.S.-led invasion of Iraq on 20 March 2003, respectively. The number of futures at a given date varies between 8 and 12. The number of options at a given date varies between 23 and 87. Each time series consists of 4082 daily observations from 2 January 1990 to 18 May 2006.

range between 1% and 4%. In contrast, the option RMSEs for the SV2gen (and SV2) specifications mainly spike during the period from the Iraqi invasion of Kuwait to the beginning of Operation Desert Storm. Outside of this turbulent period, the RMSEs are fairly stable and mostly fluctuate between 1% and 2%.

Table 4 reports the MAEs of the options within each moneyness–maturity category, when the SV1 and SV2 specifications are estimated on the entire data set (since the SV2 and SV2gen specifications have virtually the same option pricing performance; in the remainder of the section, we display results only for the SV1 and SV2 specifications). For the SV1 specification, the MAEs

Table 4
Mean absolute pricing errors of options within each moneyness–maturity category

Moneyness	Model	Contract							
		M1	M2	M3	M4	M5	M6	Q1	Q2
0.78–0.82	SV1	5.51	2.58	1.78	1.71	1.79	1.98	2.35	2.83
	SV2	2.85	1.97	1.63	1.56	1.49	1.57	1.64	2.04
0.82–0.86	SV1	3.93	2.17	1.51	1.43	1.56	1.85	2.21	2.71
	SV2	2.11	1.49	1.39	1.28	1.22	1.25	1.40	1.73
0.86–0.90	SV1	3.38	1.94	1.15	1.16	1.33	1.65	2.07	2.70
	SV2	1.79	1.24	1.07	1.00	0.98	0.99	1.12	1.52
0.90–0.94	SV1	3.13	1.66	0.93	0.93	1.19	1.50	1.98	2.54
	SV2	1.55	0.95	0.84	0.75	0.70	0.71	0.90	1.39
0.94–0.98	SV1	2.98	1.59	0.84	0.82	1.13	1.39	1.90	2.37
	SV2	1.43	0.82	0.74	0.64	0.54	0.51	0.73	1.21
0.98–1.02	SV1	3.17	1.64	0.90	0.86	1.15	1.41	1.89	2.31
	SV2	1.50	0.89	0.80	0.68	0.59	0.57	0.70	1.12
1.02–1.06	SV1	2.99	1.67	0.92	0.83	1.10	1.38	1.83	2.27
	SV2	1.34	0.97	0.84	0.73	0.64	0.61	0.72	1.09
1.06–1.10	SV1	3.23	1.81	1.05	0.91	1.15	1.39	1.81	2.25
	SV2	1.26	1.09	0.96	0.83	0.75	0.75	0.78	1.17
1.10–1.14	SV1	3.57	2.13	1.24	1.02	1.22	1.37	1.82	2.25
	SV2	1.29	1.26	1.12	0.94	0.89	0.88	0.87	1.21
1.14–1.18	SV1	4.09	2.53	1.49	1.15	1.24	1.41	1.81	2.37
	SV2	1.55	1.55	1.33	1.07	0.99	1.02	0.95	1.21
1.18–1.22	SV1	4.98	2.78	1.75	1.33	1.36	1.46	1.74	2.11
	SV2	2.17	1.79	1.54	1.24	1.14	1.09	0.96	1.19

The table reports the mean absolute pricing errors (MAEs) of the options within each moneyness–maturity category, when the SV1 and SV2 specifications are estimated on the entire data set. The pricing errors are defined as the differences between fitted and actual lognormal implied volatilities and reported in percentages. Moneyness is defined as option strike divided by the price of the underlying futures contract. Each MAE is computed on the basis of a maximum of 4082 daily observations from 2 January 1990 to 18 May 2006.

lie between 0.82% and 5.51%. It has a good fit to near-ATM options with maturities of three to four months, but the pricing performance deteriorates as we move toward the edges of the volatility surface. In particular, the specification has difficulty matching short-term OTM put and call options. For the SV2 specification, the MAEs lie in a range from 0.51% to 2.85%, and it performs much better at pricing options across maturity and moneyness.

Table 5 compares the SV1 and SV2 specifications in terms of their ability to price options within each moneyness–maturity category. It reports the mean differences in absolute pricing errors between the two specifications along with the associated *t*-statistics. We observe that, except for OTM calls on the M4 futures contract, the improvement in pricing performance for the SV2 specification is statistically significant for all options.

The improvement of SV2 over SV1 in the moneyness dimension is also clear from Figure 4, which shows the average of the actual lognormal implied volatility “smiles” for options on the different futures contracts as well as the averages of the fitted “smiles” for the SV1 and SV2 specifications. While SV1 has a fairly good fit to the average “smiles” for options on longer-term futures contracts, it is not able to match the average “smiles” for options on the shorter-term futures contracts. In contrast, SV2 has a very good fit to all

Table 5
Comparison of the SV1 and SV2 specifications for pricing options

Moneyness	Contract							
	M1	M2	M3	M4	M5	M6	Q1	Q2
0.78–0.82	–2.67*** (–3.70)	–0.61*** (–3.38)	–0.15** (–1.96)	–0.15* (–1.77)	–0.30*** (–2.86)	–0.41*** (–3.10)	–0.71*** (–4.80)	–0.79*** (–4.46)
0.82–0.86	–1.83*** (–3.93)	–0.68*** (–4.63)	–0.12** (–1.97)	–0.15** (–2.03)	–0.34*** (–3.81)	–0.60*** (–5.40)	–0.81*** (–6.31)	–0.97*** (–5.70)
0.86–0.90	–1.59*** (–5.00)	–0.70*** (–5.75)	–0.08* (–1.74)	–0.16** (–2.16)	–0.35*** (–4.64)	–0.66*** (–6.55)	–0.95*** (–7.83)	–1.18*** (–7.02)
0.90–0.94	–1.58*** (–7.24)	–0.72*** (–8.22)	–0.09** (–2.16)	–0.18** (–2.56)	–0.48*** (–6.39)	–0.79*** (–8.02)	–1.07*** (–9.16)	–1.15*** (–7.15)
0.94–0.98	–1.55*** (–9.11)	–0.77*** (–9.62)	–0.10** (–2.27)	–0.19** (–2.49)	–0.59*** (–7.43)	–0.88*** (–9.30)	–1.18*** (–9.40)	–1.16*** (–7.62)
0.98–1.02	–1.67*** (–10.03)	–0.75*** (–8.67)	–0.10** (–2.19)	–0.18** (–2.20)	–0.56*** (–7.42)	–0.83*** (–8.01)	–1.20*** (–8.84)	–1.20*** (–7.54)
1.02–1.06	–1.65*** (–8.38)	–0.70*** (–7.62)	–0.09** (–2.16)	–0.10 (–1.37)	–0.45*** (–6.34)	–0.77*** (–7.43)	–1.11*** (–7.80)	–1.18*** (–6.29)
1.06–1.10	–1.97*** (–7.78)	–0.72*** (–7.67)	–0.10*** (–2.80)	–0.08 (–1.04)	–0.40*** (–5.13)	–0.63*** (–5.49)	–1.03*** (–6.99)	–1.08*** (–5.07)
1.10–1.14	–2.28*** (–7.14)	–0.86*** (–8.25)	–0.12*** (–3.15)	–0.08 (–0.94)	–0.33*** (–3.46)	–0.49*** (–3.62)	–0.96*** (–5.99)	–1.04*** (–4.30)
1.14–1.18	–2.54*** (–5.51)	–0.99*** (–7.40)	–0.16*** (–3.11)	–0.08 (–0.91)	–0.25** (–2.54)	–0.40*** (–2.62)	–0.85*** (–4.80)	–1.16*** (–4.68)
1.18–1.22	–2.81*** (–4.06)	–0.99*** (–5.62)	–0.21*** (–2.72)	–0.10 (–0.72)	–0.22** (–2.00)	–0.37** (–2.36)	–0.78*** (–4.27)	–0.92*** (–3.74)

The table compares the SV1 and SV2 specifications in terms of their ability to price options within each moneyness–maturity category. It reports the mean differences in absolute pricing errors between the two specifications, when they are estimated on the entire data set. The pricing errors are defined as the differences between fitted and actual lognormal implied volatilities and reported in percentages. *T*-statistics, corrected for serial correlation up to 50 lags, are in parentheses. Moneyness is defined as option strike divided by the price of the underlying futures contract. Each statistic is computed on the basis of a maximum of 4082 daily observations from 2 January 1990 to 18 May 2006. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

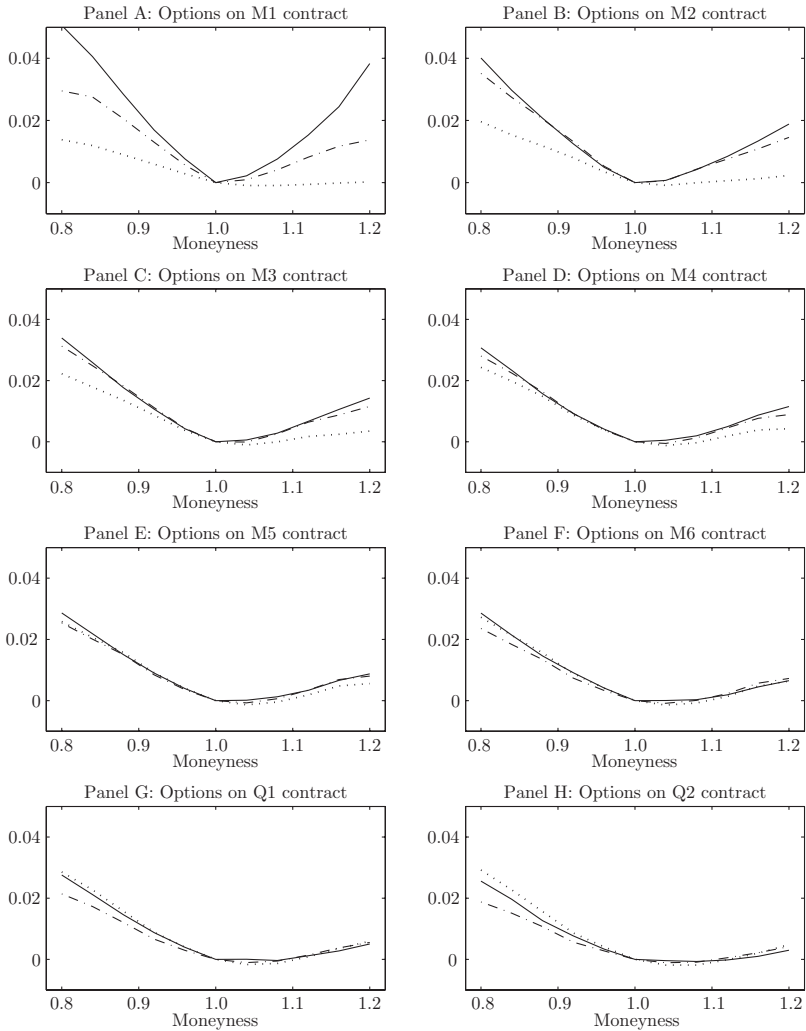


Figure 4
Average fit to lognormal implied volatility "smiles"

The — shows the average of the actual lognormal implied volatility "smiles," shows the average of the fitted "smiles" for the SV1 specification, and - - - shows the average of the fitted "smiles" for the SV2 specification. Both specifications are estimated on the entire data set. The "smiles" are the differences between the lognormal implied volatilities across moneyness and the lognormal implied volatilities of the corresponding ATM options. Moneyness is defined as option strike divided by the price of the underlying futures contract. Averages are taken over a maximum of 4082 daily observations from 2 January 1990 to 18 May 2006.

average implied volatility "smiles," except for the very short one, where the average "smile" does not exhibit sufficient curvature. This is a well-known deficiency of stochastic volatility models; as the option maturity approaches zero, the implied volatility "smile" flattens, whereas it tends to become more

pronounced in the data. We could fine-tune the model to match the very short-term “smile” by augmenting it with jumps in the spot price process along the lines of Bates (1996, 2000), Bakshi, Cao, and Chen (1997), and Duffie, Pan, and Singleton (2000), among others, for equity and FX options.³⁷ However, as Bakshi, Cao, and Chen (1997) conclude, “once stochastic volatility is modeled, adding other features [such as jumps and stochastic interest rates] will usually lead to second-order pricing improvements.”³⁸

Note that the average implied volatility “smiles” are skewed slightly toward OTM puts (that is, OTM puts are, on average, slightly more expensive than corresponding OTM calls), indicating that the skewness of the conditional risk-neutral futures return distribution is negative, on average, which is consistent with the negative signs on the estimates of ρ_{15} and ρ_{35} in the SV1 and SV2 specifications. Interestingly, for many, if not most, commodities the implied volatility “smiles” are, on average, skewed toward OTM calls. As a robustness check, we have reestimated the two specifications subject to the restriction $\rho_{15} = \rho_{35} = 0$. In this case, volatility is completely unspanned, and the conditional risk-neutral futures return distribution is symmetric. For both specifications, we observe a significant deterioration in the fit to options, particularly in the second subsample.³⁹

4.3 Hedging performance

Finally, we assess the importance of unspanned stochastic volatility when hedging options. On each date, we construct an option portfolio from all the options that meet the selection criteria discussed in Section 3.1. This portfolio consists of a large number of OTM put and call options with a wide range of moneyness and maturity. We use two weighting schemes: one where the options are equally weighted and one where the options are weighted by their open interests. The first scheme has the virtue of simplicity, while the second, by weighting with the number of outstanding contracts, produces a portfolio that is more representative of the entire option market. Since open interest tends to decrease with option maturity, options with short maturities, on average, receive higher weight in the second scheme than in the first scheme.⁴⁰

Hedge ratios are computed from the SV2gen specification, since it has the best fit to the futures term structure. In this specification, futures prices are driven by three factors, while option prices are driven by five factors. We analyze three hedges. The first hedge uses three futures contracts (M1, M4, and Q2)

³⁷ Hilliard and Reis (1998) develop a model for pricing commodity derivatives where the spot price follows a jump-diffusion process. However, their model does not account for stochastic volatility, and they make no attempt to fit their model to actual option data.

³⁸ See Bakshi, Cao, and Chen (1997), p. 2027.

³⁹ For instance, for the SV2 specification estimated on the second subsample, the mean in-sample option RMSE increases from 1.204% to 1.311% when imposing the restriction $\rho_{15} = \rho_{35} = 0$.

⁴⁰ We have experimented with other schemes, such as weighting the options with the inverse of their Black (1976) vegas, and obtain results that are very similar to those reported in this section.

Table 6
Hedging in the presence of unspanned stochastic volatility

	1990–2006 <i>IS</i>	1990–1997 <i>IS</i>	1998–2006 <i>IS</i>	1990–1997 <i>OOS</i>	1998–2006 <i>OOS</i>
Panel A: Equally-weighted portfolio					
P1	3.11	2.54	3.66	2.54	3.66
P2	2.82	2.29	3.39	2.33	3.42
P3	2.28	1.85	2.94	2.01	3.29
P4	1.65	1.24	1.95	1.43	2.14
P2–P1	–0.28*** (–13.04)	–0.25*** (–7.85)	–0.27*** (–10.34)	–0.21*** (–4.88)	–0.24*** (–10.32)
P3–P2	–0.54*** (–6.95)	–0.44*** (–3.51)	–0.44*** (–4.89)	–0.33** (–2.43)	–0.13* (–1.80)
P4–P3	–0.63*** (–6.70)	–0.61*** (–3.30)	–0.99*** (–9.06)	–0.57*** (–3.04)	–1.15*** (–9.29)
Panel B: Open interest-weighted portfolio					
P1	3.82	3.42	4.21	3.42	4.21
P2	3.32	2.94	3.70	2.98	3.82
P3	2.53	2.20	3.03	2.30	3.57
P4	2.15	1.87	2.37	2.02	2.60
P2–P1	–0.50*** (–15.93)	–0.48*** (–11.15)	–0.51*** (–9.36)	–0.45*** (–9.65)	–0.40*** (–12.68)
P3–P2	–0.79*** (–8.18)	–0.74*** (–4.25)	–0.67*** (–7.01)	–0.68*** (–4.61)	–0.24** (–2.57)
P4–P3	–0.37*** (–6.66)	–0.33*** (–3.96)	–0.66*** (–7.30)	–0.28*** (–2.91)	–0.98*** (–8.95)

The table reports the mean absolute daily profit–loss for an option portfolio that is either unhedged or hedged to a varying degree. P1 denotes the unhedged portfolio, P2 denotes the portfolio hedged with three futures contracts, P3 denotes the portfolio hedged with three futures contracts and one option, and P4 denotes the portfolio hedged with three futures contracts and two options. Hedge ratios are computed using the SV2gen specification. The individual options in the portfolio are weighted equally in Panel A and weighted by their open interests in Panel B. The three periods are 2 January 1990 to 17 May 2006 (4081 daily observations), 2 January 1990 to 31 December 1997 (2008 daily observations), and 2 January 1998 to 17 May 2006 (2073 daily observations). *IS* refers to in-sample results, while *OOS* refers to out-of-sample results, i.e., results for the first (second) subsample when the model is estimated on the second (first) subsample. The profits–losses are reported in cents. *T*-statistics, corrected for serial correlation up to 50 lags, are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

as hedge instruments, the second hedge adds one option (the closest-to-ATM option on the M1 contract) to the set of hedge instruments, while the third hedge adds one more option (the closest-to-ATM option on the Q2 contract) to the set of hedge instruments.⁴¹

We compute time series of the daily profit–loss on the unhedged option portfolio as well as the different hedged option portfolios. The profit–loss on a hedged option portfolio is given by the change in the value of the unhedged option portfolio minus the change in the value of the hedge position, where the latter is computed from the hedge ratios and the change in the value of the hedge instruments.

Table 6 shows, for both weighting schemes, the mean absolute daily profit–loss on the unhedged and the three hedged portfolios for the same three in-sample and two out-of-sample periods as in Table 3. The table also reports

⁴¹ In our data set, options on the Q2 contract are available only from 17 September 1992. Prior to this date, we use the closest-to-ATM option on the M6 contract instead.

Table 7
Assessing the conversion of American to European option prices

Moneyiness	Maturity	$v_1(t)$	SV^E	SV^A	$SD(SV^A)$	LN^E	$LN^E - SV^E$	$\frac{LN^E - SV^E}{SV^E}$
0.80	93	1	0.175	0.178	0.000	0.177	0.002	0.010
0.90	93	1	0.675	0.679	0.001	0.676	0.000	0.001
1.00	93	1	2.235	2.244	0.001	2.235	-0.001	-0.000
0.80	261	1	0.585	0.591	0.001	0.580	-0.005	-0.008
0.90	261	1	1.338	1.351	0.001	1.330	-0.008	-0.006
1.00	261	1	3.015	3.040	0.002	2.992	-0.022	-0.007
0.80	93	5	1.273	1.280	0.001	1.274	0.000	0.000
0.90	93	5	2.711	2.723	0.002	2.712	0.001	0.000
1.00	93	5	4.954	4.974	0.003	4.953	-0.001	-0.000
0.80	261	5	2.143	2.164	0.002	2.130	-0.014	-0.006
0.90	261	5	3.701	3.736	0.003	3.679	-0.022	-0.006
1.00	261	5	5.920	5.991	0.004	5.897	-0.023	-0.004

In this table, we assume that the SV1 specification, estimated on the entire sample, is the true stochastic volatility model generating market prices. SV^E denotes the price of a European put option with a given moneyiness and maturity. SV^A denotes the price of the corresponding American put option computed by simulation, using the least squares Monte Carlo (LSM) approach of Longstaff and Schwartz (2001). This price is based on 1,000,000 (500,000 plus 500,000 antithetic) paths with the European option price used as control variate. $SD(SV^A)$ denotes the standard deviation of the Monte Carlo estimate. LN^E denotes the option price obtained by converting the American option price back into its European counterpart under the (counterfactual) assumption that the underlying futures contract follows a geometric Brownian motion. The state variables are set to $S(t) = 50$, $x(t) = 0$, $\phi(t) = 0$, and either $v_1(t) = 1$ or $v_1(t) = 5$. Moneyiness is given by strike divided by the price of the underlying futures contract. Option maturity is in business days. The interest rate is set to 5%.

the mean differences in absolute daily profits–losses, along with the associated t -statistics. Hedging the option portfolio solely with futures contracts causes only small reductions in the mean absolute profits–losses, since only a small fraction of volatility is spanned by the futures contracts (and since the option portfolio is composed of both puts and calls and, therefore, its “delta” is already fairly close to zero). However, adding options to the set of hedge instruments leads to large and significant decreases in the mean absolute profits–losses, consistent with volatility being predominantly unspanned by the futures contracts. These results hold true both in-sample and out-of-sample and regardless of the weighting scheme.⁴² For instance, for the full sample and with equal weighting of the individual options, the mean absolute daily profit–loss is reduced from 3.11 cents to 2.82 cents, when hedging solely with futures contracts, and further reduced to 1.65 cents, when adding two options to the set of hedge instruments.⁴³

⁴² The results are also fairly insensitive to the exact choice of maturities for the futures and options used as hedge instruments. Furthermore, using the SV2 specification to compute hedge ratios and using two futures contracts as hedge instruments instead of three (since futures prices are driven by two factors in this specification) produce almost identical results.

⁴³ Naturally, if market prices were indeed generated by the SV2gen specification, the daily profit–loss of the hedged option portfolio would be close to zero when using three futures contracts and two options as hedge instruments (some hedging error would remain since we are hedging discretely rather than continuously). However, since the SV2gen specification is only an approximation of reality and since there is also idiosyncratic noise in option prices and, to a lesser extent, futures prices, a sizable hedging error remains even after including two options as hedge instruments.

5. Conclusion

We have developed a general model for pricing commodity derivatives in the presence of unspanned stochastic volatility. Futures prices are driven by three factors (one factor being the spot price of the commodity and two factors affecting the forward cost of carry curve) and option prices are driven by two additional volatility factors. Both volatility factors may contain a spanned and an unspanned component and both factors may affect the instantaneous volatility of the spot price and the forward cost of carry. The model is highly tractable with quasi-analytical prices of European options on futures contracts. Furthermore, the dynamics of the futures curve can be described in terms of a low-dimensional affine state vector making the model suitable for estimation and for pricing complex commodity derivatives, including real options, by simulation, where early exercise features can be handled by the least squares approach of Longstaff and Schwartz (2001).⁴⁴

We estimate the model as well as several nested, more parsimonious, specifications on NYMEX crude oil derivatives using an extensive panel data set of 45,517 futures prices and 233,104 option prices, spanning 4082 business days. We find that two volatility factors are necessary to fit options on futures contracts across the maturity and moneyness dimensions. Both volatility factors are predominantly unspanned by the futures contracts, and the first volatility factor drives virtually all of the instantaneous volatility of the spot price and the front end of the forward cost of carry curve (with the second volatility factor being more important for the instantaneous volatility of longer-term forward cost of carry rates). Therefore, a more parsimonious two-factor specification, where the second volatility factor is completely unspanned by the futures contracts and does not affect the instantaneous volatility of the spot price or the forward cost of carry, performs almost as well as the general specification in terms of pricing short-term and medium-term options.

The model can be extended along several dimensions. For instance, we might include separate processes for the forward interest rate and the forward convenience yield. Furthermore, while the model generally fits the average implied volatility “smiles,” it does not quite capture the time variation in the skew of the implied volatility “smiles,” which is indicative of time-varying skewness of the conditional risk-neutral futures return distribution. It will be interesting to investigate extensions of the model that are able to better capture this feature of the data.⁴⁵

Finally, in the empirical parts of the paper, we have used data from the crude oil derivatives market since this is the most important and most liquid commodity derivatives market. It is likely that markets related to crude oil, such as gasoline and heating oil and possibly natural gas, also exhibit unspanned

⁴⁴ See Schwartz and Trolle (forthcoming) for a real option application of the model.

⁴⁵ Carr and Wu (2007) and Bakshi, Carr, and Wu (2008) document time-varying skewness of the conditional risk-neutral distribution of currency returns and suggest several models that are able to match this feature.

stochastic volatility. An interesting question is whether unspanned stochastic volatility is an important feature of commodity markets less related to crude oil, such as the markets for base metals and agricultural products. We leave these issues for future research.

Appendix A. Proofs

Proof of Proposition 1. We introduce the process

$$Y(t, T) = \int_t^T y(t, u)du \tag{A1}$$

with dynamics given by

$$dY(t, T) = \left(-\delta(t) + \int_t^T \mu_y(t, u)du \right) dt + \sqrt{v_1(t)} \int_t^T \sigma_{y1}(t, u)du dW_3^Q(t) + \sqrt{v_2(t)} \int_t^T \sigma_{y2}(t, u)du dW_4^Q(t). \tag{A2}$$

$F(t, T)$ is given by

$$F(t, T) = S(t)e^{Y(t, T)}, \tag{A3}$$

which follows

$$\begin{aligned} \frac{dF(t, T)}{F(t, T)} &= \frac{dS(t)}{S(t)} + dY(t, T) + \frac{1}{2}(dY(t, T))^2 + \frac{dS(t)}{S(t)}dY(t, T) \\ &= \left(\int_t^T \mu_y(t, u)du + \left(\frac{1}{2} \left(\int_t^T \sigma_{y1}(t, u)du \right)^2 + \rho_{13}\sigma_{S1} \int_t^T \sigma_{y1}(t, u)du \right) v_1(t) \right. \\ &\quad \left. + \left(\frac{1}{2} \left(\int_t^T \sigma_{y2}(t, u)du \right)^2 + \rho_{24}\sigma_{S2} \int_t^T \sigma_{y2}(t, u)du \right) v_2(t) \right) dt \\ &\quad + \sqrt{v_1(t)} \left(\sigma_{S1}dW_1^Q(t) + \int_t^T \sigma_{y1}(t, u)du dW_3^Q(t) \right) \\ &\quad + \sqrt{v_2(t)} \left(\sigma_{S2}dW_2^Q(t) + \int_t^T \sigma_{y2}(t, u)du dW_4^Q(t) \right). \end{aligned} \tag{A4}$$

In the absence of arbitrage, the drift must equal zero. Imposing this condition and differentiating with respect to T yield Equation (7).

Proof of Proposition 2. With $\sigma_{yi}(t, T)$ given as Equation (8), $\mu_y(t, T)$ is given by Equation (7) as

$$\begin{aligned} \mu_y(t, T) &= v_1(t) \left(\frac{\alpha_1^2}{\gamma_1} e^{-2\gamma_1(T-t)} - \left(\frac{\alpha_1}{\gamma_1} + \rho_{13}\sigma_{S1} \right) \alpha_1 e^{-\gamma_1(T-t)} \right) \\ &\quad + v_2(t) \left(\frac{\alpha_2^2}{\gamma_2} e^{-2\gamma_2(T-t)} - \left(\frac{\alpha_2}{\gamma_2} + \rho_{24}\sigma_{S2} \right) \alpha_2 e^{-\gamma_2(T-t)} \right). \end{aligned} \tag{A5}$$

Integrating Equation (2) and using that $e^{-\gamma_i(T-u)} = e^{-\gamma_i(T-t)}e^{-\gamma_i(t-u)}$, we obtain

$$\begin{aligned} y(t, T) &= y(0, T) + \int_0^t \mu_y(u, T)du + \int_0^t \sigma_{y1}(u, T)\sqrt{v_1(u)}dW_3^Q(u) + \int_0^t \sigma_{y2}(u, T)\sqrt{v_2(u)}dW_4^Q(u) \\ &= y(0, T) + \sum_{i=1}^2 \left(\alpha_i e^{-\gamma_i(T-t)} x_i(t) + \alpha_i e^{-2\gamma_i(T-t)} \phi_i(t) \right), \end{aligned} \quad (A6)$$

where

$$x_1(t) = - \int_0^t v_1(u) \left(\frac{\alpha_1}{\gamma_1} + \rho_{13}\sigma_{S1} \right) e^{-\gamma_1(t-u)} du + \int_0^t e^{-\gamma_1(t-u)} \sqrt{v_1(u)} dW_3^Q(u), \quad (A7)$$

$$x_2(t) = - \int_0^t v_2(u) \left(\frac{\alpha_2}{\gamma_2} + \rho_{24}\sigma_{S2} \right) e^{-\gamma_2(t-u)} du + \int_0^t e^{-\gamma_2(t-u)} \sqrt{v_2(u)} dW_4^Q(u), \quad (A8)$$

$$\phi_i(t) = \int_0^t v_i(u) \frac{\alpha_i}{\gamma_i} e^{-2\gamma_i(t-u)} du, \quad i = 1, 2. \quad (A9)$$

Applying Ito's lemma to these expressions gives the dynamics stated in Equations (10)–(12).

Proof of Proposition 3. The proof is similar to those in Duffie, Pan, and Singleton (2000) and Collin-Dufresne and Goldstein (2003). We can rewrite Equation (19) as

$$\psi(u, t, T_0, T_1) = E_t^Q \left[E_{T_0}^Q \left[e^{u \log(F(T_0, T_1))} \right] \right] = E_t^Q [\psi(u, T_0, T_0, T_1)]. \quad (A10)$$

Therefore, the proof consists of showing that the process $\xi(t) \equiv \psi(u, t, T_0, T_1)$ is a martingale under Q . To this end, we conjecture that $\psi(u, t, T_0, T_1)$ is of the form of Equation (20). Applying Ito's lemma to $\xi(t)$ and setting $\tau = T_0 - t$, we obtain

$$\begin{aligned} \frac{d\xi(t)}{\xi(t)} &= \left(- \frac{dM(\tau)}{d\tau} - \frac{dN_1(\tau)}{d\tau} v_1(t) - \frac{dN_2(\tau)}{d\tau} v_2(t) \right) dt + N_1(\tau) dv_1(t) + N_2(\tau) dv_2(t) \\ &\quad + u \frac{dF(t, T_1)}{F(t, T_1)} + \frac{1}{2} N_1(\tau)^2 (dv_1(t))^2 + \frac{1}{2} N_2(\tau)^2 (dv_2(t))^2 + \frac{1}{2} (u^2 - u) \left(\frac{dF(t, T_1)}{F(t, T_1)} \right)^2 \\ &\quad + N_1(\tau) u dv_1(t) \frac{dF(t, T_1)}{F(t, T_1)} + N_2(\tau) u dv_2(t) \frac{dF(t, T_1)}{F(t, T_1)} + N_1(\tau) N_2(\tau) dv_1(t) dv_2(t). \end{aligned} \quad (A11)$$

For $\xi(t)$ to be a martingale, it must hold that

$$\begin{aligned} 0 &= \frac{1}{dt} E_t^Q \left[\frac{d\xi(t)}{\xi(t)} \right] \\ &= - \frac{dM(\tau)}{d\tau} - \frac{dN_1(\tau)}{d\tau} v_1(t) - \frac{dN_2(\tau)}{d\tau} v_2(t) + N_1(\tau)(\eta_1 - \kappa_1 v_1(t) - \kappa_{12} v_2(t)) \\ &\quad + N_2(\tau)(\eta_2 - \kappa_{21} v_1(t) - \kappa_2 v_2(t)) + \frac{1}{2} N_1(\tau)^2 \sigma_{v1}^2 v_1(t) + \frac{1}{2} N_2(\tau)^2 \sigma_{v2}^2 v_2(t) \\ &\quad + \frac{1}{2} (u^2 - u) \left(\left(\sigma_{S1}^2 + B_{x1}(T_1 - t)^2 + 2\rho_{13}\sigma_{S1} B_{x1}(T_1 - t) \right) v_1(t) \right. \\ &\quad \left. + \left(\sigma_{S2}^2 + B_{x2}(T_1 - t)^2 + 2\rho_{24}\sigma_{S2} B_{x2}(T_1 - t) \right) v_2(t) \right) \\ &\quad + N_1(\tau) u \sigma_{v1} (\rho_{15}\sigma_{S1} + \rho_{35} B_{x1}(T_1 - t)) v_1(t) + N_2(\tau) u \sigma_{v2} (\rho_{26}\sigma_{S2} + \rho_{46} B_{x2}(T_1 - t)) v_2(t) \\ &= - \frac{dM(\tau)}{d\tau} + N_1(\tau)\eta_1 + N_2(\tau)\eta_2 \\ &\quad + \left[- \frac{dN_1(\tau)}{d\tau} - N_2(\tau)\kappa_{21} + N_1(\tau)(-\kappa_1 + u\sigma_{v1}(\rho_{15}\sigma_{S1} + \rho_{35} B_{x1}(T_1 - t))) \right] + \frac{1}{2} N_1(\tau)^2 \sigma_{v1}^2 \end{aligned}$$

$$\begin{aligned}
 & + \frac{1}{2}(u^2 - u)\left(\sigma_{S1}^2 + B_{x1}(T_1 - t)^2 + 2\rho_{13}\sigma_{S1}B_{x1}(T_1 - t)\right)\Big]v(t) \\
 & + \left[-\frac{dN_2(\tau)}{d\tau} - N_1(\tau)\kappa_{12} + N_2(\tau)(-\kappa_2 + u\sigma_{v2}(\rho_{26}\sigma_{S2} + \rho_{46}B_{x2}(T_1 - t))) + \frac{1}{2}N_2(\tau)^2\sigma_{v2}^2\right. \\
 & \left. + \frac{1}{2}(u^2 - u)\left(\sigma_{S2}^2 + B_{x2}(T_1 - t)^2 + 2\rho_{24}\sigma_{S2}B_{x2}(T_1 - t)\right)\right]v_2(t). \tag{A12}
 \end{aligned}$$

Hence, $\xi(t)$ is a martingale, provided that $M(\tau)$, $N_1(\tau)$, and $N_2(\tau)$ satisfy Equations (21), (22), and (23), respectively. Furthermore, we have

$$\psi(u, T_0, T_0, T_1) = e^{u\log(F(T_0, T_1))}, \tag{A13}$$

which is true, provided that $M(0) = 0$, $N_1(0) = 0$, and $N_2(0) = 0$.

Proof of Proposition 4. Again, we follow Duffie, Pan, and Singleton (2000) and Collin-Dufresne and Goldstein (2003). Assuming $r(t)$ and $F(t, T_1)$ are uncorrelated under \mathcal{Q} , the time- t price of a European put option expiring at time T_0 with strike K on a futures contract expiring at time T_1 , $\mathcal{P}(t, T_0, T_1, K)$, is given by

$$\begin{aligned}
 \mathcal{P}(t, T_0, T_1, K) & = E_t^{\mathcal{Q}}\left[e^{-\int_t^{T_0} r(s)ds}(K - F(T_0, T_1))\mathbf{1}_{F(T_0, T_1) < K}\right] \\
 & = P(t, T_0)\left(K E_t^{\mathcal{Q}}[\mathbf{1}_{\log(F(T_0, T_1)) < \log(K)}] - E_t^{\mathcal{Q}}[e^{\log(F(T_0, T_1))}\mathbf{1}_{\log(F(T_0, T_1)) < \log(K)}]\right) \\
 & = P(t, T_0)(K G_{0,1}(\log(K)) - G_{1,1}(\log(K))), \tag{A14}
 \end{aligned}$$

where

$$G_{a,b}(y) = E_t^{\mathcal{Q}}\left[e^{a\log(F(T_0, T_1))}\mathbf{1}_{b\log(F(T_0, T_1)) < y}\right]. \tag{A15}$$

To evaluate $G_{a,b}(y)$, note that its Fourier transform is given by

$$\begin{aligned}
 \mathcal{G}_{a,b}(y) & = \int_{\mathbb{R}} e^{iy} dG_{a,b}(y) \\
 & = E_t^{\mathcal{Q}}\left[e^{(a+ib)\log(F(T_0, T_1))}\right] \\
 & = \psi(a + ib, t, T_0, T_1), \tag{A16}
 \end{aligned}$$

where $i = \sqrt{-1}$. Applying the Fourier inversion theorem, we have

$$G_{a,b}(y) = \frac{\psi(a, t, T_0, T_1)}{2} - \frac{1}{\pi} \int_0^{\infty} \frac{\text{Im}[\psi(a + ib, t, T_0, T_1)e^{-iuy}]}{u} du. \tag{A17}$$

Appendix B. Conversion of American to European Option Prices

Here, we outline the procedure used for converting the American option prices in the data set to their European counterparts. We also discuss the accuracy of the approach. The procedure is very similar to that used by Broadie, Chernov, and Johannes (2007) for estimating equity derivatives models using S&P 500 futures options. For each option, we assume that the price of the underlying futures contract follows a geometric Brownian motion, in which case a very accurate price for the American option can be obtained with the Barone-Adesi and Whaley (1987) formula. Inverting this formula for a given American option price yields a lognormal implied volatility, from which we can

price the associated European option with the Black (1976) formula.⁴⁶ Naturally, this procedure is inherently inconsistent; the whole point of the paper is to investigate stochastic volatility, yet when approximating the early exercise premium, we assume that the underlying futures contract follows a geometric Brownian motion with constant volatility. However, the procedure does implicitly take variation in volatility into account, since we view each option in isolation and, therefore, the implied volatility obtained by inverting the Barone-Adesi and Whaley (1987) formula in reality varies across time, option maturity, and moneyness.

Nevertheless, it is important to investigate the size of the bias induced by the conversion procedure. To do so, we assume that the true stochastic volatility model generating market prices is the SV1 specification with parameters equal to the full-sample estimates reported in Tables 1 and 2. For a given set of state variables, prices of European options are computed quasi-analytically, and prices of the corresponding American put options are computed by simulation, using the least squares Monte Carlo (LSM) approach developed by Longstaff and Schwartz (2001).⁴⁷ American option prices are then converted back into European option prices using the conversion procedure outlined above. We consider put options with moneyness 0.80, 0.90, and 1.00, which spans the moneyness range in the sample, and maturities of 93 and 261 business days, which are, respectively, the mean and the maximum option maturity in the sample. We assume a relatively flat futures term structure with $S(t) = 50$, $x(t) = 0$, and $\phi(t) = 0$, and consider both a high-volatility day, where $v_1(t) = 5$, and a low-volatility day, where $v_1(t) = 1$.⁴⁸ The interest rate is set to 5%. Table 7 compares the true European prices with the approximate European prices obtained from the conversion procedure. For the options with a maturity of 93 business days and, therefore, for most of the options in the sample, the approximation errors are virtually zero. Even for the longest option maturity in the sample, the approximation errors are very small both in absolute and relative terms.⁴⁹ Therefore, in our view, the conversion procedure has negligible impact on the estimation results.⁵⁰

Appendix C. Estimation Details

The extended Kalman filter

Let $\hat{X}_t = E_t[X_t]$ and $\hat{X}_{t|t-1} = E_{t-1}[X_t]$ denote expectations of X_t (respectively, including and excluding z_t), and let P_t and $P_{t|t-1}$ denote the corresponding estimation error covariance matrices. Linearizing the h -function in Equation (34) around $\hat{X}_{t|t-1}$, we have

$$z_t = \left(h(\hat{X}_{t|t-1}) - H'_t \hat{X}_{t|t-1} \right) + H'_t X_t + u_t, \quad u_t \sim \text{iid } N(0, \Omega), \quad (\text{C1})$$

⁴⁶ We have also inverted the American option prices using binomial and trinomial trees. This gives results that are very similar to using the Barone-Adesi and Whaley (1987) formula.

⁴⁷ The LSM estimates are based on 1,000,000 (500,000 plus 500,000 antithetic) paths with the European option prices used as control variates. We approximate the continuation value of the option with a linear combination of the following basis functions: a constant, $F(t, T_1)$, $F(t, T_1)^2$, $F(t, T_1)^3$, $F(t, T_1)^4$, $v(t)$, $v(t)^2$, $v(t)^3$, $v(t)^4$, $v(t)F(t, T_1)$, $v(t)^2 F(t, T_1)$, and $v(t)F(t, T_1)^2$. We have verified that alternative specifications, involving an equal or larger number of terms, yield results that are virtually indistinguishable from those obtained with this specification.

⁴⁸ In our data set, options on the Q2 contract—the longest options that we consider—are available only from 17 September 1992. Over the period that these options are available, the minimum, the mean, and the maximum of $v_1(t)$ are 0.44, 2.60, and 8.79, respectively.

⁴⁹ In absolute terms, the maximum approximation error occurs for the ATM option on a high-volatility day (2.3 cents), while in relative terms, the maximum approximation error occurs for the most OTM option on a low-volatility day (0.80%).

⁵⁰ We have experimented with other combinations of the state variables—including situations where the futures curve is strongly in contango or backwardation—as well as with the SV2 and SV2gen specifications and find results that are very similar to those reported here.

where

$$H'_t = \left. \frac{\delta h(X_t)}{\delta X'_t} \right|_{X_t = \hat{X}_{t|t-1}}. \tag{C2}$$

Assuming w_t in Equation (33) is Gaussian, we have

$$X_t = \Phi_0 + \Phi_X X_{t-1} + w_t, \quad w_t \sim \text{iid } N(0, Q_t). \tag{C3}$$

The Kalman filter applied to Equations (C1) and (C3) yields

$$\hat{X}_{t|t-1} = \Phi_0 + \Phi_X \hat{X}_{t-1}, \tag{C4}$$

$$P_{t|t-1} = \Phi_X P_{t-1} \Phi'_X + Q_t, \tag{C5}$$

and

$$\hat{X}_t = \hat{X}_{t|t-1} + P_{t|t-1} H'_t F_t^{-1} \epsilon_t, \tag{C6}$$

$$P_t = P_{t|t-1} - P_{t|t-1} H'_t F_t^{-1} H_t P_{t|t-1}, \tag{C7}$$

where

$$\epsilon_t = z_t - h(\hat{X}_{t|t-1}), \tag{C8}$$

$$F_t = H_t P_{t|t-1} H'_t + \Omega. \tag{C9}$$

The loglikelihood function is constructed from Equations (C8) and (C9) as

$$\log L = -\frac{1}{2} \log 2\pi \sum_{i=1}^T N_i - \frac{1}{2} \sum_{i=1}^T \log |F_i| - \frac{1}{2} \sum_{i=1}^T \epsilon'_i F_i^{-1} \epsilon_i, \tag{C10}$$

where T is the number of observation dates and N_i is the dimension of ϵ_i .

Approximating the true distribution of w_t in Equation (33) with a Gaussian makes this a QML procedure. While QML estimation has been shown to be consistent in many settings, it is in fact not consistent in the present context due to the linearization of the h -function, and the fact that the conditional covariance matrix Q in the recursions depends on the Kalman filter estimate(s) of the volatility state variable(s) rather than the true, but unobservable, value(s); see, e.g., Duan and Simonato (1999). In Trolle and Schwartz (2009), we investigate the small-sample properties of the QML/Kalman filter approach in a related context of estimating multi-factor stochastic volatility term structure models using interest rate derivatives. We find virtually no biases in the estimates of the parameters identified under Q and only small and insignificant biases in the estimates of the drift parameters in the P -dynamics. We expect that similar results will be obtained in the present context.

Numerical issues

To maximize the loglikelihood function, we initially apply the Nelder–Mead algorithm and later switch to the gradient-based BFGS algorithm. The optimization is repeated with several different plausible initial parameter guesses, to minimize the risk of not reaching the global optimum. The systems of ODEs in Equations (21)–(23) are solved with a standard fourth-order Runge–Kutta

algorithm, and the integral in Equation (25) is evaluated with the Gauss–Legendre quadrature formula, using 30 integration points and truncating the integral at 400.⁵¹

References

- Alizadeh, S., M. W. Brandt, and F. X. Diebold. 2002. Range-Based Estimation of Stochastic Volatility Models. *Journal of Finance* 57:1047–91.
- Amin, K., V. Ng, and S. C. Pirrong. 1995. Valuing Energy Derivatives. In *Managing Energy Price Risk*. London: Risk Publications.
- Andersen, T. G., T. Bollerslev, F. X. Diebold, and H. Ebens. 2001a. The Distribution of Realized Stock Return Volatility. *Journal of Financial Economics* 61:43–76.
- Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys. 2001b. The Distribution of Realized Exchange Rate Volatility. *Journal of the American Statistical Association* 96:42–55.
- Bakshi, G., C. Cao, and Z. Chen. 1997. Empirical Performance of Alternative Option Pricing Models. *Journal of Finance* 52:2003–49.
- Bakshi, G., P. Carr, and L. Wu. 2008. Stochastic Risk Premiums, Stochastic Skewness in Currency Options, and Stochastic Discount Factors in International Economies. *Journal of Financial Economics* 87:132–56.
- Bakshi, G., and D. Madan. 2000. Spanning and Derivative-Security Valuation. *Journal of Financial Economics* 55:205–38.
- Barone-Adesi, G., and R. Whaley. 1987. Efficient Analytic Approximation of American Option Values. *Journal of Finance* 42:301–20.
- Bates, D. S. 1996. Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options. *Review of Financial Studies* 9:69–107.
- Bates, D. S. 2000. Post-'87 Crash Fears in the S&P 500 Futures Option Market. *Journal of Econometrics* 94:181–238.
- Bhar, R., and C. Chiarella. 1997. Transformation of Heath–Jarrow–Morton Models to Markovian Systems. *European Journal of Finance* 3:1–26.
- BIS. 2007. *Quarterly Review (March): International Banking and Financial Market Developments*.
- Black, F. 1976. The Pricing of Commodity Contracts. *Journal of Financial Economics* 3:167–79.
- Bollerslev, T., and H. Zhou. 2002. Estimating Stochastic Volatility Diffusion Using Conditional Moments of Integrated Volatility. *Journal of Econometrics* 109:33–65.
- Brennan, M. J. 1991. The Price of Convenience and the Valuation of Commodity Contingent Claims. In Diderik Lund and Bernt Øksendal (eds.), *Stochastic Models and Option Values*. Amsterdam: North-Holland.
- Broadie, M., M. Chernov, and M. Johannes. 2007. Model Specification and Risk Premiums: Evidence from Futures Options. *Journal of Finance* 62:1453–90.
- Carlson, M., Z. Khoker, and S. Titman. 2007. Equilibrium Exhaustible Resource Price Dynamics. *Journal of Finance* 62:1663–703.
- Carr, P., and D. Madan. 1999. Option Valuation Using the Fast Fourier Transform. *Journal of Computational Finance* 2:61–73.
- Carr, P., and L. Wu. 2007. Stochastic Skew in Currency Options. *Journal of Financial Economics* 86:213–47.
- Casassus, J., and P. Collin-Dufresne. 2005. Stochastic Convenience Yield Implied from Commodity Futures and Interest Rates. *Journal of Finance* 60:2283–331.

⁵¹ More specifically, we use 15 points on the interval 0–50 and another 15 points on the interval 50–400. Increasing the number of points and/or the truncation of the integral does not change the likelihood value.

- Casassus, J., P. Collin-Dufresne, and B. R. Routledge. 2003. Equilibrium Commodity Prices with Irreversible Investment and Non-Linear Technologies. Working Paper, Carnegie Mellon University.
- Casassus, J., P. Collin-Dufresne, and R. Goldstein. 2005. Unspanned Stochastic Volatility and Fixed Income Derivatives Pricing. *Journal of Banking and Finance* 29:2723–49.
- Chambers, M. J., and R. E. Bailey. 1996. A Theory of Commodity Price Fluctuations. *Journal of Political Economy* 104:924–57.
- Chernov, M., R. Gallant, E. Ghysels, and G. Tauchen. 2003. Alternative Models for Stock Price Dynamics. *Journal of Econometrics* 116:225–57.
- Clewlow, L., and C. Strickland. 1999. A Multi-Factor Model for Energy Derivatives. Working Paper, University of Technology, Sydney.
- Collin-Dufresne, P., and R. Goldstein. 2002. Do Bonds Span the Fixed Income Markets? Theory and Evidence for Unspanned Stochastic Volatility. *Journal of Finance* 57:1685–730.
- Collin-Dufresne, P., and R. Goldstein. 2003. Generalizing the Affine Framework to HJM and Random Field Models. Working Paper, University of California, Berkeley.
- Cortazar, G., and E. S. Schwartz. 1994. The Valuation of Commodity Contingent Claims. *Journal of Derivatives* 1:27–39.
- Dai, Q., and K. Singleton. 2000. Specification Analysis of Affine Term Structure Models. *Journal of Finance* 55:1943–78.
- Deaton, A., and G. Laroque. 1992. On the Behaviour of Commodity Prices. *Review of Economic Studies* 59:1–23.
- Deaton, A., and G. Laroque. 1996. Competitive Storage and Commodity Price Dynamics. *Journal of Political Economy* 104:896–923.
- Diebold, F., and R. Mariano. 1995. Comparing Predictive Accuracy. *Journal of Business and Economic Statistics* 13:253–63.
- Doran, J. S., and E. I. Ronn. 2006. Computing the Market Price of Volatility Risk in the Energy Commodity Markets. Working Paper, Florida State University.
- Duan, J.-C., and J.-G. Simonato. 1999. Estimating and Testing Exponential-Effine Term Structure Models by Kalman Filter. *Review of Quantitative Finance and Accounting* 13:111–35.
- Duffie, D. 2001. *Dynamic Asset Pricing Theory*. Princeton, NJ: Princeton University Press.
- Duffie, D., J. Pan, and K. Singleton. 2000. Transform Analysis and Asset Pricing for Affine Jump-Diffusions. *Econometrica* 68:1343–76.
- Engle, R. F., and G. G. J. Lee. 1999. A Permanent and Transitory Model of Stock Return Volatility. In Robert F. Engle and Halbert White (eds.), *Cointegration, Causality, and Forecasting: A Festschrift in Honor of Clive W. J. Granger*. Oxford: Oxford University Press.
- Eydeland, A., and H. Geman. 1998. Pricing Power Derivatives. *Risk* 11:71–73.
- Gibson, R., and E. S. Schwartz. 1990. Stochastic Convenience Yield and the Pricing of Oil Contingent Claims. *Journal of Finance* 45:959–76.
- Hamilton, J. D. 1994. *Time Series Analysis*. Princeton, NJ: Princeton University Press.
- Harvey, A. C. 1989. *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge: Cambridge University Press.
- Heath, D., R. Jarrow, and A. Morton. 1992. Bond Pricing and the Term Structure of Interest Rates: A New Methodology for Contingent Claims Valuation. *Econometrica* 60:77–105.
- Heston, S. 1993. A Closed Form Solution for Options with Stochastic Volatility. *Review of Financial Studies* 6:327–43.

- Hilliard, J. E., and J. Reis. 1998. Valuation of Commodity Futures and Options under Stochastic Convenience Yields, Interest Rates and Jump Diffusions in the Spot. *Journal of Financial and Quantitative Analysis* 33:61–86.
- Kogan, L., D. Livdan, and A. Yaron. 2005. Futures Prices in a Production Economy with Investment Constraints. Working Paper, NBER.
- Litzenberger, R. H., and N. Rabinowitz. 1995. Backwardation in Oil Futures Markets: Theory and Empirical Evidence. *Journal of Finance* 50:1517–45.
- Longstaff, F., and E. Schwartz. 2001. Valuing American Options by Simulation: A Simple Least-Square Approach. *Review of Financial Studies* 14:113–47.
- Miltersen, K. 2003. Commodity Price Modelling That Matches Current Observables: A New Approach. *Quantitative Finance* 3:51–58.
- Miltersen, K., and E. S. Schwartz. 1998. Pricing of Options on Commodity Futures with Stochastic Term Structures of Convenience Yields and Interest Rates. *Journal of Financial and Quantitative Analysis* 33:33–59.
- Nelson, C., and A. Siegel. 1987. Parsimonious Modeling of Yield Curves. *Journal of Business* 60:473–89.
- Newey, W., and K. West. 1987. A Simple, Positive Semi-Definite, Heteroscedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica* 55:703–8.
- Ng, V. K., and C. Pirrong. 1994. Fundamentals and Volatility: Storage, Spreads and the Dynamics of Metal Prices. *Journal of Business* 67:203–30.
- Nielsen, M. J., and E. S. Schwartz. 2004. Theory of Storage and the Pricing of Commodity Claims. *Review of Derivatives Research* 7:5–24.
- Richter, M., and C. Sørensen. 2002. Stochastic Volatility and Seasonality in Commodity Futures and Options: The Case of Soybeans. Working Paper, Copenhagen Business School.
- Ritchken, P., and L. Sankarasubramaniam. 1995. Volatility Structures of Forward Rates and the Dynamics of the Term Structure. *Mathematical Finance* 5:55–72.
- Routledge, B. R., D. J. Seppi, and C. S. Spatt. 2000. Equilibrium Forward Curves for Commodities. *Journal of Finance* 55:1297–338.
- Schwartz, E. S. 1997. The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging. *Journal of Finance* 52:923–73.
- Schwartz, E. S., and A. B. Trolle. Forthcoming. Pricing Expropriation Risk in Natural Resource Contracts: A Real Options Approach. In William Hogan and Federico Sturzenegger (eds.), *Populism and Natural Resources*. Cambridge, MA: MIT Press.
- Schwartz, E. S., and J. E. Smith. 2000. Short-Term Variations and Long-Term Dynamics of Commodity Prices. *Management Science* 46:893–911.
- Trolle, A. B., and E. S. Schwartz. 2009. A General Stochastic Volatility Model for the Pricing of Interest Rate Derivatives. *Review of Financial Studies* 22:2007–57.