

Short-Term Variations and Long-Term Dynamics in Commodity Prices

Eduardo Schwartz • James E. Smith

Anderson Graduate School of Management, University of California, Los Angeles, Los Angeles, California 90095-1481

Fuqua School of Business, Duke University, Durham, North Carolina 27708-0120

eduardo.schwartz@anderson.ucla.edu • jes9@mail.duke.edu

In this article, we develop a two-factor model of commodity prices that allows mean-reversion in short-term prices and uncertainty in the equilibrium level to which prices revert. Although these two factors are not directly observable, they may be estimated from spot and futures prices. Intuitively, movements in prices for long-maturity futures contracts provide information about the equilibrium price level, and differences between the prices for the short- and long-term contracts provide information about short-term variations in prices. We show that, although this model does not explicitly consider changes in convenience yields over time, this short-term/long-term model is equivalent to the stochastic convenience yield model developed in Gibson and Schwartz (1990). We estimate the parameters of the model using prices for oil futures contracts and apply the model to some hypothetical oil-linked assets to demonstrate its use and some of its advantages over the Gibson-Schwartz model. (*Commodity Prices; Real Options; Stochastic Dynamic Model*)

1. Introduction

Stochastic models of commodity prices play a central role when evaluating commodity-related securities and projects. Early studies in this area typically assumed that commodity prices followed a “random walk” described by geometric Brownian motion.¹ This is the model of stock price uncertainty underlying the famous Black-Scholes option pricing formula and it leads to closed-form solutions in some interesting cases. In this model, prices are expected to grow at some constant rate with the variance in future spot prices increasing in proportion to time. If prices increase (or decrease) more than anticipated in one time period, all future forecasts are increased (or decreased) proportionally.

More recently, a number of authors have considered the use of mean-reverting price models and argued

that these models are more appropriate for many commodities.² Intuitively, when the price of a commodity is higher than some long-run mean or equilibrium price level, the supply of the commodity will increase because higher cost producers of the commodity will enter the market—new production comes on line, older production expected to go off line stays on line—thereby putting downward pressure on prices. Conversely, when prices are relatively low, supply will decrease since some of the high-cost producers will exit, putting upward pressure on prices. When these entries and exits are not instantaneous, prices may be temporarily high or low but will tend to revert toward the equilibrium level.

There are elements of truth in each of these simple models of commodity prices. For most commodities,

¹ See, for example, Brennan and Schwartz (1985), Paddock et al. (1988), and Smith and McCardle (1998).

² See, for example, Laughton and Jacoby (1993, 1995), Cortazar and Schwartz (1994), Dixit and Pindyck (1994), and Smith and McCardle (1999).

there appear to be some mean reversion in prices but there is also uncertainty about the equilibrium price to which prices revert. In this article, we develop a simple two-factor model of commodity prices that captures both of these effects; the goal is to provide a model that is more realistic than these standard models but yet simple enough to be useful for evaluating real and financial options. In our model, the equilibrium price level is assumed to evolve according to geometric Brownian motion with drift reflecting expectations of the exhaustion of existing supply, improving technology for the production and discovery of the commodity, inflation, as well as political and regulatory effects. The short-term deviations—defined as the difference between spot and equilibrium prices—are expected to revert toward zero following an Ornstein-Uhlenbeck process. These deviations may reflect, for example, short-term changes in demand resulting from variations in the weather or intermittent supply disruptions, and are tempered by the ability of market participants to adjust inventory levels in response to changing market conditions.³

Although neither of these factors is directly observable, the two factors in this model may be estimated from spot and futures prices. If there are long-maturity futures contracts for this commodity, then we can estimate these two state variables accurately over time: Intuitively, changes in the long-maturity futures prices give information about changes in the equilibrium price, and changes in the difference between near- and long-term futures prices give information about the short-term deviations. If there are no traded long-maturity futures contracts, then we must estimate the levels of these state variables and may have to treat them probabilistically. The short-term/long-term model of commodity prices developed in here is particularly convenient because it allows us to use standard Kalman filtering techniques to estimate these state variables in the same way as in Schwartz (1997).

³ Although one would intuitively expect positive price deviations to correspond to periods with low inventories (and vice versa), we do not explicitly model the role of storage decisions. See, e.g., Pindyck (1993) and Routledge et al. (1999) for more complex models that explicitly consider storage decisions and their impact on spot and futures prices.

Unlike most other recent models of commodity prices, this short-term/long-term model does not explicitly consider convenience yields—defined in Brennan (1991) as “the flow of services which accrues to the owner of a physical inventory but not to the owner of a contract for future delivery”—or stochastic convenience yields, even when valuing futures contracts or options on these futures. Nevertheless, our short-term/long-term model is exactly equivalent to the stochastic convenience yield model developed in Gibson and Schwartz (1990) in that the state variables in each model can be represented as linear combinations of the state variables in the other.

Although our short-term/long-term is formally equivalent to the Gibson-Schwartz model, we believe it has several advantages. While many find the notion of convenience yields elusive, the idea of stochastically evolving short-term deviations and equilibrium prices seems more natural and intuitive.⁴ Moreover, these factors are more “orthogonal” in their dynamics, which leads to analytic results that are more transparent and allow us to simplify the analysis of many long-term investments. The new formulation also clarifies some of the econometric issues that arise in Schwartz (1997). In particular, we identify two parameters of the model that cannot be estimated with much precision using futures prices and show that this indeterminacy is irrelevant for valuation purposes. Furthermore, the one-factor Ornstein-Uhlenbeck and geometric Brownian motion price models are simple (nested) special cases of this two-factor model, which facilitates empirical comparisons of model performance. The short-term/long-term easily outperforms these two simple models in our empirical comparisons. In addition, this article provides a more complete probabilistic analysis of the model and a deeper treatment of the use of Kalman filtering techniques (as compared to Schwartz 1997), characterizing the accuracy of the state variables estimates and relating Kalman filter estimates to implied estimates of the state variables.

⁴ Indeed others, notably Ross (1997) and Pilopovic (1998), have recently and independently developed multifactor price models where the two-factors are interpreted as short- and long-term factors.

We begin by formally defining the short-term/long-term model in §2 and deriving the distributions for future spot prices. In §3, we describe the risk-neutral version of the model and use it to derive closed-form expressions for futures prices and for European options on these futures. In §4, we establish the relationship between this model and the Gibson-Schwartz stochastic convenience yield model. In §5, we discuss the estimation of state variables and model parameters and, in §6, we present empirical estimates based on historical oil futures and forward prices as well as the model comparison results. In §7, we apply our model to some hypothetical real options problems. In §8, we describe some extensions of the model and offer concluding remarks.

2. The Short-Term/Long-Term Model

Let S_t denote the spot price of a commodity at time t . We will decompose spot prices into two stochastic factors as $\ln(S_t) = \chi_t + \xi_t$, where χ_t will be referred to as the short-term deviation in prices and ξ_t the equilibrium price level. (Seasonality can be incorporated by including time-dependent constants in this equation.) The short-run deviations (χ_t) are assumed to revert toward zero following an Ornstein-Uhlenbeck process

$$d\chi_t = -\kappa\chi_t dt + \sigma_\chi dz_{\chi_t} \tag{1}$$

and the equilibrium level (ξ_t) is assumed to follow a Brownian motion process

$$d\xi_t = \mu_\xi dt + \sigma_\xi dz_{\xi_t} \tag{2}$$

Here dz_χ and dz_ξ are correlated increments of standard Brownian motion processes with $dz_\chi dz_\xi = \rho_{\chi\xi} dt$. As indicated in the introduction, changes in the short-term deviations (χ_t) represent temporary changes in prices (resulting from, for example, unusual weather or a supply disruption) that are not expected to persist. Changes in the equilibrium level (ξ_t) represent fundamental changes that are expected to persist. The mean-reversion coefficient (κ) describes the rate at which the short-term deviations are expected to disappear and, as we will see below, $-\ln(0.5)/\kappa$ can be

interpreted as the “half-life” of the deviations—the time in which a deviation χ_t is expected to halve. The short-term/long-term model includes the standard geometric Brownian motion and Ornstein-Uhlenbeck price models as special cases when there is uncertainty about only one of the two factors.

We can write analytic forms for the distributions of the state variables and spot prices in the short-term/long-term model as follows. Given χ_0 and ξ_0 , following the derivation in the Appendix, we find that χ_t and ξ_t are jointly normally distributed with mean vector and covariance matrix:

$$E[(\chi_t, \xi_t)] = [e^{-\kappa t}\chi_0, \xi_0 + \mu_\xi t] \quad \text{and} \tag{3a}$$

$$\text{Cov}[(\chi_t, \xi_t)] = \begin{bmatrix} (1 - e^{-2\kappa t}) \frac{\sigma_\chi^2}{2\kappa} & (1 - e^{-\kappa t}) \frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} \\ (1 - e^{-\kappa t}) \frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} & \sigma_\xi^2 t \end{bmatrix}. \tag{3b}$$

Given χ_0 and ξ_0 , the log of the future spot price is then normally distributed with

$$E[\ln(S_t)] = e^{-\kappa t}\chi_0 + \xi_0 + \mu_\xi t \quad \text{and} \tag{4a}$$

$$\text{Var}[\ln(S_t)] = (1 - e^{-2\kappa t}) \frac{\sigma_\chi^2}{2\kappa} + \sigma_\xi^2 t + 2(1 - e^{-\kappa t}) \frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa}. \tag{4b}$$

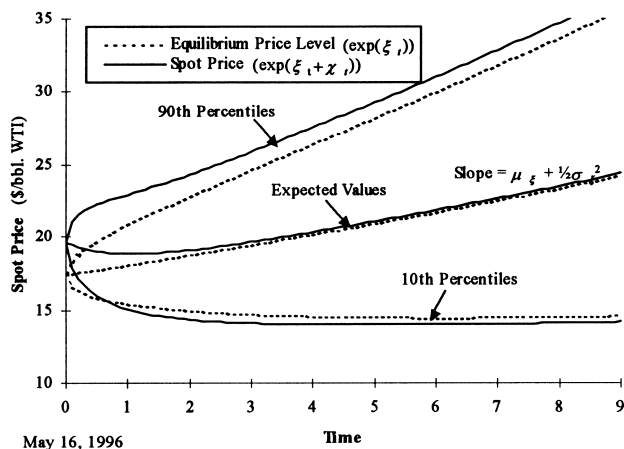
The spot price is then log-normally distributed with the expected price given by

$$E[S_t] = \exp(E[\ln(S_t)] + \frac{1}{2} \text{Var}[\ln(S_t)])$$

$$\begin{aligned} \text{or } \ln(E[S_t]) &= E[\ln(S_t)] + \frac{1}{2} \text{Var}[\ln(S_t)] \\ &= e^{-\kappa t}\chi_0 + \xi_0 + \mu_\xi t \\ &\quad + \frac{1}{2} \left((1 - e^{-2\kappa t}) \frac{\sigma_\chi^2}{2\kappa} \right. \\ &\quad \left. + \sigma_\xi^2 t + 2(1 - e^{-\kappa t}) \frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} \right). \end{aligned} \tag{5}$$

As the forecast horizon increases (i.e., as $t \rightarrow \infty$), the $e^{-\kappa t}$ and $e^{-2\kappa t}$ terms approach zero and the log of the expected spot price approaches

Figure 1 Probabilistic Forecasts for Spot and Equilibrium Prices



$$\left(\xi_0 + \frac{\sigma_\chi^2}{4\kappa} + \frac{\sigma_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} \right) + (\mu_\xi + \frac{1}{2}\sigma_\xi^2)t. \quad (6)$$

Thus, in the long run, the expected spot prices behave as if they started at an “effective long-run price” of $\exp(\xi_0 + \sigma_\chi^2/4\kappa + \rho_{\chi\xi}\sigma_\chi\sigma_\xi/\kappa)$ and grow at a rate of $(\mu_\xi + \frac{1}{2}\sigma_\xi^2)$. This effective long-run price is slightly different from the equilibrium price $\exp(\xi_0)$ with the constant difference reflecting the contribution of the short-term volatility to expected spot prices.

Figure 1 shows some example probabilistic forecasts generated by the model. Here we use parameter estimates based on the Enron data (shown in Table 2 of §6.1 and modified as described in Footnote 9) and show forecasts generated on May 16, 1996. On this date, the state variables are estimated to be $\chi_0 = 0.119$ and $\xi_0 = 2.857$, corresponding to a current spot price of \$19.61 ($=\exp(\chi_0 + \xi_0)$) and an equilibrium price of \$17.41 ($=\exp(\xi_0)$). The solid lines in Figure 1 describe forecasts for spot prices and the dashed lines describe forecasts for the equilibrium price. The center lines show the expected value forecasts for each variable and the upper and lower lines are “confidence bands” for each variable such that there is a 90% and 10% chance (respectively) that the variable will be below that amount on that particular date. Here we see that the spot price is expected to drop toward the equilibrium price, with the current deviation expected to be halved in 7 months ($= -\ln(0.5)/\kappa$). The equilibrium price is expected to grow over time at a constant rate

of 3.67% ($=\mu_\xi + \frac{1}{2}\sigma_\xi^2$). Comparing the two sets of confidence bands, we see that most of the uncertainty in near-term spot prices is due to uncertainty about the short-term deviations, but, after a few years, most of the uncertainty in spot prices is due to uncertainty about the then-prevailing equilibrium price.

3. Risk-Neutral Processes and Valuation

We now develop the “risk-neutral” version of the model that we will use to value futures contracts and options on these futures, as well as other commodity-related investments. In the risk-neutral valuation paradigm one uses risk-neutral stochastic processes to describe the dynamics of the underlying state variables, and discounts all cash flows at a risk-free rate (see, e.g., Duffie 1992 for a rigorous development of the risk-neutral valuation framework). In our model, we introduce two additional parameters, λ_χ and λ_ξ , that specify constant reductions in the drifts for each process. Specifically, we assume that these risk-neutral stochastic processes are of the form

$$d\chi_t = (-\kappa\chi_t - \lambda_\chi)dt + \sigma_\chi dz_\chi^* \quad \text{and} \quad (7a)$$

$$d\xi_t = (\mu_\xi - \lambda_\xi)dt + \sigma_\xi dz_\xi^* \quad (7b)$$

where again dz_χ^* and dz_ξ^* are increments of standard Brownian motion processes with $dz_\chi^* dz_\xi^* = \rho_{\chi\xi} dt$. The risk-neutral process for the short-term deviation (χ_t) is now an Ornstein-Uhlenbeck process reverting to $-\lambda_\chi/\kappa$ (rather than 0 as assumed in the true process), and the risk-neutral process for equilibrium prices is still a geometric Brownian motion, but now has drift $\mu_\xi^* \equiv \mu_\xi - \lambda_\xi$ (rather than μ_ξ as assumed in the true process). This form of risk adjustment is fairly standard (see, e.g., Schwartz 1995) and can be formally justified by assuming that the two state variables are priced according to the intertemporal asset pricing models developed in Merton (1973) and Cox et al. (1985). In these models, the risk premiums take the form of adjustments to the drift of the stochastic processes and, if we further assume that the correlations between changes in the state variables and

aggregate wealth in the economy are constant, then these reductions in drift would also be constant.⁵

Given χ_0 and ξ_0 , following a derivation similar to that of Equation (3), we find that under the risk-neutral process, χ_t and ξ_t are jointly normally distributed with mean vector and covariance matrix:

$$E^*[(\chi_t, \xi_t)] = [e^{-\kappa t}\chi_0 - (1 - e^{-\kappa t})\lambda_\chi/\kappa, \xi_0 + \mu_\xi^*t],$$

$$\text{Cov}^*[(\chi_t, \xi_t)] = \text{Cov}[(\chi_t, \xi_t)].$$

Here (and below) we use asterisks to denote expectations and variances taken with respect to the risk-neutral process defined by (7a) and (7b) rather than the “true” process defined by Equations (1) and (2). Under this risk-neutral process, the log of the future spot price, $\ln(S_t) = \chi_t + \xi_t$, is normally distributed with

$$E^*[\ln(S_t)] = e^{-\kappa t}\chi_0 + \xi_0 - (1 - e^{-\kappa t})\lambda_\chi/\kappa + \mu_\xi^*t, \quad (8a)$$

$$\text{Var}^*[\ln(S_t)] = \text{Var}[\ln(S_t)]. \quad (8b)$$

Comparing Equations (4) and (8), we see that the risk premiums reduce the log of the expected spot price by $(1 - e^{-\kappa t})\lambda_\chi/\kappa + \lambda_\xi t$; this premium depends on time but not the value of the state variables.

3.1. Valuing Futures Contracts

Let $F_{T,0}$ denote the current market price for a futures contract with time T until maturity. In the risk-neutral valuation framework, futures prices are equal to the expected future spot price under the risk-neutral process and, assuming that interest rates are deterministic (or, more generally, independent of spot prices), forward prices are equal to futures prices (see, e.g., Duffie 1992). By Equation (8), we can write the futures prices as

$$\begin{aligned} \ln(F_{T,0}) &= \ln(E^*[S_T]) \\ &= E^*[\ln(S_T)] + \frac{1}{2} \text{Var}^*[\ln(S_T)] \\ &= e^{-\kappa T}\chi_0 + \xi_0 + A(T) \end{aligned} \quad (9)$$


$$\begin{aligned} \text{where } A(T) &= \mu_\xi^*T - (1 - e^{-\kappa T}) \frac{\lambda_\chi}{\kappa} \\ &+ \frac{1}{2} \left((1 - e^{-2\kappa T}) \frac{\sigma_\chi^2}{2\kappa} + \sigma_\xi^2 T \right. \\ &\left. + 2(1 - e^{-\kappa T}) \frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} \right). \end{aligned}$$

The relationship between futures prices and expected spot prices is illustrated in Figure 2; here we use the same parameters as in Figure 1 and again show futures prices and forecasts from May 16, 1996. The actual futures prices for that day are marked with x 's in Figure 2 and will be discussed in §6. Comparing Equations (5) and (9), we see that the risk premium for the short-term deviations subtracts a constant amount (λ_χ/κ) from the effective long-run price—the time-0 intercept of the line supporting long-term expected spot prices in Figure 2 (this line is defined by Equation (6))—and the risk premium (λ_ξ) for the equilibrium level reduces the slope of long-term futures curve. Changes in the equilibrium price over time shift these long-term lines up and down, with the differences in intercepts and slopes remaining constant. Changes in short-term deviation lift or lower the front end of the futures curve and expected spot price curve, with the two curves sharing the same time-0 spot price.

From Equation (9), we see that the volatility of the price ($F_{T,0}$) for a futures contract maturing at time T —the instantaneous variance of $\ln(F_{T,0})$ —is given by $e^{-2\kappa T}\sigma_\chi^2 + \sigma_\xi^2 + 2e^{-\kappa T}\rho_{\chi\xi}\sigma_\chi\sigma_\xi$. The volatility is thus independent of the state variables. For near maturity futures contracts (i.e., $T = 0$), the volatility is equal to the volatility of the sum of the short- and long-term factors. As the maturity of the contract increases, the short-term deviations make less of a contribution to the volatility and, in the limit as $T \rightarrow \infty$, the instantaneous volatility approaches the volatility of the equilibrium price level (σ_ξ). This volatility relationship is illustrated in Figure 3 (again using the parameters from the Enron data) along with the empirical volatil-

⁵ One could easily extend this basic model to allow for a short-term risk premium that is a linear function of the short-term deviations; i.e., $\lambda_\chi = \beta\chi_t + \alpha$. This would allow the possibility that the short-term risk premium would be higher (or lower) in periods when spot prices are higher than the long-run equilibrium level. In this case, we could rewrite Equation (7a) with α in place of λ_χ , and $\kappa^* = \kappa + \beta$ in place of κ . Thus, with this extension, we would use the risk-neutral mean-reversion rate κ^* when working with the risk-neutral process and the true mean-reversion rate κ when working with the true process. We describe our estimates of this extended model in Footnote 10.

Figure 2 Futures Prices and Expected Spot Prices



ities calculated from the Enron data. Here we can see that the model fits the empirical volatilities well. The “option volatilities” shown there will be discussed shortly.

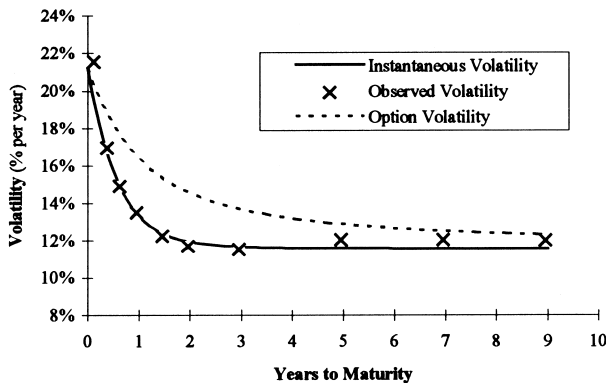
3.2. Valuing European Options on Futures Contracts

We can use the risk-neutral approach to derive analytic forms for the value of European options on futures contracts in this model. In this approach, the value of a European option on a futures contract is

given by calculating its expected future value using the risk-neutral process and discounting at the risk-free rate. Let $F_{T,t}$ denote the price at time t of a futures contract expiring at time T ($T > t$). From the formula for current futures prices (Equation (9)), we can write $F_{T,t}$ in terms of the time- t state variables as $\ln(F_{T,t}) = e^{-\kappa(T-t)}\chi_t + \xi_t + A(T-t)$. Given current state variables χ_0 and ξ_0 , because χ_t and ξ_t are jointly normally distributed under the risk-neutral process, $\phi \equiv \ln(F_{T,t})$ is also normally distributed with

$$\begin{aligned} \mu_\phi(t, T) &\equiv E^*[\ln(F_{T,t})] \\ &= e^{-\kappa(T-t)}E^*[\chi_t] + E^*[\xi_t] + A(T-t) \\ &= e^{-\kappa(T-t)}e^{-\kappa t}\chi_0 + \xi_0 + \mu_\xi^*t + A(T-t) \\ &= e^{-\kappa T}\chi_0 + \xi_0 + \mu_\xi^*t + A(T-t), \\ \sigma_\phi^2(t, T) &\equiv \text{Var}^*[\ln(F_{T,t})] \\ &= e^{-2\kappa(T-t)}\text{Var}^*[\chi_t] + \text{Var}^*[\xi_t] \\ &\quad + 2e^{-\kappa(T-t)}\text{cov}^*(\chi_t, \xi_t) \\ &= e^{-2\kappa(T-t)}(1 - e^{-2\kappa t})\frac{\sigma_\chi^2}{2\kappa} + \sigma_\xi^2 t \end{aligned}$$

Figure 3 Model and Empirical Volatilities



$$+ 2e^{-\kappa(T-t)}(1 - e^{-\kappa t}) \frac{\rho\chi\xi\sigma_\chi\sigma_\xi}{\kappa}.$$

One can verify that under the risk-neutral process, the expected futures price at time t ($E^*[F_{T,t}] = \exp(\mu_\phi(t, T) + \frac{1}{2}\sigma_\phi^2(t, T))$) is equal to the time- t current futures price $F_{T,0}$ (given by Equation (9)). This equality follows from the fact that futures prices are given by expected spot prices under the risk-neutral process ($F_{T,t} = E_t^*[S_T]$) and the law of iterated expectations ($E^*[E_t^*[S_T]] = E^*[S_T]$).

The fact that future futures prices are log-normally distributed under the risk-neutral process allows us to write a closed-form expression for valuing European put and call options on these futures. Explicitly, the value of a European call option on a futures contract maturing at time T , with strike price K , and time t until the option expires, is

$$e^{-rt}E^*[\max(F_{T,t} - K, 0)] = e^{-rt}(F_{T,0}N(d) - KN(d - \sigma_\phi(t, T))),$$

$$\text{where } d = \frac{\ln(F/K)}{\sigma_\phi(t, T)} + \frac{1}{2}\sigma_\phi(t, T)$$

and $N(d)$ indicates cumulative probabilities for the standard normal distribution (i.e., $P(Z < d)$). Similarly, the value of a European put with the same parameters is

$$e^{-rt}E^*[\max(K - F_{T,t}, 0)] = e^{-rt}(-F_{T,0}N(d) + KN(d - \sigma_\phi(t, T))).$$

These option valuation formulas are analogous to the Black-Scholes formulas for valuing European options on stocks that do not pay dividends. Here the stock price corresponds to the present value of the futures commitment ($e^{-rt}F_{T,0}$) and the equivalent annualized volatility would be $\sigma_\phi(t, T)/\sqrt{t}$. Figure 3 shows these annualized "option volatilities," assuming that the option expires at maturity of the futures contract (i.e., $t = T$). In this figure, we see that the annualized option volatility, representing an average of future futures volatilities, is greater than the instantaneous volatility of the underlying futures contract. As the maturity of the futures contract increases (i.e., as $T \rightarrow \infty$), the annualized option volatility ap-

proaches the volatility of the equilibrium level (σ_ξ) as most of the uncertainty about spot prices at maturity is a result of uncertainty about then-prevailing equilibrium prices.

4. Relationship to the Gibson-Schwartz Model

When compared to other recent models of commodity prices, this short-term/long-term model is unusual in that it makes no mention of convenience yields, let alone stochastic convenience yields. Yet, as mentioned in the introduction, the short-term/long-term model is equivalent to the stochastic convenience yield model developed in Gibson and Schwartz (1990) in that the factors in each model can be represented as linear combinations of the factors in the other.

To show this equivalence, we first briefly describe the Gibson-Schwartz model. Adopting the notation of Schwartz (1997), we let δ_t denote the time- t convenience yield and let X_t denote the log of the time- t current spot price. The stochastic convenience yield model assumes that these variables evolve according to

$$dX_t = (\mu - \delta_t - \frac{1}{2}\sigma_1^2)dt + \sigma_1dz_1, \tag{10}$$

$$d\delta_t = \kappa(\alpha - \delta_t)dt + \sigma_2dz_2, \tag{11}$$

where dz_1 and dz_2 are correlated increments of standard Brownian motion process with $dz_1dz_2 = \rho dt$. The convenience yield δ_t follows an Ornstein-Uhlenbeck process with equilibrium level α and rate of mean reversion κ , and plays a role in spot price through its appearance in the drift term in Equation (10). There is no overlap in notations for the parameters of the two models except for the mean-reversion parameter (κ); as we will see shortly, the mean-reversion parameters coincide in the two models.

The variables in the long-term/short-term model can be written in terms of the variables of the stochastic convenience yield model as follows:

$$\chi_t = \text{short-term deviation} = \frac{1}{\kappa}(\delta_t - \alpha), \tag{12}$$

$$\begin{aligned} \xi_t &= \text{equilibrium price level} \\ &= X_t - \chi_t = X_t - \frac{1}{\kappa} (\delta_t - \alpha). \end{aligned} \quad (13)$$

To establish the equivalence of the two models, we can write the stochastic process equations for the state variables of the short-term/long-term model using the equations for the stochastic convenience yield model (Equations (10) and (11)) and relate the parameters in the two models:

$$\begin{aligned} d\chi_t &= \frac{1}{\kappa} d\delta_t && \text{(using (12))} \\ &= (\alpha - \delta_t)dt + \frac{\sigma_2}{\kappa} dz_2 && \text{(using (10))} \\ &= -\kappa\chi_t dt + \frac{\sigma_2}{\kappa} dz_2 && \text{(using (12) again)} \end{aligned}$$

and

$$\begin{aligned} d\xi_t &= dX_t - \frac{1}{\kappa} d\delta_t && \text{(using (13))} \\ &= (\mu - \delta_t - \frac{1}{2}\sigma_1^2)dt + \sigma_1 dz_1 \\ &\quad - (\alpha - \delta_t)dt - \frac{\sigma_2}{\kappa} dz_2 && \text{(using (10) and (11))} \\ &= (\mu - \alpha - \frac{1}{2}\sigma_1^2)dt + \sigma_1 dz_1 - \frac{\sigma_2}{\kappa} dz_2. && \text{(rearranging)} \end{aligned}$$

Comparing these forms with Equations (1) and (2), we see that the two models are equivalent if we relate the parameters of the two models as shown in Table 1.

The risk premiums can be similarly related by equating parameters in the risk-neutralized versions of the two models. The risk-neutral version of the short-term/long-term model is described in Equations (7a) and (7b). In the stochastic convenience yield model, the risk-neutral process sets the drift of the spot price to be the risk-free rate (r) less the then-prevailing convenience yield δ_t . The convenience yield "paid" to the holder of the physical commodity is thus analogous to a dividend yield paid to the holders of common stocks. Explicitly, the risk-neutral processes are assumed to be of the form

$$\begin{aligned} dX_t &= (r - \delta_t - \frac{1}{2}\sigma_1^2)dt + \sigma_1 dz_1^*, \\ d\delta_t &= [\kappa(\alpha - \delta_t) - \lambda]dt + \sigma_2 dz_2^*, \end{aligned}$$

where dz_1^* and dz_2^* are correlated increments of standard Brownian motion process with $dz_1^* dz_2^* = \rho dt$, and r is the risk-free rate. Following an analysis similar to that used to equate parameters of the true processes, we can see that the values for λ_χ and λ_ξ shown in Table 1 make the two models equivalent. Given the equivalence of the two models, we can substitute terms from Table 1 into the equations derived earlier for valuing futures and options contracts in the short-term/long-term model and state the

Table 1 The Relationships Between Parameters in the Short-Term/Long-Term Model and the Stochastic Convenience Model of Gibson and Schwartz (1990)

Short-Term/Long-Term Model Parameter		
Symbol	Description	Definition in Terms of Stochastic Convenience Yield Model
κ	Short-term mean-reversion rate	κ
σ_χ	Short-term volatility	σ_2/κ
dz_χ	Short-term process increments	dz_2
μ_ξ	Equilibrium drift rate	$(\mu - \alpha - \frac{1}{2}\sigma_1^2)$
σ_ξ	Equilibrium volatility	$(\sigma_1^2 + \sigma_2^2/\kappa^2 - 2\rho\sigma_1\sigma_2/\kappa)^{1/2}$
dz_ξ	Equilibrium process increments	$(\sigma_1 dz_1 - (\sigma_2/\kappa) dz_2)(\sigma_1^2 + \sigma_2^2/\kappa^2 - 2\rho\sigma_1\sigma_2/\kappa)^{-1/2}$
$\rho_{\xi\chi}$	Correlation in increments	$(\rho\sigma_1 - \sigma_2/\kappa)(\sigma_1^2 + \sigma_2^2/\kappa^2 - 2\rho\sigma_1\sigma_2/\kappa)^{-1/2}$
λ_χ	Short-term risk premium	λ/κ
λ_ξ	Equilibrium risk premium	$\mu - r - \lambda/\kappa$

corresponding result in terms of the stochastic convenience yield model.

Counting the parameters involved in each of the two models, we find that the short-term/long-term model has a total of seven parameters ($\kappa, \sigma_{\chi}, \mu_{\xi}, \sigma_{\xi}, \rho_{\xi\chi}, \lambda_{\chi}, \lambda_{\xi}$) and the stochastic convenience yield model has eight parameters ($\kappa, \alpha, \sigma_2, \mu, \sigma_1, \rho, \lambda, r$). Given the formal equivalence of the two models, this suggests that one of the parameters in the stochastic convenience yield model is unnecessary or redundant. Indeed, the risk-free rate r is, in a sense, unnecessary for modeling spot prices: If we replace the risk-free rate r with $r + \Delta$ and, in compensation, replace δ_t by $\delta_t + \Delta$, α by $\alpha + \Delta$, and μ by $\mu + \Delta$, we find that X_t and the new δ_t follow the same true and risk-neutral processes and lead to the same estimates of the state variables in the short- and long-term model.⁶ The risk-free rate is thus not required for specifying the spot price dynamics (in either the true or risk-neutral process), for valuing futures or forward contracts, or for estimating the model from futures and forward prices. The risk-free rate would, however, be required to value many derivative securities (including options) and real assets using either the stochastic convenience yield model or the short-term/long-term model.

5. Estimating the State Variables and Model Parameters

As indicated in the introduction, the state variables in the short-term/long-term model are not directly observable and must be estimated from spot and/or futures prices. If both short- and long-maturity futures contracts are traded, intuitively, changes in the long-maturity futures prices give information about changes in the equilibrium price and changes in the difference between near- and long-term futures prices give information about the short-term deviations. If there are no traded long-maturity futures contracts, we may have to estimate the levels of the state variables and treat them probabilistically. In both

cases, the estimates can be generated using Kalman filtering techniques. The Kalman filter also facilitates the calculation of the likelihood of observing a particular data series given a particular set of model parameters; this allows us to estimate parameters using maximum likelihood techniques.

In this section, we briefly review the Kalman filter, discuss its use in estimating state variables and parameters from spot and futures prices, and finally talk about the use of implied methods for estimation. Detailed accounts of Kalman filtering are given in Harvey (1989) and West and Harrison (1996).

5.1. Kalman Filtering

The Kalman filter is a recursive procedure for computing estimates of unobserved state variables based on observations that depend on these state variables. Given a prior distribution on the initial value of the state variables and a model describing the likelihood of the observations as a function of the true values, the Kalman filter generates updated posterior distributions for these state variables in accordance with Bayes' rule. To formulate the short-term/long-term model for use with the Kalman filter, we will work with discrete time steps and define equations describing the evolution of the state variables and the relationship between the observed futures prices and the state variables. Here the state variables are the short-term deviation (χ_t) and equilibrium level (ξ_t), and the observations are the logs of the prices of the available futures contracts. Casting this relationship in terms of the Kalman filter, the evolution of the state variables is described by the *transition equation*, which from Equation (3) can be written as

$$\mathbf{x}_t = \mathbf{c} + \mathbf{G}\mathbf{x}_{t-1} + \boldsymbol{\omega}_t, \quad t = 1, \dots, n_T \quad (14)$$

where

$$\mathbf{x}_t \equiv [\chi_t, \xi_t], \quad \text{a } 2 \times 1 \text{ vector of state variables;}$$

$$\mathbf{c} \equiv [0, \mu_{\xi}\Delta t], \quad \text{a } 2 \times 1 \text{ vector;}$$

$$\mathbf{G} \equiv \begin{bmatrix} e^{-\kappa\Delta t} & 0 \\ 0 & 1 \end{bmatrix}, \quad \text{a } 2 \times 2 \text{ matrix;}$$

$\boldsymbol{\omega}_t$ is a 2×1 vector of serially uncorrelated, normally distributed disturbances with $E[\boldsymbol{\omega}_t] = \mathbf{0}$ and

$\text{Var}[\boldsymbol{\omega}_t] = \mathbf{W} \equiv \text{Cov}[(\chi_{\Delta t}, \xi_{\Delta t})]$ (given by Equation (3b));

⁶ If we derive the stochastic convenience yield model in terms of the short-term/long-term model, we see that the differences ($\delta_t - \alpha$), ($\mu - \alpha$), and ($r - \alpha$) are uniquely determined but the specific values of δ_t , α , μ , and r are not.

$\Delta t \equiv$ the length of the time steps; and
 $n_T \equiv$ the number of time periods in the data set.

The *measurement equation* describes the relationship between the state variables and the observed prices. From Equation (9), this is:

$$\mathbf{y}_t = \mathbf{d}_t + \mathbf{F}'_t \mathbf{x}_t + \mathbf{v}_t, \quad t = 1, \dots, n_T, \quad (15)$$

where:

$\mathbf{y}_t \equiv [\ln F_{T_1}, \dots, \ln F_{T_n}]$, a $n \times 1$ vector of observed (log) futures prices with time maturities T_1, T_2, \dots, T_n ;

$\mathbf{d}_t \equiv [A(T_1), \dots, A(T_n)]$, a $n \times 1$ vector;

$\mathbf{F}_t \equiv [e^{-\kappa T_1} 1, \dots, e^{-\kappa T_n} 1]$, a $n \times 2$ matrix; and

\mathbf{v}_t , a $n \times 1$ vector of serially uncorrelated, normally distributed disturbances with:

$$E[v_t] = 0, \quad \text{Cov}[v_t] = V.$$

Everything in this formulation is derived directly from our model with the exception of the introduction of the measurement errors (v_t). These measurement errors can be interpreted as representing errors in the reporting of prices (perhaps due to asynchronous price quotes) or, alternatively, as errors in the model's fit to observed prices.

Given these equations and a set of observed futures prices ($y_t, t = 1, \dots, n_T$), we run the Kalman filter recursively beginning with a prior distribution on the initial values of the state variables ($x_0 = [\chi_0, \xi_0]$). We assume the prior is multivariate normal with mean vector \mathbf{m}_0 and covariance matrix \mathbf{C}_0 . In each subsequent period, we use the observation y_t and the previous period's mean vector and covariance matrix to calculate the posterior mean vector and covariance matrix for the then-current state variables. Using the notation of Equations (14) and (15), the mean and covariance of the state variables conditioned on all of the information available at time t are given by:

$$E[\chi_t, \xi_t] = \mathbf{m}_t \equiv \mathbf{a}_t + \mathbf{A}_t(\mathbf{y}_t - \mathbf{f}_t), \quad (16a)$$

$$\text{Var}[\chi_t, \xi_t] = \mathbf{C}_t \equiv \mathbf{R}_t - \mathbf{A}_t \mathbf{Q}_t \mathbf{A}'_t, \quad (16b)$$

where $\mathbf{a}_t \equiv \mathbf{c} + \mathbf{G} \mathbf{m}_{t-1}$ and $\mathbf{R}_t \equiv \mathbf{G}_t \mathbf{C}_{t-1} \mathbf{G}'_t + \mathbf{W}$ are the mean and covariance of (χ_t, ξ_t) based on what is known at period $t - 1$. Similarly, $\mathbf{f}_t \equiv \mathbf{d}_t + \mathbf{F}'_t \mathbf{a}_t$ and $\mathbf{Q}_t \equiv \mathbf{F}'_t \mathbf{R}_t \mathbf{F}_t + \mathbf{V}$ are the mean and covariance of the

period- t futures prices given what is known at period $t - 1$. The matrix $\mathbf{A}_t \equiv \mathbf{R}_t \mathbf{F}_t \mathbf{Q}_t^{-1}$ defines a correction to the predicted state variables (\mathbf{a}_t) based on the difference between the (log) prices observed at time t (y_t) and the predicted time- t price vector (\mathbf{f}_t).

5.2. Estimation of State Variables from Spot and Futures Prices

The accuracy with which the state variables can be estimated depends on the kind and quality of information observed. Suppose that you observe only spot prices in each period. In this setting, there will always be some uncertainty in the estimates of the state variables because, given a change in spot price, it is impossible to tell whether the change is due to a change in the short-term deviation or a change in the equilibrium price or some combination thereof. After running the Kalman filter for a while, the variance in the state variable estimates (given by (16b)) will approach an asymptotic value that is independent of the particular price sequence observed or the assumed prior distributions (see West and Harrison 1996, p. 162). To illustrate the accuracy of the state variable estimates given spot price data, suppose we observe spot prices once each week (i.e., $\Delta t = 1$ week) and have zero measurement error associated with these observations. Using the volatility estimates from the Enron data (shown in Table 2), we find that the standard deviations for the estimates of the short-term deviations and equilibrium prices both asymptotically approach 0.07995. Thus, given only spot data, the state variables can only be estimated within approximately $\pm 8\%$.⁷

If there is uncertainty about the level of the state variables, the forecasts of spot and equilibrium prices given in §2 must be augmented to reflect this additional uncertainty. If we let the $\hat{\chi}_0$ and $\hat{\xi}_0$ denote the mean of the current state variable (given by Equation (16a) as $\mathbf{m}_t = [\hat{\chi}_0, \hat{\xi}_0]$) and $\hat{\sigma}_{\chi}^2, \hat{\sigma}_{\xi}^2$ and $\hat{\rho}_{\chi\xi}$ the corresponding variances and correlation coefficient (defined by the covariance matrix \mathbf{C}_t given by Equations

⁷ This asymptotic covariance matrix is difficult to compute analytically, but may be easily evaluated numerically. Increasing the frequency of the samples from weekly to daily or finer time steps does not significantly improve the ability to estimate the state variables: With hourly samples rather than weekly samples, the asymptotic standard deviations of this example are unchanged in the first five decimal places.

Table 2 Maximum-Likelihood Parameter Estimates

Parameter	Description	Futures Data		Enron Data			
		Estimate	Standard Error	Estimate	Standard Error		
κ	Short-term mean-reversion rate	1.49	0.03	1.19	0.03		
σ_χ	Short-term volatility	28.6%	1.0%	15.8%	0.9%		
λ_χ	Short-term risk premium	15.7%	14.4%	1.4%	8.2%		
μ_ξ	Equilibrium drift rate	-1.25%	7.28%	-3.86%	7.28%		
σ_ξ	Equilibrium volatility	14.5%	0.5%	11.5%	0.6%		
μ_ξ^*	Equilibrium risk-neutral drift rate	1.15%	0.13%	1.61%	0.12%		
$\rho_{\xi\chi}$	Correlation in increments	0.300	0.044	0.189	0.096		
	Standard deviation(s) of error for measurement equation						
		Contract Maturity		Contract Maturity			
S_1	"	1 mo.	0.042	0.002	2 mo.	0.027	0.001
S_2	"	5 mo.	0.006	0.001	5 mo.	0.006	0.001
S_3	"	9 mo.	0.003	0.000	8 mo.	0.000	
S_4	"	13 mo.	0.000		12 mo.	0.002	0.000
S_5	"	17 mo.	0.004	0.000	18 mo.	0.000	
S_6	"				2 yrs.	0.005	0.000
S_7	"				3 yrs.	0.014	0.003
S_8	"				5 yrs.	0.032	0.015
S_9	"				7 yrs.	0.043	0.036
S_{10}	"				9 yrs.	0.055	0.041

tion (16b)), then, from Equation (3), the mean and variance for the state variables at time t is given by:

$$E[(\chi_t, \xi_t)] = [e^{-\kappa t} \hat{\chi}_0, \hat{\xi}_0 + \mu_\xi t] \quad \text{and}$$

$\text{Cov}[(\chi_t, \xi_t)]$

$$= \begin{bmatrix} (1 - e^{-2\kappa t}) \frac{\sigma_\chi^2}{2\kappa} & (1 - e^{-\kappa t}) \frac{\rho_{\chi\xi} \sigma_\chi \sigma_\xi}{\kappa} \\ (1 - e^{-\kappa t}) \frac{\rho_{\chi\xi} \sigma_\chi \sigma_\xi}{\kappa} & \sigma_\xi^2 e^{-\kappa t} \end{bmatrix} + \begin{bmatrix} e^{-2\kappa t} \hat{\sigma}_\chi^2 & e^{-\kappa t} \hat{\rho}_{\chi\xi} \hat{\sigma}_\chi \hat{\sigma}_\xi \\ e^{-\kappa t} \hat{\rho}_{\chi\xi} \hat{\sigma}_\chi \hat{\sigma}_\xi & \hat{\sigma}_\xi^2 \end{bmatrix}.$$

Comparing this with Equation (3), we see that the uncertainty about the current state variables serves to increase uncertainty about their future values by adding terms to the covariance matrix. Although there may be considerable uncertainty about the values of the state variables at any time, this uncertainty has relatively little impact on forecasts and futures prices. If the state variable estimates are based on spot price observations, the state variable estimates will be per-

fectly negatively correlated (since they must sum to the observed log spot price) and the errors will cancel each other in the short-term spot price forecasts. In the longer term, the uncertainty about the short-term deviation has little impact (as $e^{-\kappa t}$ approaches 0), and the uncertainty about the current equilibrium price (with variance $\hat{\sigma}_\xi^2$) grows small compared with the uncertainty due to potential changes in equilibrium prices (with variance $\sigma_\xi^2 t$).

If, instead of observing only spot prices, we observe prices for a vector of futures contracts with varying maturities, there will typically be little uncertainty about the state variables. In fact, if we observe prices for two contracts with different maturities and have zero measurement error, we can invert Equation (15) and estimate the state variables exactly. With multiple contracts and measurement errors for all contracts, we cannot estimate the state variables exactly, but if the measurement errors are small, there will be little uncertainty in the state variable estimates. In the empirical results of the next section, there is essentially zero error in the state variable estimates.

5.3. Parameter Estimation

The Kalman filtering procedure allows us to estimate the state variables over time given particular assumptions about the parameters of the process; all of the previous probabilistic results assumed that the parameters of the process were known. The Kalman filtering paradigm also allows one to efficiently calculate the likelihood of a set of observations given a particular set of parameters (see, e.g., Harvey 1989, Chapter 3.4 for details). By varying the parameters and rerunning the Kalman filter for each set of parameters, we can identify the set of parameters that maximizes this likelihood function. In our short-term/long-term model, there are seven model parameters to be estimated ($\kappa, \sigma_{\chi}, \mu_{\xi}, \sigma_{\xi}, \rho_{\xi\chi}, \lambda_{\chi}, \mu_{\xi}^*$) plus the terms in the covariance matrix for the measurement errors (\mathbf{V}). In general, there are $(n + 1)n/2$ free variables in the covariance matrix, where n is the number of futures contracts whose prices are observed (the matrix must be symmetric). As in Schwartz (1997), we simplify the estimation problem by assuming that \mathbf{V} is diagonal with diagonal elements (s_1^2, \dots, s_n^2) . We used the "maxlik" routine in Gauss to numerically determine parameter estimates and standard errors for these estimates. To be sure that our maximum likelihood estimation routine reaches a global (rather than local) maximum, we reran the optimization problem from a variety of initial parameter values. In all cases, we started the Kalman filter with a prior mean (\mathbf{m}_0) and covariance matrix (\mathbf{C}_0) based on the observed means and covariance in the data. Although the likelihood scores vary somewhat, the estimated state variables and parameters did not appear to be very sensitive to the assumed initial mean and covariance.

5.4. Implied State Variable and Parameter Estimates

In some applications, rather than gathering historical futures data and running the Kalman filter, it may be easier to select state variables to fit the current futures curve. For example, given estimates of the model parameters, one might set up a spreadsheet that calculates model futures prices as a function of the current state variables (using Equation (9)) and then choose values for the state variables to minimize the squared differences between the observed futures prices and those given by

the model. Graphically in Figure 2, we pick an equilibrium price and deviation to fit the model's futures prices to the observed futures prices, marked with x s in the figure. The errors in the fit could be weighted to improve the accuracy of the fit to particular parts of the futures curves. These "implied" state variable estimates may differ from those generated by running the Kalman filter (notably, they do not depend on any prior state variable estimates or any previous futures prices), but if the futures prices are sufficiently informative and the errors are weighted appropriately, these implied estimates will be very similar to those generated by the Kalman filter.⁸ One could also incorporate option prices into this process using the formulas of §3.2 to value options as a function of the state variable estimates.

Implied methods could also be used to determine some of the model parameters as well as the state variables. For example, one might choose the risk-neutral equilibrium drift rate μ_{ξ}^* and the mean-reversion rate κ to fit the current futures curve. Similarly, given historical volatilities for futures prices with varying maturities or implied volatilities for options on near-term and long-term futures, we could estimate the volatility parameters ($\sigma_{\chi}, \sigma_{\xi}$, and σ_{ξ}) by choosing volatility parameters to fit the observed futures or options volatilities (i.e., fitting the

⁸ To clarify this claim, first note using the definition of \mathbf{f}_t , we can rewrite Equation (16a) as $\mathbf{m}_t = E[\chi_t, \xi_t] = (\mathbf{I} - \mathbf{A}_t \mathbf{F}_t') \mathbf{a}_{t-1} + \mathbf{A}_t (\mathbf{y}_t - \mathbf{d}_t)$. Assuming \mathbf{V} and \mathbf{R}_t are invertible, \mathbf{A}_t can be rewritten as $\mathbf{A}_t = \mathbf{R}_t' \mathbf{F}_t' (\mathbf{F}_t' \mathbf{R}_t \mathbf{F}_t + \mathbf{V})^{-1} = (\mathbf{F}_t \mathbf{V}^{-1} \mathbf{F}_t' + \mathbf{R}_t^{-1})^{-1} \mathbf{F}_t \mathbf{V}^{-1}$; the first equality follows from the definitions following Equation (16), and the second equality can be verified by postmultiplying both sides by $(\mathbf{F}_t' \mathbf{R}_t \mathbf{F}_t + \mathbf{V})$ and premultiplying by $(\mathbf{F}_t \mathbf{V}^{-1} \mathbf{F}_t' + \mathbf{R}_t^{-1})$. The futures prices will be "very informative" if $\mathbf{F}_t \mathbf{V}^{-1} \mathbf{F}_t'$ is large compared to \mathbf{R}_t^{-1} ; intuitively this requires the measurement errors associated with the futures (with covariance matrix \mathbf{V}) to be small compared to the uncertainty about the values of the state variables before seeing the futures prices (with covariance matrix \mathbf{R}_t). In this case, $\mathbf{F}_t \mathbf{V}^{-1} \mathbf{F}_t'$ will dominate \mathbf{R}_t^{-1} , \mathbf{A}_t will approach $(\mathbf{F}_t \mathbf{V}^{-1} \mathbf{F}_t')^{-1} \mathbf{F}_t \mathbf{V}^{-1}$, and $\mathbf{A}_t \mathbf{F}_t'$ will approach \mathbf{I} . The means given by Equation (16a) therefore approach $\mathbf{m}_t = E[\chi_t, \xi_t] = (\mathbf{F}_t \mathbf{V}^{-1} \mathbf{F}_t')^{-1} \mathbf{F}_t \mathbf{V}^{-1} (\mathbf{y}_t - \mathbf{d}_t)$. This expression is also the solution to the weighted least squares problem of choosing \mathbf{m}_t to minimize $(\mathbf{y}_t - \mathbf{d}_t - \mathbf{F}_t' \mathbf{m}_t)' \mathbf{V}^{-1} (\mathbf{y}_t - \mathbf{d}_t - \mathbf{F}_t' \mathbf{m}_t)$. Thus when the futures contracts are very informative, the Kalman filter estimates are equivalent to minimizing the sum of squared errors in the model's fit to the current futures curve, provided we weight the errors according to the precision matrix (\mathbf{V}^{-1}) for the measurement errors.

curves shown in Figure 3 to the observed data). While these estimates will not be identical to those given by maximum likelihood estimation, they may be easier to generate and sufficient for many purposes.


6. Empirical Results

We describe estimates of our model based on two different data sets; both were used in Schwartz (1997) and are described in more detail there. In the first data set, the observations consist of weekly observations of prices for NYMEX crude oil futures contracts maturing in the next month and in approximately 5, 9, 13, and 17 months (5 contracts total). We use futures prices from 1/2/90 to 2/17/95 with a total of 259 sets of observations of 5 futures prices. These prices are publicly available and were obtained from Knight-Ridder Financial. The second data set consists of proprietary historical crude oil forward price curves made available by Enron Capital and Trade Resources. This data set covers the time period from 1/15/93 to 5/16/96, and for each date we use prices for 10 forward contracts, maturing in approximately 2, 5, and 8 months and in 1, 1.5, 2, 3, 5, 7, and 9 years. The Enron data set includes a total of 163 sets of observations of 10 forward prices.

6.1. State Variable and Parameter Estimates

Table 2 shows maximum likelihood parameter estimates for each data set and Figure 4 shows the estimated values of the equilibrium price (given as $\exp(\xi_i)$) and spot price ($\exp(\xi_i + \chi_i)$) for the futures data; the state variable estimates for the Enron data set are similar in the time period where the two data sets overlap. Both data sets show significant mean reversion in the short-term deviations: In the futures data, the "half-life" of the short-term deviations is approximately 6 months ($= -\ln(0.5)/\kappa$) and, in the Enron data set, the half-life is about 7 months. In both data sets, we see that the spot prices are much more volatile than the equilibrium prices, reflecting the substantial short-term volatility. The spot prices were sometimes above and sometimes below the equilibrium price level with the greatest differences occurring during the Gulf War in the summer and fall of 1990 (see Figure 4), when spot prices rose above \$40 per barrel

Figure 4 Estimated Spot and Equilibrium Prices for the Futures Data



while the equilibrium price levels reached only \$25 per barrel. Spot prices were well below the equilibrium levels during the last quarter of 1993 and first quarter of 1994. This was a period when "ongoing high production levels, by both Organization of Petroleum Exporting Countries (OPEC) and other countries, were more than sufficient to satisfy stagnant global demand, resulting in a continuing increase in worldwide petroleum inventory levels" (Energy Information Administration (EIA) 1994, p. xi). The fact that equilibrium prices did not follow spot prices down suggests that market participants did not expect this excess production to continue and, in fact, it did not: "The onset of warm weather, speculation on changes in self-imposed OPEC production levels, and tight supplies all helped U.S. crude oil prices gradually climb to the year's high, \$20.72, on June 16 [1994]" (EIA 1995, p. xiii).

Table 3 shows the errors in the model's fit to the futures prices. In general, the model fits the mid-term contracts best with larger errors for the very short- and very long-term contracts. The largest errors were in fitting the near-term contract in the futures data (a mean absolute error in log prices of 0.0314) and in the 9-year contract in the Enron data (a mean absolute error of 0.0332). In both data sets, some of the mid-term futures prices were matched with essentially no error, leading to a measurement error matrix (\mathbf{V}) that is positive semidefinite but not positive definite. The accuracy of the model's fit for particular contracts is determined to a large extent by our assumptions about

Table 3 Errors in the Model Fit to the Logarithm of Futures Prices

Futures Data				Enron Data			
Contract Maturity	Mean Error	Standard Deviation of Error	Mean Absolute Error	Contract Maturity	Mean Error	Standard Deviation of Error	Mean Absolute Error
1 mo.	-0.0053	0.0414	0.0314	2 mo.	0.0104	0.0249	0.0207
5 mo.	0.0005	0.0044	0.0035	5 mo.	0.0017	0.0058	0.0049
9 mo.	-0.0002	0.0025	0.0020	8 mo.	0.0000	0.0000	0.0000
13 mo.	0.0000	0.0000	0.0000	12 mo.	-0.0006	0.0023	0.0019
17 mo.	0.0000	0.0035	0.0028	18 mo.	0.0000	0.0000	0.0000
				2 yrs.	0.0012	0.0047	0.0040
				3 yrs.	0.0040	0.0133	0.0121
				5 yrs.	0.0042	0.0320	0.0276
				7 yrs.	-0.0076	0.0429	0.0318
				9 yrs.	-0.0245	0.0493	0.0332

the measurement errors. If, for example, we want to be sure that the model perfectly replicates spot prices, we could choose a measurement error covariance matrix (\mathbf{V}) with zero variance for the spot prices: With two state variables to estimate in each period, we can perfectly match prices for up to two contracts in each period. The measurement error standard deviations reported in Table 2 were selected to maximize the likelihood of the data and can be thought of as providing the best overall fit to the data. However, for some applications, we might choose different error covariance structures to obtain better fits for particular contract maturities.

Examining the standard errors for the parameter estimates in Table 2, we see that two parameters, the long-term drift (μ_ξ) and short-term risk premium (λ_χ), are not estimated with much accuracy. This indeterminacy can be explained graphically using Figure 2. Because our observations consist of futures prices (marked with xs in Figure 2), in each period, we get good estimates of the spot price ($\chi_t + \xi_t$) and the time-0 intercept of the line supporting the long-term futures price. We also get a good estimate of the risk-adjusted growth rate ($\mu_\xi^* = \mu_\xi - \lambda_\xi$) because it is the average growth rate for long-term futures prices. The expected spot prices—represented by the upper curve in Figure 2—are, however, never directly observed, and we cannot accurately determine the precise location of this curve or its long-term slope. The risk premiums λ_χ and λ_ξ describe the differences

between the expected prices and futures prices and, because price expectations are not observed, these risk premiums are not well estimated. Errors in the estimate of λ_ξ appear in Table 2 as errors in the estimate of μ_ξ , and errors in the estimate of λ_χ shift all of the estimates of ξ_t up or down by a constant (λ_χ/κ), with the χ_t adjusting accordingly so as to preserve the sum ($\chi_t + \xi_t$) corresponding to the log of the observed spot price. In essence, using futures data, we can precisely estimate the risk-neutral process for spot prices but cannot precisely estimate the true process. We can see this formally by examining the risk-neutral distribution for spot prices defined by Equation (8). First note that, if we are given the risk-neutral drift (μ_ξ^*), the true equilibrium drift (μ_ξ) plays no role in the risk-neutral distribution for spot prices. Second, note that if we replace λ_χ by $\lambda_\chi + \Delta$ for any Δ , and in compensation replace χ_0 by $\chi_0 - \Delta/\kappa$ and ξ_0 by $\xi_0 + \Delta/\kappa$, the risk-neutral distribution for spot prices is unchanged. The two parameters that we cannot estimate thus do not affect the risk-neutral distributions for spot prices (and, more generally, the risk-neutral stochastic process for spot prices) and hence do not affect the valuation of securities and projects that depend only on spot prices.

These two dimensions of indeterminacy in the parameters in the short-term/long-term model (or the corresponding dimensions in the Gibson-Schwartz model) do not affect the robustness of the model for use in valuation problems, although it does affect its robustness for forecasting purposes. To precisely esti-

mate the risk premiums (λ_χ and λ_ξ) and true values of the state variables, we would have to use a much longer time series or include observations that depend on the true price process rather than the risk-neutral process (for example, published price forecasts). Alternatively, we might use implied techniques like those discussed in §5.4 to choose μ_ξ and λ_χ to match exogenously specified price forecasts.^{9,10}

6.2. Model Comparisons

We can also use these data sets to compare our model with commonly used single-factor models in terms of their abilities to capture the dynamics of futures prices. As indicated in §2, our model includes geometric Brownian motion and Ornstein-Uhlenbeck processes as special cases when there is no uncertainty about one of the two state variables. The geometric Brownian motion model is given by considering uncertainty in equilibrium prices only and taking χ_0 , σ_χ , and λ_χ to be zero. The geometric Ornstein-Uhlenbeck model is given by assuming a con-

⁹ Because equilibrium prices declined over this time period, we obtain negative estimates for the expected growth of the equilibrium price level (μ_ξ) in both data sets. This is not likely to be representative of investor expectations of long-run growth during that time period, and leads to a negative estimate of the long-term risk premium ($\lambda_\xi = \mu_\xi - \mu_\xi^* = -3.9\% - 1.6\% = -5.5\%$, for the Enron data). If we fix the long-term expected growth rate at a value that is more representative of investor expectations, say 3%, then the long-term risk premium becomes well estimated and positive ($\lambda_\xi = 1.4\%$ with standard error of 0.12% for the Enron data); all other parameter estimates are essentially unchanged. The plots of Figures 1 and 2 both assume $\mu_\xi = 3\%$ and use the estimate $\lambda_\xi = 1.4\%$ to generate price forecasts that are more representative of investor expectations. To improve the readability of Figure 2, we have taken $\lambda_\chi = 5\%$ and modified the estimates of χ_i and ξ_i to preserve the spot price and futures prices. These modifications change the spot price forecasts but preserve the model's estimates of futures prices and the risk-neutral process for spot prices.

¹⁰ If we allow the short-term risk premium to depend on the short-term deviations as discussed in Footnote 5, using the Enron data, we find an estimate of the risk-neutral mean-reversion rate (κ^*) of 1.19 (with a standard error of 0.03) and a true mean-reversion rate (κ) of 1.79 (with a standard error of 0.87). These point estimates suggest that short-term risk premiums are lower in periods with higher short-term deviations (i.e., $\beta = \kappa^* - \kappa = 1.19 - 1.79 = -0.60$), but given the magnitude of the standard error for κ , we cannot draw any conclusions about the sign of this effect. Thus, here again, we can estimate the parameters of the risk-neutral process well but cannot estimate corresponding parameters of the true process well.

Table 4 Log-Likelihood Scores for Model Comparison

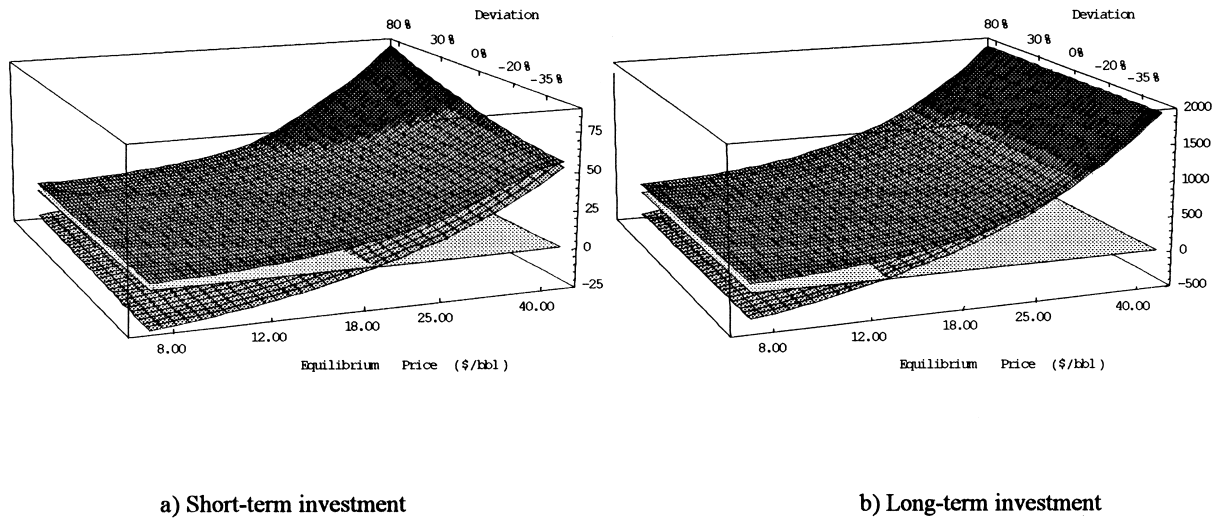
Model	Futures Data	Enron Data
Geometric Brownian Motion	3860	4995
Geometric Ornstein-Uhlenbeck Process	4331	5574
Short-term/Long-term Model	5140	6182

stant equilibrium price (which must be estimated) and taking $\mu_{\xi'}$, $\sigma_{\xi'}$, and λ_ξ to be zero.

Table 4 shows the log-likelihood scores given by maximum likelihood estimation of each model with each data set. We see that the two-factor model has the largest log-likelihood scores for each of the data sets. Since the simpler models are restrictions of our two-factor model, we can compare the differences in log-likelihood scores for each data set to see whether the additional parameters of the two-factor model provide a statistically significant improvement in the model's ability to explain the observed data. The relevant test statistic for this comparison is the chi-squared distribution with 3 degrees of freedom (the one-factor models are obtained by placing 3 restrictions on parameters in the two-factor model) and the 99th percentile of this distribution is 11.34. Given that the log-likelihood scores increase by more than 600 in all cases, we see that the improvements provided by the short-term/long-term model are quite significant.

We can understand why the short-term/long-term model outperforms these other models by considering the shapes of the futures curves generated by each of these models. Figure 2 shows an example for the short-term/long-term model. In this figure, we see a situation with a positive short-term deviation and the futures curve starts high, decreases rapidly, and then trends back up. In cases with negative short-term deviations, the initial part of the curve is below the long-term trend. If we observe the time series of futures curves, we see the short end of the curve "flapping" around the long-term trend as well as vertical shifts in the long-term trend. The futures curve implied by the simple geometric Brownian motion model would appear in Figure 2 as a straight line; while this model allows this line to shift up and down, it cannot capture the "flapping" that we see at the short-end of the futures curve. The futures curve implied by the Ornstein-Uhlenbeck process allows "flapping" at

Figure 5 Value Functions for the Short- and Long-Term Investments



the short end of the curve, but it does not allow for the vertical shifts at the long end of the curve. The short-term/long-term model allows both the short-term “flapping” and the long-term shifts and hence fits the data much better than these simpler models.

7. Illustrative Real Options Applications

To illustrate some of the possible uses of the short-term/long-term model, we consider two hypothetical investments in oil properties. The first example is representative of a long-term investment. In this case, the firm owns the right to develop a property and can exercise this option at any time. It costs \$800,000 to develop the property and there is a 3-year construction lag from the time that the development decision is made and oil production begins. After the lag, oil production starts at a rate of 5,000 barrels per year and declines exponentially at a rate of 5% per year. The second example is representative of a short-term development option. In this case, development costs \$40,000, there is no construction lag, and oil is produced at an initial rate of 1,000 barrels of oil per year, declining exponentially at a rate of 40% per year. To simplify the analysis, for both projects we assume that there are no operating costs, royalties or taxes and that, once production starts, it continues indefinitely.

The problem is to determine the optimal exercise strategy and the value of these investments in the different possible price states.

We value these example projects using the risk-neutral technique, using parameter estimates based on the Enron data (given in Table 2) and discounting at a risk-free rate of 5% per year. Figures 5a and 5b display the results for these examples, showing the value and optimal exercise policies as a function of the equilibrium and deviation state variables.¹¹ The dark gray, upper surfaces in these figures show the value of the property under the optimal exercise policy. The lighter gray, lower surfaces show the value of the property if the firm were forced to exercise immediately. The darkest regions in each figure

¹¹ These values and policies were calculated by solving a discrete-time, infinite-horizon dynamic-programming problem where, in each period, the owner can exercise the option to develop the field or hold it until the next period. The value at exercise is given as the expected discounted revenue over the infinite life of the project; these values were calculated numerically. The equilibrium and deviation price variables were discretized into 35 steps each, for a statespace with a total of 1,225 different states; time was discretized into intervals of 0.1 year. The dynamic program was formulated as a linear program (as described in Bertsekas 1995, p. 49) and solved using a commercial linear programming package (CPLEX). For finite-horizon problems, we can solve for optimal policies and values using multivariate lattice models, analogous to those developed in Boyle et al. (1989) for multivariate geometric Brownian motions.

show where the two value surfaces coincide and it is optimal to exercise the option to develop the field. The zero plane is shown to provide a reference. The equilibrium prices are labeled in dollars per barrel (given as $\exp(\xi_i)$) and the deviations are labeled as a percentage increase or decrease over the equilibrium price (given as $\exp(\chi_i) - 1$). The spot price is given as the equilibrium price plus this proportional increase or decrease ($S_t = \exp(\xi_i + \chi_i) = \exp(\xi_i) \exp(\chi_i)$).

It is interesting to compare the sensitivities of the optimal strategies to the state variables in the two examples. In the short-term investment, the values and policies are sensitive to both state variables and the value increases with both the short-term deviation and the equilibrium price. The equilibrium price at which one would choose to exercise the option changes substantially depending on the short-term deviation. For example, for a +70% deviation, the threshold equilibrium price is approximately \$23 per barrel. With negative deviations, we do not exercise for any equilibrium price under \$40 per barrel: In these cases, the optimal policy suggests waiting for the deviations to again turn positive. In some of these cases, the optimal policy suggests waiting even though the value given immediate exercise is higher than in other cases where exercise is optimal. The short-term deviations thus have an impact on the timing of the investment decision beyond their impact on the value of the investment.

In contrast, the values and policies for the long-term investment are quite insensitive to the short-term deviations. Even though the investment has positive value whenever the equilibrium price exceeds approximately \$13 per barrel, it is not optimal to actually exercise the option until prices reach approximately \$30 per barrel; these equilibrium price thresholds change only slightly with the short-term deviations. This insensitivity for the long-term investment is a result of the three-year construction lag and the long productive lifetime, both of which dampen the effect of the short-term deviations. Because of this insensitivity, when valuing long-term investments like this one, we can simplify our analysis by reducing the two-factor model to a single-factor model that considers uncertainty in equilibrium prices only. In these

applications, we would use the two-factor model to estimate the current equilibrium price (since it is not directly observed) but we would not model the short-term deviations stochastically.

The ability to simplify the analysis of long-term investments is one of the advantages of the short-term/long-term model over the Gibson-Schwartz model and other similar models. As discussed in §4, assuming no uncertainty in the short-term deviations in the short-term/long-term model is equivalent to assuming no uncertainty in convenience yields in the Gibson-Schwartz model. But if we assume no uncertainty in convenience yields in the Gibson-Schwartz model, we would use a volatility for the spot prices that reflects uncertainty in both the short-term deviations and equilibrium prices (given by $\sigma_{\chi}^2 + \sigma_{\xi}^2 + 2\rho_{\chi\xi}\sigma_{\chi}\sigma_{\xi}$ in our notation and given by σ_1^2 in the Gibson-Schwartz model) and is inappropriately large for long-term investments: Using the Enron data, the instantaneous volatility for spot prices is approximately 21% per year whereas the volatility for the equilibrium price is only 11.5%. The factorization in the short-term/long-term model can thus help simplify the evaluation of long-term commodity investments in a way that is not so easily achieved using the Gibson-Schwartz model. (See Schwartz 1998 for a discussion of how to use the Gibson-Schwartz model to value long-term commodity assets.)

8. Summary and Conclusions

In this article, we propose a new way of thinking about the stochastic behavior of commodity prices and develop a two-factor model that allows for short-term mean-reverting variations in prices, and at the same time allows uncertainty in the equilibrium level to which prices revert. Although this short-term/long-term model makes no mention of convenience yields, the model turns out to be exactly equivalent to the stochastic convenience yield model of Gibson and Schwartz (1990) with the short-term price deviation being a linear function of the instantaneous convenience yield. The short-term/long-term model thus provides an alternative interpretation of the results of the stochastic convenience yield model in which changes in short-term futures prices are interpreted as

short-term price variations rather than as changes in the instantaneous convenience yield.

Although our short-term/long-term model and the Gibson-Schwartz stochastic convenience yield model are formally equivalent, we believe that the short-term/long-term model is easier to interpret and work with for several reasons. First, whereas many find it hard to think about “convenience yields” (let alone stochastic convenience yields), the notions of short-term variations and long-term equilibrium price levels seem natural and lead to results that are more transparent. For example, in the short-term/long-term model we find that the volatility of prices for futures contracts is given by the volatility of the sum of the short- and long-term factors. As the maturity of the contract increases, the futures volatility approaches the volatility of the equilibrium price.

A second advantage of the short-term/long-term model is that the two factors in the short-term/long-term model are more “orthogonal” in their dynamics. In the Gibson-Schwartz stochastic convenience yield model, the convenience yield plays a role in the stochastic process for the spot price. In the short-term/long-term model, the only interaction between factors comes through the correlation of their stochastic increments (estimated at 0.189 and 0.300 for the two data sets), and this correlation in increments is much less than the correlation between increments in the stochastic convenience yield model (estimated at 0.845 and 0.922 for the same two data sets). This orthogonality allows us to think more clearly about the impacts of each factor when evaluating commodity-related projects and derivative securities. In particular, for many long-term investments, we may be able to safely ignore the short-term variations and evaluate investments using a one-factor model that considers uncertainty in equilibrium prices only, modeled using a standard geometric Brownian motion process.

By separating short- and long-term price components and using futures prices to distinguish between them, we provide a conceptual framework for developing richer models of commodity price movements. We have developed one such extension in which the growth rate for the equilibrium price (μ_t) is stochastic (the details are available from the authors); this third factor improves the model’s ability to fit long-term

futures prices and match changes in the growth rate for long-term futures prices. Many other extensions are possible. To improve the model’s ability to fit short-term futures prices, we might consider the possibility of allowing the deviation reversion rate (κ) to be stochastic, incorporating lagged measurement errors, or, to better fit option prices, allowing the short-term volatility (σ_λ) to be stochastic. For some commodities, like electricity, we might consider adding Poisson jumps to capture the impact of system failures. To improve the modeling of long-term price uncertainty, we might consider alternative models of equilibrium price movements. For example, one might consider the use of a simple mean-reverting model for equilibrium prices or an equilibrium price model, like that discussed in Pindyck (1997), that allows “u-shaped” price trajectories for long-run prices.

Although these extensions may improve the performance of the model in terms of its ability to describe the stochastic evolution of spot and futures prices, we must balance our desire for fidelity in the price models with the need for parsimony in the models used to evaluate complex real or financial options. Given the difficulty of valuing options on several state variables (especially American-style options) and the fact that price is typically one of many relevant uncertainties in real options applications, the more complex models may not be suitable for use in these applications. Although the simple two-factor model that we develop in this article is more complex than the commonly used one-factor models, we believe that this additional complexity provides a much more realistic model of the short- and long-term dynamics of commodity prices that can improve the quality of the valuations with a minimum of additional effort.¹²

¹² Jim Smith’s work on this project was supported in part by the National Science Foundation under grant no SBR-9809176. The authors are grateful for the helpful comments provided by an associate editor and two anonymous referees.

Appendix

Derivation of Equation (3). We proceed by first finding the mean vector and covariance matrix for a discrete-time approximation of the process based on the stochastic differential Equations (1) and (2), and then take the limit as the time steps are made infinitesimally small. The discrete-time approximation of the pro-

cess with time steps of length $\Delta t = t/n$ can be written as $\mathbf{x}_t = \mathbf{c} + \mathbf{Q}\mathbf{x}_{t-1} + \boldsymbol{\eta}_t$, where $\mathbf{x}_t \equiv [\chi_t, \xi_t]$, $\mathbf{c} \equiv [0, \mu_\xi \Delta t]$,

$$\mathbf{Q} \equiv \begin{bmatrix} \phi & 0 \\ 0 & 1 \end{bmatrix},$$

$\phi \equiv 1 - \kappa \Delta t$, and $\boldsymbol{\eta}_t$ is a 2×1 vector of serially uncorrelated, normally distributed disturbances with $E[\boldsymbol{\eta}_t] = 0$, and

$$\text{Var}[\boldsymbol{\eta}_t] = \mathbf{W} \equiv \begin{bmatrix} \sigma_\chi^2 \Delta t & \rho_{\chi\xi} \sigma_\chi \sigma_\xi \Delta t \\ \rho_{\chi\xi} \sigma_\chi \sigma_\xi \Delta t & \sigma_\xi^2 \Delta t \end{bmatrix}.$$

With this process, the n -step ahead mean vector (\mathbf{m}_n) and covariance matrix (\mathbf{V}_n) are given recursively by $\mathbf{m}_n = \mathbf{c} + \mathbf{Q}\mathbf{m}_{n-1}$ and $\mathbf{V}_n = \mathbf{Q}\mathbf{V}_{n-1}\mathbf{Q}' + \mathbf{W}$, with $\mathbf{m}_0 = \mathbf{x}_0 \equiv [\chi_0, \xi_0]$ and $\mathbf{V}_0 = 0$ (see, for example, Harvey 1989, p. 109). Applying this recursion, we find

$$\mathbf{m}_n = [\phi^n \chi_0, \xi_0 + \mu_\xi n \Delta t],$$

$$\mathbf{V}_n = \begin{bmatrix} \sigma_\chi^2 \Delta t \sum_{i=0}^{n-1} \phi^{2i} & \rho_{\chi\xi} \sigma_\chi \sigma_\xi \Delta t \sum_{i=0}^{n-1} \phi^i \\ \rho_{\chi\xi} \sigma_\chi \sigma_\xi \Delta t \sum_{i=0}^{n-1} \phi^i & n \Delta t \sigma_\chi^2 \end{bmatrix}.$$

(A symbolic processor like Mathematica or Maple is useful for checking these recursive calculations.) We can rewrite the geometric series in \mathbf{m}_n and \mathbf{V}_n , using

$$\sum_{i=0}^{n-1} \phi^i = \frac{1 - \phi^n}{1 - \phi} \quad \text{and} \quad \sum_{i=0}^{n-1} \phi^{2i} = \frac{1 - \phi^{2n}}{1 - \phi^2}.$$

The errors in these discrete time approximation are of an order smaller than Δt (see Karlin and Taylor 1981, p. 160). Thus, if we take the limit as $n \rightarrow \infty$ and $\Delta t = t/n \rightarrow 0$, then $\phi^n = (1 - \kappa t/n)^n$ approaches $e^{-\kappa t}$, ϕ^{2n} approaches $e^{-2\kappa t}$, and

$$\frac{1 - \phi^{n-1}}{1 - \phi} \Delta t \rightarrow \frac{(1 - e^{-\kappa t})}{\kappa} \quad \text{and} \quad \frac{1 - \phi^{2(n-1)}}{1 - \phi^2} \Delta t \rightarrow \frac{(1 - e^{-2\kappa t})}{2\kappa}.$$

Substituting these limiting forms into the expressions for \mathbf{m}_n and \mathbf{V}_n , we arrive at the mean vector and covariance matrix given in Equation (3).

References

Bertsekas, D. P. 1995. *Dynamic Programming and Optimal Control*. Athena-Scientific, Belmont, MA.
 Boyle, P. B., J. Evnine, S. Gibbs. 1989. Numerical evaluation of multivariate contingent claims. *Rev. Financial Stud.* **2** 241–250.
 Brennan, M. J. 1991. The price of convenience and the valuation of commodity contingent claims. D. Lund, B. Oksendal, eds. *Stochastic Models and Option Values: Applications to Resources, Environment, and Investment Problems*. North-Holland, New York.
 —, E. S. Schwartz. 1985. Evaluating natural resource investments. *J. Bus.* **58** 135–157.

Cortazar, G., E. S. Schwartz. 1994. The evaluation of commodity contingent claims. *J. Derivatives* **1** 27–39.
 Cox, J. C., J. E. Ingersoll, S. A. Ross. 1985. An intertemporal general equilibrium model of asset prices. *Econometrica* **53** 363–384.
 Dixit, A. K., R. S. Pindyck. 1994. *Investment Under Uncertainty*. Princeton University Press, Princeton, NJ.
 Duffie, D. 1992. *Dynamic Asset Pricing Theory*. Princeton University Press, Princeton, NJ.
 Energy Information Administration. 1994. *Petroleum Marketing Annual 1993*. Office of Oil and Gas, U.S. Department of Energy, Washington, D.C.
 Energy Information Administration. 1995. *Petroleum Marketing Annual 1994*. Office of Oil and Gas, U.S. Department of Energy, Washington, D.C.
 Gibson, R., E. S. Schwartz. 1990. Stochastic convenience yield and the pricing of oil contingent claims. *J. Finance* **45** 959–976.
 Harvey, A. C. 1989. *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press, Cambridge, U.K.
 Karlin, S., H. M. Taylor. 1981. *A Second Course in Stochastic Processes*. Academic Press, New York.
 Laughton, D. G., Henry D. Jacoby. 1993. Reversion, timing options, and long-term decision making. *Financial Management* **33** 225–240.
 —, —. 1995. The effects of reversion on commodity projects of different length. L. Trigeorgis, ed. *Real Options in Capital Investment: Models, Strategies and Applications*. Praeger, Westport, CT. 185–205.
 Merton, R. C. 1973. An intertemporal capital asset pricing model. *Econometrica* **41** 867–887.
 Paddock, J. L., D. R. Siegel, J. L. Smith. 1988. Option valuation of claims on real assets: The case of offshore petroleum leases. *Quart. J. Econom.* **103** 479–503.
 Pilopovic, D. 1998. *Energy Risk*. McGraw-Hill, New York.
 Pindyck, R. S. 1993. Inventories and the short-run dynamics of commodity prices. *Rand J. Econom.* **25** 141–159.
 —. 1997. The long-run dynamics of commodity prices. Working Paper, Massachusetts Institute of Technology, Cambridge, MA.
 Ross, S. 1997. Hedging long run commitments: Exercises in incomplete market pricing. *Banca Monte Econom. Notes* **26** 99–132.
 Routledge, B. R., D. J. Seppi, C. S. Spatt. 1999. Equilibrium forward curves for commodities. *J. Finance* **55** 1297–1338.
 Schwartz, E. S. 1995. Review of *Investment Under Uncertainty*. A. K. Dixit, R. S. Pindyck. *J. Finance* **50** 1924–1928.
 —. 1997. The stochastic behavior of commodity prices: Implications for valuation and hedging. *J. Finance* **52** 923–973.
 —. 1998. Valuing long term commodity assets. *Financial Management* **27** 57–66.
 Smith, J. E., K. F. McCardle. 1998. Valuing oil properties: Integrating option pricing and decision analysis approaches. *Oper. Res.* **46** 198–217.
 —, —. 1999. Options in the real world: Lessons learned in evaluating oil and gas investments. *Oper. Res.* **47** 1–15.
 West, M., J. Harrison. 1996. *Bayesian Forecasting and Dynamic Models*, 2nd ed. Springer-Verlag, New York.

Accepted by Phelim P. Boyle; received in April, 1999. This paper was with the authors 6 months for 2 revisions.