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

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

and

James E. Smith*
Fuqua School of Business
Duke University
Durham, NC 27708-0120

First Version: June 13, 1997
This Version: May 26, 1998

*Jim Smith's work on this project was supported in part by the National Science Foundation under grant no SBR-9511364. We are grateful to Enron Capital and Trade Resources for making the long-term forward price data available.

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

ABSTRACT

In this paper, 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. Though these two factors are not directly observable, they may be estimated from futures prices. Intuitively, movements in prices for long-maturity futures contracts provides information about the equilibrium price level and differences between the prices for the short- and long-term contracts provides information about short-term variations in prices. We estimate this model using prices for oil futures contracts and show that, even though 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 also develop and estimate a three-factor extension of this model that considers uncertainty in equilibrium growth rates and provides an improved fit to long-term futures prices. This three-factor model is new contribution to literature on commodity price modeling and illustrates the power of the basic framework developed in the paper.

1. Introduction

The stochastic behavior of commodity prices plays a central role in models for evaluating commodity-related securities and projects. Recently, a number of authors have considered the impact of mean reversion in commodity prices on the values and optimal policies for these securities and projects.[1] To illustrate the phenomenon of mean reversion, suppose a prolonged winter cold spell hits the East Coast of the United States. This cold spell increases the demand for oil and natural gas for heating and depletes oil and gas inventories in that area. This leads to an increase in spot prices and this increase in spot prices, in turn, leads to an increase in supply as producers transport oil and gas to the East Coast markets that would normally have been sent elsewhere. As the weather returns to normal and the extra supplies arrive, inventories are restored and prices return to their normal levels.

In general, in an equilibrium setting, we would expect that when prices are higher than the equilibrium level, the supply of the commodity will increase since higher cost producers of the commodity will enter the market, thereby putting downward pressure on prices. Conversely, when prices are relatively low, supply will decrease since some of the higher cost producers will exit, thereby putting upward pressure on prices. When these entries and exits take time, the adjustment is not instantaneous and prices may be temporarily high or low, but will tend to revert to the equilibrium level. These equilibrium prices may also change over time, in response changes in the supply of or demand for the commodity.

[1] See, for example, Brennan (1991), Gibson and Schwartz (1990), Laughton and Jacoby (1993, 1995), Cortazar and Schwartz (1994), Dixit and Pindyck (1994), Schwartz (1997, 1998), Smith and McCardle (1997).

In this paper, we develop a new two-factor model of commodity prices that allows mean-reversion in prices and uncertainty in the equilibrium level to which prices revert.[2] In this 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, and inflation. Changes in this relationship -- perhaps due to changes in production or exploration technology, a change in political structure or policies, or major new discoveries -- may result in unanticipated and persistent shifts in the equilibrium price level. The second factor in the model is the difference between spot prices and the equilibrium price level and can be interpreted as representing short-term deviations from the equilibrium price level. These short-term deviations might be caused by, for example, the weather (as in the earlier example) affecting demand or a temporary supply disruption. The short-term deviations are not expected to persist and are expected to revert towards zero following an Ornstein-Uhlenbeck process.

Though neither of these factors is directly observable, the two factors may be estimated using information from the prices of futures contracts. Intuitively, changes in the long-maturity futures prices gives information about changes in the equilibrium price and changes in the difference between near- and long-term futures prices gives information about the short-term deviations. The short-term/long-term model of commodity prices developed in this paper is particularly convenient because we can write closed-form expressions for futures prices and the logarithm of these futures prices turns out to be a linear function of the model factors. This allows us to use standard Kalman filtering techniques to estimate these factors over time and use maximum likelihood methods to estimate the parameters of the stochastic processes governing their behavior. In addition, we can derive analytic expressions for valuing European options written on these futures contracts.

---

[2] Our model is descriptive in nature in that we do not formally model the equilibrium process leading to mean-reversion in prices. For equilibrium models leading to mean-reversion, see Zhou (1998), Lund (1993), and

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, this short-term/long-term model turns out to be 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. In particular, the instantaneous convenience yield, modeled as a mean-reverting stochastic process in the Gibson-Schwartz model, can be represented as a linear function of the short-term deviation in the short-term/long-term model. In the short-term/long-term model, the difference between the true price process and "risk-adjusted" price process used to value futures contracts is captured by constant adjustments to the drift terms associated with the short- and long-term factors. The short-term/long-term model thus provides an alternative and, in many respects, simpler interpretation of the results of the stochastic convenience yield model where changes in the slope of the futures curve are interpreted as short-term variations in prices rather than changes in the instantaneous convenience yield.

We begin by formally defining the short-term/long-term model in Section 2 and deriving the distributions for future spot prices. In section 3, we describe the "risk-neutralized" version of the model and use it to derive closed-form expressions for futures prices and for European options on these futures. In Section 4, we establish the relationship between this model and the Gibson-Schwartz stochastic convenience yield model and, in section 5, we briefly describe the Kalman filtering procedure used to estimate the model. In section 6, we estimate the model using two data sets. The first data set consists of prices for publicly traded oil futures contracts from 1990 to 1995 and includes contracts with maturities up to 17 months. The second data set consists of proprietary forward price data made available by Enron

---

Routledge, Seppi and Spatt (1998). Pindyck (1980) and Litzenberger and Rabinowitz (1995) also develop

Capital and Trade Resources and includes contracts from 1993 to 1996 with maturities of up to 9 years. In section 7, we describe a three-factor extension of the basic model that considers uncertainty in the growth rate for equilibrium prices and provides an improved the fit to long-term futures prices. This three-factor model is new contribution to literature on commodity price modeling and illustrates the power of the basic framework developed in the paper. We offer some concluding remarks in section 8.

---

equilibrium models of oil price determination.

2. The Short-Term/Long-Term Model

Let S_t denote the spot price of a commodity at time t and let $X_t = \ln(S_t)$. We will decompose X_t into two factors as

$$X_t = \chi_t + \xi_t \tag{1}$$

where χ_t represents the time-t short-term deviation in log prices and ξ_t represents the time-t equilibrium level for log prices. (Seasonality can be incorporated by including time-dependent constants in this equation.) Both χ_t and ξ_t are unobservable state variables that can be estimated by a Kalman filtering process as described in Section 5 below. The short-run deviations (χ_t) are assumed to revert towards zero following an Ornstein-Uhlenbeck process

$$d\chi_t = -\kappa\,\chi_t\,dt + \sigma_\chi\,dz_\chi\ , \tag{2}$$

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

$$d\xi_t = \mu_\xi\ dt + \sigma_\xi\,dz_\xi\ . \tag{3}$$

Here dz_χ and dz_ξ are correlated increments of standard Brownian motion process 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 (due to, for example, unusual weather or a supply disruption) that are not expected to persist. Changes in the equilibrium level (ξ_t) represents more fundamental market factors with 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 in that any deviation χ_t is expected to halved in this time period.

Combining equations (1) through (3), we can write a stochastic differential equation for log prices as

$$dX_t = -\kappa(X_t - \xi_t')\,dt + \sigma_X\,dz_X\ , \tag{4}$$

with $\xi_t' = \xi_t + \mu_\xi/\kappa$, $\sigma_X^2 = \sigma_\chi^2 + \sigma_\xi^2 + 2\,\rho_{\chi\xi}\sigma_\chi\sigma_\xi$, and $dz_X = (1/\sigma_X)\,(\sigma_\chi\, dz_\chi + \sigma_\xi\, dz_\xi)$. Thus we can interpret the model as having mean-reverting spot prices with a stochastic "mean." This short-term/long-term model includes the standard geometric Brownian motion model of spot prices as a special case where $\sigma_\chi = 0$ and $\chi_0 = 0$; in this case, the only uncertainty is the equilibrium level (ξ_t) and spot prices are equal to equilibrium prices. Also included is the simple model of mean reversion referred to as Model I in Schwartz (1997); this model has no uncertainty or growth in the equilibrium level and takes $\sigma_\xi = 0$ and $\mu_\xi = 0$.

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 appendix, we find that χ_t and ξ_t are jointly normally distributed with mean vector and covariance matrix:

$$\mathrm{E}[(\chi_t, \xi_t)] = [e^{-\kappa t}\chi_0\,,\ \xi_0 + \mu_\xi t\,] \qquad \text{and} \tag{5a}$$

$$\mathrm{Cov}[(\chi_t, \xi_t)] = \begin{bmatrix} (1-e^{-2\kappa t})\dfrac{\sigma_\chi^2}{2\kappa} & (1-e^{-\kappa t})\dfrac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} \\ (1-e^{-\kappa t})\dfrac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} & \sigma_\xi^2 t \end{bmatrix} \tag{5b}$$

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

$$\mathrm{E}[X_t] = e^{-\kappa t}\chi_0 + \xi_0 + \mu_\xi t \qquad \text{and} \tag{6a}$$

$$\mathrm{Var}[X_t] = (1-e^{-2\kappa t})\frac{\sigma_\chi^2}{2\kappa} + \sigma_\xi^2 t + 2\,(1-e^{-\kappa t})\frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} \tag{6b}$$

The spot price (S_t) is then log-normally distributed with expected price given by

$$\mathrm{E}[S_t] = \exp\left(\mathrm{E}[X_t] + \tfrac{1}{2}\mathrm{Var}[X_t]\right)$$

or

$$\ln(\mathrm{E}[S_t]) = \mathrm{E}[X_t] + \frac{1}{2}\mathrm{Var}[X_t] \tag{7}$$

$$= e^{-\kappa t}\chi_0 + \xi_0 + \mu_\xi t + \frac{1}{2}\left((1-e^{-2\kappa t})\frac{\sigma_\chi^2}{2\kappa} + \sigma_\xi^2 t + 2\,(1-e^{-\kappa t})\frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa}\right)$$

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

$$\left(\xi_0 + \frac{\sigma_\chi^2}{4\kappa} + \frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa}\right) + \left(\mu_\xi + \frac{1}{2}\sigma_\xi^2\right) t \ . \tag{8}$$

Thus, the expected long-run price is determined *as if* the current spot price were $\exp(\xi_0 + \sigma_\chi^2/4\kappa + \rho_{\chi\xi}\sigma_\chi\sigma_\xi/\kappa)$ and the spot price is expected to grow at $(\mu_\xi + \frac{1}{2}\sigma_\xi^2)$. Thus we can interpret $\exp(\xi_t + \sigma_\chi^2/4\kappa + \rho_{\chi\xi}\sigma_\chi\sigma_\xi/\kappa)$ as a "long-run mean price" summarizing time-t expectations about long-run prices.[3]

To illustrate the implications of the short-term/long-term model, Figure 1 shows probabilistic forecasts generated using the model. We use parameter estimates based on the Enron forward data (shown in Table II below and modified as described in footnote 6) and show forecasts generated on May 16, 1996. On this date, the state variables are estimated to be $\chi_0 = .119$ and $\xi_0 = 2.857$, corresponding to a current spot price of \$19.61 ($=\exp(\chi_0 + \xi_0)$) and an equilibrium price level of \$17.41 ($=\exp(\xi_0)$). The solid lines in Figure 1 represent forecasts for spot prices and the dashed lines represent forecasts for the equilibrium price level. The center lines represent the expected value forecasts for each variable and the upper and lower lines represent "confidence bands" for each variable such that there is a 10 percent or 90 percent chance (respectively) that the variable will be below that level on that particular date. The forecasts for the spot price show that prices are expected to drop towards the long-run mean (given by equation 8), with the current deviation expected to be halved in 7 months ($=-\ln(0.5)/\kappa$). The equilibrium price level is expected to increase 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, in the longer term, most of the uncertainty in spot prices is due to

[3] The terms involving σ_χ and σ_ξ appearing in (8) reflect the contribution of uncertainty in the short-term deviations and equilibrium level to the expectations in (7); these appear because we are taking logarithms of the expected price rather than expectations of log price. The expected log spot price approaches $\xi_0 + \mu_\xi t$ as $t \to \infty$ and the median spot price approaches $\exp(\xi_0 + \mu_\xi t)$.

uncertainty about the then-prevailing equilibrium price level.

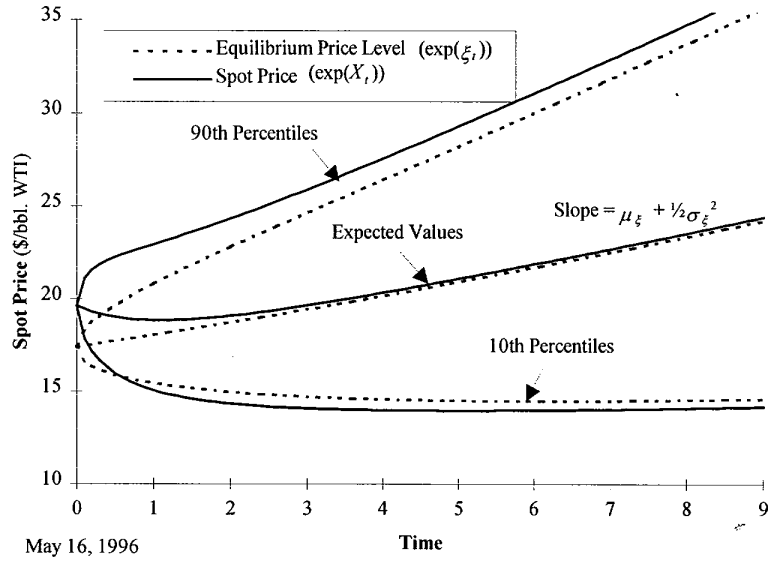

Figure 1: Probabilistic Forecasts for Spot and Equilibrium Prices

3. Risk-Neutral Processes

We now develop the "risk-neutral" (or "risk-adjusted") version of the model that will be used to value futures contracts and options on these futures. Let λ_χ and λ_ξ denote risk premiums for χ and ξ respectively. Using asterisks to denote the risk-neutralized versions of their "true" counterparts, we assume that these risk-neutral stochastic processes are of the form:

$$d\chi_t^* = (-\kappa\chi_t^* - \lambda_\chi)\, dt + \sigma_\chi dz_\chi^* \tag{9}$$

$$d\xi_t^* = (\mu_\xi - \lambda_\xi)\, dt + \sigma_\xi dz_\xi^* \tag{10}$$

where, again, dz_χ^* and dz_ξ^* are increments of standard Brownian motion process with $dz_\chi^*\, dz_\xi^* = \rho_{\chi\xi}\, dt$. Thus, the risk premiums are assumed to be constant reductions in the drifts of each process. With this assumption, the risk-neutral short-term process (χ_t^*) follows an Ornstein-Uhlenbeck process reverting to

$-\lambda_\chi/\kappa$ (rather than 0 as assumed in the true process) and the risk-neutralized long-term process has drift $\mu_\xi^* \equiv \mu_\xi - \lambda_\xi$ (rather than μ_ξ as assumed in the true process).

Given $\chi_0^* = \chi_0$ and $\xi_0^* = \xi_0$, following the same derivation as in equation (5), we find χ_t^* and ξ_t^* are jointly normally distributed with mean vector and covariance matrix:

$$\mathrm{E}[(\chi_t^*, \xi_t^*)] = [e^{-\kappa t}\chi_0 - (1 - e^{-\kappa t})\lambda_\chi/\kappa,\ \xi_0 + \mu_\xi^* t] \tag{11a}$$

$$\mathrm{Cov}[(\chi_t^*, \xi_t^*)] = \mathrm{Cov}[(\chi_t, \xi_t)] \quad . \tag{11b}$$

Under this risk-neutral distribution, the log of the future spot price (X_t^*) is normally distributed with:

$$\mathrm{E}[X_t^*] = e^{-\kappa t}\chi_0 + \xi_0 - (1 - e^{-\kappa t})\lambda_\chi/\kappa + \mu_\xi^* t \tag{12a}$$

$$\mathrm{Var}[X_t^*] = \mathrm{Var}[X_t] \quad . \tag{12b}$$

Comparing equations (6) and (12), we see that the net effect of the risk-premiums in the model is a reduction in the log of the expected spot prices of $(1 - e^{-\kappa t})\lambda_\chi/\kappa + \lambda_\chi t$; this premium depends on time but is independent of the value of the state variables.[4]

2.1 Valuing Futures Contracts

Let $F_{T,0}$ denote the current market price for a futures contract with time T until maturity. Using the theory of risk-neutral valuation, 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., Cox, Ingersoll and Ross 1981 or Duffie 1992). Since future spot prices are log-normally distributed, we have

$$\ln(F_{T,0}) = \ln(\mathrm{E}[S_t^*]) \tag{13}$$

$$= \mathrm{E}[X_T^*] + \frac{1}{2}\mathrm{Var}[X_T^*]$$

$$= e^{-\kappa T}\chi_0 + \xi_0 + A(T)$$

where:

$$A(T) = \mu_\xi^* T - (1 - e^{-\kappa T})\frac{\lambda_\chi}{\kappa} + \frac{1}{2}\left((1 - e^{-2\kappa T})\frac{\sigma_\chi^2}{2\kappa} + \sigma_\xi^2 T + 2(1 - e^{-\kappa T})\frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa}\right)$$

As the maturity of the contract increases (i.e., as $T \to \infty$), the short-term deviations are expected to have reverted, and from equation (13) we see that the log of the value of the futures contract approaches

$$\left(\xi_0 - \frac{\lambda_\chi}{\kappa} + \frac{\sigma_\chi^2}{4\kappa} + \frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa}\right) + \left(\mu_\xi^* + \frac{1}{2}\sigma_\xi^2\right)T \quad . \tag{14}$$

Thus, according to this model, prices for the long-maturity futures contracts are determined *as if* the current spot price were $\exp(\xi_0 - \lambda_\chi/\kappa + \sigma_\chi^2/4\kappa + \rho_{\chi\xi}\sigma_\chi\sigma_\xi/\kappa)$ and the futures prices grow at $(\mu_\xi^* + \frac{1}{2}\sigma_\xi^2)$; this "as if" current spot price is referred to as the "shadow spot price" in Schwartz (1998).

Comparing equation (14) to the long-run expectations for the true process (equation 8 above), we see that the risk premium for the short-term deviations subtracts a constant amount (λ_χ/κ) from the effective "long-run mean" price and the risk-premium (λ_ξ) for the equilibrium level reduces the growth in the long-term futures curve below the long-term expected growth in spot prices (since $\mu_\xi^* = \mu_\xi - \lambda_\xi$). This 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 model futures prices and forecasts from May 16, 1996. We also show actual futures prices for that day and see that the model fits the observations reasonably well; we will discuss the fit in more detail in Section 6 below.

---

[4] 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 (9) with α in place of λ_χ and $\kappa^* = \kappa + \beta$ in place of κ. Thus, with this extension, we would use a risk-adjusted 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 8 below.

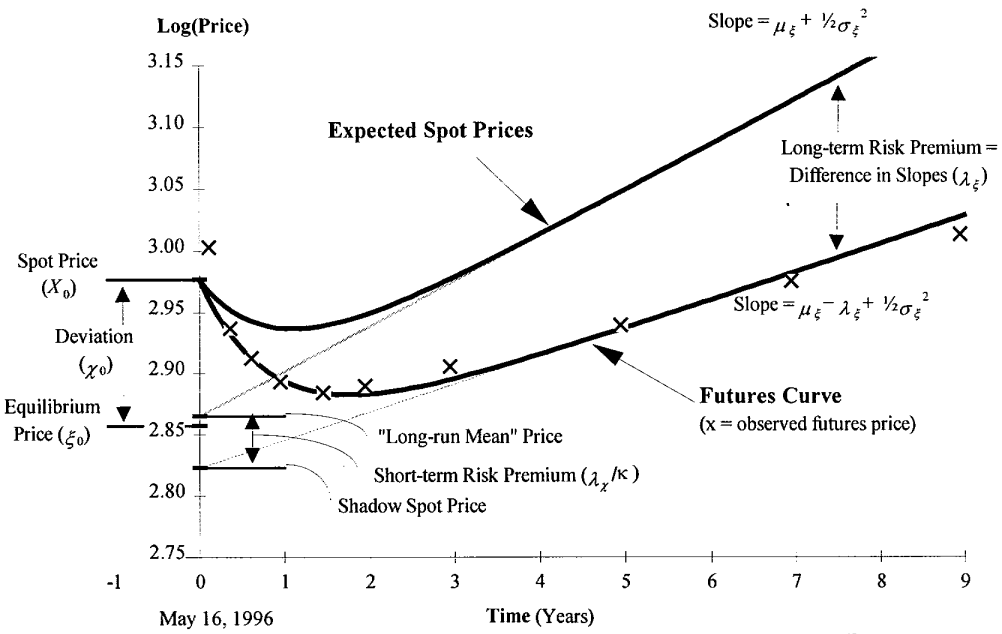

Figure 2: Futures Prices and Expected Spot Prices

In this model, the instantaneous volatility of the futures price ($\sigma(F_{T,0})$) is independent of the state variables and, from equation (13), is given as

$$\sigma^2(F_{T,0}) = e^{-2\kappa T}\ \sigma_\chi^2 + \sigma_\xi^2 + 2\, e^{-\kappa T}\, \rho_{\chi\xi}\, \sigma_\chi\, \sigma_\xi \ . \tag{15}$$

Thus, the volatility in prices for near maturity futures contracts (i.e., $T = 0$) 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 and less of a contribution to the volatility of the futures prices. In the limit, as $T \to \infty$, the instantaneous volatility approaches the volatility of the equilibrium price level (σ_ξ) as prices for futures contracts with long maturities are affected only by changes in the equilibrium level. This volatility relationship is illustrated in Figure 3 (again using the parameters from the Enron data) along with the empirical volatilities calculated from the Enron data. Here we can see that the model fits the empirical volatilities reasonably well. The "option volatilities" for the model are also shown and will be discussed shortly.

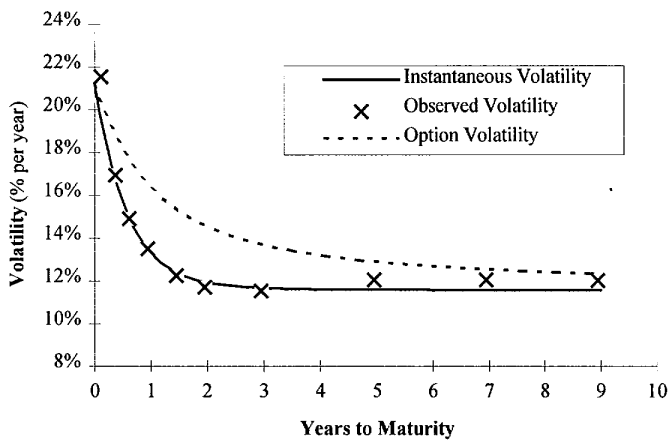

Figure 3: Model and empirical volatilities

2.2 Valuing European Options on Futures Contracts

We can also derive analytic forms for the value of European options on futures contracts in this model. We begin by deriving the risk-neutral distribution for future futures prices, as the value of European options on futures will be given by taking expectations with respect to this distribution. Let $F^*_{T,t}$ denote the price at time t of a futures contract expiring at time T ($T > t$); again the asterisk indicates that we will be working with the risk-neutral distribution for futures prices rather than the true distribution. From the formula for current futures prices (equation 13 above), we have:

$$\ln(F^*_{T,t}) = e^{-\kappa(T-t)}\chi^*_t + \xi^*_t + A(T-t) \ . \tag{16}$$

Since χ^*_t and ξ^*_t are jointly normally distributed, $\phi \equiv \ln(F^*_{T,t})$ is normally distributed with

$$\begin{aligned}\mu_\phi(t,T) \equiv \mathrm{E}[\ln(F^*_{T,t})] &= e^{-\kappa(T-t)}\,\mathrm{E}[\chi^*_t] + \mathrm{E}[\xi^*_t] + A(T-t) \qquad (17a)\\ &= e^{-\kappa(T-t)}\,(e^{-\kappa t}\chi_0 - (1-e^{-\kappa t})\,\lambda_\chi/\kappa) + \xi_0 + \mu^*_\xi t + A(T-t)\end{aligned}$$

$$\sigma^2_\phi(t,T) \equiv \mathrm{Var}[\ln(F^*_{T,t})] = e^{-2\kappa(T-t)}\mathrm{Var}[\chi^*_t] + \mathrm{Var}[\xi^*_t] + 2\,e^{-\kappa(T-t)}\,\mathrm{cov}(\chi^*_t, \xi^*_t) \tag{17b}$$

$$= e^{-2\kappa(T-t)}\,(1-e^{-2\kappa t})\,\frac{\sigma_\chi^2}{2\kappa} \;+\; \sigma_\xi^2 t \;+ 2\,e^{-\kappa(T-t)}\,(1-e^{-\kappa t})$$

$$\frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa}\;.$$

As a log-normal random variable, the expected futures price at time t (under the risk-neutral process) is given as

$$\mathrm{E}[F^*_{T,t}] = \exp\left(\mu_\phi(t,T) + \frac{1}{2}\,\sigma_\phi^2(t,T)\right)\;. \tag{18}$$

After some algebra, we find that this expected futures price at time t is equal to the current futures price (i.e., $\mathrm{E}[F^*_{T,t}] = F_{T,0}$ where $F_{T,0}$ is given by equation 13). Thus the futures price is a martingale under the risk-neutral measure, as required in the theory of risk-neutral valuation of futures contracts (see, e.g., Duffie 1992, pg. 122).

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. In the risk-neutral valuation framework, the value of a European option is given by determining its expected present value at expiration where expectations are calculated using the risk-neutral processes and discounting is done at the risk-free rate (r). 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 given as

$$e^{-rt}\;\mathrm{E}[\max(F^*_{T,t} - K,\,0)] = \;e^{-rt}\;\left(F_{T,0}\,\mathrm{N}(d) - K\,\mathrm{N}(d - \sigma_\phi(t,T))\right) \tag{19}$$

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

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

$$e^{-rt}\;\mathrm{E}[\max(K - F^*_{T,t},\,0)] = \;e^{-rt}\;\left(-F_{T,0}\,\mathrm{N}(d) + K\,\mathrm{N}(d - \sigma_\phi(t,T))\right). \tag{20}$$

These option valuation formulas are analogous to the Black-Scholes formula for options on common stocks. Here the "stock price" for the Black-Scholes formula corresponds to the present value of the futures commitment ($e^{-rt} F_{T,0}$) and the equivalent annualized volatility would be $\sigma_\phi(t,T)/\sqrt{t}$. These annualized "option volatilities" are shown in Figure 3; for this plot, we assume that the option expires at maturity of the futures contract (i.e., $t = T$). Here 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 and/or the time until option expiration increases (i.e. , as $t \to \infty$ and/or $T \to \infty$, assuming $T > t$), the annualized option volatility approaches the volatility of the equilibrium level (σ_ξ) as most of the uncertainty in value for a long-term option is due to uncertainty in the equilibrium level at the time the option expires.

4. Relationship to the Gibson-Schwartz Stochastic Convenience Yield 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 turns out to be 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 stochastic convenience yield 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 (as in Section 1 above). The stochastic convenience yield model assumes that:

$$dX_t = (\mu - \delta_t - \tfrac{1}{2}\sigma_1^2)\, dt + \sigma_1\, dz_1 \tag{21}$$

$$d\delta_t = \kappa(\alpha - \delta_t)\, dt + \sigma_2\, dz_2 \tag{22}$$

where dz_1 and dz_2 are correlated increments of standard Brownian motion process with $dz_1\, dz_2 = \rho\, dt$. Thus μ and σ_1 represent drift and volatility terms for the spot price diffusion process. The convenience yield δ_t follows an Ornstein-Uhlenbeck process with equilibrium level α and rate of mean reversion κ. In equation (21), the convenience yield also affects the drift of the spot price process. There is no overlap in notations for the parameters in two models except for the mean-reversion parameter (κ). As we will see shortly, this parameter is the same 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:

$$\chi_t = \text{short-term deviation} = \frac{1}{\kappa}(\delta_t - \alpha) \tag{23}$$

$$\xi_t = \text{equilibrium price level} = X_t - \chi_t = X_t - \frac{1}{\kappa}(\delta_t - \alpha) \tag{24}$$

To establish the equivalence of the two models, we can write the stochastic process equations (equations 2 and 3 above) for the short and long-term model using the stochastic process assumptions for the stochastic convenience yield model (equations 21 and 22) and relate the parameters in the two models:

$$\begin{aligned} d\chi_t &= \frac{1}{\kappa} d\delta_t \\ &= (\alpha - \delta_t)\, dt + \frac{\sigma_2}{\kappa} dz_2 \\ &= -\kappa \chi_t\, dt + \frac{\sigma_2}{\kappa} dz_2 \\ &= -\kappa \chi_t\, dt + \sigma_\chi dz_\chi\ , \end{aligned}$$

and

$$\begin{aligned} d\xi_t &= dX_t - \frac{1}{\kappa} d\delta_t \\ &= (\mu - \delta_t - \tfrac{1}{2}\sigma_1^2)dt + \sigma_1\, dz_1 - (\alpha - \delta_t)\, dt - \frac{\sigma_2}{\kappa} dz_2 \\ &= (\mu - \alpha - \tfrac{1}{2}\sigma_1^2)dt + \sigma_1\, dz_1 - \frac{\sigma_2}{\kappa} dz_2 \\ &= \mu_\xi\, dt + \sigma_\xi dz_\xi \end{aligned}$$

Thus the two models have equivalent true (as opposed to risk-neutral) processes if we define the parameters of the short-term/long-term model as in Table I.

Short-Term/Long-Term Model Parameter		Definition in Terms of
Symbol	Description	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	Long-term drift rate	$(\mu - \alpha - \frac{1}{2}\sigma_1^2)$
σ_ξ	Long-term volatility	$(\sigma_1^2 + \sigma_2^2/\kappa^2 - 2\rho\sigma_1\sigma_2/\kappa)^{1/2}$
dz_ξ	Long-term 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	Long-term risk premium	$\mu - r - \lambda/\kappa$

Table I: The relationships between parameters in the short-term/long-term model and the stochastic convenience model of Gibson and Schwartz (1990)

The risk parameters (λ_χ and λ_ξ) are determined by equating parameters in the risk-neutralized versions of the two models. In the Gibson-Schwartz stochastic convenience yield model, the risk-neutral process is assumed to be of the form:

$$dX_t^* = (r - \delta_t - \tfrac{1}{2}\sigma_1^2)\, dt + \sigma_1\, dz_1^* \tag{25}$$

$$d\delta_t^* = [\kappa(\alpha - \delta_t^*) - \lambda]dt + \sigma_2\, dz_2^* \tag{26}$$

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. We can then write the risk-neutralized version of the short-term/long-term model in terms of the stochastic convenience yield model as follows:

$$\begin{aligned} d\chi_t^* &= \frac{1}{\kappa} d\delta_t^* \\ &= (\alpha - \delta_t^* - \frac{\lambda}{\kappa})\, dt + \frac{\sigma_2}{\kappa} dz_2^* \\ &= (-\kappa\chi_t^* - \frac{\lambda}{\kappa})\, dt + \frac{\sigma_2}{\kappa} dz_2^* \\ &= (-\kappa\chi_t^* - \lambda_\chi)\, dt + \sigma_\chi\, dz_\chi^* \end{aligned}$$

and,

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

$$= (r - \alpha - \tfrac{1}{2}\sigma_1^2 + \frac{\lambda}{\kappa})\, dt \;\; + \sigma_1\, dz_1^* \; - \frac{\sigma_2}{\kappa}\, dz_2^*$$
$$= \quad (\mu_\xi - \lambda_\xi)\; dt \qquad\qquad + \sigma_\xi\, dz_\xi^*$$

From this, we can see that the values for λ_χ and λ_ξ shown in Table I make the two models equivalent. The risk premium for the convenience yield (λ) in the Gibson-Schwartz model thus translates directly into the risk-premium for the short-term variations ($\lambda_\chi = \lambda/\kappa$). The equilibrium convenience yield (α) in the Gibson-Schwartz model affects the drift of the long-term mean (μ_ξ) and the levels of the short-term deviation (χ_t) and long-term mean (ξ_t) (through equations 23 and 24), but does not affect the risk premiums associated with the equilibrium price level (λ_ξ).

Given the equivalence between the two models, we can substitute terms from Table I into the equations derived earlier for the short-term/long-term model and state the corresponding result in terms of the stochastic convenience yield model. For example, substituting terms into the equation for futures prices in the short-term/long-term model (equation 13), we obtain the corresponding equation for futures prices in the stochastic convenience yield model (equation 19 in Schwartz 1997). Similarly, by substituting the terms into the equations for options prices in the short-term/long-term model (equations 19 and 20), we get the corresponding equations for the Gibson-Schwartz model (equations 23 and 24 in Schwartz 1998; see also Shimko 1994).

5. Estimation Procedures

As indicated in the introduction, the state variables in the short-term/long-term model are not directly observable and must be estimated using information contained in futures prices and/or contemporaneous price forecasts. Intuitively, the long-maturity futures contracts give information about the current estimate of the long-term mean or equilibrium level (ξ_t) and the difference between the long- and short-term contracts gives us information about the current short-term deviation (χ_t). The Kalman filter is a recursive procedure for computing estimates of the state variables given observations that

depend on these state variables. Standard Kalman filter techniques require that (a) the state variables evolve according to a linear model with jointly normally distributed disturbances or "innovations" and (b) the observations be a linear function of the state variables with jointly normally distributed measurement errors. If we take the observations to be the logarithms of futures prices and assume that errors in these observations are jointly normally distributed, the short-term/long-term model satisfies the conditions required to apply standard Kalman filtering techniques. The estimated state variables at time t can then be interpreted as expectations conditional on all of the information available up to and including time t. Detailed accounts of the Kalman filter are given in Harvey (1989; Chapter 3) and West and Harrison (1996).

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 log of the prices of futures contracts with varying maturities. Casting this relationship in terms of the Kalman filter, the evolution of the state variables is described by the *transition equation* which from equation (5) can be written as:

$$\boldsymbol{x}_t = \boldsymbol{c} + \boldsymbol{Q}\,\boldsymbol{x}_{t-1} + \eta_t \quad , \quad t = 1, \ldots, n_T \tag{27}$$

where: $\boldsymbol{x}_t \equiv [\chi_t, \xi_t]$, a 2×1 vector of state variables

$\boldsymbol{c} \equiv [0, \mu_\xi \Delta t]$, a 2×1 vector

$\boldsymbol{Q} \equiv \begin{bmatrix} e^{-\kappa\Delta t} & 0 \\ 0 & 1 \end{bmatrix}$, a 2×2 matrix

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

$$\mathrm{Var}[\eta_t] = \begin{bmatrix} (1 - e^{-2\kappa\Delta t})\dfrac{\sigma_\chi^2}{2\kappa} & (1 - e^{-\kappa\Delta t})\dfrac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} \\ (1 - e^{-\kappa\Delta t})\dfrac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa} & \sigma_\xi^2\Delta t \end{bmatrix}$$

$\Delta t \equiv$ the length of the time steps.

The *measurement equation* describes the relationship between the state variables and the observed prices. From equation (13), this is:

$$y_t = \boldsymbol{d}_t + \boldsymbol{Z}_t \boldsymbol{x}_t + \varepsilon_t \ , \ t = 1, ..., n_T \tag{28}$$

where:

$\boldsymbol{y}_t \equiv [\ln F_{Ti}]$, a $n \times 1$ vector of observables

$\boldsymbol{d}_t \equiv [A(T_i)]$, a $n \times 1$ vector

$\boldsymbol{Z}_t \equiv [e^{-\kappa Ti}, 1]$, a $n \times 2$ matrix

ε_t , a $n \times 1$ vector of serially uncorrelated, normally distributed disturbances with:

$$E[\varepsilon_t] = \mathbf{0}, \ \mathrm{Cov}[\varepsilon_t] = \boldsymbol{H}$$

$n_T \equiv$ the number of time periods in the data set .

Given these equations and a set of observed futures prices ($\boldsymbol{y}_t \equiv [\ln F_{ti}]$, $t = 1, ..., n_T$), we "run" the Kalman filter recursively: we start with initial estimates of the state variables ($\boldsymbol{x}_0 = [\chi_0, \xi_0]$) and, in each period ($t = 1, ..., n_T$), we use the observation $\boldsymbol{y}_t$ and the previous period's estimates of the state variables ($\boldsymbol{x}_{t-1} = [\chi_{t-1}, \xi_{t-1}]$)to generate estimates of the current state variables ($\boldsymbol{x}_t = [\chi_t, \xi_t]$). Using the notation of equations (27 and 28), the estimated state variables are given by:

$$\boldsymbol{x}_t = \boldsymbol{c} + \boldsymbol{Q}\,\boldsymbol{x}_{t-1} + \boldsymbol{K}_t\left(\boldsymbol{y}_t - \boldsymbol{d}_t - \boldsymbol{Z}_t\,(\boldsymbol{c} + \boldsymbol{Q}\,\boldsymbol{x}_{t-1})\right)$$

The first part of the equation ($\boldsymbol{c} + \boldsymbol{Q}\,\boldsymbol{x}_{t-1}$) can be interpreted as the one-period ahead predicted value of the state variable, given the previous period's state variable estimate. The second term is a correction to this prediction based on the difference between the observed price vector ($\boldsymbol{y}_t$) and the predicted price vector ($\boldsymbol{d}_t + \boldsymbol{Z}_t\,(\boldsymbol{c} + \boldsymbol{Q}\,\boldsymbol{x}_{t-1})$. In addition, these state variable estimates are accompanied by a covariance matrix that describes the accuracy of the estimates. (See, e.g., Harvey 1989 or West and Harrison 1996 for details.)[5]

---

[5] Following Schwartz (1997), we start the Kalman filter assuming initial state variable values of $(\chi_0, \xi_0) = (0, \bar{x})$ where $\bar{x}$ is the average near-term futures price in the data set. The estimated parameters and state variables do not appear to be sensitive to the assumed initial values. The initial covariance matrix is similarly based on observed standard deviations and the results also appear to be insensitive to these assumptions.

The Kalman filtering procedure thus allows us to estimate the state variables over time given particular assumptions about the parameters of the process. We can also calculate the likelihood of the observations given a particular set of parameters. 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; see Harvey (1989; Chapter 3.4) for details. With the short-term/long-term model, there are 7 model parameters to be estimated (κ, σ_χ, μ_ξ, σ_ξ, $\rho_{\xi\chi}$, λ_χ, μ_ξ^*) plus the terms in the covariance matrix for the measurement errors (H). 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 H is diagonal with diagonal elements (s_1^2, ..., s_n^2). We used the "maxlik" routine in Gauss to numerically determine parameter estimates and standard errors for these estimates.

6. Empirical Results

We estimate our model using two different data sets; both data sets 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 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.

Table II shows parameter estimates for each data set and Figures 4(a) and (b) shows the estimated values of the equilibrium price (given as $\exp(\xi_t)$) and spot price ($\exp(\xi_t + \chi_t)$) for the two data sets. 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 short-term 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 4a) when spot prices rose above $40 per barrel 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 where "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, pg. 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, pg. Xiii).

Parameter	Description	Futures Data Estimate	Standard Error	Enron Data 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	.300	.044	.189	.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
NT	Number of time periods (weeks)	259		163	
N	Number of futures contracts	5		10	
	Log Likelihood	5140.0		6182.5	
	Time Frame	1/2/90 to 2/17/95		1/15/93 to 5/16/96	

Table II: Maximum-likelihood parameter estimates

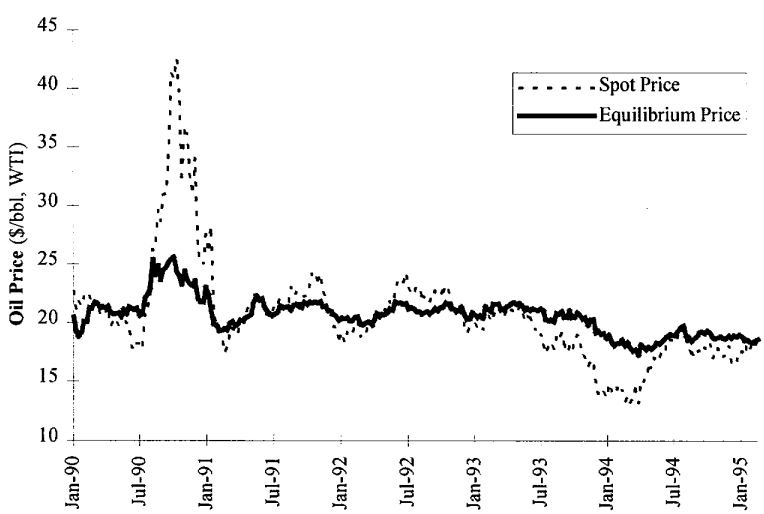

Figure 4a: Estimated spot and equilibrium prices for the futures data

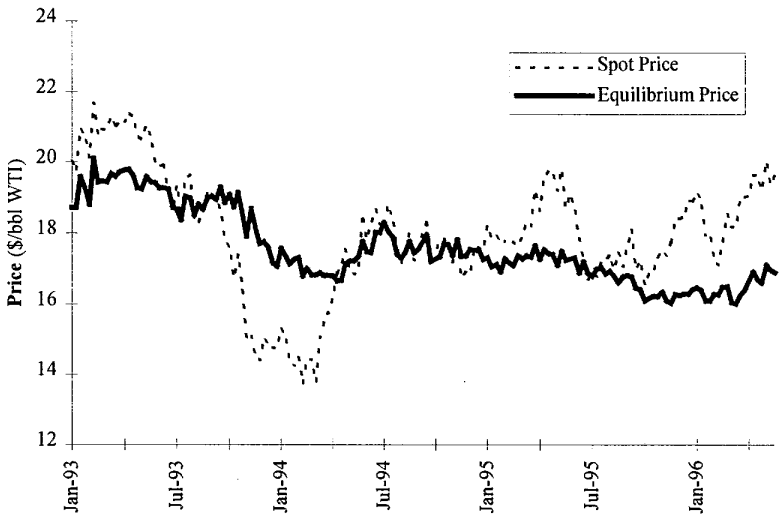

Figure 4b: Estimated spot and equilibrium prices for Enron data

Table III shows the errors in the model's fit to the futures prices. In general, the model appears to fit the mid-term contracts best, with the greatest errors in the prices for 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 .0314) and in the 9-year contract in the Enron data (a mean absolute error of .0332). In both data sets, some of the mid-term futures prices were matched with essentially no error. The standard deviations for the observed errors correspond well to the estimated standard deviations for the measurement equations shown in Table II.

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

Table III: Errors in the model fit to the logarithm of futures prices

Examining the standard errors in Table II, we see that two of the model's parameters, the long-term drift (μ_ξ) and short-term risk premium (λ_χ), are not very accurately estimated.[6] This indeterminacy can be illustrated graphically using Figure 2. Because our observations consist of futures prices (marked with Xs in Figure 2), we get a good estimate of the risk-adjusted growth rate ($\mu_\xi^* = \mu_\xi - \lambda_\xi$) because it determines the growth rate for long-term futures prices. 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 ($\xi_0 - \lambda_\chi/\kappa + \sigma_\chi^2 /4\kappa + \rho_{\chi\xi}\sigma_\chi\sigma_\xi/\kappa$); this is the log of what we earlier referred to earlier as the "shadow spot price." The expected spot prices – represented by the uppermost curve in Figure 2 – are, however, never directly observed and we cannot accurately determine precise location of this curve at any given time. The risk premiums λ_χ and λ_ξ describe the differences between the expected prices and futures prices and, because of the uncertainty about expected prices, these risk premiums are not well estimated. Errors in the estimate of λ_ξ appear in Table II 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 amount (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 spot price. Because these errors shift the state variable estimates by a constant, we can detect changes in the state variables even though we cannot

[6] The negative estimates for the expected growth of the equilibrium price level (μ_ξ) are likely to be artifacts of the time series used for estimation and not representative of investor expectations of long-run growth during that time period. This leads to a counterintuitive 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 .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_t and ξ_t to preserve the spot price and futures prices.

identify their precise level; this allows us to obtain reasonably precise estimates of the volatilities for the two factors.[7]

In summary, though most of the parameter estimates for the model are sensible and appear to be accurately estimated, there are two "degrees of freedom" in the model that are poorly estimated from futures prices. Though these degrees of freedom may be parameterized in different ways, this freedom fundamentally reflects our inability to estimate the risk premiums (λ_χ and λ_ξ) using futures data alone. To estimate these risk-premiums precisely, we would have to use a much longer time series (which is impossible given the relatively short history of long maturity oil futures contracts) or include observations (perhaps in the form of published price forecasts) contingent upon the true, as opposed to the risk-neutral, price process. While these risk premiums may be of academic interest, they are not practically relevant: it is, after all, the risk-neutral process that is relevant for valuing commodity-related securities and projects and the parameters of the risk-neutral process are well-estimated in this framework.[8]

---

[7] To clarify the role of the risk premiums in the model and their irrelevance for valuing futures contracts, we can define "risk-neutralized" versions of the equilibrium and deviation state variables as $\xi_t^+ \equiv \xi_t - \lambda_\chi/\kappa$ and $\chi_t^+ \equiv \chi_t + \lambda_\chi/\kappa$; the spot price ($X_t$) is still given as the sum of the two state variables, $\chi_t^+ + \xi_t^+$. From equations (9) and (10), we can write the risk-neutral process for these risk-neutral state-variables as $d\chi_t^+ = -\kappa\,\chi_t^+\,dt + \sigma_\chi\,dz_\chi^*$ and $d\xi_t^+ = \mu_\xi^*\,dt + \sigma_\xi\,dz_\xi^*$. From equation (13), we can write the log of the futures price as

$$\ln(F_{T,0}) = e^{-\kappa T}\chi_0^+ + \xi_0^+ + \mu_\xi^* T + \frac{1}{2}\left((1-e^{-2\kappa T})\frac{\sigma_\chi^2}{2\kappa} + \sigma_\xi^2 T + 2\,(1-e^{-\kappa T})\frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa}\right).$$

Thus, given these risk-neutral state variables (χ_t^+ and ξ_t^+) and the risk-adjusted drift (μ_ξ^*), the futures prices are independent of the values of the true state variables (χ_t and ξ_t) and the risk-premiums (λ_χ and λ_ξ). Moreover, since these risk-neutral processes are independent of the risk-premiums, the market value of any claim dependent on the spot price (X_t) – given by taking expectations using the risk-neutral processes – are independent of these risk premiums as well. The risk-premiums would, however, play a role in the true process for these new state variables.

[8] If we allow the short-term risk premium to depend on the short-term deviations as discussed in footnote 4, using the Enron data, we find an estimate of the risk-neutral mean reversion rate (κ^*) of 1.19 (with a standard error of .03) and a true mean-reversion rate (κ) of 1.79 (with a standard error of .87). These point estimates suggest that short-term risk premiums are lower in periods with higher short-term deviations (i.e., $\beta = \kappa^* - \kappa = 1.19 - 1.79 = -.60$), 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.

7. Incorporating A Stochastic Growth Rate

Looking at the errors shown in Table III for the model fit to the Enron data, we see that the greatest errors are at the long-end of the futures curve. Examining the errors more closely, we find that the reason for this poor fit is that the "slope" at the long-end of the futures curve has apparently changed over time, while the model assumes this to be constant. For example, in Figure 2 we see that the slope of the model's fit to the long-term futures prices exceeds that of the actual futures price. A simple way to accommodate these changes in slope would be to use the short-term/long-term model to determine futures prices, but allow the equilibrium growth rate (μ_ξ) or its risk-adjusted counterpart (μ_ξ^*) to vary from period to period to fit then-current futures prices. This approach is easy to implement[9] and provides an improved fit to futures curves, but it is theoretically inconsistent in that it allows parameters to vary that are not treated as stochastic when valuing the futures and options on these futures. In this section, we describe an extension of the short-term/long-term model in which the growth rate for the equilibrium price (μ_ξ) is modeled as stochastic and futures and options are valued reflecting this additional source of uncertainty. This uncertainty in equilibrium growth rates may reflect uncertainty about the rate of discovery or depletion of new reserves, uncertainty about demand growth over time, and/or uncertainty about inflation. As we will see, incorporating this third factor greatly improves the model's ability to fit long-term futures prices.

7.1 The Extended Model

In this extension, we assume that the short-term deviations (χ_t) and equilibrium prices (ξ_t) follow the stochastic differential equations (2) and (3) but with the equilibrium growth rate (μ_ξ) in equation (3) being replaced by and stochastic factor μ_t that follows a stochastic process described by

[9] Since μ_ξ^* enters equation for future prices linearly (equation 13), we could use standard Kalman filtering techniques to estimate equation μ_ξ^* as well as χ_t and ξ_t.

$$d\mu_t = -\eta(\mu_t - \bar{\mu})\, dt + \sigma_\mu\, dz_\mu \ . \tag{29}$$

Here dz_μ denotes increments of standard Brownian motion process that are correlated with the increments in equations (2) and (3) with pairwise correlations given by $\rho_{\chi\mu}$, and $\rho_{\xi\mu}$, respectively. Thus we assume that the long-term growth rate follows a mean-reverting process with a "natural home" or "long-term mean" equal to $\bar{\mu}$. If, for example, you believe that prices should grow with interest rates, as in Hotelling's classic model of prices for exhaustible resources, equation (29) is equivalent to assuming that interest rates evolve as in Vasicek (1977). More generally, interest rates and equilibrium growth rates may follow distinct processes, but we might expect them to possess similar dynamics.[10]

In valuing futures contracts, we assume that the risk-neutral version of (29) is of the form

$$d\mu_t^* = (-\eta(\mu_t - \bar{\mu}) - \lambda_\mu)\, dt + \sigma_\mu\, dz_\mu^* \ , \tag{30}$$

with $\mu_0^* = \mu_0$ so that the risk premiums again take the form of a constant reduction in drift. Thus, the risk-neutral process for μ_t is an Ornstein-Uhlenbeck process reverting to $\bar{\mu}^* \equiv \bar{\mu} - \lambda_\mu/\eta$ rather than $\bar{\mu}$. Given this assumption, we can derive the risk-neutral joint distribution for the three-factor model following a derivation similar to that of equation (5); see the appendix. Given $\chi_0^* = \chi_0$, $\xi_0^* = \xi_0$, and $\mu_0^* = \mu_0$, we find that χ_t^*, ξ_t^*, and μ_t^* are jointly normally distributed with mean vector and covariance matrix:

$$E[(\chi_t^*, \xi_t^*, \mu_t^*)] = \left[\ e^{-\kappa t}\chi_0 - (1-e^{-\kappa t})\lambda_\chi/\kappa,\ \xi_0 + (\bar{\mu}^* - \lambda_\xi)t + (\mu_0 - \bar{\mu}^*)\frac{(1-e^{-\eta t})}{\eta},\ \mu_0 - (\mu_0 - \bar{\mu}^*)(1-e^{-\eta t})\ \right] \tag{31a}$$

$$\mathrm{Cov}[(\chi_t^*, \xi_t^*, \mu_t^*)] = \begin{bmatrix} \sigma_{11}(t) & \sigma_{12}(t) & \sigma_{13}(t) \\ \sigma_{12}(t) & \sigma_{22}(t) & \sigma_{23}(t) \\ \sigma_{13}(t) & \sigma_{23}(t) & \sigma_{33}(t) \end{bmatrix} \tag{31b}$$

---

[10] Schwartz (1997) develops a three-factor commodity price model where the three factors are spot prices, convenience yields, and interest rates. The only impact of interest rates in that model is through discounting in the valuation of the futures and forward contracts and, in the empirical analysis, interest rates are estimated independently of spot prices and convenience yields. Here, the third factor relates directly to the futures curve and thereby provides better fits to futures prices.

where $$\sigma_{11}(t) = \sigma_\chi^2 \frac{(1-e^{-2\kappa t})}{2\kappa} ,$$

$$\sigma_{12}(t) = \rho_{\chi\xi}\sigma_\chi\sigma_\xi \frac{(1-e^{-\kappa t})}{\kappa} + \frac{\rho_{\chi\mu}\sigma_\chi\sigma_\mu}{\eta}\left(\frac{(1-e^{-\kappa t})}{\kappa} - \frac{(1-e^{-(\kappa+\eta)t})}{(\kappa+\eta)}\right) ,$$

$$\sigma_{13}(t) = \rho_{\chi\mu}\sigma_\chi\sigma_\mu \frac{(1-e^{-(\kappa+\eta)t})}{(\kappa+\eta)} ,$$

$$\sigma_{22}(t) = \sigma_\xi^2 t + \frac{\rho_{\xi\mu}\sigma_\xi\sigma_\mu}{\eta}\left(t - \frac{(1-e^{-\eta t})}{\eta}\right) + \frac{\sigma_\mu^2}{\eta^2}\left(t - 2\frac{(1-e^{-\eta t})}{\eta} + \frac{(1-e^{-2\eta t})}{2\eta}\right) ,$$

$$\sigma_{23}(t) = \rho_{\xi\mu}\sigma_\xi\sigma_\mu \frac{(1-e^{-\eta t})}{\eta} + \frac{\sigma_\mu^2}{\eta}\left(\frac{(1-e^{-\eta t})}{\eta} + \frac{(1-e^{-2\eta t})}{2\eta}\right) ,$$

$$\sigma_{33}(t) = \sigma_\mu^2 \frac{(1-e^{-2\eta t})}{2\eta} .$$

(The joint distribution for the true, as opposed to risk-neutral, stochastic process may be found by substituting zeros for the risk premiums.)

Under this risk-neutral distribution, the log of the future spot price (X_t^*) is normally distributed with:

$$\mathrm{E}[X_t^*] = e^{-\kappa t}\chi_0 - (1-e^{-\kappa t})\lambda_\chi/\kappa + \xi_0 + (\bar{\mu}-\lambda_\xi)t + (\mu_0-\bar{\mu}^*)\frac{(1-e^{-\eta t})}{\eta}$$

$$\mathrm{Var}[X_t^*] = = \sigma_{11}(t) + \sigma_{22}(t) + 2\sigma_{12}(t) .$$

Following the same analysis as in the two-factor model, we find a futures price $F_{T,0}$ satisfying

$$\begin{aligned} \ln(F_{T,0}) &= \ln(\mathrm{E}[S_t^*]) \\ &= \mathrm{E}[X_T^*] + \frac{1}{2}\mathrm{Var}[X_T^*] \\ &= e^{-\kappa T}\chi_0 + \xi_0 + (\mu_0 - \bar{\mu}^*)\frac{(1-e^{-\eta T})}{\eta} + B(T) \end{aligned} \tag{32}$$

where $B(T)$ depends on the time to maturity but is independent of the state variables (χ_0, ξ_0, μ_0):

$$B(T) = -(1-e^{-\kappa T})\lambda_\chi/\kappa + (\bar{\mu}^* - \lambda_\xi)T + \frac{1}{2}(\sigma_{11}(T) + \sigma_{22}(T) + 2\sigma_{12}(T)) .$$

Thus, as in the two-factor model, the log of the futures price is a linear function of the state variables. This allows us to estimate state variables over time using standard Kalman filtering techniques and estimate model parameters using maximum likelihood methods. We can also derive analytic formulas for European options following a derivation analogous to that of section 3.

As in the two-factor model, the instantaneous volatility of futures prices depends on the time to maturity but is independent of the state variables. From equation (32), this volatility is given as

$$\sigma^2(F_{T,0}) = e^{-2\kappa T}\sigma_\chi^2 + \sigma_\xi^2 + \sigma_\mu^2\frac{(1-e^{-\eta T})^2}{\eta^2} + 2e^{-\kappa T}\rho_{\chi\xi}\sigma_\chi\sigma_\xi + 2e^{-\kappa T}\frac{(1-e^{-\eta T})}{\eta}\rho_{\chi\mu}\sigma_\chi\sigma_\mu + 2\frac{(1-e^{-\eta T})}{\eta}\rho_{\xi\mu}\sigma_\xi\sigma_\mu$$

.

This volatility relationship is illustrated in Figure 5, using the parameter estimates from the Enron data described below. Here, as with the two-factor model, the volatility in prices for near maturity futures contracts (i.e., $T = 0$) is equal to the volatility of the sum of the short-term deviation and equilibrium levels ($\sigma^2(F_{0,0}) = \sigma_\chi^2 + \sigma_\xi^2 + 2\rho_{\chi\xi}\sigma_\chi\sigma_\xi$). As the maturity of the contract increases, the short-term deviations make less and less of a contribution to the volatility and the volatility decreases. As maturity increases more, the volatility begins to increase as the uncertainty about the equilibrium growth rate begins to play a larger role (the $(1-e^{-\eta T})$ term in equation 15 is increasing in T). As $T \to \infty$, $\sigma^2(F_{T,0})$ approaches a constant $\sigma_\xi^2 + \sigma_\mu^2/\eta^2 + 2\rho_{\xi\mu}\sigma_\xi\sigma_\mu/\eta$ (= 13.6% per year with the Enron data). Comparing this volatility curve to that of Figure 3, we see that the two- and three-factor lead to similar volatilities, with the three-factor model leading to slightly higher volatility estimates for very long-term futures contracts (13.6% vs. 11.5%). The option volatilities and observed volatilities are also shown and may be interpreted like those in Figure 3.

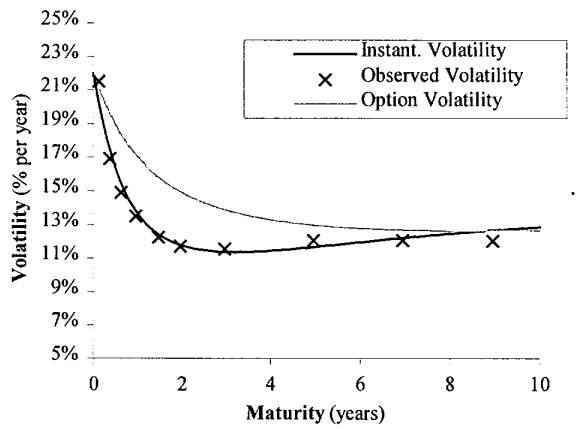

Figure 5: Volatility estimates for three-factor model.

7.2 Empirical Results

We estimate this extended model using the Kalman filtering and maximum likelihood approach described in section 5 with the Enron data. We chose not to use the futures data in this context because we felt that their relatively short maturities (up to 18 months) would not allow us to accurately identify changes in the equilibrium growth rate over time. In contrast, the Enron data includes contracts with maturities up to 9 years and changes in the expected growth rate are more transparent.

The parameter estimates for this extended model are shown in Table IV. Here we see that the model fits the observed futures prices quite well: the standard errors for the measurement equation are less than 1 percent for all but the near-term contract which has a standard deviation of error of about 2%. Overall, the likelihood function has increased from 6,182 for the two-factor model with the same data set to 7,464.[11] Examining the estimated values of new state variable (μ_t) (see Figure 6), we see that the equilibrium growth rate has changed substantially over the time horizon covered by the data set, starting around -9% and moving up to around -4% in late 1993. Because the futures prices do not directly depend

[11] We can compare these results to those obtained using Schwartz's (1997) three-factor model on this same data set (see his Table 9). There the overall likelihood function is 6,287 and the standard errors are similar to those obtained using the two-factor model.

on the values of μ_t (as discussed in the previous section), we do not place much confidence in the levels of the state variables μ_t presented in Figure 5. The values of the risk-adjusted growth rate ($\mu_t - \lambda_\xi$) are more reliably estimated and are also shown in Figure 5. The estimates for the other two state variables are qualitatively similar to those shown in Figure 4 and are not shown here.

			Enron Data	
Parameter	**Description**		**Estimate**	**Standard Error**
κ	Short-term mean-reversion rate		1.26	0.03
σ_χ	Short-term volatility		14.5%	1.0%
λ_χ	Short-term risk premium		1.4%	6.0%
σ_ξ	Equilibrium volatility		13.3%	0.7%
η	Mean-reversion rate for eq. growth rate		.226	.014
$\bar{\mu}$	Mean eq. growth rate		-4.9%	8.3%
$\bar{\mu}^*$	Risk-adjusted mean eq. growth rate		-8.6	8.8%
$(\bar{\mu}^* - \lambda_\xi)$	Risk-adjusted mean eq. growth rate		0.1%	0.1%
σ_μ	Volatility in eq. growth rate		3.3%	0.2%
$\rho_{\chi\xi}$	Correlation between dev. and eq.		.267	.113
$\rho_{\chi\mu}$	Correlation between dev. and growth		-.138	.100
$\rho_{\xi\mu}$	Correlation between eq. and growth		-.524	.060
	Standard deviation(s) of error for measurement equation	Contract Maturity		
s_1	"	2 mo.	0.021	0.001
s_2	"	5 mo.	0.004	0.000
s_3	"	8 mo.	0.000	.
s_4	"	12 mo.	0.002	0.000
s_5	"	18 mo.	0.002	0.000
s_6	"	2 yrs.	0.003	0.000
s_7	"	3 yrs.	0.005	0.000
s_8	"	5 yrs.	0.006	0.001
s_9	"	7 yrs.	0.000	.
s_{10}	"	9 yrs.	0.008	0.001
NT	Number of time periods (weeks)		163	
N	Number of futures contracts		10	
	Log Likelihood		7,463.7	

Table IV: Parameter estimates for the three factor model

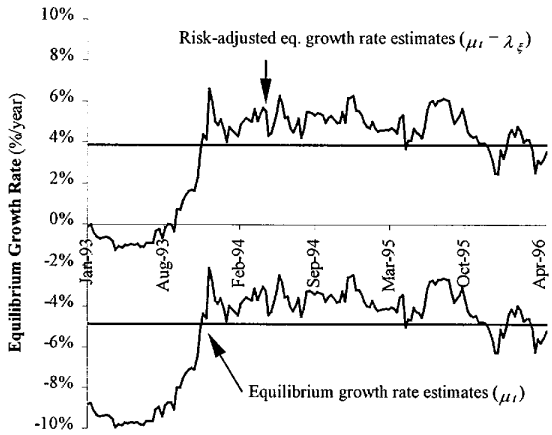

Figure 6: Estimates of the equilibrium growth rate over time

Examining the parameter estimates for three-factor in table IV and comparing them to the corresponding estimates for the two-factor model in Table II, we see that those parameters that appear in both models have similar estimates. Examining the new parameters, we see that the equilibrium growth rate has a volatility (σ_μ) of approximately 3.3% per year and reverts more slowly than the short-term deviations ($\eta = .226$, corresponding to a half-life of approximately 3 years). Given the relatively short time series, we find that we cannot accurately estimate the mean growth rate ($\bar{\mu}$) or its risk-adjusted counterpart ($\bar{\mu}^* \equiv \bar{\mu} - \lambda_\mu / \eta$). We can, however, obtain a reasonably precise estimate of ($\bar{\mu}^* - \lambda_\xi$), because of the role this term plays in determining the long-term futures prices (see equation (32)). Taking differences of these estimates, we find point estimates for the risk premiums of $\lambda_\xi = 8.7\%$ and $\lambda_\mu =$ 0.85%, with relatively large standard errors (approximately 8%). As in the previous section, this reflects the fact that the risk-neutral spot price process does not depend on these risk premiums and estimates of these parameters rely on the dynamics of the implied state variable process.

In summary, this three-factor model provides much improved fits to futures prices over time. Though the formulas for valuing futures and European options are analytic, they are substantially more complicated than the corresponding formulas for the two-factor model. This additional complexity would

seem worthwhile when valuing futures or options on futures with long maturities or long-lived real options.

8. Summary and Conclusions

In this paper, we propose a new way of thinking about the stochastic behavior of commodity prices. We 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. Though 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 the risk premiums are non-stochastic and changes in the slope of the futures curve are interpreted as short-term price variations rather than changes in the instantaneous convenience yield.

Though our two-factor short-term/long-term model and the Gibson-Schwartz stochastic convenience yield model are formally equivalent, we believe that this new short-term/long-term model is easier to interpret and work with for several reasons. First, while 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 fairly natural and may 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 by the sum of the volatilities for the short- and long-term factors. As the maturity of the contract increases, the futures volatility approaches the volatility of the equilibrium level.

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. Moreover, this correlation in increments (estimated at .189 and .300 for the two data sets) is much less than the correlation between increments in the stochastic convenience yield model (estimated at .845 and .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. For example, consider the problem of evaluating an (American) option to develop an oil field that takes several years to develop and 20 or 30 years to produce. Given the strong mean-reversion in the short-term variations, we may be able to safely ignore the short-term variations and evaluate the field using a one-factor model that considers uncertainty in equilibrium prices only. In this case, we would use our model to estimate the current equilibrium price[12] and then use an annual volatility that reflects this one uncertainty without being contaminated by the short-term variations in spot prices. To value other projects and securities, we may need to consider both short- and long-term dynamics. For example, to value an oil or gas production facility that can temporarily curtail production, we may need to consider both short- and long-term factors. In these cases, we can use multivariate lattice models, analogous to those developed in Boyle, Envine, and Gibbs (1989), to value the project and determine optimal exercise policies.

Perhaps most important, 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. In this paper, we have developed one such extension in which the growth rate for the equilibrium price is stochastic; this third factor greatly improves the model's ability to fit 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 errors in the fit to the futures curves, or, to better fit option prices,

allowing the short-term volatility (σ_χ) to be stochastic as well. 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. This short-term/long-term modeling framework, coupled with estimation procedures based on futures prices, thus seems like a promising approach for modeling commodity price uncertainty and valuing commodity related projects and derivative securities.

Appendix

Derivation of Equations (5) and (11): We derive equation (11); equation (5) is given as a special case in which $\lambda_\chi = \lambda_\xi = 0$. 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 (2) and (3) and then take the limit as the time steps are made infinitesimally small. The discrete-time approximation of the process with time steps of length $\Delta t = t/n$ can be written as $\boldsymbol{x}_t = \boldsymbol{c} + \boldsymbol{Q}\,\boldsymbol{x}_{t-1} + \eta_t$ where $\boldsymbol{x}_t \equiv [\chi_t, \xi_t]$, $\boldsymbol{c} \equiv [-\lambda_\chi \Delta t, \mu_\xi^* \Delta t]$,

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

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

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

Note that this is an approximation of the state transition equation given by equation (27). With this process, the n-step ahead mean vector ($\boldsymbol{m}_n$) and covariance matrix ($\boldsymbol{V}_n$) are given recursively by $\boldsymbol{m}_n = \boldsymbol{Q}\,\boldsymbol{m}_{n-1}$ and $\boldsymbol{V}_n = \boldsymbol{Q}\,\boldsymbol{V}_{n-1}\,\boldsymbol{Q}' + \boldsymbol{W}$ with $\boldsymbol{m}_0 = \boldsymbol{x}_0 \equiv [\chi_0, \xi_0]$ and $\boldsymbol{V}_0 = \mathbf{0}$ (see, for example, Harvey 1989, pg. 109). Applying this recursion, we find

$$\boldsymbol{m}_n = [\phi^n \chi_0 - \lambda_\chi \Delta t \textstyle\sum_{i=0}^{n-1} \phi^i,\ \xi_0 + \mu_\xi^* n \Delta t]$$

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

---

[12] As suggested in Schwartz (1998), an easy way to estimate the current values of the state variables is to choose state values to fit the current futures curve to minimize the sum of squared errors (or perhaps a weighted sum of squared errors) in this fit. These "implied estimates" can be interpreted as Kalman filter estimates in the limiting case where the measurement error covariance matrix (H in equation (28)) approaches 0.

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

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

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

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

Substituting these limiting forms into the expressions for $\boldsymbol{m}_n$ and V_n, we arrive at the time-t mean vector and covariance matrix given in equations (5) and (11).///

Derivation of Equation (31): This derivation is analogous to the derivation of equations (5) and (11) given above. Here, from equations (9), (10) and (30), the transition equation is given as $\boldsymbol{x}_t = \boldsymbol{c} + \boldsymbol{Q}\,\boldsymbol{x}_{t-1} + \eta_t$ with $\boldsymbol{x}_t \equiv [\chi_t, \xi_t, \mu_t]$, $\boldsymbol{c} \equiv [-\lambda_\chi\Delta t, -\lambda_\xi\Delta t, \eta\bar{\mu}^*\Delta t]$,

$$\boldsymbol{Q} \equiv \begin{bmatrix} \phi_1 & 0 & 0 \\ 0 & 1 & \Delta t \\ 0 & 0 & \phi_2 \end{bmatrix},$$

$\phi_1 \equiv 1-\kappa\Delta t$, $\phi_2 = 1-\eta\Delta t$, η_t is a vector of serially uncorrelated, normally distributed disturbances with $E[\eta_t] = \mathbf{0}$ and

$$\mathrm{Var}[\eta_t] = \boldsymbol{W} \equiv \begin{bmatrix} \sigma_\chi^2\Delta t & \rho_{\chi\xi}\sigma_\chi\sigma_\xi\Delta t & \rho_{\chi\mu}\sigma_\chi\sigma_\mu\Delta t \\ \rho_{\chi\xi}\sigma_\chi\sigma_\xi\Delta t & \sigma_\xi^2\Delta t & \rho_{\xi\mu}\sigma_\xi\sigma_\mu\Delta t \\ \rho_{\chi\mu}\sigma_\chi\sigma_\mu\Delta t & \rho_{\xi\mu}\sigma_\xi\sigma_\mu\Delta t & \sigma_\mu^2\Delta t \end{bmatrix} .$$

Applying the same recursive procedure as before ($\boldsymbol{m}_n = \boldsymbol{Q}\,\boldsymbol{m}_{n-1}$ and with $\boldsymbol{m}_0 = \boldsymbol{x}_0 \equiv [\chi_0, \xi_0]$), the n-step ahead mean vector ($\boldsymbol{m}_n$) is given as:

$$\boldsymbol{m}_n = \begin{bmatrix} \phi_1^n\chi_0 - \lambda_\chi\sum_{i=0}^{n-1}\phi_1^i \\ \xi_0 - \lambda_\xi n\Delta t + \mu_0\Delta t\sum_{i=0}^{n-1}\phi_2^i + \eta\bar{\mu}^*\Delta t^2\sum_{i=0}^{n-2}\sum_{j=0}^{i}\phi_1^j \\ \phi_2^n\mu_0 + \eta\bar{\mu}^*\Delta t\sum_{i=0}^{n-1}\phi_2^i \end{bmatrix} .$$

Most of these expressions were encountered in the derivation of equations (5) and (11) and have similar limiting forms here. The one new expression is the double summation in the second entry. Recognizing the nested geometric series and taking the limit as $n \to \infty$ and $\Delta t = t/n \to 0$, this can be rewritten as

$$\Delta t^2\sum_{i=0}^{n-2}\sum_{j=0}^{i}\phi_2^j = \Delta t^2\sum_{i=0}^{n-2}\left(\frac{1-\phi_2^i}{1-\phi_2}\right) = \frac{\Delta t^2}{1-\phi_2}\left((n-1) + \frac{1-\phi_2^{n-2}}{1-\phi_2}\right) \to \frac{1}{\eta}\left(t + \frac{1-e^{-\eta t}}{\eta}\right).$$

Thus the mean vector approaches

$$\boldsymbol{m}_T = \begin{bmatrix} e^{-\kappa t}\chi_0 - \lambda_\chi(1-e^{-\kappa t})/\kappa \\ \xi_0 - \lambda_\xi t + \mu_0(1-e^{-\eta t})/\eta + \bar{\mu}^*(t + (1-e^{-\eta t})/\eta) \\ e^{-\eta t}\mu_0 + \bar{\mu}^*(1-e^{-\eta t}) \end{bmatrix},$$

which with some rearrangement leads to the form in equation (31a).

To derive the covariance matrix, we proceed term by term through the matrix, beginning in each case with the terms given by applying discrete time recursion ($\boldsymbol{V}_n = \boldsymbol{Q}\,\boldsymbol{V}_{n-1}\,\boldsymbol{Q}' + \boldsymbol{W}$ with $\boldsymbol{V}_0 = \boldsymbol{0}$) and taking the limits as $n \to \infty$ and $\Delta t = t/n \to 0$. Let σ_{ijn} denote the entry in the ith row and jth column of $\boldsymbol{V}_n$. The first term, σ_{11n}, is similar that encountered in the derivation of equation (5b) above:

$$\sigma_{11n} = \sigma_\chi^2 \Delta t \sum_{i=0}^{n-1} \phi_1^{2i} \;\to\; \sigma_{11}(t) = \sigma_\chi^2 \frac{(1-e^{-2\kappa t})}{2\kappa}$$

The recursion for the second term yields

$$\sigma_{12n} = \rho_{\chi\xi}\sigma_\chi\sigma_\xi \Delta t \sum_{i=0}^{n-1} \phi_1^i + \rho_{\chi\mu}\sigma_\chi\sigma_\mu \Delta t^2 \sum_{i=1}^{n-1}\left(\phi_1^i \sum_{j=0}^{i-1} \phi_2^j\right) .$$

The first part of this expression is familiar. We can handle the second term by recognizing the nested geometric series:

$$\Delta t^2 \sum_{i=1}^{n-1}\left(\phi_1^i \sum_{j=0}^{i-1} \phi_2^j\right) = \Delta t^2 \sum_{i=1}^{n-1} \phi_1^i \left(\frac{1-\phi_2^{i-1}}{1-\phi_2}\right) = \frac{\phi_1 \Delta t^2}{1-\phi_2}\left(\frac{1-\phi_1^{n-2}}{1-\phi_1} - \frac{1-(\phi_1\phi_2)^{n-2}}{1-\phi_1\phi_2}\right) .$$

Taking the limit, we then find

$$\sigma_{12n} \to \sigma_{12}(t) = \rho_{\chi\xi}\sigma_\chi\sigma_\xi \frac{(1-e^{-\kappa t})}{\kappa} + \frac{\rho_{\chi\mu}\sigma_\chi\sigma_\mu}{\eta}\left(\frac{(1-e^{-\kappa t})}{\kappa} - \frac{(1-e^{-(\kappa+\eta)t})}{(\kappa+\eta)}\right) .$$

For σ_{13n}, we find a familiar form,

$$\sigma_{13n} = \rho_{\chi\mu}\sigma_\chi\sigma_\mu \Delta t \sum_{i=0}^{n-1}(\phi_1\phi_2)^i \to \sigma_{12}(t) = \rho_{\chi\mu}\sigma_\chi\sigma_\mu \frac{(1-e^{-(\kappa+\eta)t})}{(\kappa+\eta)} .$$

For σ_{22n}, the recursion yields,

$$\sigma_{22n} = \sigma_\xi^2 n\Delta t + \rho_{\xi\mu}\sigma_\xi\sigma_\mu \Delta t^2 \sum_{i=0}^{n-2}\sum_{j=0}^{i} \phi_2^j + \sigma_\mu^2 \Delta t^3 \sum_{i=0}^{n-2}\left(\left(\sum_{j=0}^{i}\phi_2^j\right)\left(\sum_{j=0}^{i}\phi_2^j\right)\right)$$

The limit of the first term ($\sigma_\xi^2 n\Delta t$) is straightforward. We encountered a form similar to the second term in the derivation of $\boldsymbol{m}_n$. The third term can be rewritten using

$$\Delta t^3 \sum_{i=0}^{n-2}\left(\left(\sum_{j=0}^{i}\phi_2^j\right)\left(\sum_{j=0}^{i}\phi_2^j\right)\right) = \Delta t^3 \sum_{i=0}^{n-2}\left(\left(\frac{1-\phi_2^i}{1-\phi_2}\right)\left(\frac{1-\phi_2^i}{1-\phi_2}\right)\right) = \frac{\Delta t^3}{(1-\phi_2)^2}\sum_{i=0}^{n-2}(1-2\phi_2^i+\phi_2^{2i}) .$$

Taking the limit of this expression and the others in σ_{22n}, we have

$$\sigma_{22n} \to \sigma_{22}(t) = \sigma_\xi^2 t + \frac{\rho_{\xi\mu}\sigma_\xi\sigma_\mu}{\eta}\left(t - \frac{(1-e^{-\eta t})}{\eta}\right) + \frac{\sigma_\mu^2}{\eta^2}\left(t - 2\frac{(1-e^{-\eta t})}{\eta} + \frac{(1-e^{-2\eta t})}{2\eta}\right) .$$

For σ_{23n}, the recursion yields

$$\sigma_{23n} = \rho_{\xi\mu}\sigma_\xi\sigma_\mu\Delta t\sum_{i=0}^{n-1}\phi_2^i + \sigma_\mu^2\Delta t^2 \sum_{i=1}^{n-1}(\phi_2^i \sum_{j=0}^{i-1} \phi_2^j) .$$

The first term is, by now, familiar. We encountered a term similar to the second in our derivation of $\sigma_{12}(t)$, albeit with ϕ_1 in place of one of the ϕ_2. Taking the limit gives,

$$\sigma_{23n} \rightarrow \sigma_{23}(t) = \rho_{\xi\mu}\sigma_\xi\sigma_\mu\frac{(1-e^{-\eta t})}{\eta} + \frac{\sigma_\mu^2}{\eta}\left(\frac{(1-e^{-\eta t})}{\eta} + \frac{(1-e^{-2\eta t})}{2\eta}\right)$$

Finally, the expression for σ_{33} and $\sigma_{33}(t)$ are analogous to the expressions for σ_{11} and $\sigma_{11}(t)$.///

References

Bessembinder, H., J. F. Coughenour, P. J. Seguin, and M. M. Smoller, 1995, Mean reversion in equilibrium assset prices: Evidence from the futures term structure, *Journal of Finance* 50, 361-375.

Boyle, P. B., J. Evnine, S. Gibbs, 1989, "Numerical evaluation of multivariate contingent claims," *Review of Financial Studies* 2, 241-250.

Brennan, M. J., 1991, The price of convenience and the valuation of commodity contingent claims, in D. Lund and B. Oksendal, eds., *Stochastic Models and Option Values: Applications to Resources, Environment, and Investment Problems*, North-Holland.

Brennan, M. J., and N. Crew, 1995, Hedging long maturity commodity commitments with short-dated futures contracts, Working paper, Anderson Graduate School of Management, U.C.L.A.

Cortazar, G., and E. S. Schwartz, 1994, The evaluation of commodity contingent claims, *Journal of Derivatives* 1, 27-39.

Cox, J., J. Ingersoll, and S. Ross, 1981, The relation between forward prices and futures prices, *Journal of Financial Economics* 9, 321-346.

Dixit, A. K., and R. S. Pindyck, 1994, *Investment Under Uncertainty*, Princeton University Press.

Duffie, D., 1992, *Dynamic Asset Pricing Theory*, Princeton University Press.

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., and E. S. Schwartz, 1990, Stochastic convenience yield and the pricing of oil contingent claims, *Journal of Finance* 45, 959-976.

Harvey, A. C., 1989, *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge University Press.

Karlin, S. ,and H. M Taylor, 1981, *A Second Course in Stochastic Processes*, Academic Press.

Laughton, D. G., and Henry D. Jacoby, 1993, Reversion, timing options, and long-term decision making, *Financial Management* 33, 225-240.

Laughton, D. G., and Henry D. Jacoby, 1995, The effects of reversion on commodity projects of different length, in L. Trigeorgis, ed., *Real Options in Capital Investment: Models, Strategies and Applications,* 185-205.

Litzenberger, R. H. and N. Rabinowitz 1995, Backwardation in oil futures markets: theory and empirical evidence, *Journal of Finance* 50, 1517-1545.

Pindyck, R.S. 1980, Uncertainty and exhaustible resource markets, *Journal of Political Economy* 88, 1203-1225.

Pindyck, R.S. 1994, Inventories and the short-run dynamics of commodity prices, *The RAND Journal of Economics* 25, 141-159.

Pindyck, R.S. 1997, The long-run dynamics of commodity prices, Working paper, Massachusetts Institute of Technology.

Routledge, B., D. Seppi, and C. Spatt 1998, Equilibrium forward curves for commodities, Working paper, Graduate School of Industrial Administration, Carnegie-Mellon University.

Schwartz, E. S. 1997, The stochastic behavior of commodity prices: implications for valuation and hedging, *Journal of Finance* 52, 923-73.

Schwartz, E. S. 1998, Valuing long term commodity assets, *Financial Management* 27, 57-66.

Shimko, D. C. 1994, Options on futures spreads: hedging, speculation, and valuation, *Journal of Futures Markets* 14, 183-214.

Smith, J. E., and K. F. McCardle 1997, Options in the Real World: Lessons learned in evaluating oil and gas investments, forthcoming in *Operations Research.*

Vasicek, O. 1977, An equilibrium characterization of the term structure, *Journal of Financial Economics* 5, 177-188.

West, M. and J. Harrison 1996, *Bayesian Forecasting and Dynamic Models,* Second Edition, Springer-Verlag, New York.

Zhou, Z. 1998, An equilibrium analysis of hedging with liquidity constraints, speculation, and government price subsidy in a commodity, forthcoming *Journal of Finance.*