



Eduardo S. Schwartz
President of the American Finance Association
1996

The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging

EDUARDO S. SCHWARTZ*

ABSTRACT

In this article we compare three models of the stochastic behavior of commodity prices that take into account mean reversion, in terms of their ability to price existing futures contracts, and their implication with respect to the valuation of other financial and real assets. The first model is a simple one-factor model in which the logarithm of the spot price of the commodity is assumed to follow a mean reverting process. The second model takes into account a second stochastic factor, the convenience yield of the commodity, which is assumed to follow a mean reverting process. Finally, the third model also includes stochastic interest rates. The Kalman filter methodology is used to estimate the parameters of the three models for two commercial commodities, copper and oil, and one precious metal, gold. The analysis reveals strong mean reversion in the commercial commodity prices. Using the estimated parameters, we analyze the implications of the models for the term structure of futures prices and volatilities beyond the observed contracts, and for hedging contracts for future delivery. Finally, we analyze the implications of the models for capital budgeting decisions.

THE STOCHASTIC BEHAVIOR OF commodity prices plays a central role in the models for valuing financial contingent claims on the commodity, and in the procedures for evaluating investments to extract or produce the commodity. Earlier studies, by assuming that interest rates and convenience yields are constant allowed for a straight forward extension of the procedures developed for common stock option pricing to the valuation of financial and real commodity contingent claims. The assumption, however, is clearly not very satisfactory since it implies that the volatility of future prices is equal to the volatility of spot prices, and that the distribution of future spot prices under the equivalent martingale measure has a variance that increases without bound as the horizon increases. In an equilibrium setting we would expect that when prices are relatively high, supply will increase since higher cost producers of the commodity will enter the market putting a downward pressure on prices. Conversely, when prices are relatively low, supply will decrease since some of

* Anderson Graduate School of Management, University of California Los Angeles. This article was prepared for the Presidential Address of the American Finance Association meetings in New Orleans, January 1997. I thank Richard Roll, Alan Kraus, Ernest Greenwood, and Kristian Miltersen for suggestions and Olivier Ledoit, Piet de Jong and Julian Franks for helpful discussions. I especially thank Michael Brennan, not only for his comments on this article, but for 25 years of inspiration, advice, and friendship. Part of this research was done while I was a visiting scholar at the University of British Columbia during the summer of 1996.

the higher cost producers will exit the market, putting upward pressure on prices. The impact of relative prices on the supply of the commodity will induce mean reversion in commodity prices.¹

In this article we compare three models of the stochastic behavior of commodity prices that take into account mean reversion, in terms of their ability to price existing futures contracts and their implication with respect to the valuation of other financial and real assets. The first model is a simple one-factor model in which the logarithm of the spot price of the commodity is assumed to follow a mean reverting process of the Ornstein–Uhlenbeck type. The second model we consider is a variation of the two-factor Gibson and Schwartz (1990) model. The second factor in this model is the convenience yield of the commodity and it is assumed to follow a mean reverting process. Finally, we extend the Gibson and Schwartz model to include stochastic interest rates. In this three-factor model, the instantaneous interest rate is also assumed to follow a mean reverting process as in Vasicek (1977).

For these three models, closed form solutions for the prices of futures and forward contracts² can be obtained, which greatly simplifies the comparative statics and empirical estimation. In addition, for all three models the logarithm of the futures price is linear in the underlying factors, a property that turns out to be very useful in view of the econometric technique used to estimate the parameters of the models.

One of the main difficulties in the empirical implementation of commodity price models is that frequently the factors or state variables of these models are not directly observable. In many cases the spot price of a commodity is so uncertain that the corresponding futures contract closest to maturity is used as a proxy for the spot price. The instantaneous convenience yield is even more difficult to estimate. Futures contracts, however, are traded on several exchanges and their prices are more easily observed.

A tool that is especially well suited to deal with situations in which the state variables are not observable, but are known to be generated by a Markov process, is the state space form. Once a model has been put in state space form, the Kalman filter may be applied to estimate the parameters of the model and the time series of the unobservable state variables.

We apply the Kalman filter method to estimate the parameters of the three models for two commercial commodities, copper and oil, and for one precious metal, gold. The analysis reveals strong mean reversion in the commercial commodity prices, but not for the precious metal. Using the estimated parameters we analyze the implications of the models for the term structure of

¹ The mean reverting nature of commodity prices has been considered in a series of recent articles. See for example Brennan (1991), Gibson and Schwartz (1990), Cortazar and Schwartz (1994), Bessembinder, Coughenour, Seguin, and Smoller (1995), and Ross (1995).

² Since the first two models assume that the interest rate is constant, for these models prices of futures and forward contracts are the same (see Cox, Ingersoll, and Ross (1981)).

futures prices and volatilities beyond the observed contracts, and for hedging contracts for future delivery.

The “real options” approach to investment under uncertainty (for an excellent recent survey of the literature see Dixit and Pindyck (1994)) emphasizes the importance of uncertainty for the value of a project and for determining when the project should be undertaken. The valuation of natural resource investment projects and the rule for determining when it is optimal to invest depend significantly on the stochastic process assumed for the underlying commodity price.³ We compare the value and the investment rule for simple projects under the different assumptions about the commodity price process implied by the models, using realistic estimated parameters.

The remainder of the article is organized as follows. The valuation models are developed in Section I and Section II delineates their empirical counterparts. Section III describes the data, and Section IV reports the empirical estimates of the models and a comparison of their relative performance. In Section V the implications of the models for the volatility of futures returns are discussed. Section VI considers futures contracts with longer maturities than the available data, and Section VII the hedging of contracts for future delivery. Section VIII looks at the implications of the models for investment under uncertainty and compares their predictions with two benchmarks: the discounted cash flow criterion and a real option model with no mean reversion. Section IX concludes.

I. Valuation Models

In this section we present three models of commodity prices and derive the corresponding formulas for pricing futures contracts in each model. The first model, which is a one-factor model, assumes that the logarithm of the spot price of the commodity follows a mean reverting process of the Ornstein–Uhlenbeck type. The second model includes a second stochastic factor, the convenience yield, which is mean reverting and positively correlated with the spot price.⁴ The third model extends the second one by allowing for stochastic interest rates. The three models are very tractable, since they allow for closed form solutions for futures prices and for a linear relation between the logarithm of futures prices and the underlying factors. These properties will be extensively used in the empirical work that follows.

³ See for example Ingersoll and Ross (1992).

⁴ The positive correlation between changes in the spot price and changes in the convenience yield of the commodity is induced by the level of inventories. When inventories of the commodity decrease, the spot price should increase since the commodity is scarce and the convenience yield should also increase since futures prices will not increase as much as the spot price, and vice versa when inventories increase.

A. Model 1

To develop the one-factor model we first assume that the commodity spot price follows the stochastic process:⁵

$$dS = \kappa(\mu - \ln S)Sdt + \sigma Sdz \quad (1)$$

Defining $X = \ln S$ and applying Ito's Lemma, this implies that the log price can be characterized by an Ornstein–Uhlenbeck stochastic process:

$$dX = \kappa(\alpha - X)dt + \sigma dz \quad (2)$$

$$\alpha = \mu - \frac{\sigma^2}{2\kappa} \quad (3)$$

The magnitude of the speed of adjustment $\kappa > 0$ measures the degree of mean reversion to the long run mean log price, α . The second term in equation (2) characterizes the volatility of the process, with dz being an increment to a standard Brownian motion.

In this model, the commodity is not an asset in the usual sense⁶ and the spot price, or equivalently the log of the spot price, plays the role of an underlying state variable upon which contingent claims can be written. Under standard assumptions, the dynamics of the Ornstein–Uhlenbeck process under the equivalent martingale measure can be written as:⁷

$$dX = \kappa(\alpha^* - X)dt + \sigma dz^* \quad (4)$$

where $\alpha^* = \alpha - \lambda$, λ is the market price of risk (assumed constant)⁸ and dz^* is the increment to the Brownian motion under the equivalent martingale measure.

From equation (4), the conditional distribution of X at time T under the equivalent martingale measure is normal with mean and variance:

$$\begin{aligned} E_0[X(T)] &= e^{-\kappa T}X(0) + (1 - e^{-\kappa T})\alpha^* \\ \text{Var}_0[X(T)] &= \frac{\sigma^2}{2\kappa} (1 - e^{-2\kappa T}). \end{aligned} \quad (5)$$

Since $X = \ln S$, the spot price of the commodity at time T is log-normally distributed under the martingale measure with these same parameters.

Assuming a constant interest rate, the futures (or forward) price of the commodity with maturity T is the expected price of the commodity at time T

⁵ This model is similar to the one proposed by Ross (1995).

⁶ See Ross (1995).

⁷ See for example Bjerk Sund and Ekern (1995).

⁸ More generally we would expect the market price of risk to be related to the business cycle and to be correlated with the level of inventories.

under the equivalent martingale measure. Then, from the properties of the log-normal distribution, we have:

$$F(S, T) = E[S(T)] = \exp(E_0[X(T)] + \frac{1}{2} \text{Var}_0[X(T)]) \quad (6)$$

Then:

$$F(S, T) = \exp \left[e^{-\kappa T} \ln S + (1 - e^{-\kappa T}) \alpha^* + \frac{\sigma^2}{4\kappa} (1 - e^{-2\kappa T}) \right] \quad (7)$$

Or, in log form:

$$\ln F(S, T) = e^{-\kappa T} \ln S + (1 - e^{-\kappa T}) \alpha^* + \frac{\sigma^2}{4\kappa} (1 - e^{-2\kappa T}) \quad (8)$$

This last equation is the one used in the empirical tests.

It is easy to verify that equation (7) is the solution of the partial differential equation:

$$\frac{1}{2} \sigma^2 S^2 F_{SS} + \kappa (\mu - \lambda - \ln S) S F_S - F_T = 0 \quad (9)$$

with terminal boundary condition $F(S, 0) = S$.

B. Model 2

The two factor model is based on the one developed by Gibson and Schwartz (1990). The first factor is the spot price of the commodity and the second is the instantaneous convenience yield, δ .⁹ These factors are assumed to follow the joint stochastic process:

$$dS = (\mu - \delta) S dt + \sigma_1 S dz_1 \quad (10)$$

$$d\delta = \kappa(\alpha - \delta) dt + \sigma_2 dz_2 \quad (11)$$

where the increments to standard Brownian motion are correlated with:

$$dz_1 dz_2 = \rho dt. \quad (12)$$

Equation (10) is a standard process for the commodity price allowing for a stochastic convenience yield, which follows a Ornstein-Uhlenbeck stochastic process described in equation (11). Note that if δ , instead of being stochastic as in equation (11), is a deterministic function of S , $\delta(S) = \kappa \ln S$, Model 2 reduces to Model 1, and if δ is constant it reduces to the model of Brennan and Schwartz (1985).

⁹ The convenience yield can be interpreted as the flow of services accruing to the holder of the spot commodity but not to the owner of a futures contract.

Defining once again $X = \ln S$ and applying Ito's Lemma, the process for the log price can be written as

$$dX = (\mu - \delta - \frac{1}{2}\sigma_1^2)dt + \sigma_1 dz_1. \quad (13)$$

In this model the commodity is treated as an asset that pays a stochastic dividend yield δ . Thus, the risk adjusted drift of the commodity price process will be $r - \delta$. Since convenience yield risk cannot be hedged, the risk-adjusted convenience yield process will have a market price of risk associated with it. The stochastic process for the factors under the equivalent martingale measure can be expressed as:¹⁰

$$dS = (r - \delta)Sdt + \sigma_1 S dz_1^* \quad (14)$$

$$d\delta = [\kappa(\alpha - \delta) - \lambda]dt + \sigma_2 dz_2^* \quad (15)$$

$$dz_1^* dz_2^* = \rho dt \quad (16)$$

where now λ is the market price of convenience yield risk, which is assumed constant. Futures prices must then satisfy the partial differential equation:

$$\begin{aligned} \frac{1}{2}\sigma_1^2 S^2 F_{SS} + \sigma_1 \sigma_2 \rho S F_{S\delta} + \frac{1}{2}\sigma_2^2 F_{\delta\delta} \\ + (r - \delta)S F_S + (\kappa(\alpha - \delta) - \lambda)F_\delta - F_T = 0 \end{aligned} \quad (17)$$

subject to the terminal boundary condition $F(S, \delta, 0) = S$.

Jamshidian and Fein (1990) and Bjerksund (1991) have shown that the solution to (17) is:¹¹

$$F(S, \delta, T) = S \exp \left[-\delta \frac{1 - e^{-\kappa T}}{\kappa} + A(T) \right] \quad (18)$$

Or, in log form:

$$\ln F(S, \delta, T) = \ln S - \delta \frac{1 - e^{-\kappa T}}{\kappa} + A(T) \quad (19)$$

where

$$\begin{aligned} A(T) = \left(r - \hat{\alpha} + \frac{1}{2} \frac{\sigma_2^2}{\kappa^2} - \frac{\sigma_1 \sigma_2 \rho}{\kappa} \right) T + \frac{1}{4} \sigma_2^2 \frac{1 - e^{-2\kappa T}}{\kappa^3} \\ + \left(\hat{\alpha} \kappa + \sigma_1 \sigma_2 \rho - \frac{\sigma_2^2}{\kappa} \right) \frac{1 - e^{-\kappa T}}{\kappa^2} \end{aligned} \quad (20)$$

$$\hat{\alpha} = \alpha - \frac{\lambda}{\kappa}.$$

¹⁰ See Gibson and Schwartz (1990).

¹¹ Brennan and Crew (1995) use this formulation in their article.

C. Model 3

Model 3 is a three-factor model of commodity contingent claims. The stochastic factors or state variables in the model are the spot price of the commodity, the instantaneous convenience yield, and the instantaneous interest rate. By assuming a simple mean reverting process for the interest rate, it is possible to obtain a closed form solution for futures prices.

Assuming that the instantaneous risk free interest rate follows a Ornstein–Uhlenbeck stochastic process (as in Vasicek (1977)), Model 2 can easily be extended to a three-factor model. Using (14) and (15), the joint stochastic process for the factors under the equivalent martingale measure can be expressed as:

$$dS = (r - \delta)Sdt + \sigma_1 S dz_1^* \quad (21)$$

$$d\delta = \kappa(\hat{\alpha} - \delta)dt + \sigma_2 dz_2^* \quad (22)$$

$$dr = a(m^* - r)dt + \sigma_3 dz_3^* \quad (23)$$

$$dz_1^* dz_2^* = \rho_1 dt, \quad dz_2^* dz_3^* = \rho_2 dt, \quad dz_1^* dz_3^* = \rho_3 dt. \quad (24)$$

Where a and m^* are, respectively, the speed of adjustment coefficient and the risk adjusted mean short rate of the interest rate process. Futures prices must then satisfy the partial differential equation:

$$\begin{aligned} & \frac{1}{2}\sigma_1^2 S^2 F_{SS} + \frac{1}{2}\sigma_2^2 F_{\delta\delta} + \frac{1}{2}\sigma_3^2 F_{rr} + \sigma_1\sigma_2\rho_1 S F_{S\delta} + \sigma_2\sigma_3\rho_2 F_{\delta r} + \sigma_1\sigma_3\rho_3 S F_{Sr} \\ & + (r - \delta)S F_S + \kappa(\hat{\alpha} - \delta)F_{\delta} + a(m^* - r)F_r - F_T = 0 \end{aligned} \quad (25)$$

subject to the terminal boundary condition $F(S, \delta, r, 0) = S$.

The solution to partial differential equation (25) subject to its terminal boundary condition can be shown to be:¹²

$$F(S, \delta, r, T) = S \exp \left[\frac{-\delta(1 - e^{-\kappa T})}{\kappa} + \frac{r(1 - e^{-aT})}{a} + C(T) \right] \quad (26)$$

Or, in log form:

$$\ln F(S, \delta, r, T) = \ln S - \frac{\delta(1 - e^{-\kappa T})}{\kappa} + \frac{r(1 - e^{-aT})}{a} + C(T), \quad (27)$$

¹² This can easily be verified by substitution.

where

$$\begin{aligned}
 C(T) = & \frac{(\kappa \hat{\alpha} + \sigma_1 \sigma_2 \rho_1)((1 - e^{-\kappa T}) - \kappa T)}{\kappa^2} \\
 & - \frac{\sigma_2^2(4(1 - e^{-\kappa T}) - (1 - e^{-2\kappa T}) - 2\kappa T)}{4\kappa^3} \\
 & - \frac{(am^* + \sigma_1 \sigma_3 \rho_3)((1 - e^{-aT}) - aT)}{a^2} \\
 & - \frac{\sigma_3^2(4(1 - e^{-aT}) - (1 - e^{-2aT}) - 2aT)}{4a^3} \\
 & + \sigma_2 \sigma_3 \rho_2 \left(\frac{(1 - e^{-\kappa T}) + (1 - e^{-aT}) - (1 - e^{-(\kappa+a)T})}{\kappa a (\kappa + a)} \right. \\
 & \left. + \frac{\kappa^2(1 - e^{-aT}) + a^2(1 - e^{-\kappa T}) - \kappa a^2 T - a \kappa^2 T}{\kappa^2 a^2 (\kappa + a)} \right). \quad (28)
 \end{aligned}$$

Since in this model interest rates are stochastic, futures prices are not equal to forward prices. With the assumed risk-adjusted stochastic process for the instantaneous interest rate given in equation (23), the present value of a unit discount bond payable at time T when the interest rate is r is given by (see Vasicek (1977)):

$$\begin{aligned}
 B(r, T) = \exp \left[-r \frac{1 - e^{-aT}}{a} + \frac{m^*((1 - e^{-aT}) - aT)}{a} \right. \\
 \left. - \frac{\sigma_3^2(4(1 - e^{-aT}) - (1 - e^{-2aT}) - 2aT)}{4a^3} \right]. \quad (29)
 \end{aligned}$$

To obtain the present value of a forward commitment to deliver one unit of the commodity, $P(S, \delta, r, T)$, we need to solve a partial differential equation and boundary conditions identical to equation (25) except that in the right-hand side we have rP , instead of zero. The solution to this modified equation is

$$P(S, \delta, r, T) = S \exp \left[\frac{-\delta(1 - e^{-\kappa T})}{\kappa} + D(T) \right] \quad (30)$$

where

$$\begin{aligned}
 D(T) = & \frac{(\kappa \hat{\alpha} + \sigma_1 \sigma_2 \rho_1)((1 - e^{-\kappa T}) - \kappa T)}{\kappa^2} \\
 & - \frac{\sigma_2^2(4(1 - e^{-\kappa T}) - (1 - e^{-2\kappa T}) - 2\kappa T)}{4\kappa^3}. \quad (31)
 \end{aligned}$$

Given the present value of a forward commitment in equation (30) and the present value of a unit discount bond in equation (29), the forward price

implied by Model 3 can be easily obtained by dividing $P(S, \delta, r, T)$ by $B(r, T)$. Note that the present value of a forward commitment in Model 3 is independent of the interest rate r and is identical to the corresponding one in Model 2. Forward prices in both models, however, are different.

II. Empirical Models

As mentioned in the introduction, one of the difficulties in the empirical implementation of commodity price models is that frequently the factors or state variables of these models are not directly observable. For some commodities the spot price is hard to obtain, and the futures contract closest to maturity is used as a proxy for the spot price. The problems of estimating the instantaneous convenience yield are even more complex; normally two futures prices with different maturities are used to compute the convenience yield.¹³ The instantaneous interest rate is also not directly observable. Futures contracts, however, are widely traded in several exchanges and their prices are more easily observed.

The state space form is the appropriate procedure to deal with situations in which the state variables are not observable, but are known to be generated by a Markov process. Once a model has been cast in state space form, the Kalman filter may be applied to estimate the parameters of the model and the time series of the unobservable state variables.

The general state space form applies to a multivariate time series of observable variables, in this case futures prices for different maturities, related to an unobservable vector of state variables (state vector), in this case the spot price alone or both the spot price and the instantaneous convenience yield, via a *measurement equation*. In our context the measurement equations are obtained from equations (8), (19), and (27) for the one, two, and three-factor models, respectively, by adding serially and cross-sectionally uncorrelated disturbances with mean zero to take into account bid-ask spreads, price limits, nonsimultaneity of the observations, errors in the data, etc. This simple structure for the measurement errors is imposed so that the serial correlation and cross correlation in the log prices is attributed to the variation of the unobservable state variables. The unobservable state variables are generated via the *transition equation*, which in our context is a discrete time version of the stochastic process for the state variables: equation (2) for the one-factor model and equations (11) to (13) for the two and three-factor models.¹⁴ The Kalman filter is a recursive procedure for computing the optimal estimator of the state vector at time t , based on the information available at time t , and it enables the estimate of the state vector to be continuously updated as new information becomes available. When the disturbances and the initial state

¹³ See for example Gibson and Schwartz (1990).

¹⁴ As we shall explain later, we estimate a simplified version of Model 3 in which the interest rate process is estimated separately.

vector are normally distributed the Kalman filter enables the likelihood function to be calculated, which allows for the estimation of any unknown parameters of the model and provides the basis for statistical testing and model specification. For a detailed discussion of state space models and the Kalman filter see Chapter 3 in Harvey (1989).

A. Model 1

From equation (8) the measurement equation can be written as:

$$y_t = d_t + Z_t X_t + \varepsilon_t, \quad t = 1, \dots, NT$$

where

$$\begin{aligned} y_t &= [\ln F(T_i)], \quad i = 1, \dots, N, \quad N \times 1 \text{ vector of observables} \\ d_t &= \left[(1 - e^{-\kappa T_i}) \alpha^* + \frac{\sigma^2}{4\kappa} (1 - e^{-2\kappa T_i}) \right], \quad i = 1, \dots, N, \quad N \times 1 \text{ vector} \\ Z_t &= [e^{-\kappa T_i}], \quad i = 1, \dots, N, \quad N \times 1 \text{ vector} \\ \varepsilon_t, \quad &N \times 1 \text{ vector of serially uncorrelated disturbances with} \\ &E(\varepsilon_t) = 0, \quad \text{Var}(\varepsilon_t) = H, \end{aligned} \quad (32)$$

and from equation (2) the transition equation can be written as:¹⁵

$$X_t = c_t + Q_t X_{t-1} + \eta_t, \quad t = 1, \dots, NT$$

where

$$\begin{aligned} c_t &= \kappa \alpha \Delta t \quad Q_t = 1 - \kappa \Delta t \\ \eta_t, \quad &\text{serially uncorrelated disturbances with} \\ &E(\eta_t) = 0, \quad \text{Var}(\eta_t) = \sigma^2 \Delta t. \end{aligned} \quad (33)$$

B. Model 2

From equation (19) the measurement equation can be written as:

$$y_t = d_t + Z_t [X_t, \delta_t]' + \varepsilon_t, \quad t = 1, \dots, NT$$

¹⁵ The exact transition equation is:

$$X_t = \alpha(1 - e^{-\kappa \Delta t}) + e^{-\kappa \Delta t} X_{t-1} + \eta_t$$

Using weekly data the linear approximation gives the identical parameters estimates up to the fourth significant figure and has been used in all the estimations.

where

$$\begin{aligned}
 y_t &= [\ln F(T_i)], \quad i = 1, \dots, N, \quad N \times 1 \text{ vector of observables} \\
 d_t &= [A(T_i)], \quad i = 1, \dots, N, \quad N \times 1 \text{ vector} \\
 Z_t &= \left[1, -\frac{1 - e^{-\kappa T_i}}{\kappa} \right], \quad i = 1, \dots, N, \quad N \times 2 \text{ matrix} \\
 \varepsilon_t, \quad N \times 1 \text{ vector of serially uncorrelated disturbances with} \\
 E(\varepsilon_t) &= 0, \quad \text{Var}(\varepsilon_t) = H,
 \end{aligned} \tag{34}$$

and from equations (11) to (13) the transition equation can be written as:

$$[X_t, \delta_t]' = c_t + Q_t[X_{t-1}, \delta_{t-1}]' + \eta_t, \quad t = 1, \dots, NT$$

where

$$\begin{aligned}
 c_t &= [(\mu - \frac{1}{2}\sigma_1^2)\Delta t, \kappa\alpha\Delta t]', \quad 2 \times 1 \text{ vector} \\
 Q_t &= \begin{vmatrix} 1 & -\Delta t \\ 0 & 1 - \kappa\Delta t \end{vmatrix} \\
 \eta_t, \quad &\text{serially uncorrelated disturbances with} \\
 E(\eta_t) &= 0, \quad \text{Var}(\eta_t) = \begin{vmatrix} \sigma_1^2\Delta t & \rho\sigma_1\sigma_2\Delta t \\ \rho\sigma_1\sigma_2\Delta t & \sigma_2^2\Delta t \end{vmatrix}.
 \end{aligned} \tag{35}$$

C. Model 3

We estimate a simplified version of Model 3. Ideally, the commodity spot price process, the convenience yield process, and the interest rate process should be estimated simultaneously from a time series and cross-section of futures prices and discount bond prices. To simplify the estimation we first estimate the parameters of the interest rate process, and then we use Model 3 to estimate the parameters of the spot price and convenience yield processes. We are essentially assuming that the parameters of the interest rate process are not affected by commodity futures prices, which seems to be a reasonable assumption.

Once we have estimated the interest rate process, we have only to estimate the parameters and state variables from the spot price and the convenience yield processes. From equations (27) and (28) the measurement equation is then:

$$y_t = d_t + Z_t[X_t, \delta_t]' + \varepsilon_t, \quad t = 1, \dots, NT$$

Table I
Oil Data

Futures Contract	Mean Price (Standard Error)	Mean Maturity (Standard Error)
Panel A: From 1/2/85 to 2/17/95: 510 Weekly Observations		
F1	\$19.99 (4.52)	0.043 (0.024) years
F3	19.65 (4.08)	0.210 (0.025)
F5	19.45 (3.74)	0.377 (0.024)
F7	19.31 (3.51)	0.543 (0.024)
F9	19.21 (3.35)	0.710 (0.025)
Panel B: From 1/2/90 to 2/17/95: 259 Weekly Observations		
F1	\$20.41 (4.13)	0.043 (0.024) years
F3	20.26 (3.54)	0.210 (0.025)
F5	20.09 (3.02)	0.376 (0.025)
F7	19.94 (2.62)	0.543 (0.025)
F9	19.84 (2.32)	0.709 (0.025)
Panel C: From 1/2/90 to 2/17/95: 259 Weekly Observations		
F1	\$20.41 (4.13)	0.043 (0.024) years
F5	20.09 (3.02)	0.376 (0.025)
F9	19.84 (2.32)	0.709 (0.025)
F13	19.76 (1.95)	1.041 (0.025)
F17	19.76 (1.74)	1.374 (0.025)
Forward Maturity	Mean Price (Standard Error)	Mean Maturity (Standard Error)
Panel D: From 1/15/93 to 5/16/96: 163 Weekly Observations (Enron Data)		
2 Months	18.16 (1.54)	0.122 (0.024)
5 Months	18.00 (1.31)	0.372 (0.024)
8 Months	18.00 (1.23)	0.621 (0.024)
1 Year	18.05 (1.15)	0.955 (0.024)
1½ Years	18.20 (1.09)	1.457 (0.024)
2 Years	18.38 (1.03)	1.955 (0.024)
3 Years	18.81 (0.95)	2.955 (0.024)
5 Years	19.67 (0.87)	4.955 (0.024)
7 Years	20.34 (0.79)	6.955 (0.024)
9 Years	20.92 (0.71)	8.955 (0.024)

where

$$y_t = [\ln F(T_i)], \quad i = 1, \dots, N, \quad N \times 1 \text{ vector of observables}$$
$$d_t = \left[\frac{r_t(1 - e^{-aT_i})}{a} + C(T_i) \right], \quad i = 1, \dots, N, \quad N \times 1 \text{ vector}$$
$$Z_t = \left[1, -\frac{1 - e^{-\kappa T_i}}{\kappa} \right], \quad i = 1, \dots, N, \quad N \times 2 \text{ matrix}$$




Figure 1. The figure shows for each week the time to maturity for the five oil futures contracts used in the estimation, starting from 1/2/85 to 2/17/95. This pattern of time to maturity is representative of all the data used.

$$\varepsilon_t, \quad N \times 1 \text{ vector of serially uncorrelated disturbances with} \\ E(\varepsilon_t) = 0, \quad \text{Var}(\varepsilon_t) = H. \quad (36)$$

Since we are using the Kalman filter to estimate the same state variables as in Model 2, the transition equation for this model is also (35).

III. Data

The data used to test the models consist of weekly observations¹⁶ of futures prices for two commercial commodities, oil and copper, and one precious metal, gold. In every case five futures contracts (i.e., $N = 5$) were used in the estimation.¹⁷ For different commodities and different time periods, however, different specific futures contracts had to be used since they vary across commodities and through time for a particular commodity. The interest rate

¹⁶ The data obtained from Knight-Ridder Financial consists of daily observations. To approximately transform it into weekly data every fifth observation from the original data was used in the empirical tests.

¹⁷ Except for the Enron oil data, where 10 contracts were used in the estimation.

Table II
Copper Data

Panel A: From 7/29/88 to 6/13/95: 347 Weekly Observations		
Futures Contract	Mean Price (Standard Error)	Mean Maturity (Standard Error)
F1	110.04 (18.05) cents	0.109 (0.065) years
F3	105.45 (13.54)	0.504 (0.084)
F5	102.42 (10.95)	0.899 (0.065)
F7	100.46 (9.34)	1.299 (0.085)
F9	99.78 (8.79)	1.663 (0.111)

Table III
Gold Data

Futures Contract	Mean Price (Standard Error)	Mean Maturity (Standard Error)
Panel A: From 1/2/85 to 6/13/95: 527 Weekly Observations		
F1	\$379.27 (40.95)	0.084 (0.048) years
F3	386.12 (42.46)	0.417 (0.048)
F6	397.55 (44.96)	0.917 (0.048)
F9	409.95 (47.79)	1.416 (0.048)
F11	418.87 (49.79)	1.749 (0.049)
Panel B: From 11/21/90 to 6/13/95: 230 Weekly Observations		
F1	\$365.50 (19.57)	0.084 (0.048) years
F3	370.12 (20.56)	0.417 (0.048)
F6	378.02 (22.30)	0.917 (0.048)
F9	386.70 (24.11)	1.413 (0.048)
F11	393.03 (25.32)	1.745 (0.049)
Panel C: From 11/21/90 to 6/13/95: 230 Weekly Observations		
F1	\$365.50 (19.57)	0.084 (0.048) years
F5	375.31 (21.70)	0.750 (0.048)
F9	386.70 (24.11)	1.413 (0.048)
F13	403.02 (27.05)	2.237 (0.144)
F18	460.85 (33.93)	4.703 (0.145)

data consisted in yields on 3-Month Treasury Bills. These data was used in the models requiring variable interest rates.

The oil data used are described in Table I. From 1/2/85 to 2/17/95, complete data on the first nine contracts were available, so the first set of tests used contracts F1, F3, F5, F7, and F9 (see Panel A), where F1 is the contract closest to maturity, F2 is the second contract closest to maturity, and so on. Since the contracts have a fixed maturity date, the time to maturity changes as time progresses. Figure 1 shows that time to maturity remains within a narrow

Table IV
One Factor Model: Oil

Period	1/2/85 to 2/17/95	1/2/90 to 2/17/95	1/2/90 to 2/17/95	1/15/93 to 5/16/96
Contracts	F1, F3, F5, F7, F9	F1, F3, F5, F7, F9	F1, F5, F9, F13, F17	Enron Data
NOBS	510	259	259	163
κ	0.301 (0.005)	0.694 (0.010)	0.428 (0.008)	0.099 (0.003)
μ	3.093 (0.346)	3.037 (0.228)	2.991 (0.280)	2.857 (0.635)
σ	0.334 (0.005)	0.326 (0.008)	0.257 (0.007)	0.129 (0.007)
λ	-0.242 (0.346)	-0.072 (0.228)	0.002 (0.279)	-0.320 (0.636)
ξ_1	0.049 (0.003)	0.045 (0.005)	0.080 (0.006)	0.079 (0.012)
ξ_2	0.018 (0.001)	0.017 (0.002)	0.031 (0.004)	0.046 (0.033)
ξ_3	0	0	0.010 (0.001)	0.029 (0.025)
ξ_4	0.012 (0.002)	0.009 (0.002)	0	0.014 (0.005)
ξ_5	0.022 (0.003)	0.015 (0.003)	0.007 (0.001)	0
ξ_6				0.007 (0.001)
ξ_7				0.018 (0.003)
ξ_8				0.031 (0.015)
ξ_9				0.035 (0.029)
ξ_{10}				0.035 (0.019)
Log-likelihood function	8130	4369	4345	5146

*(Standard errors in parentheses)

NOBS = number of observations.




Figure 2. The figure shows for each week the Model 1 oil futures prediction errors for the five futures contracts used in the estimation, starting from 1/2/85 to 2/17/95.




Figure 3. The figure shows for each week the Model 1 estimated state variable (the logarithm of the spot price) and the logarithm of the oil futures price for the contract closest to maturity, starting from 1/2/85 to 2/17/95.

Table V
One Factor Model: Copper

Period	7/29/88 to 6/13/95
Contracts	F1, F3, F5, F7, F9
NOBS	347
κ	0.369 (0.009)
μ	4.854 (0.230)
σ	0.233 (0.007)
λ	-0.339 (0.230)
ξ_1	0.064 (0.002)
ξ_2	0.023 (0.001)
ξ_3	0
ξ_4	0.015 (0.001)
ξ_5	0.021 (0.002)
Log-likelihood function	5482

* (Standard errors in parentheses)

NOBS = number of observations.

range for each one of the contracts during the whole sample period. The figure is representative of the maturity structure for all the data used in this study. Starting in 1/2/90, complete data on 17 oil futures contracts were available extending the maximum maturity of the contracts from an average of 0.71

Table VI
Two Factor Model: Oil

Period	1/2/85 to 2/17/95	1/2/90 to 2/17/95	1/2/90 to 2/17/95	1/15/93 to 5/16/96
Contracts	F1, F3, F5, F7, F9	F1, F3, F5, F7, F9	F1, F5, F9, F13, F17	Enron Data
NOBS	510	259	259	163
μ	0.142 (0.125)	0.244 (0.150)	0.238 (0.160)	0.082 (0.120)
κ	1.876 (0.024)	1.829 (0.033)	1.488 (0.027)	1.187 (0.026)
α	0.106 (0.025)	0.184 (0.110)	0.180 (0.126)	0.090 (0.086)
σ_1	0.393 (0.007)	0.374 (0.011)	0.358 (0.010)	0.212 (0.011)
σ_2	0.527 (0.015)	0.556 (0.024)	0.426 (0.017)	0.187 (0.012)
ρ	0.766 (0.013)	0.882 (0.013)	0.922 (0.006)	0.845 (0.024)
λ	0.198 (0.166)	0.316 (0.203)	0.291 (0.190)	0.093 (0.101)
ξ_1	0.022 (0.001)	0.020 (0.001)	0.043 (0.002)	0.027 (0.001)
ξ_2	0.001 (0.001)	0	0.006 (0.001)	0.006 (0.001)
ξ_3	0.003 (0.001)	0.004 (0.000)	0.003 (0.000)	0
ξ_4	0	0	0	0.002 (0.000)
ξ_5	0.005 (0.000)	0.006 (0.000)	0.004 (0.000)	0
ξ_6				0.004 (0.000)
ξ_7				0.014 (0.003)
ξ_8				0.032 (0.015)
ξ_9				0.043 (0.036)
ξ_{10}				0.055 (0.039)
Log-likelihood function	10267	5256	5139	6182

* (Standard errors in parentheses)

NOBS = number of observations.

years to 1.34 years. Longer maturity contracts are of most interest in this study, since we will be concerned with investment decisions in real assets with much longer maturities. The second set of tests (see Panel C), then used contracts F1, F5, F9, F13, and F17. To be able to distinguish whether the possible differences between these two tests were due to the time period used or to the contracts used, a third set of tests were done using the contracts of the first set of tests over the period of the second set of tests (see Panel B).

The High Grade copper futures contract started trading in 1988, so complete data on the first nine contracts were available from 7/29/88 to 6/13/95 (see Table II). The only set of tests performed for copper used contracts F1, F3, F5, F7, and F9. In the case of copper, however, the last contract had an average maturity of 1.66 years (contract F17 for oil has an average maturity of only 1.34).

The gold data are shown in Table III. Complete data for the first eleven contracts were available from 1/2/85 to 6/13/95, so the first set of tests used contracts F1, F3, F6, F9, and F11 (see Panel A). The average maturity of the F11 contracts was 1.75 years. Starting on 11/21/90, complete data on the first 18 contracts were available so the second set of tests used contracts F1, F5, F9, F13, and F18, extending the average maturity of the longest contract to 4.70




Figure 4. The figure shows for each week the Model 2 estimated state variables (the logarithm of the spot price and the instantaneous convenience yield) and the logarithm of the oil futures price for the contract closest to maturity, starting from 1/2/85 to 2/17/95. A fixed amount of 2 has been added to each value of the convenience yield to make its scale comparable to the other variables shown.

years (see Panel C). As was done for the case of oil, a third set of tests was performed using the same contracts as in the first tests during the period of the second tests (see Panel B).

In addition to the publicly available futures data described above, for the purposes of this study Enron Capital and Trade Resources made available some proprietary historical crude oil forward price curves from 1/15/93 to 5/16/96. From these data ten forward prices were used in the estimation, ranging in maturities from two months to nine years (see Panel D of Table I). The great advantage of these data is the longer maturities of the contracts. The disadvantage is that we do not know exactly how the crude oil forward curves were constructed.

IV. Empirical Results

A. Model 1

Table IV presents the results for the one-factor model applied to the four data sets for oil described in Table I. In all cases the speed of adjustment coefficient is highly significant, and the market price of risk is not significantly different from zero. The main difference between Columns 2 and 3, which use the same contracts, is that the later period has much stronger mean reversion (0.7 as opposed to 0.3 for the whole period). The other parameters are very




Figure 5. The figure shows for each week the Model 2 oil futures prediction errors for the five futures contracts used in the estimation, starting from 1/2/85 to 2/17/95.

Table VII
Two Factor Model: Copper

Period	7/29/88 to 6/13/95
Contracts	F1, F3, F5, F7, F9
NOBS	347
μ	0.326 (0.110)
κ	1.156 (0.041)
α	0.248 (0.098)
σ_1	0.274 (0.012)
σ_2	0.280 (0.017)
ρ	0.818 (0.020)
λ	0.256 (0.114)
ξ_1	0.033 (0.001)
ξ_2	0.003 (0.001)
ξ_3	0.006 (0.000)
ξ_4	0.005 (0.000)
ξ_5	0.009 (0.001)
Log-likelihood function	6473

* (Standard errors in parentheses)

NOBS = number of observations.

similar both in magnitude and significance. In all cases, one of the standard deviations of the measurement errors goes to zero. This is a common phenomenon in this type of analysis.

Table VIII
Two Factor Model: Gold

Period	1/2/85 to 6/13/95	11/21/90 to 6/13/95	11/21/90 to 6/13/95
Contracts	F1, F3, F6, F9, F11	F1, F3, F6, F9, F11	F1, F5, F9, F13, F18
NOBS	527	230	230
μ	0.039 (0.044)	0.033 (0.054)	0.030 (0.054)
κ	0.011 (0.008)	0.114 (0.015)	0.298 (0.018)
α	-0.002 (0.322)	0.018 (0.052)	0.019 (0.023)
σ_1	0.135 (0.003)	0.106 (0.004)	0.107 (0.004)
σ_2	0.016 (0.001)	0.0124 (0.0007)	0.015 (0.001)
ρ	0.056 (0.034)	0.113 (0.066)	0.250 (0.068)
λ	0.0067 (0.0036)	0.0076 (0.0060)	0.008 (0.007)
ξ_1	0.003 (0.000)	0.002 (0.000)	0.004 (0.000)
ξ_2	0	0	0
ξ_3	0.001 (0.000)	0.001 (0.000)	0.001 (0.001)
ξ_4	0	0	0.001 (0.000)
ξ_5	0.001 (0.000)	0.001 (0.000)	0.012 (0.001)
Log-likelihood function	14437	6662	5660

* (Standard errors in parentheses)

NOBS = number of observations.

The prediction errors¹⁸ for the first oil data set are shown in Figure 2. Though the average error is quite small, 0.0032, the average absolute deviation is 0.033, which is around one percent of the log of the price of the futures contract closest to maturity, reflecting some very large deviations in the figure. Note also that there seems to be some negative autocorrelation of the errors, which could imply the existence of some errors in the data. Figure 3 presents the value of the state variable and the log of the first futures contract for the same data set. Here we can see that the state variable (the log of the spot price) follows closely, but is not identical to, the log of the price of the futures contract closest to maturity (F1).

Comparing now Columns 3 and 4 of Table IV, which use the same time period but different futures contracts, we see that the main effect of extending the maturity of the contracts is to reduce the mean reversion parameter (from 0.7 to 0.4). This can have important implications in the application of this model to long term oil investment, since in this case the relevant futures contracts would be much longer in maturity. Even though the Enron data covers a different time period (with some overlap), from Column 5 we can see that the mean reversion parameter is even smaller (0.1).

Table V shows the results of applying the one factor model to the copper data. These results are similar to those for oil. There is strong and significant mean reversion (of 0.37) and the market price of risk is positive, but not significantly different from zero.

¹⁸ Also called *innovations* since they represent the new information in the latest observation.

Table IX
Three Factor Model: Oil

Period	1/2/90 to 2/17/95	1/15/93 to 5/16/96
Contracts	F1, F5, F9, F13, F17	Enron Data
NOBS	259	163
μ	0.315 (0.125)	0.008 (0.109)
κ	1.314 (0.027)	0.976 (0.022)
α	0.249 (0.093)	0.038 (0.077)
σ_1	0.344 (0.009)	0.196 (0.009)
σ_2	0.372 (0.014)	0.145 (0.008)
ρ_1	0.915 (0.007)	0.809 (0.027)
λ	0.353 (0.123)	0.013 (0.075)
ξ_1	0.045 (0.002)	0.028 (0.001)
ξ_2	0.007 (0.001)	0.006 (0.001)
ξ_3	0.003 (0.000)	0
ξ_4	0	0.002 (0.000)
ξ_5	0.004 (0.000)	0.000 (0.001)
ξ_6		0.005 (0.000)
ξ_7		0.013 (0.002)
ξ_8		0.024 (0.008)
ξ_9		0.032 (0.014)
ξ_{10}		0.053 (0.023)
Log-likelihood function	5128	6287
σ_3	0.0081	0.0073
a	0.2	0.2
$R(\infty)$	0.07	0.07
ρ_2	-0.0039	0.0399
ρ_3	-0.0293	-0.0057

* (Standard errors in parentheses)

NOBS = number of observations.

The one-factor model could not be fitted to the gold data, giving us a first indication that there is no detectable mean reversion in gold prices for the period considered.

B. Model 2

The risk-free rate of interest, r , which is assumed constant, is a parameter in Model 2. It appears in the measurement equation (23) through $A(T)$ defined in equation (20). For purposes of estimating this two-factor model, we assumed a constant risk-free interest rate of 0.06, which was approximately the average interest rate over the period considered.¹⁹ The risk-free rate enters into the analysis through the risk-adjusted process for the spot price described in equation (14). From this equation it can be seen that any variation in the interest rate through time will be absorbed by variations in the convenience

¹⁹ For the Enron oil data we used an interest rate of 0.05, since in this latter period interest rates were lower.

Table X
Three Factor Model: Copper

Period	7/29/88 to 6/13/95
Contracts	F1, F3, F5, F7, F9
NOBS	347
μ	0.332 (0.094)
κ	1.045 (0.030)
α	0.255 (0.078)
σ_1	0.266 (0.011)
σ_2	0.249 (0.014)
ρ_1	0.805 (0.022)
λ	0.243 (0.082)
ξ_1	0.032 (0.001)
ξ_2	0.004 (0.001)
ξ_3	0.005 (0.000)
ξ_4	0.005 (0.000)
ξ_5	0.007 (0.000)
Log-likelihood function	6520
σ_3	0.0096
a	0.2
$R(\infty)$	0.07
ρ_2	0.1243
ρ_3	0.0964

* (Standard errors in parentheses)

NOBS = number of observations.

yield. The estimated instantaneous convenience yield will then be a composite of the actual convenience yield and the deviations of the interest rate from 0.06.

Table VI reports the results of the two-factor model for the four oil data sets. In all cases the speed of adjustment coefficient in the convenience yield equation and the correlation coefficient between the two factors are large and highly significant; the total expected return on the spot commodity (μ), the average convenience yield, and the market price of convenience yield risk are positive but not always significant at standard levels. It is interesting to note that the estimations that use longer term futures have a somewhat lower mean reversion coefficient (1.2 for the Enron data versus 1.5 for the longer term data and 1.8 for the shorter term oil data). Figure 4 shows the two state variables for the first data set, which covers all the sample period²⁰ and the log price of the futures contract closest to maturity. The strong correlation between the two state variables (0.77) and the closeness between the log spot price and the log of F1 can be observed from the figure. Figure 5 displays the prediction errors. Both the mean error of 0.0016 and the average absolute deviation of 0.029 are smaller than those in Model 1 (see Figure 2).

²⁰ A fixed amount of 2 has been added to each value of the convenience yield to make its scale more comparable with the log of the spot price.

Table XI
Three Factor Model: Gold

Period	11/21/90 to 6/13/95
Contracts	F1, F5, F9, F13, F18
NOBS	230
μ	0.023 (0.054)
κ	0.023 (0.023)
α	0.021 (0.189)
σ_1	0.106 (0.004)
σ_2	0.009 (0.001)
ρ_1	0.208 (0.069)
λ	0.002 (0.004)
ξ_1	0.003 (0.000)
ξ_2	0
ξ_3	0.001 (0.000)
ξ_4	0.001 (0.000)
ξ_5	0.015 (0.001)
Log-likelihood function	5688
σ_3	0.0082
a	0.2
$R^{(\infty)}$	0.07
ρ_2	-0.4005
ρ_3	-0.0260


* (Standard errors in parentheses)

NOBS = number of observations.


The results for copper, presented in Table VII, also show very strong and significant correlation between the state variables and mean reversion in the convenience yield. The total expected return on copper, the average convenience yield, and the market price of risk are also positive and significant, but not as strongly as the other parameters. Note that the average convenience yield is high (25 percent per year) because we are estimating instantaneous convenience yields, that is for futures contracts maturing in the next instant of time.

The results for gold, displayed in Table VIII, differ importantly from those for the commercial commodities. The mean reversion in the convenience yield and the correlation between the state variables are significant only in the later period and of much smaller magnitude, 0.3 and 0.25, respectively, for the longer term data. The total expected return, average convenience yield, and market price of risk are insignificant in all cases. As we shall see when we discuss the three-factor model, even the mean reversion in the convenience yield becomes insignificant when stochastic interest rates are considered, which suggests that in the two-factor model, mean reversion in convenience yield is proxying for mean reversion in interest rates (which are assumed constant in this model).

A



B



C




Figure 6. (a) The figure shows the term structure of oil futures prices for 1/2/90, the starting date for the long term oil data, and the term structure implied by the three models. (b) The figure shows the term structure of oil futures prices for 12/22/93, a date on which oil futures prices were in contango, and the term structure implied by the three models. (c) The figure shows the term structure of oil futures prices for 2/17/95, the last date for the oil data, and the term structure implied by the three models.

Table XII

Cross-Section Comparison Between Models 1, 2, and 3 Out of Sample Oil Data 1/2/90 to 2/17/95

Model Contract	RMSE			Mean Error		
	1	2	3	1	2	3
Panel A: In Dollars						
F2	1.537	0.538	0.577	0.453	-0.057	-0.032
F3	1.215	0.325	0.363	0.376	-0.019	-0.003
F4	0.952	0.183	0.22	0.3	0.004	0.013
F6	0.55	0.061	0.074	0.163	0.01	0.011
F7	0.41	0.058	0.056	0.108	0.006	0.005
F8	0.296	0.054	0.044	0.063	0	-0.001
F10	0.138	0.045	0.037	0.009	-0.005	-0.006
F11	0.082	0.035	0.03	-0.001	-0.003	-0.004
F12	0.039	0.02	0.019	-0.003	-0.001	-0.002
F14	0.039	0.02	0.019	0.009	0.001	0.001
F15	0.075	0.036	0.034	0.021	0.001	0.001
F16	0.109	0.053	0.051	0.037	0.001	0
All	0.668	0.193	0.21	0.128	-0.005	-0.001
Panel B: In Percentage						
F2	6.203	2.574	2.795	1.566	-0.273	-0.188
F3	4.94	1.546	1.737	1.386	-0.091	-0.035
F4	3.92	0.862	1.032	1.159	0.012	0.045
F6	2.369	0.272	0.329	0.671	0.042	0.046
F7	1.814	0.253	0.238	0.45	0.021	0.018
F8	1.35	0.256	0.204	0.268	-0.003	-0.008
F10	0.664	0.218	0.179	0.04	-0.024	-0.03
F11	0.406	0.169	0.146	-0.002	-0.016	-0.02
F12	0.194	0.099	0.089	-0.013	-0.006	-0.008
F14	0.192	0.1	0.093	0.038	0.005	0.005
F15	0.366	0.18	0.171	0.095	0.007	0.006
F16	0.536	0.265	0.257	0.171	0.007	0.004
All	2.74	0.92	1.011	0.486	-0.027	-0.014

C. Model 3

The parameters for the risk-adjusted interest rate process (23) were obtained using 3-month Treasury Bill yields. The standard deviation of changes in r , σ_3 , for each data set were computed using contemporaneous yields. As reported in Tables IX, X, and XI, this standard deviation varied from 0.0073 to 0.0096 for different time periods. The speed of adjustment coefficient of the




Figure 7. (a) The figure shows the term structure of copper futures prices for 7/29/88, the starting date for the copper data, and the term structure implied by the three models. (b) The figure shows the term structure of copper futures prices for 7/13/93, a date on which copper futures prices were in contango, and the term structure implied by the three models. (c) The figure shows the term structure of copper futures prices for 6/13/95, the last date for the copper data, and the term structure implied by the three models.

Table XIII
Cross-Section Comparison Between Models 1, 2, and 3 Out of Sample
Copper Data 7/29/88 to 6/13/95

Model Contract	RMSE			Mean Error		
	1	2	3	1	2	3
Panel A: In Cents						
F2	4.617	1.311	1.27	1.75	-0.155	-0.142
F4	1.17	0.548	0.539	0.327	0.116	0.119
F6	1.133	0.616	0.621	-0.194	-0.105	-0.098
F8	1.824	0.534	0.482	0.257	-0.042	-0.064
All	2.612	0.819	0.794	0.535	-0.046	-0.046
Panel B: In Percentage						
F2	3.929	1.103	1.073	1.404	-0.151	-0.135
F4	1.063	0.483	0.474	0.284	0.093	0.094
F6	1.064	0.568	0.572	-0.188	-0.102	-0.097
F8	1.78	0.512	0.462	0.213	-0.036	-0.058
All	2.284	0.713	0.692	0.428	-0.049	-0.049


process, α , was assumed to be equal to 0.2, which implies that one half of any deviation from the average interest rate is expected to be corrected in 3.5 years.²¹ The risk-adjusted drift of the process, m^* , was computed so that the infinite maturity discount yield be 7 percent. In the Vasicek model this infinite maturity yield is:

$$R(\infty) = m^* - \frac{\sigma_3^2}{2\alpha^2}. \quad (37)$$


The instantaneous correlations between the interest rate process and the process for the log spot price and the convenience yield defined in equation (24) were approximated by the correlations obtained using weekly data between the three-month Treasury Bill yield and the values of these state variables obtained from Model 2. The estimated value of these correlations are shown in

²¹ The mean reversion coefficient can be obtained by running a regression of changes in interest rate on lagged interest rates. Since there is a lot of measurement error on this parameter, we took an average value.

A



B



C




Figure 8. (a) The figure shows the term structure of oil futures prices for 1/15/93, the starting date for the Enron oil data, and the term structure implied by the three models. (b) The figure shows the term structure of oil futures prices for 2/15/94, a date on which the Enron oil futures prices were in contango, and the term structure implied by the three models. (c) The figure shows the term structure of oil futures prices for 5/16/96, the last date for the Enron oil data, and the term structure implied by the three models.

Table XIV

**Cross-Section Comparison Between Models 1, 2, and 3 Out of Sample
Enron Oil Data 1/15/93 to 5/16/96**

Model Contract	RMSE			Mean Error		
	1	2	3	1	2	3
Panel A: In Dollars						
1 Month	1.86	0.905	0.942	0.839	0.332	0.427
4 Months	0.934	0.179	0.181	0.413	0.05	0.089
7 Months	0.588	0.028	0.028	0.254	0.005	0.012
10 Months	0.358	0.034	0.033	0.153	-0.005	-0.014
15 Months	0.108	0.033	0.031	0.043	-0.004	-0.013
21 Months	0.078	0.046	0.043	-0.027	0.009	0.021
2½ Years	0.245	0.183	0.162	-0.058	0.039	0.093
4 Years	0.505	0.493	0.439	0.095	0.185	0.283
6 Years	0.713	0.802	0.575	0.191	0.092	0.135
8 Years	0.756	1.027	0.792	0.221	-0.238	-0.395
All	0.789	0.532	0.459	0.212	0.046	0.064
Panel B: In Percentage						
1 Month	9.428	4.305	4.46	3.958	1.701	2.182
4 Months	5.247	0.967	0.987	2.113	0.303	0.506
7 Months	3.347	0.155	0.156	1.335	0.034	0.068
10 Months	2.036	0.184	0.181	0.819	-0.034	-0.079
15 Months	0.611	0.177	0.172	0.235	-0.03	-0.077
21 Months	0.435	0.249	0.242	-0.148	0.057	0.12
2½ Years	1.316	0.962	0.885	-0.31	0.242	0.522
4 Years	2.584	2.506	2.257	0.461	0.997	1.483
6 Years	3.499	3.925	2.813	0.883	0.47	0.666
8 Years	3.586	4.944	3.827	0.996	-1.154	-1.933
All	4.072	2.582	2.224	1.034	0.259	0.346

Tables IX, X, and XI. Apart from the correlation between the interest rate and the convenience yield for copper (0.12) and gold (0.4), the other correlations were very close to zero.²²

²² Since the correlation between the interest rate process and the commodity spot process is zero, futures prices are equal to forward prices (see Cox, Ingersoll, and Ross (1981)); in the estimation we treated the forward prices in the Enron data as futures prices.

Table XV
Time Series Comparison Between Models 1, 2, and 3: Last 50
Observations of Oil Data

Model	RMSE (of Log Prices)			Mean Error (of Log Prices)		
	1	2	3	1	2	3
Panel A: In-Sample Parameter Estimation						
F1	0.0830	0.0538	0.0540	0.0628	0.0291	0.0279
F5	0.0390	0.0240	0.0248	0.0260	0.0049	0.0040
F9	0.0230	0.0195	0.0194	0.0096	0.0003	-0.0007
F13	0.0171	0.0170	0.0170	0.0012	0.0012	0.0001
F17	0.0157	0.0161	0.0160	-0.0033	0.0034	0.0016
All	0.0435	0.0300	0.0299	0.0193	0.0077	0.0066
Panel B: Out-of-Sample Parameter Estimation						
F1	0.0919	0.0551	0.0541	0.0770	0.0315	0.0289
F5	0.0378	0.0249	0.0246	0.0273	0.0045	0.0030
F9	0.0198	0.0195	0.0195	0.0020	-0.0004	-0.0017
F13	0.0210	0.0170	0.0170	-0.0125	0.0014	0.0002
F17	0.0262	0.0166	0.0162	-0.0212	0.0050	0.0028
All	0.0477	0.0303	0.0299	0.0145	0.0084	0.0066

Tables IX and X present the parameter estimates for Model 3 for oil and copper, respectively.²³ For both oil and copper these results are quite close to those of Model 2. The parameters of the processes for the spot price and convenience yield seem to be robust to the specification of the interest rate process. This does not mean, however, that the value of a futures contract is insensitive to the interest rate used in the computation.

Table XI reports the parameter estimates of Model 3 for gold. The fact that the mean reversion in the convenience yield becomes insignificant when stochastic interest rates are included and that the correlation between changes in the interest rate and the convenience yield is so high in absolute terms (-0.4) suggests that the models are misspecified for gold. Mean reversion in prices induced by a mean-reverting convenience yield does not seem to hold for gold.

D. Comparing the Three Models

For the purpose of comparing the relative performance of the three models, we will concentrate on three data sets: the long term oil futures data, the copper futures data, and the Enron oil forward data. Figures 6, 7, and 8 illustrate how the three models fit the data for three dates in the sample period. These dates were the first observation in the sample, the last observation in the sample, and an intermediate observation chosen so that futures

²³ The results for the shorter maturity data sets are similar. For this reason, in what follows we report only results for the longer term data.

Table XVI
Time Series Comparison Between Models 1, 2, and 3: Last 50
Observations of Copper Data

Model	RMSE (of Log Prices)			Mean Error (of Log Prices)		
	1	2	3	1	2	3
Panel A: In-Sample Parameter Estimation						
F1	0.0453	0.0430	0.0412	0.0046	-0.0253	-0.0232
F3	0.0294	0.0216	0.0217	0.0061	0.0029	0.0024
F5	0.0192	0.0194	0.0190	0.0027	0.0070	0.0062
F7	0.0207	0.0180	0.0180	-0.0000	-0.00126	-0.0016
F9	0.0255	0.0187	0.0186	0.0140	-0.00021	-0.0003
All	0.0296	0.0260	0.0253	0.0055	-0.0034	-0.0033
Panel B: Out-of-Sample Parameter Estimation						
F1	0.0455	0.0467	0.0434	0.0058	-0.0301	-0.0264
F3	0.0295	0.0218	0.0216	0.0068	0.0030	0.0023
F5	0.0193	0.0199	0.0193	0.0031	0.0080	0.0068
F7	0.0207	0.0182	0.0181	0.0001	-0.0013	-0.0015
F9	0.0256	0.0188	0.0188	0.0141	-0.0018	-0.0012
All	0.0297	0.0273	0.0261	0.0060	-0.0044	-0.0040

prices were in contango, since the first and last observations were in backwardation.²⁴ Figures 6a, 6b, and 6c correspond to the long term oil data, Figures 7a, 7b, and 7c to the copper data, and Figures 8a, 8b, and 8c to the Enron oil data. From these figures we can observe that Model 1 is very often incapable of adequately describing the data. From equation (7) we can see that when the maturity of the futures contract tends to infinity, in Model 1 the futures price converges to:

$$F(S, \infty) = \exp \left[\alpha^* + \frac{\sigma^2}{4\kappa} \right] \quad (38)$$

which is independent of the spot price. For the estimated parameters of Model 1 the infinite maturity futures prices are 20.13 dollars for the long-term oil data, 22.99 dollars for the Enron oil data, and 88.08 cents for the copper data. When the spot price is above these infinite maturity futures price, Model 1 will be in backwardation and when the spot price is below it will be in contango. This feature does not allow for much flexibility in the term structure of futures prices.

From Figures 6 and 7 we can also observe that the term structure of futures prices implied by Models 2 and 3 for the long-term oil data and the copper data are sometimes indistinguishable and that they are always very close to each other. This holds true for all the in-sample term structures. As we shall see in

²⁴ The illustrations are representative of the rest of the observations in the sample.

Table XVII
Time Series Comparison Between Models 1, 2, and 3: Last 50
Observations of Enron Oil Data

Model	RMSE (of Log Prices)			Mean Error (of Log Prices)		
	1	2	3	1	2	3
Panel A: In-Sample Parameter Estimation						
2 Months	0.1184	0.0480	0.0522	0.1024	0.0287	0.0340
5 Months	0.0646	0.0213	0.0220	0.0563	0.0047	0.0071
8 Months	0.0422	0.0173	0.0173	0.0353	0.0010	0.0019
1 Year	0.0239	0.0149	0.0149	0.0168	-0.0001	0.0002
1½ Years	0.0129	0.0131	0.0132	-0.0001	0.0006	0.0013
2 Years	0.0151	0.0144	0.0152	-0.0075	0.0040	0.0056
3 Years	0.0201	0.0155	0.0173	-0.0145	0.0080	0.0101
5 Years	0.0270	0.0193	0.0160	-0.0154	0.0103	0.0065
7 Years	0.0280	0.0215	0.0183	-0.0086	0.0085	-0.0077
9 Years	0.0221	0.0177	0.0358	-0.0002	0.0009	-0.0318
All	0.0483	0.0225	0.0251	0.0164	0.0067	0.0027
Panel B: Out-of-Sample Parameter Estimation						
2 Months	0.1391	0.0540	0.0622	0.1258	0.0364	0.0450
5 Months	0.0810	0.0221	0.0242	0.0744	0.0075	0.0117
8 Months	0.0534	0.0172	0.0172	0.0485	0.0012	0.0021
1 Year	0.0297	0.0150	0.0149	0.0241	-0.0010	-0.0031
1½ Years	0.0128	0.0131	0.0132	-0.0002	0.0008	-0.0034
2 Years	0.0186	0.0157	0.0126	-0.0135	0.0067	0.0017
3 Years	0.0319	0.0215	0.0165	-0.0288	0.0163	0.0110
5 Years	0.0429	0.0301	0.0225	-0.0360	0.0263	0.0183
7 Years	0.0392	0.0338	0.0196	-0.0274	0.0297	0.0133
9 Years	0.0267	0.0303	0.0142	-0.0131	0.0264	-0.0025
All	0.0594	0.0279	0.0258	0.0154	0.0150	0.0094

Section VI, this does not mean that both models have the same implications for futures prices with longer maturities than the observable ones. For the Enron oil data, however, Models 2 and 3 can imply quite different term structure of futures prices, as shown in Figure 8.

The empirical models we have estimated are not nested, so a statistical comparison between them is not straightforward.²⁵ Model 1 is not nested in Model 2 nor 3. Even though Model 2 is nested in Model 3, we estimate the same number of parameters in both models since in Model 3 we take as exogenous the parameters of the interest rate process. This complicates the interpretation of the likelihood function value.

To compare the relative performance of the models we performed two types of tests. The first were cross-section tests using all the available futures prices

²⁵ There are procedures for testing nonnested models. See, for example, Cox (1961, 1962), Atkinson (1970), Pesaran and Deaton (1978), and Davidson and MacKinnon (1981). These procedures, however, are not easily adapted to our framework.

Table XVIII
Volatilities of Futures Returns Implied by Models 1, 2, and 3


Model	Oil Futures		Copper Futures		Enron Oil Forwards	
	Zero Maturity	Infinite Maturity	Zero Maturity	Infinite Maturity	Zero Maturity	Infinite Maturity
1	0.257	0	0.233	0	0.129	0
2	0.358	0.145	0.274	0.159	0.212	0.159
3	0.344	0.146	0.266	0.166	0.196	0.166

not used in the estimation of the parameters of the models. Recall that only five futures prices (and ten for the Enron data) were used to estimate the parameters of the models. The second were time series tests using the last 50 observations in the sample.


In the cross-section tests we used the parameters and the value of the state variables estimated over the whole sample period to price at every observation date all the futures contracts that were not used in the estimation. In some sense, this is an out-of-sample test because we are only pricing contracts not used in the estimation. The parameters used, however, have been estimated over the whole sample period so in strict sense only the last observation is truly out-of-sample. Also, prices of futures contracts are correlated. The procedure, anyhow, should give us an indication of the relative performance of the models. Tables XII, XIII, and XIV show the results for the long term oil, copper, and Enron oil data, respectively. For oil we had 12 futures prices not used in the estimation, but for copper we only had 4. For the Enron oil data we used 10 forward prices not used in the estimation. We report the root mean square error and the mean error in monetary terms in Panel A of the tables and in percentage in Panel B for the contracts available and also for all the contracts together. Models 2 and 3 clearly outperform Model 1 in every dimension and for the three data sets. For example, for the oil data the RMSE is 2.7 percent for Model 1, whereas it is around 1 percent for Models 2 and 3. The relative performance of Models 2 and 3 is not so clear when we consider only the shorter maturity futures data, Model 2 outperforms Model 3 for oil, and vice-versa for copper. For the futures contracts also the copper data fit all the models better than the oil data. This is possibly due to the fact that the oil data period includes the Gulf War in August of 1990, which had a dramatic impact in the oil markets. For the longer maturity Enron oil data, however, Model 3 appears as the clear winner: the RMSE is 4.1 percent for Model 1, 2.6 percent for Model 2, and 2.2 percent for Model 3. Note also that for this data, volatilities are smaller since it does not include the Gulf War period.

A truly out-of-sample time series test of the models would compute the prediction errors for period $t + 1$ using all the information available up to period t . This would require the estimation of the new parameters of the model at every period t . Instead, we computed an upper and lower bound of these prediction errors in the following manner. First, we computed the prediction

A



B



C




Figure 9. (a) The figure shows the term structure of volatilities of oil futures returns and the term structure implied by the three models. (b) The figure shows the term structure of volatilities of copper futures returns and the term structure implied by the three models. (c) The figure shows the term structure of volatilities of oil futures returns and the term structure implied by the three models, using the parameters for the Enron oil data.

errors for the last 50 observations of the sample using the parameters estimates for the whole sample period. This is a lower bound of the error, since we are using the last observations also to estimate the parameters. Second, we estimated the parameters over the period that did not include the last 50 observations and used them to estimate the prediction errors for the last 50 observations of the sample. This is an upper bound of the error since we did not update the parameters of the models as new data became known. Tables XV, XVI, and XVII present the results of these tests for oil, copper, and Enron oil data, respectively. Panel A of this table gives the root mean square errors and the mean errors of the log prices²⁶ using the in-sample estimation of parameters, whereas Panel B gives the same information but using the out-of-sample parameter estimates. The period over which the parameters are estimated does not make a significant difference in the results, justifying the procedure used. As before, Models 2 and 3 clearly outperform Model 1. Model 3 marginally outperforms Model 2 for the oil and copper data. In the out-of-sample parameter estimation, Model 3 outperforms Model 2 for the Enron oil data, but the reverse occurs for the in-sample parameter estimation.

V. Volatility of Futures Returns

When we use a model to value a financial or real asset contingent on a commodity price, we are interested in modeling not only the term structure of futures prices, but also the term structure of volatilities. Each model considered has different implications for the term structure of the volatilities of commodity futures returns. One property that will be common to the three models is that volatilities will be independent of the state variables of each model and will only depend on time to maturity of the futures contracts.


Applying Ito's Lemma to equation (7) we see that the term structure of proportional futures volatilities in Model 1 is given by

$$\sigma_F^2(T) = \sigma^2 e^{-2\kappa T}. \quad (39)$$


A feature of this model is that as the time to maturity of the futures contract approaches infinity, the volatility of its price converges to zero.

²⁶ The reason to deal with log prices comes from their use in the Kalman filter estimation.

A



B



C




Figure 10. (a) The figure shows the term structures of oil futures prices for 1/2/90 implied by the three models for maturities up to ten years. (b) The figure shows the term structure of oil futures prices for 12/22/93 implied by the three models for maturities up to ten years. (c) The figure shows the term structure of oil futures prices for 2/17/95 implied by the three models for maturities up to ten years.

From equation (18) we can obtain the volatility of futures returns in Model 2:

$$\sigma_F^2(T) = \sigma_1^2 + \sigma_2^2 \frac{(1 - e^{-\kappa T})^2}{\kappa^2} - 2\rho\sigma_1\sigma_2 \frac{(1 - e^{-\kappa T})}{\kappa}. \quad (40)$$

When time to maturity approaches infinity, the volatility of futures returns in this model converges to:

$$\sigma_F^2(\infty) = \sigma_1^2 + \frac{\sigma_2^2}{\kappa^2} - \frac{2\rho\sigma_1\sigma_2}{\kappa}. \quad (41)$$

Finally, the volatility in Model 3 can be obtained from equation (26):

$$\begin{aligned} \sigma_F^2(T) = & \sigma_1^2 + \sigma_2^2 \frac{(1 - e^{-\kappa T})^2}{\kappa^2} + \sigma_3^2 \frac{(1 - e^{-aT})^2}{a^2} - 2\rho_1\sigma_1\sigma_2 \frac{(1 - e^{-\kappa T})}{\kappa} \\ & + 2\rho_3\sigma_1\sigma_3 \frac{(1 - e^{-aT})}{a} - 2\rho_2\sigma_2\sigma_3 \frac{(1 - e^{-aT})(1 - e^{-\kappa T})}{a\kappa}. \end{aligned} \quad (42)$$


When time to maturity approaches infinity this expression converges to:

$$\sigma_F^2(\infty) = \sigma_1^2 + \frac{\sigma_2^2}{\kappa^2} + \frac{\sigma_3^2}{a^2} - \frac{2\rho_1\sigma_1\sigma_2}{\kappa} + \frac{2\rho_3\sigma_1\sigma_3}{a} - \frac{2\rho_2\sigma_2\sigma_3}{a\kappa}. \quad (43)$$


Figures 9a, 9b, and 9c plot the volatility of futures returns implied by the three models (using equations (39), (40), and (42)) for the Oil data, the Copper data, and the Enron oil data, respectively. The figures also show the actual volatility of futures returns of the contracts used in the estimation of the parameters of the models. It is surprising to see how well Models 2 and 3 fit the volatility of the data. It should be noted that only futures prices were used in the estimation; the volatility of futures returns is not an input in the estimation. The only volatility that enters into the estimation procedure is the volatility of the unobserved state variables. Finally, the figures show that Model 1 implies volatilities that are always smaller than the volatility of the data, with the difference being smaller for midmaturities and then increasing with increasing maturity of the futures contracts.

Figures 9 also indicate that, for the data considered, Models 2 and 3 imply very similar futures volatilities. The reason for this is that the volatility of the interest rate process in Model 3 (see Table IX) is of an order of magnitude smaller than the volatilities of the other state variables (around 1/25). If we

A



B



C




Figure 11. (a) The figure shows the term structure of copper futures prices for 7/29/88 implied by the three models for maturities up to ten years. (b) The figure shows the term structure of copper futures prices for 7/13/93 implied by the three models for maturities up to ten years. (c) The figure shows the term structure of copper futures prices for 6/13/95 implied by the three models for maturities up to ten years.

Table XIX

Cross-Section Comparison Between Models 1, 2, and 3 Out-of-Sample Enron Oil Data 1/15/93 to 5/16/96

Model	RMSE in Dollars				RMSE in Percentage			
	1	2	3 (7%)	3 (6%)	1	2	3 (7%)	3 (6%)
2 Years	0.35	0.09	0.16	0.20	1.91	0.49	0.86	1.12
3 Years	0.65	0.29	0.37	0.38	3.42	1.45	1.91	1.98
5 Years	1.29	0.80	0.96	0.65	6.42	3.94	4.89	3.24
7 Years	1.79	1.40	1.93	0.87	8.59	6.84	9.53	4.24
9 Years	2.24	2.14	3.27	1.33	10.53	10.22	15.68	6.38
All	1.44	1.20	1.76	0.79	6.95	5.81	8.54	3.85

compare equation (42) with (40), we see that the volatility of futures returns in Model 3 converges to the one in Model 2 when the volatility of interest rates goes to zero. Table XVIII reports the limiting volatilities (when time to maturity is zero and infinity) for the three models. Note that for the three data sets the volatilities for Model 2 start at a higher level and end at a lower level than for Model 3. The crossover is barely observed in Figures 9.

In summary, Model 1, which has very different implications about the volatility of futures returns as the maturity of the contract increases than Models 2 and 3, is incapable of describing the volatility of the futures data. Models 2 and 3 give similar implications because the volatility of the convenience yield overshadows the volatility of interest rates. This feature has important implications for valuation when the models are used in situations that involve longer term assets.

VI. Long Maturity Futures Contracts

The futures contracts available to estimate the parameters of the models discussed in this article have maximum maturities that are less than two years. Only the proprietary Enron oil forward curves have longer term maturities. It is of great interest to find out what are the implications of the models estimated with short maturity futures contracts with respect to longer term contracts, since many of the potential applications would involve assets with maturities longer than two years. In the previous section we examined the implications of the models with respect to volatility, and here we look at the implications with respect to price.

Figures 10 and 11 show the futures prices implied by the models for contracts up to ten years to maturity for the same observations as in Figures 6

and 7, respectively. Here we observe that even though Models 2 and 3 give very similar values in the range of observed contracts, they can diverge substantially as maturity increases. As mentioned in Section IV.D (see equation (38)), in Model 1 the futures price converges to a fixed value, so the rate of change in price as maturity increases converges to zero. For the estimated parameters in Models 2 and 3, however, even if initially the term structure of futures prices is decreasing, it eventually turns upward. In Models 2 and 3, the rate of change in price as maturity increases can easily be obtained by taking the derivative of the futures price with respect to time to maturity, dividing by the price, and taking the limit when time to maturity goes to infinity. For both models this "forward cost of carry"²⁷ converges to a rate that is independent of the initial value of the state variables. In Model 2 this rate is

$$\frac{1}{F} \frac{\partial F}{\partial T} (T \rightarrow \infty) = r - \hat{\alpha} + \frac{\sigma_2^2}{2\kappa^2} - \frac{\rho\sigma_1\sigma_2}{\kappa} \quad (44)$$

which translates, given the estimated parameters, into 2.71 percent per year for oil futures and 0.85 percent per year for copper futures. The corresponding rate in Model 3 is

$$\frac{1}{F} \frac{\partial F}{\partial T} (T \rightarrow \infty) = m^* - \hat{\alpha} + \frac{\sigma_2^2}{2\kappa^2} - \frac{\rho_1\sigma_1\sigma_2}{\kappa} + \frac{\sigma_3^2}{2a^2} + \frac{\rho_3\sigma_1\sigma_3}{a} - \frac{\rho_2\sigma_2\sigma_3}{\kappa a}, \quad (45)$$

which gives substantially higher values, 4.19 percent per year for oil futures and 2.70 percent per year for copper futures. This difference of 1.48 percent per year for oil futures and 1.95 percent per year for copper futures can make a big difference for futures prices with ten years to maturity, as can be observed in the figures.

The infinite maturity discount bond yield assumed in Model 3 has significant impact on the value to which the forward cost of carry converges. If the model is reestimated with an infinite maturity discount yield of 6 percent instead of 7 percent, there is practically no change in the estimated parameters, but the forward cost of carry converges to 3.56 percent for oil (instead of 4.19 percent) and 2.00 percent for copper (instead of 2.70 percent). These values are still higher, but much closer to those in Model 2.

Even when the models are estimated with the longer term Enron oil data, the forward cost of carry converges to a higher value in Model 3 (3.38 percent per year) than in Model 2 (2.25 percent per year). When Model 3 is reestimated assuming an infinite maturity discount bond yield of 6 percent instead of 7 percent, the obtained value of 3.03 percent per year is closer to the one in Model 2.

To give some insight into the performance of the models on longer term data when they are estimated on short term data, we estimate the parameters of the three models using the first five forward prices of the Enron oil data with maturities of 2 months to 1.5 years and then analyze how well they price the

²⁷ See Cortazar and Schwartz (1994).

last five forward contracts with maturities of 2 to 9 years. Table XIX reports the root mean square error (RMSE), both in dollars and in percentage, for Models 1, 2, and two versions of Model 3, one assuming a 7 percent infinite maturity discount bond yield and the other a 6 percent. The table shows that Model 2 always outperforms Model 1, with a RMSE on all contracts of 5.8 percent versus 7 percent. Model 3 performs the best when a 6 percent infinite maturity discount bond yield is assumed (with a RMSE on all contracts of 3.9 percent), but performs the worst when a 7 percent infinite maturity discount yield is assumed (with a RMSE on all contracts of 8.5 percent). This confirms the importance of the interest rate process parameters in Model 3 for the valuation of long term contracts. The main difference between Models 2 and 3 remains the valuation of long term futures contracts.

VII. Hedging Contracts for Future Delivery

An issue that has received increased attention in the literature²⁸ is the feasibility of hedging long term forward commitments in commodities with the existing short term futures contracts. The three models discussed in this article have implications for hedging strategies, which we now briefly discuss.

To properly hedge a forward commitment in a particular commodity, the sensitivity of the present value of the commitment with respect to each one of the underlying factors must equal the sensitivity of the portfolio of futures contracts used to hedge the commitment with respect to the same factors. This implies that the number of futures contracts required for the hedge is equal to the number of factors in the model used. Since Models 1 and 2 assume constant interest rates, and therefore futures prices are equal to forward prices, the present value of the forward commitment per unit of the commodity can simply be obtained by discounting the future (forward) price in equations (7) and (18).²⁹ For Model 3 the present value of a forward commitment, $P(S, \delta, r, T)$, is given in equation (30).

The number of long positions w_i in future contract with maturity t_i required to hedge a forward commitment to deliver one unit of a commodity at time T is then obtained in each one of the models by solving the following system of equations:

Model 1 (from equation (7)):

$$w_1 F_S(S, t_1) = e^{-rT} F_S(S, T). \quad (46)$$

Model 2 (from equation (18)):

$$\begin{aligned} w_1 F_S(S, \delta, t_1) + w_2 F_S(S, \delta, t_2) &= e^{-rT} F_S(S, \delta, T) \\ w_1 F_\delta(S, \delta, t_1) + w_2 F_\delta(S, \delta, t_2) &= e^{-rT} F_\delta(S, \delta, T). \end{aligned} \quad (47)$$

²⁸ See for example Brennan and Crew (1995), Culp and Miller (1994, 1995), Edwards and Canter (1995), and Ross (1995).

²⁹ Note that for Model 2 this present value is independent of the interest rate r .

Model 3 (from equations (26) and (30)):³⁰

$$\begin{aligned} w_1 F_S(S, \delta, r, t_1) + w_2 F_S(S, \delta, r, t_2) + w_3 F_S(S, \delta, r, t_3) &= P_S(S, \delta, r, T) \\ w_1 F_\delta(S, \delta, r, t_1) + w_2 F_\delta(S, \delta, r, t_2) + w_3 F_\delta(S, \delta, r, t_3) &= P_\delta(S, \delta, r, T) \\ w_1 F_r(S, \delta, r, t_1) + w_2 F_r(S, \delta, r, t_2) + w_3 F_r(S, \delta, r, t_3) &= P_r(S, \delta, r, T). \end{aligned} \quad (48)$$

Note that in Models 2 and 3 the hedge ratios w_i are independent of the spot price S .

Figure 12 illustrates the positions required to hedge a ten-year forward commitment to deliver one barrel of oil implied by the three models, using the parameter estimates from the Enron data. Panel A presents the number of one-month oil futures contracts required in Model 1 as a function of the spot price. Due to the strong mean reversion implied by this model, the hedging positions are unreasonably low. For example, for a spot price of \$20 the hedge ratio is only 0.25. Panel B shows the number of one-month futures contracts short and the number of one-year futures contract long required in Model 2 to hedge the ten-year oil forward commitment as a function of the convenience yield. Even though the hedge ratios do not depend on the spot price we need two futures contracts to hedge the two risk factors in the model.³¹ For example, for a convenience yield of 0.10, the hedge ratios are 1.09 long in the one-year contract and 0.34 short in the one-month contract. Finally, Panel C displays the number of six-month futures contracts long (top surface), one month futures contracts short (bottom surface), and one-year unit discount bonds long (middle surface) required in Model 3 as a function of the convenience yield and the instantaneous interest rate. For example, for a convenience yield of 0.10 and an interest rate of 0.05, the hedge ratios are 1.96 long in the six-month contract, 1.29 short in the one-month contract, and 0.94 long in the one-year discount bond. In this case we take positions in the discount bond to hedge against changes in interest rates. Note that in this model, as well and in Model 2, the hedge ratios are not very sensitive to the factors.


VIII. Investment under Uncertainty

The stochastic behavior of commodity prices has important implications for the valuation of projects related to the prices of those commodities (mines, oil deposits) and for the determination of the optimal investment rule, i.e., the commodity price above which it is optimal to undertake the project immediately. In this section we analyze a simple investment project and evaluate it using the three models developed in this article and two other benchmark procedures. The first benchmark is the traditional discounted cash flow (DCF) criteria and the second is a real option approach based on the assumption that


³⁰ Note that in this model the interest risk could be hedged using a bond instead of a third futures contract. This is what we do in the illustration that follows.

³¹ The results here are very similar to those presented for oil in Brennan and Crew (1995).


HEDGING TEN YEAR OIL FORWARD COMMITMENTS WITH FUTURES CONTRACTS



A. Model 1: 1 month futures



B. Model 2: 1 month and one year futures



C. Model 3: 1 and 6 month futures and 1 year discount bond

Figure 12. (a) The figure shows the number of one-month oil futures contracts required in Model 1 to hedge a ten-year forward commitment to deliver one barrel of oil, as a function of the current spot price. (b) The figure shows the number of one-month and one-year oil futures contracts required in Model 2 to hedge a ten-year forward commitment to deliver one barrel of oil, as a function of the current instantaneous convenience yield. (c) The figure shows the number of one-month and six-month oil futures contracts and one-year unit discount bonds required in Model 3 to hedge a five-year forward commitment to deliver one barrel of oil, as a function of the current instantaneous interest rate and convenience yield.

commodity prices follow a geometric random walk, that is, neglecting mean reversion.

The objective of this section is to show that different methods of analysis can give quite different values for an investment project and for the optimal investment rule, even if the assumptions made in each methodology are as realistic as possible. To be able to judge which methodology is superior, however, we would need to apply them to situations in which we have valid transaction prices. The conclusion to a recent article in the *Engineering and Mining Journal* on the capital budgeting methods used by mining companies directly addresses the issue we want to raise:

The use of DCF techniques to project valuations appears to be the industry standard. One disturbing result of the study, however, is the inability to explain the high premium that market values command over DCF valuations. The well known and used DCF analysis does not allow for placing premium values on projects under consideration. Perhaps the newer techniques such as option pricing methods of valuation may provide more accurate market value results (Bhappu and Guzman (1995), page 38).

We want to analyze a project as simply as possible, but that retains the main features that we want to highlight. Consider a copper mine that can produce one ounce of copper at the end of each year for ten years. Suppose that the initial investment required is $K = \$2$ and that the unit cost of production is $C = \$0.40$ (constant for the ten years). Assume that once the investment is done production will go ahead for the following ten years; that is, we neglect in this analysis the options to close and open the mine and the option to abandon it,³² and concentrate entirely on the option to invest. The first step in all the procedures we will discuss consists in the determination of the net present value of the project once it has been decided to go ahead with the investment (this is the "boundary condition" of the second step), and the second step consists in the evaluation of the option to invest. The net present value of the project once the investment has been decided is

$$NPV = \sum_{T=1}^{10} P(r, T, \cdot) - C \sum_{T=1}^{10} B(r, T) - K \quad (49)$$

where $P(r, T, \cdot)$ is the present value of the commodity to be received at time T when the interest rate is r (it also depends on the specific factors of a particular model), and similarly $B(r, T)$ is the present value of \$1 (which is equal to $\exp(-rT)$ when the interest rate is constant). It is important to note that the net present value (NPV) in equation (49) will be different for the different approaches since each one of them implies a different present value for the

³² For a detailed discussion of these options see Brennan and Schwartz (1985). Since the procedures to evaluate the mine are numerical, it would be trivial to incorporate them in the analysis. We have chosen to leave them out to simplify the analysis and presentation.

Table XX
Investment Criteria and Project Value in Mine Example

Row	Model	δ	r	NPV = 0	S^*	$S_1 = 0.50$	$S_2 = 1.00$	$S_3 = 1.50$
1	DCF	**	0.10	0.73	0.73	-1.40	1.61*	4.61*
2	DCF	**	0.12	0.76	0.76	-1.45	1.29*	4.03*
3	DCF	**	0.15	0.82	0.82	-1.52	0.88*	3.28*
4	Model 0	0.118	0.06	0.89	1.30	0.11	0.99	3.38*
5	Model 1	**	0.06	0.26	0.97	1.23	1.75*	2.52*
6	Model 2	0.10	0.06	0.70	1.12	0.27	2.38	5.94*
7	Model 2	0.25	0.06	0.80	1.20	0.16	1.68	4.68*
8	Model 2	0.40	0.06	0.90	1.32	0.09	1.15	3.58*
9	Model 3	0.10	0.03	0.69	0.96	0.19	2.29*	5.99*
10	Model 3	0.25	0.03	0.79	1.10	0.17	1.42	4.60*
11	Model 3	0.40	0.03	0.90	1.24	0.15	0.83	3.38*
12	Model 3	0.10	0.06	0.66	1.02	0.25	2.55	6.25*
13	Model 3	0.25	0.06	0.75	1.14	0.21	1.69	4.86*
14	Model 3	0.40	0.06	0.86	1.26	0.18	1.05	3.64*
15	Model 3	0.10	0.09	0.62	1.18	0.32	2.81	6.48*
16	Model 3	0.25	0.09	0.71	1.24	0.27	1.96	5.09*
17	Model 3	0.40	0.09	0.82	1.36	0.23	1.29	3.87*

* For copper spot prices above S^* it is optimal to invest immediately.

** This model does not use the convenience yield.

DCF = Discounted Cash Flow.

commodity to be received sometime in the future (and also for the interest rate in Model 3).

We will first evaluate our simple project using the traditional DCF criteria and the constant convenience yield model of Brennan and Schwartz (1985). Then, we will evaluate it using the three models developed and estimated in this article. All the results are reported in Table XX, which gives the copper spot price above which it is optimal to invest, S^* , the value of the project for copper spot prices of $S_1 = \$0.50$, $S_2 = \$1.00$, and $S_3 = \$1.50$, and the copper spot price at which the net present value of the project is zero (NPV = 0).

A. Discounted Cash Flow Criteria

In this approach the expected net cash flows are discounted at a rate that reflects the risk of these cash flows, so we need to specify the discount rate and the expected future copper spot prices for the next ten years. In practice, it is common to assume that spot prices will remain constant or use some industry prediction of future spot prices, and use discount rates that vary between 0.10 and 0.15.³³ In rows 1, 2, and 3 we report the results using risk-adjusted discount rates of 0.10, 0.12 and 0.15 respectively, and a flat copper spot price.

³³ Moyen, Slade, and Uppal (1996) report that most firms use a long-run commodity price, that there is a substantial agreement concerning this price, and that the most common hurdle rate used is fifteen percent.

At a discount rate of 0.12 the net present value is zero for a copper price of \$0.76 and the value of the mine is \$1.29 for a copper price of \$1.00. Note that the value of the project is very sensitive to the discount rate used, almost doubling from \$0.88 for a discount rate of 0.15 to \$1.61, for a discount rate of 0.10. There is always a discount rate that gives the same value as the other approaches discussed below, but as we shall see, all option-based methods give spot prices above which it is optimal to invest significantly higher than in the discounted cash flow criteria. The assumption of constant copper spot prices is also quite arbitrary; small deviations from this assumption can give very different project values.

B. Constant Convenience Yield: Model 0

In the methods that use the real options approach to valuation, instead of discounting expected cash flows at a risk-adjusted discount rate, certainty equivalent cash flows are discounted at the risk-free interest rate. For commodities this certainty equivalent cash flow is related to the forward price of the underlying commodity, which is equal to the futures price if the interest rate is constant. The different models make different assumptions about the stochastic behavior of commodity prices and therefore imply a different valuation of forward and futures contracts.

In the constant convenience yield model, which we shall call Model 0, the risk-adjusted process for the spot commodity price is assumed to follow a geometric Brownian motion:

$$\frac{dS}{S} = (r - c)dt + \sigma dz^* \quad (50)$$

where we use c for the convenience yield to distinguish it from δ used in the stochastic convenience yield models. In this model the net present value (49) becomes:

$$\text{NPV}(S) = S \sum_{T=1}^{10} e^{-cT} - C \sum_{T=1}^{10} e^{-rT} - K = S\beta_1 - \beta_2 \quad (51)$$

and the option to invest, $V(S)$, satisfies the ordinary differential equation:

$$1/2\sigma^2 S^2 V_{SS} + (r - c)SV_S - rV = 0 \quad (52)$$

subject to the boundary condition:

$$V(S) \geq \max[S\beta_1 - \beta_2, 0] \quad (53)$$

The solution to this equation is

$$V(S) = (S^* \beta_1 - \beta_2) \left(\frac{S}{S^*} \right)^d$$

where

$$S^* = \frac{\beta_2 d}{\beta_1 (d - 1)}$$

$$d = \frac{1}{2} - \frac{r - c}{\sigma^2} + \sqrt{\left(\frac{1}{2} - \frac{r - c}{\sigma^2} \right)^2 + \frac{2r}{\sigma^2}}. \quad (54)$$

S^* is the commodity price above which it is optimal to invest in the project.

Row 4 of Table XX shows the results for Model 0. The future contract closest to maturity (F_1) is used as a proxy for the spot to compute the standard deviation of the return on the spot copper price ($\sigma = 0.266$) and the convenience yield is computed as the average of the weekly convenience yields calculated from the first (F_1) and last (F_9) futures contract assuming an interest rate of 0.06 ($c = 0.118$). The critical copper price is \$1.30, which is, as expected, substantially higher than those obtained using the DCF criteria, and the value of the option to invest when the price of copper is \$1.00 is \$0.99; only for a discount rate of 0.15 does the discounted cash flow criteria give a lower value of the project at this price. The reason for this is that with the estimated convenience yield for this period of approximately 0.12 and an interest rate of 0.06, the risk-adjusted drift is close to -0.06 , which implies that the term structure of futures prices is declining at a rate of 6 percent per year.³⁴

C. Mean Reverting Spot Price: Model 1

The net present value in Model 1 is obtained from equation (49) by discounting the futures (or forward) prices as given by equation (7). To obtain the value of the investment option, $V(S)$, we need to solve a partial differential equation (PDE) identical to equation (9), except that in the right hand side we have rV instead of zero (also, if the investment option has infinite maturity the partial derivative with respect to time to maturity disappears). The boundary condition is the maximum of the net present value for this case and zero.

The PDE is solved by numerical methods and the results using the estimated parameters from Table 5 are reported in Row 5 of Table XX.³⁵ Since in this model futures prices converge to \$0.88 with zero volatility whatever the initial copper spot price, the value of the project is less sensitive to the spot price than

³⁴ A more typical long term convenience yield is 0.06. The example, however, shows the high variability of the convenience yield and its significant impact on valuation.

³⁵ For the three cases that involved the numerical solution of the PDE, we assumed that the investment option had a maturity of ten years.

any of the other methods discussed and the net present value becomes negative only when the spot price is below \$0.26, the lowest of any other case considered. In spite of this the critical spot price is quite high at \$0.97 since the spot price has a relatively high volatility.

D. Stochastic Convenience Yield: Model 2

For this model the NPV in equation (49) depends both on the spot price and the convenience yield, and the present value of one unit of the commodity is obtained by discounting the future (or forward) price given in equation (18). The value of the option to invest, $V(S, \delta)$, satisfies a PDE identical to (17), except that the right-hand side is rV instead of zero.

The corresponding PDE is solved numerically for the estimated parameters in Table VII and the results are given in Rows 6, 7, and 8 of Table XX for instantaneous convenience yields of 0.10 (a low convenience yield), 0.25 (an average convenience yield),³⁶ and 0.40 (a high convenience yield), respectively. Since the convenience yield is highly correlated with the spot price and strongly mean reverting, a low convenience yield indicates that spot prices will tend to go up, so the commodity spot price above which it is optimal to invest is lower and the value of the option to invest higher than when the convenience yield is high, which indicates that spot prices will tend to go down. For comparison with the other models is best to focus on the case with average convenience yield: the critical spot price of \$1.20 is lower than in Model 0, where mean reversion is neglected but higher than Model 1, where mean reversion seems to be too strong for longer maturities.

E. Stochastic Convenience Yield and Interest Rates: Model 3

In Model 3 the NPV of the project equation (49) depends also on the interest rate, in addition to the spot price of the commodity and the convenience yield. The present value of a unit of the commodity deliverable at time T must be computed using equation (30) since now forward prices are not equal to futures prices, and the present value of a unit discount bond must be computed using equation (29). The value of the option to invest, $V(S, \delta, r)$, satisfies a PDE identical to (25) except that the right hand side is rV instead of zero.

The PDE is solved by numerical methods for the estimated parameters in Table X and the results are reported in Rows 9 to 17 of Table XX for three different values of the convenience yield (0.10, 0.25, and 0.40) and three different interest rates (0.03, 0.06, and 0.09). Copper spot prices that trigger investment and the value of the investment option are very sensitive to the initial value of the convenience yield and interest rate. Note that for an initial convenience yield of 0.25 and an interest rate of 0.06 the copper spot price that triggers investment is lower than in Model 2 (\$1.14 versus \$1.20) but the value

³⁶ The reason this average convenience yield is much larger than the one used in Model 0 is that the former is an instantaneous convenience yield, whereas the latter is a convenience yield obtained between the first and last future contracts available.




Figure 13. The figure shows the net present value of a simple copper project, assuming that the investment is immediately undertaken, implied by the three models and two benchmarks: the discounted cash flow criteria and an option model that assumes constant convenience yield (Model 0), for different copper spot prices.

of the investment option is very similar when the spot copper price is \$1.00 (\$1.69 versus \$1.68).

F. Discussion

To be able to compare the investment implications of the different models we will concentrate the discussion of Model 2, assuming an average initial convenience yield (Row 7 of Table XX) and of Model 3, assuming an average initial convenience yield and interest rate (Row 13 of Table XX). For the DCF criteria we will assume a risk adjusted discount rate of 0.12 (Row 2 of Table XX). Figure 13 shows, for the different models considered, the net present value of the project once the investment has been decided. The investment will actually be undertaken, however, only when this value is positive for the DCF criteria, or when the investment option is optimally exercised in the other (option) models; this value, then, represents the boundary condition for the investment option problem. From equation (49) and the formulas for forward prices it is easy to see that in every case, with the exception of Model 1, the investment value is a linear function of the spot price. The strong mean reversion to a fixed price makes the investment value in Model 1 very insensitive to the spot price. The other models give similar values for low spot prices, but their values diverge when the spot price increases, Model 3 giving the highest value, followed by Model 2, the NPV criteria, and finally Model 0. The valuation

results for Models 1, 2, and 3 are consistent with the long term futures prices shown in Figure 11.

When the option element of the investment is considered, the values obtained under the different models will be nonlinear functions of the spot price (and also of the other factors in the particular model) and the investment rule will be determined by the optimal exercise of the option. From Table XX we observe that the NPV criteria give the lowest copper spot price above which it is optimal to invest and Model 0, which does not consider mean reversion, gives the highest. The three models discussed in this article, which take into account mean reversion, give intermediate copper prices that trigger investment (\$0.97 for Model 1, \$1.20 for Model 2, and \$1.14 for Model 3).

IX. Conclusion

The pricing and hedging of commodity contingent claims and the valuation of natural resource investments depend critically on the assumed stochastic process of the underlying commodity. In this article we develop three models that in different ways take into account the mean reverting nature of commodity prices, and estimate them using recent data on oil and copper futures prices.

The major difficulty we encounter in the analysis is that publicly available futures contracts have maturities shorter than two years, whereas many of the assets we wish to price and hedge have maturities much longer than two years. A particular model could fit very well the available short term futures contracts and do a poor job in predicting longer term futures prices that are essential for hedging long term commitments in the commodity or valuing mines or oil deposits. The proprietary oil data provided by Enron has allowed us to gain some insight into this issue.

Each one of models we consider has implications with respect to the volatility of futures returns and with respect to the behavior of long term futures prices. Model 1 implies that the volatility of futures returns will converge to zero and futures prices will converge to a fixed value as maturity increases. Models 2 and 3, however, for the estimated data imply that futures volatility will decrease but converge to a fixed value different from zero and the term structure of futures prices will eventually turn upward and converge to a fixed rate of growth even if initially is in strong backwardation. The evidence from the Enron oil forward curves imply that these properties are more desirable.

The real options approach to capital budgeting is gaining support both in the academic community and in the practice of finance. The analysis in this article suggests that it is very important to consider mean reversion in prices in evaluating projects. The DCF criteria induces investment too early (i.e., when prices are too low), but the real options approach induces investment too late (i.e., when prices are too high) when it neglects mean reversion in prices.

REFERENCES

- Atkinson, A. C., 1970, A method for discrimination between models, *Journal of the Royal Statistical Society* B32, 323–353.
- Bessembinder, H., J. F. Coughenour, P. J. Seguin, and M. M. Smoller, 1995, Mean reversion in equilibrium asset prices: evidence from the futures term structure, *Journal of Finance* 50:1, 361–375.
- Bhappu, R. R., and J. Guzman, 1995, Mineral investment decision making: A study of mining company practices, *Engineering and Mining Journal* 70, 36–38.
- Bjerk Sund, P., 1991, Contingent claims evaluation when the convenience yield is stochastic: analytical results, Working paper, Norwegian School of Economics and Business Administration.
- Bjerk Sund, P., and S. Ekern, 1995, Contingent claims evaluation of mean-reverting cash flows in shipping, in L. Trigeorgis, Ed.: *Real Options in Capital Investment: Models, Strategies, and Applications*, Praeger.
- 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*, North Holland.
- Brennan, M. J., and N. Crew, 1995, Hedging long maturity commodity commitments with short-dated futures contracts, Working paper, UCLA.
- Brennan, M. J., and E. S. Schwartz, 1985, Evaluating natural resource investments, *Journal of Business* 58, 133–155.
- Cortazar, G., and E. S. Schwartz, 1994, The evaluation of commodity contingent claims, *Journal of Derivatives* 1, 27–39.
- Cox, R. C., 1961, Tests of separate families of hypotheses, in Proceedings of the fourth Berkeley symposium on mathematical statistics and probability, Vol. 1, University of California Press, 105–123.
- Cox, R. C., 1962, Further results on tests of separate families of hypotheses, *Journal of the Royal Statistical Society* B24, 406–424.
- Cox, J. C., J. E. Ingersoll, and S. A. Ross, 1981, The relation between forward prices and futures prices, *Journal of Financial Economics* 9, 321–346.
- Culp, C. L., and M. H. Miller, 1994, Hedging a flow of commodity derivatives with futures: Lessons from Metallgesellschaft, *Derivatives Quarterly* 1, 7–15.
- Culp, C. L., and M. H. Miller, 1995, Metallgesellschaft and the economics of synthetic storage, *Journal of Applied Corporate Finance* 7, 62–76.
- Davidson, R., and J. G. MacKinnon, 1981, Several tests for model specification in the presence of alternative hypotheses, *Econometrica* 49, 781–793.
- Dixit, A. K., and R. S. Pindyck, 1994, *Investment under uncertainty*, Princeton University Press.
- Edwards, F., and M. Canter, 1995, The collapse of Metallgesellschaft: Unhedgeable risks, poor hedging strategy, or just bad luck?, *Journal of Futures Markets*, Forthcoming.
- 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.
- Ingersoll, J. E., and S. A. Ross, 1992, Waiting to invest: Investment under uncertainty, *Journal of Business* 65, 1–29.
- Jamshidian, F. and M. Fein, 1990, Closed-form solutions for oil futures and European options in the Gibson–Schwartz model: A note, Working paper, Merrill Lynch Capital Markets.
- Moyen, N., M. Slade, and R. Uppal, 1996, Valuing risk and flexibility: A comparison of methods, Working paper, University of British Columbia.
- Pesaran, M. H., and A. S. Deaton, 1976, Testing non-nested nonlinear regression models, *Econometrica* 46, 677–694.
- Ross, S. A., 1995, Hedging long run commitments: Exercises in incomplete market pricing, preliminary draft.
- Vasicek, O., 1977, An equilibrium characterization of the term structure, *Journal of Financial Economics* 5, 177–188.