

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

An Investigation of Alternative Optimization Parameters
with the Kalman Filter

Abstract

In this paper, the authors report on the use of a Kalman Filter to model oil price movements. The two-factor model utilized was originally developed by Schwartz and Smith (2000) and allows mean reversion in short-term price movements and uncertainty in the equilibrium level to which prices revert. The parameters of this model are estimated using prices for oil contracts over the time period covered by the original article. An alternative set of optimization parameters are then developed and compared with Schwartz and Smith's findings to determine whether the results are substantially different. Finally, the data set is updated to the present date and a determination of the pricing errors over the most recent period is reported.

Keywords: Kalman Filter, futures,, likelihood function, commodity prices

1. Introduction

The pricing of futures and forward contracts for most financial instruments is relatively straight forward. As explained in Hull (1993), the forward price of any security with a known payout is given by $F = Se^{(r-q)t}$ where S is the current spot price, r is the current interest rate, q is the known payout rate, and t is the time to expiration. This equation needs to be somewhat modified when analyzing consumable commodities. As Hull indicated for commodity forwards, the appropriate valuation formula is $F = Se^{(r-u-y)t}$ where u is the annualized storage costs and y is termed the convenience yield.

The empirical implementation of this pricing formula contains two difficulties. First, the spot price of a commodity can sometimes be so uncertain and illiquid that the first futures contract is used as a proxy. Second, while the convenience yield can be viewed as the benefit of ownership of the physical commodity, the actual measurement of this yield is problematic. In fact, Hull stated that the convenience yield measures the extent to which the left-hand side of the pricing equation is less than the right hand side. Clearly, this definition would make pricing a forward contract in isolation quite difficult.

Fama and French (1987) empirically tested a model of commodity futures prices based on interest rates, warehousing costs, and convenience yields on 21 commodities. To measure the effect of convenience yields, they employed a seasonal dummy variable. They found statistically significant variation in the futures prices based on changes in interest rates and this seasonality variable. Crude oil was not tested in their study.

Gibson and Schwartz (1990) developed a two-factor model for oil futures prices based on the spot price of oil and the instantaneous convenience yield. The spot price was assumed to undertake a lognormal-stationary distribution while the convenience yield was assumed to undertake mean reverting Brownian motion. In general, their results indicated that the two-factor model was a reasonable tool to model oil-linked assets such as futures prices.

Schwartz and Smith (2000) used a somewhat different approach in specifying a simple two-factor model. Specifically, the equilibrium price level was modeled to evolve according to geometric Brownian motion with drift. The short-term deviations, as defined as the difference between spot and equilibrium prices, were modeled to revert to zero according to an Ornstein-Uhlenbeck process. Though this may seem like a dramatic departure from the convenience yield models, Schwartz and Smith demonstrated that this model is in fact equivalent to Gibson and Schwartz.

The main difficulty in empirically testing this and other commodity models is that frequently the underlying variables are not directly observable. 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. The Kalman Filter generates estimates for the state variables and also facilitates the calculation of the likelihood of observing a certain data series given a particular set of model parameters. Maximum likelihood estimation can then be used to determine the optimal set of parameters. Hamilton (1994) provides a complete treatment of the application of Kalman Filter to time series.

Curt Wells (1996) examined the use of Kalman Filters in finance. He identified several uses of Kalman Filter including modeling term structure, forward exchange rate premium, real interest rates, and arbitrage pricing theory. His conclusion is that many more models than those few mentioned should be estimated with time-varying coefficients using the Kalman Filter technique.

Our goal is to use the Kalman Filter technique to determine the appropriate parameters for the Schwartz and Smith two-factor oil price model. We define the data under Section 2 – Data Set. We explain the short-term/ long-term model in more detail and the Kalman filter and optimization techniques we employed in Section 3 – Implementation of the Kalman Filter. In Section 4 – Results and Discussion, we examine

the estimated parameters and associated pricing errors. Finally in section 5, we offer some thoughts for additional research and concluding remarks.

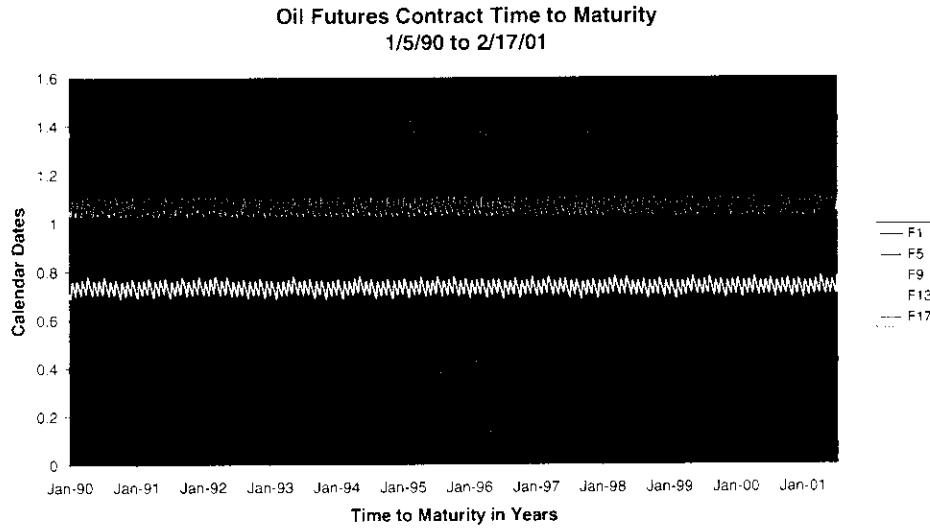
2. Data Set

To simulate the results reported in Schwartz and Smith (2000) we attempted to obtain as close a data set as possible. As detailed in Schwartz (1997), the data set used to test the model consisted of weekly observations for Nymex Crude oil futures contracts maturing in the next month and in approximately 5, 9, 13, and 17 months (5 contracts total). We obtained end of week observations from Bloomberg Financial Markets for the period used by Schwartz covering 1/5/90 to 2/17/95 (n=268). We also obtained the most recent data set from 2/24/95 to 07/17/01 (n=335).

An additional data set was obtained to more closely replicate the Schwartz and Smith study. As detailed in Schwartz (1997), the pricing data consisted of daily observations that were approximately transformed into weekly data by retaining every fifth observation. We obtained daily observations from Bloomberg Financial Markets for the period 1/2/90 to 2/15/95 for the oil futures contracts maturing in the next month and in approximately 5, 9, and 13 months. Every fifth observation was then taken (n=257). Unfortunately, while a complete set of weekly observations for the 17-month contract was available, there were 11 daily observations missing. As a 16-month contract was available for these observations, the average basis of the prior end of week and later end of week observation was applied to estimate these observations. This data set closely resembled that of Schwartz (n=258).

Since the contracts have a fixed maturity date, the time to maturity changes as time progresses. As detailed by the NYMEX web site, crude oil contract trading terminates at the close of business on the third business day prior to the 25th calendar day of the month. Figure 1 shows how time to maturity fluctuates as each contract matures.

Figure 1. Time to Maturity of Various Futures Contracts



3. Implementation of the Kalman Filter

Stochastic model

In the Schwartz and Smith (2000) model, the spot price of a commodity can be decomposed into two stochastic factors $\ln(S_t) = \chi_t + \xi_t$ where χ_t will be referred to as the short term deviation in prices and ξ_t is the equilibrium price level. The short run deviations (χ_t) are assumed to revert toward zero following an Ornstein-Uhlenbeck process:

$$d\chi_t = -\kappa\chi_t dt + \sigma_\chi dz_\chi$$

The equilibrium level (ξ_t) is assumed to follow a Brownian motion process:

$$d\xi_t = \mu_\xi dt + \sigma_\xi dz_\xi$$

In this model dz_χ and dz_ξ are correlated increments of standard Brownian motion processes. Putting these relationships in terms of the Kalman Filter, the evolution of the state variables is described using the following transition equation:

$$\bar{x}_t = \bar{c} + \bar{G}\bar{x}_{t-1} + \bar{\omega}_t, t = 1, \dots, n_T$$

where

$$\bar{x}_t \equiv \begin{bmatrix} \chi_t \\ \xi_t \end{bmatrix},$$

$$\bar{c} \equiv \begin{bmatrix} 0_t \\ \mu_\xi \Delta t \end{bmatrix},$$

$$\bar{G} \equiv \begin{bmatrix} e^{-\kappa \Delta t} & 0 \\ 0 & 1 \end{bmatrix}$$

and ω_t is a 2 x 1 vector of serially uncorrelated normally distributed disturbances with $E[\omega_t] = 0$ and $\text{Var}[\omega_t] = \text{Cov}[(\chi_{\Delta t}, \xi_{\Delta t})]$, $\Delta t \equiv$ the length of time steps and $n_T \equiv$ the number of time periods in the data set.

Kalman Filter Model

The measurement equation relates the state variables to the observed prices. In our case this is:

$$\bar{y}_t = \bar{d}_t + \bar{F}_t' \bar{x}_t + \bar{v}_t, t = 1, \dots, n_T$$

where $\bar{y}_t \equiv [\ln F_{T_1}, \dots, F_{T_n}]$, is a vector of the observed futures prices with time maturities T_1, T_2, \dots, T_n .

$$\bar{d}_t \equiv [A(T_1), \dots, A(T_n)],$$

$$\bar{F}_t \equiv [e^{-\kappa T_1} 1, \dots, e^{-\kappa T_n} 1],$$

and v_t is a vector of serially uncorrelated normally distributed disturbances with $E[v_t] = 0$ and $\text{Cov}[v_t] = V$.

Given these equations and a set of observed futures prices, Hamilton (1994) derived an iterative recursive algorithm as follows:

$$\begin{aligned} z_{t+1} &= Fz_t + v_{t+1} \\ y_t &= A'x_t + H'z_t + w_t \end{aligned}$$

Based on Hamilton (above), we can then derive an iterative algorithm for the Kalman Filter:

$$\text{IK (1)} \quad z_{|0} \text{ \& } P_{|0}$$

$$\begin{aligned} \text{IK (2)} \quad y_{t|t-1} &= A'x_t + H'z_{t|t-1} \\ K_t &= F\Pi_t = FP_{t|t-1}H(H'P_{t|t-1}H + R)^{-1} \end{aligned}$$

$$\begin{aligned} \text{IK (3)} \quad z_{t|t} &= z_{t|t-1} + \Pi_t(y_t - y_{t|t-1}) \\ P_{t|t} &= P_{t|t-1} - \Pi_t H' P_{t|t-1} \end{aligned}$$

$$\begin{aligned} \text{IK (4)} \quad z_{t+1|t} &= Fz_{t|t} = Fz_{t|t-1} + K_t(y_t - y_{t|t-1}) \\ P_{t+1|t} &= FP_{t|t}F' + Q \end{aligned}$$

In order, to achieve faster computation, we calculated the iterative $P_{t|t}$ using the steady state value of $P_{t|t}$ as t tends to infinity. To obtain the P steady state, we solved the Riccati equation associated to the Kalman Filter:

$$\text{(RK)} \quad FPF' - P - K_t H' P F' + Q = 0$$

Model Parameters

The model parameters (K) is defined by 7 parameters:

$$\begin{matrix} \sigma_x & \mu_x & \lambda_x \\ \sigma_\xi & \mu_\xi & \mu_\xi^* & \rho_{\xi,x} \end{matrix}$$

as well as a matrix \mathbf{V} that represents the variance and covariance of the observations. In general there are $(n+1)*n/2$ free variables in the covariance matrix where n is the number of futures contracts whose prices are observed. In Schwartz (1997) and Schwartz and Smith (2000) this estimation problem is simplified by assuming that \mathbf{V} is diagonal with diagonal elements (s_1^2, \dots, s_n^2) . We term this set of 12 parameters the Diagonal Parameters.

We then defined an enhanced version of the problem where the observation noise covariance values matrix is defined in a generic way using the Cholesky Decomposition:

$$\begin{bmatrix} \sigma_{1,1} & 0 & 0 & 0 & 0 \\ \sigma_{2,1} & \sigma_{2,2} & 0 & 0 & 0 \\ \sigma_{3,1} & \sigma_{3,2} & \sigma_{3,3} & 0 & 0 \\ \sigma_{4,1} & \sigma_{4,2} & \sigma_{4,3} & \sigma_{4,4} & 0 \\ \sigma_{5,1} & \sigma_{5,2} & \sigma_{5,3} & \sigma_{5,4} & \sigma_{5,5} \end{bmatrix}$$

This decomposition allows us to define the observation error matrix by 15 scalar parameters. This transformation will ensure the symmetry of \mathbf{V} and the semi-definite positive nature of the matrix. We term this set of 22 parameters the Full Parameters.

Parameters values and Maximum likelihood

To estimate the 22 Full Parameters of (K), we used the maximum likelihood estimation defined by Hamilton:

$$(\text{MLF}) L(\alpha) = \max_{(K)} \sum_t \text{Log}(f_{y_t|x_t, y_{t-1}}(y_t|x_t, y_{t-1}))$$

where

$$f_{y_t|y_{:t-1}}(y_t|x_t, y_{t-1}) = (2\pi)^{-\frac{1}{2}} |H'PH + R|^{-\frac{1}{2}} \times e^{-\frac{1}{2}} \left\{ (y_t - y_{t|t-1}) (H'PH + R)^{-1} (y_t - y_{t|t-1}) \right\}$$

We can define the space where the parameters are defined as:

(SMLF) σ_x σ_ξ are positives

μ_x λ_x μ_ξ μ_ξ^* are real

$\rho_{\xi,x}$ is between 0 and 1

$\text{Chol}(V)'$. $\text{Chol}(V)$ is semi-definite positive

Numerical optimization

We optimized the likelihood function defined above to find the best estimate of our model parameters (K). Due to the high number of dimensions, we discarded the grid search method. Instead, we realized an optimization by computing the gradient of (MLA) where (MLA) is defined as:

$$(MLA) \quad L(\alpha) = L(\alpha_0) + g(\alpha_0)(\alpha - \alpha_0) + (\alpha - \alpha_0)' H(\alpha_0) (\alpha - \alpha_0)$$

$$g(\alpha_0) = \frac{\partial L}{\partial \alpha_i} \Big|_{\alpha=\alpha_0} \quad \text{where } \alpha = \begin{pmatrix} \alpha_1 \\ \vdots \\ \alpha_n \end{pmatrix}$$

$$H(\alpha_0) = \frac{\partial^2 L}{\partial \alpha_i \partial \alpha_j} \Big|_{\alpha=\alpha_0}$$

We first picked an initial random value as a parameter set in the space (SMLF). From this point we computed the gradient that maximizes the ascending slope in all directions. We adjusted the parameters in that direction with a dynamically adjusted step size that preserves the gain on the likelihood function. When the algorithm reached a flatter region we gradually adjusted the step smaller until we reached a local maximum value. If the algorithm could not make further progress, we randomly threw rays in the 22


dimensions scaled by the gradient sensitivity in order to reach a higher point that may reinitiate a new search to a global maximum.

We tried several algorithms among them the matlab “fminsearch” and the Newton-Raphson method:

$$(NRA) \alpha_{k+1} = \alpha_k - h(\alpha_k)^{-1} g(\alpha_k)$$

Figure 2 demonstrates the optimization routine where each circle represents a change in the parameters and a recalculation of the likelihood function. The gridded surface is the likelihood function projected on the two-dimensional space sigma-ksi, rho.

Figure 2. Likelihood Function Optimization for Sigma-Ksi and Rho



Implementation of Optimization Procedure

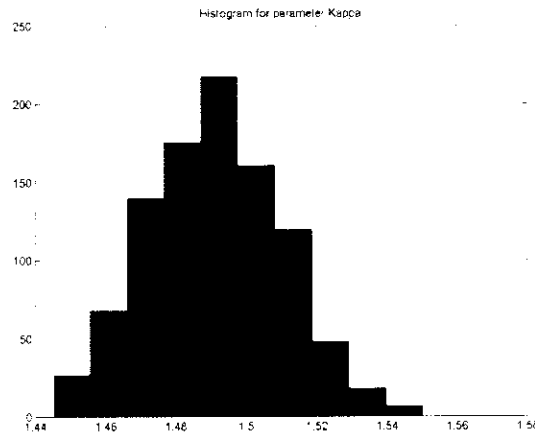
To be sure that our estimation procedure reaches a global (rather than a local) maximum, we reran the optimization procedure from several thousand different sets of randomly

selected initial parameters. In addition we generated randomness in the state variable noise matrices. As this procedure is very computationally intensive it was necessary to utilize the computing power of the Berkeley Research Center.

Due to numerical instability for ill-conditioned matrices, it was necessary to compute a constraint on the condition number of the Cholesky Decomposition. We then prevented the system from initially selecting values in this region.


In order to test the procedure, we first computed the matrix norm of the difference between the riccati solution and the iterative matrix P_{it} for several initial conditions. We systematically noticed an excellent convergence. We then considered the Kalman system as a function of the state noise matrix with 22 output parameters. We drew several state noise matrices, and examined the output of the maximum likelihood Kalman system for the 22 output parameters. We analyzed the distribution of these parameters around their means as in Figure 3 and compared this value to results found in Schwartz (2000).

Figure 3. Histogram for parameter distribution analysis



The following figure demonstrates the likelihood function vs. the number of iterations of the optimization routine. It also illustrates how quickly the likelihood function converges towards its maximum value.

Figure 4. Likelihood Function vs. Number of Iterations



In comparing the Diagonal Parameter optimization with the Full Parameter optimization, we noticed that the Diagonal optimization was not giving results as good in terms of the likelihood function level. We know that the Diagonal optimization is done on the matrix space restricted to the diagonal matrices and that this space is by construction not as dense as the space of the symmetric semi-definite positive function. Therefore, it is normal that the maximum likelihood function will not reach the same level. However we noticed from the results that the Diagonal Parameter estimates were very close to the Full Parameter estimates.

To compare the Full and Diagonal Parameterization it was necessary to estimate the pricing error. More precisely, we represented the average estimation of the error across all five futures prices each time. From these results, we estimated this average for several runs, and then ran several simulations of the following estimator:

$$\text{Probability } (\mu_{Full} > \mu_{Partial})$$

Where

$$\mu^{Full} = \frac{1}{m} \sum_{j=1}^m \frac{1}{5} \sum_k \mathcal{E}_t^k$$

$$\mu^{Partial} = \frac{1}{m} \sum_{j=1}^m \frac{1}{5} \sum_k \mathcal{E}_t^k$$

$$\mathcal{E}_t^k = |y_t^k - y_{t|t}^k| \quad k \text{ is across future prices}$$

For a thousand runs, the above estimator gives a histogram that could be associated with a normal distribution with a mean of 0.9 and a standard deviation of 0.1. The appropriate result would show that the Full Parameter optimization gives better estimates in probability than the Diagonal Parameter optimization.

4. Results and Discussion

Based on the Schwartz daily data and using the Full Parameter method of optimization we derived the following results. These are compared to Table 2 as provided in Schwartz (2000):

Table 1 Maximum Likelihood Parameter Estimates
Every fifth daily data - 01/02/90 to 02/15/95


Parameter	Description	Our Estimate	Schwartz Estimate	Schwartz Std. Error
κ	Short-term mean reversion	1.50	1.49	.03
σ_γ	Short-term volatility	27.6%	28.6%	1.0%
λ_γ	Short-term risk premium	15.7%	15.7%	14.4%
μ_ξ	Equilibrium drift rate	-2.25%	-1.25%	7.28%
σ_ξ	Equilibrium volatility	13.5%	14.5%	0.5%
μ_ξ^*	Equilibrium risk-neutral drift	1.25%	1.15%	.13%
$\rho_{\xi\gamma}$	Correlation in increments	0.301	0.300	.044

The optimization of the maximum likelihood function derives very similar results when the data set is restricted to that used by Schwartz. The same procedure was applied to the weekly data rather than every fifth observation of the daily data. Very similar estimates were derived and are reported in Table 2.

To determine the sensitivity of the likelihood function to small changes in each of these parameters, gradient analysis was performed at the maximum likelihood estimate. The results are shown graphically in Figure 5 where every component of the gradient is


represented as a vector. Each vector emanates from the origin with a length proportionate to the gradient component value.

Figure 5. Magnitude of Parameter Changes on Likelihood Function



Clearly the drifts as denoted by the mean reversion rate κ and the long equilibrium drift rate μ_ϵ^* , have the largest effect on the likelihood function. The effects of the Cholesky parameters can be seen in Figure 6 where the scale of the magnitude is related to κ .

Figure 6. Magnitude of Parameter Changes on Likelihood Function Including Cholesky Factors



Clearly the Cholesky factors affect the likelihood function to a larger degree than the other parameters. This is most likely due to the fact that a linear change in the Choleski parameter has a non-linear effect on the likelihood function.

We then re-ran the optimization with a more restricted set of parameters as detailed under Implementation of the Kalman Filter. The resulting parameter estimates and maximum likelihood estimates are presented in Table 2.


Table 2 Maximum Likelihood Parameter Estimates
 Diagonal Optimization Function vs. Full Optimization Function
 Weekly data - 01/05/90 to 02/17/95

Parameter	Descriptions	Diagonal Parameter Estimate	Full Parameter Estimate
κ	Short-term mean reversion	1.49	1.50
σ_{χ}	Short-term volatility	27.6%	27.6%
λ_{χ}	Short-term risk premium	20.7%	15.7%
μ_{ξ}	Equilibrium drift rate	-2.25%	-2.25%
σ_{ξ}	Equilibrium volatility	14.5%	13.5%
μ_{ξ}^*	Equilibrium risk-neutral drift	1.15%	1.25%
$\rho_{\xi\chi}$	Correlation in increments	0.30	0.301
	Maximum Likelihood Estimate	4757	5112

Though there were certain differences in the estimates, on the whole both the Diagonal and Full Parameter optimization methods derived similar results. Interestingly, the Diagonal Parameter optimization provided results closer to Schwartz for both the equilibrium volatility and the equilibrium risk-neutral drift.

The time series of pricing errors over several runs for the Diagonal and Full Parameter optimization were compared. Despite finding parameters that are substantially equivalent, the pricing error of the Diagonal Parameter optimization was consistently higher as shown in Figure 7. More precisely the results represent the average estimation of the error across futures prices for every time and in probability. This results shows that the Full Parameter optimization gives less pricing error in probability than the Diagonal Parameter optimization.

Figure 7. Pricing Error Comparison



Finally we updated the model until the current date using the Full Parameter optimization method. The following table displays the results for the last six years and for the entire estimate period.

Table 3 Maximum Likelihood Parameter Estimates
 Full Optimization Function
 Weekly data - 02/24/95 to 07/17/01

Symbol	Description	02/24/95 to 07/17/01 Estimate	01/05/90 to 07/17/01 Estimate
κ	Short-term mean reversion	1.495	1.495
σ_x	Short-term volatility	28.1%	28.1%
λ_x	Short-term risk premium	18.2%	18.2%
μ_ξ	Equilibrium drift rate	-1.75%	-1.75%
σ_ξ	Equilibrium volatility	14.0%	14.0%
μ_ξ^*	Equilibrium risk-neutral drift	1.15%	1.20%
$\rho_{\xi x}$	Correlation in increments	0.301	.3005
	Maximum Likelihood Estimate	5867	10656

Interestingly, the estimates are quite consistent using this new set of data. It is helpful to now look at the results of this analysis and compare the estimates of short-term versus long-term price movements over various periods. The estimates over the same time period as Schwartz(2000) appears very similar to Figure 4 in his paper. Our results are presented in Figure 8.

Figure 8 Estimated Spot and Equilibrium Prices, Jan. 5/90 to Feb. 17/95

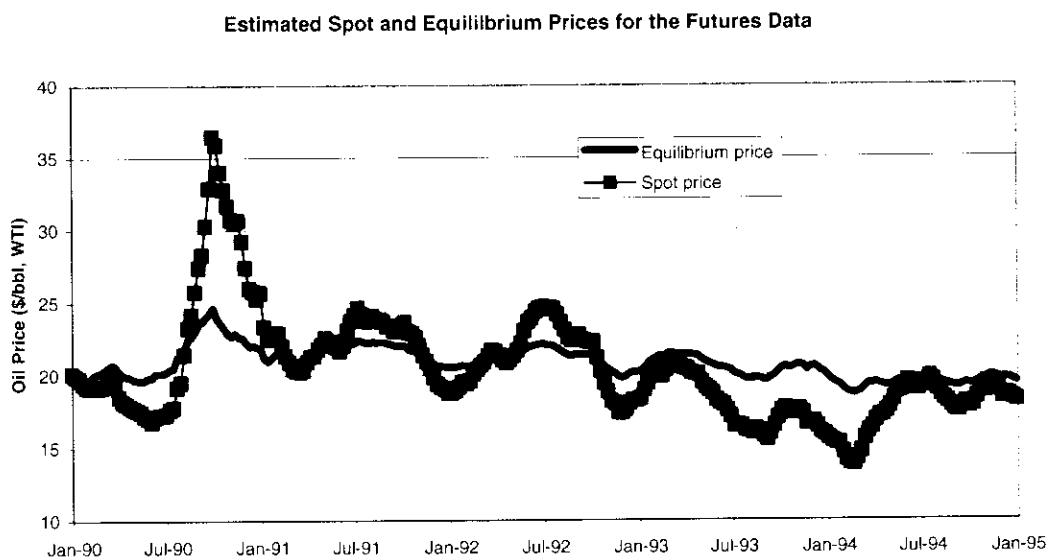




Figure 9 shows the same results for the new data period from 02/24/95 to 07/17/01.

Figure 9 Estimated Spot and Equilibrium Prices, Feb. 24/95 to Jul. 17/01



The following two graphs show the estimated spot and equilibrium prices for two very important time periods.


Figure 10 Estimated Spot and Equilibrium Prices, Jun. 27/97 to Jun. 25/99



In early 1999, oil prices fell to their lowest level in over a decade. The Kalman Filter reveals that this fall had two components. First in early 98 the short-term price fell

significantly. Later in early 99 the long-term price fell substantially. It is interesting that the long-term decline in early 99 quickly became more positive after the short-term price reversed.


Figure 11 Estimated Spot and Equilibrium Prices, Jun. 25/99 to Jun. 29/01



In contrast, this period was punctuated with significantly rising prices as OPEC curtailed production and inventories fell to historically low levels. Again the Kalman Filter was able to disaggregate the price rise into two components. First, in late 1999, the short-term price movement became positive. Later, in mid-2000, the equilibrium price increased consistently over a six month period.

Finally, an analysis of the pricing errors found that, consistent with Schwartz and Smith (2000), the shortest maturity futures contracts were the hardest to value. The mean squared pricing errors for the first and second futures contracts are presented in Figure 12.

Figure 12 Mean Squared Pricing Error for
1 Month and 5 Month Futures Contracts



Two conclusions can be drawn. First, the average error over the Schwartz data set was much lower than the average error over the last year and a half. Despite the consistency in the parameters, it appears that the model is having difficulty adjusting to the substantial price volatility in oil prices over the most recent period. Second, the mean squared errors are demonstrating heteroscedasticity of the volatility. As the pricing error is largest for the shortest term to maturity futures contracts, this indicates that an ARCH or GARCH model may more accurately reflect the standard deviation in the short-term factor.

5. Conclusion

This paper accomplished three main results. First, the Kalman Filter appears to have successfully disaggregated oil price movements into a short and long-term price component. For oil analysts, these results may be helpful in modeling oil price movements. Short-term changes may be determined by such factors as variations in weather or short-term supply disruptions. Intuitively, when the price of a commodity is consistently higher than the long-run equilibrium level, additional marginal supply will

be brought on, putting downward pressure on prices. A determination of the short-term and long-term price is critical in this analysis.

As stated in Schwartz [2000], these results can also be helpful in capital allocation by providing a more consistent estimator of equilibrium prices and volatilities for real option analysis of capital spending projects. The ability to simplify the analysis of long-term investments is one of the advantages of the short-term/long-term model.

Second, we defined an enhanced version of the optimization parameters where the observation noise covariance values matrix was defined in a generic way using the Cholesky Decomposition. Previous articles have simplified the problem by assuming that the noise matrix has only diagonal elements. While no substantial difference was found in the optimized parameters, this broader set was found to reduce the error in the pricing equation.

Third, we extended the original Schwartz data set by six years. The optimized parameters for our new data set were found to be quite similar to the Schwartz data set. However, the squared pricing errors were found to increase substantially over the past year and a half relative to the prior eight years. In addition, the errors indicated a degree of heteroscedasticity. Future research should investigate an alternative characterization of the short-term pricing equation. This characterization could involve volatility dependence on lagged squares of the short-term factor as well as lagged values of the volatility itself.

References

- Fama, E., and K. French (1987), Commodity Futures Prices: Some Evidence on Forecast Powers, Premiums, and the Theory of Storage, *Journal of Business*, 60 (1), 55-73.
- Gibson, R., and E. Schwartz (1990), Stochastic Convenience Yield and the Pricing of Oil Contingent Claims, *Journal of Finance*, 45 (3), 959-975.
- Hamilton, J., (1994), *Time Series Analysis*, Princeton University Press, Princeton, New Jersey.
- Hull, J., (1993), *Options, Futures, and Other Derivative Securities, Second Edition*, Prentice-Hall, Englewood Cliffs, New Jersey.
- Schwartz, E., (1997), The Stochastic Behavior of Commodity Prices, *Journal of Finance* 52 (3), 923-973.
- Schwartz, E., and J. Smith, (2000), Short-Term Variation and Long-Term Dynamics in Commodity Prices, *Management Science*, 46 (7), 893-911.
- Wells, C., (1996), *The Kalman Filter in Finance, Advanced Studies in Theoretical and Applied Econometrics Volume 32*, Kluwer Academic Publishers, Dordrecht, Netherlands.