

A Seasonal Stochastic Volatility Model for Futures Price Term Structure

Yanjun Tian and Paul L. Fackler

August 7, 2000

Abstract

Pricing and estimation issues of exponential affine stochastic volatility models are discussed. One specific model is estimated with Chicago Board of Trade futures price data, where the instantaneous mean and volatility of commodity spot price are allowed to be time varying. Model performance is evaluated based on its fit to the futures price term structure and the model implied state variable behavior.

1 Introduction

Increasingly, continuous time diffusion factor models are being used to model commodity price behavior. In these models, arbitrage relationships are directly imposed by expressing futures prices in terms of an underlying, but unobserved, factor model. Such models provide an integrated framework for examining dynamic commodity price behavior and its implications for pricing options and other derivatives, as well as for analyzing commodity related management decision problems including production, storage, and hedging decisions.

By specifying the dynamics of the underlying factors, rather than working directly with the dynamics of the futures prices, the diffusion factor models can address the problem that futures prices do not form standard time series. The dynamics of futures price depends critically on the delivery time and the time-to-maturity of a particular contract. The practice of forming a continuous futures price time series by splicing together series with different maturity dates ignores this dependence. Furthermore, the practice is not useful in modeling the term structure of futures prices.

Most diffusion factor models for futures prices to date have assumed a Gaussian (constant volatility) underlying factor model (Brennan and Schwartz, Gibson and Schwartz, Schwartz). However, the constant volatility assumption is known to be inadequate in commodity market as in many other markets. The evidence for this is two fold. First, implied volatilities from commodity options across different strike prices are not constant. Second, the time series data of commodity prices exhibit intertemporal volatility clustering (Baillie and Myers, Kenyon, et al., Streeter and Tomek).

In addition, commodity prices are known to be highly seasonal. Seasonal patterns in the volatility of individual price series and the covariance among them are important for managing commodity price risks. To date, however, seasonality has not been incorporated into the diffusion factor models.

The goal of this paper is to examine the role of stochastic volatility (SV) and seasonality in modeling the term structure of commodity futures prices. However, introducing SV itself is nontrivial. As a result, derivative asset prices often need to be solved numerically as a function of the underlying state variables. Moreover, the state transition density is not known in closed-form, which leads to difficulties in econometric analysis. Inevitably, time-varying parameters causes additional complications.

The computational and econometric considerations favor a specific class of SV

models, the affine SV models. In the affine models, the drift terms and instantaneous covariance matrix of the factor model are both affine in the underlying factors. This particular diffusion form provides computational tractability in multivariate settings because it leads to nearly closed-form, easily evaluated expressions for futures and (European) options prices (Schwartz, Heston, Miltersen and Schwartz). Furthermore, estimation methodologies for this class of models using futures price data are relatively well developed.

The remainder of the paper begins with a description of affine diffusions and the pricing of futures contracts based on affine diffusion and then discusses model specification issues and introduces a specific three factor model. Section 4 presents a GMM approach in which the moment conditions are written in terms of the conditional characteristic function. Section 5 reports empirical results for Chicago Board of trade corn futures. A summary concludes the paper.

2 Affine Diffusion Models for Futures Prices

This section provides some background knowledge that will be used in later discussion. It first introduces the affine class of diffusion models and then discusses futures pricing when the spot asset price is characterized as an affine diffusion.

An affine diffusion model can be described by a stochastic differential equation (SDE), in which both the instantaneous mean and variance are affine functions of the process:

$$dx = (a + Ax) dt + C\sqrt{b + Bx}dW, \tag{1}$$

where a and b are N -vectors; A , B , and C are $N \times N$ matrices; W is a N -vector of independent Brownian motions.¹

¹Parameterizing the instantaneous volatility as $C\sqrt{b + Bx}$ ensures that the covariance is nonneg-

A cornerstone of modern finance theory is the equivalence between the lack of arbitrage opportunities and the existence of a probability measure under which the value of an asset is equal to the expected value of the returns generated by the asset, discounted at the risk free rate. Such a probability measure, here denoted the Q measure, will in general differ from the data generating mechanism (denoted the P measure) to the extent that the market assigns non-zero price to the risks that drive the asset price.

In this study, it is assumed that the underlying state process is affine under both the P and Q measures. The two measures are related through a risk premium Λ which is also affine in x , and of the form $\Lambda = C \text{diag}(a + Bx) d$, where d is a N -vector of constants. The state diffusion under the Q measure is described by

$$dX = (\hat{a} + \hat{A}x) dt + C \sqrt{\text{diag}(b + Bx)} d\hat{W}, \quad (2)$$

where $\hat{a} = a - C \text{diag}(d) b$, $\hat{A} = A - C \text{diag}(d) B$, and \hat{W} is a N -vector of independent standard Brownian motions under the Q measure.

Given the state diffusion defined by *SDE* (2), the relationship between futures prices and the underlying state variables is determined by the intertemporal arbitrage relationship and appropriate boundary conditions. The intertemporal arbitrage relationship is that a futures price is a martingale under the Q measure. The associated terminal condition is that futures price at a contract's maturity date, T , is equal to the spot price.

Let $F(x, \tau; T)$ denote the price of a futures contract with a time-to-maturity τ , traded at time $t = T - \tau$. The arbitrage relationship can be expressed as

$$E^Q dF(x, \tau; T)/dt = 0, \quad (3)$$

ative definite so long as $b + Bx(t)$ is nonnegative. The nonnegative definiteness of $b + Bx(t)$ is the essential admissibility condition of an affine diffusion. For more details, see Duffie and Kan.

which can be expanded into a partial differential equation (*PDE*) by applying Ito's Lemma:

$$F_\tau = F_x (\hat{a} + \hat{A}x) + \frac{1}{2} \text{tr} (C \text{diag}(b + Bx) C^\top F_{xx}), \quad (4)$$

with the associated boundary condition is

$$F(x; 0; T) = S(x, T), \quad (5)$$

where S is the underlying spot price, which is a function of the underlying state process.

It can be verified that, for an affine state process, with

$$\ln S = h_0 + hx, \quad (6)$$

the futures price has the form

$$F(x; \tau, T) = e^{\beta_0(\tau) + \beta(\tau)x(t)} \quad (7)$$

where $\beta_0(\tau)$ and $\beta(\tau)$ satisfy the following *ODEs*,

$$\frac{d\beta}{d\tau} = \beta \hat{A} + \frac{1}{2} \beta C \text{diag}(\beta C) B, \quad (8)$$

$$\frac{d\beta_0}{d\tau} = \beta \hat{a} + \frac{1}{2} \beta C \text{diag}(\beta C) b. \quad (9)$$

with boundary conditions $\beta(0) = h$ and $\beta_0(0) = h_0$. These *ODEs* are easily solved using standard methods such as Runge-Kutta (Press, et al.) or collocation (chapter 11, Judd).

The affine diffusion model makes it possible to specify and solve explicitly multi-factor models for futures price. In particular, it is possible to specify models of spot prices that are affected by stochastic convenience yield and stochastic volatility. The affine structure does come at a cost, however, as it imposes strong restrictions on the behavior of futures prices at different maturities. In particular, for an n -factor model, the affine diffusion model implies that the log of any futures price can be expressed as an affine function of the logs of exactly n other futures prices.

3 Model Specification

Two features of the affine diffusion model are critical in determining the qualitative nature of the state process. First, the eigenvalues of A determine the stationarity properties. If all of the eigenvalues are negative, the process has a stationary invariant distribution (although it is generally difficult to express this in closed form).

Dai and Singleton have shown that the behavior of affine diffusion models also depends critically on the rank of B , which defines how many independent stochastic volatility terms affect the system. If $m = \text{rank}(B) = 0$ (i.e., $B = 0$) the model has constant volatility and is Gaussian. For $m > 0$ the volatility of the process is dependent on the level of x .

The number of stochastic volatility terms affects both the admissibility of an affine diffusion and the identification of model parameters. An affine diffusion is said to be admissible if $b + Bx$ is non-negative with probability one. A model is identified if no affine transformation of the latent state model exists that results in the same observable variables. The rank of B affects the number of parameters that can be identified; specifically, if the observable variables can be written as an affine function of the state, then the number of parameters that can be identified can be shown to be $n(n + 1) + m + 1$.²

The state variables in this framework are latent and therefore can be specified in a variety of ways. We will focus on a 3-factor specification in which the factors are interpreted in a natural way. The first factor is interpreted as the log of the spot price. This means that the initial condition for a futures price is that $F(x, 0; T) = \exp(0 + [1 \ 0 \ 0]x) = e^{x_1 T}$. The second and third factors will be interpreted as the rate of change (drift) and the instantaneous proportional variance of the spot

²This assumes that A is full rank and $m > 0$.

price:

$$dS = x_2 S dt + \sqrt{x_3} S d\tilde{W},$$

where \tilde{W} is a one-dimensional standard Brownian motion. Applying Ito's Lemma,

$$d \ln(S) \equiv dx_1 = (x_2 - x_3/2)dt + \sqrt{x_3}d\tilde{W}.$$

We further assume that only one stochastic volatility term effects the system ($\text{rank}(B)=1$). This model can be written as

$$dx = \left(\begin{array}{c} \begin{bmatrix} 0 \\ a_2 \\ a_3 \end{bmatrix} + \begin{bmatrix} 0 & 1 & -1/2 \\ A_{21} & A_{22} & A_{23} \\ 0 & 0 & A_{33} \end{bmatrix} x \end{array} \right) dt + \begin{bmatrix} 1 & C_{12} & C_{13} \\ C_{21} & 1 & C_{23} \\ 0 & 0 & 1 \end{bmatrix} \text{diag} \left(\sqrt{\begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix} + \begin{bmatrix} 0 & 0 & B_{13} \\ 0 & 0 & B_{23} \\ 0 & 0 & B_{33} \end{bmatrix} x} \right) dW, \quad (10)$$

subject to the two restrictions that

$$b_1 = -(C_{12}^2 b_2 + C_{13}^2 b_3)$$

and

$$B_{13} = 1 - (C_{12}^2 B_{23} + C_{13}^2 B_{33}).$$

These restrictions ensure that element (1,1) of the $Cov(x) = C \text{diag}(b + Bx) C^\top$ is equal to x_3 .

The model has 14 free parameters to be estimated, which can be expressed as a vector, θ :

$$\theta = [a_2 \ a_3 \ A_{21} \ A_{22} \ A_{23} \ A_{33} \ C_{12} \ C_{13} \ C_{21} \ C_{23} \ b_2 \ b_3 \ B_{32} \ B_{33} \ d_1 \ d_2 \ d_3]^\top,$$

with d_1, d_2, d_3 being the market price of risk of the three factors. It is a maximal model, in the sense of Dai and Singleton, in that it contains the maximum number of parameters that can be identified.

A discretized version of this model makes clear the relationship between it and more familiar time series models:

$$x_{t+\Delta} \approx a\Delta + (I_3 + A\Delta)x_t + e_{t+\Delta},$$

where $e_{t+\Delta}$ is approximately Gaussian with mean 0 and

$$\text{Cov}(e_{t+\Delta}) = C \text{diag}(b + Bx_t) C^\top \Delta = \Omega_0 + \Omega_1 x_{3t}.$$

Thus the model is a first order vector autoregression with a standard deviation that follows a first order autoregressive process. Thus the underlying state process is akin to an ARCH-M process (Engle).

The qualitative behavior of this stochastic drift and volatility model is summarized by the fact that it has a single stochastic volatility ($m = 1$) and its drift is governed by the eigenvalues of A which equal $(A_{22} \pm \sqrt{A_{22}^2 + 4A_{21}})/2$ and A_{33} . Stationarity requires that A_{21} , A_{22} and A_{33} all be negative. Given stationarity, the volatility process is expected to converge exponentially to its long-run mean of $-a_3/A_{33}$. Also, the spot price process will exhibit cyclic behavior if $A_{21} < -A_{22}^2/4$.

To make the model seasonal, we allow a_2 and a_3 to be seasonal. This allows the level of the drift and volatility to be seasonal without causing an explosion of model parameters. The assumption that A , B , and C are not seasonal means that the responses of price to the factors (the elasticities) do not depend on the time-of-year. In this paper, seasonal parameters are specified as Fourier functions of time; i.e., the seasonality function $c(t)$ is represented by

$$c(t) = c^0 + \sum_{i=1}^k c^{is} \sin(2\pi it) + c^{ic} \cos(2\pi it), \quad (11)$$

where k is an integer number that specifies the number of sin and cos terms needed, and t is the calendar time.

4 Model Estimation

All observed futures prices are used in estimation. The data can be viewed as having a panel structure, with observations indexed by the time-to-maturity (τ) and trade date (t). Two main issues are encountered in the estimation of affine SV models. First, the state transition probability density, which measures how the states evolve from one trade date to the next, does not, in general, have an analytic expression, which makes exact maximum likelihood estimation impractical, especially for multivariate models.³ Second, the state variables are latent in most cases, so some method of filtering the state variables from the observable data is necessary. Moreover, the number of observed variables is generally greater than the number of state variables, so the model cannot fit all of the observed variables exactly and hence assumptions need to be made about the cross-sectional error structure.

Denote e^m as the measurement error, then it follows that, on each day, t ,

$$F(x, \tau; T) - \beta_0(T - t) = \beta(T - t)x(t) + e^m(t), \quad (12)$$

where F , β_0 , and $e^m(t)$ are $n(t)$ vectors; β is an $n(t) \times N$ matrix. When the model is correctly specified, measurement error can nonetheless arise. Causal factors include bid/ask spreads and other transactions frictions and non-equilibrium prices due to non-synchronous measurement or illiquidity. In addition, errors can arise from incorrect model assumptions. Clues concerning model inadequacies can be obtained by examining the size and behavior of the measurement error. This paper assumes that the measurement error arises from transaction frictions alone and that it is serially and cross-sectionally independent, i.e., $e^m(t)$ is normally distributed with mean 0 and variance $H(t)$. Given this assumption, the state variables at time t can be filtered as

³When $B = 0$, the model is Gaussian and the transition distribution can be written in closed form. Also, the (multi-factor) CIR model (A , B and C diagonal and $b = 0$) is a special case in which each factor has a non-central χ^2 conditional distribution.

GLS estimator of $x(t)$ based on futures price at time t alone, i.e.,

$$\hat{x}(t) = \left(\beta (T - t)' H (t)^{-1} \beta (T - t) \right)^{-1} \beta (T - t)' H (t)^{-1} (\ln F (x, \tau; T) - \beta_0 (T - t)), (13)$$

Prediction errors are calculated as the differences between the market observed log of futures prices and the model predicted ones. For a given set of parameters, a time series of filtered state variables is first calculated, then the one-step ahead predictions of the state variables can be calculated with Euler approximation,

$$E (x (t + h) | F (x, \tau; T)) = \hat{x}(t) + (a + A\hat{x}(t)) h, (14)$$

and the one-step ahead prediction error of the log of the futures prices can then be calculated as following,

$$e^p(t) = \ln F (x, \tau - h; T) - \beta_0 (T - t - h) + \beta (T - t - h) E (x (t + h) | F (x, \tau; T)). (15)$$

Affine diffusion models have implications about the evolution of the state variables over time and the futures prices term structure across time-to-maturity. The former is defined under the actual probability measure, the later under the risk neutral probability measure. Measurement error alone only provides information about the risk neutral probability measure; prediction error, as a combination of state evolution uncertainty and measurement error, incorporates information under both probability measures. This paper applies a nonlinear least squares estimation approach, minimizing the sum of squared prediction errors of the log of the futures prices. This ensures that the model is fitted to both the cross-section and time series aspects of the futures price data. This nonlinear least squares estimator is unbiased.

5 Data and Empirical Results

5.1 Data

The approach was applied to Chicago Board of Trade corn futures prices over the period 1/1/1975 through 3/10/1998. Wednesday prices were used except in 10 case when the market was closed, in which case Tuesday prices were substituted, for a total of 1209 trading dates.⁴

A total of 8158 price observations were used, averaging 6.7 contracts per trade date. The number of contract prices used per day ranged between 4 and 10, with most of the dates recording between 6 and 8 contract maturities.⁵ The maximum time to maturity was nearly 3 years. Prices during the delivery month were excluded and the delivery date was taken to be the 15th of the month.

5.2 Empirical Results

The model given by (10) is estimated with a_2 and a_3 each having 3 sin and cos terms and $H(t) = \sigma^2 I$. Estimation results can be summarized in four aspects: the prediction error, the measurement error, dynamics of the filtered state variables, and fit of the term structures.

The model is able to achieve a root mean squared prediction error of 2.5%. This represents a considerable improvement over the naive, no-change model, which has a root mean squared prediction error of 3.3%. Figures 1 display the values of the prediction errors across time. The prediction error appears to be serially correlated and heteroskedastic, which is consistent with the model. Figure 2 shows square of prediction error as a function of time-to-maturity, which comes from a fourth order

⁴The missing Wednesday dates were 7/4/1979, 7/4/1984, 12/25/1985, 1/1/1986, 7/4/1990, 12/25/1991, 1/1/1992, 4/27/1994, 12/25/1996, and 1/1/1997.

⁵Details of the number of contracts traded per day are:

| Number of Maturities | 4 | 5 | 6 | 7 | 8 | 9 | 10 | Total |
|------------------------|---|-----|-----|-----|-----|-----|----|-------|
| Number of Observations | 4 | 140 | 416 | 394 | 123 | 116 | 16 | 1209 |

polynomial fit. The figure shows that squared prediction error decrease as time-to-maturity increases up to about 2 years, which is consistent with the notion that current economic situation has less impact on longer time-to-maturity contract. The slightly increase in squared prediction error for time-to-maturity longer than 2 years may be related to the lack of data for long time-to-maturity contracts.

The root mean squared measurement error is 0.97%. Measurement errors are plotted in Figures 3. Figure 3 shows that the measurement error is essentially unbiased, although the model does not fit the long time-to-maturity contracts (above one and a half year) well, which may also be attributable to the lack of data for those contracts. Square of measurement error is fitted as a linear function of a fifth order polynomial of the time-to-maturity, which is plotted in figure 4. The figure indicates that a heteroscedastic measurement error structure may be more appropriate.

Filtered state variables are shown in Figure 5. Log of spot price has a sample average of 0.96; the average instantaneous drift rate is 0.035; the average instantaneous volatility is 0.26. The correlation between log of the spot price and the instantaneous drift rate is -0.38, between log of the spot price and the instantaneous volatility is 0.29. The signs of these correlations are consistent with storage theory, which suggests that the spot price is high when stock levels are low. At such times the speed of mean reversion is slow and market price is more sensitive to demand and supply shocks. Filtered spot price and the nearby futures price is presented in Figure 6. Spot price does follow nearby futures price very closely. The eigenvalues of the A matrix are -0.5142, -2.5201, and -3.6557; as all are strictly negative, the model is stationary.

Figure 7 shows the impact of each factor in determining futures prices at different time-to-maturities (τ); these functions are interpretable as the partial elasticities of the factors for futures prices maturing in τ periods. The factor loading function associated with the spot price drops monotonically whereas those associated with the instantaneous drift and volatility first increase and then decrease. The former is most

influential on futures maturing in about one and a half years, the latter in about 6-7 months. Figure 8 displays the $\beta_0(t)$ functions for the 5 contract months. The seasonality of the futures prices are evident in these functions.

Depending on the season and the underlying economic condition, the term structure of futures prices can take on various shapes. Term structures on the first Wednesday of every even month in 1976 and 1988 are plotted in Figure 9, 10, 11, 12. Those term structures are representative. From the pictures, it is clear that there is a strong seasonal pattern in the term structures, which is due to the fact that agricultural commodities are seasonally produced, and hence the inventory level and the demand/supply relationship vary seasonally.

Normally, the term structure for contracts mature earlier in the year (Mar., May, and Jul.) tends to be upward sloping, since inventory level keeps dropping between two harvest season, and an upward sloping term structure supports the economic activity of holding stock; Sep. and Dec. contract usually have lower prices, which is due to the fact that new harvest starts from Aug. and lasts till Dec. During this period, inventory level keep increasing.

However, in some extreme cases, e.g., in 1988, this seasonal pattern can be violated. Early in 1988, the price level started very low (the nearby futures price is 2 dollars per bushel in Feb, but the sample average of nearby futures price is about 2.6 dollars per bushel) and the whole term structure is upward sloping (no drop in the Sep. and Dec. contract) indicating that the market expected the price level to increase over time. At this time, with extreme low price and high inventory, cost-of-carry has huge impact in determining the term structure of futures prices. Upward sloping term structure is consistent with the storage theory. Then, as summer time approached, knowing that it was going to be a big drought year, the term structure for futures contracts mature before the next harvest season became even more upward sloping, indicating that the market expected the price level to rise dramatically. Con-

sistent with this market expectation, the nearby futures price rises to nearly 3 dollars per bushel in Aug. Comparing the term structures in 1976 with those in 1988, it is clear that demand and supply relationship is the fundamental factor that determines the term structure of futures prices, and seasonality in futures price term structure reflects the seasonality in the underlying economic condition. This again explains why it is better to model a whole commodity market with some common underlying factors rather than directly model individual futures contract.

Comparing the observed vs. model implied futures price termstructure, figures 9, 10, 11, 12 show that the model is able to model various shapes of the term structures reasonably well, which indicates that the latent factor approach is effective and the time varying parameters is successful in capturing the strong seasonality exhibited in the market.

6 Summary

This paper takes a latent factor approach in modeling commodity futures price term structure. This approach assumes that all commodity prices including spot, futures, and options, etc., are determined by some common, underlying factors, and hence it provides an integrated framework in modeling the cross-sectional relationships among them. Furthermore, by modeling the dynamics of the underlying factors, it overcomes a typical problem faced by the traditional time series approach (which do not apply the inter-temporal arbitrage relationship) that futures and options prices are all double indexed, consequently, they do not form standard time series and the time-to-maturity effect can not modeled in a natural manner.

To improve upon previous work on latent factor models in the agricultural economics literature, this paper considers seasonality and stochastic volatility in the latent factor model framework. They allow the models to address the strong annual

seasonality and non-constant volatility phenomena observed in the commodity markets. Particularly, a special case of affine diffusion models is considered in which the three underlying factors are interpreted as the log of the spot price, its instantaneous drift rate and volatility. Applied to the Chicago Board of Trade corn futures price data, the model is able to fit the futures price term structures under various economic situations. Particularly, the seasonal pattern in the term structures is captured well. The advantage of allowing stochastic volatility, however, is hard to evaluate in the context of pricing commodity futures. Further research is necessary in comparing the effectiveness of this model vs. the Gaussian three-factor models in pricing commodity options and in real option analysis.

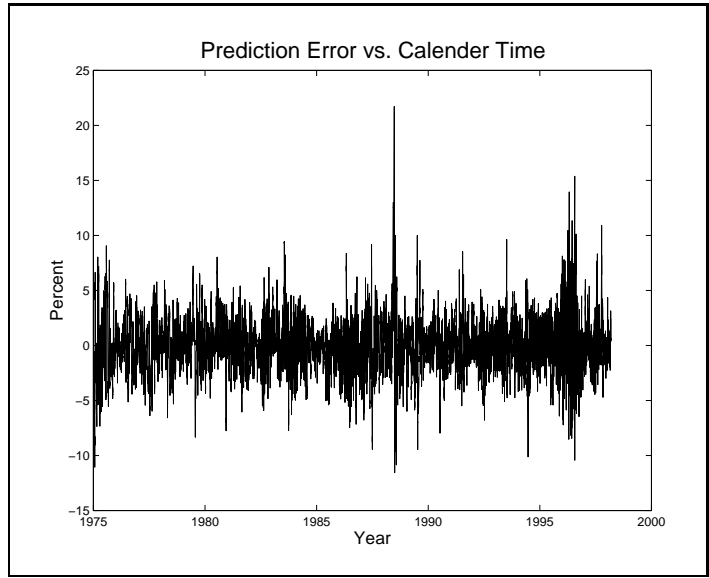


Figure 1




Figure 2




Figure 3

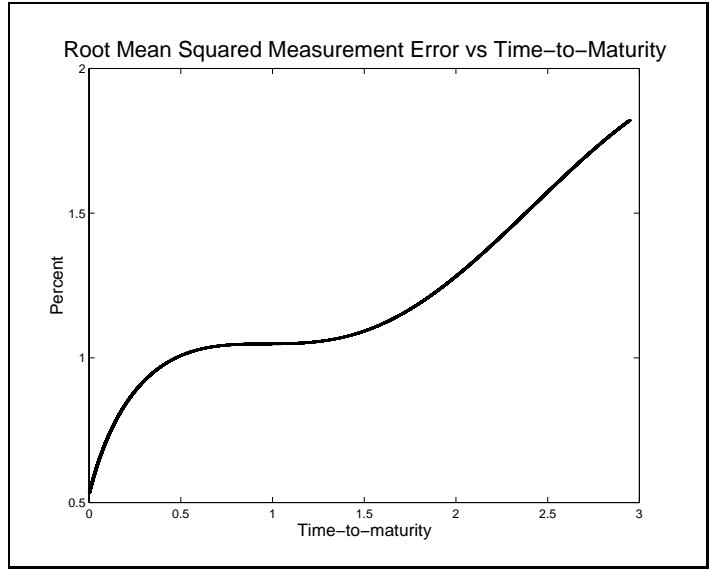


Figure 4

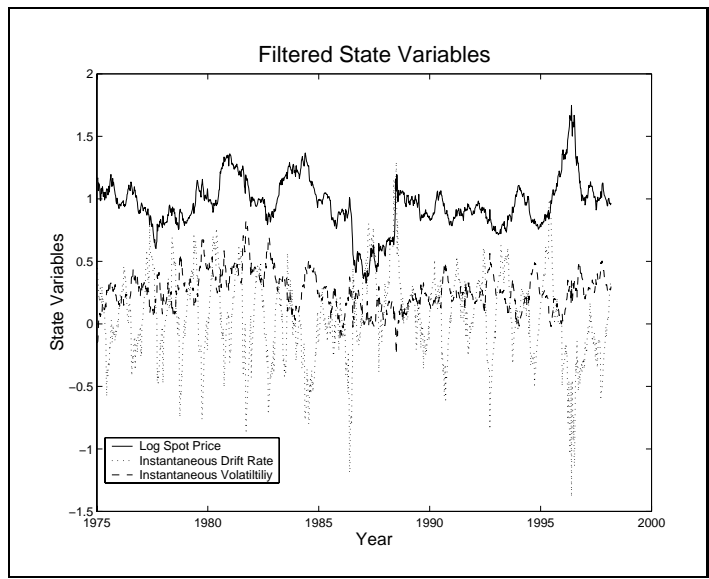


Figure 5

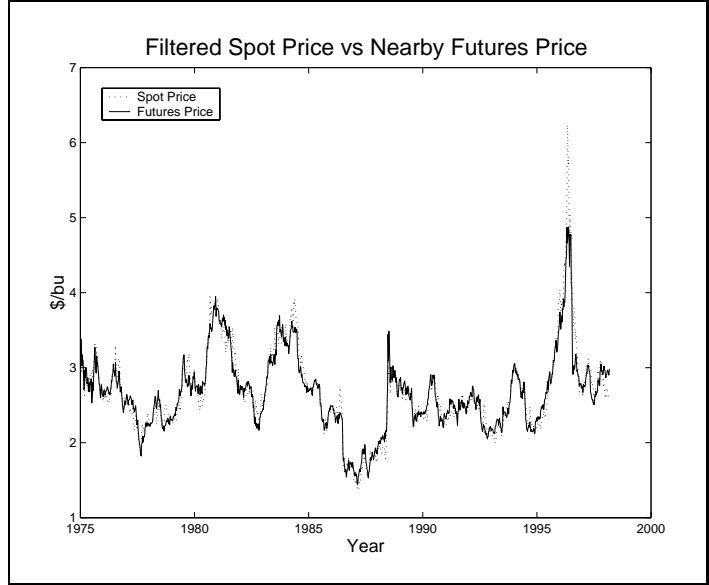


Figure 6

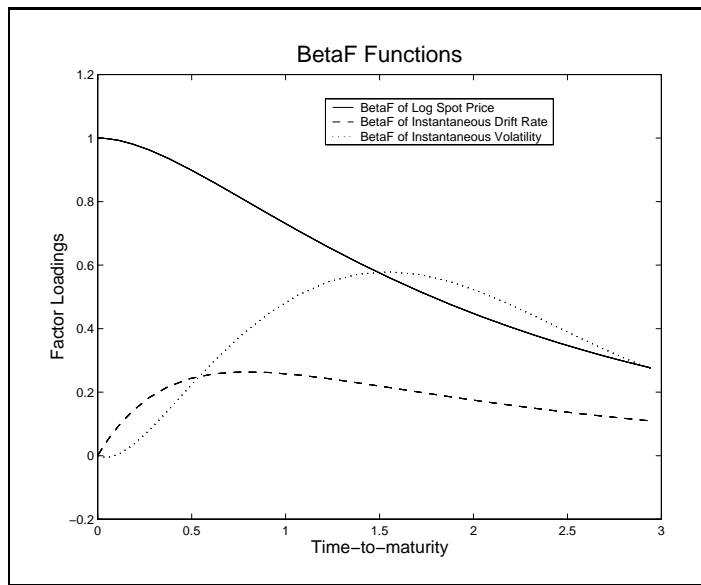


Figure 7

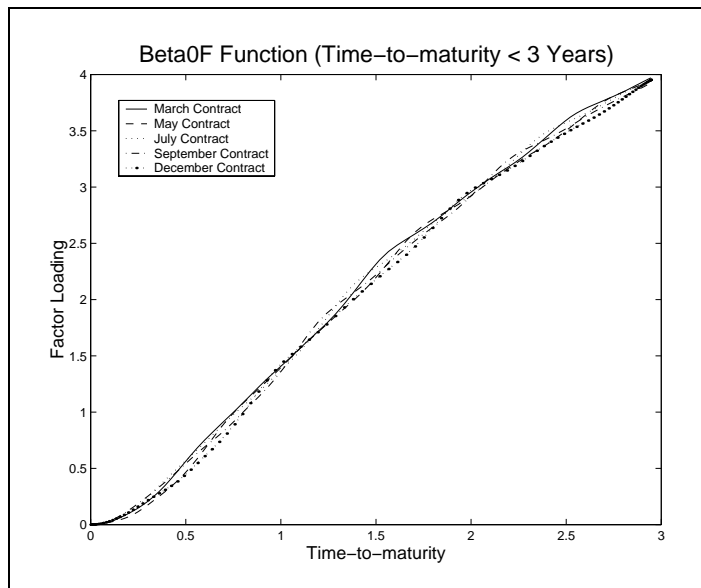


Figure 8

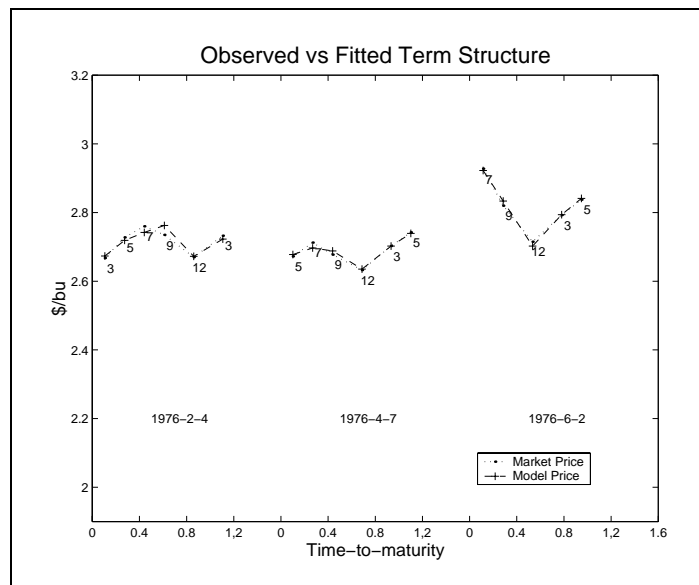


Figure 9




Figure 10

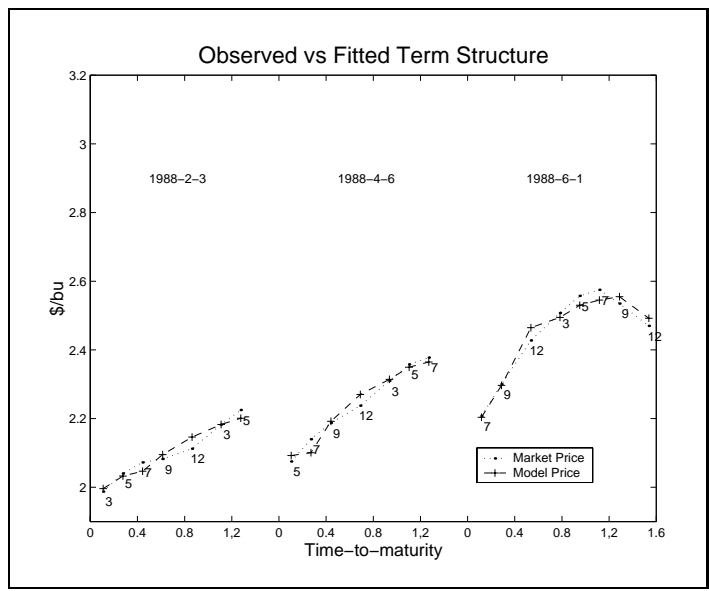


Figure 11

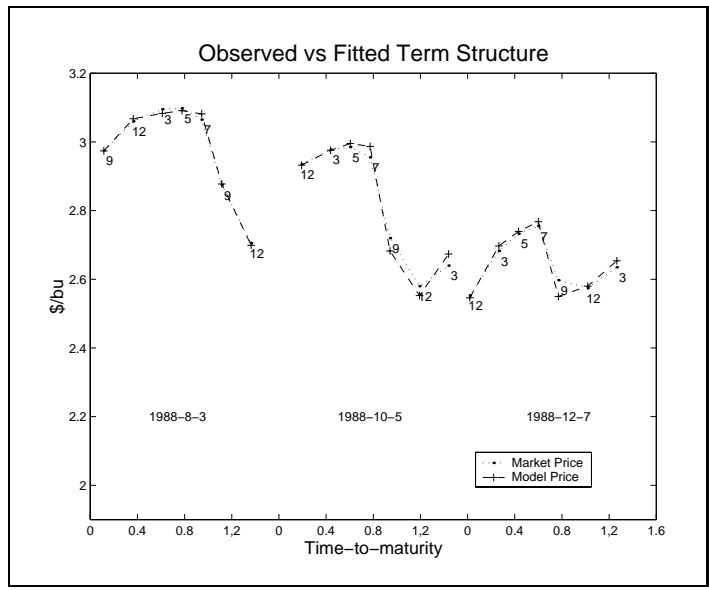


Figure 12

Table 1: Parameter Estimates

| Parameter | Coefficient |
|------------|-------------|
| a_2^0 | 6.15018 |
| a_2^{1s} | 0.3535 |
| a_2^{1c} | -0.3014 |
| a_2^{2s} | 0.6510 |
| a_2^{2c} | -0.2576 |
| a_2^{3s} | 0.0682 |
| a_2^{3c} | 1.4010 |
| a_3^0 | 1.0165 |
| a_3^{1s} | 0.3126 |
| a_3^{1c} | 0.6385 |
| a_3^{2s} | -0.1284 |
| a_3^{2c} | 0.1018 |
| a_3^{3s} | 0.0448 |
| a_3^{3c} | -0.3312 |
| A_{21} | -1.2959 |
| A_{22} | -3.0343 |
| A_{23} | 5.5672 |
| A_{33} | -3.6557 |
| C_{12} | -0.7755 |
| C_{13} | -0.1068 |
| C_{21} | 1.2361 |
| C_{23} | 0.7102 |
| b_2 | 0.0140 |
| b_3 | -1.3658 |
| B_{23} | 0.6280 |
| B_{33} | 3.9757 |
| d_1 | 0.2207 |
| d_2 | -0.2006 |
| d_3 | -0.0830 |

References

- Baillie, R. T. and R. J. Myers. "Bivariate Garch Estimation of the Optimal Commodity Futures Hedge." *Journal of Applied Econometrics* 6(1991):109–124.
- Brennan, Michael J. and Eduardo S. Schwartz. "Evaluating Natural Resource Investments." *Journal of Business* 58(1985):135–157.
- Chen, Ren-raw and Louis Scott. "Maximum-Likelihood Estimation of a Multi-Factor Equilibrium Model of the Term Structure of Interest Rates." *Journal of Fixed Income* 3(1993):14–31.
- Cortazar, Gonzalo and Eduardo S. Schwartz. "The Valuation of Commodity-Contingent Claims." *Journal of Derivatives* 1(1994):27–39.
- Cox, John C., Jonathan E. Ingersoll, and Stephen A. Ross. "The Relation Between Forward and Futures Prices." *Journal of Financial Economics* 9(1981):321–346.
- Cox, John C., Jonathan E. Ingersoll, and Stephen A. Ross. "A Theory of the Term Structure of Interest Rates." *Econometrica* 53(1985):385–407.
- Dai, Qiang and Kenneth J. Singleton. "Specification Analysis for Affine Term Structure Models." *Journal of Finance* (1999):forthcoming.
- Duffie, Darrell and Rui Kan. "A Yield-Factor Model of Interest Rates." *Mathematical Finance* 6(1996):379–406.
- Duffie, Darrell. *Dynamic Asset Pricing Theory*, 2nd ed. Princeton, NJ: Princeton University Press, 1996.
- Gibson, Rajna and Eduardo Schwartz "Valuation of Long Term Oil-Liked Assets," *Stochastic Models and Option Values*, ed. D. Lund and B. Oksendal, pp.73–101. New York: North Holland, 1991.
- Heston, Steven L. "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options." *Review of Financial Studies* 6(1993):327–343.
- Ho, Thomas S.Y. and S.B. Lee. "Term Structure Movements and Pricing Interest Rate Contingent Claims." *Journal of Finance* 41(1986):1011–1029.
- Hull, John and Alan White. "Pricing Interest Rate Derivative Securities." *Review of Financial Studies* 3(1990):573–592.

- Kenyon, David, Kenneth Kling, Jim Jordan, William Seale, and Nancy McCabe. “Factors Affecting Agricultural Futures Price Variance.” *Journal of Futures Markets* 7(1987):73–92.
- Miltersen, Kristian R. and Eduardo S. Schwartz. “Pricing of Options on Commodity Futures with Stochastic Term Structures of Convenience Yields and Interest Rates.” *Journal of Financial and Quantitative Analysis* 33(1998):33–59.
- Press, William H., Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes*, 2nd ed. Cambridge: Cambridge University Press, 1992.
- Schwartz, Eduardo S. “The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging.” *Journal of Finance* 52(1997):923–973.
- Streeter, Deborah H. and William G. Tomek. “Variability In Soybean Futures Prices: An Integrated Framework.” *Journal of Futures Markets* 12(1992):705–730.
- Vasicek, Oldrich. “An Equilibrium Characterization of the Term Structure.” *Journal of Financial Economics* 5(1977):177–188.