

2015

Price Mean Reversion, Seasonality, and Options Markets

Chad Hart

Iowa State University, chart@iastate.edu

Sergio H. Lence

Iowa State University, shlence@iastate.edu

Dermot J. Hayes

Iowa State University, dhayes@iastate.edu

Na Jin

Discover Financial Services

Follow this and additional works at: http://lib.dr.iastate.edu/econ_las_pubs

 Part of the [Agricultural and Resource Economics Commons](#), [Growth and Development Commons](#), and the [Regional Economics Commons](#)

The complete bibliographic information for this item can be found at http://lib.dr.iastate.edu/econ_las_pubs/70. For information on how to cite this item, please visit <http://lib.dr.iastate.edu/howtocite.html>.

This Article is brought to you for free and open access by the Economics at Iowa State University Digital Repository. It has been accepted for inclusion in Economics Publications by an authorized administrator of Iowa State University Digital Repository. For more information, please contact digirep@iastate.edu.

Price Mean Reversion, Seasonality, and Options Markets

Abstract

Options on agricultural commodities with maturities exceeding one year seldom trade. One possible reason to explain this lack of trading is that we do not have an accurate option pricing model for products where mean reversion in spot-price levels can be expected. Standard option pricing models assume proportionality between price variance and time to maturity. This proportionality is not a valid assumption for commodities whose supply response brings prices back to production costs. The model proposed here incorporates mean reversion in spot-price levels and includes a correction for seasonality. Mean reversion and seasonality are both observed in the soybean market. The empirical analysis lends strong support to the model.

Keywords

Bayesian statistics, commodity markets, mean reversion, options, seasonality

Disciplines

Agricultural and Resource Economics | Growth and Development | Regional Economics

Comments

This is a pre-copyedited, author-produced PDF of an article accepted for publication in American Journal of Agricultural Economics following peer review. The version of record is available online at: <http://dx.doi.org/10.1093/ajae/aav045>

PRICE MEAN REVERSION, SEASONALITY, AND OPTIONS MARKETS

Chad E. Hart, Sergio H. Lence, Dermot J. Hayes, and Na Jin¹

Abstract

Options on agricultural commodities with maturities exceeding one year seldom trade. One possible reason to explain the lack of trading is that we do not have an accurate option pricing model for products where mean reversion in spot price levels can be expected. Standard option pricing models assume proportionality between price variance and time to maturity. This proportionality is not a valid assumption for commodities whose supply response brings prices back to production costs. The model proposed here incorporates mean reversion in spot price levels and includes a correction for seasonality. Mean reversion and seasonality are both observed in the soybean market. The empirical analysis lends strong support to the model.

Keywords: Bayesian statistics, commodity markets, mean reversion, options, seasonality.

JEL: G13, Q11

¹ Chad E. Hart is an associate professor in the Department of Economics at Iowa State University. Sergio H. Lence is a Marlin Cole Chair of International Agricultural Economics and a professor in the Department of Economics at Iowa State University. Dermot J. Hayes is a Pioneer Chair in Agribusiness and a professor in the Departments of Economics and Finance at Iowa State University. Na Jin is a model risk manager at Discover Financial Services. We thank the editor and anonymous referees for their guidance and thoughtful comments. Correspondence should be sent to: chart@iastate.edu or shlence@iastate.edu. This project is an extension of research originally supported by a USDA Risk Management Agency grant funded in 2003.

Black (1976) followed the procedure first used by Black and Scholes (1973) to obtain a closed-form solution for the premium on futures options. Analogous to the Black-Scholes formula, Black's futures option pricing model assumes that futures price volatility increases in proportion to the square root of time. This assumption is reasonable for stocks and currencies, but is inconsistent with mean reversion in spot commodity prices. Most agricultural commodity markets demonstrate mean reversion to production costs (Bessembinder et al. (1995)), which suggests that spot prices will eventually revert to this cost because of supply response.¹ If this is true, and if price volatility is incorrectly assumed to increase in proportion to the square root of time, the value of long-term options will be overestimated.

Analytical formulas can play an important role in bringing closer together valuations of long and short traders. The formula proposed here can help in this regard by reducing uncertainty and increasing demand.² As of January 24, 2014, the percentage share of all outstanding options with maturity greater to or equal to one year was as follows: Corn 1.4%, Soybeans 0.1%, Crude Oil 18.1%, Intel Corp. 37.4%, Google Inc. 38.2%, IBM 35.5%, Ford 45.1%, and Apple Inc. 37.5%. As this data shows, long-term options on stocks are commonly traded, whereas options on key agricultural commodity futures such as corn and soybeans are seldom traded. It is not possible to replicate a distant-maturity option by rolling over short-maturity options, because the payoffs of the latter are generally different from the values of the distant maturity options at the rollover times. The popularity of long-term options on individual stocks shows that these options provide a unique contribution to option buyers. This contribution may include improved hedging of long-term risks or improved speculation on long-term price movements.

Lo and Wang (1995) noticed that the predictability of an asset's returns may greatly affect the value of an option on the asset. Counterintuitively, the drift of an asset's returns does not enter the option pricing formula, even though return predictability is usually driven by the drift. Lo and Wang resolved this apparent paradox by noting that the drift parameters associated with return predictability enter the option formula indirectly through the volatility. Lo and Wang

found that “even small amounts of predictability [in the drift term] can have a large impact on option prices, **especially for longer maturity options**” (p. 90) [emphasis ours].³

Schwartz (1997) recognized one aspect of this problem in the context of commodity futures. He had the insight that price imbalances caused by temporary shortages and surpluses would eventually disappear without affecting the long-run volatility level. For example, an oil shortage can make the convenience yield greater than the storage cost, and this can cause nearby futures prices to exceed the prices of more distant contracts.⁴ Several researchers (Miltersen and Schwartz (1998) and Hilliard and Reis (1998)) have proposed closed-form option-pricing models that incorporate reversion to the mean in the convenience yield. However, the spot price in Schwartz's (1997) two-factor model is assumed to be trending rather than mean reverting. If convenience yield is a constant, the spot price in the Schwartz model is assumed to follow a geometric Brownian motion. Therefore, such models are most likely to be relevant to exhaustible commodity markets such as gold and oil, for which Hotelling's Principle might be expected to hold.

A number of studies report evidence of mean reversion in spot commodity prices (e.g., Peterson, Ma, and Ritchie (1992); Allen, Ma, and Pace (1994); Walburger and Foster (1995); Irwin, Zulauf, and Jackson (1996); Casassus and Collin-Dufresne (2005); Wang and Tomek (2007)). Cassasus and Collin-Dufresne (2005) developed a three-factor futures model for non-seasonal commodities. Their model assumes convenience yield is a linear function of spot prices and interest rates, which induces mean-reversion in spot prices under the risk-neutral measure. Although closed-form solutions for the corresponding option valuation formulas are not available based on their model setup, the authors were able to report European option values using a Fourier inversion approach. The authors estimated the model using futures prices for silver, gold, and copper. They found empirical evidence of mean reversion in the spot prices, thereby implying lower estimates of option prices.⁵ Interestingly, recent theoretical work by Bobenrieth, Bobenrieth, and Wright (2014) demonstrates that spot prices can exhibit bubble-like

movements and paths consistent with mean reversion, when in fact neither is a feature of the underlying dynamic rational expectations storage model.

We argue that mean reversion in the spot price level as well as in the convenience yield is a key feature of agricultural commodity markets such as grains. For example, when prices are relatively high, supply will increase, which will in turn put downward pressure on prices. On the demand side, when prices are high, the quantity demanded will decrease, which will also induce prices to decrease. A similar story can be told when prices are relatively low.

In the two-factor option pricing model derived below, spot prices and convenience yields are assumed to follow the stochastic process advocated by Lence, Ott, and Hart (2013), which is a generalization of Jin et al. (2012). The first factor is the log-spot price, and the second factor is the instantaneous risk-neutral expected change in the log-spot price. As shown in Subsection 2.1 below, the second factor equals a constant minus the convenience yield. The implied convenience yield captures the fact that convenience yield is usually positively correlated with the spot price. In addition, the proposed specification allows both the spot price and the convenience yield to behave as an Ornstein-Uhlenbeck process. Allowing convenience yield to be a function of the spot commodity price leads to mean reversion of spot prices under both the historical and risk-neutral measures. This model structure also incorporates the Schwartz model as a special case.

Seasonality is known to be an empirical characteristic of most commodity markets. It is especially important for agricultural commodities with seasonal production or demand patterns. Given the same data set, models that ignore seasonality will likely induce a higher estimated volatility of the factors than their counterparts augmented with seasonality. This in turn will lead to a prediction of larger option prices.

Sørensen (2002) modeled seasonality in agricultural commodity futures. He introduced seasonality by adding a deterministic seasonal component to the commodity spot price. A closed-form futures pricing formula was derived based on his one-factor model with seasonality. Richter and Sørensen (2002) proposed a three-factor model to explore seasonal patterns in both the spot

price level and the volatility in commodity markets. They allowed a parameter in the drift term of the convenience yield to be a trigonometric function of time. Seasonality in the volatility term was incorporated by adding a deterministic trigonometric function of time. However, closed-form solutions for futures and option pricing formulas are not available based on their model setup.

Seasonality is introduced into our option pricing model by allowing the parameters in the drift terms of the two factors (spot price and convenience yield) to be a periodic function of calendar time. The evaluation of option pricing expressions can be reduced to the problem of solving ordinary differential equations (Duffie, Pan, and Singleton (2000)). Adding seasonality into the model makes the solution more complicated, since the resulting stochastic differential equations are inhomogeneous in time because the drift coefficients are functions of calendar time. However, we still get closed-form expressions for option pricing formulas, which greatly facilitate the empirical work.

A negative relationship between supply/inventories and convenience yields is predicted by the theory of storage. Thus, the convenience yield from marginal storage is high when inventory is low or supply is scarce, and the opposite is true when inventory is high or supply is large. Since commodity supply exhibits seasonality, the convenience yield is also assumed to behave as a mean-reverting process with seasonality.

We apply a Bayesian Markov Chain Monte Carlo (MCMC) algorithm to estimate our model using futures prices for soybeans. At each sample date we observe a cross section of futures prices with fixed maturity months, which are related to the latent variables through the futures and option pricing formulas. Following the ideas in Chen and Scott (1993), we assume that all but two futures contract prices are observed with measurement error. The latent value of the state variables can be filtered out at each sample date using the futures pricing formula by inversion based on the two futures contract prices observed without error. The empirical results support the hypothesis that spot prices in commodity markets are mean reverting and exhibit seasonality. The empirical work suggests that Schwartz's model has limitations for renewable

commodity markets. The impact of how the basic assumption of mean-reverting spot prices and seasonality affects the model's prediction about the price of commodity derivatives with short and long times to maturity is shown by comparing Black's, Schwartz's, and our models.

The rest of this paper is organized as follows. Section 1 uses graphical examples to intuitively show how the basic assumption of mean reversion will affect the value of option contracts. In section 2, we discuss the advocated generalization of Schwartz's two-factor model. Seasonality is also introduced into the proposed model and four two-factor models are defined. In section 3, option pricing formulas are derived. Section 4 describes the empirical model, the data set, and the estimation method. The econometric analysis and estimation results are discussed in section 5. The last section concludes the paper.

Graphical Examples

Figures 1, 2, and 3 are graphical representations of the assumptions that underlie the three models (Black, Schwartz, and our model), and they show how our model differs from those of Schwartz and Black. All three figures depict the same simulated time series of the logarithm of prices ($\ln(S_t)$), and all contain the upper and lower confidence intervals for these log-prices at two points in time.

Figure 1 shows log-prices under the standard Black assumption of Brownian motion. It can be observed that the confidence interval for log-prices increases in proportion to the square root of time. The heavy black line shows the expected log-price path as of time t_0 , demonstrating a small amount of growth as might be expected for the cash prices for some commodities or stocks. If futures markets existed for this asset, this heavy line would reflect the temporal basis. At time t_1 the realized cash price is lower than was expected at time t_0 and the heavy gray line shows the expected log-price path from this lower point. All of the price reduction from time t_0 through time t_1 is viewed as permanent in this model. Therefore, the updated expected log-price path runs parallel to the original, but at a level that reflects the underperformance of the spot price between times t_0 and t_1 .

Figure 2 illustrates Schwartz's two-factor model, and is otherwise identical to Figure 1. A key difference between Figures 1 and 2 is that when the log-price path is updated at time t_1 , Schwartz's two-factor model recognizes that the price drop that occurred just before time t_1 was in part due to a temporary reduction in the convenience yield, reflecting a temporary surplus of the commodity. The model assumes that this temporary component will gradually disappear, and therefore it adjusts the expected time path of log-cash prices for this expected price recovery. However, once this temporary adjustment vanishes, Schwartz's two-factor model behaves very much like Black's.

Figure 3 represents the model advocated here. The log-spot price path after times t_0 and t_1 contain adjustments to the temporary imbalances as in Schwartz's two-factor model. However, the model also contains one additional piece of information. It recognizes that the generally high (low) level of spot prices observed at time t_0 (t_1) is well above (below) the production costs for this commodity. This suggests an increase (a reduction) in supply until prices recover to these expected production costs. Therefore, the heavy black line approaches the heavy gray line as the model implicitly adjusts supply and demand so that expected future prices lie on the path representing expected production costs. This additional piece of information has a dramatic effect on the upper and lower confidence levels, because the model recognizes that all price deviations around these expected production costs are of a temporary nature and it therefore tightens the confidence interval around this log-spot-price path.

The upper and lower confidence intervals are directly related to the value of option contracts. Therefore, this intuitive evidence suggests that when mean reversion in the log-price level is added to mean reversion in the convenience yield, the value of option contracts will be lower. The magnitude of the option premiums' upward bias is a function of the models' parameters. But it is clear that the magnitude of the bias will increase with the options' time to expiration.

A Generalization of Schwartz's Two-Factor Model

The proposed model of premiums for futures options is based on the theoretical model of commodity price behavior recently used by Lence, Ott, and Hart (2013). Commodity prices are modeled in continuous time as a system of stochastic differential equations in an affine term structure class (see, e.g., Chapter 17 in Pennacchi (2008)). The key advantage of the affine class is that it allows for a rich variety of tractable specifications of asset pricing models.

In the advocated model, the entire futures curve is driven by two "factors" (Y_1 and Y_2) which follow a bivariate Gaussian process. Ignoring seasonality for the time being to simplify the presentation, the "historical" or "true" process for the two factors is given by

$$(1) \quad d\underline{Y}(t) = [\underline{\kappa}_0 - \underline{\kappa} \underline{Y}(t)] dt + \underline{\Sigma}^{0.5} d\underline{W}(t),$$

Where vector $\underline{Y}(t) \equiv [Y_1(t) Y_2(t)]^T$, vector $\underline{\kappa}_0 \equiv [\kappa_{10} \kappa_{20}]^T$ and matrix $\underline{\kappa} \equiv [\kappa_{11} \kappa_{12}; \kappa_{21} \kappa_{22}]$ comprise historical parameters, the positive definite matrix $\underline{\Sigma} \equiv [\sigma_1^2 \rho_{12} \sigma_1 \sigma_2; \rho_{12} \sigma_1 \sigma_2 \sigma_2^2]$ is the variance-covariance matrix comprising the standard deviations σ_i ($i = 1, 2$) and the correlation coefficient ρ_{12} , and vector $d\underline{W}(t) \equiv [dW_1(t) dW_2(t)]^T$ contains historical independent pure Brownian motions.

The model assumes that the market prices of risk of the two factors are affine functions of the underlying factors,

$$(2) \quad \underline{\Lambda}(t) = \underline{\Sigma}^{-0.5} [\underline{\lambda}_0 - \underline{\lambda} \underline{Y}(t)],$$

where $\underline{\Lambda}(t) \equiv [\Lambda_1(t) \Lambda_2(t)]^T$ is the vector comprising the market prices of risk, and vector $\underline{\lambda}_0 \equiv [\lambda_{10} \lambda_{20}]^T$ and matrix $\underline{\lambda} \equiv [\lambda_{11} \lambda_{12}; \lambda_{21} \lambda_{22}]$ contain risk-premium parameters. With the equality $d\tilde{\underline{W}}(t) = \underline{\Lambda}(t) dt + d\underline{W}(t)$ establishing the link between the risk-neutral pure Brownian motions $d\tilde{\underline{W}}(t) \equiv [d\tilde{W}_1(t) d\tilde{W}_2(t)]^T$ and the historical Brownian motions $d\underline{W}(t)$, the risk-neutral process (3) follows immediately from equations (1) and (2),

$$(3) \quad d\underline{Y}(t) = [\tilde{\underline{\kappa}}_0 - \tilde{\underline{\kappa}} \underline{Y}(t)] dt + \underline{\Sigma}^{0.5} d\tilde{\underline{W}}(t),$$

where $\tilde{\underline{\kappa}}_0 \equiv [\tilde{\kappa}_{10} \tilde{\kappa}_{20}]^T$, $\tilde{\underline{\kappa}} \equiv [\tilde{\kappa}_{11} \tilde{\kappa}_{12}; \tilde{\kappa}_{21} \tilde{\kappa}_{22}]$, and $\tilde{\kappa}_{ij} \equiv \kappa_{ij} - \lambda_{ij}$.

No-arbitrage restrictions imply that the date- t futures price for a contract maturing at time $T \equiv t + \tau \geq t$ is the risk-neutral expectation at time t of the date- T spot price, $F(t, \tau) = \tilde{E}_t[S(t + \tau)]$. Therefore, by assuming that the logarithm of the spot price ($\ln[S(t)]$) is an affine function of the factors as in

$$(4) \quad \ln[S(t)] = \phi_0 + \phi_1 Y_1(t) + \phi_2 Y_2(t),$$

for some constants ϕ_0 , ϕ_1 , and ϕ_2 , the Gaussian risk-neutral distribution of the factors (3) yields an affine structure for the log-futures curve (see, e.g., Duffie, Pan, and Singleton (2000) or Chapter 17 in Pennacchi (2008)). The logarithms of all futures prices are also affine functions of the date- t factor values, so that

$$(5) \quad F(t, \tau) = \tilde{E}_t\{\exp[\ln(S(t + \tau))]\},$$

$$(5') \quad = \tilde{E}_t\{\exp[\phi_0 + \phi_1 Y_1(t + \tau) + \phi_2 Y_2(t + \tau)]\},$$

$$(5'') \quad = \exp[A(\tau) + \underline{B}(\tau)^T \underline{Y}(t)],$$

where coefficients $A(\tau)$ and $\underline{B}(\tau) \equiv [B_1(\tau) B_2(\tau)]^T$ solve the ordinary differential equations

$$(6) \quad \frac{d\underline{B}(\tau)}{d\tau} = -\tilde{\underline{\kappa}}^T \underline{B}(\tau), \text{ and}$$

$$(7) \quad \frac{dA(\tau)}{d\tau} = \tilde{\kappa}_0^T \underline{B}(\tau) + \frac{1}{2} \underline{B}(\tau)^T \underline{\Sigma} \underline{B}(\tau),$$

subject to the boundary conditions $\underline{B}(0) = [\phi_1 \ \phi_2]^T$ and $A(0) = \phi_0$.

To be able to identify the parameters involved in processes (1) and (3), factors Y_1 and Y_2 must be normalized. The need for normalization stems from the results in Dai and Singleton (2000), which imply that there is an infinite number of *invariant affine transformations* of the factors that, if accompanied by appropriate changes in the process parameters, leave the futures curve unchanged. Collin-Dufresne, Goldstein, and Jones (2008) proposed to normalize factors in terms of theoretically observed state variables with meaningful economic interpretations, and proved that their advocated representation has a unique global maximum.

For the case of futures prices, the factor normalization proposed by Collin-Dufresne, Goldstein, and Jones (2008) consists of setting the first factor equal to the logarithm of the spot price ($Y_1(t) = \ln[S(t)]$), and the second factor equal to the expected risk-neutral instantaneous change in the log-spot price ($Y_2(t) = \tilde{E}_t\{d\ln[S(t)]\}$). This normalization implies the restrictions $[\tilde{\kappa}_{10} \ \tilde{\kappa}_{11} \ \tilde{\kappa}_{12} \ \phi_0 \ \phi_1 \ \phi_2] = [0 \ 0 \ -1 \ 0 \ 1 \ 0]$. Therefore, the normalized system's set of identifiable parameters is $\{\tilde{\kappa}_{20}, \tilde{\kappa}_{21}, \tilde{\kappa}_{22}, \kappa_{10}, \kappa_{20}, \kappa_{11}, \kappa_{12}, \kappa_{21}, \kappa_{22}, \sigma_1, \sigma_2, \rho_{12}\}$.⁶ Alternatively, the set of identifiable parameters may be stated as $\{\tilde{\kappa}_{20}, \tilde{\kappa}_{21}, \tilde{\kappa}_{22}, \lambda_{10}, \lambda_{20}, \lambda_{11}, \lambda_{12}, \lambda_{21}, \lambda_{22}, \sigma_1, \sigma_2, \rho_{12}\}$ because $\kappa_{ij} = \tilde{\kappa}_{ij} + \lambda_{ij}$.

Given the aforementioned restrictions $[\tilde{\kappa}_{10} \ \tilde{\kappa}_{11} \ \tilde{\kappa}_{12}] = [0 \ 0 \ -1]$, and assuming that $\tilde{\kappa}_{21} \neq 0$ (so that matrix $\tilde{\kappa}$ is invertible), the instantaneous drift in process (3) can be written as (8')

$$(8) \quad \tilde{E}_t \begin{bmatrix} dY_1(t) \\ dY_2(t) \end{bmatrix} = \left(\begin{bmatrix} 0 \\ \tilde{\kappa}_{20} \end{bmatrix} - \begin{bmatrix} 0 & -1 \\ \tilde{\kappa}_{21} & \tilde{\kappa}_{22} \end{bmatrix} \begin{bmatrix} Y_1(t) \\ Y_2(t) \end{bmatrix} \right) dt,$$

$$(8') \quad = \begin{bmatrix} Y_2(t) \\ \tilde{\kappa}_{21} \left[\frac{\tilde{\kappa}_{20}}{\tilde{\kappa}_{21}} - Y_1(t) \right] - \tilde{\kappa}_{22} Y_2(t) \end{bmatrix} dt.$$

Intuitively, the right-hand side of equation (8') implies that if $Y_1 = \tilde{\kappa}_{20} / \tilde{\kappa}_{21}$, $Y_2 = 0$, and there were no shocks to the system, neither factor would be expected to change over time under the risk-neutral distribution. In other words, $\tilde{\kappa}_{20} / \tilde{\kappa}_{21}$ and zero represent the long-term risk-neutral means of the first and second factors, respectively. It can also be seen that if $\tilde{\kappa}_{21} > 0$, the larger the amount by which Y_1 is above (below) its long-term risk-neutral mean $\tilde{\kappa}_{20} / \tilde{\kappa}_{21}$, the greater the risk-neutral expected reduction in Y_2 . Since Y_2 is the expected risk-neutral instantaneous change in Y_1 , this implies that Y_1 tends to revert back to the long-term risk-neutral mean ($\tilde{\kappa}_{20} / \tilde{\kappa}_{21}$) if $\tilde{\kappa}_{21} > 0$. Further, the greater the coefficient $\tilde{\kappa}_{21}$, the faster the risk-neutral mean reversion characterizing Y_1 . This is why parameter $\tilde{\kappa}_{21}$ is also labeled Y_1 's speed of mean reversion. The opposite case of $\tilde{\kappa}_{21} < 0$ implies an explosive process, in the sense that Y_1 tends to move away from $\tilde{\kappa}_{20} / \tilde{\kappa}_{21}$ if $Y_1 \neq \tilde{\kappa}_{20} / \tilde{\kappa}_{21}$. Moreover, the greater the magnitude of $|\tilde{\kappa}_{21}|$, the faster Y_1 shifts away from $\tilde{\kappa}_{20} / \tilde{\kappa}_{21}$ if $Y_1 \neq \tilde{\kappa}_{20} / \tilde{\kappa}_{21}$ and $\tilde{\kappa}_{21} < 0$. Following a similar reasoning, one can interpret parameter $\tilde{\kappa}_{22}$ as the speed of mean reversion of factor Y_2 (i.e., Y_1 's expected risk-neutral drift).

Lence, Ott, and Hart (2013) reported closed-form solutions for the futures curve corresponding to the advocated model. They consist of

$$(9) \quad A(\tau) = \alpha_0(\tau) - \alpha_0(0) + \alpha_{\underline{x}}(\tau) - \alpha_{\underline{x}}(0),$$

$$(10) \quad \underline{B}(\tau) = \begin{bmatrix} \sum_{i=1}^2 b_{1i} \exp(R_i \tau) \\ \sum_{i=1}^2 b_{2i} \exp(R_i \tau) \end{bmatrix},$$

$$\text{where: } \alpha_0(\tau) \equiv \tilde{\kappa}_{20} \sum_{i=1}^2 \frac{b_{2i}}{R_i} \exp(R_i \tau),$$

$$\alpha_{\underline{z}}(\tau) \equiv \frac{h_{12}}{(R_1 + R_2)} \exp[(R_1 + R_2) \tau] + \frac{1}{4} \sum_{i=1}^2 \frac{h_{ii}}{R_i} \exp(2 R_i \tau),$$

$$R_1 \equiv (-\tilde{\kappa}_{22} + \sqrt{\tilde{\kappa}_{22}^2 - 4\tilde{\kappa}_{21}})/2, R_2 \equiv (-\tilde{\kappa}_{22} - \sqrt{\tilde{\kappa}_{22}^2 - 4\tilde{\kappa}_{21}})/2, b_{1i} \equiv (R_i + \tilde{\kappa}_{22}) b_{2i}, b_{2i} \equiv 1/(2 R_i + \tilde{\kappa}_{22}), \text{ and } h_{ij} \equiv \sigma_1^2 b_{1i} b_{1j} + \sigma_2^2 b_{2i} b_{2j} + \rho_{12} \sigma_1 \sigma_2 (b_{1i} b_{2j} + b_{1j} b_{2i}) \text{ for } i, j \in \{1, 2\}.$$

Schwartz's Two-Factor Model

Based on Kaldor (1939)'s fundamental insight that commodity prices are characterized by convenience yields, Schwartz (1997) postulated a path-breaking two-factor model with the log-spot price as the first factor, and the convenience yield net of storage cost ($c(t)$) as the second factor. Schwartz derived futures prices based on this model. The corresponding option prices were obtained by Miltersen and Schwartz (1998) and Hilliard and Reis (1998).

Although Schwartz's second factor is different from the second factor advocated here, it is straightforward to prove that Schwartz's model is a special case of the present one. To see this point, note that the instantaneous risk-neutral expected rate of return to the commodity holder consists of the risk-neutral expected relative price change plus the net convenience yield, and it must be equal to the risk-free rate (r). That is, equality $\tilde{E}_t[dS(t)]/S(t) + c(t) = r$ must hold. Further, by application of Itô's lemma one obtains $Y_2(t) = \tilde{E}_t\{d\ln[S(t)]\} = \tilde{E}_t[dS(t)]/S(t) - \sigma^2/2$, where $\sigma^2 \equiv \sigma_1^2 + 2\rho_{12}\sigma_1\sigma_2 + \sigma_2^2$. Therefore, the net convenience yield can be expressed as $c(t) = r - Y_2(t) - \sigma^2/2$, which implies that Schwartz's two-factor model can be alternatively written in terms of the present model's second factor $Y_2(t) = \tilde{E}_t\{d\ln[S(t)]\}$ instead of the net convenience yield.

More specifically, Schwartz's model is isomorphic to the advocated one (i.e., equations (2.1)-(2.4) with $Y_1(t) = \ln[S(t)]$, $Y_2(t) = \tilde{E}_t\{d\ln[S(t)]\}$, and $[\tilde{\kappa}_{10} \tilde{\kappa}_{11} \tilde{\kappa}_{12} \phi_0 \phi_1 \phi_2] = [0 \ 0 \ -1 \ 0 \ 1 \ 0]$), under the additional restrictions $\tilde{\kappa}_{21} = \lambda_{11} = \lambda_{12} = \lambda_{21} = \lambda_{22} = 0$ (or, alternatively, $\tilde{\kappa}_{21} = \kappa_{11} = \kappa_{21} = 0$, $\kappa_{12} = -1$, and $\kappa_{22} = \tilde{\kappa}_{22}$). The constraints $\tilde{\kappa}_{21} = \lambda_{21} = 0$ imply that neither the risk-neutral nor the historical log-spot price process is allowed to be mean reverting. Given (8') and the

relationship $c(t) = r - Y_2(t) - \sigma^2/2$, it is also clear that such restrictions prevent net convenience yields from being positively associated with spot prices. This is an undesirable feature of Schwartz's model because when a commodity is in relatively short (plentiful) supply, its price is typically high (low) and its net convenience yield is high (low), as well. Ultimately, however, whether the restrictions are warranted or not is an empirical matter. Testing whether $\tilde{\kappa}_{21} = \lambda_{21} = 0$ holds allows us to determine whether the spot price is mean reverting or not in a given market.⁷

Seasonality

So far we have assumed that all parameters are constant throughout the year. However, most commodity markets differ from the markets for stocks, bonds, and other conventional financial assets, in that they usually exhibit seasonal patterns. For example, prices for annual crops are typically high in the pre-harvest season and low at peak-harvest. To capture this feature, periodicity is introduced in selected parameters by a truncated Fourier series. We incorporate seasonality by adding periodic functions of calendar time to parameters $\tilde{\kappa}_{20}$ (which determines the long-term risk-neutral means of the spot log price and its instantaneous drift), and λ_{10} and λ_{20} (which determine the long-term historical means). The specific expressions are

$$(11) \quad \tilde{\kappa}_{20}(t) = \tilde{\kappa}_{20} + \sum_{h=1}^2 \left[\tilde{\kappa}_{20,h,\sin} \sin\left(\frac{2\pi h}{\delta} t\right) + \tilde{\kappa}_{20,h,\cos} \cos\left(\frac{2\pi h}{\delta} t\right) \right] \text{ and}$$

$$(12) \quad \lambda_{i0}(t) = \lambda_{i0} + \sum_{h=1}^2 \left[\lambda_{i0,h,\sin} \sin\left(\frac{2\pi h}{\delta} t\right) + \lambda_{i0,h,\cos} \cos\left(\frac{2\pi h}{\delta} t\right) \right],$$

for $i = \{1, 2\}$, where $\tilde{\kappa}_{20} \equiv \{ \tilde{\kappa}_{20}, \tilde{\kappa}_{20,1,\sin}, \tilde{\kappa}_{20,1,\cos}, \tilde{\kappa}_{20,2,\sin}, \tilde{\kappa}_{20,2,\cos} \}$ and $\lambda_{i0} \equiv \{ \lambda_{i0}, \lambda_{i0,1,\sin}, \lambda_{i0,1,\cos}, \lambda_{i0,2,\sin}, \lambda_{i0,2,\cos} \}$ are sets of parameters, and δ denotes the length of the periodic time interval (e.g., $\delta = 12$ if time is measured in months and the focus is on annual seasonality). Note

that the model does not exhibit seasonality if $\tilde{\kappa}_{20,h,\sin} = \tilde{\kappa}_{20,h,\cos} = \lambda_{i0,h,\sin} = \lambda_{i0,h,\cos} = 0$ for $i = \{1, 2\}$ and $h = \{1, 2\}$.

The ordinary differential equations (6) and (7) imply that the proposed seasonality structure only affects the $A(\cdot)$ coefficients. Given (11), the solution to (7) is

$$(13) \quad A(t, \tau) = \alpha_0(\tau) - \alpha_0(0) + \alpha_{seas}(t, \tau) - \alpha_{seas}(t + \tau, 0) + \alpha_{\underline{x}}(\tau) - \alpha_{\underline{x}}(0),$$

$$\text{where: } \alpha_{seas}(t, \tau) \equiv \sum_{h=1}^2 [\tilde{\kappa}_{20,h,\sin} \alpha_{h,\sin}(t, \tau) + \tilde{\kappa}_{20,h,\cos} \alpha_{h,\cos}(t, \tau)],$$

$$\alpha_{h,\sin}(t, \tau) \equiv \sum_{i=1}^2 \frac{b_{2i}}{(2\pi h / \delta)^2 + R_i^2} [R_i \sin(\frac{2\pi h}{\delta} t) + \frac{2\pi h}{\delta} \cos(\frac{2\pi h}{\delta} t)] \exp(R_i \tau),$$

$$\alpha_{h,\cos}(t, \tau) \equiv \sum_{i=1}^2 \frac{b_{2i}}{(2\pi k / \delta)^2 + R_i^2} [R_i \cos(\frac{2\pi h}{\delta} t) - \frac{2\pi h}{\delta} \sin(\frac{2\pi h}{\delta} t)] \exp(R_i \tau),$$

for $h \in \{1, 2\}$. Comparison of expressions (9) and (13) reveals that they are the same, except for the extra $\alpha_{seas}(\cdot)$ terms in the latter. Such terms are a function of both time to maturity and calendar time, and collapse to zero in the absence of seasonality.

The risk-neutral process (2.3) with factors $[Y_1(t) Y_2(t)] = [\ln[S(t)] \tilde{E}_t\{d\ln[S(t)]\}]$ incorporating seasonality provides us the basic foundation for pricing options on commodity futures contracts, which we address in the following section.

Options on Futures Contracts

Let $C[F(t, \tau_0 + \Delta T), K, \tau_0]$ for $\tau_0 \in [0, \tau]$ denote the price at time t of a European call option with a strike price of K expiring at time $T_O \equiv t + \tau_0$ on a futures contract that matures at time $T = t + \tau_0 + \Delta T$, where $\Delta T \equiv T - T_O = \tau - \tau_0$. At expiration, the payoff of such an option is $\max[F(t +$

$\tau_0, \Delta T) - K, 0]$. Standard no-arbitrage arguments can be applied to show that the call premium at time t is given by

$$(14) \quad C[F(t, \tau_0 + \Delta T), K, \tau_0] = \exp(-r \tau_0) \tilde{E}_t\{\max[F(t + \tau_0, \Delta T) - K, 0]\}.$$

The analytical solution for the call option premium can be computed by applying the method proposed by Duffie, Pan, and Singleton (2000), as we do next.

To solve for the risk-neutral expectation on the right-hand side of (14), note that the moment-generating function of the log-futures price at date $(t + \tau_0)$ under the equivalent martingale measure is defined by

$$(15) \quad M_{\ln[F(t+\tau_0, \Delta T)]}(z) \equiv \tilde{E}_t\{\exp[z \ln(F(t + \tau_0, \Delta T))]\},$$

$$(15') \quad = \tilde{E}_t\{\exp[z (A(t + \tau_0, \Delta T) + \underline{B}(\Delta T)^T \underline{Y}(t + \tau_0))]\},$$

$$(15'') \quad = \exp\{\text{Mean}[\underline{Y}(t), t, \tau_0, \Delta T] z + \frac{1}{2} \text{Var}(\tau_0, \Delta T) z^2\},$$

$$\begin{aligned} \text{where: } \text{Mean}[\underline{Y}(t), t, \tau_0, \Delta T] &\equiv A(t, \tau_0 + \Delta T) + \underline{B}(\tau_0 + \Delta T)^T \underline{Y}(t), \\ &= \ln[F(t, \tau_0 + \Delta T)] - 0.5 \text{Var}(\tau_0, \Delta T), \\ &= F(t, \tau_0 + \Delta T), \end{aligned}$$

$$\text{Var}(\tau_0, \Delta T) \equiv 2 [\alpha_{\underline{z}}(\tau_0 + \Delta T) - \alpha_{\underline{z}}(\Delta T)],$$

and z is a real number in a neighborhood around zero. Expression (15') is obtained by using the futures price formula (5'') to substitute for $F(t + \tau_0, \Delta T)$ in equation (15). The expectation term on the right-hand side of equation (15') is of the same form as equation (3) in Duffie, Pan, and

Singleton (2000), so their method can be applied to derive the analytical expression (15") (see the supplemental appendix online for details).

Moment-generating function (15") reminds us that, given the information set at date t , the log-futures price $\ln[F(t + \tau_0, \Delta T)]$ is risk-neutrally distributed as a normal random variable with mean $Mean[\underline{Y}(t), t, \tau_0, \Delta T] = F(t, \tau_0 + \Delta T)$ and variance $Var(\tau_0, \Delta T)$.^{8,9} Armed with such a function, it is straightforward to derive from equation (3.1) the following analytical solution for the call price (see the supplemental appendix online)

$$(16) \quad C[F(t, \tau), K, \tau_0] = \exp(-r \tau_0) [F(t, \tau) N(d_1) - K N(d_2)],$$

where $N(\cdot)$ is the standard normal cumulative distribution, $d_1 \equiv \{\ln[F(t, \tau)/K] + 0.5 Var(\tau_0, \tau - \tau_0)\}/[Var(\tau_0, \tau - \tau_0)]^{0.5}$, and $d_2 \equiv \{\ln[F(t, \tau)/K] - 0.5 Var(\tau_0, \tau - \tau_0)\}/[Var(\tau_0, \tau - \tau_0)]^{0.5}$.

Notably, as discussed by Lo and Wang (1995, p. 95), formula (16) for the premium of a call option is of the same form as Black's option pricing equation. Given a strike price K , the underlying futures price $F(t, \tau)$, and the time to option maturity τ_0 , the option price depends on $Var(\tau_0, \Delta T)$, which is the variance of the logarithm of $F(t + \tau_0, \Delta T)$ as of time t . However, the variance depends on the risk-neutral speed of mean reversion of the log-spot price $\tilde{\kappa}_{21}$. For a mean-reverting process (i.e., $\tilde{\kappa}_{21} > 0$), the variance is bounded as τ_0 goes to infinity, whereas the variance grows without limit with τ_0 if $\tilde{\kappa}_{21} = 0$ (i.e., the cases of geometric Brownian motion and Schwartz's model).¹⁰

The premium for an European futures put (obtained in a similar manner) can be written as

$$(17) \quad P[F(t, \tau), K, \tau_0] = C[F(t, \tau), K, \tau_0] + \exp(-r \tau_0) [K - F(t, \tau)],$$

$$(17') \quad = \exp(-r \tau_0) [K N(-d_2) - F(t, \tau) N(-d_1)].$$

Formula (17') follows from the futures call premium (16) and the fact that $[1 - N(d)] = N(-d)$.

Estimation

Recalling that $\kappa_{ij} = \lambda_{ij} + \tilde{\kappa}_{ij}$, under the proposed normalization (i.e., $[\tilde{\kappa}_{10} \ \tilde{\kappa}_{11} \ \tilde{\kappa}_{12}] = [0 \ 0 \ -1]$) and allowing for seasonality, the historical process (1) can be written in matrix form as

$$(18) \quad \begin{bmatrix} dY_1(t) \\ dY_2(t) \end{bmatrix} \sim N \left(\left(\begin{bmatrix} \lambda_{10}(t) \\ \lambda_{20}(t) + \tilde{\kappa}_{20}(t) \end{bmatrix} - \begin{bmatrix} \lambda_{11} & -1 + \lambda_{12} \\ \lambda_{21} + \tilde{\kappa}_{21} & \lambda_{21} + \tilde{\kappa}_{21} \end{bmatrix} \begin{bmatrix} Y_1(t) \\ Y_2(t) \end{bmatrix} \right) dt, \ \underline{\Sigma} dt \right).$$

In practice, it is impossible to estimate model (18) because it requires continuous-time data.

Therefore, for estimation purposes we resort to its first-order discretized version

$$(19) \quad \begin{bmatrix} Y_1(t+1) - Y_1(t) \\ Y_2(t+1) - Y_2(t) \end{bmatrix} = \left(\begin{bmatrix} \lambda_{10}(t) \\ \tilde{\kappa}_{20}(t) + \lambda_{20}(t) \end{bmatrix} - \begin{bmatrix} \lambda_{11} & -1 + \lambda_{12} \\ \tilde{\kappa}_{21} + \lambda_{21} & \tilde{\kappa}_{22} + \lambda_{22} \end{bmatrix} \begin{bmatrix} Y_1(t) \\ Y_2(t) \end{bmatrix} \right) + \begin{bmatrix} \varepsilon_1(t+1) \\ \varepsilon_2(t+1) \end{bmatrix},$$

$$\begin{bmatrix} \varepsilon_1(t+1) \\ \varepsilon_2(t+1) \end{bmatrix} \text{ i.i.d. } N(\underline{0}_{(2 \times 1)}, \underline{\Sigma}),$$

where $\underline{0}_{(2 \times 1)}$ is a (2×1) vector of zeroes and the time interval used for the discretization is set equal to one. The strategy to estimate equation (19) using observed futures prices and option premiums is explained next.

Suppose we have a data set consisting of historical observations on date- t futures prices for $m > 2$ maturities and n_{Dates} , i.e., $F_{t,\tau}$ for $t \in \{1, \dots, n_{Dates}\}$ and $\tau \in \{1, \dots, m\}$. Among the m futures contracts with distinct maturity dates, it is assumed that two of the log-futures are perfectly correlated with the vector of state variables $\underline{Y}(t)$, and the remaining $(m - 2)$ log-futures are observed with normally distributed errors $\underline{e}_t^F \equiv [e_t^{F_1} \dots e_t^{F_{m-2}}]^T \sim N(\underline{0}_{((m-2) \times 1)}, \sigma_F^2 \underline{\Omega})$, where $\sigma_F^2 > 0$ is a scalar, and $\underline{\Omega}$ is an $((m - 2) \times (m - 2))$ matrix with element (i, j) equal to $\rho_F^{|i-j|}$ for $\rho_F \in (-1, 1)$. Then, the estimation system consists of

$$(20) \quad \underline{X}_t^\bullet = \underline{A}^\bullet + \underline{B}^\bullet \underline{Y}(t) + \underline{e}_t^F, \text{ and}$$

$$(21) \quad \underline{X}_t^\circ = \underline{A}^\circ + \underline{B}^\circ \underline{Y}(t),$$

where $\underline{X}_t^\bullet \equiv [\ln(F_{t,\tau_1^\bullet}) \dots \ln(F_{t,\tau_{m-2}^\bullet})]^\top$ is the vector with the $(m-2)$ imperfectly correlated log-futures, $\underline{X}_t^\circ \equiv [\ln(F_{t,\tau_{m-1}^\circ}) \ln(F_{t,\tau_m^\circ})]^\top$ denotes the vector with the two perfectly correlated log-futures prices, and matrices \underline{A}^\bullet , \underline{B}^\bullet , \underline{A}° , and \underline{B}° are defined as

$$(22) \quad \underline{A}^\bullet \equiv \begin{bmatrix} A(t, \tau_1^\bullet) \\ \vdots \\ A(t, \tau_{m-2}^\bullet) \end{bmatrix},$$

$$(23) \quad \underline{B}^\bullet \equiv \begin{bmatrix} B_1(\tau_1^\bullet) & B_2(\tau_1^\bullet) \\ \vdots & \vdots \\ B_1(\tau_{m-2}^\bullet) & B_2(\tau_{m-2}^\bullet) \end{bmatrix},$$

$$(24) \quad \underline{A}^\circ \equiv \begin{bmatrix} A(t, \tau_{m-1}^\circ) \\ A(t, \tau_m^\circ) \end{bmatrix}, \text{ and}$$

$$(25) \quad \underline{B}^\circ \equiv \begin{bmatrix} B_1(\tau_{m-1}^\circ) & B_2(\tau_{m-1}^\circ) \\ B_1(\tau_m^\circ) & B_2(\tau_m^\circ) \end{bmatrix}.$$

Direct estimation of the historical evolution equation (19) is not feasible because the vector of latent factors $\underline{Y}(t)$ is not observed.¹¹ However, the factors can be computed from equation (21) as $\underline{Y}(t) = (\underline{B}^\circ)^{-1} (\underline{X}_t^\circ - \underline{A}^\circ)$, provided the (2×2) matrix \underline{B}° is invertible. That is, the value of the state variables can be exactly filtered out at each sample date by inversion based on the two log-futures prices observed without error.

The estimation system (19)-(25) can be easily augmented to incorporate data on option premiums. As with the imperfectly correlated log-futures, option prices are also assumed to be observed with errors. Given the call and put formulas (16) and (17'), the following equations are added to the empirical model to make use of option data in the estimation

$$(26) \quad \ln(C_{t,\tau_0,\tau}^K) = -r \tau_0 + \ln[F_{t,\tau} N(d_{1,t,\tau}^K) - K N(d_{2,t,\tau}^K)] + e_t^{C_{\tau_0,\tau}^K},$$

$$(27) \quad \ln(P_{t,\tau_0,\tau}^K) = -r \tau_0 + \ln[K N(-d_{2,t,\tau}^K) - F_{t,\tau} N(-d_{1,t,\tau}^K)] + e_t^{P_{\tau_0,\tau}^K},$$

where $C_{t,\tau_0,\tau}^K$ ($P_{t,\tau_0,\tau}^K$) is the actual premium observed at date t for a call (put) with strike price K and maturity date $(t + \tau_0)$ on a futures expiring at time $(t + \tau)$, $d_{1,t,\tau}^K \equiv [\ln(F_{t,\tau})/K + 0.5 \text{Var}(\tau_0, \tau - \tau_0)]/[\text{Var}(\tau_0, \tau - \tau_0)]^{0.5}$, $d_{2,t,\tau}^K \equiv [\ln(F_{t,\tau})/K - 0.5 \text{Var}(\tau_0, \tau - \tau_0)]/[\text{Var}(\tau_0, \tau - \tau_0)]^{0.5}$, and $e_t^{C_{\tau_0,\tau}^K} \sim N(0, \sigma_c^2)$ and $e_t^{P_{\tau_0,\tau}^K} \sim N(0, \sigma_p^2)$ are regression residuals. The serially and cross-sectionally uncorrelated mean-zero disturbances $e_t^{C_{\tau_0,\tau}^K}$ and $e_t^{P_{\tau_0,\tau}^K}$ are added into the logarithms of the put and call option formulas to take into account non-simultaneity of observations, errors in the data, and other potential sources of errors.

Description of the Data

The futures data employed to estimate the models consist of monthly observations of futures prices for soybeans.¹² Soybean futures price data are obtained from the CME Group from January 2006 through December 2012, for a total of $n_{Dates} = 84$ observation dates. The futures prices used, denoted in cents per bushel, are settlement prices for the 15th calendar day of each month. If the 15th of the month falls on a weekend or a holiday, the nearest trading day's settlement price is used. Settlement prices observed with zero trading volume are discarded, because they are set by the CME Group administration for the purpose of calculating margins. In other words, such prices are not actual trading prices.

Since the longest maturity in the soybean futures sample is $\tau = 46$ months, the ideal data set would consist of a balanced panel of $84 \times 46 = 3,864$ observations. However, futures for some maturities are not traded. Soybean futures currently have only seven maturity months: January, March, May, July, August, September, and November. In addition, data with far-away maturities are often missing because they are not traded. For example, on September 15, 2006, only seven prices are observed for soybean futures, corresponding to the expiration dates of November 2006, January 2007, March 2007, May 2007, July 2007, November 2007, and November 2008. Letting the $(i, j)^{\text{th}}$ element of our data set be the price of the futures contract that expires j months after date i , this means that all elements of the 9th row of our soybean futures data set are missing, except for the 2nd, 4th, 6th, 8th, 10th, 14th, and 26th column. As a result, there are only 963 total observations available on soybean futures.

Soybean futures options are observed over the same period as futures (i.e., from January 2006 until December 2012), for 84 observation dates. The data on futures options was constructed in a similar manner as the futures data, i.e., option premiums are mid-month settlement prices for contracts with positive trading volume. Futures option contracts expire around three-quarters of a month prior to the expiration of the underlying futures contracts, so that $\Delta T = \tau - \tau_o = 0.75$. At any particular date, a variety of option contracts with different underlying futures contracts and/or strike prices are traded. The longest maturity date for the option's underlying futures contract is 25 months in our sample.

Our model applies to European options, while the options data correspond to American options. American options are more valuable than their European counterparts, because of the existence of early exercise opportunity for the former. However, after calculating values for both American- and European-type options, Plato (1985) concluded that the difference between them for near-the-money options is negligible. Hence, we estimate the model using only premiums for options with strike prices immediately above and immediately below the corresponding observed futures price. We also discarded option prices whenever the strike price was more than \$3.00 per bushel away from the underlying futures price. Given the underlying futures data available, the

ideal premium data sets for call and put options would consist of 1,926 ($= 2 \times 963$) observations each. However, only 883 (810) call (put) option prices are observed during the sample period because of limited trading.

Estimation Method

Bayesian MCMC methods are employed to estimate the parameters associated with the advocated model, because they are better suited for the present purposes than more traditional methods (Johannes and Polson, 2010). The MCMC environment allows us to easily incorporate seasonality and the uneven distribution of maturity dates throughout the year that characterizes the futures contracts under analysis. More importantly, with the advocated method it is straightforward to bridge the gap between the theory (based on constant-time-to-maturity futures prices) and the observed data (consisting of constant-maturity-date futures prices).

Succinctly, the proposed method consists of constructing a Markov chain with the desired probability distribution as its equilibrium distribution, and using an MCMC algorithm to draw samples from it. Upon convergence, the samples drawn are taken to be samples from the equilibrium distribution. The MCMC method allows us to draw the samples without having to know the exact form of the equilibrium probability distribution at any point. Intuitively, the method works because the Markov chain (provided it is correctly set up) spends in each location an amount of time proportional to the height of the equilibrium distribution.

The MCMC method described next is designed to estimate the proposed model using both futures and options data. Schwartz's model and/or estimation using futures data only can be easily retrieved by imposing the corresponding parameter restrictions into the procedure. The full set of model parameters to be estimated is $\{\tilde{\kappa}_{20}, \tilde{\kappa}_{21}, \tilde{\kappa}_{22}, \underline{\lambda}_{10}, \underline{\lambda}_{20}, \underline{\lambda}, \underline{\Sigma}\}$, plus the model's goodness-of-fit parameters. The latter consist of $\{\sigma_F, \rho_F\}$ for the estimation using futures data only, and $\{\sigma_F, \rho_F, \sigma_C, \sigma_P\}$ for the estimation relying on both futures and options data. The risk-free interest rate r (which enters the closed-form option pricing formulas (16) and (17)) is a constant in our two-factor model, and is fixed at $r = 0.0025$ per month (or 3% per year).¹³ We use

non-informative priors for the parameters in the set $\{\tilde{\kappa}_{20}, \tilde{\kappa}_{21}, \tilde{\kappa}_{22}, \underline{\lambda}_{10}, \underline{\lambda}_{20}, \underline{\lambda}, \underline{\Sigma}, \rho_F\}$. The proposed priors for the lack-of-fit parameters $\{\sigma_F, \sigma_C, \sigma_P\}$ are scaled inverse-chi-squared distributions

$$(28) \quad \sigma_i^2 \sim \text{Inv-}\chi^2(\bar{v}_i, \bar{\sigma}_i^2),$$

for $i \in \{F, C, P\}$ and prior parameters \bar{v}_i and $\bar{\sigma}_i^2$.

The MCMC iteration steps are outlined in the supplemental appendix online. In the present study, four chains of the Bayesian MCMC procedure were run for each model specification. The chains were tested for convergence using the Gelman and Rubin (1992) tests. By comparing the estimates for the common parameter set $\{\tilde{\kappa}_{20}, \tilde{\kappa}_{21}, \tilde{\kappa}_{22}, \underline{\lambda}_{10}, \underline{\lambda}_{20}, \underline{\lambda}, \underline{\Sigma}, \sigma_F, \rho_F\}$ obtained using both futures and options data with the one computed by means of futures data only, we can assess whether option prices are driven by the same set of underlying parameters as futures.

Results and Discussion

To explore the model space, we estimated four different variations of the model. Model 1 is the Schwartz model with the addition of seasonality. Model 2 is the advocated model. Model 3 is the Schwartz model with the risk premium parameters ($\underline{\lambda}$) set equal to zero. Model 4 is the advocated model with the risk premium parameters set equal to zero. All four models are first estimated using only futures data. Then all four models are estimated again, incorporating options data as well. To judge convergence, Gelman and Rubin's R statistic is utilized. Chains were estimated until all parameters had an R statistic below 1.3. To quickly summarize the results from the various estimations, Table 1 displays the median values for the parameters and the log-likelihood for each model when only futures data were used in the estimation. Table 2 displays the same for when futures and options data were used in the estimation. Full results for all of the model estimations are provided in the supplemental appendix online.

To more formally compare the models, the log-likelihood based on (29) and (30)

$$(29) \quad \begin{bmatrix} Y_1^{(j)}(t+1) - Y_1^{(j)}(t) \\ Y_2^{(j)}(t+1) - Y_2^{(j)}(t) \end{bmatrix} - \left(\begin{bmatrix} \lambda_{10}^{(j)}(t) \\ \tilde{\kappa}_{20}^{(j)}(t) + \lambda_{20}^{(j)}(t) \end{bmatrix} - \begin{bmatrix} \lambda_{11}^{(j)} & -1 + \lambda_{12}^{(j)} \\ \tilde{\kappa}_{21}^{(j)} + \lambda_{21}^{(j)} & \tilde{\kappa}_{22}^{(j)} + \lambda_{11}^{(j)} \end{bmatrix} \begin{bmatrix} Y_1^{(j)}(t) \\ Y_2^{(j)}(t) \end{bmatrix} \right) \sim N(\mathbf{0}_{(2 \times 1)}, \underline{\Sigma}^{(j)}) \text{ and}$$

$$(30) \quad \underline{X}_t^{\bullet(j)} - \underline{A}^{\bullet(j)} - \underline{B}^{\bullet(j)} \underline{Y}^{(j)}(t) \sim N(\mathbf{0}_{((m-2) \times 1)}, (\sigma_F^{(j)})^2 \underline{Q}^{(j)})$$

was tracked during the estimation. The log-likelihood values favor the versions of the advocated model (Models 2 and 4) over the Schwartz versions. For both the Schwartz and advocated models, the removal of the risk premiums had minimal impact on the log-likelihood values. Thus, Model 4 is the preferred specification as it is more parsimonious.

For all four specifications, the parameters related to the risk premiums, with one exception, are indistinguishable from zero. And for all four specifications, two of the four seasonality parameters are significantly different than zero. For Models 2 and 4, the risk-neutral intercept and the speed of mean reversion in both factor 1 (the log-spot price) and factor 2 (the expected risk-neutral instantaneous change in the log-spot price) are significantly different than zero, along with the correlation parameters. As both $\tilde{\kappa}_{21}$ and $\tilde{\kappa}_{22}$ are greater than zero, the system is reverting back to the long run means for both factors, which are $\tilde{\kappa}_{20} / \tilde{\kappa}_{21}$ and zero for factors 1 and 2 respectively. In comparing the parameter estimates between the Schwartz and the advocated model, there is considerable overlap between the estimates. However, there are two significant differences. The risk-neutral intercept ($\tilde{\kappa}_{20}$) changes signs from negative to positive, and the speed of mean reversion for factor 2 ($\tilde{\kappa}_{22}$) increases by almost 50%.

Table 2 shows that the relative fit across structural models remained unchanged as options were added to the underlying data set. Model 4 is the preferred specification. For the

parameter estimates, the posterior probability bands overlap between the “Futures Only” estimation and the “Futures and Options” estimation for all parameters, with the exception of σ_2 , which is consistently smaller once options data is included in the estimation. Thus, option prices are driven by the same set of underlying parameters as futures.

Another approach to explore the model specifications is to compare out-of-sample projections for futures and options prices. As with the data for estimation, the futures data employed for the out-of-sample tests consist of monthly observations of futures settlement prices for soybeans on the 15th calendar day of each month. If the 15th of the month falls on a weekend or a holiday, the nearest trading day's settlement price is used. Two time periods were utilized for the out-of-sample examination, 2004-2005 and 2013. The 2004-2005 time period allows for an examination of the model fit in a period before agricultural commodity prices rose significantly. The 2013 time period allows one to examine the model fit after most commodity markets had receded from record highs.

The converged chains from each model specification were used to simulate futures and options prices for the out-of-sample time periods. The mean value for simulated futures and options prices from each model specification was then compared to the actual value from the time period. As Table 3 shows, the advocated model structure outperformed the Schwartz model structure for the out-of-sample futures simulation in both time periods and with both data sets (“Futures Only” vs. “Futures and Options”). Also, the simplification of the models by setting the risk premiums equal to zero did not diminish the out-of-sample performance for either time period or with either data set.¹⁴

As the option premiums are dependent on the underlying futures prices, the ability of the models to project the futures prices is crucial. Based on the above results, the advocated models perform those projections better than currently used model structures. For the out-of-sample options simulations, the focus was narrowed to Models 2 and 4, based on the model estimation and out-of-sample futures simulation results. As with the futures simulations, data from two time periods (2004-2005 and 2013) were utilized to examine the ability of each model to project the

actual premiums for options traded during the time periods. For this analysis, the actual premiums for calls and puts with strike prices immediately above and immediately below the corresponding observed futures prices were gathered. Following the procedures for the futures data, the options data employed for the out-of-sample tests consist of monthly observations of option premium settlement prices on the 15th calendar day of each month. If the 15th of the month falls on a weekend or a holiday, the nearest trading day's settlement price is used.

To narrow the focus of this test to concentrate on the ability of the models to project option premiums, the futures prices for the simulations were set equal to the actual futures data. This removes the impact of the imprecision of the futures forecast (already measured in the futures out-of-sample test) from the options out-of-sample test and limits the discrepancies to the estimates of the implied volatilities. It is also here where one would expect to see the impact of the addition of the options data in the estimation process. In projecting futures prices, both the “Futures Only” and “Futures and Options” estimations provided, as expected, very similar futures price projections. If anything, the addition of the options data added a little noise to futures price projections. However, the addition of the options data and accompanying equations to the model should have provided information for more accurate projection of option premiums. And as Table 4 shows, that is the case. The option premium projections are markedly better from the “Futures and Options” estimations. Also, whereas the restrictions of zero risk premiums have dramatically improved the futures price projections and the option premium projections in the “Futures Only” estimation, that same set of restrictions had very little impact for the “Futures and Options” option premium projections.

Figure 4 truly captures the issues raised by the present study. The thick black line shows the historical futures price variance for November soybean futures that have actively traded at least 25 months before maturity.¹⁵ The thin black line shows the implied variance from the traditional Black model for these same futures. The dashed line displays the implied variance from the Schwartz model (Model 3 from the current analysis). And the gray line displays the implied variance from the advocated model (Model 4). As options premiums depend on futures

price variance and options trading depends on appropriately priced option premiums, the differences among the lines highlight the discrepancies on the variance estimates and the resulting impact on option trading. For longer-term options, if premiums are being determined under Black or Schwartz style models, but mean reversion in spot prices holds; then variance estimates and option premiums are too high and trading will be severely limited. As the figure shows, while the variance estimates are similar for the three models for maturities under one year, the discrepancies increase as the time to maturity increases. For a three-year time to maturity, the variance gap between Black and Model 4 is 13.6%. For a five year time to maturity, the gap increases to 24.5%. And the larger the variance gap is, the larger the discrepancy in option premiums becomes.

Conclusions

We generalize Schwartz's two-factor model by allowing mean reversion in spot prices, which is a key feature of agricultural commodity markets. Agricultural commodity markets also exhibit seasonal patterns. We introduce seasonality into our model by adding periodic functions to the parameters associated with commodity prices. Closed-form futures and option pricing formulas are derived. We show that Schwartz's model is a special case of our model.

Soybean futures price data from the CME Group are employed to estimate the models by means of a Bayesian MCMC algorithm. Estimates for the Schwartz model are obtained by imposing the corresponding restrictions on our model. Our option pricing model incorporates spot price mean reversion to long-run production costs in the soybean market, and as a result, it provides option premiums that are lower than in the Black or Schwartz models. This is the key finding of our work and it suggests that the inappropriate use of these other models will result in overpricing of long-term options. And, in fact, these markets suffer from a lack of liquidity for long-term options.

References

- Allen, M., C. Ma, and D. Pace. 1994. "Over-Reactions in U.S. Agricultural Commodity Prices." *Journal of Agricultural Economics* 45(2):240-251.
- Bessembinder, H., C. Jay, P. Seguin, and M. Smoller. 1995. "Mean Reversion in Equilibrium Asset Prices: Evidence from the Futures Term Structure." *Journal of Finance* 50(1):361-375.
- Black, F. 1976. "The Pricing of Commodity Contracts." *Journal of Financial Economics* 3(2):167-179.
- Black, F. and M. Scholes. 1973. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy* 81(3):637-654.
- Bobenrieth, E. S. A., J. R. A. Bobenrieth, and B. D. Wright. 2014. "Bubble Troubles? Rational Storage, Mean Reversion, and Runs in Commodity Prices." In Chavas, J.-P., D. Hummels, and B. D. Wright, eds., *The Economics of Food Price Volatility*, Chapter 5, pp. 193-208. Chicago, IL: University of Chicago Press.
- Casassus, J. and P. Collin-Dufresne. 2005. "Stochastic Convenience Yield Implied from Commodity Futures and Interest Rates." *Journal of Finance* 60(5):2283-2331.
- Chen, R. and L. Scott. 1993. "Maximum Likelihood Estimation for a Multifactor Equilibrium Model of the Term Structure of Interest Rates." *Journal of Fixed Income* 3(3):14-31.
- Collin-Dufresne, P., R. S. Goldstein, and C. S. Jones. 2008. "Identification of Maximal Affine Term Structure Models." *Journal of Finance* 63(2):743-795.
- Dai, Q. and K. Singleton. 2000. "Specification Analysis of Affine Term Structure Models." *Journal of Finance* 55(5):1943-1978.
- Deaton, A., and G. Laroque. 1992. "On the Behaviour of Commodity Prices." *Review of Economic Studies* 59(1):1-23.
- Duffie, D., J. Pan, and K. Singleton. 2000. "Transform Analysis and Asset Pricing for Affine Jump-Diffusions." *Econometrica* 68(6):1343-1376.

- Egelkraut, T. M., P. Garcia, and B. J. Sherrick. 2007. "The Term Structure of Implied Forward Volatility: Recovery and Informational Content in the Corn Options Market." *American Journal of Agricultural Economics* 89(1):1-11.
- Gelman, A. and D. B. Rubin. 1992. "Inference from Iterative Simulation Using Multiple Sequences." *Statistical Science* 7(4):457-472.
- Hilliard, J., and J. Reis. 1998. "Valuation of Commodity Futures and Options under Stochastic Convenience Yields, Interest Rates, and Jump Diffusions in the Spot." *Journal of Financial and Quantitative Analysis* 33(1):61-86.
- Irwin, S. H., C. R. Zulauf, and T. E. Jackson. 1996. "Monte Carlo Analysis of Mean Reversion in Commodity Futures Prices." *American Journal of Agricultural Economics* 78(2):387-399.
- Jin, N., S. H. Lence, C. Hart, and D. Hayes. 2012. "The Term Structure of Long-Run Commodity Futures." *American Journal of Agricultural Economics* 94(3):718-735.
- Johannes, M., and N. Polson. 2010. "MCMC Methods for Continuous-Time Financial Econometrics." In Ait-Sahalia, Y., and L. P. Hansen, Eds., *Handbook of Financial Econometrics, Vol. 2-Applications*, Chapter 13, pp. 1-72. Oxford, UK: Elsevier Science Ltd.
- Kaldor, N. 1939. "Speculation and Economic Stability." *Review of Economic Studies* 7(1):1-27.
- Kim, H. S., B. W. Brorsen, and K. B. Anderson. 2010. "Profit Margin Hedging." *American Journal of Agricultural Economics* 92(3):638-653.
- Koekebakker, S., and G. Lien. 2004. "Volatility and Price Jumps in Agricultural Futures Prices - Evidence from Wheat Options." *American Journal of Agricultural Economics* 86(4):1018-1031.
- Lence, S. H., H. G. Ott, and C. E. Hart. 2013. "Long-Term Futures Curves and Seasonal Structures of Wheat in the European Union and the United States." *Journal of Futures Markets* 33(12):1118-1142.

- Lo, A. and J. Wang. 1995. "Implementing Option Pricing Models when Asset Returns are Predictable." *Journal of Finance* 50(1):87-129.
- Miltersen, K. and E. Schwartz. 1998. "Pricing of Options on Commodity Futures with Stochastic Term Structure of Convenience Yields and Interest Rates." *Journal of Financial and Quantitative Analysis* 33(1):33-59.
- Pennacchi, G. 2008. *Theory of Asset Pricing*. Boston, MA: Pearson Education.
- Peterson, R., C. Ma, and R. Ritchey. 1992. "Dependence in Commodity Prices." *Journal of Futures Markets* 12(4):429-446.
- Plato, G. 1985. "Valuing American Options on Commodity Futures Contracts." *Agricultural Economics Research* 37(2):1-14.
- Richter, M. and C. Sørensen. 2002. "Stochastic Volatility and Seasonality in Commodity Futures and Options: The Case of Soybeans." Unpublished working paper, Department of Finance, Copenhagen Business School.
- Sanders, D. R., and S. H. Irwin. 2012. "A Reappraisal of Investing in Commodity Futures Markets." *Applied Economic Perspectives and Policy* 34(3):515-530.
- Schwartz, E. 1997. "The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging." *Journal of Finance* 52(3):923-973.
- Sørensen, C. 2002. "Modeling Seasonality in Agricultural Commodity Futures." *Journal of Futures Markets* 22(5):393-426.
- Trolle, A. and E. Schwartz. 2009. "Unspanned Stochastic Volatility and the Pricing of Commodity Derivatives." *Review of Financial Studies* 22(11):4423-4461.
- Walburger, A. and K. Foster. 1995. "Mean Reversion as a Test for Inefficient Price Discovery in U.S. Regional Cattle Markets." Published abstract. *American Journal of Agricultural Economics* 77(5):1370.
- Wang, D. and W. G. Tomek. 2007. "Commodity Prices and Unit Root Tests." *American Journal of Agricultural Economics* 89(4):873-889.

Williams, J. C. and B. D. Wright. 1991. *Storage and Commodity Markets*. Cambridge, MA:
Cambridge University Press.

Table 1. Medians from the Posterior Probability Band Estimated Using Futures Data Only

	Model 1	Model 2	Model 3	Model 4
Risk-Neutral Parameters				
$\tilde{\kappa}_{20}$	-0.0003*	0.0040	-0.0003*	0.0041
$\tilde{\kappa}_{201,sin}$	-0.0003	-0.0001	-0.0003	-0.0001
$\tilde{\kappa}_{201,cos}$	0.0030*	0.0031*	0.0030*	0.0031*
$\tilde{\kappa}_{202,sin}$	-0.0048*	-0.0049*	-0.0045*	-0.0048*
$\tilde{\kappa}_{202,cos}$	-0.0004	-0.0003	-0.0002	-0.0002
$\tilde{\kappa}_{21}$		0.0007*		0.0007*
$\tilde{\kappa}_{22}$	0.0852*	0.1255*	0.0846*	0.1252*
Risk-Premiums				
λ_{10}	0.0146	0.2103		
$\lambda_{101,sin}$	0.0117	0.0130		
$\lambda_{101,cos}$	0.0069	0.0065		
$\lambda_{102,sin}$	-0.0169	-0.0170		
$\lambda_{102,cos}$	0.0111	0.0131		
λ_{20}	-0.0002	0.0290		
$\lambda_{201,sin}$	-0.0011	-0.0005		
$\lambda_{201,cos}$	-0.0003	-0.0001		
$\lambda_{202,sin}$	0.0014	0.0014		
$\lambda_{202,cos}$	-0.0023	-0.0017		
λ_{11}		0.0282		
λ_{12}		0.0151		
λ_{21}		0.0041		
λ_{22}		0.1144		
Covariance Matrix				
σ_1	0.0915*	0.1009*	0.0905*	0.0959*
σ_2	0.0202*	0.0169*	0.0101*	0.0144*
ρ_{12}	-0.7215*	-0.7066	-0.7285*	-0.7187
Residual Errors				
σ_F	0.0202*	0.0193*	0.0202*	0.0193*
ρ_F	0.1314*	0.1556*	0.1320*	0.1564*
Log-Likelihood	3132	3161	3136	3171

* indicates that the 95% posterior probability band does not include zero.

Note: Model 1 is the Schwartz model with the addition of seasonality. Model 2 is the advocated model. Model 3 is the Schwartz model with the risk premium parameters ($\underline{\lambda}$) set equal to zero. Model 4 is the advocated model with the risk premium parameters set equal to zero. The simulations are based on 963 futures observations.

Table 2. Medians from the Posterior Probability Band Estimated Using Futures and Options Data

	Model 1	Model 2	Model 3	Model 4
Risk-Neutral Parameters				
$\tilde{\kappa}_{20}$	-0.0004*	0.0042*	-0.0004*	0.0042*
$\tilde{\kappa}_{201,sin}$	-0.0003	-0.0001	-0.0003	-0.0001
$\tilde{\kappa}_{201,cos}$	0.0030*	0.0030*	0.0030*	0.0030*
$\tilde{\kappa}_{202,sin}$	-0.0047*	-0.0048*	-0.0041*	-0.0045*
$\tilde{\kappa}_{202,cos}$	-0.0005	-0.0003	0.0000	0.0000
$\tilde{\kappa}_{21}$		0.0007*		0.0007*
$\tilde{\kappa}_{22}$	0.0821*	0.1209*	0.0824*	0.1211*
Risk-Premiums				
λ_{10}	0.0153	0.1619		
$\lambda_{101,sin}$	0.0116	0.0109		
$\lambda_{101,cos}$	0.0070	0.0059		
$\lambda_{102,sin}$	-0.0167	-0.0166		
$\lambda_{102,cos}$	0.0111	0.0110		
λ_{20}	-0.0002	0.0251		
$\lambda_{201,sin}$	-0.0011	-0.0005		
$\lambda_{201,cos}$	-0.0003	-0.0001		
$\lambda_{202,sin}$	0.0014	0.0013		
$\lambda_{202,cos}$	-0.0022*	-0.0017		
λ_{11}		0.0211		
λ_{12}		-0.0229		
λ_{21}		0.0037		
λ_{22}		0.1027		
Covariance Matrix				
σ_1	0.0824*	0.0841*	0.0840*	0.0841*
σ_2	0.0064*	0.0081*	0.0064*	0.0081*
ρ_{12}	-0.4785*	-0.4401*	-0.4793*	-0.4399*
Residual Errors				
σ_F	0.0204*	0.0195*	0.0204*	0.0195*
ρ_F	0.1360*	0.1620*	0.1356*	0.1588*
σ_C	0.1275*	0.1275*	0.1275*	0.1275*
σ_P	0.1360*	0.1360*	0.1360*	0.1360*
Log-Likelihood	-1015	-981	-1015	-979

* indicates that the 95% posterior probability band does not include zero.

Note: Model 1 is the Schwartz model with the addition of seasonality. Model 2 is the advocated model. Model 3 is the Schwartz model with the risk premium parameters ($\underline{\lambda}$) set equal to zero. Model 4 is the advocated model with the risk premium parameters set equal to zero. The simulations are based on 963 futures, 883 call option, and 810 put option observations.

Table 3. Out-of-Sample Root-Mean-Squared-Error (RMSE) Results for Futures

Model	RMSE for Model Estimated Using only Futures Data	RMSE for Model Estimated Using Futures and Options data
Out-of-Sample Period: 2004-2005		
1	25.675	25.651
2	25.206	25.291
3	25.733	25.776
4	25.252	25.360
Out-of-Sample Period: 2013		
1	36.705	36.692
2	33.395	33.465
3	36.811	36.855
4	33.378	33.500

Note: Model 1 is the Schwartz model with the addition of seasonality. Model 2 is the advocated model. Model 3 is the Schwartz model with the risk premium parameters (λ) set equal to zero. Model 4 is the advocated model with the risk premium parameters set equal to zero. The out-of-sample data set contain 296 futures observations.

Table 4. Out-of-Sample Root-Mean-Squared-Error (RMSE) Results for Options

Model	RMSE for Model Estimated Using only Futures Data	RMSE for Model Estimated Using Futures and Options data
Out-of-Sample Period: 2004-2005		
2	23.702	7.189
4	9.881	7.188
Out-of-Sample Period: 2013		
2	66.204	23.349
4	37.256	23.344

Note: Model 2 is the advocated model. Model 4 is the advocated model with the risk premium parameters set equal to zero. The out-of-sample data set contain 427 call option and 371 put option observations.

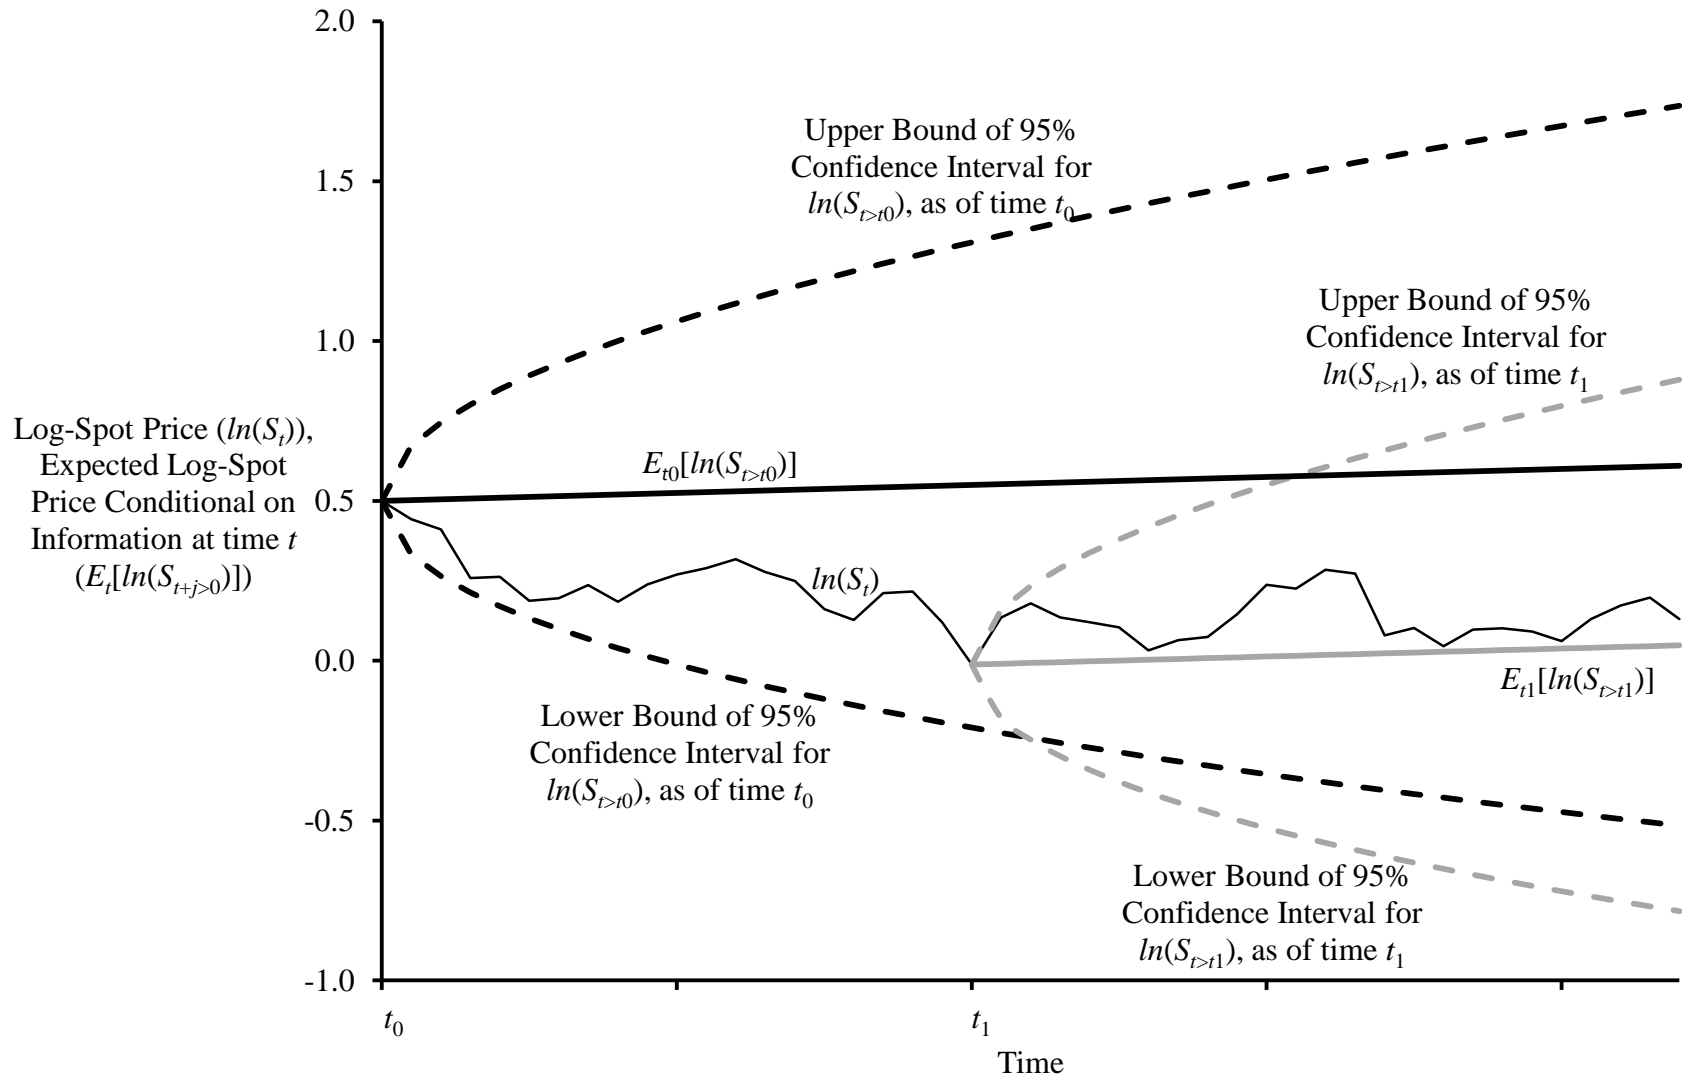


Figure 1. Hypothetical behavior of the logarithm of the spot price ($\ln(S_t)$), conditional expectation of the logarithm of the spot price ($E_t[\ln(S_{T>t})]$), and 95% confidence interval under Black's model

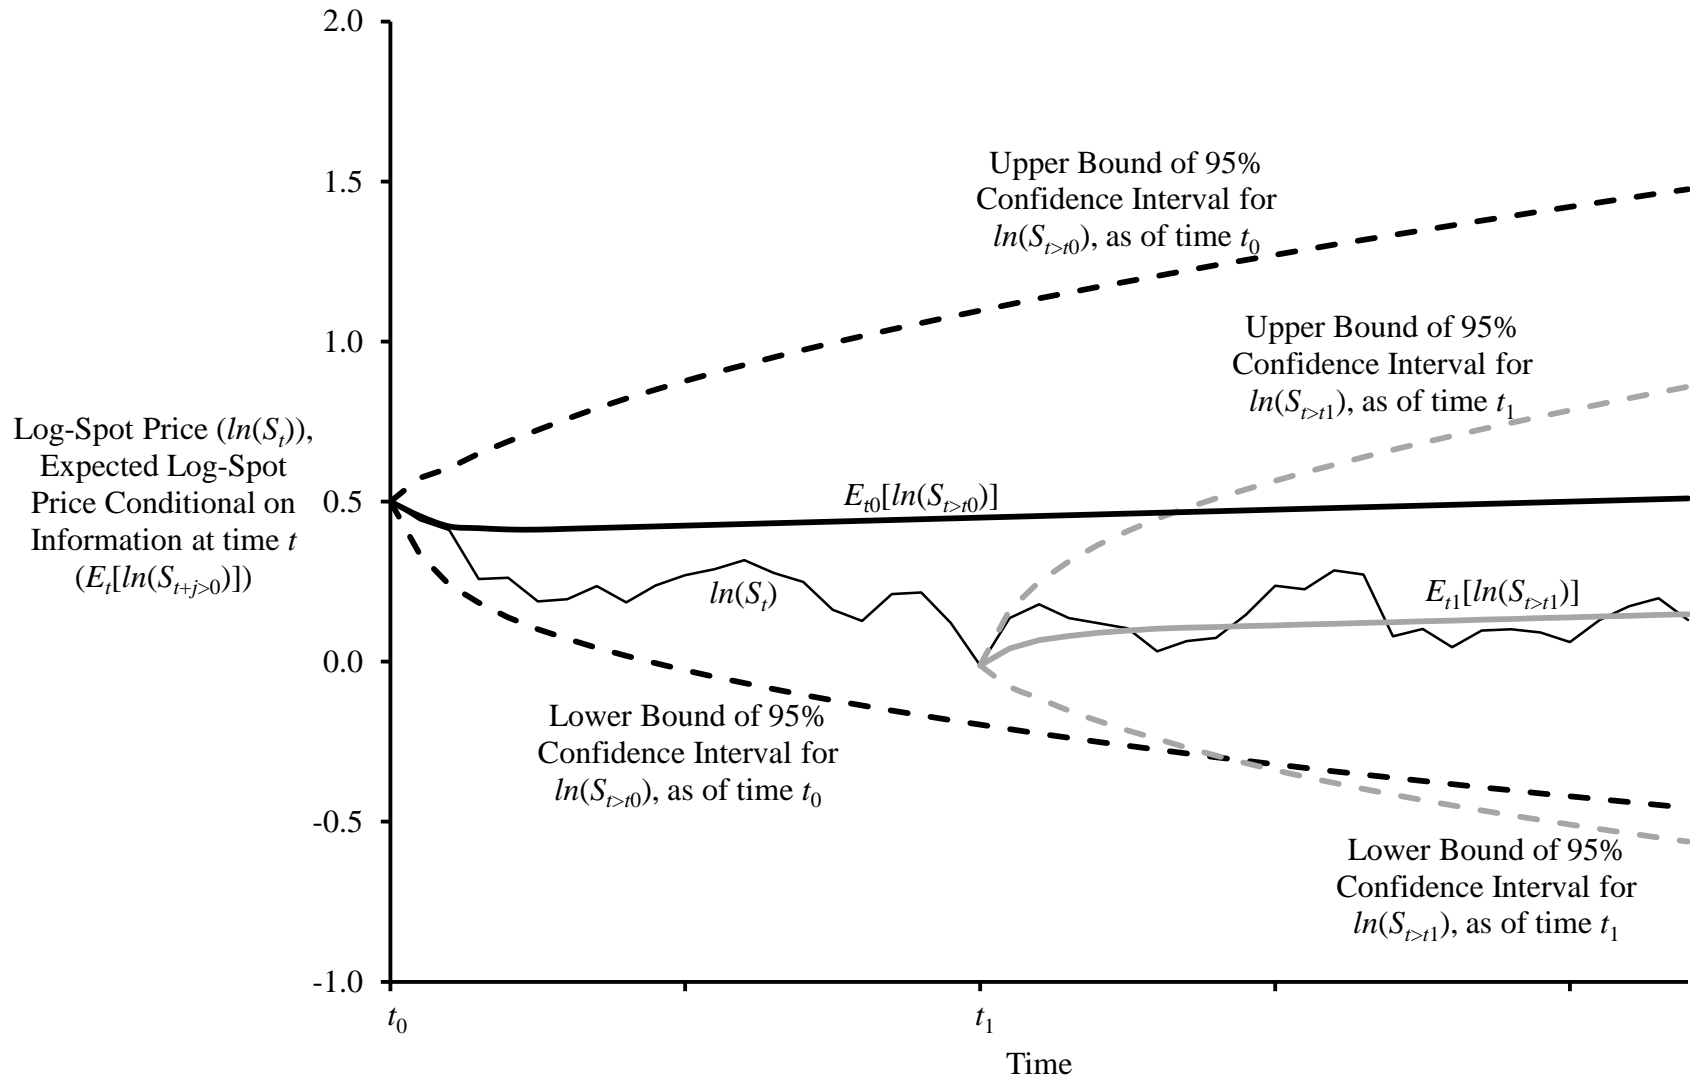


Figure 2. Hypothetical behavior of the logarithm of the spot price ($\ln(S_t)$), conditional expectation of the logarithm of the spot price ($E_t[\ln(S_{T>t})]$), and 95% confidence interval under Schwartz's two-factor model

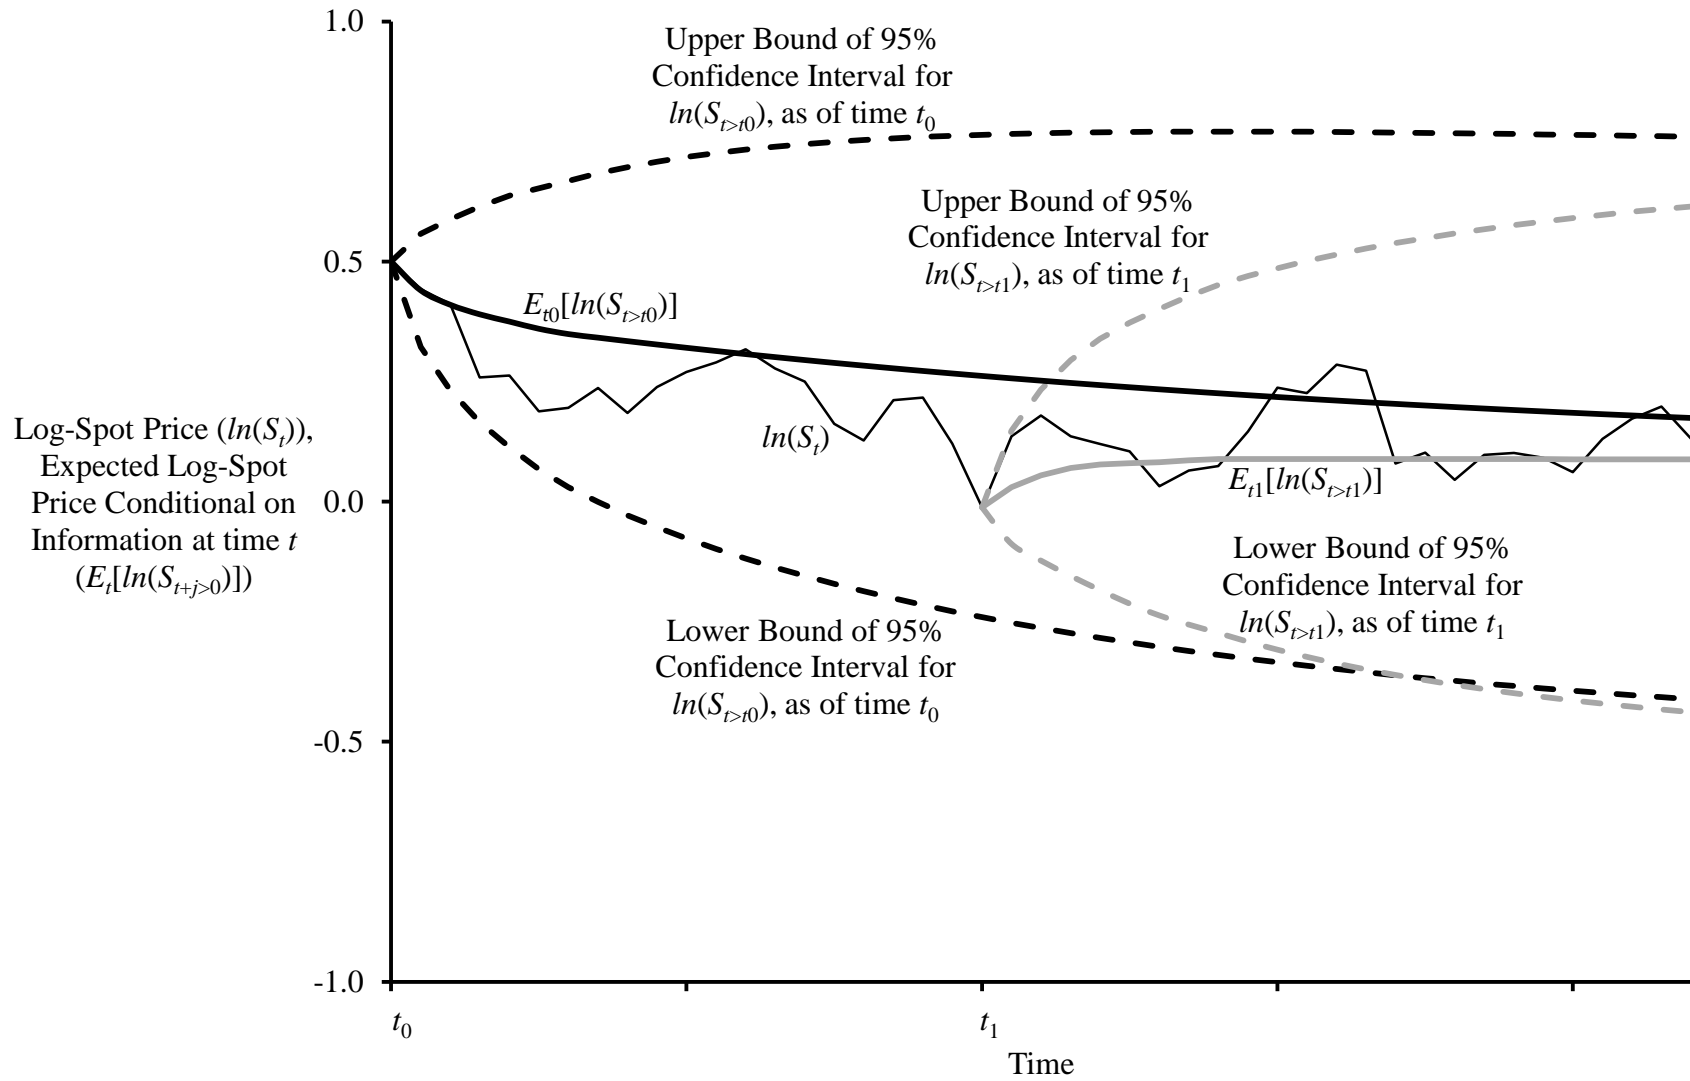


Figure 3. Hypothetical behavior of the logarithm of the spot price ($\ln(S_t)$), conditional expectation of the logarithm of the spot price ($E_t[\ln(S_{T>t})]$), and 95% confidence interval under advocated model

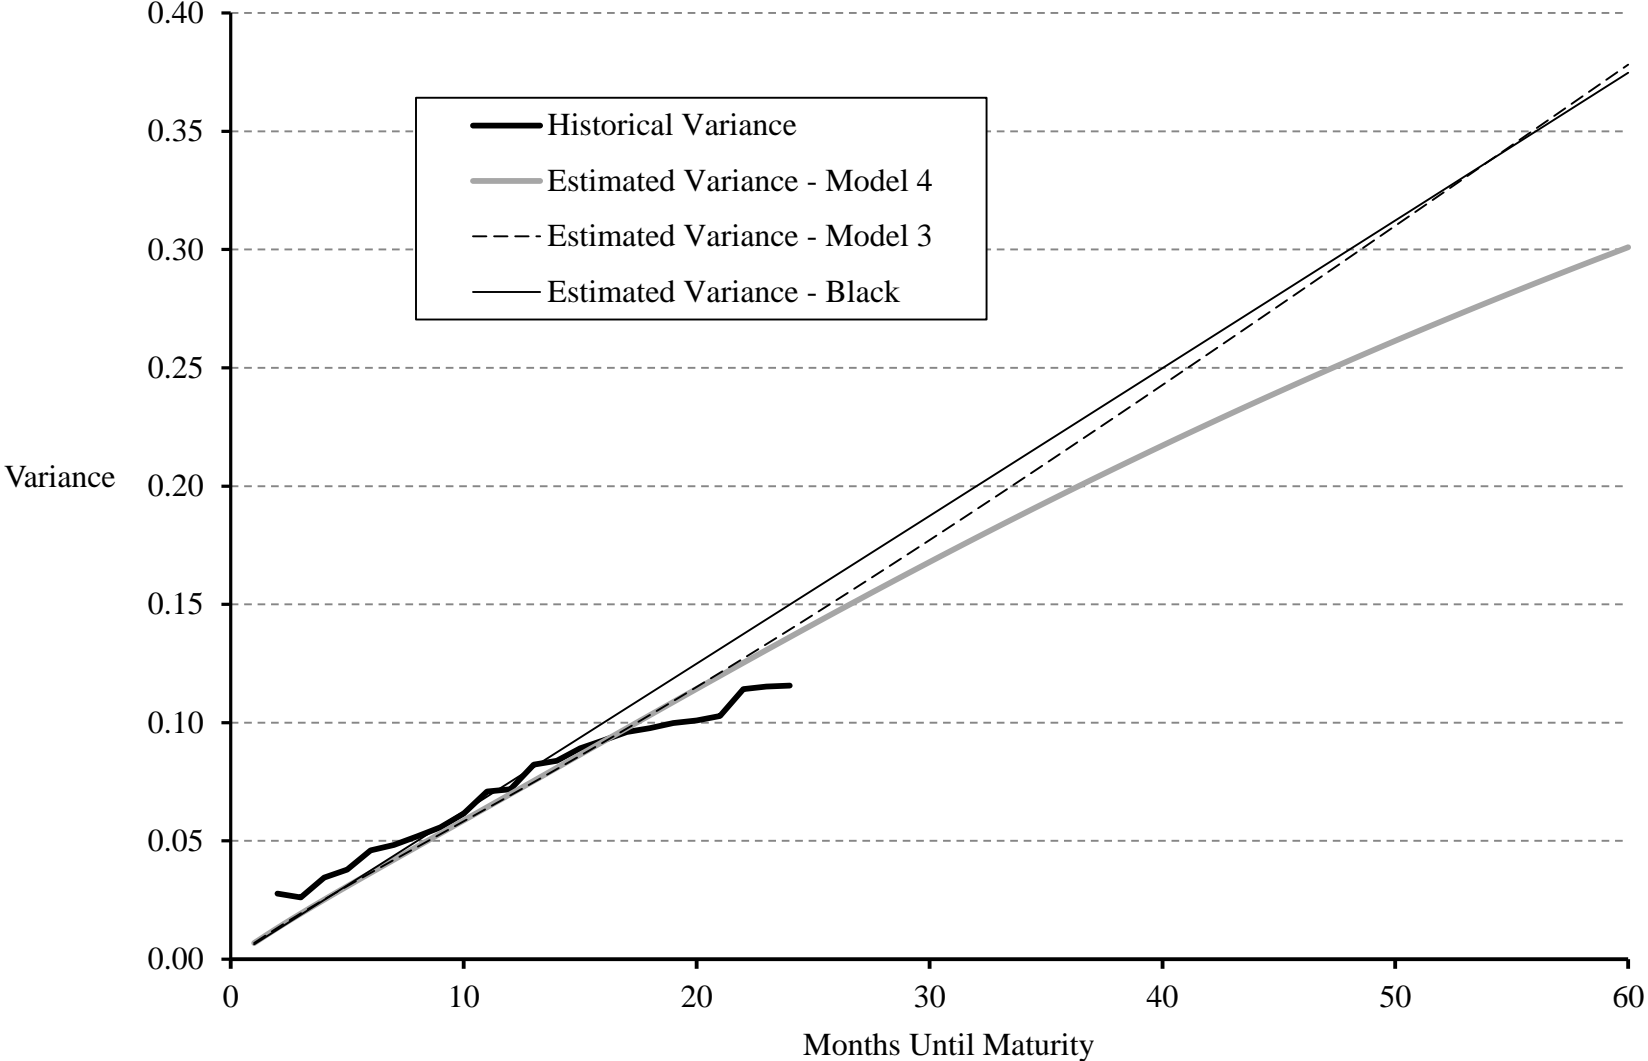


Figure 4. Historical and estimated futures price variance for soybeans

¹The present article focuses on mean reversion and seasonality in spot prices. It is important to point out, however, that other studies (e.g., Egelkraut, Garcia, and Sherrick (2007), Koekebakker and Lien (2004)) have uncovered mean reversion and seasonality in volatility, which also have important implications for option pricing.

²This motivation was provided to us by an anonymous reviewer.

³According to Lo and Wang (p. 89), "...if the drift depends only on exogenous time-varying economic factors, then an increase in predictability unambiguously decreases option values. But if the drift also depends upon lagged prices, then an increase in predictability can either increase or decrease option values, depending on the particular specification of the drift." The specification used in the present study belongs the latter class, and it is such that predictability reduces option values.

⁴It is important to note that the existence and behavior of convenience yield in storable commodity markets has been the focus of studies addressing the theory of rational commodity market storage, including the work on dynamic rational expectations storage models (e.g., Williams and Wright (1991), and Deaton and Laroque (1992)).

⁵Lower premiums for options on long-term futures follow from the fact that mean reversion in spot prices does reduce the volatility of long-term futures, even if futures themselves are not mean reverting. In fact, the absence of arbitrage opportunities rules out mean reversion in futures prices, because in such instances they should be martingales under the risk-neutral measure. Previous research has not found evidence of mean reversion patterns in commodity futures (Irwin, Zulauf, and Jackson (1996) and Kim, Brorsen, and Anderson (2010)). Therefore, the theoretical model used here assumes that futures prices are not mean reverting.

⁶The set comprises 12 parameters, which is the maximum number of parameters that can be identified in a bivariate Gaussian system.

⁷The model recently proposed by Jin et al. (2012) is also isomorphic to the present one, under the restrictions $\lambda_{11} = \lambda_{12} = \lambda_{21} = \lambda_{22} = 0$ (or, alternatively, $\kappa_{11} = 0$, $\kappa_{12} = -1$, $\kappa_{21} = \tilde{\kappa}_{21}$, and $\kappa_{22} = \tilde{\kappa}_{22}$). That is, Jin et al. allow for mean reversion in spot prices, but their specification is slightly more restrictive than the one used here.

⁸Recall that if parameters $E[\ln(x)]$ and $Var[\ln(x)]$ denote respectively the mean and variance of $\ln(x)$, where x is an arbitrary random variable distributed as log-normal, then $E(x) = \exp\{E[\ln(x)] + 0.5 Var[\ln(x)]\}$. Hence, if $E[\ln(x)] = \ln(F_x) - 0.5 Var[\ln(x)]$ for some arbitrary real number F_x , it follows that $E(x) = F_x$. This is why $Mean[\underline{Y}(t), t, \tau_O, \Delta T] = F(t, \tau_O + \Delta T)$.

⁹The normal distribution yields tractable solutions and facilitates the empirical analysis.

However, normality is only approximately satisfied empirically, due to non-zero skewness and excess kurtosis. Allowing for distributions that better match the empirical data is an important area for future research.

¹⁰Black's formula, which is derived under the assumption of a geometric Brownian motion, can be obtained from expression (16) as the limiting when $\tilde{\kappa}_{20} = \tilde{\kappa}_{21} = \tilde{\kappa}_{22} = \sigma_2 = 0$ (i.e., Y_2 is a constant).

¹¹Although the first factor (i.e., the log-spot price) may be observable, the second factor (i.e., the instantaneous risk-neutral drift of the log-spot price) cannot be observed.

¹²The use of monthly data implies that the time interval used for the discretization is set equal to one month in the empirical model (19), and that $\delta = 12$ in expressions (11)-(12) (because the periodic interval is one year).

¹³The model can be extended by explicitly modeling the stochastic interest rate. However, the pricing error for commodity futures that arises from ignoring the stochastic nature of interest rate is negligible (see, e.g., the discussion in Schwartz (1997) and Trolle and Schwartz (2009)).

¹⁴According to an anonymous reviewer, the accuracy of models that impose zero risk premiums is another confirmation that risk premiums in agricultural futures markets are negligible (see, e.g., results in Sanders and Irwin (2012) on expected returns to investments in commodity futures).

¹⁵More specifically, the data underlying Figure 4 consist of the monthly futures prices for the contracts maturing in November 2004, 2007, 2008, 2009, 2010, 2011, 2012, and 2013. For each of those contracts, we calculated the realized monthly rates of return as the first differences in the log-futures prices, and used them to estimate the variance for each horizon (e.g., the 24-months-until-maturity variance for the November 2004 contract was calculated as 24 times the monthly variance computed from the realized monthly rates of return from November 2002 until October 2004). The average of the x-months-until-maturity variances across the aforementioned contracts is the value depicted in Figure 4 as the x-months-until-maturity variance (e.g., the variance of the 24-month horizon in Figure 4 is the average of the 24-month horizon variances for the November 2004, 2007, 2008, 2009, 2010, 2011, 2012, and 2013 contracts).