High Frequency Price Dynamics and Derivatives Prices For Continuously Produced, Storable Commodities

Craig Pirrong
John M. Olin School of Business
Washington University

September 14, 1998

1 Introduction

Commodity price dynamics, the economics of commodity storage, and commodity options pricing models have been the subject numerous recent studies, but the research in these areas has been largely disjoint.[1] This is unfortunate, because this article demonstrates that these subjects are inextricably

[1]For theoretical analyses of commodity price dynamics with applications to option pricing see Brennan (1986), Brennan and Schwartz (1985), Gibson and Schwartz (1991), Cortazar and Schwartz (1992), Amin, Ng, and Pirrong (1995), and Schwartz (1997). These papers take commodity price processes as exogenous. Fama and French (1992), Ng and Pirrong (1994, 1996) and Duffie and Gray (1995) present empirical evidence on commodity price volatility; Ng-Pirrong and Fama-French document empirically relations between price volatility and spot-forward spreads. Danthine (1977), Sheinkman and Schectman (1983), Williams and Wright (1991), Deaton and Laroque (1992, 1996), Chambers and Bailey (1996), and Routledge, Seppi, and Spatt (1997) analyze the theory of storage and undertake some empirical tests. These storage theory papers discuss some aspects of commodity price dynamics in low frequency data, but not high frequency data. The empirical tests in the storage-theory articles (especially those by Deaton-Laroque) are primarily limited to determining whether storage can explain the high autocorrelations observed in annual commodity price data. Zhou (1997) presents a theory of commodity price volatility and futures price mean reversion based on incomplete hedging. Zhou's work is quite distinct from the storage-theory genre, but has implications for volatility skews and the

linked. Solution of a dynamic recursive model of the market for a storable, continuously produced commodity demonstrates that the economics of storage and production exert a decisive influence on the short-term dynamics of commodity prices. When marginal costs are convex and increasing (due, for instance, to a capacity constraint) and demand is stochastic, commodity prices exhibit complex dynamics characterized by spot price variances that vary significantly and systematically with demand conditions and the amount of inventory on hand. In brief, the prices of storable, continuously produced commodities are highly volatile when demand and supply conditions are tight (due to high demand and/or low inventories) but exhibit little volatility when supply is abundant (due to low demand and/or high inventories). Moreover, when there are both transitory and persistent demand shocks, spot-forward price correlations also vary systematically with market conditions.

These results have important implications for derivatives pricing. Traditional commodity options pricing models (see especially Schwartz, 1997) do not incorporate state-dependent volatilities and correlations.[2] In the storage economy simulated herein, such constant volatility models generate large options mispricings. The constant volatility models misprice short dated, out-of-the-money calls by as much as 700 percent. Mispricings are somewhat smaller in magnitude for in-the-money, longer dated calls. The direction of the mispricings varies with market conditions. The constant volatility benchmark model grossly underprices options when demand is high and/or inventories are very short; this occurs when the commodity price is high. When demand is low or inventories are abundant, however, the benchmark model overprices commodity options. Thus, standard commodity option pricing models generate large, state-dependent pricing errors in the simulated commodity market because commodity price volatility is strongly state-dependent.

This analysis also implies that options prices for a storable commodity with stochastic demand and convex increasing marginal costs should exhibit

relation between price levels and price volatility that are similar to those derived below.

[2]See Brennan (1986), Brennan and Schwartz (1987), Gibson and Schwartz (1991), Cortazar and Schwartz (1992), Amin, Ng, and Pirrong (1996), and Schwartz (1997) for multi-factor, constant volatility models of commodity prices. These models use the spot price and the "convenience yield" as the two factors. Although these multi-factor models generate a richer set of dynamics for the commodity term structure, they do not permit the state-dependence in volatilities implied by the storage model.

volatility "skews" or "smiles."[3] Moreover, the shape of the volatility skew depends on current supply and demand conditions in the market. The relation between implied volatility and strike is strictly increasing for strikes surrounding the current forward price when this forward price is well above average. The relation between implied volatility and strike is U-shaped with a minimum occuring at a strike somewhat below the current forward price when this forward price is slightly above average or below average. Thus, if the storage model accurately describes real world pricing processes, diffusion models that specify a relation between fundamental conditions (as measured by the spot price or spot-forward spread since the state variables in the storage problem are likely unobservable to the econometrician) and volatility should produce option pricing results superior to those generated by models which do not embed these features.

Finally, the storage model has important implications for the hedging of commodity positions. Because the variance of commodity prices and the covariance between spot and forward prices are state-dependent, hedge ratios will also be state-dependent if there is a mis-match between the maturities of the hedging instrument and the claim being hedged. These variations in hedge ratios in response to changes in fundamental market conditions can be very large even if the maturity mismatch is relatively small.

This article focuses commodities that are produced continuously.[4] It is thus most relevant for the analysis of goods such as industrial metals (e.g., aluminum) or energy products (e.g., crude oil). Given that these are among the most heavily traded commodities in the world, even this focused analysis is of considerable import. There is some recent empirical evidence that is consistent with the basic predictions of the model. For example, as the theory predicts, energy and industrial metals prices are more volatile when spot-forward spreads indicate that supply conditions are tight than when

[3] See Rubenstein (1994), Derman and Kani (1995), and Chriss (1996) for a discussion of volatility smiles.

[4] Seasonally produced commodities, such as grains and oilseeds, present additional complexities. In particular, the price functions must vary by season; the pricing function in the spring differs from that in the fall. This substantially increases the computational costs. Moreover, expected harvest is likely to be a state variable in the seasonal model. This increases the dimensionality of the problem and therefore increases the computational burden even further. Pirrong (1998b) presents ananalysis of price behavior in the market for a seasonally produced commodity. Interestingly, volatility patterns are considerably different for seasonal and continuously produced commodities.

spreads indicate abundant supply.[5]

The remainder of this article is organized as follows. Section 2 presents a storage equilibrium model that incorporates convex, increasing marginal production costs and both persistent and transitory demand shocks. Section 3 reviews the main results concerning the relation between fundamental conditions, prices, spreads, volatility, and volatility skews. It also discusses the properties of volatility estimated using ARCH and GARCH models. Section 4 analyzes the ability of standard option pricing models to value options accurately in the storage economy. Section 5 summarizes the work.

2 Equilibrium in a Commodity Market with Storage

This section presents a model of prices and spreads in a competitive market for a storable commodity. Formally, time is indexed by $t \in \{1, \ldots, \infty\}$. To focus on short-term price dynamics, the model envisions short time intervals on the order of a month or week, if not less. This model assumes that the commodity is produced in every time period.

At each point in time, there is a competitive market for a commodity. The inverse demand for the commodity is given by $D(c_t, z_t, \eta_t)$ where c_t is the amount of the commodity consumed at t, $D_1(c, z, \eta) < 0$, $D_{11}(c, z, \eta) \geq 0$; z_t is a "persistent" demand shock with $D_2(c, z, \eta) > 0$; and η is a "transitory" demand shock, with $D_3(c, z, \eta) > 0$, where subscript i indicates the partial derivative with respect to argument i. All agents observe the demand shocks prior to trade at t and $D(c, z, \eta)$ is public information. The persistent demand disturbances have compact support $Z = \{z \in \mathbf{R} | -\infty < \underline{z} \leq z \leq \bar{z} < \infty\}$. $\mathcal{Z}$ denotes the Borel sets of Z. The sequence of random demand shocks $\{z_t\}$ is a first order Markov process generated by the transition function $Q(z, z')$ on $(Z, \mathcal{Z})$, where $Q(z, z')$ has the Feller property.[6] In the analysis that follows, it is assumed that Q is monotone. Montonicity can be interpreted as meaning that demand shocks are positively autocorrelated, hence the "persistent" nomenclature (Stokey and Lucas, 1989, Hopenhayn and Prescott, 1992,

[5]Ng and Pirrong (1994, 1996), Litzenberger and Rabinowitz (1995), Pirrong (1997).

[6]Stokey and Lucas (1989, p. 237) show that the assumption of a first-order Markov process results in no loss of generality because any higher order Markov process can be expressed as a first-order process with an expanded state space.

Chambers and Bailey, 1996.) The existence of a persistent component of demand is plausible for continuously produced industrial commodities (e.g., copper) because the documented persistence of real income shocks should result in persistent demand shocks for such commodities. State-dependent distributions of demand shocks also provide more interesting time series behavior of commodity prices, and are likely needed to explain the high autocorrelations observed in commodity price data.[7]

The transitory shocks η have compact support $H = \{\eta \in \mathbf{R} | -\infty < \underline{\eta} \leq \eta \leq \bar{\eta} < \infty\}$. $\mathcal{H}$ denotes the Borel sets of H. The demand shocks $\{\eta_t\}$ are i.i.d. with a probability measure $\mu(\eta_t)$ on $(H, \mathcal{H})$. The z and η are independent.

Earlier articles in this literature, including Schienkman and Schectman (1983), Danthine (1977), Deaton and Laroque (1992), and most of Williams and Wright (1991) assume transitory, i.i.d. demand shocks. More recent papers, including Chambers and Bailey (1996) and Deaton and Laroque (1996) include persistent demand shocks. No extant work includes both persistent and transitory demand shocks. Both permanent and transitory shocks are included in the model because a single shock model cannot explain salient features of empirical commodity price behavior. Specifically, a single shock model cannot explain the relation between spot-forward spreads and the correlation between spot and forward price changes documented by Ng-Pirrong (1994, 1996) and Pirrong (1997).

The commodity may be produced at each point in time. If the amount produced at time t is q_t, total production costs are $C(q_t)$, where $MC(q_t) \equiv C'(q_t) > 0$, $MC'(q_t) = C''(q_t) > 0$, and $MC''(q_t) = C'''(q_t) > 0$. That is, marginal costs are increasing and convex in q_t. This cost structure reflects limits on production capacity at any point in time. Output is non-negative, i.e., $q_t \geq 0$. This production technology is more general than that analyzed in the received literature; extant articles typically assume perfectly inelastic supply.

The commodity is storable and storage is costly. If amount s_t is stored in period t, only $(1-\delta)s_t$ remains at the beginning of period $t+1$. Moreover, storage is non-negative: $s_t \geq 0$. This non-negativity constraint is crucial to the understanding the dynamics of storables prices. If s_{t-1} was stored at $t-1$, then the following condition must hold in equilibrium:

$$(1-\delta)s_{t-1} + q_t - s_t = c_t.$$

[7] See especially Deaton and Laroque (1996).

In words, consumption at t equals the sum of production and carry-in less carry-out at t.

All producers and consumers in the market are price takers. That is, the market is perfectly competitive. Moreover, all market participants are risk neutral. By the Second Welfare Theorem the competitive equilibrium in this market is Pareto optimal. Therefore, due to risk neutrality, in a competitive equilibrium $\{s_t\}$ and $\{q_t\}$ solve:

$$\sup_{s_t \geq 0, q_t \geq 0} E_t\{\Sigma_{t=0}^{\infty} \beta^t G(s_t, q_t, s_{t-1}, z_t, \eta_t)\} \tag{1}$$

where $\beta = (1-\delta)/(1+R)$, R is the riskless rate of interest, and

$$G(s_t, q_t, s_{t-1}, z_t, \eta_t) = \int_0^{q_t+(1-\delta)s_{t-1}-s_t} D(y, z_t, \eta_t)dy - C(q_t)$$

That is, $G(.)$ equals the value of consumption net of production costs. It is readily verified that $G(.)$ is increasing in s_{t-1} and concave in its first three arguments.

Expression (1) can be rewritten in functional equation form:

$$\begin{aligned} v(x_t, z_t, \eta_t) = \sup_{s_t \geq 0, q_t \geq 0} \quad & G(s_t, q_t, s_{t-1}, z_t, \eta_t) \\ & + \beta \int \int v(x_{t+1}, z_{t+1}, \eta_{t+1}) Q(z_t, dz_{t+1}) \mu(d\eta_{t+1}) \end{aligned} \tag{2}$$

where $x_t = (1-\delta)s_{t-1}$.

Given the assumptions made here, there exist an equilibrium value function $v(x_t, z_t, \eta_t)$ and policy functions $s(x_t, z_t, \eta_t)$ and $q(x_t, z_t, \eta_t)$.[8] The first-order conditions for (2) are useful in solving for the value and policy functions. Defining $P(x, z, \eta) = D[x + q(x, z, \eta) - s(x, z, \eta), z, \eta]$ these conditions are:

$$P(x, z, \eta) \leq MC[q(x, z, \eta)] \tag{3}$$

with equality when $q(x, z, \eta) > 0$ and

$$P(x, z, \eta) \geq \beta \int \int P[(1-\delta)s(x, z, \eta), z', \eta'] Q(z, dz') \mu(d\eta') \equiv \beta F_1(x, z, \eta) \tag{4}$$

[8]See Pirrong (1998a) for proofs. These proofs exploit the contraction mapping theorem.

with equality when $s(x, z, \eta) > 0$. In words, expression (3) states that in equilibrium price equals marginal cost and expression (4) states that if storage is positive, price equals the discounted expected price (where discounting encompasses both storage cost δ and interest rate R.) When $s(x, z, \eta) = 0$ the "spot" price of the commodity may exceed the discounted expected spot price for next period. In the risk neutral economy, this expected spot price is the one-period forward price represented by $F_1(x, z, \eta)$.

Given these functions, it is possible to determine how spot and forward prices vary with the state variables. It is not possible to solve for the value and policy functions in closed-form, so it is not possible to determine the behavior of price variances and correlations in closed form. Instead, numerical methods are required. The next section describes the method employed to determine the equilibrium in a storage economy. I then proceed to discuss the implications of the storage equilibrium for commodity price variances and correlations and commodity option prices.

3 A Numerical Analysis of Short-term Price Dynamics for Continuously Produced Storable Commodities

Due to the difficulty of deriving analytically the properties of equilibria in recursive dynamic systems like that studied here, it is standard to employ numerical techniques to understand how these systems behave. For example, numerous studies of stochastic growth models employ numerical techniques.[9] Similarly, it is standard to solve storage problems numerically. Williams and Wright (1991), Deaton and Laroque (1992, 1996), Chambers and Bailey (1996), and Routledge, Seppi, and Spatt (1997) all utilize numerical techniques to study storage equilibria in models. However, these studies do not focus on high frequency/short term price dynamics that are of great interest and which must be understood to price contingent claims on storable commodities. In addition, with the exception of Routledge, Seppi, and Spatt they do not incorporate both transitory and persistent demand shocks. Moreover, none study the implications of increasing convex marginal costs despite the

[9] Eleven articles in the January, 1990 issue of the *Journal of Business and Economic Statistics* present numerical solutions to a non-linear rational expectations stochastic growth model. Taylor and Uhlig (1990) summarize these studies.

fact that these cost conditions are plausibly highly relevant for commodities such as industrial metals (e.g., copper), crude oil, heating oil, gasoline, and natural gas because capacity constraints are likely to create such cost structures.[10] Finally, they do not analyze the implications of state-dependent volatility for the option pricing and the implied volatilities of commodity options.

This section addresses these issues. I first describe briefly the numerical methodology. I then show how the price functions generated by the solution of the storage equilibrium value and policy functions can can be employed to determine spot and forward price variances and spot-forward correlations as a functions of the state variables. I then discuss numerical results which show how spot and forward prices structures, spot and forward percentage price change variances, and the probability distributions of spot and forward percentage price changes change as demand and initial supply vary.

3.1 Methodology

The basic contours of the numerical solution of recursive dynamic economic models are fairly well understood, so only a brief description of the methodology is required. There are three state variables: a persistent shock z, a transitory shock η, and initial inventory x. The process for the persistent shock is assumed to be:

$$\Delta z_t = (\rho - 1) z_{t-1} + \Delta v_t \tag{5}$$

where and Δv_t is an i.i.d. normal shock with mean zero and variance Δt. Furthermore, η_t is a white noise transitory shock:

$$\Delta \eta_t = -\eta_{t-1} + \Delta w_t \tag{6}$$

where Δw_t is another i.i.d. process with zero mean and variance Δt. Assume that Δv_t and Δw_t are uncorrelated.

The first step of the analysis is to discretize the problem by establishing a grid in z, η, and x.

The values of the grid in z are bounded by -2.4 and +2.4 and are equally spaced with increments Δz. There are N_z points along the grid,

$$\mathcal{A} = \{-2.4, -2.4 + \Delta z, -2.4 + 2\Delta z, \ldots, 2.4 - \Delta z, 2.4\}.$$

[10] Bresnehan and Suslow (1989), Pindyck (1992).

Define $z_j = -2.4 + (j-1)\Delta z$. The transition probability matrix for z is constructed as follows:

$$\pi(z_j, z_k) \equiv \Pr[z_{t+1} = z_k | z_t = z_j] = N(\epsilon_1) - N(\epsilon_2) \quad (7)$$

where

$$\epsilon_1 = .5[z_{k+1} + z_k] - \rho z_j$$

and

$$\epsilon_2 = .5[z_k + z_{k-1}] - \rho z_j$$

where ϵ is an i.i.d. unit normal variate. That is, if $z_t = z_j$, (i.e., the initial value of z is at the j'th point on the z grid) the probability that the value of z in the next period is $z_{t+1} = z_k$ (i.e., is at the k'th point on the grid) equals the probability that a standard normal variate falls in an interval of length Δz centered on $z_k - \rho z_j$.[11]

The values of η are also bounded by -2.4 and +2.4 and are equally spaced with increments $\Delta\eta$. There are N_η points in the η dimension of the grid. Call $\eta_j = -2.4 + (j-1)\Delta\eta$. In the discretization,

$$p(\eta_j, \eta_k) \equiv \Pr[\eta_{t+1} = \eta_k | \eta_t = \eta_j] = N(e_1) - N(e_2) \quad (8)$$

where

$$e_1 = .5(\eta_{k+1} + \eta_k)$$

and

$$e_2 = .5(\eta_k + \eta_{k-1})$$

That is, the probability that the demand shock will equal η_k given that its previous value was η_j is equal to the probability that an i.i.d. standard normal variate will fall in an interval of length $\Delta\eta$ centered on η_k.

For each value of z and η in the grid, values of x are spread equally along the interval $[0, x_{max}]$ where x_{max} is determined in a trial and error process that ensures that equilibrium storage never exceeds x_{max} in long simulation runs.

Demand functions are linear in both consumption and the demand shocks:

$$D(x+q-s, z, \eta) = a - b(x+q-s) + \sigma_z z + \sigma_\eta \eta.$$

[11] See Deaton and Laroque (1996) for a similar discretization of an AR1 process.

The σ_z and σ_η are parameters that measure the sensitivity of demand to persistent and transitory shocks.[12] Cost functions are of the form:

$$C(q) = \theta + \frac{\phi}{\gamma - q}.$$

This cost function exhibits increasing and convex marginal costs for $q < \gamma$. γ represents a production capacity constraint.

Given the grid and demand and cost functions, an initial guess for the functions $P(x, z, \eta)$ is formed in three steps. (There is one such function for each pair $\{z, \eta\}$.) First, for each value of x it is assumed that $s(x, z, \eta) = .8x$. Second, given this guess, first order condition (3) is solved for q. This determines the price at each of the x points on the grid. A fourth order polynomial in x is fit to these prices using OLS, and the resulting polynomial function $P(x, z, \eta)$ is used as the initial guess for the pricing function for this $\{z, \eta\}$ pair.

Given an initial guess for the price functions, for each z, η and x in the grid, the following equations are solved (using the Newton-Raphson method) for s and q:

$$D(x + q - s, z, \eta) = C'(q)$$

$$D(x + q - s, z, \eta) = \beta \sum_{k=1}^{N_z} \sum_{i=1}^{N_\eta} \pi(z, z_k) p(\eta_i) P[(1 - \delta)s, z_k, \eta_i] = \beta F_1(x, z, \eta).$$

The first equation equates price and marginal cost. The second requires price to equal expected price. If $s(x, z_j, \eta_j) > 0$, then price at this grid point is set to $D[x + q(x, z_j, \eta) - s(x, z_j, \eta)]$. If $s(x, z_j, \eta) \leq 0$, $s(x, z_j, \eta) = 0$ and price at this grid point equals $D[x + q(x, z_j, \eta)]$. After determining prices at each node of the grid, for each $\{z, \eta\}$ a fourth order polynomial in x is fit to these prices using OLS. These new polynomials are used as the $P(x, z, \eta)$ functions in the next iteration and the process is repeated. The process stops when the average absolute percentage price change between iterations is small (e.g., .001 percent).

Upon convergence this process defines a spot price surface $P^*(x, z, \eta)$. Using this surface, it is possible to study how expected prices, spot and forward price volatilities, and spot and forward price distributions vary with

[12] It is equivalent to assume that (a) demand is $a - b(x + q - s) + z + \eta$ and (b) the variance of the z process in (5) is $\sigma_z^2 \Delta t$ and that the variance of the η process in (6) is $\sigma_\eta^2 \Delta t$.

fundamental supply and demand conditions measured by the state variables x, z, and η.

3.2 Price Variances and Correlations in the Storage Economy

The equilibrium price surface $P^*(x, z, \eta)$ can be used to determine the characteristics of commodity price dynamics. The equilibrium price surface and transition probabilities also imply a forward price function $F(x, z, \eta, t, T)$ that gives the forward price of the commodity as of time t for delivery at any time $T > t$. This surface can be employed to determine the variance of the forward price and the correlation between spot and forward prices.

Applying a Taylor expansion to $P^*(x, z, \eta)$ and $F(x, z, \eta, t, T)$, substituting from (5) and (6), and ignoring terms of $o(\Delta t)$ implies:

$$\Delta P^* = P_z^*(\rho - 1)z_{t-1} - P_\eta^* \eta_{t-1} + P_z^* \Delta v_t + P_\eta^* \Delta w_t \tag{9}$$

$$\Delta F_t = F_z(\rho - 1)z_{t-1} - F_\eta \eta_{t-1} + F_z \Delta v_t + F_\eta \Delta w_t. \tag{10}$$

The variances of the spot and forward price changes (ignoring terms of $o(\Delta t)$) are:

$$\sigma_{P^*}^2 \equiv \frac{E(\Delta P^* - E\Delta P^*)^2}{\Delta t} = P_z^{*2} + P_\eta^{*2} \tag{11}$$

$$\sigma_F^2 \equiv \frac{E(\Delta F - E\Delta F)^2}{\Delta t} = F_z^2 + F_\eta^2. \tag{12}$$

The covariance between spot and forward price changes is:

$$cov(\Delta F_t, \Delta P_t^*) = F_z P_z^* + F_\eta P_\eta^*. \tag{13}$$

The correlation between the spot and forward price changes is therefore:

$$corr(\Delta F_t, \Delta P_t^*) = \frac{F_z P_z^* + F_\eta P_\eta^*}{\sqrt{(F_z^2 + F_\eta^2)(P_z^{*2} + P_\eta^{*2})}}. \tag{14}$$

Therefore, by taking the partial derivatives of the spot and forward price functions generated by the solution to the storage problem it is possible to estimate the variances of spot and forward prices and the correlation between them.[13]

[13] An examination of (14) shows that multiple shocks are required to generate correlations other than -1, 0, or +1.

Since spot and forward prices are a function of the state of the economy (as measured by carry-in x, and demand shocks η and z), the foregoing expressions imply that variances and correlations should be state-dependent as well. The next section demonstrates that this is indeed the case in the numerical solution of the storage problem. I then show that this has important implications for derivatives pricing.

3.3 Fundamentals and Commodity Price Dynamics

Figures 1-5 present typical results for a numerical implementation of the model. The parameters used to construct the Figures were: $\rho = .5$, $a = 40$, $b = .2$, $\sigma_z = .5$, $\sigma_\eta = .5$, $\theta = 14$, $\phi = 3$, and $\gamma = 120$. Although the exact results are unique to these parameters, all of the numerous parameter combinations studied produce similar qualitative results so those presented are representative. The parameters are calibrated through a trial-and-error process to give reasonable weekly volatilities (e.g., weekly volatilies of the magnitudes observed in the copper and oil markets). Therefore, the values of R and δ are reasonable estimates of weekly interest and storage charges; specifically, $R = .002$ and $\delta = .001$. Computational considerations lead to a choice of $N_z = N_\eta = 9$.

Figure 1 illustrates how equilibrium spot prices P^* vary with the state variables. The "Demand Index" axis gives an ordinal measure of the demand state. The demand index is $(j-1)N_z + i$ where $j \leq N_z$ is the permanent demand shock index value and $i \leq N_\eta$ is the transitory demand shock index value. This structure explains the sawtooth configuration of the graph. As one moves away from the origin along the demand index axis, the transitory shock value increases while the persistent shock value remains constant until the transitory shock reaches η_{N_η}. The persistent demand shock index then increases by one and the transitory shock value falls to η_1. The process repeats until the last point on this axis, which measures prices for the maximum values of the persistent and transitory demand shocks. The "Storage Index" is an ordinal measure of carry-in, with a value of 1 corresponding to zero carry-in and a value of 12 corresponding to x_{max}.

The figures show that for given values of η and z, prices are decreasing and convex in x. Moreover, for given x and z (η) prices are increasing and convex in η (z). Thus, greater scarcity implies higher prices, and an increase in scarcity (a rise in z or η or a fall in x) raises prices at an increasing rate. This curvature reflects in large part the convexity of marginal costs.

Figure 2 depicts the spread between the spot price and the two-period forward price. The two-period forward price is determined using iterated expectations. The solution to the original dynamic program generates both $P(x,z,\eta)$ and $F_1(x,z,\eta)$ where P is the spot price and F_1 the one-period forward price. The two-period forward price is the expectation of the one period forward price one period hence:

$$F_2(x,z,\eta) = \sum_{j=1}^{N_z} \pi_z(z,z_j)p(\eta)F_1[(1-\delta)s(x,z,\eta),z,\eta].$$

Note that the spot-forward spread, defined as $F_2 - P$, exhibits what is sometimes referred to as a "supply of storage" relation between spot-forward spreads and carry-in. Specifically, for a given z and η, $F_2 - P$ is an increasing, concave function of x. Moreover, this spread is also a decreasing, concave function of z and η. This relation is observed in many markets. Figure 3 depicts an empirical supply of storage curve relating the spot-three month spread and exchange warehouses for lead. There is a clear increasing, concave relation between the spread and carry-in. There is some scatter in the points. This reflects the fact that demand also influences the spread. As a consequence, for a given level of carry-in the spread is lower when demand is low than when it is high.

Figure 4 is the crucial exhibit in the analysis because it depicts the relation between fundamentals and percentage price change variance. Expression (11) implies that

$$var(\frac{\Delta P_t^*}{P_t^*}) = \frac{P_z^{*2} + P_\eta^{*2}}{P_t^{*2}}. \qquad (15)$$

The partial derivatives in the expression are evaluated at each point on the grid using an implicit finite difference approximation.[14] That is, for $N_z >$

[14]Using partial derivatives to calculate the variances and correlations is substantially faster computationally than calculating the moments directly from the price function and the transition probability matrix. Moreover, the partial derivative approach provides some intuition underlying the results. For instance, it is quite intuitive that prices are more sensitive to demand shocks when demand is high and/or stocks are low than when demand is low and/or stocks are high. The partial derivative representation readily translates this result into an understanding of the effect of demand and supply conditions on variances and correlations. Direct calculation of moments was performed for several implementations of the model. These moments were compared to those derived using the partial derivative approach. The results are virtually identical.

$j > 1$,

$$P_z^*(x_k, z_j, \eta_i) = \frac{P^*(x_k, z_{j+1}, \eta_i) - P^*(x_k, z_{j-1}, \eta_i)}{2\Delta z}.$$

For $j = 1$,

$$P_z^*(x_k, z_j, \eta_i) = \frac{P^*(x_k, z_{j+1}, \eta_i) - P^*(x_k, z_j, \eta_i)}{\Delta z}.$$

and for $j = N_z$,

$$P_z^*(x_k, z_j, \eta_i) = \frac{P^*(x_k, z_j, \eta_i) - P^*(x_k, z_{j-1}, \eta_i)}{\Delta z}.$$

The partial derivative with respect to η, P_η, is similarly defined at each point on the grid.

Figure 4 shows that the spot return variance increases with η (for given z and x), declines with x (for given z and η), and increases with z (for given η and x). Variance increases at an increasing rate when z or η rises and when x falls. When demand is high and initial supplies low, volatility is very high. Indeed, volatility is many orders of magnitude larger when $x \approx 0$, $z \approx 2.4$, and $\eta \approx 2.4$ than when $x \approx x_{max}$, $z \approx -2.4$, and $\eta \approx -2.4$. The variance when demand is at its maximum and carry-in is zero is 250 times greater than when demand is at its minimum and carry-in is larger than its mean in long simulation runs. Empirical variances for industrial metals presented in Ng-Pirrong (1994) exhibit such wide variations.

A comparison of Figures 1, 2, and 4 reveals a close relation between volatility, prices, and spreads. Volatility is high when prices are high and spot-forward spreads are wide. The relation between spreads and volatility exhibited in the figures is virtually identical to the empirical pattern found by Ng and Pirrong for four industrial metals (copper, aluminum, lead, and zinc, 1995) and refined petroleum products (heating oil and gasoline, 1996). Litzenberger and Rabinowitz (1995) demonstrate that crude oil implied volatility varies closely with backwardation: volatility implied by crude oil futures options is higher, the greater the backwardation in the market. This finding is consistent with their theory's implications, although Litzenberger and Rabinowitz's explanation of the phenomenon differs from that advanced here.

The numerical analysis also demonstrates that the volatility of the spot price is never smaller than the volatility of the forward price, that these volatilities are almost identical when x is large, and that the difference between spot and forward volatility is increasing in price.[15]

Finally, Figure 5 shows that the relation between carry-in, demand, and the correlation between spot and forward returns is consistent with the conjectures made in section 3 and some existing empirical evidence. Figure 5 depicts $corr(\Delta P^*/P^*, \Delta F_2/F_2)$. As before, the partial derivatives needed to evaluate this correlation using (13) are approximated using an implicit finite difference scheme. Note that correlation is nearly 1.00 when carry-in is large, and/or demand is low. However, when carry-in is low, the correlation tends to fall as demand rises. Similarly, for a given level of demand, the correlation tends to fall as carry-in declines. This decline in correlation becomes precipitous as carry-in approaches 0. A comparison of figures 1, 2, and 5 reveals that the spot-forward correlation is typically high when prices are low and backwardation is small. Conversely, the spot-forward correlation is low when prices are high and backwardation is severe. This is consistent with the empirical evidence from metals markets presented by Ng-Pirrong (1994), and for energy markets by Ng-Pirrong (1996) and Pirrong (1997).

The intuition behind the behavior of the correlation is readily understood. When $s_t(x_t, z_t, \eta_t) = 0$ (i.e., the optimal solution to the storage problem at t dictates zero carryover), $F_\eta = 0$ and $P_\eta > 0$. $F_\eta = 0$ when there is no storage because a purely transitory shock can have no effect on forward prices; storage is the only link between transitory demand shocks and prices, and this linkage is broken when there is a stockout. Thus, by (14) the correlation is less than 1.00 in this case. Conversely, when carry-in is immense, $F_z/F \approx P^*_z/P^* \approx F_\eta/F \approx P^*_\eta/P^*$. This is true because with storage, the distinction between transitory and persistent shocks is trivial; storage causes transitory shocks to have a persistent effect. Moreover, when storage is high, arbitrage ensures that spot and forward prices tend to move in similar proportions when demand changes. This generates a spot-forward correlation of nearly 1.00.

Repeated experimentation with the model shows that supply conditions are crucial in determining the behavior of volatility and spot-forward corre-

[15] The relevant graphs that illustrate these points are omitted, but are available on request from the author. By volatility I mean the standard deviation in the *percentage* price changes. It is possible that the variance of forward price changes exceeds the variance of spot price changes because the forward prices sometimes exceed spot prices. Note that when stocks are very large $F \approx P^*/\beta$. Therefore, $dF/F \approx dP^*/\beta/P^*/\beta = dP^*/P$ while $dF = dP^*/\beta > dP^*$.

lations. Holding mean demand and the volatility of demand shocks constant, the variance and correlation functions depicted in Figures 4 and 6 are very sensitive to the choice of production capacity γ. Increasing γ causes volatility to fall, and reduces the sensitivity of the volatility function to demand shocks. That is, the sawtooths in the figures become flatter.[16] Not surprisingly, increasing σ_z and σ_η also causes volatility to increase.

In sum, this analysis demonstrates that a dynamic storage model predicts particular relations between inventory, demand, prices, spot-forward spreads, spot return variances, and spot-forward return correlations. These predictions are consistent with existing empirical evidence on the dynamics the prices of continuously produced commodities.

Since variance exerts a decisive influence on options pricing, the foregoing analysis strongly suggests that models that do not embody the relation betweeen fundamentals and commodity price volatility may severely misprice commodity options. Section 4 demonstrates that this is indeed the case in the storage model studied here.

3.4 Hedging

The variances and correlations discussed above have important implications for hedging. Consider a firm that uses inventories (or a short-dated futures contract) to hedge a forward delivery commitment.[17] It is well known that the variance minimizing hedge ratio in this case is $-cov(\Delta P^*, \Delta F)/var(\Delta P^*)$. Since both the numerator and denominator in this expression are state dependent, the hedge ratio will be so as well. Figure 6 depicts the hedge ratio for a two-period forward delivery commitment as a function of demand and inventory. For simplicity, the permanent demand shock z is held constant at its expected value of 0 and the transitory demand shock increases as one moves down the demand index axis. Note that the hedge ratio is nearly 1.00 when inventories are high, but that the hedge ratio drops sharply when carry-in falls to near zero. Moreover, for such small values of inventory the absolute value of the hedge ratio declines as demand increases, i.e., as η

[16] A similar result obtains when holding γ constant but increasing the demand intercept a.

[17] The actions of Metalgesellschaft present an extreme example of this phenomenon. This firm hedged oil delivery commitments extending from one month to 10 years forward using one month futures contracts. See Ross (1995), Pirrong (1996), and Schwartz (1997) for analysis of the firm's activities.

increases. For example, when the transitory demand shock is 1.2 standard deviations above 0 and inventory is about 1.5 standard deviations below its average value in long simulations, the hedge ratio is .88. When transitory demand is 1.2 standard deviations above zero and carry-in is zero, the hedge ratio falls to a mere .14. In conjunction with figures 1 and 2, this implies that the hedge ratio is well below 1.00 when prices are high and the spot forward spread is negative, but is close to 1.00 when prices are low and spot and forward prices are at full carry. The analysis also implies that holding inventory constant, the hedge ratio declines in absolute value as the permanent demand shock increases.[18]

It should be noted that these deviations from a one-for-one hedging strategy occur in Figure 6 even though the timing mismatch between the hedging instrument and the thing being hedged is small. Ross (1995) and Schwartz (1997) show that a one-for-one hedging strategy is typically not variance minimizing if the maturity difference is large. This is true because the stationarity of the storage economy causes the variance of the deferred obligation to fall as its maturity increases. The variance minimizing hedger should therefore choose a smaller hedge ratio for longer maturity instruments. Hedge ratios may fall well below one even if the timing mismatch is small (as is the case in Figure 6) when inventories are low or demand is very high because spot and forward prices are not highly correlated under these circumstances. Thus, even a one- or two-month mismatch should induce a hedger to choose a hedge ratio well below one when prices are high and the market is inverted (i.e., $P^* > F$).

The correlation analysis presented in the previous subsection implies that hedging effectiveness is also state dependent. Hedging effectiveness, that is, the fraction of variance that a hedger can eliminate through the choice of the variance minimizing hedge ratio, is equal to the squared correlation between the spot and forward returns. Since this correlation declines as inventory falls and demand increases, hedging effectiveness declines in these circumstances as well.

In sum, the state dependence of commodity price variances and correlations has important implications for hedgers. Hedgers should choose different hedge ratios when supply/demand conditions are tight than when they are

[18] A similar analysis holds when the hedger uses a forward contract to hedge inventory (or a shorter dated forward delivery commitment). In brief, in these circumstances a hedger should choose a hedge ratio well above 1.00 when supply/demand conditions are tight and a hedge ratio of close to 1.00 when supply/demand conditions are slack.

slack. In particular, this analysis implies that hedging in a high price, inverted market is very difficult. In such a market variance minimizing hedge ratios may change dramatically over short time periods. Moreover, hedging effectiveness is typically low during these periods.

3.5 Time Series Properties of Commodity Volatility

Ng-Pirrong (1994, 1996) document that after controlling for scarcity (as measured by price or the squared spot-forward spread), industrial metal and energy price volatilities still exhibit GARCH effects. That is, when estimating the following system of equations the parameters γ_1, γ_2, and γ_3 are all positive and significant:

$$\Delta \ln P_t = \alpha_0 + \sum_{j=1}^{10} \alpha_j \Delta \ln P_{t-j} + \alpha_{11} \ln P_{t-1} + \epsilon_t \tag{16}$$

$$\sigma_t^2 = \omega + \gamma_1 \sigma_{t-1}^2 + \gamma_2 \epsilon_{t-1}^2 + \gamma_3 P_{t-1} \tag{17}$$

where σ_t is the standard deviation of the residual return ϵ_t. Moreover, $\gamma_1 + \gamma_2$ is near 1.00 in these studies. This section examines whether similar results obtain for the storage economy.

To investigate this issue, I use a random number generator and the solution to the storage model described above to create 1000 series of 2500 observations of the spot price. For each series, I fit the following model: where the variance of ϵ_t is given by (18). The inclusion of the lagged log spot price in the mean equation (19) reflects the fact that the storage model predicts that the spot price should be mean reverting. Indeed, in the simulations this coefficient is always negative and significant.

Table 1 presents the mean values of selected coefficients from (18) and (19). Note that the lagged price coefficient is positive in the variance equation, as expected given the earlier analysis; it is highly significant in each of the simulations. That is, the return variance is greater, the higher the price. Moreover, the ARCH coefficient γ_2 is also positive. This too is unsurprising because as noted earlier the spot price does not capture all of the effects of both state variables on volatility. The ARCH term ϵ_{t-1}^2 apparently proxies for the factors affecting scarcity imperfectly measured by price. Finally, the coefficient on lagged volatility, γ_1, is positive but far less than .9; this coefficient is sometimes insignificant in the individual simulated estimates.

This last finding varies from the empirical evidence presented in Ng-Pirrong, who find values for the lagged variance coefficient on the order of .9. γ_1 measures the persistence of volatility disturbances. The coefficient implied by the dynamic programming model is substantially smaller (and the ARCH coefficient γ_2 is larger) than found in real world data. Although volatility persists in the storage model, this occurs in large part because the spot price is autocorrelated due to storage and the persistence of demand shocks. Controlling for price, volatility shocks in the simulations exhibit less persistence than found in metals data by Ng-Pirrong.[19]

This last finding is the main deviation between the empirical evidence on commodity price dynamics and the implications of the storage model. One possible explanation for this result is that the storage model analyzed assumes that the volatility of demand shocks is constant over time. If the demand shocks exhibit GARCH behavior, this may be translated into the behavior of price volatilities.

Since the nature of demand disturbances affects the economics of the storage decision, verification of this conjecture requires solution of a modified storage problem that permits GARCH demand disturbances. This problem involves more state variables and is therefore substantially more numerically demanding. Solution of this problem awaits future research.

3.6 Summary

Numerical solution of a recursive dynamic storage model in a market for a continuously produced good with increasing, convex marginal production costs makes strong predictions about the relation between the dynamics of spot and forward prices on the one hand and inventories and demand conditions on the other. Specifically, the model predicts: (a) a "supply of storage" relation between spot-forward spreads and inventories, (b) a strong relation between spot price and spot-forward spread levels on the one hand and the volatility of spot prices on the other, and (c) a strong relation between spot prices and spot-forward spreads and the spot-forward return correlation.

There is some empirical evidence from the industrial metals and energy markets consistent with these implications. These results therefore have im-

[19]Given the strong mean reversion in price, volatility shocks dampen out quickly in this model. This is difficult to reconcile with the evidence of fractional integration of volatility in metals markets (Teyssiere *et al*, 1997). Fractional integration of the demand shock volatility may be necessary to explain this result.

portant consequences for our understanding of the dynamics of commodity prices which merit detailed empirical examination. Since volatility is central to option pricing, these results also have important implications for the pricing of commodity contingent claims, as the next section shows.

4 State-dependent Volatility and Derivatives Pricing

Pricing models based on lognormal diffusion processes are standard in both the literature and practice on commodity options.[20] The analysis in sections 3.1-3.3 suggests that the standard models are severely mis-specified. As a consequence they will misprice commodity options. This section demonstrates that these mispricings are likely to be extremely large and to be strike and maturity dependent.

4.1 A Two Factor Constant Volatility Model of Commodity Option Price Dynamics

I use Model II of Schwartz (1997) as a benchmark options pricing model. The Schwartz model posits two factors, a spot price S_t and a convenience yield δ_t. The stochastic processes for these variables are given by:

$$dS = (r - \delta)Sdt + \sigma_1 S dw_1 \tag{18}$$

$$d\delta_t = \kappa(\alpha - \delta) + \sigma_2 dw_2 \tag{19}$$

where dw_1 and dw_2 are standard Brownian motions with correlation ρdt. The volatility parameters σ_1 and σ_2 and the correlation ρ are constants. Thus, both the spot price and the convenience yield have constant volatility. This model does permit "twists" in the term structure of futures prices, but does not allow the variance of spot or forward prices to depend on the level of price or convenience yield. This raises the question: Can such a model price commodity derivatives accurately if commodity prices are generated by a storage economy like that modeled above?

[20]See Brennan (1986), Brennan and Schwartz (1987), Gibson and Schwartz (1991), Cortazar and Schwartz (1992), Amin, Ng, and Pirrong (1996), and Schwartz (1997) for multi-factor, constant volatility models of commodity prices.

To answer this question, I estimate the relevant parameters for this model in a data sample generated by simulation of the storage model solved above, estimate the parameters of the S and δ processes from this sample, and then compare the prices of options on the commodity generated by these parameters and Schwartz's Model II to the "true" prices of these options in the storage economy; the true option prices are implied by the pricing function P^* and the transition probabilities $\pi(z_1, z_2)$ and $p(\eta)$.

I first simulate a time series of 1000 observations using the numerically solved pricing function for the storage economy described in sections 3.1-3.3. For each observation, the spot price, the quantity stored, and forward prices for delivery in one, two, three, and four periods hence are generated. Storage induces path-dependence in the spot price, so estimation of forward prices requires calculation of expected values over all future possible paths. For example, given an initial carry-in, a four period forward price requires calculation of the expected value over each path of shocks that can occur over the next 4 periods. In the discretized framework, with N_z possible values of the demand shock z and N_η possible values of the demand shock η, there are $(N_\eta N_z)^\tau$ possible paths of demand shocks over τ periods. The τ period forward price is the expected value over these paths of the spot price in τ periods.

Since the spot price is observable in this synthetic time series, it is only necessary to estimate one state variable in the sample–the convenience yield.[21] This is done interatively. In the first stage of the process, for each observation $t = 1, \ldots, 1000$ I estimate the convenience yield as:

$$\hat{\delta}_t = \ln(1 + R) + .25 \sum_{\tau=1}^{4} \frac{1}{\tau} \ln \frac{S_t}{F_t^\tau}$$

where S_t is the simulated spot price at t, and F_t^τ is the simulated τ period forward price as of t.[22] Given this initial estimate of the convenience yield series, I estimate the following regression:

$$\Delta\hat{\delta}_t = a - b\hat{\delta}_{t-1} + e_t \tag{20}$$

[21]In contrast, in Schwartz's (1997) empirical analysis both the spot price and the convenience yield are non-observable. He uses a Kalman filtering technique to extract these unobserved state variables from the data.

[22]This is motivated by the following standard definition of the convenience yield: $F_t^\tau = e^{(r-\delta_t)\tau} S_t$ where r is the continuously compounded interest rate $r = \ln(1 + R)$.

where a and b are parameters and e_t is an error term. κ is then set equal to b and α is then set equal to a/b. Moreover, σ_2^2 is set equal to $var(e_t)$. Again using the $\hat{\delta}_t$ series, I then calculate $u_t = \Delta \ln S_t - \hat{\delta}_t - R$ and set $\sigma_1^2 = var(u_t)$. Finally, $\rho = corr(e_t, u_t)$.

Given this set of parameters, the convenience yield is reestimated using expression (19) in Schwartz. Specifically:

$$\ln F_t^\tau - \ln S_t - A(\tau) = -\delta_t \frac{1 - e^{-\kappa\tau}}{\kappa} \tag{21}$$

where $A(\tau)$ is a function of τ that depends on α, κ, σ_1, σ_2, and ρ; this function is given by expression (20) in Schwartz (1997). In the next iteration, for each $t = 1, \ldots 1000$, $\hat{\delta}_t$ is the slope coefficient in a regression of $\ln F_t^\tau - \ln S_t - A(\tau)$ against $(1 - exp(-\kappa\tau))/\kappa$. Given the new estimates of the convenience yields for each t, (20) is re-estimated and κ, α, and σ_2 are determined as before. Now $\sigma_1^2 = var(u_t)$ where $u_t = \Delta \ln S_t - \hat{\delta}_t - .5\sigma_1^2$ and $\rho = corr(e_t, u_t)$.[23] This process continues until the average absolute percentage change in the estimated $\hat{\delta}_t$ between iterations is smaller than .001 percent. The estimates of σ_1, σ_2, κ, α, and ρ are then used to price call options on the commodity.

In the Schwartz model, futures/forward prices with τ periods to expiration are lognormally distributed with instantaneous variance:

$$\sigma_F^2(\tau) = \sigma_1^2 + \sigma_2^2 \frac{(1 - e^{-\kappa\tau})^2}{\kappa^2} - 2\rho\sigma_1\sigma_2 \frac{(1 - e^{-\kappa\tau})}{\kappa} \tag{22}$$

Since the future is a traded asset with a variance that is a function of time alone, it is possible to price simple options on the commodity using the Black model with:

$$d_1 = \frac{\ln(\frac{F_t^\tau}{K}) + .5 \int_0^\tau \sigma_F^2(u) du}{\sqrt{\int_0^\tau \sigma_F^2(u) du}} \tag{23}$$

where K is the strike price of a call option. Given the parameters estimated using the process described above, it is therefore possible to determine the price of any European call or put option on the commodity.

The "true" value of the option is derived as follows. For given starting values of x, z, and η, and thus of spot price P^*, the terminal price P_p and the path probability π_p determined for each of the $(N_\eta N_z)^\tau$ paths p_τ. One,

[23]The σ_1 used to calculate u_t is the value generated in the previous iteration.

two, and three period options are priced, so $\tau = 1, 2, 3$.[24] The value of the call with strike K equals that expires in τ periods is:

$$c(x, z, \eta, K) = \sum_{p=1}^{(N_\eta N_z)^\tau} \pi_p \max[P_p^* - K, 0]. \tag{24}$$

The forward price for maturity in τ periods is:

$$F^\tau(x, z, \eta, K) = (1 + R)^{-\tau} \sum_{p=1}^{(N_\eta N_z)^\tau} \pi_p P_p^*. \tag{25}$$

Options are priced using (24) and (25) for 9 different strikes for each expiration. These strikes are spaced evenly around the forward price in increments of .05. That is, an at-the-money, four out-of-the-money, and four in-the-money calls are priced for each maturity for a variety of starting values for x and z.

4.2 Results

A comparison of option prices generated by the two models implies that the constant volatility model generates serious mis-pricing of options due to the mis-specifications discussed above. Call $C(K, \tau)$ the "true" price of an option expiring in τ periods with strike price K as given by (24). Call $C_S(K, \tau)$ the price of the option implied by the Schwartz model. Table 2 reports

$$B(K, \tau) = \frac{C(K, \tau)}{C_S(K, \tau)} - 1 \tag{26}$$

for several values of initial carry-in and permanent demand shock; for simplicity, the temporary demand shock is assumed to equal 0 in all cases studied. That is, $B(K, \tau)$ is a measure of the bias in the option prices generated by the Schwartz model. $B(K, \tau) > 0$ implies that the Schwartz model underestimates the option price. $B(K, \tau) < 0$ implies that it overestimates the option price.

Panel A of Table 2 assumes that the economy is in a low price state with high carry-in and a permanent demand shock 1.2 standard deviations below

[24]Computational considerations limit τ. The number of computations increases exponentionally with τ.

its mean. Panel B assumes that the economy is in a medium price state with carry-in equal to its average in long simulations and a permanent demand shock equal to its expected value of 0. Panel C assumes a high price state with zero carry-in and a permanent demand shock 1.2 standard deviations above the mean. In each panel, the strike price increases as one proceeds down any column, and time to expiration increases as one proceeds to the right along any row. Row 5 gives the price of at-the-money calls; rows 6-9 give the prices of out-of-the-money calls and rows 1-4 give the prices of in-the-money calls. Columns 1, 2, and 3 give the prices of 1 period, 2 period, and 3 period options, respectively.

An examination of the various panels gives clear evidence of strike price and time-to-expiration biases. Moreover, the size and sign of these biases vary depending on whether the economy is in a high price, low price, or medium price state. In a high price state (Panel C) the Schwartz model systematically underprices call options. The pricing biases are more severe, the higher the strike price and the shorter the time to expiration. That is, the constant volatility model seriously underprices shortest-dated out-of-the-money calls. The mispricing for the deepest-out-of-the-money, short-dated calls is over 700 percent. Even the deepest-in-the-money, longest-dated calls are underpriced by 25 percent. In a medium price state (Panel B), the Schwartz model tends to *overprice* short-dated in-the-money and at-the-money calls and underprice short-dated out-of-the-money calls. The model underprices longer-dated calls. The pricing bias is increasing in strike for all maturities. In the low price state (Panel A), the Schwartz model overprices all calls studied. In this state, the underpricing is more severe, the higher the strike. For a given strike, mispricing is not always monotonic in time to expiration, although it does decrease (in absolute value) in maturity for at-the-money and out-of-the-money strikes.

These mispricings are readily understood based on the analysis of the behavior of volatility in the storage economy provided in section 3.3. When applied to the simulated data generated by the storage economy model, the Schwartz model essentially estimates an average spot volatility. This volatility is smaller than the true volatility when prices are high, and is larger than the true volatility when prices are low because volatility generally increases in price. This explains the underpricing of options in the high price environment and the overpricing in the low price environment.

Moreover, the lognormal distribution underlying the Schwartz model does not capture the skewness and kurtosis in return distributions in the storage economy. These distributions are highly right-skewed and leptokurtotic in high and medium price environments. Figure 7 presents histograms of the frequency of P_2^* for high-price and moderate-price initial conditions respectively. Note the highly right-skewed distribution in Figure 8a (corresponding to a high initial price) and the fat-tailed, somewhat less skewed distribution in Figure 8b (corresponding to a moderate initial price). The skewness and kurtosis of the return distributions depend on the initial spot price. Spot returns exhibit greater skewness and smaller kurtosis when the initial spot price is high. For instance, when the spot price is high (as illustrated in Figure 7a), the coefficient of skewness of the return distribution is .5844 and the coefficient of excess kurtosis of returns is .2755. When the spot price is moderate (as illustrated in Figure 7b), the coefficient of skewness of returns is .5504 and the coefficient of excess kurtosis of returns is .4017. Consequently, the symmetric distributions implied by the Schwartz model underestimate the probability of upward price spikes and therefore underestimate the value of out-of-the-money calls. Biases become smaller as time-to-expiration increases (for at-the-money options) because the mean reversion inherent in the storage economy (and captured by the estimated parameters of the Schwartz model) becomes dominant as the pricing horizon increases. Volatility in both models shrinks monotonically in both models as the horizon increases.

In brief, these results show that neglecting the state-dependent price volatility leads to serious option pricing biases. The state-dependence in volatility documented in the analyses of the storage economy in section 3.3 leads to state-dependent pricing biases. Pricing biases are huge for short-dated, out-of-the-money calls when spot prices are high.

These features are reflected in volatility skews or smiles. Figure 8 depicts two examples of the resulting volatility skews for $\tau = 2$. It plots the volatility that sets the option price given by the Black model to the true option price. The x-axis in these figures measures strike ordinally; strike 5 is the at-the-money strike, strikes with numbers greater than 5 are out-of-the-money, and strikes with numbers less than 5 are in-the-money. Figure 8a depicts the skew when price is high (i.e., x is low and z is large) and Figure 8b depicts the skew when price is moderate (i.e., for moderate values of x and z). Note that there is an increasing relation between implied volatility when price is high, and a U-shaped relation when price is moderate. The minimum volatility point in the U-shaped case occurs two strikes out-the-money. The shape of these "smiles" corresponds to the pricing biases identified in Table 2.

In sum, standard commodity option pricing models fare poorly in the storage economy because they are severly mis-specified. The time and state dependence in volatility causes option prices to differ substantially from the values implied by a two factor model with constant volatilities with parameters calibrated to simulated data generated by the storage model. Although the two factor model permits more variation in the shape of the commodity price term structure, it does not capture the state dependent skewness and kurtosis in terminal price distributions. These pricing biases are especially severe for short-dated, high-strike call price options in high price environments. Computational considerations preclude evaluation of pricing biases for very long-dated options (including real investment options), but the decline of these biases in high and medium price environments as option maturities increase suggests that they may not be as severe for long dated options. Thus, standard constant volatility pricing models may be adequate for pricing long dated options and investment projects, but are highly questionable for pricing short dated commodity options, especially in high price environments.

4.3 Implications for Commodity Option Pricing Models

The foregoing analysis and the supporting empirical evidence in Ng-Pirrong (1994, 1996) and Pirrong (1996) imply that even multi-factor models of commodity price dynamics are mis-specified if spot price variances are assumed to be constant, as is typically the case. This implies that such models cannot price commodity options consistently. This raises the question: Is there an alternative reduced form model that can do better?

One alternative to a constant-spot variance model is to allow the variance to depend on the level of the spot price:

$$\frac{dP}{P} = \alpha(P)dt + \sigma(P)dW$$

A model of this type can be fit to option quotes using the method of Bodurtha-Jermakyan (1997).

Although this type of model should outperform constant volatility models in pricing options on the spot and options on futures (because of its ability to capture the relation between the level of prices and volatility), its one factor structure is unsatisfactory for pricing correlation dependent options, such as options on spreads (e.g., an option on the spread between a one month and a three month forward), commodity swaptions, or real options to undertake investments in commodity production capacity that can produce output for multiple periods. Multi-factor models are necessary to produce instantaneous spot-forward correlations that differ from one. Since the storage model implies that the spot-forward correlation is state dependent, at least two factors are required to generate a model that captures this empirical regularity.

Another alternative is a stochastic volatility model of the form studied by Heston (1993) with a positive correlation between the Brownian motion in the spot process and the Brownian motion in the volatility process. This form of model could capture the skewness and kurtosis of the spot price distribution. This model cannot generate a spot-forward correlation less than one, however, because the stochastic component of forward prices in this model does not depend upon the volatility.

The foregoing suggests that "off-the-shelf" diffusion models fail to capture salient features of commodity price dynamics. A two-factor model of the type studied by Schwartz (1997), but with the spot price-convenience yield variance-covariance matrix dependent upon the spot price may be necessary to capture fully the features documented herein. Given functions relating the spot price to (a) the spot price variance, (b) convenience yield variance, and (c) the spot price-convenience yield correlation, it is fairly straightforward to solve the "direct problem" of determining the price of a contingent claim with given boundary conditions using numerical methods. In theory, the methods of Bodurtha-Jermakyan can be applied to solution of the "inverse problem," i.e., the determination of the variance and correlation functions from a set of option quotes. Actual solution of this problem faces daunting numerical problems, however.[25]

In sum, the analysis demonstrates that conventional constant volatility option pricing techniques are likely to be inadequate for valuing commodity options due to the complex, non-linear relations between prices, variances, and spot-forward correlations. Some more advanced reduced form option pricing models may prove suitable for pricing non-correlation dependent options, such as options on forwards or options on the spot. The complex relation between prices and spot-forward correlations presents difficult challenges to any attempt to derive a multi-factor diffusion process suitable for pricing correlation-sensitive contingent claims such as commodity swaptions or real commodity investment options. If reduced form models are inadequate or im-

[25]Personal communication with Martin Jermakyan.

practical, it may prove necessary to value options using a structural model of the commodity market along the lines of the analyzed in sections 3.1-3.3. One difficulty is readily evident in this approach, however. Specifically, the crucial state variables z and η, and perhaps x are not observable. Instead they must be inferred. This requires joint econometric estimation of the state variables and crucial supply and demand paramters including ρ, σ_z, σ_η, a, b, ϕ, and γ. Given the inherent non-linearity of the model and the necessity of solving it numerically this is a forbidding task.

5 Summary and Conclusion

This article demonstrates that storage and the characteristics of production cost functions have important implications for the price dynamics of continuously produced, storable commodities. In particular, if marginal production costs are convex, increasing functions of output and there are both persistent and transitory demand shocks, commodity price volatility varies markedly with changes in underlying fundamental conditions. When market conditions are "tight" due to a lack of inventories or high demand, volatility can be orders of magnitude larger than when market conditions are less constrained. Similarly, spot-forward correlations are nearly one when supplies are abundant, but far below one when market conditions are tight.

These findings have important implications. First, they imply that standard commodity option pricing models based on mean reverting log normal spot commodity price processes are likely to generate severe mispricings. In the simulated storage economy these price biases are strike price and maturity dependent. Moreover, the size and sign of the biases depend on whether commodity prices are high or low. Standard option pricing models typically undervalue commodity options when the spot price is high and overvalue them when the commodity price is low. Second, these results imply that commodity option implied volatilities should be skewed, and that the shape of the skew should depend upon the level of prices (which measures the tightness of market conditions).

This analysis has numerous testable implications, some of which have already been verified (Ng and Pirrong, 1994, 1996). One empirical finding at variance from the model's predictions is the weaker GARCH effect in simulated storage model prices after controlling for scarcity (as measured by price) than is found in real world. This disparity may arise because the model assumes that the volatility of demand shocks is equal over time, whereas real world demand shocks may exhibit GARCH effects.

Some predictions of the model are observationally equivalent to those generated by the model of Zhou (1997). Zhou shows that incomplete hedging in a commodity market can cause volatility skews. In Zhou's framework, hedging considerations produce a systematic relation between market price levels and risk premia, which in turn creates a systematic relation between price levels and price volatility. The storage theory and the incomplete hedging model are completely different in spirit and motivation, so future research should try to distinguish between their empircal predictions. Some differences are apparent. Most important, the incomplete hedging approach implies that futures/forward prices should be mean reverting: the storage approach does not. It is not clear whether the incomplete hedging approach predicts that the shape of the volatility skew should depend upon the level of the commodity price as is the case in the storage model studied here. It is evident that the existing incomplete hedging model of Zhou does not predict the documented relation between commodity price levels (or spread levels) and the correlation between spot and forward returns because it contains a single factor. Perhaps multi-factor models of this type can generate such a pattern, but this has not been proven to date.

Future work in this area should include extension of the analysis to encompass periodically produced commodities (such as agricultural products).[26] Allowing investment in productive capacity is also worth exploring. Investment, like storage, is a way of smoothing out demand shocks. Integration of a real options investment model like that of Dixit and Pindyck (1995) with the richer economic environment of this article is an interesting, and certainly challenging, possibility. Generalizing the demand function to permit intertemporal substitutability is another possible extension of the model.

In sum, when marginal production costs for continuously produced commodities are increasing and convex, the volatilities are time varying. In particular, volatilities should be high when prices are high and spot-forward spreads wide. Moreover, options implied volatilities should be skewed, with the shape of the skew depending upon the existing scarcity in the market. Empirical data for continuously produced commodities exhibit many of these features. Consequently, models of commodity price dynamics which omit

[26] Pirrong (1998b) solves and analyzes a four-season recursive dynamic model of a market for a seasonally-produced commodity.

these features are mis-specified and are likely unreliable for pricing commodity options.

Table 1–Average Simulated GARCH Coefficients	
Coefficient	Value
γ_1	.2439
γ_2	.3059
γ_3	.00327
α_{11}	-.0412

Table 2
Constant Volatility Model Option Pricing Biases
Panel A–Low Price Environment

Strike Index	Periods to Expiration		
	1	2	3
1	-.0727	-.0976	-.1011
2	-.1443	-.1592	-.2024
3	-.2475	-.2358	-.2024
4	-.3812	-.3204	-.2552
5	-.5296	-.4042	-.3020
6	-.6753	-.4794	-.3377
7	-.7624	-.5434	-.3578
8	-.8315	-.5879	-.3601
9	-.8993	-.6182	-.3449

Table 2
Constant Volatility Model Option Pricing Biases
Panel B–Moderate Price Environment

Strike Index	Periods to Expiration		
	1	2	3
1	-.0523	-.0204	.0182
2	-.0691	-.0121	.0476
3	-.0835	.0116	.0961
4	-.0991	.0549	.1688
5	-.0788	.1230	.2730
6	-.0537	.2220	.4191
7	-.0207	.3693	.6199
8	.1214	.5779	.8972
9	.2194	.8593	1.2816

Table 2 Constant Volatility Model Option Pricing Biases Panel C–High Price Environment			
Strike Index	Periods to Expiration		
	1	2	3
1	.4670	.3437	.2574
2	.6837	.4868	.3600
3	.9711	.6684	.4968
4	1.3582	.9093	.6911
5	1.8876	1.2563	.9420
6	2.6233	1.7566	1.2757
7	3.6629	2.4391	1.7278
8	5.1567	3.3956	2.4010
9	7.3374	4.7560	3.3319

References

[1] Amin, Kaushik, Victor Ng, and Craig Pirrong. 1995. "Valuing Energy Derivatives," in R. Jamieson, ed., *Managing Energy Price Risk.* London: Risk Publications.

[2] Bodurtha, James, and Martin Jermakyan. 1997. "Non-Parametric Estimation of an Implied Volatility Surface." Working paper, Georgetown University.

[3] Brennan, Michael. 1991. "The Price of Convenience and the Valuation of Commodity Contingent Claims," in D. Lund and B. Oksendal, eds., *Stochastic Models and Option Values.* New York: North Holland.

[4] Brennan, Michael, and Eduardo Schwartz. 1985. "Evaluating Natural Resource Investments." 58 *Journal of Business*: 133-155.

[5] Cortazar, G., and Eduardo Schwartz. 1992. "The Valuation of Commodity Contingent Claims." Working paper, UCLA.

[6] Chambers, Marcus, and Roy Bailey. 1996. "A Theory of Commodity Price Fluctuations." 104 *Journal Of Political Economy*: 924-957.

[7] Deaton, Angus, and Guy Laroque. 1992. "On the Behavior of Commodity Prices." 59 *Review of Economic Studies*: 1-23.

[8] Deaton, Angus, and Guy Laroque. 1996. "Competitive Storage and Commodity Price Dynamics." 104 *Journal of Political Economy*: 896-923.

[9] Dixit, Avinash and Robert Pindyck. 1994. *Investment Under Uncertainty.* Princeton: Princeton University Press.

[10] Duffie, Darrell and Stephen Gray. 1995. "Volatility in Energy Prices," in R. Jamieson, ed., *Managing Energy Price Risk.* London: Risk Publications.

[11] Gibson, Rajna, and Eduardo Schwartz. 1990. "Stochastic Convenience Yield and the Pricing of Oil Contingent Claims." 45 *Journal of Finance*: 959-976.

[12] Grossman, Sanford, and Zhongquan Zhou. 1996. "Equilibrium Analysis of Portfolio Insurance." 51 *Journal of Finance*: 1379-1403.

[13] Heston, Steven. "A Closed-Form Solution of Options with a Stochastic Volatility, with Applications to Bond and Currency Options." *Review of Financial Studies*:

[14] Litzenberger, Robert, and Nir Rabinowitz. 1995. "Backwardation in Oil Futures Markets: Theory and Empirical Evidence." 50 *Journal of Finance*: 1517-1545.

[15] Ng, Victor, and Craig Pirrong. 1994. "Fundamentals and Volatility: Storage, Spreads, and the Dynamics of Metals Prices." 67 *Journal of Business*: 203-230.

[16] Ng, Victor, and Craig Pirrong. 1996. "Price Dynamics in Refined Petroleum Spot and Futures Markets." 2 *Journal of Empirical Finance*: 359-388.

[17] Pirrong, Craig. 1997. "Metalgesellschaft: A Prudent Hedger Ruined, or a Wildcatter on NYMEX?" 17 *Journal of Futures Markets*: 543-578.

[18] Pirrong, Craig. 1998a. "High Frequency Price Dynamics and Volatility Skews for Continuously Produced Commodities." Working paper, Washington University.

[19] Pirrong, Craig. 1998b. "High Frequency Price Dynamics and and Autocorrelations for Periodically Produced Commodities." Working paper, Washington University.

[20] Routledge, Bryan, Duane Seppi, and Chester Spatt. 1997. "Equilibrium Forward Curves for Commodities." Working paper, Carnegie Mellon University.

[21] Rubinstein, Mark. 1994. "Implied Binomial Trees." 49 *Journal of Finance*: 771-818.

[22] Scheinkman, Jose, and Jack Schectman. 1983. "A Simple Competitive Model of Production with Storage." 50 *Review of Economic Studies*: 427-441.

[23] Schwartz, Eduardo. 1997. "The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging." 52 *Journal of Finance*: 923-974.

[24] Teyssiere, Gilles, Christopher Gilbert, and Celso Brunetti. 1997. "Non-Ferrous Metals Price Volatility: A Fractionally Integrated Process?" in Robert Jameson, ed., *Managing Metals Price Risk.* London: Risk Publications.

[25] Williams, Jeffrey, and Brian Wright. 1991. *Storage and Commodity Prices.* Cambridge: Cambridge University Press.

[26] Working, Holbrook. 1948. "The Theory of the Price of Storage." 39 *American Economic Review*: 1254-1262.

[27] Zhou, Zhongquan. 1997. "Incomplete Hedging and Commodity Price Volatility." Working paper, Washington University.

Figure 1-Spot Price v. Demand State and Carry-in

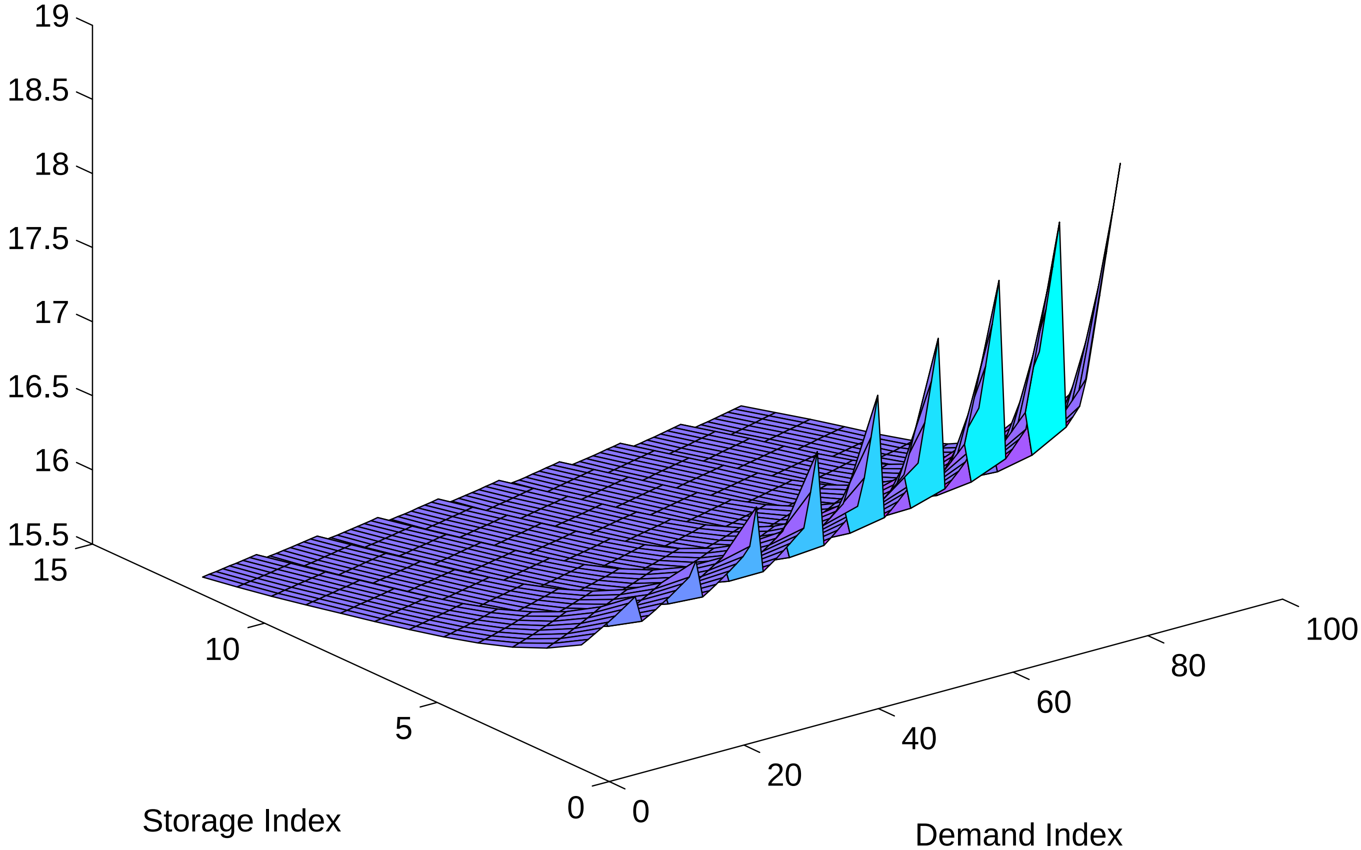

Figure 2-Spot-Forward Spread v. Demand State and Carry-in

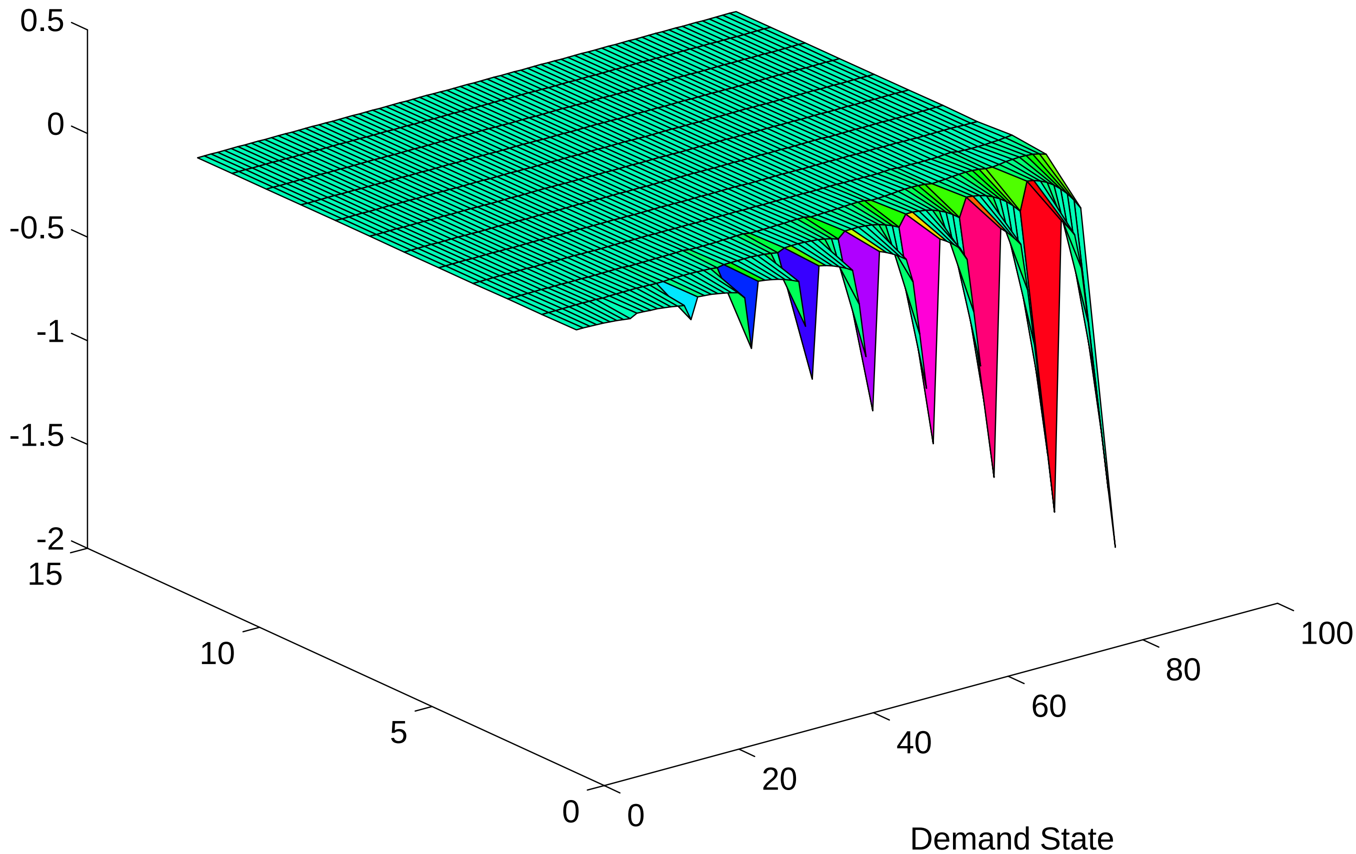

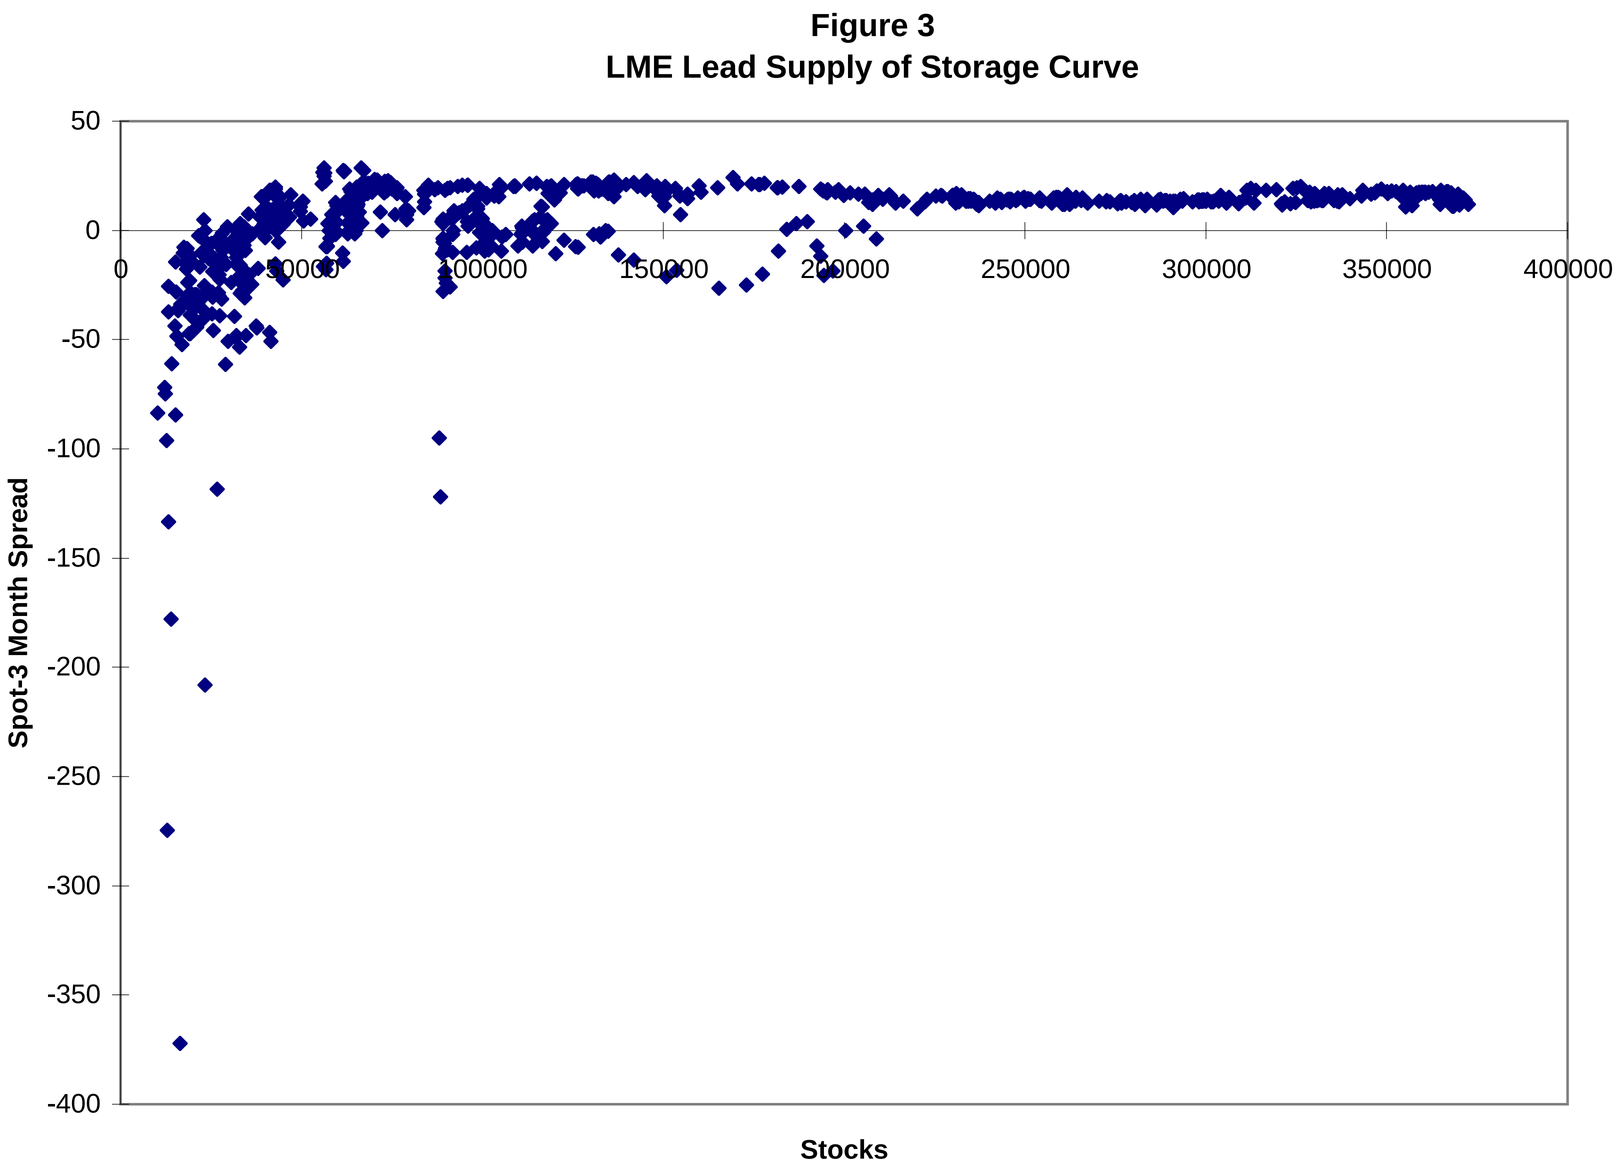
Figure 3
LME Lead Supply of Storage Curve
Spot-3 Month Spread
Stocks
50
0
-50
-100
-150
-200
-250
-300
-350
-400
0
50000
100000
150000
200000
250000
300000
350000
400000

Figure 4-Spot Return Variance v. Demand State and Carry-in

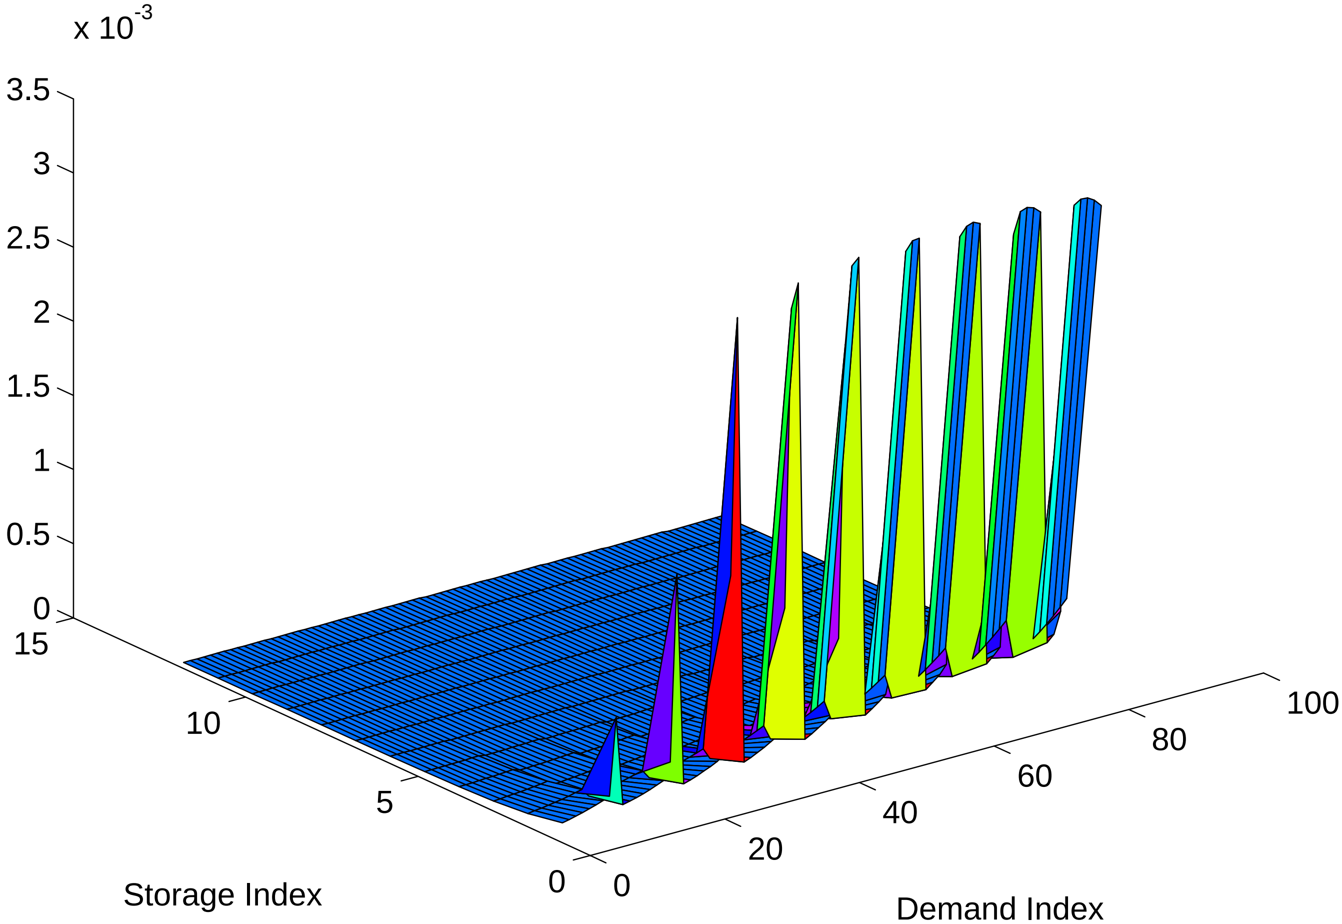

Figure 5-Spot-Forward Correlation v. Demand State and Carry-in

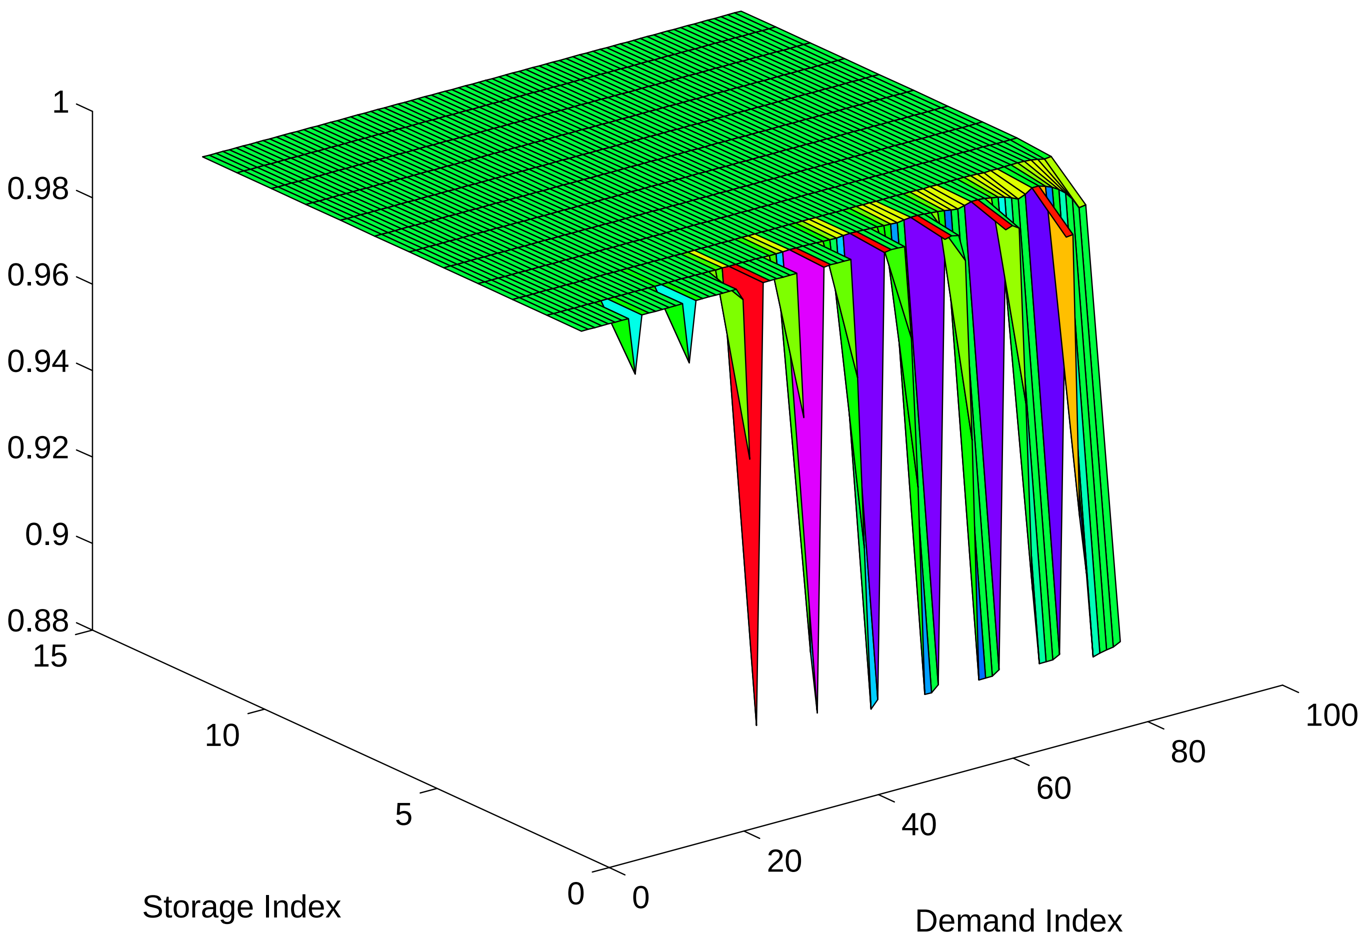

Figure 6-Hedge Ratios as a Function of Carry-in and Demand

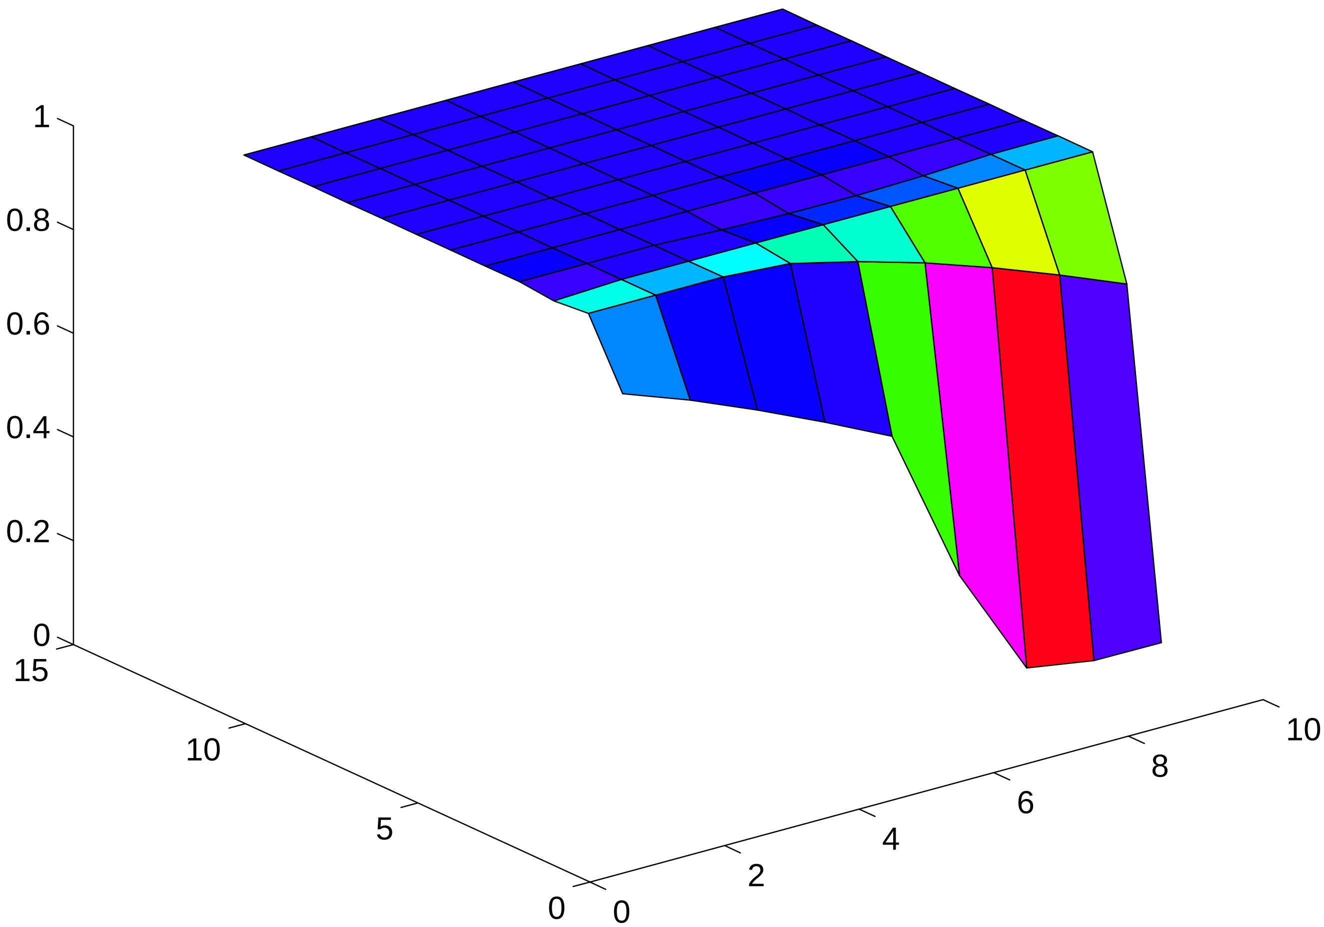

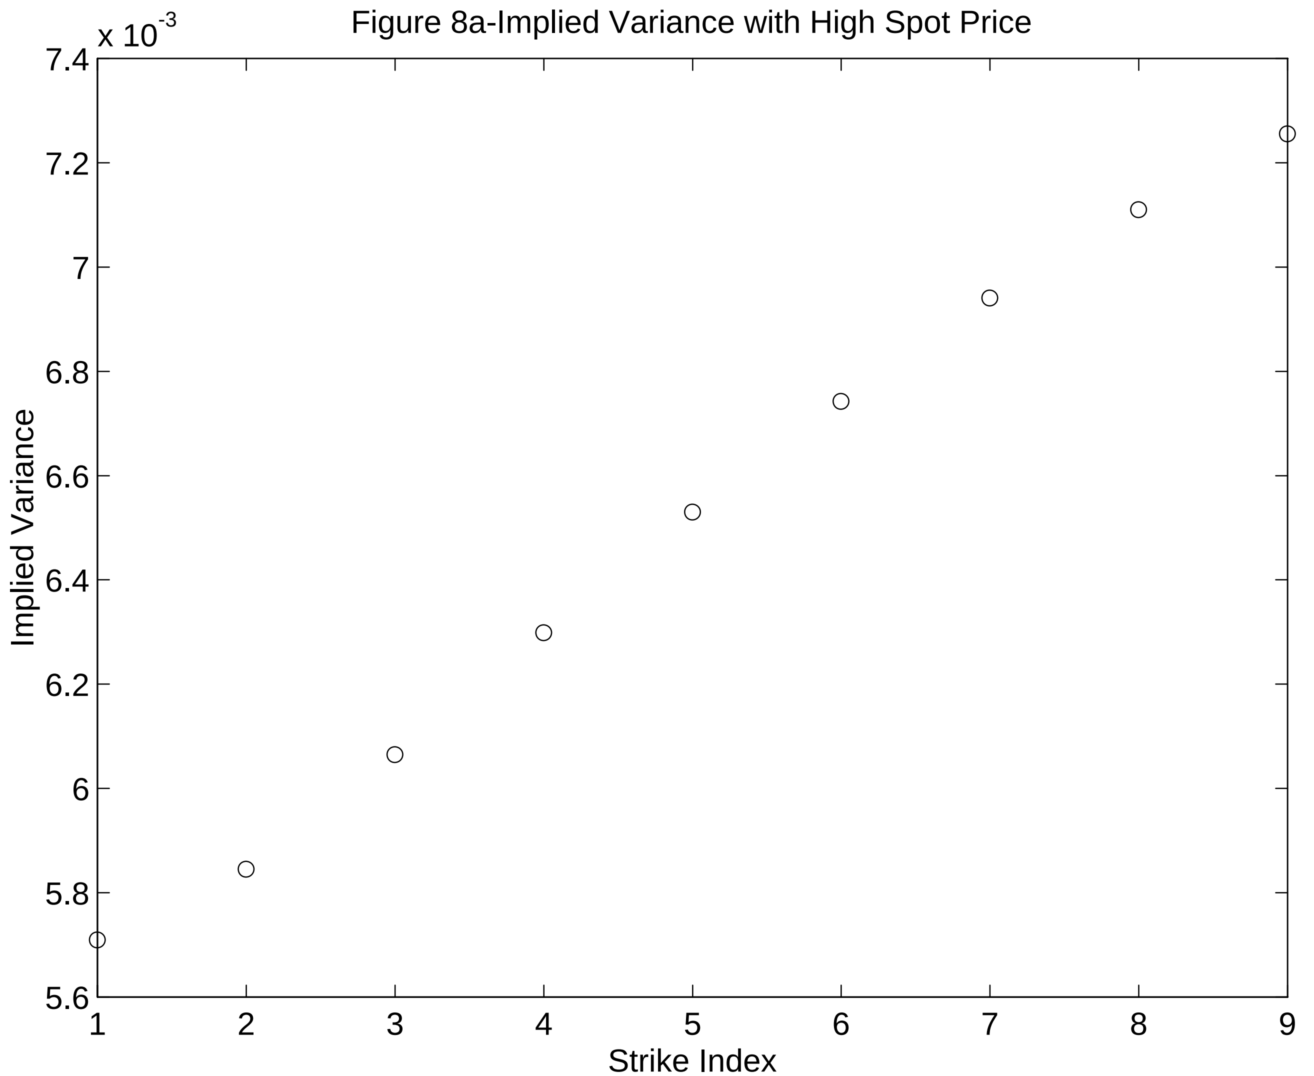

Figure 8a-Implied Variance with High Spot Price
x 10^-3
7.4
7.2
7
6.8
6.6
6.4
6.2
6
5.8
5.6
Implied Variance
1
2
3
4
5
6
7
8
9
Strike Index

Figure 8b-Implied Variance with Moderate Spot Price

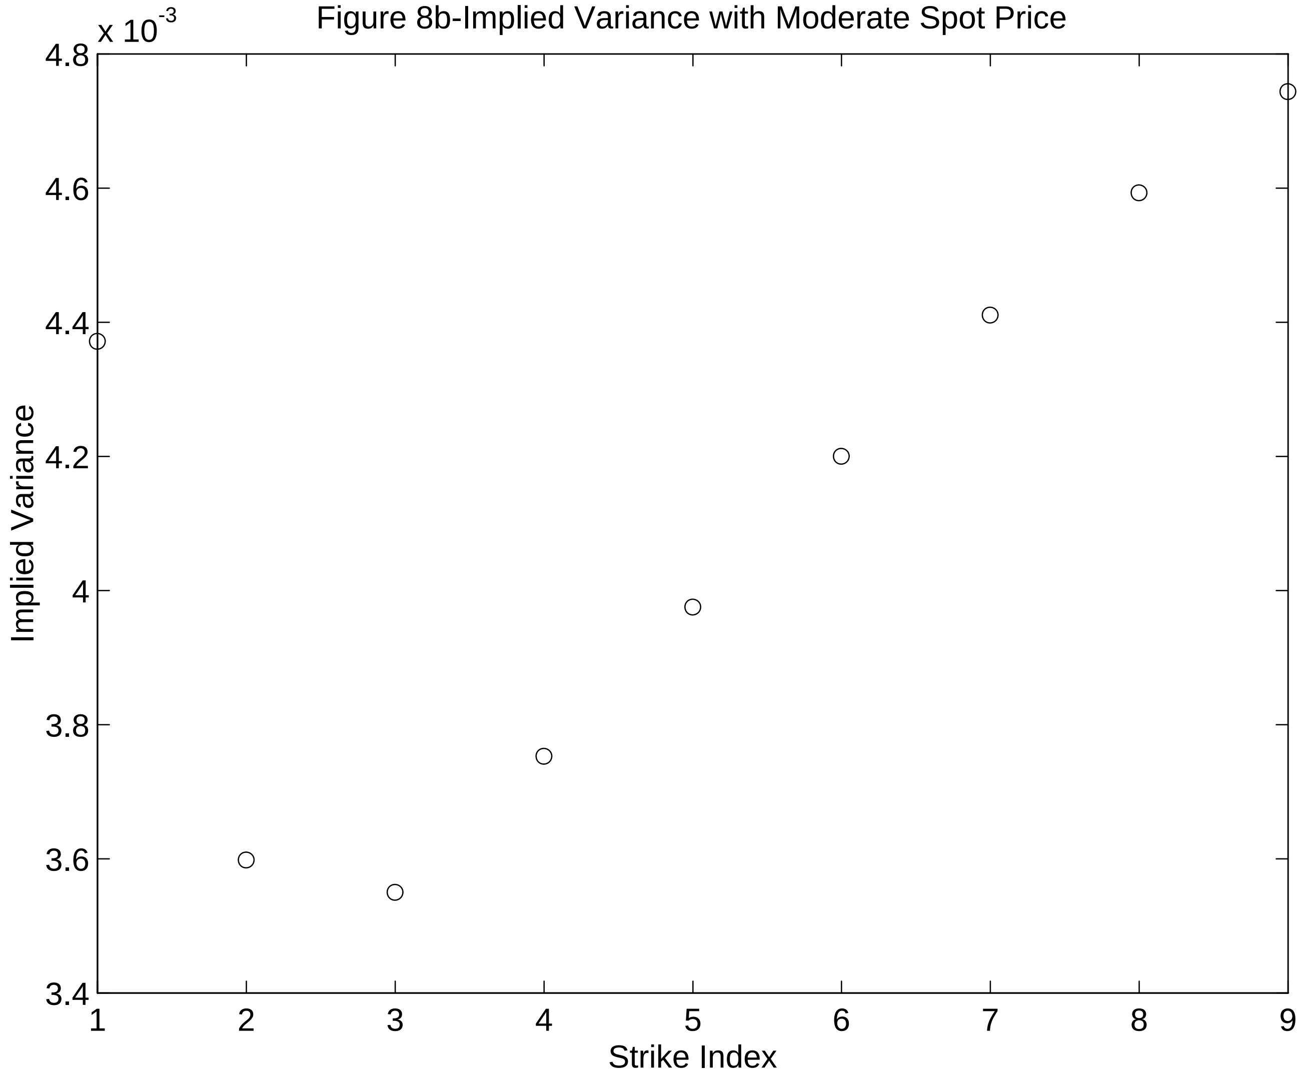

Figure 7a-Price Distribution with High Spot Price

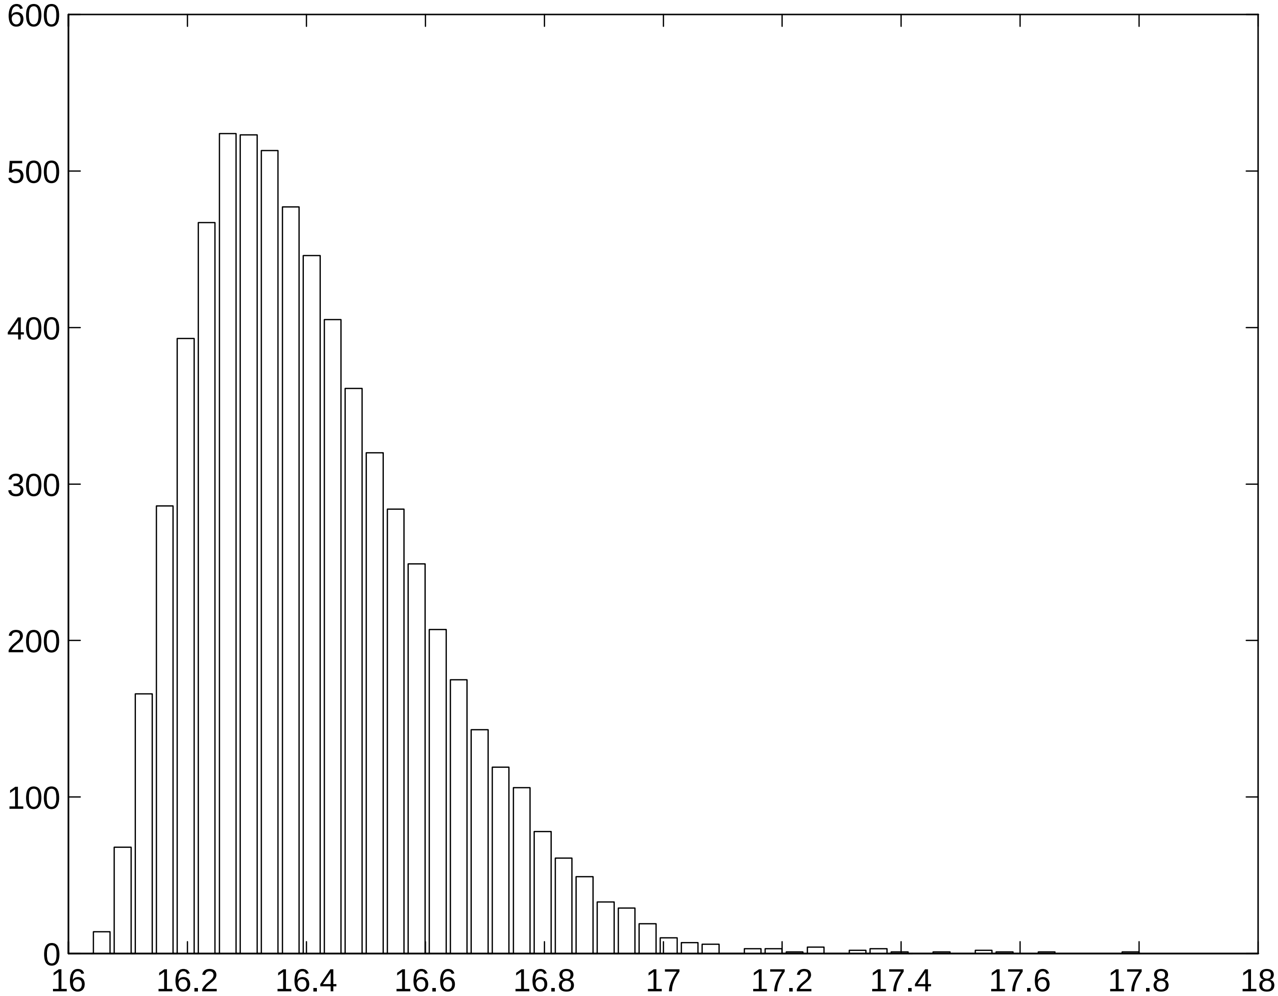

Figure 7b-Price Distribution with Moderate Spot Price