Victor K. Ng
Stephen Craig Pirrong
University of Michigan

Fundamentals and Volatility: Storage, Spreads, and the Dynamics of Metals Prices*

The theory of storage implies that inventory and demand conditions affect (*a*) the variances and correlations of commodity spot and forward prices and (*b*) the spread between spot and forward prices. For four industrial metals and one precious metal over the 1986–92 period, the observed relations between the spread and the variances and correlations of spot and forward prices are consistent with the theory. Since the close connection between spreads and real supply-and-demand conditions is well documented, the results strongly suggest that fundamental factors determine the dynamics of metals prices.

I. Introduction

What determines the volatility of asset prices? Most economists have traditionally argued that fundamental factors determine volatility. Others assert, however, that asset prices are driven by "animal spirits" and other random forces which induce excess volatility.[1]

In this article we investigate whether fundamentals determine the return variances and correlations for an important group of commodities—the metals. To do so, we exploit the implications of the theory of storage. This theory implies that fundamental supply-and-demand conditions determine the spread between spot and forward prices. Since the spread is observable on a daily basis, it is possible to test whether the spread and the dynamics of metals prices are related in the way the theory predicts. A close correspondence between the predicted and observed behavior of the spreads, variances, and correlations of metals returns is consistent with the hypothesis that fundamentals determine price dynamics.

* The suggestions of an anonymous referee contributed to the clarity of the analysis and exposition. We are grateful to Kaushik Amin, James Bodurtha, and the participants at the National Bureau of Economic Research Summer Institute Workshop on Asset Pricing for their useful comments.

1. See especially Shiller (1989) for a statement of the view that asset prices are excessively volatile.

(*Journal of Business*, 1994, vol. 67, no. 2)

0021-9398/94/6702-0002$01.50

The theory of storage implies that the interest-and-storage-adjusted spread (the "adjusted spread" hereafter) widens as (*a*) inventories fall relative to demand and (*b*) the industry marginal cost of production function becomes less elastic.[2] These conditions make the supply curve less elastic. Since microeconomic theory predicts that commodity prices become more variable as the supply curve becomes less elastic, the volatility of spot and forward prices should increase as the adjusted spread widens.

The storage theory makes other predictions about the dynamics of commodity prices. Specifically, supply is more elastic in the long run than in the short run. Therefore, the spot return volatility should exceed the forward-return volatility when the adjusted spread is wide. Conversely, when supplies are abundant, the adjusted spread is nearly zero, and spot and forward volatilities should both be small and nearly equal. In addition, as inventories decline and the spread widens, a stock out becomes more likely. Since a stock out breaks the arbitrage link between spot and forward prices, the correlation between these prices should decline as the adjusted spread widens.

To test these implications, we employ a bivariate dynamic model that allows past spreads to affect (*a*) the volatility of spot and forward returns and (*b*) the correlation between them. We first estimate this model by using data on spot and forward prices for industrial metals traded on the London Metal Exchange (LME) over the 1986–92 period. The metals studied are aluminum, copper, lead, and zinc. After estimating the model parameters, we determine how the spread is related to the ratio between spot- and forward-return volatilities, the forward price elasticity, the variance of the spread, and hedge ratios.

The results for the industrial metals are consistent with the predictions of the theory. The lagged-squared-adjusted spread has a statistically significant effect on the variances of both spot and forward returns and on the correlation between these returns. Indeed, spreads are highly correlated with spot- and forward-return volatilities. Moreover, variations in lagged-squared spreads explain between 50% and 70% of the innovations in industrial metal spot-return variances and between 50% and 62% of the innovations in industrial metal forward-return volatilities. The other attributes of industrial metal price dynamics also behave in accordance with the theory. This decisive empirical confirmation of the implications of the theory of storage strongly supports the view that fundamentals drive the price dynamics of industrial metals.

For contrast, we also estimate the model for a precious metal, silver.

2. For analyses of the theory of storage and evidence relating to it, see Working (1948, 1949), Telser (1958), Bresnahan and Suslow (1985), Bresnahan and Spiller (1986), Williams (1986), and Williams and Wright (1989, 1991).

This metal is widely held as a store of value, rather than for industrial use. Therefore, the theory predicts that the adjusted spread should be very close to zero and that it should not vary much. This was indeed the case in the 1986–92 period. Moreover, the spread explains little of the variation in silver volatility over time.

The findings for silver sharpen the interpretation of the very strong results for industrial metals. Since the theory of storage predicts marked differences between the behavior of silver prices and industrial metal prices, our documentation of such differences further bolsters the conclusion that fundamentals determine the behavior of metals prices.

This work makes several substantive contributions to the literature. The most important of these is the demonstration of the strong link between fundamentals and industrial metal price volatility. Since the adjusted spread is a parsimonious summary of supply-and-demand conditions, and since we observe it on a daily basis, our tests of the relation between fundamentals and commodity price volatility are far more powerful than those of Anderson (1985) and Kenyon et al. (1987). These studies rely on crude measures of supply-and-demand conditions (e.g., seasonals), observe volatility at low frequency (monthly), and employ less powerful statistical techniques. Moreover, our maximum likelihood estimators of return variances are more efficient than the absolute-value-of-the-return measure of daily variance employed by Roll (1984) in his study of the variability of orange juice futures prices. As a result, we estimate the relation between fundamentals and volatility with greater precision. Finally, our findings imply that previous analyses that use an ad hoc bivariate generalized autoregressive conditional heteroscedasticity (GARCH) framework to model return variances (e.g., Baillie and Myers 1991) but which ignore the effect of the spread are misspecified and inconsistent with the theory of storage.

This article also presents new tests of the theory of storage. Although we obviously build on the seminal article of Fama and French (1988), our methods offer several advantages. Specifically, we estimate the effect of the spread on spot- and forward-return variances and on the correlation between spot and futures returns. Fama and French examine only the variance of the basis (which is a function of these three parameters) and the ratio of spot- and forward-return volatilities. Moreover, our methodology allows us to determine the marginal effect of the spread on variances, correlations, and elasticities. Fama and French cannot do so because they compare statistics from subsamples in which the spread is positive to those in which the spread is negative. The ability to estimate volatility at each point in the sample also permits us to execute more detailed tests of Samuelson's (1965) hypothesis that spot prices should vary more than forward prices. Finally, if our distributional assumptions and model specification are correct, our maximum likelihood approach provides more efficient estimates of the relevant parameters than the ordinary least squares (OLS) estimates of Fama and French. (We have estimated the model by using alternative distributions and specifications and have found that the empirical results are robust to these changes.)

The remainder of this article is organized as follows. The next section reviews the theory of storage; outlines the predicted relations between inventories, spreads, volatility, and correlations; and describes our estimation and testing approach. Section III describes the data used to test the hypotheses and then reports and interprets the results. Section IV provides a brief summary of the work.

II. Fundamentals and Price Dynamics: Inventory, the Spread, Volatility, and Correlation

A. The Theory of Storage

There are two versions of the theory of storage. Kaldor (1939) proposed the first, and better-known, version. Working (1948, 1949), Telser (1958), Williams (1986), and Brennan (1991) have elaborated on it. This theory asserts that processors and consumers of a commodity receive a stream of implicit benefits when they hold inventories of the good. This benefit is called the "convenience yield." Firms earn the convenience yield because stocks-on-hand allow them to respond more flexibly and efficiently to unexpected supply-and-demand shocks. The theory posits that the marginal value of convenience declines as inventory increases. Empirical evidence produced by Working (1948, 1949), Telser (1958), and Brennan (1991) is consistent with this prediction. This evidence also suggests that the convenience yield is a convex function of stocks. The top panel of figure 1 depicts the convenience-yield function.

Arbitrage ensures that the convenience yield affects the relation between spot and forward prices. Since holders of stocks earn the convenience yield but owners of forward contracts do not, a positive convenience yield depresses the forward price relative to the spot price. This effect of the convenience yield on the futures price is analogous to the effect of a dividend yield on the price of a stock-index futures contract. Formally, let F_t be the forward (or futures) price at time for delivery of a commodity at time $T > t$, and call S_t the spot price of the commodity at t. Moreover, let $w_{t,T}$ equal the cost of physically storing a unit of the commodity from t to T, and define $r_{t,T}$ as the yield at t on a discount bond that matures at T. Finally, assume that the convenience yield equals $c_{t,T}$. Then, it is well known that the no-arbitrage relation between the spot and forward prices is

$$F_t - w_{t,T} = S_t e^{(r_{t,T} - c_{t,T})(T-t)}. \tag{1}$$

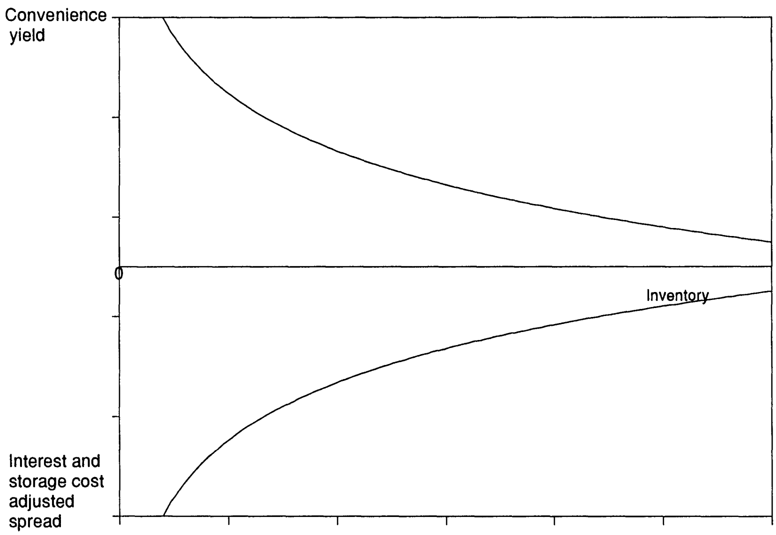

FIG. 1.—The relation between inventories, the convenience yield, and the interest and storage cost adjusted spread between forward and spot prices.

Thus, a rise in the convenience yield causes the forward price to decline relative to the spot price. It is also evident that $c_{t,T} \geq 0$. If not, speculators could earn an arbitrage profit by shorting a forward contract, buying the commodity, and storing it until expiration of the forward contract.

The relation between spot and forward prices is frequently expressed in terms of the interest and storage adjusted spread. This is defined as

$$z_t \equiv \frac{\ln(F_t - w_{t,T}) - \ln S_t}{T - t} - r_{t,T} = -c_{t,T} \leq 0. \quad (2)$$

The adjusted spread equals the annualized percentage difference between the forward and spot prices at t, net of storage and interest costs incurred to hold inventory from t to T. Given the behavior of the convenience yield, z_t varies directly with inventories. That is, z_t becomes more negative as inventories decline. Moreover, since $c_{t,T}$ is a convex function of stocks, z_t is concave. The bottom panel of figure 1 illustrates this relation.

The second version of the theory of storage does not rely on the construct of a convenience yield but produces identical implications for the relation between z_t and inventories. The articles of Bresnahan and Spiller (1986), Williams and Wright (1989, 1991), and Deaton and Laroque (1991) imply that z_t is an increasing, concave function of inventory even when processors and marketers receive no implicit benefit from holding inventories. This is true because the probability of stock out prior to the expiration of the forward contract varies inversely with inventories. Spot prices exceed forward prices (adjusting for carrying costs) when a stock out occurs, because under these circumstances it is impossible to undertake intertemporal arbitrage transactions; arbitrageurs cannot augment current consumption with future production. Instead, the spot price rises as high as is necessary to equilibrate supply and demand in the spot market. Thus, when stocks decrease, the probability of a stock out increases, and z_t declines.[3] This alternative theory of storage therefore implies that intertemporal consumption optimization and a nonnegativity constraint on storage, rather than an explicit convenience yield, explain the relation between spot-forward spreads and inventories.

Which of the two versions of the theory of storage is more plausible is irrelevant to the present study, because the models are observationally equivalent for our purposes. The conclusions that we discuss below hold as long as there is an increasing, concave relation between z_t and stocks, regardless of the structural relation that produces this reduced form.

B. *Testable Implications of the Theory of Storage*

French (1986), Fama and French (1987, 1988), and Williams and Wright (1991) derive the implications of a concave, increasing relation be-

3. Keynes (1930) first noted the linkage between negative spreads ("backwardation") and stock outs. It may be argued that, for commodities such as metals (or grains or petroleum products), a stock out is an unlikely event, and this version of the theory is therefore unrealistic. It is indeed the case that it is extremely unlikely that *worldwide* (or even U.S.) stocks of a commodity would ever be exhausted. However, Williams and Wright (1989, 1991) demonstrate that the spot-forward spread is determined, not by worldwide or national stocks, but by the inventories at the delivery point for the spot and forward contracts. It is quite possible for stocks to be consumed completely at a single point (such as the delivery point) even when there are inventories at other locations. Moreover, Bresnahan and Spiller (1986) prove that stock outs at a particular location must occur with positive probability in a no-bubbles economy. This is an intuitive result. If a stock out at a particular point were a zero-probability event, then some of the inventory would never be consumed. This is wasteful. Thus, optimality requires that stock outs occur with positive probability. There is some empirical evidence that is consistent with this theory. Using data from the Australian wheat trade, Brennan, Wright, and Williams (1992) provide an empirical example of a situation where (*a*) the adjusted spread is negative in a central market, (*b*) stock outs occur in the central market, and (*c*) inventories outside the central market are always positive. This situation occurs because the costs of transporting more wheat from outlying areas to the central market immediately are higher than the expected costs of shipping it in the future. This makes it economic to delay shipping the wheat, even though the present value of the price for deferred delivery is lower than the price for immediate delivery. As a result, no wheat is stored in the central market despite the backwardation. Deaton and Laroque (1991) demonstrate that certain attributes of commodity prices behave as the stock-out model predicts. Specifically, actual and stimulated prices exhibit autocorrelation with occasional price spikes. It is also worth noting that rigorous formal models of the stock-out version of the theory have been derived. In contrast, derivations of the more familiar convenience-yield theory are primarily intuitive and informal.

tween adjusted spreads and inventories for spot- and forward-return volatilities, the ratio of these volatilities, and the correlation between spot and forward returns. There are three primary implications; three other implications follow immediately.

IMPLICATION 1. If current and permanent shocks predominate,[4] the variances of spot and forward returns vary inversely with z_{t-1}.[5] That is, spot and forward returns are more volatile, the wider the adjusted spread.

The reasoning behind this result is as follows. A decline in stocks has two effects. First, it causes z_{t-1} to decrease. Second, supply conditions become more constrained. This reduces the elasticity of supply. Since, for a given distribution of demand shocks, prices are more variable when supply is less elastic, this second effect causes spot and forward prices to become more volatile. Thus, there is a negative relation between z_{t-1} and the variance of spot and forward returns.[6] Figure 2 illustrates this prediction.

This implication speaks directly to the issue of the importance of fundamentals in determining spot- and forward-return variances. Since z_{t-1} summarizes real supply-and-demand conditions, an empirical demonstration that movements of this variable explain a large fraction of the movements in return variances is consistent with the hypothesis that fundamentals primarily determine volatility.

IMPLICATION 2. If current and permanent supply-and-demand shocks predominate, then (*a*) when $z_{t-1} = 0$, the variance of a commodity's spot return equals the variance of its forward return, and (*b*) as z_{t-1} decreases, the variance of the spot return increases relative to the variance of the forward return. Thus, the ratio of the forward-return variance to the spot-return variance decreases as the spread widens.[7]

To see the basis for this result, consider the effect of a positive current or permanent demand shock. Such a shock causes the spot

4. Current shocks affect demand and supply at t only; permanent ones affect demand at both t and T in the same fashion.

5. Note that (*a*) $\Delta \ln S_t = \ln S_t - \ln S_{t-1}$ and (*b*) z_{t-1} are related to inventory conditions immediately prior to the shock that generates the return at t, whereas z_t includes the effects of the shock. Thus, variations in z_{t-1} measure variations in initial supply-and-demand conditions. Put another way, z_{t-1} is in the relevant information set, whereas z_t is not.

6. It should also be noted that a decrease in the elasticity of the marginal cost of production curve similarly affects the elasticity of the supply curve and therefore leads to an increase in price volatility. Williams and Wright (1991, pp. 131–35) demonstrate, moreover, that such a change in the elasticity of the marginal cost of production causes the spread to widen. Ceteris paribus, an increase in demand has the same effect. In order to simplify the exposition, we focus our attention on how changes in inventories affect the spread and volatility; it is a straightforward exercise to generalize our implications to encompass cost elasticity and demand changes.

7. Samuelson's (1965) theory also predicts that spot volatility should exceed forward volatility.

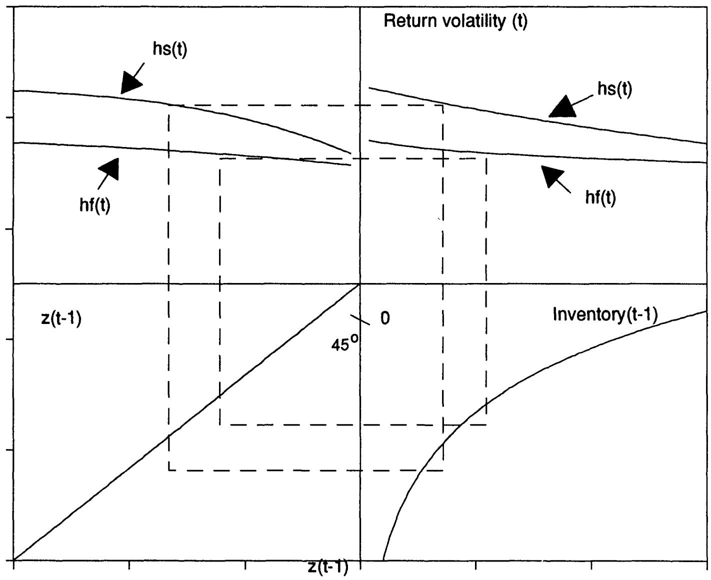

FIG. 2.—The relation between spot- and forward-return volatilities and the interest and storage cost adjusted spread between forward and spot prices. The axis labeled "*hs*(*t*)" measures the spot-return conditional volatility at time *t* (given all information available at time $t - 1$). The axis labeled "*hf*(*t*)" measures the forward-return volatility at time *t*. The axis labeled "$z(t - 1)$" measures the interest and storage cost adjusted spread at time $t - 1$. The axis labeled "*Inventory*($t - 1$)" measures the level of inventory at $t - 1$.

price to rise and inventories to fall. Because of the convexity of the convenience-yield function, the decline in inventories has virtually no effect on the convenience yield when stocks are large. Thus, equation (1) implies that the spot and forward prices move by nearly equal amounts in response to the shock.

Now consider the effect of the demand shock when stocks are small. In this case, the shock again causes the spot price to rise and inventories to decline. Since the convenience yield is a convex function of inventories, a decline in inventories causes this yield to increase substantially when stocks are small. In these circumstances, equation (1) implies that the forward price moves less than the spot price in response to the demand shock. Thus, the spot return is more volatile than the forward return when stocks are small, whereas these returns are equally volatile when inventories are large. Figure 2 depicts this relation.

It is also possible to derive this result from the stock-out version of the theory. When stocks are low, spot prices must change dramatically in response to supply-and-demand shocks because inventories and production cannot adjust immediately to accommodate them. Over a long time horizon, however, agents can adjust real variables in response to a shock. Therefore, a smaller forward-price response equilibrates the market for future delivery. Thus, when the adjusted spread is wide, the spot volatility exceeds the forward volatility. Conversely, when the adjusted spread is close to zero, a stock out is a very remote possibility. As a result, speculators will almost always be able to execute cash-and-carry arbitrage transactions which ensure that the forward price exceeds the spot price by the cost-of-carry. This guarantees that when the spread is nearly zero, the spot and forward prices move in equal amounts when a shock occurs.[8]

IMPLICATION 3. The correlation between spot and forward returns equals one when $z_{t-1} = 0$. As z_{t-1} decreases from zero, this correlation declines.

This implication is most easily derived from the stock-out version of the theory of storage. As just noted, when stocks are large and thus $z_{t-1} = 0$, spot and forward prices move in lockstep in response to supply-and-demand shocks. In this case, spot and forward returns are nearly perfectly correlated.

As stocks decline, the probability of a stock out increases. If some shocks are temporary, spot prices can move independently of forward prices during a stock out. This occurs because agents cannot undertake the cash-and-carry arbitrage transactions that normally link them. Therefore, when the spread is wide, the spot and forward prices are imperfectly correlated because there is an appreciable probability that these prices will move independently in the near future. As a result, the spot-forward return correlation should decline systematically as the spread widens. Figure 3 illustrates this relation.

These primary predictions of the theory of storage lead directly to three further implications.

IMPLICATION 4. Define the change in the log difference between forward and spot prices as $V_t \equiv \Delta(\ln F_t - \ln S_t)$. The variance of V_t equals zero when $z_{t-1} = 0$. Moreover, the volatility of V_t is decreasing in z_{t-1}.

When the spread is zero, the variances of spot and forward returns are equal, and these returns are perfectly correlated. This implies that the variance of V_t equals zero when $z_{t-1} = 0$. When the spread widens, the variances of spot and forward returns increase, the spot-return variance increases more than the forward-return variance, and the correlation between the returns declines. Together, these effects imply that the variance of the forward-spot relative price increases as z_{t-1} declines. Figure 4 illustrates this pattern.

8. This long-run adjustment process should also cause the spread to be mean reverting.

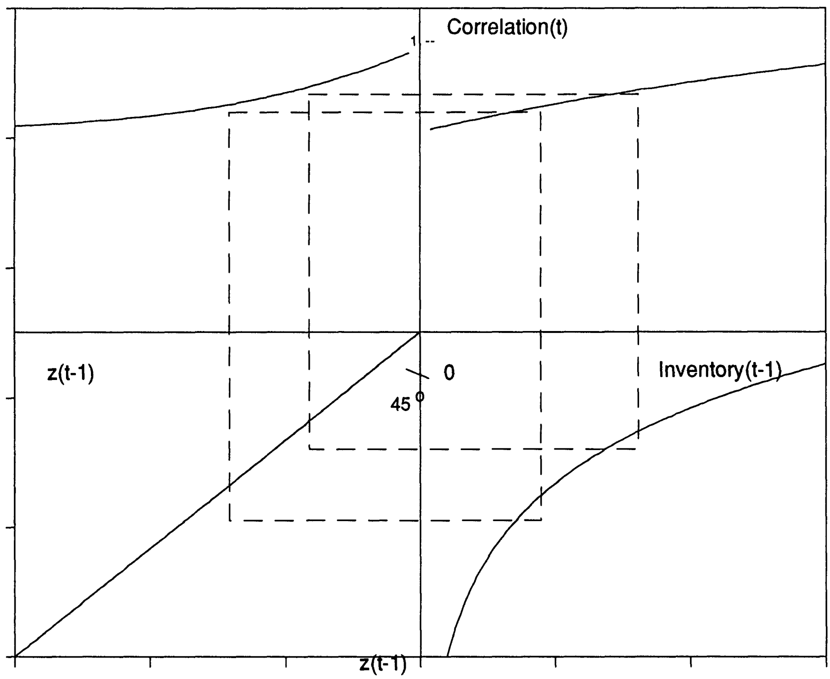

FIG. 3.—The relation between the spot-forward return correlation and the interest and storage cost adjusted spread between the forward and spot prices. The axis labeled "*Correlation*(t)" measures the conditional correlation between forward and spot returns at time t (given all information available at time $t - 1$). The axis labeled "z(t − 1)" measures the interest and storage cost adjusted spread at time $t - 1$. The axis labeled "*Inventory*(t − 1)" measures the level of inventory at $t - 1$.

IMPLICATION 5. Define the elasticity of the forward price with respect to the spot price as $e_t = \Delta \ln F_t / \Delta \ln S_t$. If the forward price equals the expected spot price,[9] then $e_t = 1$ when $z_{t-1} = 0$. This elasticity is increasing in z_{t-1}. That is, the elasticity falls below one as the spread widens.

French (1986) proves this result formally. The intuition behind it is similar to that underlying implication 2. When stocks are large, the spot and forward prices move by nearly the same amount in response to a demand or supply shock. When stocks are small, these prices tend to move in the same direction, but the spot price moves more than the forward price in response to a given permanent or current shock.

9. This is strictly true only if all agents are risk neutral, or under the equivalent martingale measure.

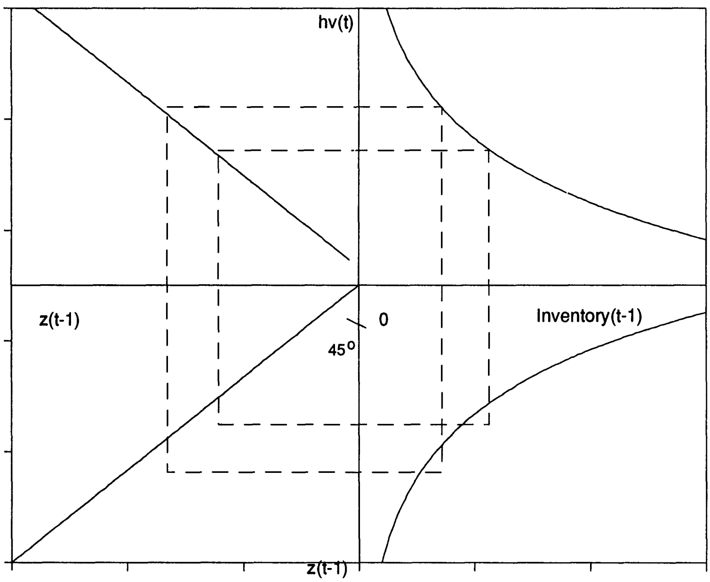

FIG. 4.—The relation between the variance of the change in the basis and the interest and storage cost adjusted spread between the forward and spot prices. The axis labeled "$hv(t)$" measures the conditional volatility of the change in the adjusted basis at time t. The axis labeled "$z(t-1)$" measures the interest and storage cost adjusted spread at time $t-1$. The axis labeled "*Inventory*$(t-1)$" measures the level of inventory at $t-1$.

IMPLICATION 6. The spot-forward hedge ratio should vary with the adjusted spread.

The number of forward contracts to short to hedge a spot position depends on the variances and covariance of the spot and forward returns. Since these depend on the adjusted spread, hedge ratios should depend on it as well.

To summarize, the theory of storage implies that fundamental supply-and-demand conditions should systematically affect the dynamics of commodity prices. Moreover, since it is extremely well documented empirically that these fundamental factors determine the spot-forward spread, this variable is an excellent proxy for these factors.[10] Empirical confirmation of these implications concerning the relation between the adjusted spread and the full array of commodity price dynamics would thus provide strong evidence that fundamental conditions are the primary determinants of these dynamics.

10. Examples of empirical tests of the relation between inventories and spreads include Working (1948, 1949), Howell (1956), Telser (1958), Weymar (1974), Gray and Peck (1981), Thompson (1986), Williams (1986), Pindyck (1990), and Brennan (1991).

C. Model Specification

We test these six predictions empirically by using an error-correction model with time-varying means, variances, and covariances. The conditional means of the spot and futures returns are specified as

$$\Delta \ln S_t = \alpha_s + \sum_{i=1}^{5} \beta_{i,s} \Delta \ln S_{t-i} + \sum_{i=1}^{5} \gamma_{i,s} \Delta \ln F_{t-i} + \mu_s z_{t-1} + \epsilon_t \quad (3)$$

and

$$\Delta \ln F_t = \alpha_f + \sum_{i=1}^{5} \beta_{i,f} \Delta \ln S_{t-i} + \sum_{i=1}^{5} \gamma_{i,f} \Delta \ln F_{t-i} + \mu_f z_{t-1} + \eta_t, \quad (4)$$

where the α's, β's, γ's, and μ's are parameters and ϵ_t and η_t are random-error terms. The inclusion of the lagged-adjusted spread terms reflects the fact that spot and forward prices should obey a long-term equilibrium relationship. Therefore, when spot and forward prices diverge widely, agents should reduce current consumption and increase production in order to drive them back together again. That is, the spot and forward prices should be cointegrated.

The conditional variances and covariances of the spot and futures returns are specified as an augmented bivariate GARCH model. Formally, the equations for the conditional variance of the spot return, $h_{s,t}$, and the conditional variance of the forward return, $h_{f,t}$, are as follows:

$$h_{s,t} = \omega_s + \delta_1 h_{s,t-1} + \delta_2 \epsilon_{t-1}^2 + \delta_3 z_{t-1}^2, \quad (5)$$

and

$$h_{f,t} = \omega_f + \phi_1 h_{f,t-1} + \phi_2 \eta_{t-1}^2 + \phi_3 z_{t-1}^2. \quad (6)$$

The inclusion of the z_{t-1}^2 terms allows the adjusted spread to affect volatility. We use z_{t-1}^2 instead of z_{t-1} or $|z_{t-1}|$ in the empirical work because it produces uniformly superior results. The log likelihood is always largest for the squared spread specification, although the qualitative results are very similar regardless of which spread variable is employed. This contrasting performance is of some interest, as it implies that widening the spread increases volatility at an increasing rate.

The theory discussed in Section II*B* implies $\delta_3 > \phi_3 > 0$. The first inequality reflects the prediction that current supply conditions have a more pronounced effect on spot volatility than forward volatility. The second inequality formalizes the prediction that prices are more volatile when inventories are low and the spread is wide, and hence when z_{t-1}^2 is large. The lagged variance and lagged-squared-error terms reflect the possibility that other factors may introduce time-varying elements into volatility. For example, the theories of Kyle (1985), Black (1986), and Ross (1989) imply that informed trading increases volatility and that informed individuals exploit their advantage by spreading out their trading over time. If information flows vary over time, therefore, volatility varies as well, and volatility shocks should persist. Alternatively, erratic speculative trading can also induce time-varying volatility.

The conditional covariance of the spot and forward returns is specified by using a variant of the model of Kroner and Sultan (1991). In order to capture the relation between spreads and the covariance, we assume the following functional form:

$$\sigma_{s,f,t} = \rho\sqrt{h_{s,t}h_{f,t}} + \theta z_{t-1}^2, \tag{7}$$

where $\sigma_{s,f,t}$ is the covariance between spot and forward returns. The parameter ρ is the correlation between spot and forward returns when $z_{t-1} = 0$. Theory predicts $\rho \approx 1$ because spot and forward returns are nearly perfectly correlated when the market is at full carry (i.e., when $z_{t-1} = 0$). The parameter θ captures the effect of the lagged-adjusted spread on the covariance. Since the correlation between spot and forward returns should decline as the lagged-adjusted spread widens, the theory predicts $\theta < 0$.

This model of the covariance captures important relations between spot and forward price correlations and the spread. These relations are predicted by theory and have been ignored heretofore. For instance, the bivariate GARCH specification employed in the work of Baillie and Myers (1991) is misspecified if spot-forward correlations depend on supply conditions, and hence on the spread.

We estimate equations (3) and (4) by using OLS to obtain ϵ_t and η_t. We then estimate equations (5)–(7) as a system by using the method of maximum likelihood, when it is assumed that

$$\begin{pmatrix}\epsilon_t\\ \eta_t\end{pmatrix} | I_{t-1} \sim \text{Student-}t\left(\begin{pmatrix}0\\0\end{pmatrix}\begin{pmatrix}h_{s,t} & \sigma_{s,f,t}\\ \sigma_{s,f,t} & h_{f,t}\end{pmatrix}, \Theta_{s,f}\right), \tag{8}$$

where $\Theta_{s,f}$ is the inverse of the number of the degrees of freedom of the t-distribution, and I_{t-1} is the information set as of $t - 1$. Equation (8) states that distribution of the spot and forward returns at t conditional on information available at $t - 1$ is a bivariate Student-t with $1/\Theta_{s,f}$ degrees of freedom. The use of the bivariate conditional t-distribution (where the number of degrees of freedom is estimated), as opposed to the normal, is intended to capture the fat tails of the joint distribution of the returns. Since the t-distribution contains the normal as a limiting case, it is more general than the bivariate normal distribution typically used in other studies of this nature.[11]

11. Autoregressive conditional heteroscedasticity is frequently advanced as an explanation for the kurtosis observed in many return series. It is not necessarily the case, however, that the conditional distribution of the returns is normal. As a result, there may be two sources of excess unconditional kurtosis: time-varying volatility and excess

Two considerations justify this two-step procedure and the use of OLS to estimate equations (3) and (4). First, since the information matrix of the general error-correction model with time-varying volatility is block diagonal with respect to the mean and variance equation parameters, the separate execution of the two steps does not reduce the efficiency of the variance estimates. Second, since the explanatory variables in equations (3) and (4) are identical, the estimated coefficients (and hence the residuals) are the same regardless of whether the equations are estimated singly or jointly. We employ a two-step procedure instead of a full maximum likelihood estimation of both mean and variance equations because of the large number of parameters involved, which makes convergence problematic. Furthermore, we are primarily interested in the dynamics of the variances and covariances.

Given our estimates of $h_{s,t}$, $h_{f,t}$, and $\sigma_{s,f,t}$, we determine the time-series behavior of the variance of V_t and its relation to the behavior of the adjusted spread. Denoting the variance of V_t by $h_{v,t}$, we use the estimates of variances and the covariance and the relation

$$h_{v,t} = h_{s,t} + h_{f,t} - 2\sigma_{s,f,t} \tag{9}$$

to calculate the variance of the change in the forward-spot log difference.

We also use the parameter estimates to calculate the forward price elasticity, e_t. Assume that at t, the relation between the spot and forward returns is

$$\Delta \ln F_t = b(t)\Delta \ln S_t + \nu_t,$$

where ν_t is a random-error term. Then, $e_t = b(t) = \sigma_{s,f,t}/h_{s,t}$. This method of calculating the elasticity has two advantages over that of Fama and French (1988), who regress the contemporaneous forward return against the contemporaneous spot return in subsamples in which the interest adjusted spread is positive and subsamples in which it is negative. First, correlations between the error term and the independent variable may bias coefficients in a regression of one endogenous variable on another. Second, this approach allows us to calculate

kurtosis in the conditional distribution of the error term. Hence, the conditional t-distribution is more general. Moreover, the large t-statistics for the inverse of the degree-of-freedom parameter in our results imply that the coefficient is estimated with considerable precision and that conditional returns exhibit appreciable kurtosis. In order to determine whether our results are sensitive to the more general distribution, we have also estimated the model by using a conditional bivariate normal distribution and have computed robust t-statistics by using the method of Bollerslev and Wooldridge (in press). The results are virtually identical to those we report, with the exception that the t-distribution outperforms the normal distribution (as measured by the value of the log likelihood).

the marginal effect of changes in z_{t-1}^2 on e_t over its entire observed range.

Finally, the estimated conditional variances and covariances determine optimal hedge ratios. The instantaneous variance-minimizing hedge ratio is the number of forward contracts to sell per unit of the spot commodity in order to minimize the variance of the combined position over the next day. It is well known that this ratio equals $\sigma_{s,f,t}/h_{f,t} \equiv H$. From equations (6) and (7) it follows that

$$H = \rho \sqrt{\frac{h_{s,t}}{h_{f,t}}} + \frac{\theta z_{t-1}^2}{h_{f,t}}.$$

III. Empirical Evidence

A. Data

To test the hypotheses concerning the joint dynamics of futures and spot prices discussed in the previous section, we employ data on spot and 3-month forward prices for copper, lead, silver, and zinc from the London Metal Exchange for the period from September 1, 1986, to September 15, 1992, and for aluminum prices for the period from August 27, 1987, to September 15, 1992. As Fama and French (1988) note, these data have desirable properties. Specifically, both the spot and forward prices are determined nearly simultaneously in ring trading, that is, in an open-outcry mechanism. Moreover, there are no limit prices at the LME, so observed prices are market clearing prices. For these reasons, unlike spot prices for other commodities, the spot prices for metals are reliable transactions prices. Finally, unlike agricultural products, metals production is nonseasonal, and therefore the assumptions utilized above that current or permanent shocks predominate are plausible. It should also be noted that $T - t$ (i.e., the time to expiration) is a constant throughout this sample, as the 3-month forward contract traded on any day t becomes "prompt" (i.e., delivery occurs) on day $t + 90$.

There are two open-outcry "rings" each day; each metal trades in the ring twice in each session. For each trading day and for each of the five metals, we collect the second-session, second-ring closing spot and forward prices. We use these prices to calculate close-to-close spot and forward returns.

We also use these prices to calculate $\ln F_t$ and $\ln S_t$. Then, using equation (2), we adjust this difference by the relevant short-term interest rate and the cost of storage to calculate z_t. Copper and lead trade in pounds sterling, whereas aluminum, zinc, and silver are quoted in dollars. We therefore use the Eurosterling rate to calculate financing costs for the first two commodities, and we use the Eurodollar rate to do so for the latter three (after converting each to a continuously compounded equivalent rate). The London Metal Exchange, Ltd., kindly provided weekly warehousing fees for the four industrial metals—aluminum, copper, lead, and zinc. Since delivery can occur in 38 warehouses in 12 countries in Asia, Europe, and North America, and since different warehouses charge different storage fees, there is no unique value of $w_{t,T}$. We use the median value of the warehousing fee (multiplied by 13 in order to determine the cost of storage over the relevant 3-month period) as our estimate of storage costs. We obtained silver storage charges from several warehousing firms in the United States and the United Kingdom. After adjusting for the time value of money and storage fees, the z_t's for all days and all metals are negative. Thus, there were no apparent cash-and-carry arbitrage opportunities.[12]

B. *Exploratory Data Analysis*

An exploratory analysis of the data supports the hypotheses advanced in Section II. Table 1 reports the autocorrelation coefficients for $\Delta \ln S_t$, $\Delta \ln F_t$, $(\Delta \ln S_t)^2$, $(\Delta \ln F_t)^2$, and $(\Delta \ln S_t)(\Delta \ln F_t)$; the squared returns at t are measures of the variances of spot and forward returns at t, and the cross product of spot and forward returns is a measure of their covariance. Although the autocorrelations provide only weak evidence of a time-varying mean, there is strong evidence of time-varying volatility; the autocorrelations of both $(\Delta \ln S_t)^2$ and $(\Delta \ln F_t)^2$ are positive and significant.

Table 2 presents the summary statistics for returns, squared returns, the product of spot and forward returns, and the adjusted spread. Note that the forward return variance is consistently smaller than the spot return variance for the industrial metals; this difference is significant

12. Fama and French (1988) do not adjust for storage costs. Thus, they sometimes find $z_t > 0$.

TABLE 1 **Metals Autocorrelations**

Lag	$\Delta \ln S_t$	$(\Delta \ln S_t)^2$	$\Delta \ln F_t$	$(\Delta \ln F_t)^2$	$\Delta \ln F_t \Delta \ln S_t$
Aluminum:					
1	.11*	.33*	.06*	.21*	.24*
2	.01	.10*	.04	.12*	.05*
3	.03	.19*	−.04	.16*	.11*
4	−.01	.17*	.01	.14*	.09*
5	.07*	.12*	−.01	.09*	.01*
6	−.04	.24*	.00	.13*	.09*
7	−.04	.20*	.02	.14*	.14*
8	.03	.15*	−.04	.09*	.09*
9	.02	.18*	.07*	.11*	.12*
10	.00	.19*	−.01	.18*	.18*
$Q(12)$	30.45*	516.00*	22.87*	288.48*	210.64*

TABLE 1 *(Continued)*

Lag	$\Delta \ln S_t$	$(\Delta \ln S_t)^2$	$\Delta \ln F_t$	$(\Delta \ln F_t)^2$	$\Delta \ln F_t \Delta \ln S_t$
Copper:					
1	.01	.28*	.00	.25*	.26*
2	−.01	.18*	−.07*	.09*	.12*
3	.02	.19*	−.01	.12*	.15*
4	−.08*	.18*	.10*	.14*	.14*
5	.07*	.15*	.07*	.08*	.10*
6	−.08*	.15*	−.04	.09*	.11*
7	−.03	.15*	−.03	.13*	.15*
8	−.01	.16*	.02	.12*	.12*
9	.00	.13*	.02	.08*	.10*
10	−.02	.13*	−.00	.09*	.10*
$Q(12)$	46.19*	558.80*	40.75*	320.67*	390.42*
Lead:					
1	.02	.19*	−.05*	.09*	.11*
2	−.03	.13*	−.06*	.06*	.06*
3	−.06*	.20*	−.06*	.13*	.14*
4	.00	.10*	.03	.07*	.14*
5	.06*	.13*	.05	.05	.05
6	.02	.16*	−.03	.10*	.05
7	.03	.19*	.00	.10*	.11*
8	.04	.13*	.07*	.07*	.10*
9	.04	.15*	.03	.02	.05
10	.04	.14*	.01	.07*	.03
$Q(12)$	29.93*	392.45*	34.56*	102.48*	117.18*
Zinc:					
1	.10*	.28*	.01	.19*	.14*
2	−.06*	.07*	−.01	.06*	.07*
3	−.03	.03	−.03	.05	.03
4	.01	.08*	.03	.05	.09*
5	.04	.09*	.04	.09*	.07*
6	−.01	.13*	−.01	.04	.04
7	−.01	.06*	−.03	.06*	.12*
8	.01	.08*	.03	.13*	.07*
9	.08*	.05	.09*	.06*	.17*
10	.00	.04	.04	.08*	.17*
$Q(12)$	38.22*	201.25*	24.37*	149.55*	113.59*
Silver:					
1	−.03	.09*	−.02	.09*	.09*
2	.00	.16*	.00	.17*	.16*
3	−.06	.14*	−.06*	.13*	.14*
4	.02	.03*	.01	.03	.03
5	−.03	.10	−.02	.09	.09*
6	.01	.01	.00	.01	.01
7	−.01	.01	.00	.02	.01
8	−.03	.02	−.03	.02	.02
9	−.02	.03	−.02	.03	.03
10	.01	.01	−.01	.01	.01
$Q(12)$	12.96	100.88*	11.75	98.84*	99.91*

NOTE.—This table reports the autocorrelation coefficients for daily spot and forward returns, squared daily spot and forward returns, and the product of daily spot and forward returns, for each of the five metals in our sample. The spot return equals $\Delta \ln S_t$, and the forward return equals $\Delta \ln F_t$. The sample period is from September 1 to September 15, 1992, for copper, lead, silver, and zinc. Sample size equals 1,522 for these four metals. The sample period is from August 27, 1987, to September 15, 1992, for aluminum. Sample size equals 1,272 for this metal. The standard error for each coefficient for each metal equals .03. The row labeled "$Q(12)$" reports that Ljung-Box statistic for twelfth-order serial correlation, which is distributed χ^2, with 21 degrees of freedom. The critical value at the 5% level equals 21.

* Coefficients are significant at the 5% level. The $Q(12)$ statistics are significant at the 5% level.

TABLE 2 **Summary Statistics**

	Metal				
	AL	CU	PB	ZN	AG
VAR($\Delta \ln S_t$)	4.19E − 4	3.59E − 4	3.58E − 4	2.89E − 4	2.65E − 4
VAR[$(\Delta \ln S_t)^2$]	1.66E − 6	6.89E − 7	7.18E − 7	5.66E − 7	2.88E − 6
VAR($\Delta \ln F_t$)	2.25E − 4	2.21E − 4	1.98E − 4	1.77E − 4	2.65E − 4
VAR[$(\Delta \ln F_t)^2$]	3.15E − 7	3.16E − 7	2.28E − 7	1.80E − 7	2.72E − 6
VAR[$(\Delta \ln F_t \Delta \ln S_t)$]	5.41E − 7	3.82E − 7	3.15E − 7	1.84E − 7	2.65E − 6
Kurtosis($\Delta \ln S_t$)	8.85	4.29	5.94	5.07	7.21
Kurtosis($\Delta \ln F_t$)	9.99	8.58	6.95	5.31	7.19
MEAN(z_t)	−.0425	−.0551	−.0554	−.0484	−.0011
VAR(z_t)	.0316	.00198	.00240	.00101	1.50E − 5

NOTE.—This table reports the summary statistics for spot and forward returns, squared spot and forward returns, and the product of spot and forward returns and the adjusted spread. MEAN(X) gives the sample mean of variable X. VAR(X) gives the sample variance of variable X. Kurtosis(X) is the coefficient of kurtosis of variable X. The spot return is $\Delta \ln S_t$, whereas the forward return is $\Delta \ln F_t$. The adjusted spread is z_t. AL, CU, PB, ZN, and AG are abbreviations for aluminum, copper, lead, zinc, and silver, respectively. The sample period (number of observations) is from August 27, 1987, to September 15, 1992 (1,272) for AL and from September 1, 1986, to September 15, 1992 (1,522) for the other metals.

at the 1% level for each metal. Conversely, spot and forward return variances are identical (to the fourth significant digit) for silver. Unsurprisingly, (*a*) the absolute value of the average adjusted spread and (*b*) the variance of the adjusted spread are both much larger for the industrial metals than for silver. This is consistent with the hypotheses that (*a*) investors/speculators hold the marginal inventories of silver as a store of value, so these inventories generate no convenience value and/or the stock-out probability is nil, and (*b*) industrial metal inventories generate a convenience yield and/or stock-out probabilities are appreciable for these commodities. Finally, note the significant kurtosis in the returns. This highlights the potential advantages of using the conditional t-distribution in our analysis.

Table 3 reports the correlations between squared spot and forward returns and z_{t-1}^2, the correlation between the product of spot and forward returns and z_{t-1}^2, and the partial correlation (holding lagged-squared spot and forward returns constant) between this product and z_{t-1}^2. First, consider the results for the four industrial metals—aluminum, copper, lead, and zinc. For these commodities, the lagged-squared-adjusted spread is positively correlated with both spot and forward squared returns, which is consistent with the hypothesis that returns are more variable when the spread is wide. Moreover, z_{t-1}^2 is more strongly correlated with the squared spot return than with the squared forward return. This supports the conjecture that current supply conditions have a more pronounced effect on spot than forward volatilities. Finally, the partial correlation between the lagged-squared-adjusted spread and the product of spot and forward returns is strongly

TABLE 3 Spread-squared Return Correlations

	Metal				
	AL	CU	PB	ZN	AG
CORR[$(\Delta \ln S_t)^2$, z_{t-1}^2]	.332	.356	.342	.351	−.094
CORR[$(\Delta \ln F_t)^2$, z_{t-1}^2]	.2047	.207	.246	.097	−.093
CORR($\Delta \ln F_t \Delta \ln S_t$, z_{t-1}^2)	.2039	.236	.248	.128	−.093
PCORR($\Delta \ln F_t \Delta \ln S_t$, z_{t-1}^2)	−.335	−.246	−.252	−.310	.08

NOTE.—This table reports the correlations between squared daily returns and the squared spread, as well as the correlation and partial correlation between the product of daily spot and forward returns and the lagged-squared-adjusted spread. The variable $\Delta \ln S_t$ is the spot return, $\Delta \ln F_t$ is the forward return, and z_{t-1}^2 is the lagged-squared-adjusted spread. CORR(X, Y) reports the correlation between variables X and Y. PCORR($\Delta \ln F_t \Delta \ln S_t$, z_{t-1}^2) reports the partial correlation between the product of spot and forward returns and the lagged-squared adjusted spread, when lagged-squared spot and squared futures returns are held constant. The sample period (number of observations) is from August 27, 1987, to September 15, 1992 (1,272) for AL, and from September 1, 1986, to September 15, 1992 (1,522) for the other metals. Metals are abbreviated as in table 2.

negative. This supports the prediction that the correlation between spot and forward returns declines as the spread widens.

Next, consider the results for silver. The correlations between the lagged-adjusted-squared spread and the squared spot return, squared forward return, and the product of spot and forward returns are negative and small in absolute value. Moreover, the correlation between the spot return and the squared spread and the correlation between the forward return and the squared spread are nearly equal. Also, the partial correlation between the lagged-squared-adjusted spread and the product of the spot and forward returns (holding the lagged-squared spot and forward returns constant) is small and positive. Thus, the relation between the adjusted spread and spot and forward returns for silver differs dramatically from that observed for the industrial metals. The inability of the spread to explain the variances and covariances of silver returns is unsurprising. The spread varies little (as is seen in table 2) and cannot explain changes in silver spot and forward return volatility over time. This reflects the fact that speculators/investors hold large buffer stocks of precious metals.

Finally, we cannot reject the hypotheses that (a) spot and forward prices have a unit root and that (b) z_t is stationary. The t-statistics on the relevant coefficient in five lag-augmented Dickey-Fuller regressions for the spot and forward returns range between −.98 and −1.47. This is not significant at the 5% level in samples over 500, indicating that the spot and forward prices are nonstationary. Moreover, the augmented Dickey-Fuller statistic for z_t equals approximately −3.5 for all the metals. This is significant at the 5% level, which indicates that the interest and storage adjusted spread is mean reverting, as theory suggests.

C. Results and Interpretation: Industrial Metals

Table 4 reports the results from the estimation of equations (5)–(8).[13] Most importantly, the values of δ_3 and ϕ_3 are positive and significant. This is consistent with the theory of storage, as it implies that spot and forward prices are more volatile, the greater the adjusted spread (in absolute value); that is, these prices are more volatile, the lower the level of inventories. Furthermore, for all the metals, $\delta_3 > \phi_3$. As expected, therefore, variations in adjusted spreads induce greater changes in spot return variances than forward return variances. This reflects the greater elasticity of long-run supply curves.

Changes in z_{t-1}^2 explain a large portion of the innovations—shocks—in spot- and forward-return volatilities. There are two sources of innovations to volatility in equations (6) and (7), the lagged-squared spread and lagged-squared residual returns.[14] To measure the fraction of the innovations in volatility attributable to each source, define

$$g_s(z_{t-1}) = \delta_3 z_{t-1}^2,$$

$$q_s(\Delta \ln S_{t-1}) = \delta_2 \epsilon_{t-1}^2,$$

$$x_s(z_{t-1}, \Delta \ln S_{t-1}) = \frac{g_s(z_{t-1})}{g(z_{t-1}) + q_s(\Delta \ln S_{t-1})},$$

$$g_f(z_{t-1}) = \phi_3 z_{t-1}^2,$$

$$q_f(\Delta \ln F_{t-1}) = \phi_2 \eta_{t-1}^2,$$

and

$$x_f(z_{t-1}, \Delta \ln F_{t-1}) = \frac{g_f(z_{t-1})}{g_F(z_{t-1}) + q_f(\Delta \ln F_{t-1})}.$$

In words, $x_s(\cdot, \cdot)$ and $x_f(\cdot, \cdot)$ measure the proportion of spot and forward return volatility innovations, respectively, at t, attributable to variations in the interest and storage adjusted spread at $t - 1$. The values of $x_s(\cdot, \cdot)$ and $x_f(\cdot, \cdot)$ in these samples are typically large. The average value of $x_s(\cdot, \cdot)$ equals .62 for aluminum, .57 for copper, .69 for lead, and .65 for zinc. The average value of $x_f(\cdot, \cdot)$ equals .56 for aluminum, .53 for copper, .59 for lead, and .57 for zinc. Therefore, variations in spreads explain a large fraction of the innovations in both

13. Since our emphasis is on variances and covariances, we do not report the estimates from eqq. (4) and (5). Unsurprisingly, the adjusted R^2s for these regressions are quite low, ranging between .011 and .034. Moreover, no coefficients are statistically significant (when heteroscedastic-consistent standard errors are used) in the spot-return equations, and only a few are even marginally significant in the forward-return regressions. The F-statistics for the regressions are invariably insignificant.

14. The $h_{f,t-1}$ and $h_{s,t-1}$ terms characterize the persistence of the effects of these innovations.

TABLE 4 Spot and Forward Return Volatility Model Estimates

$$\text{Model: } h_{s,t} = \omega_s + \delta_1 h_{s,t-1} + \delta_2 \epsilon_{t-1}^2 + \delta_3 z_{t-1}^2$$
$$h_{f,t} = \omega_f + \phi_1 h_{f,t-1} + \phi_2 \eta_{t-1}^2 + \phi_3 z_{t-1}^2$$
$$\sigma_{s,f,t} = \rho(h_{s,t}, h_{f,t})^{.5} + \theta z_{t-1}^2$$

	Metal				
	AL	CU	PB	ZN	AG
ω_s	.000012	.000013	.000046	.000026	.000035
	(5.27)	(4.11)	(4.35)	(4.72)	(6.40)
δ_1	.7953	.8152	.6541	.7108	.8540
	(31.37)	(31.63)	(10.24)	(16.98)	(46.55)
δ_2	.07630	.08574	.06505	.08291	.2597
	(4.67)	(5.09)	(3.88)	(4.60)	(9.96)
δ_3	.01720	.007016	.01359	.01024	−.000113
	(5.12)	(4.51)	(4.37)	(4.73)	(−1.43)
ω_F	.000008	.000002	.000028	.000014	.000036
	(4.85)	(3.91)	(4.11)	(4.35)	(6.27)
ϕ_1	.8394	.8786	.7322	.7892	.8502
	(39.21)	(44.48)	(13.43)	(23.03)	(45.07)
ϕ_2	.06608	.06380	.06619	.07816	.2645
	(4.96)	(4.99)	(4.02)	(5.13)	(10.00)
ϕ_3	.006235	.002384	.003855	.003098	−.000156
	(4.54)	(4.04)	(3.65)	(3.89)	(−1.66)
ρ	.9888	.9850	.9907	.9569	.997
	(614.39)	(436.32)	(579.40)	(193.28)	(175.69)
θ	−.005747	−.003174	−.004300	−.003130	.00130
	(−9.76)	(−10.12)	(−10.52)	(−7.19)	(1.50)
Θ	.2264	.2131	.2561	.2140	.4330
	(13.45)	(13.07)	(12.25)	(11.36)	(50.06)
Log L	8,858.69	10,379.43	10,392.49	10,310.33	12,532.60
N	1,267	1,517	1,517	1,517	1,517

NOTE.—This table reports the maximum likelihood estimate of the variance and covariance model parameters. The row labeled "Log L" reports value of the log likelihood. The row labeled "N" provides sample size. t-statistics are in parentheses. The variable $h_{s,t}$ is the conditional spot-return variance at t; $h_{f,t}$ is the conditional forward-return variance at t; ϵ_{t-1}^2 is the squared residual spot return at $t - 1$; η_{t-1}^2 is the squared residual forward return at $t - 1$; z_{t-1}^2 is the squared-adjusted spread at $t - 1$. Metals are abbreviated as in table 2.

spot and forward volatility. This is again consistent with the theory of storage, and it implies that fundamental factors are important determinants of spot and forward price volatility.

Indeed, these figures tend to understate the importance of spreads in determining volatility innovations, because squared spreads are strongly and positively correlated with $x_s(\cdot, \cdot)$ and $x_f(\cdot, \cdot)$ for all metals. The correlations between z_{t-1}^2 and $x_s(\cdot, \cdot)$ range between .22 and .29, whereas the correlations between z_{t-1}^2 and $x_f(\cdot, \cdot)$ lie between .33 and .52. Not surprisingly, then, there are also positive correlations between $x_s(\cdot, \cdot)$ and $h_{s,t}$ (ranging between .15 and .28) and between $x_f(\cdot, \cdot)$ and $h_{f,t}$ (ranging between .13 and .26). Thus, spreads explain a larger fraction of volatility innovations when spreads and volatility are large than when they are small. Put another way, spreads are an even more important determinant of spot and forward return variance innovations when returns are particularly volatile.

The correlations between z_{t-1}^2 and $h_{s,t}$ and $h_{f,t}$ are alternative measures of the closeness of the relation between spread and price volatility. The correlation between the lagged-squared-adjusted spread and spot volatility equals .92 for aluminum, .90 for copper, .96 for lead, and .75 for zinc. The correlations between the lagged-squared-adjusted spread and forward volatility equals .89 for aluminum, .82 for copper, .92 for lead, and .60 for zinc. These correlations are considerably higher than the correlations between the squared-lagged residual and conditional spot- and forward-return variances, which range between .35 and .45. Thus, spot- and forward-return conditional volatilities vary more closely with z_{t-1}^2 than with the other source of volatility innovations.

The estimates of the conditional volatilities, $h_{s,t}$ and $h_{f,t}$, also permit a direct test of Samuelson's (1965) hypothesis that spot returns are more volatile than forward returns. In the industrial metals samples, the fitted values for $h_{s,t}$ almost always exceed those for $h_{f,t}$. For aluminum, $h_{s,t} > h_{f,t}$ for 1,208 of the 1,267 observations, whereas the average of $h_{f,t}/h_{s,t}$ equals .73. For copper, the ratio is less than one for 1,516 of the 1,517 observations, and its average value equals .68. For lead, the ratio is less than one for 1,515 of the 1,517 observations and has an average value of .66. Finally, for zinc, the ratio is less than one for 1,492 of the 1,517 observations and averages .68. These results are consistent with Samuelson's hypothesis. Moreover, the ratio of forward- to spot-return volatilities varies inversely with the lagged-squared-adjusted spread. The correlation between z_{t-1}^2 and $h_{f,t}/h_{s,t}$ equals $-.49$ for aluminum, $-.54$ for copper, $-.61$ for lead, and $-.52$ for zinc. Thus, as theory predicts, spot-return volatility rises relative to forward-return volatility as the market becomes more inverted (i.e., as z_{t-1}^2 becomes larger).

Figures 5–8 illustrate these various results clearly. The figures plot $-.05z_{t-1}^2$ rather than z_{t-1}^2 in order to improve their clarity and scaling. The relation between the spread and conditional volatilities is clear in these figures; spot- and forward-return volatilities peak and trough simultaneously with z_{t-1}^2.

These pictures give striking visual evidence of the extremely close relation between volatility and spreads. They also reveal that the conditional spot volatility almost always exceeds the conditional forward volatility, and this difference is most pronounced when the squared spread is large. Combined with the quantitative results, they provide clear evidence of the primacy of fundamentals in determining price volatility, since the return variances behave exactly as the fundamentals-based theory of storage implies.

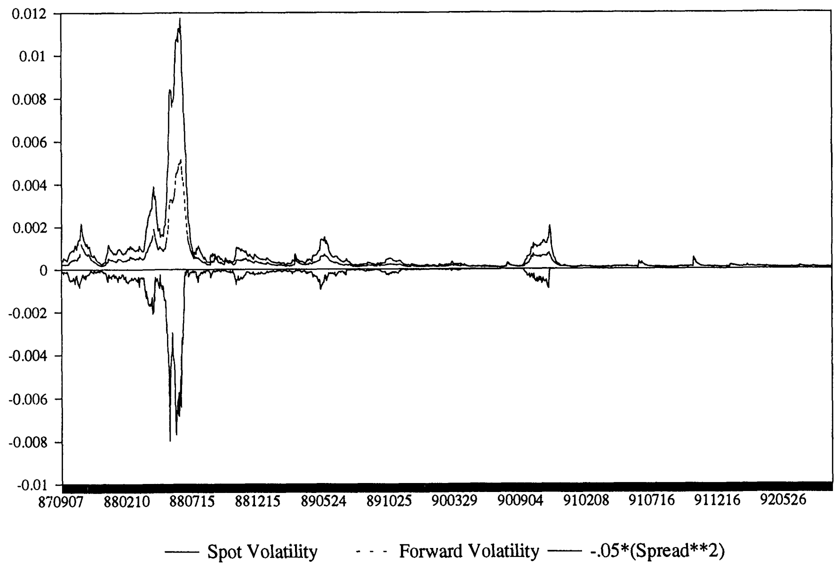

FIG. 5.—Aluminum spot and forward volatility, and the squared-adjusted spread.

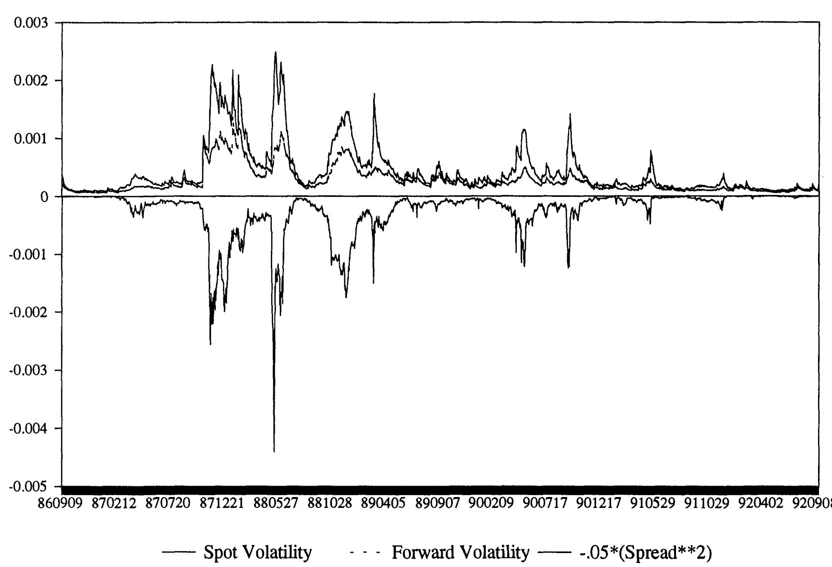

FIG. 6.—Copper spot and forward volatility, and the squared-adjusted spread.

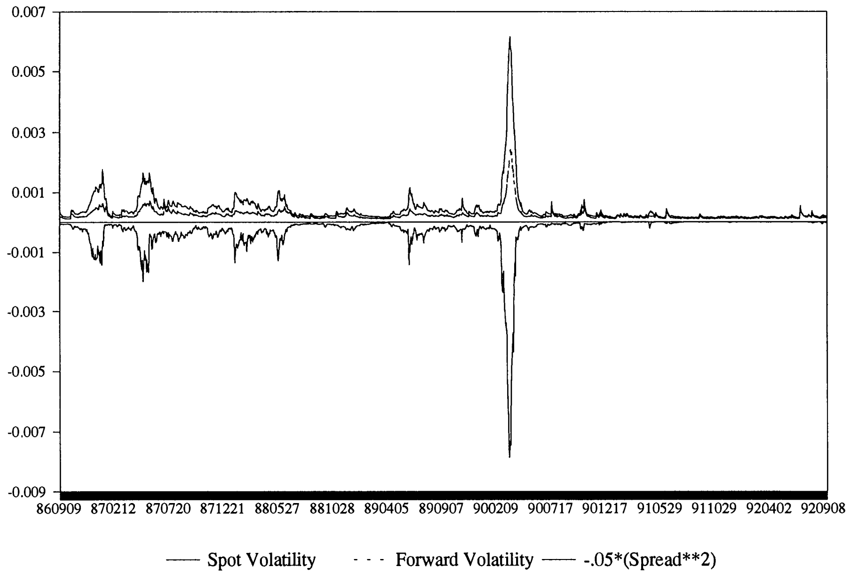

FIG. 7.—Lead spot and forward volatility, and the squared-adjusted spread

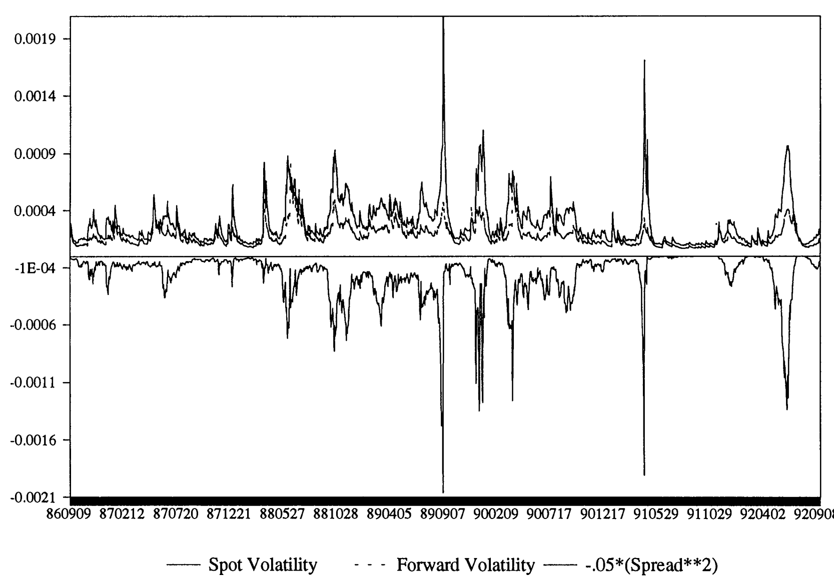

FIG. 8.—Zinc spot and forward volatility, and the squared-adjusted spread

The results in table 4 also reveal that $\theta < 0$ for all metals. Recall that

$$\sigma_{s,f,t} = \rho \sqrt{h_{s,t} h_{f,t}} + \theta z_{t-1}^2.$$

Thus, defining $\rho^*_{s,f,t}$ as the conditional correlation between $\Delta \ln F_t$ and $\Delta \ln S_t$,

$$\rho^*_{s,f,t} = \rho + \frac{\theta z_{t-1}^2}{\sqrt{h_{s,t} h_{f,t}}}. \tag{10}$$

Note that z_{t-1}^2 affects the conditional correlation between the spot and forward returns through both the numerator and the denominator of the final term of equation (9); the overall effect of the lagged-squared-adjusted spread on the contemporaneous correlation between spot and forward returns is therefore ambiguous a priori. In our sample, however, increasing the squared spread systematically reduces this correlation. Specifically, the correlation between $\rho^*_{s,f,t}$ and z_{t-1}^2 equals $-.55$ for aluminum, $-.64$ for copper, $-.79$ for lead, and $-.87$ for zinc. Thus, as expected, correlations between spot and forward returns decline when adjusted spreads widen. It is also of interest to note that the estimates of ρ are very close to one, ranging between .9659 (zinc) and .9907 (lead). Taken together, these results are consistent with the theory of storage. When spreads are zero, spot and forward returns are nearly perfectly correlated, because the ability to execute cash-and-carry and reverse cash-and-carry arbitrage ensures that spot and futures prices move nearly in lockstep. These correlations decline systematically, however, as spreads widen. This result obtains because when stocks are low, the arbitrage link between spot and forward prices is attenuated, and they vary more independently.

The results for the relation between volatilities, covariances, and the level of the spreads have implications for the dynamic properties of V_t. The theory of storage predicts a positive correlation between h_{vt} and z_{t-1}^2, and this is indeed observed. This correlation equals .99, .96, .99, and .84 for aluminum, copper, lead, and zinc, respectively. Our results also produce the predicted relation between the forward price elasticity, e_t, and z_{t-1}^2. Note that $e_t = \sigma_{s,f,t}/h_{s,t} = \rho^*_{s,f,t} \sqrt{h_{f,t}/h_{s,t}}$. Since both the volatility ratio and the correlation are approximately one when $z_{t-1}^2 = 0$, $e_t \approx 1$ in this case as well. Moreover, since the correlation and volatility ratio both fall as z_{t-1}^2 increases, the elasticity also falls, as the French (1986) model predicts.

Finally, our estimates imply that there are significant variations in the variance-minimizing hedge ratio, H. For aluminum, H varies between .91 and 1.72; for copper, it lies between .97 and 1.82; for lead, it is bounded by .92 and 1.60; while for zinc, it varies between .80 and 2.10. Moreover, there is a strong positive correlation between H and z_{t-1}^2. This correlation equals .43 for aluminum, .39 for copper, .49 for lead, and .41 for zinc. Those who use forward contracts to hedge short-term changes in the value of metals inventory should therefore alter their hedge ratios as spreads change.[15]

D. Results and Interpretations: Silver

The results for silver reported in table 4 support the conclusions derived from the preliminary data analysis. The squared-adjusted spread coefficient is negative and not statistically significant in the spot and forward return equations; according to the theory of storage, this coefficient should be positive. Moreover, the spread does not have a statistically significant effect on the correlation between spot and forward returns and again is of the wrong sign. The conditional volatilities of spot and forward returns are nearly equal throughout the sample; the average value of $h_{f,t}/h_{s,t}$ equals .994. Finally, the forward price elasticity is very close to 1.00 throughout the sample.

These results differ strikingly from those found for the industrial metals. This differential performance provides additional support for the theory of storage. Agents hold large inventories of this precious metal as a store of value. The theory predicts that as a result of these large buffer stocks, the spread should be small and vary little. Moreover, since stock outs and convenience considerations are largely irrelevant for silver, the spread should have little power to explain the behavior of silver prices. Thus, the striking differences observed between the behavior of silver prices and industrial metals prices are exactly as the theory predicts.

IV. Summary and Conclusions

This article exploits several implications of the theory of storage in order to quantify the role of supply and demand fundamentals in determining metal price volatility. Several of these implications of this theory have not been tested previously, and our methods allow far more powerful tests of some previously explored implications.

The results for industrial metals provide clear support for the theory and are thus consistent with the hypothesis that spot-and-forward-return dynamics are strongly related to variations in fundamental supply and demand conditions. As predicted, (1) spot-return volatility varies directly with the square of the spread; (2) forward-return volatility varies directly with the square of the spread, but not by as much as spot-return volatility; (3) forward returns are less volatile than spot returns; (4) the volatility of forward returns declines relative to the volatility of spot returns as the square of the interest and storage ad-

15. The hedge ratio depends on the horizon of the hedge. For a 3-month horizon, the hedge ratio is, of course, equal to one regardless of the relation between the spread and daily variances and covariances. Thus, the importance of the spread in determining the hedge ratio declines as the hedging horizon increases.

justed spread increases; (5) correlations between spot and forward returns vary inversely with the square of the spread; (6) the volatility of the changes in the forward-spot logarithmic price spread varies directly with the square of the level of this spread; (7) forward price elasticities increase as the adjusted spread narrows and are approximately equal to one when the market is at full carry; and (8) hedge ratios vary directly with the squared spread.

In contrast, the adjusted spread does not explain the dynamics of silver prices. This finding strengthens the conclusion that fundamentals drive metal price dynamics. Marginal storers hold silver as a store of value, rather than to smooth consumption and production of the commodity over time (as is the case for industrial metals). Therefore, one would expect the adjusted spread to be small and nearly invariant. As a result, it should not be an important determinant of the dynamics of silver prices. The data support these predictions. This finding demonstrates that the industrial metal results are not attributable to misspecification or statistical error or to some difference in the way "animal spirits" affect spot and futures prices. Thus, the results for silver make the argument for fundamentals even more compelling.

These results do more than simply provide support for the theory of storage or show that spot-and-forward-return volatilities depend on the interest and storage adjusted spread. They also demonstrate that, for the industrial metals in the sample period, variations in this spread explain most of the innovations in spot-and-forward-return variances and that the contribution of the spread to volatility is even larger when volatility is high. Since it is well known that spreads vary systematically with market fundamentals (which include stocks relative to demand and the elasticity of supply), these results strongly suggest that variations in market fundamentals explain a large proportion of volatility changes.

This conclusion is strengthened when one recognizes that the other source of volatility innovations in our model, lagged-return shocks, may also reflect the arrival of information related to fundamental demand and supply conditions. Thus, our findings suggest that, at least for the industrial metals, variations in volatility are largely attributable to variations in fundamental demand and supply conditions rather than to speculative noise trading.

References

Anderson, R. 1985. Some determinants of the volatility of futures prices. *Journal of Futures Markets* 5 (Fall): 331–48.

Baillie, R., and Myers, R. 1991. Bivariate GARCH estimation of the optimal commodity futures hedge. *Journal of Applied Econometrics* 6 (April–June): 109–24.

Black, F. 1986. Noise. *Journal of Finance* 41 (December): 307–27.

Bollerslev, T., and Wooldridge, J. 1992. Quasi-maximum likelihood estimation and infer-

ence in dynamic models with time varying covariances. *Econometric Review* 11, no. 2:143–72.
Brennan, D.; Wright, B.; and Williams, J. 1992. Is convenience yield a necessary hypothesis? "The supply of storage" in a wheat market. Mimeographed. Stanford, Calif.: Food Research Institute.
Brennan, M. 1991. The price of convenience and the pricing of commodity contingent claims. In D. Lund and B. Oksendal (eds.), *Stochastic Models and Options Values*. New York: Elsevier.
Bresnahan, T., and Spiller, P. 1986. Futures market backwardation under risk neutrality. *Economic Inquiry* 24 (July): 429–41.
Bresnahan, T., and Suslow, V. 1985. Inventories as an asset: The volatility of copper prices. *International Economic Review* 26 (June): 409–24.
Deaton, A., and Laroque, G. 1991. On the behavior of commodity prices. *Review of Economic Studies* 59 (January): 1–23.
Fama, E., and French, K. 1987. Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage. *Journal of Business* 60 (January): 55–74.
Fama, E., and French, K. 1988. Business cycles and the behavior of metals prices. *Journal of Finance* 43 (December): 1075–93.
French, K. 1986. Detecting spot price forecasts in futures prices. *Journal of Business* 59, suppl.: S39–S54.
Gray, R., and Peck, A. 1981. The Chicago wheat market: Recent problems in historical perspective. *Food Research Institute Studies* 18:89–115.
Howell, L. 1956. Influence of certified stocks on spot-futures price relationships for cotton. U.S.D.A. Technical Bulletin 1151.
Kaldor, N. 1939. Speculation and economic stability. *Review of Economic Studies* 7 (January): 1–27.
Kenyon, D.; Kling, K.; Jordan, J.; Seale, W.; and McCabe, N. 1987. Factors affecting agricultural futures price variance. *Journal of Futures Markets* 7 (February): 73–92.
Keynes, J. 1930. *A Treatise on Money*. New York: Harcourt Brace & Co.
Kroner, K., and Sultan, J. 1991. Foreign currency futures and time varying hedge ratios. In S. G. Rhee and R. Change (eds.), *Pacific-Basin Capital Markets Research*. Vol. 2. Amsterdam: North Holland.
Kyle, A. 1985. Continuous auctions and insider trading. *Econometrica* 53 (November): 1315–35.
Pindyck, R. 1990. Inventories and the short-run dynamics of commodity prices. Mimeographed. Cambridge, Mass.: Massachusetts Institute of Technology.
Roll, R. 1984. Orange juice and weather. *American Economic Review* 74 (December): 861–80.
Ross, S. 1989. Information and volatility: The no-arbitrage martingale approach to timing and resolution irrelevancy. *Journal of Finance* 44 (March): 1–17.
Samuelson, P. 1965. Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review* 6 (Spring): 41–49.
Shiller, R. 1989. *Market Volatility*. Cambridge, Mass.: MIT Press.
Telser, L. 1958. Futures trading and the storage of cotton and wheat. *Journal of Political Economy* 66 (June): 233–55.
Thompson, S. 1986. Returns to storage in coffee and cocoa futures markets. *Journal of Futures Markets* 6 (Winter): 541–64.
Weymar, F. 1974. The effects of cocoa inventories on price forecasting. In R. Kotey, C. Oklai, and B. Rourke (eds.), *The Economics of Cocoa Production and Marketing*. Legon: University of Ghana.
Williams, J. 1986. *The Economic Function of Futures Markets*. Cambridge: Cambridge University Press.
Williams, J., and Wright, B. 1989. A theory of negative prices for storage. *Journal of Futures Markets* 9 (February): 1–13.
Williams, J., and Wright, B. 1991. *Storage and Commodity Markets*. Cambridge: Cambridge University Press.
Working, H. 1948. The theory of the inverse carrying charge in futures markets. *Journal of Farm Economics* 30 (February): 1–28.
Working, H. 1949. The theory of the price of storage. *American Economic Review* 39 (December): 1254–62.