THE ECONOMIC JOURNAL

MAY 1993

The Economic Journal, **103** (*May*), 511–530. Published by Blackwell Publishers, 108 Cowley Road, Oxford OX4 1JF, UK and 238 Main Street, Cambridge, MA 02142, USA.

THE PRESENT VALUE MODEL OF RATIONAL COMMODITY PRICING*

Robert S. Pindyck

The present value model is the most basic description of rational asset pricing. It says that price, P_t, equals the sum of current and discounted expected future payoffs, or benefits, from ownership of the asset:

$$P_t = \delta \sum_{i=0}^{\infty} \delta^i \, \mathrm{E}_t \, \psi_{t+i}. \tag{1}$$

Hence the model explains changes in asset prices in terms of 'fundamentals', i.e. changes in expected future payoffs (ψ_{t+i}) or discount rates (δ). Most tests of the model have used data for stocks, where the payoffs are dividends, or bonds, where the payoffs are interest and principal payments. These tests have had mixed outcomes, due in part to statistical and data problems.[1]

This paper explores the limits of the present value model by testing its ability to explain the pricing of storable commodities. Applying the present value model to commodities is useful for a number of reasons. First, the model is helpful in understanding price movements, and lets us test the rationality of commodity pricing in a way that is very different from earlier tests. Second, these tests provide evidence of the robustness of the present value model. (If the model is valid, it should explain the pricing of any asset that yields a payoff stream.) Third, if the commodity is traded on a futures market, the model can be written entirely in terms of spot and futures prices, and provides a parsimonious description of rational price dynamics.

For a storable commodity, the payoff stream ψ_t is the convenience yield that accrues from holding inventories, i.e. the value of any benefits that inventories provide, including the ability to smooth production, avoid stockouts, and facilitate the scheduling of production and sales. Convenience yield is the

* This research was supported by M.I.T.'s Center for Energy Policy Research, and by the National Science Foundation under Grant No. SES 90-22823. My thanks to Mark Cooper, Steven Lotwin, and Prabhat Mehta for their research assistance, and to Terence Agbeyegbe, Ken Froot, Philip Verleger, Jeffrey Williams, and two anonymous referees for helpful comments.

[1] Most of these tests attempt to show excess volatility or predictability of returns. For a discussion of such tests, see Mankiw *et al.* (1991). Campbell and Shiller (1987) test restrictions implied by the model for the joint dynamics of P_t and ψ_t, and Pindyck and Rotemberg (1990*b*) develop tests based on the correlations of returns. One problem when applying these tests to stocks is with the measurement of payoffs; dividends and earnings are paid and announced quarterly, but firms often make statements in advance about these variables.

reason that firms hold inventories even when the expected capital gain is below the risk-adjusted rate, or negative. While economists have debated the relative importance of these different benefits, for many commodities convenience yield is quantitatively important. As shown below, firms sometimes incurred expected costs of 5 to 10% per month – plus interest and storage costs – to maintain stocks of copper, lumber, and heating oil.

The convenience yield that accrues to the owner of a commodity is directly analogous to the dividend on a stock. If the commodity is well defined and easily traded, and if aggregate storage is always positive, then (1) always holds, and price must equal the present value of the flow of expected future convenience yields. The present value model thus provides a compact explanation for changes in a commodity's price; they are due to changes in expected future convenience yields. We usually try to explain commodity price movements in terms of changes in current and future demand and supply, but changes in demand and supply in turn cause changes in current and expected future convenience yields. Hence the present value model can be viewed as a highly reduced form version of a dynamic supply and demand model.

For some commodities, such as gold, the convenience yield is almost always very small, and often insignificantly different from zero. The reason is that inventories, which are held mostly for 'investment' purposes, are very large relative to production (for gold, about 50 times annual production). But the present value model also applies to such commodities, and provides a fundamentals-based explanation of why rational investors would hold them. Investors should hold these commodities if they think there is a large enough probability that convenience yield will rise substantially in the future. With gold, this could occur if the metal were some day monetised, which would cause inventories to fall dramatically and convenience yield to rise.

For commodities traded on futures markets, convenience yield can be measured directly and (if the futures market is efficient in the sense that there are no arbitrage opportunities) without error from the relation between spot and futures prices. As a result, the present value model is also parsimonious in terms of data; tests can rely on data only for spot and futures prices. One does not, for example, need data on inventories, production costs, or other variables that affect supply, demand, or convenience yield.

I exploit futures price data to test the ability of the present value model to explain the prices of four commodities – copper, lumber, heating oil, and gold. To do this, I draw extensively on work by Campbell and Shiller (1987), who showed that the present value model implies that the price of an asset and its payoff stream are cointegrated, and derived testable implications for the joint dynamics of the two. I show that the present value model imposes similar restrictions for the joint dynamics of the spot and futures prices of a storable commodity.

The basic theory is presented in the next section. I first review the arbitrage relation that determines a commodity's convenience yield from its spot and futures prices. I then discuss the restrictions on the joint dynamics of spot and futures prices implied by (1), and a set of tests that follow from those restrictions. Finally, I also derive a present value relation for the ratio of convenience yield to price. This relation is similar to one derived by Campbell and Shiller (1989) for the log dividend-price ratio of a stock, and when combined with a model for the commodity's expected return, can be tested in the same way that (1) is.

Section II discusses the data and reviews the behaviour of prices, convenience yields, and excess returns for the four commodities. Tests of the present value model are presented in Section III. The results are mixed. Heating oil prices conform closely to the model, and none of the constraints implied by (1) are rejected. Gold, however, does not conform to the model, and copper and lumber are in between. Given these results, it is useful to see whether other tests of market efficiency result in similar patterns across commodities. Section IV examines the serial dependence of excess returns. Cutler *et al.* (1990) studied the serial correlation of returns for a broad range of assets, including gold, silver, and an index of industrial metals, but ignored convenience yield. This can lead to significant errors for industrial commodities, where convenience yield is often a large component of returns. I find that the extent of serial correlation in excess returns parallels conformance with the present value model; there is no significant serial correlation for heating oil, there is some for copper and lumber, and there is a considerable amount for gold.

I. THE PRESENT VALUE MODEL

The present value model is given by (1), where ψ_t is the 1-period per unit net marginal convenience yield, i.e. the benefit flow from holding a marginal unit of the commodity from the beginning to the end of period t, net of storage and insurance costs. Here, $\delta = 1/(1+\mu)$, where μ is the commodity-specific 1-period discount rate, i.e. the expected rate of return an investor would require to hold a unit of the commodity. (Note that (1) is the solution to the difference relation, $E_t P_{t+1} = (1+\mu) P_t - \psi_t$.) For the time being I will assume that μ is constant, and can be written as $\mu = r + \rho$, where r is the risk-free rate and ρ is a risk premium.

Futures Prices, Spot Prices, and Convenience Yield

For commodities with actively traded futures contracts, we can use futures prices to measure the net marginal convenience yield. Let $\psi_{t,T}$ be the (capitalised) flow of marginal convenience yield net of storage costs over the period t to $t+T$, per unit of commodity. Then, to avoid arbitrage opportunities, $\psi_{t,T}$ must satisfy:

$$\psi_{t,T} = (1+r_T) P_t - f_{T,t}, \tag{2}$$

where P_t is the spot price, $f_{T,t}$ is the *forward* price for delivery at $t+T$, and r_T is the risk-free T-period interest rate. To see why (2) must hold, note that the (stochastic) return from holding a unit of the commodity from t to $t+T$ is $\psi_{t,T} + (P_{t+T} - P_t)$. If one also shorts a forward contract at time t, one receives a total return of $\psi^*_{t,T} + f_{T,t} - P_t$. No outlay is required for the forward contract and this total return is non-stochastic, so it must equal $r_T P_t$, from which (2) follows.

For most commodities, futures contracts are much more actively traded than forward contracts, and futures price data are more readily available. A futures contract differs from a forward contract only in that it is 'marked to market', i.e. there is a settlement and transfer of funds at the end of each trading day. As a result, the futures price will be greater (less) than the forward price if the risk-free interest rate is stochastic and is positively (negatively) correlated with the spot price.[2] However, for most commodities the difference in the two prices is extremely small. (See French (1983) and Pindyck (1990).) Thus I use the futures price, $F_{T,t}$, in place of the forward price in (2). Also, I work with the 1-month convenience yield, which I denote as ψ_t, and futures price $F_{1,t}$.

Note that for the present value model to hold, inventories must always be positive, i.e. stockouts must not occur.[3]. We never observe aggregate inventories falling to zero in the data, but as Kahn (1991) points out for inventories of manufactured goods, one could argue that stockouts still occur. First, stockouts might occur with very low probability (but at very high cost to the firm), so they are simply not observed in a sample of 20 or so years. Second, the data aggregate inventories for different products and firms, so stockouts might occur for some products and/or firms. But these are not likely to be problems for the commodities studied here. First, the products are homogeneous and very clearly defined. Second, futures (and forward) markets are extremely liquid and have low transactions costs; any firm can easily buy or sell inventories through these markets, and therefore need never experience a stockout. Finally, I have shown elsewhere (1990) that at least for copper, heating oil, and lumber, convenience yield is highly convex in the aggregate level of inventories, and becomes very large as that level becomes small, so that firms would never allow stockouts to occur.

Implications of the Present Value Model for Spot and Futures Prices

As Campbell and Shiller (1987) have shown, if P_t and ψ_t are both integrated of order 1, the present value relation of (1) implies that they are cointegrated, and the cointegrating vector is $(1 \ -1/\mu)'$. One can therefore define a 'spread',

$$S_t' = P_t - (1/\mu)\,\psi_t, \tag{3}$$

which will be stationary. Hence, in principle, one could estimate the expected return on a commodity, μ, by running a cointegrating regression of P_t and ψ_t.

In addition, it is easily shown that (1) and (3) imply that:

$$S_t' = (1/\mu)\,\mathrm{E}_t\,\Delta P_{t+1}. \tag{4}$$

[2] If the interest rate is non-stochastic, the present value of the expected daily cash flows over the life of the futures contract equals the present value of the expected payment at termination of the forward contract, so the futures and forward prices must be equal. If the interest rate is stochastic and positively correlated with the price of the commodity (as for most industrial commodities), daily payments from price increases will on average be more heavily discounted than payments from price decreases, so the initial futures price must exceed the forward price.

[3] Deaton and Laroque (1992) developed a model of commodity prices in which stockouts play a key role – prices are usually stable, and sudden price flares are accompanied by inventory falling to near zero. People hold inventory because price goes up more when there is a shortfall than when there is a glut, making storage profitable. But there is no convenience yield in their model; inventories are held only as a speculation against price shocks.

Hence P_t and ψ_t contain all information necessary to forecast P_{t+1} optimally. If the futures market is efficient, this is equivalent to saying that P_t and $F_{1,t}$ are sufficient to forecast P_{t+1} optimally. Substituting (2) and (3) into (4) gives the standard result:

$$\mathrm{E}_t P_{t+1} = F_{1,t} + (\mu - r) P_t, \tag{5}$$

i.e. the futures price is a biased predictor of the future spot price, and the bias is equal to the commodity's expected excess return. Thus either (4) or (5) can be used to forecast P_{t+1} if μ is known. Finally, Campbell and Shiller also show that (1) and (3) together imply that:

$$\mu S'_t = \mathrm{E}_t \sum_{i=1}^{\infty} \delta^i \Delta\psi_{t+i} \tag{6}$$

so that $\mu S'_t$ is the present value of expected future changes in the convenience yield.

We can use (4) and (6) to see how futures and spot prices describe the market's expectation of how ψ_t and P_t will evolve. Assume for simplicity that $\mu = r$, so that $S'_t = (1/r)(F_{1,t} - P_t)$. (This is approximately the case for most agricultural commodities, as well as gold.) First, suppose that $\psi_t = 0$, so that $\mathrm{E}_t(P_{t+1}) = F_{1,t} = (1+r) P_t$, and $S'_t = P_t$. Although convenience yield is currently zero in this case, people hold stocks of the commodity and rationally expect price to rise at the rate of interest because they expect the convenience yield to rise in the future. (In fact, $P_t = \mathrm{E}_t \Sigma_i \delta^i \Delta\psi_{t+i}$, the present value of expected future increases in convenience yield.) This is typically the case for gold, where stocks are very large relative to production. If holdings of gold are based on 'rational fundamentals' (as opposed to a rational bubble, in which case (1) includes a term b_t satisfying $b_t = \delta \mathrm{E}_t b_{t+1}$), it must be because there is a chance that gold's convenience yield will rise sharply in the future.

Now suppose that $P_t < F_{1,t} < (1+r) P_t$. Then $S'_t > 0$, and both price and convenience yield are expected to rise. Note that $S'_t < 0$ only if ψ_t is large enough so that $F_{1,t} < P_t$; then the present value of expected future changes in ψ_t is negative. This would mean that price and convenience yields are expected to fall as supply and demand adjust towards long-run equilibrium levels and inventories rise. These patterns for P_t and ψ_t can be seen in the data for copper, where sharp increases in the spot price occurred in 1974, 1979–80, and 1988–9 as a result of strikes and other disruptions to supply that were expected (correctly) to be temporary.

As Campbell and Shiller (1987) have shown, (4) and (6) can be used to test the present value model. First, suppose μ has been estimated (e.g. from the cointegrating regression), and consider a vector of variables $\mathbf{z}_t$ (e.g. production, inventories, etc.) that might be expected to affect future spot prices. Then (4) implies that in regressions of the form:

$$\Delta P_t = \alpha_0 + \alpha_1 S'_{t-1} + \Sigma_i b_i z_{i,t-1} + \epsilon_t \tag{7}$$

the b_is should be groupwise insignificant. Second, (6) implies that Granger causality tests should show causality from S'_t to future $\Delta\psi_{t+i}$s. Finally, (6) also implies a set of cross-equation restrictions on a vector autoregression of S'_t and $\Delta\psi_t$.

One problem is that if P_t and ψ_t are in nominal terms, the nominal expected return μ will fluctuate, even if the real return is constant. Campbell and Shiller deal with this for stocks and bonds by deflating the variables, but this can introduce measurement noise. With futures market data, however, we can avoid this problem altogether by using (2), with the futures price replacing the forward price. Define a new spread $S_t = \mu S_t'$, and substitute (2) for ψ_t:

$$S_t = F_{1,t} - (1-\rho)\, P_t, \tag{8}$$

where $\rho = \mu - r$ is the expected excess return on the commodity. Thus S_t is the futures-spot spread, adjusted for the forecast bias in the futures price. Also, (8) implies that the futures and spot prices are cointegrated, with cointegrating vector $(1 \;\; \rho - 1)'$. Hence a simple regression of $F_{1,t}$ on P_t can be used to estimate the expected excess return, ρ. If real expected returns are constant, the expected excess return should likewise be constant, and can be estimated from this regression without recourse to the CAPM or some related model of asset pricing.

Equation (4) can also be written in terms of S_t, and then becomes:

$$S_t = \mathrm{E}_t\, \Delta P_{t+1}, \tag{9}$$

i.e. the spread S_t is an unbiased forecast of the change in the spot price. (This can also be derived directly from (5).) Again, the current futures and spot prices must be sufficient for the optimal prediction of future spot prices. This condition is sometimes used to test the efficiency of future markets, but its failure need not imply that the futures market is inefficient. It could instead mean that the spot price deviates from the present value relation (1), so that the bias between the futures price and the expected spot price differs from ρP_t, and (9) does not hold.

Tests of the Model

Given ρ, (6) and (9), with S_t replacing $\mu S_t'$ in (6), can be used to test (1). First, (9) implies that any variables in the information set at $t-1$ should be uncorrelated with the residuals of a regression of ΔP_t on S_{t-1}. Hence we can run regressions of the form:

$$\Delta P_t = \alpha_0 + \alpha_1 S_{t-1} + \Sigma_i\, b_i\, z_{i,\,t-1} + \epsilon_t, \tag{10}$$

where the z_is are any variables that might affect price, including commodity-specific ones such as production and inventories, and economy-wide ones such as GNP growth and inflation. We then test whether b_1, b_2, etc. are significantly different from zero.

This requires an estimate of ρ to construct S_t; I first use the estimate from the regression of $F_{1,t}$ on P_t, and then the sample mean of ρ (i.e. the sample mean of $(P_{t+1} - F_{1,t})/P_t$). A failure of the test could mean that (9) does not hold, or that the estimate of ρ used in S_t is very different from the true value. This second possibility can be ruled out by also running the regression:

$$\Delta P_t = \alpha_0 + \alpha_1 P_{t-1} + \alpha_2 F_{1,t-1} + \Sigma_i\, b_i\, z_{i,\,t-1} + \epsilon_t \tag{11}$$

and again testing that the b_is are zero.

Second, since $S_t = \mu S_t'$, (6) implies that S_t should Granger-cause $\Delta\psi_t$. I run Granger causality tests between S_t and $\Delta\psi_t$, again, constructing S_t first using the estimate of ρ from the cointegrating regression, and then using the sample mean.

Finally, as Campbell and Shiller show, (1) implies constraints on the parameters of a vector autoregression of S_t and $\Delta\psi_t$. Specifically, consider the pth-order vector autoregression:

$$\Delta\psi_t = \sum_{k=1}^{p} \gamma_{11k}\Delta\psi_{t-k} + \sum_{k=1}^{p} \gamma_{12k} S_{t-k}, \tag{12a}$$

$$S_t = \sum_{k=1}^{p} \gamma_{21k}\Delta\psi_{t-k} + \sum_{k=1}^{p} \gamma_{22k} S_{t-k}. \tag{12b}$$

Note from (6) that S_t is the present discounted value of the expected future $\Delta\psi_t$s. This in turn implies that the parameters γ_{ijk} must satisfy the following set of cross-equation restrictions:

$\gamma_{21k} = -\gamma_{11k}, k = 1,\dots,p, \gamma_{221} = 1/\delta - \gamma_{121},$ and $\gamma_{22k} = -\gamma_{12k}, k = 2,\dots,p.$[4]
These restrictions provide another test of the present value model.

The Dynamics of the Percentage Net Basis

The tests above follow from constraints that (1) imposes on spot and futures prices. Alternatively, one can work with the differential form of (1) and study the components of commodity returns. By imposing some structure on expected returns (e.g. the CAPM), one can constrain the ratio of the net convenience yield to price. This ratio, called the *percentage net basis*, is analogous to the dividend-price ratio for a stock.[5] Campbell and Shiller (1989) have derived an approximate present value relation for the *log* dividend-price ratio, and have shown that it implies parameter restrictions on a vector autoregression of this ratio and the difference between the expected return and the dividend growth rate. Because the net convenience yield is sometimes negative, I work with a simple ratio, and derive a similar approximate present value relation. This yields parameter constraints on a vector autoregression of the percentage net basis and the difference between the risk-free rate and the change in convenience yield.

Write the return on the commodity from the beginning of period t to the beginning of period $t+1$ as:

$$q_t = (P_{t+1} - P_t + \psi_t)/P_t. \tag{13}$$

Let y_t denote the percentage net basis, i.e. $y_t \equiv \psi_{t-1}/P_t$. Then we can rewrite (13) as:

$$q_t = \psi_t y_t/\psi_{t-1} + \psi_t y_t/\psi_{t-1} y_{t+1} - 1. \tag{14}$$

[4] These constraints are derived as in Campbell and Shiller (1987) as follows. Define $\mathbf{x}_t = [\Delta\psi_t, \dots, \Delta\psi_{t-p+1}, S_t, \dots, S_{t-p+1}]'$. Then (12) can be written in the form $\mathbf{x}_t = \mathbf{A}\mathbf{x}_{t-1} + \mathbf{v}_t$, where $\mathbf{A}$ is a $2p$ by $2p$ matrix. Also, forecasts from this VAR are given by $\mathrm{E}_t\mathbf{x}_{t+k} = \mathbf{A}^k\mathbf{x}_t$. Let $\mathbf{g}$ be a column vector whose $p+1$th element is 1 and whose remaining elements are 0, and let $\mathbf{h}$ be a column vector whose first element is 1 and whose remaining elements are 0. Then from (6), $S_t = \mathbf{g}'\mathbf{x}_t = \Sigma_1^\infty \delta^i \mathbf{h}'\mathbf{A}^i\mathbf{x}_t = \mathbf{h}'\delta\mathbf{A}(\mathbf{I}-\delta\mathbf{A})^{-1}\mathbf{x}_t$. This must hold for any $\mathbf{x}_t$, so $\mathbf{g}'(\mathbf{I}-\delta\mathbf{A}) = \mathbf{h}'\delta\mathbf{A}$, from which the constraints follow.

[5] The percentage net basis is $(1+r) - F_{1,t}/P_t$, but note from (2) that this is just ψ_t/P_t. In what follows, I work with the ratio ψ_{t-1}/P_t.

Now linearise q_t around the sample means $\bar{\psi}$ and $\bar{y}$:

$$q_t \approx y_t(1+1/\bar{y}) + \Delta\psi_t(1+\bar{y})/\bar{\psi} - y_{t+1}/\bar{y}. \tag{15}$$

Finally, define $\beta = 1/(1+\bar{y})$, and define the normalised variables $y'_t = y_t/\bar{y}$, and $\psi'_t = \psi_t/\bar{\psi}$.[6] Then (15) can be rewritten as:

$$\beta q_t \approx y'_t - \beta y'_{t+1} + \Delta\psi'_t. \tag{16}$$

The solution to (16) is a present value relation for the normalised percentage net basis y'_t:

$$y'_t \approx \sum_{j=0}^{\infty} \beta^j (\beta q_{t+j} - \Delta\psi'_{t+j}), \tag{17}$$

i.e. the normalised percentage net basis is approximately the present value of the future stream of returns from holding the commodity net of changes in the normalised convenience yield.

This is simply an approximate accounting relationship, but as Campbell and Shiller (1989) have shown, it can be combined with an economic model for expected returns. I will assume that the expected return is the risk-free rate plus a constant risk premium ρ, i.e. $\mathrm{E}_t\, q_{t+j} = \mathrm{E}_t\, r_{t+j} + \rho$. Then (17) becomes:

$$y'_t \approx \mathrm{E}_t \sum_{j=0}^{\infty} \beta^j (\beta r_{t+j} - \Delta\psi'_{t+j}) + \frac{\beta\rho}{1-\beta}. \tag{18}$$

Equation (18) provides another description of a commodity's price in terms of fundamentals. It says that in an equilibrium where r_t is constant and $\mathrm{E}_t \Delta\psi_{t+j} = 0$ for all j, the expected return on a commodity ($\mu = r + \rho$) equals the percentage net basis y_t. (Note that if $r_t = r$ and $\mathrm{E}_t \Delta\psi_{t+j} = 0$, (18) reduces to $y'_t = \beta\mu/(1-\beta) = \mu/\bar{y}$, or $\mu = y_t$.) In this case, $\mathrm{E}_t \Delta P_{t+1} = 0$ (which also follows from (4) and (6)). Hence unless the discount rate is expected to change, expected price changes are always due to expected changes in convenience yield.

Equation (18) imposes restrictions on the dynamics of the percentage net basis. Define $\phi_t = \beta r_t - \Delta\psi'_t + \beta\rho$, so that $y'_t = \mathrm{E}_t \Sigma_j \beta^j \phi_{t+j}$, and consider the pth-order vector autoregression:

$$y'_t = \sum_{k=1}^{p} \gamma_{11k} y'_{t-k} + \sum_{k=1}^{p} \gamma_{12k} \phi_{t-k-1}, \tag{19a}$$

$$\phi_{t-1} = \sum_{k=1}^{p} \gamma_{21k} y'_{t-k} + \sum_{k=1}^{p} \gamma_{22k} \phi_{t-k-1}. \tag{19b}$$

Then (18) implies the following restrictions: $\gamma_{211} = 1 - \beta\gamma_{111}$, $\gamma_{21k} = -\beta\gamma_{11k}$, $k = 2, \ldots, p$, and $\gamma_{22k} = -\beta\gamma_{12k}$, $k = 1, \ldots, p$. These restrictions are analogous to,

[6] For most commodities, $\bar{y}$ is 1% or less, so β is less than but close to 1. Campbell and Shiller (1989) obtain a present value relation for the log dividend-price ratio on a stock by first writing a log-linear approximation to the stock's log gross return, and then assuming that the ratio of the stock price to the sum of price plus dividend is constant. That ratio (which they denote by ρ) is analogous to β in my model. I work with the arithmetic ratio of convenience yield to price, so the only approximation required is that q_t be linearised around $\bar{\psi}$ and $\bar{y}$.

and are derived in the same way, as the restrictions on the VAR of (12*a*) and (12*b*), and can be used to test (18).

II. THE BEHAVIOUR OF SPOT AND FUTURES PRICES

In this section I discuss the data set and the calculation of the one-month convenience yield ψ_t. I also discuss the behaviour of spot and futures prices, ψ_t, and the spread S_t for the four commodities, and present estimates of the expected excess return ρ and expected total return μ.

Data

All of the tests use futures price data for the first Wednesday of each month. In all cases, that day's settlement price is obtained from the *Wall Street Journal.* Occasionally a contract price will be constrained by exchange-imposed limits on daily price moves. In those cases I use prices for the preceding Tuesday. If those prices are likewise constrained by limits, I use prices for the following Thursday, or if those are constrained, the preceding Monday.

To obtain a spot price P_t, whenever possible I use the price on the spot futures contract, i.e. the contract expiring in month t. Thus the spot and futures prices pertain to exactly the same good, and the time interval between the two delivery dates is known.[7] However, a spot contract does not trade in every month for every commodity. For months when a spot contract does not trade, I inferred a spot price from the nearest active futures contract (i.e. the active contract next to expire, typically a month or two ahead), and the next-to-nearest active contract. This is done by extrapolating the spread between these contracts backwards to the spot month:

$$P_t = F_{1,t}(F_{1,t}/F_{2,t})^{n_{01}/n_{12}}, \tag{19}$$

where $F_{1,t}$ and $F_{2,t}$ are the prices on the nearest and next-to-nearest futures contracts, and n_{01} and n_{12} are, respectively, the number of days between t and the expiration of the nearest contract, and between the nearest and next-to-nearest contract.

This provides spot prices for every month of the year, but errors can arise if the term structure of spreads is very nonlinear. To check that such errors are small, I calculated spot prices using (19) and compared them to actual spot contract prices for copper (available for 200 out of 223 observations), for lumber (114 out of 226 observations), and for gold (173 out of 194 observations). In all three cases, I found little discrepancy between the two series.[8]

[7] Alternatively, one could use data on cash prices, purportedly reflecting actual transactions over the month. But this results in an average price over the month, not a beginning-of-month price. A second and more serious problem is that a cash price can apply to a different grade or specification of the commodity (e.g. copper or gold of a different purity), and can include discounts and premiums that result from longstanding relationships between buyers and sellers.

[8] The RMS percentage error and mean percentage error for the three series are, respectively, 1·21% and −0·12% for copper, 3·99% and 0·39% for lumber, and 3·40% and 0·12% for gold. The simple correlations are 0·998 for copper, 0·983 for lumber, and 0·999 for gold. No spot contract prices were available for heating oil.

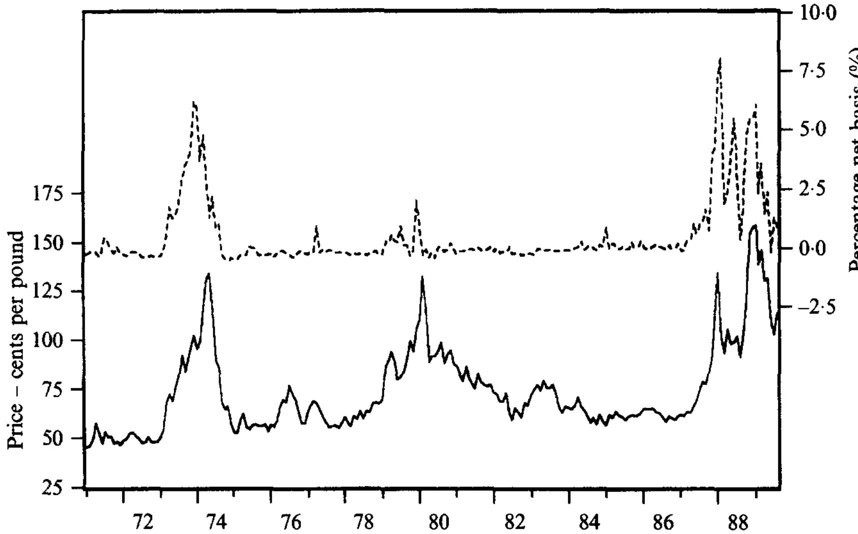

Fig. 1. Copper: price and percentage net basis; (——) price, (-----) percentage net basis.

Given a series for P_t, I then calculate the one-month net marginal convenience yield, ψ_t, using the nearest futures contract and the Treasury bill rate that applies to the same day for which the futures prices are measured. In some cases the nearest futures contract has an horizon greater than one month; I then infer a one-month futures price using the spot contract and the nearest contract if the spot contract exists, or else using the nearest and next-to-nearest contracts. (For example, if in January the nearest futures prices are for March and May and there is no January spot contract, I infer a February price using (19) with $n_{01} = 28$ and $n_{12} = 61$.)

To test the sufficiency of P_t and F_t in forecasting P_{t+1}, I use the following set of variables in the vector $\mathbf{z}_t$: the change in the exchange value of the dollar against ten other currencies, and the growth rates of the Index of Industrial Production, the Index of Industrial Commodity Prices, and the S & P 500 Index. For copper, heating oil, and lumber, $\mathbf{z}_t$ also includes the level and change of monthly U.S. production and inventories of that commodity. All of these variables are measured at the end of the month preceding the date for which prices are measured.

Prices and Convenience Yields

Figs. 1–4 show spot prices and the percentage net basis for each commodity. Note that for copper, heating oil, and lumber, price and convenience yield tend to move together. For example, copper prices rose sharply in 1973, 1979–80, and late 1987 to 1989; each time convenience yield also rose sharply, even as a percentage of price. The same was true when lumber prices rose in early 1973, 1977–9, 1983, and 1986–7. For heating oil the comovement is smaller (and much of it is seasonal), but the percentage net basis still tends to move with price. This suggests that these high prices were expected to be temporary, i.e.

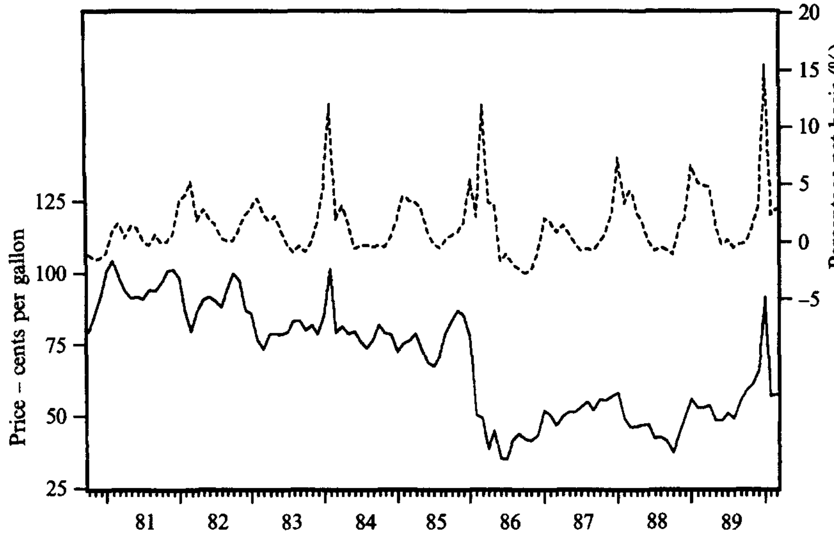

Fig. 2. Heating oil: price and percentage net basis; (——) price, (-----) percentage net basis.

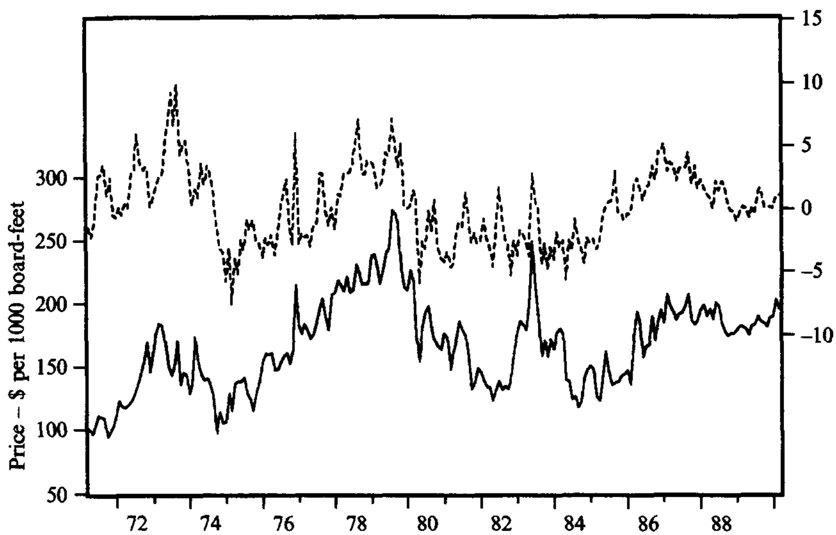

Fig. 3. Lumber: price and percentage net basis; (——) price, (-----) percentage net basis.

price (and convenience yield) would fall as supply and demand adjust towards long-run equilibrium levels.

These figures also show that for these three commodities, convenience yield is a quantitatively important part of the commodity's return. There were periods, for example, when the monthly net convenience yield was 5 to 10 % of price. Hence firms were paying 5–10 % *per month* – plus interest and direct storage costs – to maintain stocks.

Gold is very different. Monthly net convenience yield has always been less

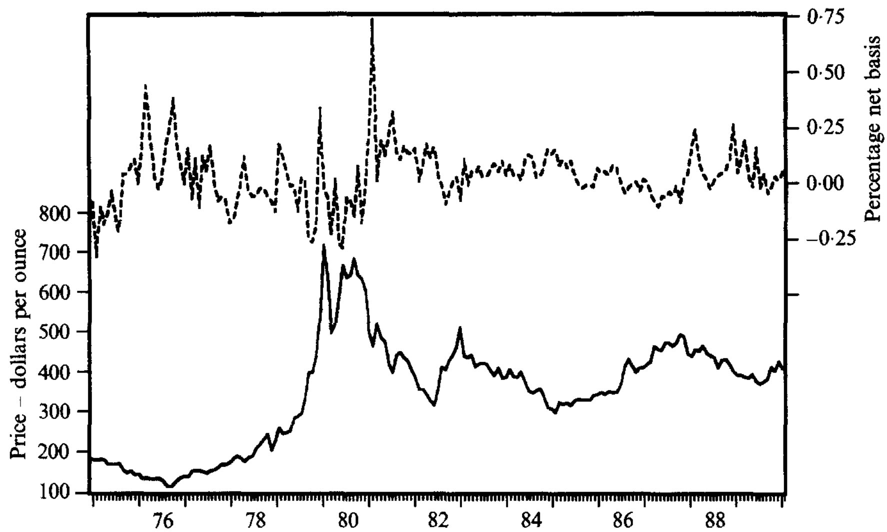

Fig. 4. Gold: price and percentage net basis; (——) price, (-----) percentage net basis.

than 1 % of price, and usually less than 0·2 %. Moreover, except for the brief spike in convenience yield in 1981, there is little comovement with price. This suggests that sharp increases in price (as in 1980 and 1982–3) were expected to persist. This is consistent with the view that the price of gold follows a speculative bubble, or alternatively that it is based on fundamentals and rose because of an expectation that convenience yield would rise in the future.

Table 1 shows results of unit root tests for spot and futures prices, convenience yield, and the spreads S_t and S_t'. (These tests include Δx_{t-1} and Δx_{t-2} on the right-hand side, but no time trend; significance levels are the same with a time trend, or with one or three lags of Δx_t.) Spot and futures prices are integrated of order 1 for all four commodities, and at least for copper, heating oil, and lumber, are clearly cointegrated. The table also shows estimates of the expected monthly excess return, ρ, from the cointegrating regression of $F_{1,t}$ on P_t, and the sample means of ρ. These estimates of ρ are close to the sample means for heating oil and gold, and imply expected annual excess returns of 11 % for heating oil, and -12 % for gold. However, they are unreasonably large for copper and lumber. This casts some doubt on the validity of the model for these two commodities, but may also reflect the relatively poor finite sample properties of the cointegrating estimator. Also, except for gold, S_t is stationary when calculated using either value of ρ. (For gold, we reject a unit root at the 5 % level when S_t is calculated using the cointegrating estimate of ρ, but not when using the sample mean.) These results are consistent with the cointegration of the futures and spot prices, with cointegrating vector $(1 \; \rho - 1)'$.

On the other hand, we strongly reject a unit root in ψ_t for all four commodities, and a regression of P_t on ψ_t yields extremely large estimates of the expected total return μ. Also, when S_t' is computed using the estimated value

Table 1

Unit Root Tests and Estimates of ρ

	Copper (1/71–8/89)	Heating oil (10/80–3/90)	Lumber (4/71–3/90)	Gold (1/75–3/90)
P_t	−2·55	−1·83	−2·92*	−1·55
ΔP_t	−9·21**	−6·89**	−11·87**	−8·24**
F_t	−2·40	−1·86	−2·90*	−1·55
ΔF_t	−8·95**	−6·93**	−11·60**	−8·25**
ψ_t	−4·27**	−5·58**	−3·77**	−5·38**
$\Delta\psi_t$	−12·52**	−7·15**	−11·08**	−13·29**
$S_t(\hat{\rho})$	−4·38**	−5·50**	−3·73**	−2·95*
$S_t(\bar{\rho})$	−4·38**	−5·50**	−3·67**	−2·22
$S'_t(\hat{\mu})$	−2·88*	−1·92	−3·25*	−1·56
$S'_t(\bar{\mu})$	−4·46**	−5·66**	−3·80**	−2·41
$\hat{\rho}$	0·04728	0·00892	0·06100	−0·01080
$\bar{\rho}$	0·00579	0·00926	0·00136	0·00011
$\hat{\mu}$	0·10863	0·53080	0·31263	−0·05962
$\bar{\mu}$	0·01237	0·01673	0·00800	0·00711

Note: Unit root tests are t-statistics on β in the regression $\Delta x_t = \alpha_0 + \alpha_1 \Delta x_{t-1} + \alpha_2 \Delta x_{t-2} + \beta x_{t-1}$. Significance levels are based on MacKinnon's (1990) critical values; * denotes significance at 5% level, ** at 1%. $\hat{\rho}$ is estimate of expected monthly excess return ρ from cointegrating regression: $F_t = \alpha_0 + (1-\rho) P_t$; $\bar{\rho}$ is sample mean. $S_t = F_t - (1-\rho) P_t$. $\hat{\mu}$ is estimate of expected monthly return μ from cointegrating regression $P_t = (1/\mu)\psi_t$; $\bar{\mu}$ is sample mean. $S'_t(\mu) = P_t - (1/\mu)\psi_t$.

of μ, we fail to reject a unit root for two commodities, and reject at only 5% for the other two. This is inconsistent with the present value relation of (1), but may reflect problems of sample size, and is similar to results of Campbell and Shiller (1987) for interest rates and the stock market. Note that if S'_t is computed using the sample mean of μ, we reject a unit root at the 1% level for every commodity but gold. It is likely that either P_t is mean-reverting but at a rate too slow to be detected in samples spanning less than 20 years, or that both P_t and ψ_t are integrated of order 1. (Agbeyegbe (1991) studies prices of pig iron, copper, lead, and zinc using data for 1871–1973, and in each case finds strong evidence of a unit root.) The sample mean of μ is clearly the preferred estimate given its better finite sample properties, and especially given the rejection of a unit root in ψ_t.

III. TEST RESULTS

Tests of the present value relation (1) are based on (9), which implies that S_t and P_t are sufficient to forecast P_{t+1}, on (6), which implies that S_t should Granger-cause $\Delta\psi_t$, and on the cross-equation restrictions on the VAR of (12*a*) and (12*b*). The second and third of these tests require a series for the futures-spot spread, S_t. I calculate S_t first using $\hat{\rho}$ estimated from the cointegrating regression, and then using the sample mean $\bar{\rho}$.

Table 2 shows F-statistics for Wald tests of the restrictions $b_i = 0$ in (10) and (11). These restrictions are never rejected for copper, heating oil, and gold. However, with lumber they are rejected at the 1% level both for (11) and for (10) when S_t is calculated using the sample mean of ρ. This result for lumber

Table 2
Sufficiency of F_t and P_t in Forecasting P_{t+1}

Eqn.		Copper	Heating oil	Lumber	Gold
(1), $\rho=\hat{\rho}$	$F(b_i)$	0·55	1·82	1·97	1·14
	$F(a_i)$	9·80*	8·45**	0·71	2·72
	R^2	0·054	0·404	0·112	0·034
(1), $\rho=\bar{\rho}$	$F(b_i)$	0·65	1·83	2·75**	1·48
	$F(a_i)$	2·22	8·42**	0·34	5·16**
	R^2	0·032	0·404	0·161	0·057
(2)	$F(b_i)$	1·27	1·78	2·92**	1·47
	$F(a_i)$	6·34**	7·17**	1·42	3·14*
	$\tilde{\rho}$	1·226	−0·932	−0·077	3·943
	R^2	0·112	0·431	0·179	0·057

Note: The F-statistics $F(b_i)$ test the restrictions $b_i = 0$ in the regressions (1) $\Delta P_t = a_0 + a_1 S_{t-1}(\rho) + \Sigma_i b_i z_{i,t-1}$, where $S_t = F_t - (1-\rho) P_t$, and (2) $\Delta P_t = a_0 + a_1 F_{t-1} + a_2 P_{t-1} + \Sigma_i b_i z_{i,t-1}$. The statistics $F(a_i)$ test the restrictions $a_0 = 0$ and $a_1 = 1$ in (1) and $a_0 = 0$, $a_1 = 1$, and $a_2 = -(1-\bar{\rho})$ in (2). A * denotes significance at the 5% level, ** at 1%. Also shown for regression (2) is $\tilde{\rho}$, the value of ρ implied by the estimate of a_2. For all commodities, $\mathbf{z}_t$ includes the change in the exchange value of the dollar against ten other currencies, and the growth rates of the Index of Industrial Production, the Index of Industrial Materials Prices, and the S & P 500 Index. For copper, heating oil, and lumber, $\mathbf{z}_t$ also includes the level and change of monthly U.S. production and inventories of that commodity.

could reflect a failure of (1), or else inefficiency in the futures market. (The latter is more plausible for lumber than the other commodities because lumber futures are more thinly traded.) The table also shows F-statistics for tests of $a_0 = 0$ and $a_1 = 1$ in (10), and $a_0 = 0$, $a_1 = 1$, and $a_2 = -(1-\bar{\rho})$ in (11), as well as the value of ρ implied by $\hat{a}_2$. These restrictions are generally rejected, and except for lumber, the implied ρs are unrealistic. This is not a rejection of the present value model (which only requires that S_{t-1} be a sufficient predictor of ΔP_t), but shows that with 10–20 years of data, 'structural' parameters cannot be recovered from these equations.

Table 3 shows the results of Granger causality tests between S_t and $\Delta\psi_t$. Equation (9) implies unidirectional causality from S_t to $\Delta\psi_t$, i.e. that we can reject the hypothesis that S_t does not cause $\Delta\psi_t$, but fail to reject the hypothesis that $\Delta\psi_t$ does not cause S_t. The Akaike Information Criterion for the number of lags was inconclusive; the AIC (and FPE) suggest between 2 and 8 lags, but are fairly flat within this range. Hence I report results for 2, 4, 6, and 8 lags.

For copper, heating oil, and lumber these results are consistent with the present value model; in each case we can clearly reject the hypothesis that S_t does not cause $\Delta\psi_t$. Furthermore, for heating oil and lumber, the causality is unidirectional; we fail to reject the noncausality of $\Delta\psi_t$ to S_t. For gold, the results are more ambiguous. We reject the hypothesis that S_t does not cause $\Delta\psi_t$ with 4, 6, or 8 lags, but not with 2 lags. Also, with any number of lags there is always a much stronger rejection of the hypothesis that $\Delta\psi_t$ does not cause S_t.

Table 3 also shows chi-square statistics for Wald tests of the cross-equation restrictions implied by (1) on the vector autoregression of S_t and $\Delta\psi_t$. (The results shown are for a 4th-order VAR, but are qualitatively the same for 2nd-

Table 3
Causality Tests and Tests of VAR Restrictions

Lags	H_0	Copper	Heating oil	Lumber	Gold
		(A) *Causality Tests*			
2	$S \nrightarrow \Delta\psi$	13·90**	14·89**	10·39**	1·88
	$\Delta\psi \nrightarrow S$	2·42	1·59	0·81	5·66**
4	$S \nrightarrow \Delta\psi$	7·55**	9·04**	4·77**	2·83*
	$\Delta\psi \nrightarrow S$	4·07**	1·02	0·33	8·41**
6	$S \nrightarrow \Delta\psi$	6·32**	5·00**	2·97**	3·74**
	$\Delta\psi \nrightarrow S$	3·39**	0·26	0·27	12·53**
8	$S \nrightarrow \Delta\psi$	4·93**	2·23*	3·00**	2·97**
	$\Delta\psi \nrightarrow S$	4·02**	0·44	0·96	9·88**
		(B) *Tests of Restrictions on VAR*			
4	Restrictions on VAR of $S(\hat{\rho})$ and $\Delta\psi$	40·58**	12·29	40·74**	67·54**
4	Restrictions on VAR of $S(\bar{\rho})$ and $\Delta\psi$	65·03**	12·85	27·61**	43·27**

Note: (A) In causality tests denoted $y \nrightarrow x$, F-statistics are shown for tests of restrictions $b_i = 0$ in regressions of $x_t = a_0 + \Sigma_i a_i x_{t-i} + \Sigma_i b_i y_{t-i}$. S is computed using $\hat{\rho}$ from cointegrating regression. (Results are qualitatively the same when $\bar{\rho}$ is used.) (B) χ^2 statistics are shown for Wald tests of restrictions on 4-period VAR of $S_t(\rho)$ and $\Delta\psi_t$. A * denotes significance at 5% level, ** at 1%.

order and 6th-order VARs.) These restrictions are strongly rejected for copper, lumber, and gold, irrespective of whether $\hat{\rho}$ or $\bar{\rho}$ is used to calculate S_t. The restrictions are accepted, however, for heating oil.

These results provide mixed evidence on the ability of the present value model to explain commodity prices. The model fits the data well for heating oil. In fact, as Figs. 5 and 6 show, the unrestricted VAR of (12*a*) and (12*b*) predicts

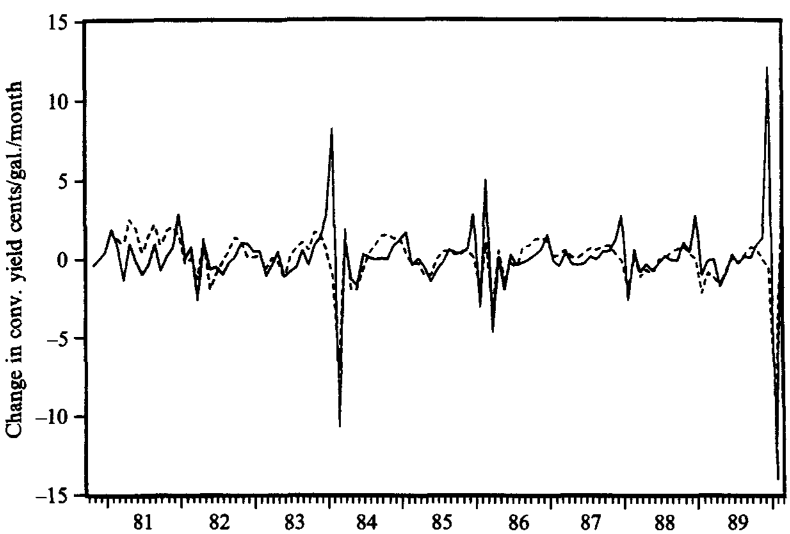

Fig. 5. Heating oil: change in convenience yield; fitted *vs.* actual; (——) actual, (-----) fitted.

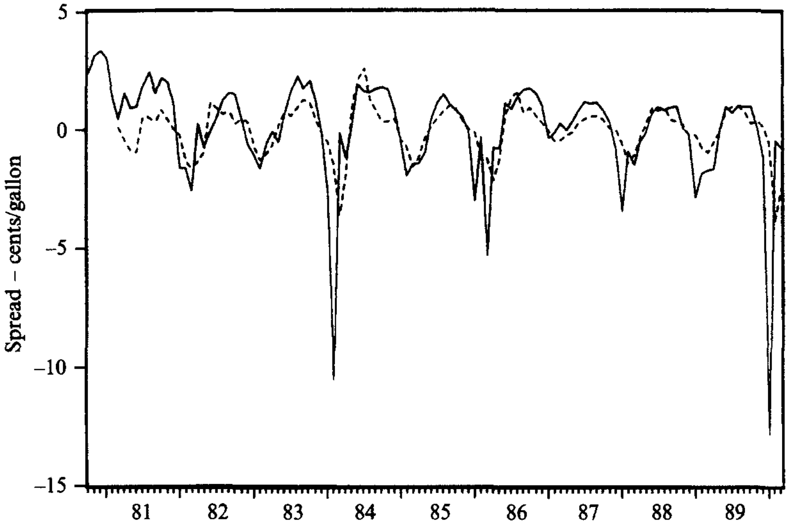

Fig. 6. Heating oil: future-spot spread; (——) actual, (-----) fitted.

monthly changes in convenience yield, $\Delta\psi_t$, and the spread, S_t, reasonably well. Some of the model's implications, however, are rejected by the data for copper and lumber. This may be because on average, convenience yield is a larger percentage of price for heating oil than for the other commodities. Hence price movements for heating oil will be tied more closely to expected near-term changes in convenience yield, rather than changes that might occur in the more distant future.

The strongest rejections are for gold; it is not even clear that futures and spot prices are cointegrated, and there is no evidence that the spot price and convenience yield are co-integrated. But if the present value model holds for gold, investors must believe that there is always a small probability that convenience yield will rise sharply. Throughout the 15 year sample, gold's convenience yield has been very small relative to price, so the present value model can only explain price movements in terms of changing market perceptions of either the mean arrival rate of an event, or the probability distribution for its size. Since such changes in perceptions are unobservable and do not affect current convenience yields, these results are not surprising.

Table 4 shows statistics for the percentage net basis, $y_t = \psi_{t-1}/P_t$, and the variable $\phi_t = \beta r_t - \Delta\psi'_t + \beta\rho$. Note that $\bar{y}$ is largest for heating oil (about 1·5% per month), and extremely small for gold. Also, we can clearly reject a unit root for both y_t and ϕ_t. The table also shows chi-square statistics for Wald tests of the cross-equation restrictions imposed by the present value relation (18) on the 4th-order VAR of $y'_t = y_t/\bar{y}$ and ϕ_{t-1}. (Again, results are qualitatively the same for 2nd- and 6th-order VARs.) These restrictions are strongly rejected for all four commodities. This is not a rejection of (1), but is troubling because it can be viewed as a rejection of a constant risk premium (recall that (18) was derived by assuming that the expected return $E_t q_{t+j} = E_t r_{t+j} + \rho$), and (1)

Table 4
Behaviour of Percentage Net Basis

	Copper	Heating oil	Lumber	Gold
		(A) *Basic Statistics*		
$\bar{\psi}$	0·5972	1·0807	0·6850	0·0873
$\bar{P}$	73·488	69·464	166·64	341·51
$\bar{y} = \overline{\psi_{-1}/P}$	0·00493	0·01468	0·00153	0·00030
$\beta = 1/(1+\bar{y})$	0·9951	0·9855	0·9985	0·9997
$\bar{\phi}$	0·00436	−0·00439	−0·01516	−0·01508
		(B) *Unit Root Tests*		
y_t	−3·87**	−4·81**	−3·76**	−5·44**
ϕ_t	−12·52**	−7·15**	−11·08**	−13·29**
		(C) *Tests of VAR Restrictions*		
$\chi^2(8)$	130·94**	95·69**	62·01**	96·48**

Note: $y_t = \psi_{t-1}/P$, $\phi_t = \beta r_t - \Delta\psi_t' + \beta\rho$, and $\psi_t' = \psi_t/\bar{\psi}$. For unit root test, see note to Table 1. χ^2 statistics are Wald tests of restrictions on 4-period VAR of y_t' and ϕ_{t-1}.

includes a constant discount rate. Alternatively, this result could be a rejection of the linear approximation of (15) used to derive (18).

IV. SERIAL CORRELATION OF EXCESS RETURNS

Given these results, I examine the serial correlation of excess returns as an additional test of market efficiency. This is useful because we can look for patterns of results across commodities that are similar to the results above for the present value model.

Cutler *et al.* (1991) found serial correlation of excess returns that is positive in the short run and negative in the long run for a broad range of assets that included gold, silver, and an index of industrial metals. However, they ignored convenience yield when measuring returns, which can lead to measurement errors. I calculate autocorrelations for excess returns that include convenience yields, and are measured relative to the three-month Treasury bill rate. I follow Cutler *et al.* and examine both individual autocorrelations and the averages of autocorrelations 1–12, 13–24, 25–36, and 37–48.

Table 5 shows autocorrelations of excess returns (corrected for small sample bias by adding $1/(T-j)$ to the jth correlation, where T is the sample size), and Box-Pierce Q statistics that test the significance of the first K autocorrelations. Observe that we can reject a non-zero first-order autocorrelation at the 5% level for copper and gold. Of the first 12 autocorrelations, 2 are significant at the 5% level for copper, four are for lumber, and five are for gold. Gold exhibits the greatest serial dependence of returns; in addition to large individual autocorrelations, all of the Q statistics are highly significant.[9] Copper and

[9] I find much greater serial dependence in excess returns for gold than do Cutler *et al.* Their estimate of ρ_1, for example, is only 0·020. However, their sample period is 1974 to 1988, while mine is 1975 to the first three months of 1990.

Table 5
Autocorrelations of Excess Returns

Autocorrel.	Copper	Heating oil	Lumber	Gold
ρ_1	0·192	−0·046	0·090	0·182
ρ_2	0·037	−0·055	−0·019	−0·155
ρ_3	0·058	−0·062	−0·030	0·021
ρ_4	0·057	0·019	0·054	0·197
ρ_5	0·080	−0·035	0·176	0·255
ρ_6	0·053	−0·044	0·160	−0·057
ρ_7	−0·094	−0·061	0·053	0·022
ρ_8	−0·013	0·088	0·054	0·146
ρ_9	0·034	0·062	0·044	−0·023
ρ_{10}	0·128	0·037	0·180	0·033
ρ_{11}	0·141	−0·053	0·143	0·147
ρ_{12}	0·102	0·071	0·092	0·064
ρ_{1-12}	0·065	−0·007	0·083	0·069
ρ_{13-24}	−0·023	0·011	0·001	−0·033
ρ_{25-36}	0·022	0·007	−0·019	0·005
ρ_{37-48}	−0·012	0·002	−0·006	−0·023
s.e.(ρ_j)	0·067	0·094	0·066	0·074
Q(12)	23·04	4·49	29·03	37·16
	(P = 0·027)	(P = 0·973)	(P = 0·004)	(P < 0 ·001)
Q(24)	36·13	11·11	40·74	59·89
	(P = 0·053)	(P = 0·988)	(P = 0·018)	(P < 0·001)
Q(48)	50·11	33·71	60·78	98·94
	(P = 0·390)	(P = 0·941)	(P = 0·102)	(P < 0·001)

Note: Autocorrelations ρ_j are bias-corrected by adding $1/(T-j)$. ρ_{1-12} is the average of the first 12 autocorrelations, ρ_{13-24} is the average of the next 12, etc. Q(K) is the Box-Pierce Q statistic for the first K autocorrelations and P is the associated probability level.

lumber show weaker evidence of serial dependence. Fewer individual autocorrelations are significant (especially for copper), and Q is significant for the first 12 or 24 autocorrelations, but not the first 48. Also, for all three commodities the serial dependence is positive for short horizons, but negative for longer horizons. This is similar to results of Fama and French (1988) for stock returns, and is consistent with the notion that prices temporarily drift away from fundamentals.

For heating oil, however, there is no serial dependence of returns. Every individual autocorrelation is within one standard deviation of zero, and the probability levels for the three Q statistics are all above 0·9. This pattern across commodities parallels that in the previous section for tests of the present value model; the strongest rejections were for gold, results for copper and lumber were mixed, but heating oil closely conformed to the model.

V. CONCLUSIONS

The present value model of rational commodity pricing can be viewed as a highly reduced form of a dynamic supply and demand model, and when the commodity is traded on a futures market, it can be tested through the constraints it imposes on the joint dynamics of spot and futures prices. I found a close conformance to the model for heating oil, but not for copper or lumber, and especially not for gold. The pattern is the same when one looks at the serial dependence of excess returns. For three of the four commodities, these results are consistent with the notion that prices temporarily drift away from fundamentals, perhaps because of 'fads'.

Earlier studies provide different evidence that commodity prices are not always based on fundamentals. For example, Roll (1984) found that only a small fraction the price movements for frozen orange juice can be explained by 'fundamentals', i.e. by variables such as the weather that in principle should explain a good deal of the variation in price. And Pindyck and Rotemberg (1990*a*) found high levels of unexplained price correlation across commodities that is also inconsistent with prices following fundamentals. However, both the Roll and Pindyck and Rotemberg results may be suspect because of the possibility that one or more key variables (that affect orange juice supply or demand, or supplies or demands for a broad range of commodities) have been omitted. The present value model, on the other hand, is based entirely on a payoff stream that can be measured from futures market data. The rejections of some of the implications of that model (together with the finding of serially dependent returns) provides additional evidence that the prices of some commodities may be partly driven by fads.

Heating oil prices, however, conform closely to the present value model, and there is no evidence of serial dependence in excess returns. Why does heating oil seem to differ from the other commodities in this respect? It may be that its high average convenience yield makes speculation too costly. A long position in heating oil costs 1·5 % per month on average in convenience yield; the odds for a speculator are much more favourable for other commodities. This result may also reflect the relatively large amount of (partly predictable) short-run variation in the convenience yield of heating oil, and (compared to lumber and copper) the relatively more active futures market.

Massachusetts Institute of Technology

Date of receipt of final typescript: August 1992

References

Agbeyegbe, Terence D. (1991). 'The stochastic behaviour of mineral-commodity prices.' In *Models, Methods, and Applications of Econometrics* (eds. P. C. B. Phillips and V. B. Hall). Basil Blackwell.

Campbell, John Y. and Shiller, Robert J. (1987). 'Cointegration and tests of present value models.' *Journal of Political Economy*. vol. 95, pp. 1062–88.

Campbell, John Y. and Shiller, Robert J. (1989). 'The dividend-price ratio and expectations of future dividends and discount factors.' *Review of Financial Studies*, vol. 1, pp. 195–228.

Cutler, David M., Poterba, James M. and Summers, Lawrence H. (1991). 'Speculative dynamics.' *Review of Economic Studies*, vol. 58, pp. 529–46.

Deaton, Angus and Laroque, Guy (1992). 'On the behavior of commodity prices.' *Review of Economic Studies*, vol. 59, pp. 1–23.

Fama, Eugene F. and French, Kenneth R. (1988). 'Permanent and transitory components of stock prices.' *Journal of Political Economy*, vol. 96, pp. 246–73.

French, Kenneth R. (1983). 'A comparison of futures and forward prices.' *Journal of Financial Economics*, vol. 12, pp. 311–42.

Kahn, James A. (1991). 'Why is production more volatile than sales? theory and evidence on the stockout-avoidance motive for inventory holding.' *Quarterly Journal of Economics* (forthcoming).

MacKinnon, James G. (1990). 'Critical values for cointegration tests.' In *Long-run Economic Relationships: Readings in Cointegration* (eds. R. F. Engle and C. W. Granger). Oxford: Oxford University Press.

Mankiw, N. Gregory, Romer, David and Shapiro, Matthew (1991). 'Stock market forecastability and volatility: a statistical appraisal.' *Review of Economic Studies*, vol. 58, pp. 455–78.

Pindyck, Robert S. (1990). 'Inventories and the short-run dynamics of commodity markets.' NBER Working Paper No. 3295.

Pindyck, Robert S. and Rotemberg, Julio J. (1990*a*). 'The excess comovement of commodity prices.' ECONOMIC JOURNAL, vol. 100, pp. 1173–89.

Pindyck, Robert S. and Rotemberg, Julio J. (1990*b*). 'The comovement of stock prices.' MIT Sloan School of Management Working Paper.

Roll, Richard (1984). 'Orange juice and the weather.' *American Economic Review*, vol. 74, pp. 861–80.