NBER WORKING PAPER SERIES

INVENTORIES AND THE SHORT-RUN DYNAMICS OF COMMODITY PRICES

Robert S. Pindyck

Working Paper No. 3295

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 1990

This research was supported by MIT's Center for Energy Policy Research, and by the National Science Foundation under Grant No. SES-8618502. I am grateful to Columbia University's Center for the Study of Futures Markets and the Commodity Research Corporation in Princeton, NJ, for providing futures market data. My thanks to Patricia Craig and John Simpson for their superb research assistance, and to Michael Brennan, Zvi Eckstein, Jeffrey Miron, Ariel Pakes, Julio Rotemberg, Menachem Sternberg, and seminar participants at MIT, Yale, and the NBER for helpful comments. This paper is part of NBER's research program in Financial Markets and Monetary Economics. Any opinions expressed are those of the author and not those of the National Bureau of Economic Research.

NBER Working Paper #3295
March 1990

INVENTORIES AND THE SHORT-RUN DYNAMICS OF COMMODITY PRICES

ABSTRACT

I examine the behavior of inventories and their role in the short-run dynamics of commodity production and price. Competitive producers of a storable commodity react to price changes by balancing costs of changing production with costs of changing inventory holdings. To determine these costs, I estimate a structural model of production, sales, and storage for copper, heating oil, and lumber. I than examine the implications of these costs for inventory behavior, and for the behavior of spot and futures prices. I find that inventories may serve to smooth production during periods of low or normal prices, but during periods of temporarily high prices inventories have a more important role in facilitating production and delivery scheduling and avoiding stockouts.

This paper differs from earlier studies of inventory behavior in three respects. First, I focus on homogeneous and highly fungible commodities. This helps avoid aggregation problems, simplifies the meaning of marginal convenience yield, and allows the use of direct measures of units produced, rather than inferences from dollar sales. Second, I estimate Euler equations, and allow marginal convenience yield to be a convex function of inventories. This is more realistic, and better explains the value of storage and its role in the dynamics of price. Third, I use futures prices to directly measure marginal convenience yield. This produces tighter estimates of the parameters of the convenience yield function.

Robert S. Pindyck
MIT
Room E52-454
50 Memorial Drive
Cambridge, MA 02139

1. Introduction.

The markets for many commodities are characterized by periods of sharp changes in prices and inventory levels.[1] This paper examines the behavior of inventories, and their role in the short-run dynamics of production and price. It also seeks to determine whether observed fluctuations in spot and futures prices can be explained in terms of rigidities in production and desired inventory holdings.

In a competitive market for a storable commodity, producers and consumers react to stochastic price fluctuations by balancing costs of adjusting consumption and production with costs of increasing or decreasing inventory holdings. These costs affect the extent to which prices fluctuate in the short run. For example, if adjustment costs are small and there are no substantial short-run diseconomies of scale, there will be little need to adjust inventories, and shocks to demand or supply that are expected to be short lived will have little impact on price. Such shocks will likewise have little impact if adjustment costs are high but the cost of drawing down inventories is low.

To determine these costs, I model the short-run dynamics of production, sales, and storage for three commodities: copper, heating oil, and lumber. I assume optimizing behavior on the part of producers, and estimate structural parameters that measure adjustment costs, costs of producing, and costs of drawing down inventories. I then examine the implications of these costs for inventory behavior, and for the behavior of spot and futures prices.

Because of its importance in the business cycle, inventory behavior in manufacturing industries has been studied extensively. Recent work has shown that there is little support for the traditional production smoothing model of inventories; in fact, the variance of production generally exceeds the variance of sales in manufacturing.[2] Instead, the evidence seems to favor models of production cost smoothing, in which inventories are used to shift production to periods in which costs are low, and models in which inventories are used to avoid stockouts and reduce scheduling costs.[3]

The role of inventories seems to be different in commodity markets. At least for two of the three commodities studied here, the variance of production is substantially less than the variance of sales, and there is evidence that inventories are used to smooth production. On the other hand, the cost of drawing down inventories rises rapidly as inventory levels fall, because inventories are also needed to maintain production and delivery schedules. This limits their use for production or production cost smoothing, particulary during periods of high prices following shocks.

Besides their focus on manufactured goods, most earlier studies of inventory behavior and short-run production dynamics rely on a linear-quadratic model to obtain an analytical solution to the firm's optimization problem. Examples include Eichenbaum's (1984, 1989) studies of finished goods inventories, the studies of automobile production and inventories by Blanchard (1983) and Blanchard and Melino (1986), and Eckstein and Eichenbaum's (1985) study of crude oil inventories. All of these models include a target level of inventory (proportional to current or anticipated next-period sales) and a quadratic cost of deviating from that level.[4]

Although convenient, the linear-quadratic specification is a major limitation of these models. First, marginal production cost might not be linear. Second and more important, a quadratic cost of deviating from a target inventory level implies that the net marginal convenience yield from holding inventory (the negative of the cost per period of having one less unit of inventory) is linear in the stock of inventory. Aside from permitting negative inventories, this is simply a bad approximation. Early studies of the theory of storage have demonstrated,[5] and a cursory look at the data in the next section confirms that at least for commodities, marginal convenience yield is a highly convex function of the stock of inventory, rising rapidly as the stock approaches zero, and remaining close to zero over a wide range of moderate to high stocks. There is no reason to expect a linear approximation to be any better for manufactured goods.

The alternative approach to the study of inventory behavior and production dynamics is to abandon the linear-quadratic framework, adopt a more general specification, and estimate the Euler equations (i.e., first-order conditions) that follow from intertemporal optimization. This approach is used in recent studies of manufacturing inventories by Miron and Zeldes (1988), who show that the data strongly reject a general model of production smoothing that takes into account unobservable cost shocks and seasonal fluctuations in sales, and Ramey (1988), who, by introducing a cubic cost function, shows that declining marginal cost may help explain the excess volatility of production. However, both of these studies maintain the assumption that the cost of deviating from the target inventory level is quadratic, so that the marginal convenience yield function is linear.

This study differs from earlier ones in three major respects. First, I focus on homogeneous and highly fungible commodities. This helps avoid aggregation problems, and simplifies the meaning of marginal convenience yield and its role in the dynamics of production and price. Also, it allows me to use direct measures of units produced, rather than inferences from dollar sales and inventories.[6] Second, as in Miron and Zeldes and Ramey, I estimate Euler equations, but I allow the marginal convenience yield to be a convex function of the stock of inventory. This is much more realistic, and better explains the value of storage and its role in the short-run dynamics of price. Third, I utilize futures market data, together with a simple arbitrage relation, to obtain a direct measure of marginal convenience yield. This makes it possible to obtain tighter estimates of the parameters of the marginal convenience yield function.[7]

The next section discusses the characteristics and measurement of marginal convenience yield, presents basic data for the three commodities, and explores the behavior of price, production and inventories. Section 3 lays out the model, and Section 4 discusses the data and estimation method. Estimation results are presented in Section 5, and Section 6 concludes.

2. Spot Prices, Futures Prices, and the Value of Storage.

This section shows how the marginal value of storage can be inferred from futures prices, and discusses its dependence on inventories and sales, and its role in short-run price formation. I also review the general behavior of price, inventories, and the marginal value of storage, together with some summary statistics, for each of the three commodities.

I will assume, as have most studies, that firms have a target level of inventory which is a function of expected next-period sales. An inventory level below this target increases the risk of stock-outs and makes it more difficult to schedule and manage production and shipments, thereby imposing a cost on the firm. Furthermore, this cost can rise dramatically as inventories fall substantially below the target level.[8] I will also assume that there is a (constant) cost of physical storage of a dollars per unit per period. Letting N_t and P_t denote the inventory level and price, and Q_{t+1} denote expected next-period sales, I can therefore write the total per period cost of storage at time t as $\Phi(N_t, Q_{t+1}, P_t) + aN_t$, with $\Phi_N < 0$, $\Phi_{NN} > 0$, $\Phi(0,Q,P) = \infty$, $\Phi \to 0$ for N large, and Φ_Q, $\Phi_P > 0$.

The marginal cost of storage is $\Phi_N(N_t, Q_{t+1}, P_t) + a$, so the cost of drawing down inventories by one unit is $-\Phi_N - a$. Thus $-\Phi_N - a$ is the net (of storage cost) per period marginal value of storage. The term $-\Phi_N$ is more commonly referred to as the marginal convenience yield from storage, and $\psi = -\Phi_N - a$ as the net marginal convenience yield.[9]

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

$$\tilde{\psi}_{t,T} = (1+r_T)P_t - f_{T,t} \qquad (1)$$

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 (1) must hold, note that the (stochastic) return from holding a unit of the commodity from t to t+T is $\tilde{\psi}_{t,T} + (P_{t+T} - P_t)$. If one also shorts a forward contract at time t, one receives a total return by the end of the period of $\tilde{\psi}_{t,T} + f_{T,t} - P_t$. 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 (1) follows.

In keeping with the literature on inventory behavior (see the references in Footnotes 2, 3, and 4), we will work with the net marginal convenience yield valued at time t. Denote this by $\psi_{t,T} = \tilde{\psi}_{t,T}/(1 + r_T)$, so that (1) becomes:[10]

$$(1 + r_T)\psi_{t,T} = (1 + r_T)P_t - f_{T,t} \qquad (1')$$

For most commodities, futures contracts are much more actively traded than forward contracts, and good 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 corresponding 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 <u>and</u> is positively (negatively) correlated with the spot price.[11] However, for most commodities the difference in the two prices is very small. In the Appendix, I estimate this difference for each commodity, using the sample variances and covariance of the interest rate and futures price, and I show that it is negligible.[12] I therefore use the futures price, $F_{T,t}$, in place of the forward price in eqn. (1').

Figures 1a, 1b, and 1c show spot prices for copper, lumber, and heating oil, together with the one-month net marginal convenience yield, $\psi_t = \psi_{t,1}$. (My data for copper and lumber extend from October 1972 through December 1987. Futures contracts for heating oil only began trading in late 1978, so data for this commodity cover November 1978 to June 1988. A discussion of the data and construction of ψ_t appears in Section 4.) Observe that price and convenience yield tend to move together. For example, there were three periods in which copper prices rose sharply: 1973, 1979-80, and the end of 1987. On each occasion (and especially the first and third), the convenience yield also rose sharply. Likewise, when lumber prices rose in early 1973, 1977-79, 1983, and 1986-87, the net convenience yield also rose. For heating oil the co-movement is smaller (and much of what there is is seasonal), but there has still been a tendency for price and convenience yield to move together.

These figures also show that firms are willing to hold inventories at substantial cost. In December 1987, the net convenience yield for copper was about 10 cents per pound per month - about 8 percent of the price. This means that firms were paying 8 percent <u>per month</u> - plus interest and direct storage costs - in order to maintain stocks.[13] The net convenience yield for lumber and heating oil also reached peaks of 8 to 10 percent of price. During these periods of high prices and high convenience yields, inventory levels were lower than normal, but still substantial. This suggests that production is rigid in the short run, and cannot be adjusted quickly in response to higher prices. But it also suggests that an important function of inventories is to avoid stockouts and facilitate the scheduling of production and sales. This function probably dominates during periods when prices are high and inventory levels are low. During more normal period, inventories may also serve to smooth production.

Table 1 compares the variances of detrended production, sales, and inventories. The first row shows the ratio of the variance of production to the variance of sales for each commodity. For copper and heating oil, the variance of production is much less than that of sales. One explanation is that demand shocks tend to be larger and more frequent than cost shocks. One might expect this to be the case with heating oil, where seasonal fluctuations in demand are considerable, and to a lesser extent for lumber. The second row shows the ratios of the nonseasonal components of the variances (obtained by first regressing each variable against a set of monthly dummies and time). As expected, this ratio is much larger for heating oil, and slightly larger for lumber. However, for copper and heating oil, the variance of production still exceeds that of sales.

Also shown in Table 1 is the ratio of the variance of production to that of inventories, normalized by the squared means. Note that for copper and heating oil, there is much more variation in inventory than in production, whether or not the variables have been deseasonalized. This suggests that for these two commodities, one important use of inventories is to smooth production. The picture is somewhat different, however, for lumber. The variances of production and sales are about the same, and production varies much more than inventories, especially after deseasonalizing the variables. Also, production and sales track each other very closely. Hence production smoothing is probably not an important role for inventories of lumber. Instead, large maintenance levels of inventories seem to be needed to maintain scheduling and avoid stockouts.

Finally, what do the data tell us about the dependence of the marginal convenience yield on the level of inventories? One would expect the marginal value of storage to be proportional to the price of the commodity, and to depend on anticipated sales. In the model presented in the next section, I use the following functional form for ψ_t, which is reasonably general but easy to estimate:[14]

$$\psi_t = \beta P_t(N_t/Q_{t+1})^{-\phi} - a \qquad (2)$$

Figures 2a, 2b, and 2c show ψ_t plotted against the inventory-sales ratio, N_t/Q_{t+1}, for each commodity. These figures suggest that ψ_t is likely to be well represented by eq. (2), with $\beta,\phi > 0$, and that the linear relationship that has been used in most studies of inventories is likely to be a poor approximation to what is in fact a highly convex function.[15]

Table 2 shows simple nonlinear least squares estimates of eqn. (2), with monthly dummy variables included for a. (These monthly dummies can capture seasonal shifts the cost of storage, as well as seasonal shifts in the gross marginal convenience yield.) For all three commodities, the fit is good, and we can easily reject $\phi = -1$, i.e., that ψ is linear in N. In addition, the monthly dummy variables are significant as a group for lumber and heating oil. As expected, there are strong seasonal fluctuations in the use of these two commodities, so that the benefit from holding inventory is likewise seasonal.

3. The Model.

Intertemporal optimization by producers requires balancing three costs: the cost of producing itself, which may vary with the level of output and over time as factor costs change; the cost of changing production, i.e., adjustment cost; and the cost of drawing down inventories. Our objective is to estimate all three of these costs, and determine their dependence on output, sales, and inventory levels. To do this, I make use of fact that in the U.S. markets for copper, heating oil, and lumber, producers can be viewed as price takers. This, together with the fact that futures prices provide a direct measure of the marginal value of storage, allows me to estimate absolute costs, rather than relative ones as in other studies (such as those of Blanchard (1983), Ramey (1988), and Miron and Zeldes (1988)).

I model the direct cost of production as quadratic in output, and I assume that there is a quadratic cost of adjusting output. For all three commodities, most inventories (and all of the inventories included in our data) are held by producers. Hence $\Phi(N_t, Q_{t+1}, P_t)$ is the cost that the firm bears from production and scheduling inefficiencies, stockouts, etc., when the inventory level is N_t and expected sales is Q_{t+1}. (This excludes physical storage costs.) Both production cost and the net benefit from storage are likely to fluctuate seasonally, so I introduce monthly dummy variables to account for each. Allowing for unobservable shocks to the cost of storage, the total cost of production can be written as:[16]

$$C_t = (c_0 + \sum_{j=1}^{11} c_j D_{jt} + \sum_{j=1}^{m} \gamma_j w_{jt} + \eta_t)y_t + (1/2)by_t^2 + (1/2)\beta_1(\Delta y_t)^2 +$$

$$\Phi(N_t, Q_{t+1}, P_t) + (a_0 + \sum_{j=1}^{11} a_j D_{jt} + \nu_t)N_t \qquad (3)$$

Here, the γ_{jt}'s are a set of factor prices; a wage index and a materials cost index for all three commodities, and in addition the price of crude oil for heating oil. These γ_{jt}'s and the error term η_t allow for both observable and unobservable cost shocks.

Inventories must satisfy the following accounting identity:

$$N_t = N_{t-1} + y_t - Q_t \qquad (4)$$

Taking price as given, firms must find a contingency plan for production and sales that maximizes the present value of the flow of expected profits, subject to eqn. (4):

$$\max E_t \sum_{\tau=t}^{\infty} R_{t,\tau}(P_\tau Q_\tau - C_\tau) \qquad (5)$$

where E_t denotes the expectation conditional on information available at t, and $R_{t,\tau}$ is the τ-period discount factor at time t. All prices and costs in this model are in nominal terms, so $R_{t,\tau} = 1/(1 + r_{t,\tau})$, where $r_{t,\tau}$ is the τ-period nominal interest rate at t. The maximization is subject to the additional constraint that $N_\tau \geq 0$ for all τ, but because $\Phi \to \infty$ as $N \to 0$, this constraint will never be binding.

To obtain first-order conditions for this problem, use eqn. (4) to eliminate y_τ, and then maximize with respect to Q_t and N_t. Maximizing with respect to Q_t yields:

$$P_t = c_0 + \sum_{j=1}^{11} c_j D_{jt} + \sum_{j=1}^{m} \gamma_j w_{jt} - by_t + \beta_1(\Delta y_t - R_{1t}E_t\Delta y_{t+1}) + \eta_t \qquad (6)$$

Maximizing with respect to N_t and using $\Phi_N = -\beta P_t(N_t/Q_{t+1})^{-\phi}$ yields:

$$0 = c_0(1 - R_{1t}) + \sum_{j=1}^{11} c_j(D_{jt} - R_{1t}D_{j,t+1}) + E_t\Big[\sum_{j=1}^{m} \gamma_j(w_{jt} - R_{1t}w_{j,t+1})$$

$$+ b(y_t - R_{1t}y_{t+1}) + \beta_1(\Delta y_t - 2R_{1t}\Delta y_{t+1} + R_{2t}\Delta y_{t+2})$$

$$+ a_0 + \sum_{j=1}^{11} a_j D_{jt} - \beta P_t(N_t/Q_{t+1})^{-\phi}\Big] + \eta_t - R_{1t}E_t\eta_{t+1} + \nu_t \qquad (7)$$

Eqn. (6) equates price with full marginal cost, where the latter includes the effect of producing an extra unit today on current and discounted expected future adjustment costs. Perturbing an optimal production plan by increasing this period's output by one unit (holding N fixed) increases the current cost of adjustment (by $\beta_1\Delta y_t$), but reduces the expected cost of adjustment next period (by $\beta_1 E_t\Delta y_{t+1}$). The equation also contains an error term, but note that this is not an expectational error; it simply represents the unexplained part of marginal cost.

Eqn. (7) describes the tradeoff between selling out of inventory versus producing, holding Q fixed. To see this, move $a_0 + \Sigma_j a_j D_{jt} - \beta P_t(N_t/Q_{t+1})^{-\phi}$ to the left-hand side. The equation then says that net marginal convenience yield (the cost over the coming period of having one less unit of inventory) must equal the expected change in full marginal cost (the increase in cost this period minus the discounted decrease next period) from producing one more unit now, rather than selling it from inventory and producing it next period instead. This expected change in marginal cost may be due to expected changes in factor prices ($R_{1t}E_t w_{j,t+1}$ may be larger or smaller than w_{jt}), expected increases in cost due to convexity of the cost function, and changes in expected adjustment costs. Again, the error terms in eqn. (7) represent the unexplained parts of marginal production and storage costs.

Eqn. (7) includes the the marginal value of storage, $\Phi_N(N_t, Q_{t+1}, P_t)$, so estimation of that equation will provide estimates of the parameters of Φ_N, β and ϕ. Miron and Zeldes and Ramey estimate the parameters of Φ_N (which they constrain to be linear) just this way. However, we can use the fact that the net marginal convenience yield, $\psi_t = -\Phi_N - a_0 - \Sigma_j a_j D_{jt}$, can be inferred from futures prices. Using eqn. (1') with a one-month futures price replacing the forward price gives the following additional equation:

$$R_{1,t}F_{1,t} - P_t = a_0 + \Sigma_j a_j D_{jt} - \beta P_t E_t (N_t/Q_{t+1})^{-\phi} \qquad (8)$$

The basic model therefore contains three equations: (6), (7) and (8). These are estimated as a system, subject to cross-equation parameter constraints.[17,18] A number of issues regarding data and estimation are discussed in the next section.

One possible problem with this model is that I have arbitrarily specified the net marginal convenience yield function, ψ_t. Of course, this is also a problem with every earlier study that includes a cost of storage. However, in this case, if the primary interest is to estimate the parameters of the production cost function and the parameter β_1 that measures the cost of adjustment, we can use eqn. (8) to eliminate ψ_t altogether. Substituting the left-hand side of (8) for the terms that represent ψ_t in (7) gives the following alternative Euler equation:

$$\begin{aligned} - R_{1t}F_{1,t} + P_t = {} & c_0(1 - R_{1t}) + \sum_{j=1}^{11} c_j (D_{jt} - R_{1t}D_{j,t+1}) + \\ & E_t\Big[\sum_{j=1}^{m} \gamma_j (w_{jt} - R_{1t}w_{j,t+1}) + b(y_t - R_{1t}y_{t+1}) + \\ & \beta_1(\Delta y_t - 2R_{1t}\Delta y_{t+1} + R_{2t}\Delta y_{t+2})\Big] + \eta_t - R_{1t}E_t\eta_{t+1} \end{aligned} \qquad (7')$$

Note that this also eliminates inventories, N_t, as a variable in the model. Estimation of eqns. (6) and (7') will yield values for β_1, b, and the other parameters describing production cost that are unaffected by possible errors in the specification of ψ_t or the measurement of N_t.

4. Estimation Method and Data.

This section discusses the method of estimating the two versions of the model (eqns. (6), (7), and (8), and eqns. (6) and (7')), and the data set.

Estimation.

A natural estimator for an Euler equation model is an instrumental variables procedure that minimizes the correlation between variables known at time t and the equation residuals. Hence I simultaneously estimate eqns. (6), (7), and (8) using iterative three-stage least squares. The choice of instruments for this procedure deserves some comment.

Recall that the error terms η_t and ν_t represent unobserved shocks to production cost, storage cost, and demand. When estimating these equations, actual values for variables at time t+1 and t+2 are used in place of their expectations, which introduces expectational errors. For example, eqn. (6) becomes:

$$P_t = c_0 + \sum_{j=1}^{m} c_j w_{jt} - by_t + \beta_1(\Delta y_t - R_{1t}\Delta y_{t+1}) + \eta_t + \epsilon_{1,t+1} \qquad (9)$$

Similarly, eqn. (7) will have a composite error term $\eta_t - R_{1t}\eta_{t+1} + \nu_t + \epsilon_{2,t+1} + \epsilon_{2,t+2}$.

Under rational expectations, the errors $\epsilon_{1,t+1}$ and $\epsilon_{2,t+2}$ (and the corresponding errors for eqn. (8)) are by definition uncorrelated with any variable known at time t. However, this need not be the case for η_t, η_{t+1}, and ν_t, which may be correlated with endogenous variables. In addition, errors may be serially correlated. Hence, I use as instruments only variables which can reasonably be viewed as exogenous. The instrument list includes the set of seasonal dummy variables, and the following variables unlagged and lagged once: M1, the Index of Industrial Production, housing starts, the rate of inflation of the PPI, the rate of growth of the S&P 500 Common Stock Index, the rate of growth of labor hours, the three-month Treasury bill rate, and the weighted exchange value of the dollar against other G-10 currencies. For copper and lumber, I also include the price of crude oil. This gives a total of 30 instruments for copper and lumber, and 28 instruments for heating oil.

If the errors are conditionally homoscedastic, the minimized value of the objective function of this procedure provides a test statistic, J, which is distributed as χ^2 with degrees of freedom equal to the number of instruments times the number of equations minus the number of parameters.[19] This statistic is used to test the model's overidentifying restrictions, and hence the hypothesis that agents are optimizing with rational expectations.

<u>Data</u>.

The model is estimated using monthly data covering the period November 1972 through December 1987 for copper and lumber, and November 1978 through June 1988 for heating oil. Leads and lags in the equations reduce the actual time bounds by two months at the beginning and end of each period.

Production and inventory levels for each commodity are measured as follows. <u>Copper</u>: y_t is U.S. production of refined copper over the month, regardless of origin (ore or recycled scrap), and N_t is end-of-month stocks of refined copper at refineries and in Comex warehouses, both measured in short tons.[20] <u>Lumber</u>: y_t is monthly production and N_t is end-of-month inventories of softwood lumber. Units are millions of board feet.[21] <u>Heating oil</u>: y_t is monthly production and N_t is end-of-month inventories of distillate (No. 2) fuel oil. Units are millions of barrels.[22]

Unit sales for each commodity is calculated from unit production and end-of-month inventories using eqn. (4). The resulting series were compared to data from the same sources that are purportedly a direct measure of unit sales. The series were mostly identical, but occasionally data points will differ by up to one percent.

The production cost model includes variables that account for observable cost shocks. For all three commodities, I use average hourly nonagricultural earnings (w_{1t}), and the producer price index for intermediate materials, supplies and components (w_{2t}). For heating oil, I include as an additional cost variable the producer price index for crude petroleum (w_{3t}).

Some issues arise with respect to the choice of discount factor and the measurement of spot price, which I discuss in turn. Some studies have used a constant (real) discount factor, but in commodity markets, changes in nominal interest rates can have important effects on inventory holdings and price. Hence it is important to let the discount factor vary across time.

The choice of $R_{1,t}$ should reflect the rate actually used to discount nominal cash flows at time t. In the case of eqn. (8), which is an arbitrage relationship, this should clearly be the risk-free rate, e.g., the nominal Treasury bill rate. In the case of eqns. (6) and (7), however, the rate should include a premium that reflects the systematic risk associated with the various components of production cost. Unfortunately, this risk is likely to vary across the components of cost (in the context of the CAPM, it will depend on the beta of the commodity as well as the betas of the individual factor inputs), so there is no simple premium that can be easily measured.[23] I therefore ignore sytematic risk and use the nominal Treasury bill rate, measured at the end of each month, to calculate $R_{1,t}$ and $R_{2,t}$.

The measurement of the spot price requires a choice among three alternative approaches. First, one can use data on cash prices, purportedly reflecting actual transactions over the month. One problem with this is that it results in an average price over the month, as opposed to an end-of-month price. (The futures prices and inventory levels apply to the end of the month.) A second and more serious problem is that a cash price can include discounts and premiums that result from longstanding relationships between buyers and sellers, and hence is not directly comparable to a futures price when calculating convenience yields.

A second approach is to use the price on the "spot" futures contract, i.e., the contract that is expiring in month t. This approach also has problems. First, the spot contract often expires before the end of the month. In addition, open interest in the spot contract (the number of contracts outstanding) falls sharply as expiration approaches and longs and shorts close out their positions. As a result, by the end of the month there may be nothing resembling a spot transaction. Second, for most commodities, active contracts do not exist for each month.[24] (With copper, for example, there are active futures only for delivery in March, May, July, September, and December.)

The third approach, which I use here, is to infer 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 as follows:

$$P_t = F_{1t}(F_{1t}/F_{2t})^{(n_{01}/n_{12})} \qquad (10)$$

where P_t is the end-of-month spot price, F_{1t} and F_{2t} are the end-of-month prices on the nearest and next-to-nearest futures contracts, and n_{01} and n_{02} are, respectively, the number of days between t and the expiration of the nearest contract, and between the nearest and next-to-nearest contract. The advantage of this approach is that it provides spot prices for every month of the year. The disadvantage is that errors can arise if the term structure of spreads is nonlinear. However, a check against prices on some spot contracts indicates that such errors are likely to be small.

Finally, eqn. (10) is used to infer the one-month net marginal convenience yield. This simply involves replacing $F_{t+1,t}$ on the left-hand side of eqn. (8) with $P_t(F_{1t}/P_t)^{(1/n_{01})}$.

5. Results.

Tables 3 and 4 show, respectively, the results of estimating eqns. (6), (7), and (8), and eqns. (6) and (7'), for each commodity. Each model was first estimated without any correction for serial correlation, but the residuals of eqns. (6), (7) and (7') appeared to be AR(1). These equations were therefore quasi-differenced, and each model was re-estimated.

For lumber, the fit of the model is poor, at least as gauged by the J statistics, which test the overidentifying restrictions. The values for J indicate a rejection of the restrictions at the 5 percent level for both versions of the model. These rejections may indicate that producers of lumber do not optimize (at least on a month-to-month basis) with rational expectations, or that there is a failure in the model's specification. The overidentifying restrictions are not rejected, however, for copper and heating oil.

As for the estimates themselves, several points stand out. First, for all three commodities, the estimated marginal convenience yield function is strongly convex -- the cost of drawing down inventories rises rapidly as levels fall. Thus while production smoothing may indeed by an important role for inventories (as the numbers in Table 1 for copper and heating oil suggest), that role is limited to periods when inventories are at normal to high levels. The use of inventory to smooth prices is likewise limited. As Figures 1A-1C show, sharp price increases are usually accompanied by sharp increases in marginal convenience yield.

These estimates of β and ϕ also constitute a strong rejection of the quadratic target inventory model that is central to most studies of manufacturing inventories. This throws into question the findings of those studies, and absent a priori reasons for believing otherwise, suggests that the role of inventories in those industries is likely to vary dramatically as demand and aggregate inventory stocks fluctuate with the business cycle.

A second point is that β_1, the adjustment cost parameter, is insignificantly different from zero and/or negative for every commodity. This is the case both for the full model, and for eqns. (6) and (7'). Also, for copper and lumber, both versions of the model yield estimates for b, the slope of the marginal cost curve, that are insignificantly different from zero. It is hard to reconcile this with a production smoothing role for inventories (even during periods when inventories are large), as suggested (at least for copper) by the numbers in Table 1.

The results for heating oil do provide evidence of rising marginal costs, and the estimates are also economically meaningful. For example, this component of cost accounts on average for over 15 percent of the price of heating oil. Over the sample period, temporary increases in output added 3 to 6 cents to marginal cost because of the convexity of the cost function. This is further evidence that heating oil inventories are used to smooth production.

Several alternative versions of the model were also estimated. First, eqns. (6), (7) and (8), and eqns. (6) and (7') were estimated using quarterly data, on the grounds that intertemporal optimization may feasible only over time horizons longer than one month. The results were not very different. Second, cubic terms were added to the production cost function, on the grounds that the rejections of the overidentifying restrictions may be due to nonlinearities in marginal cost. However, those terms were uniformly insignificant, and left the J statistics almost unchanged. Finally, a risk premium parameter was added to the discount factor in eqns. (6), (7) and (7'), but estimates of this parameter were insignificant and/or not economically meaningful.

6. Conclusions.

Unlike models of manufacturing inventories, I have stressed the convex nature of the marginal convenience yield function, and used futures market data to infer values for this variable. But this also means estimating Euler equations, with the difficulties that this necessarily entails. The greatest difficulty is that estimation of structural parameters hinges on capturing intertemporal optimization by producers over periods corresponding to the frequency of the data - one month in this case. This may be too much to expect from the data, and may explain the rejection of the over-identifying restrictions for lumber, and the failure to find any evidence of adjustment costs or, for copper and lumber, a positively sloped marginal cost curve.

Of course there may also be problems with the specification of the model. A symmetric, convex adjustment cost function ignores important irreversibilities in production. Copper is a good example of this. There are sunk costs of building mines, smelters, and refineries, and sunk costs of temporarily shutting down an operation or restarting it. Such costs can induce firms to maintain output in the face of large fluctuations in price or sales. And such costs imply that it is the size of a price change, rather than the amount of time that elapses, that is the key determinant of the change in output.[25]

Nonetheless, the data reported in Table 1 and the parameter estimates for heating oil do indicate some production smoothing role for inventories. Although this may be important during periods of low or normal prices, it is probably not the primary role of inventories during periods of temporarily high prices. The very high net marginal convenience yields that are observed at such times, and the convex convenience yield functions that are estimated for all three commodities, are evidence that inventories may then have a more important role as an input to production. That role may be to facilitate production and delivery schedules and to avoid stockouts. That it is necessary is made clear by the fact that producers are willing to keep inventories on hand at an effective cost that is sometimes very high.

Table 1 - Variance Ratios

	Copper	Lumber	Heating Oil
Var(y)/Var(Q)	0.701	0.976	0.380
$Var(y^*)/Var(Q^*)$	0.680	1.011	0.744
$(\bar{N}/\bar{y})^2 Var(y)/Var(N)$	0.191	3.187	0.263
$(\bar{N}/\bar{y})^2 Var(y^*)/Var(N^*)$	0.149	9.035	0.391
$Var(y)/Var(y^*)$	1.287	1.333	1.530
$Var(N)/Var(N^*)$	1.005	3.793	2.277
Correl(y,Q)	0.728	0.964	0.198
$Correl(y^*,Q^*)$	0.698	0.962	0.399

Note: y = production, Q = sales, N = inventory. * indicates variable is deseasonalized.

Table 2 - NLS Estimates of Eq. (2)

	$\hat{\beta}$	$\hat{\phi}$	$F(a_j)$	R^2	DW
Copper	.0112 (.0015)	.9070 (.0935)	1.01	.849	0.58
Lumber	.0796 (.0050)	2.3618 (.5129)	2.04^*	.653	0.81
Heating Oil	.0826 (.0147)	1.3947 (.3337)	3.21^*	.507	1.14

Note: Asymptotic standard errors in parentheses. $F(a_j)$ is F statistic for significance of monthly dummy variables; * indicates significant at 5% level.

Table 3 - Estimation of (6), (7), and (8)

Parameter	Copper	Lumber	Heating Oil
γ_1	-.2316 (.1077)	-.9803 (1.160)	1.122 (1.399)
γ_2	-.8866 (1.378)	-23.289 (13.609)	5.182 (7.002)
γ_3			.0276 (.1640)
b	-.0000005 (.0000072)	-.0020 (.0052)	.1682 (.0876)
β_1	-.0000004 (.0000028)	.0006 (.0018)	-.0731 (.0305)
β	.0107 (.0020)	.2492 (.0671)	.0985 (.0194)
ϕ	.8891 (.1561)	3.317 (.8669)	1.484 (.2940)
a_0	.3593 (.1232)	6.653 (1.633)	1.171 (1.036)
a_1	.0664 (.0972)	2.636 (1.050)	0.029 (.5801)
a_2	-.0158 (.0904)	1.615 (.9931)	1.829 (.6180)
a_3	.0266 (.0987)	1.358 (1.147)	1.596 (.6761)
a_4	-.0001 (.1098)	0.505 (1.164)	2.604 (.6080)
a_5	-.0248 (.0985)	-1.515 (1.034)	1.835 (.6276)
a_6	-.0993 (.1071)	-0.514 (1.192)	1.747 (.6165)
a_7	-.0501 (.1198)	-2.070 (1.244)	2.096 (.7024)
a_8	.1407 (.1036)	1.065 (1.255)	1.971 (.6349)
a_9	.1493 (.0922)	-1.168 (1.131)	2.105 (.6351)
a_{10}	.1380 (.0977)	-0.181 (1.102)	1.411 (.6202)

Table 3 - Cont'd

Parameter	Copper	Lumber	Heating Oil
a_{11}	-.0647	-0.154	1.421
	(.0934)	(1.100)	(.5668)
c_0	100.858	507.832	-165.689
	(16.770)	(273.926)	(253.498)
c_1	1.735	-2.788	-2.581
	(1.607)	(4.652)	(2.217)
c_2	2.548	-1.883	-5.597
	(2.136)	(6.217)	(2.989)
c_3	1.668	-9.859	-1.667
	(2.462)	(7.128)	(3.446)
c_4	-1.226	-2.381	-0.447
	(2.648)	(7.667)	(3.779)
c_5	-2.394	-8.102	-1.842
	(2.759)	(7.995)	(4.003)
c_6	-1.268	-10.519	-3.845
	(2.853)	(8.155)	(4.096)
c_7	-3.524	-4.712	-1.047
	(2.797)	(8.034)	(4.005)
c_8	-3.254	-20.872	1.352
	(2.718)	(7.745)	(3.854)
c_9	-3.523	-21.438	2.858
	(2.519)	(7.178)	(3.598)
c_{10}	-2.542	-9.038	4.424
	(2.161)	(6.241)	(3.166)
c_{11}	-3.854	-15.019	1.692
	(1.607)	(4.664)	(2.336)
ρ_1	.9408	.9863	.9936
	(.0378)	(.0194)	(.0186)
ρ_2	.9968	.9602	.8890
	(.0507)	(.0657)	(.1439)
J	59.15	107.11*	61.59

Note: ρ_1 and ρ_2 are AR(1) coefficients for Eqns. (6) and (7). J is the minimized value of the objective function, distributed as $\chi^2(59)$ for copper and lumber, and $\chi^2(52)$ for heating oil. A * indicates significant at 5%. Asymptotic standard errors are in parentheses.

Table 4 - Estimation of (6) and (7')

Parameter	Copper	Lumber	Heating Oil
γ_1	-.2002 (.0897)	.2799 (1.059)	1.972 (1.453)
γ_2	-1.545 (.7391)	-14.282 (9.560)	1.550 (5.665)
γ_3			-.0365 (.1733)
b	.0000027 (.0000048)	-.0032 (.0044)	.1917 (.0893)
β_1	-.0000012 (.0000017)	.0012 (.0014)	-.0956 (.0290)
c_0	102.449 (13.307)	1288.47 (467.04)	-233.873 (267.754)
c_1	1.559 (1.618)	-3.234 (4.041)	-2.601 (2.040)
c_2	2.149 (2.119)	-0.737 (5.272)	-5.591 (2.740)
c_3	1.600 (2.442)	-8.524 (6.123)	-1.911 (3.189)
c_4	-1.230 (2.608)	-5.662 (6.620)	-0.059 (3.544)
c_5	-2.202 (2.720)	-10.209 (6.861)	-1.353 (3.751)
c_6	-1.306 (2.807)	-16.285 (7.062)	-3.285 (3.818)
c_7	-3.337 (2.770)	-9.569 (7.064)	-0.570 (3.735)
c_8	-3.139 (2.715)	-22.229 (6.887)	2.477 (3.631)
c_9	-3.570 (2.529)	-22.513 (6.622)	3.765 (3.455)
c_{10}	-2.766 (2.178)	-10.190 (5.777)	4.938 (3.044)
c_{11}	-4.007 (1.639)	-13.208 (4.343)	2.386 (2.197)
ρ_1	.9404 (.0373)	.9989 (.0014)	.9930 (.0179)
ρ_2	.9792 (.0404)	.9738 (.0326)	.9387 (.1325)
J	42.29	66.35*	24.72

Note: ρ_1 and ρ_2 are AR(1) coefficients for Eqns. (6) and (7'). J is the minimized value of the objective function, distributed as $\chi^2(43)$ for copper and lumber, and $\chi^2(38)$ for heating oil. A * indicates significant at 5%. Asymptotic standard errors are in parentheses.

APPENDIX - THE FUTURES PRICE/FORWARD PRICE BIAS

This appendix shows that the futures price can be used as a proxy for the forward price in eqn. (1') with negligible measurement error. Ignoring systematic risk, the difference between the futures price, $F_{T,t}$, and the forward price, $f_{T,t}$, is

$$F_{T,t} - f_{T,t} = - \int_t^T F_{T,w} \mathrm{cov}[(dF_{T,w}/F_{T,w}),(dB_{T,w}/B_{T,w})]dw \qquad (A.1)$$

where $B_{T,w}$ is the value at time w of a discount bond that pays $1 at T, and cov[ ] is the local covariance at time w between percentage changes in F and B. (See Cox, Ingersoll and Ross (1981) and French (1983).) Let r_w be the yield to maturity of the bond. Then approximating dB/B by rdt - (T-w)dr, and $F_{T,w}$ by its mean value over (w,T), the average percentage bias, (F-f)/F, for a one-month contract is roughly:

$$\% \text{ Bias} \approx \bar{r}[\widehat{\mathrm{cov}}(\Delta r/r, \Delta F/F)] \qquad (A.2)$$

where $\bar{r}$ is the mean monthly bond yield, and $\widehat{\mathrm{cov}}$ is the sample covariance.

Using the three-month Treasury bill rate for r and the nearest active contract price for F, I obtain the following estimates for this bias: copper, .0030%; lumber, -.0032%; and heating oil, .0077%. The largest bias is for heating oil, but even this represents less than a hundreth of a cent for a one-month contract.

REFERENCES

Blanchard, Olivier J., "The Production and Inventory Behavior of the American Automobile Industry," Journal of Political Economy, June 1983, 91, 365-400.

Blanchard, Olivier J., and Angelo Melino, "The Cyclical Behavior of Prices and Quantities," Journal of Monetary Economics, 1986, 17, 379-407.

Blinder, Alan S., "Inventories and Sticky Prices: More on the Microfoundations of Macroeconomics," American Economic Review, June 1982, 72, 334-348.

Blinder, Alan S., "Can the Production Smoothing Model of Inventory Behavior Be Saved?" Quarterly Journal of Economics, August 1986, 101, 431-453.

Brennan, Michael J., "The Supply of Storage," American Economic Review, March 1958, 48, 50-72.

Brennan, Michael J., "The Cost of Convenience and the Pricing of Commodity Contingent Claims," Center for the Study of Futures Markets, Working Paper No. 130, June 1986.

Brennan, Michael J., and Eduardo S. Schwartz, "Evaluating Natural Resource Investments," Journal of Business, April 1985, 58, 135-157.

Bresnahan, Timothy F., and Valerie Y. Suslow, "Inventories as an Asset: The Volatility of Copper Prices," International Economic Review, June 1985, 26, 409-424.

Cox, John C., Jonathan Ingersoll, Jr., and Stephen A. Ross, "The Relation Between Forward Prices and Futures Prices," Journal of Financial Economics, 1981, 9, 321-346.

Eckstein, Zvi, and Martin S. Eichenbaum, "Inventories and Quantity Constrained Equilibria in Regulated Markets: The U.S. Petroleum Industry, 1947-72," in T. Sargent, ed., Energy Foresight and Strategy, Resources for the Future, Washington, D.C., 1985.

Eichenbaum, Martin S., "A Rational Expectations Equilibrium Model of Inventories of Finished Goods and Employment," Journal of Monetary Economics, August 1983, 12, 259-278.

Eichenbaum, Martin S., "Rational Expectations and the Smoothing Properties of Inventories of Finished Goods," Journal of Monetary Economics, 1984, 14, 71-96.

Eichenbaum, Martin S., "Some Empirical Evidence on the Production Level and Production Cost Smoothing Models of Inventory Investment," American Economic Review, September 1989, 79, 853-864.

Fair, Ray C., "The Production Smoothing Model is Alive and Well," NBER Working Paper, 1989.

Fama, Eugene F., and Kenneth R. French, "Business Cycles and the Behavior of Metals Prices," Journal of Finance, December 1988, 43, 1075-93.

French, Kenneth R., "A Comparison of Futures and Forward Prices," Journal of Financial Economics," September 1983, 12, 311-42.

Kahn, James A., "Inventories and the Volatility of Production," American Economic Review, Sept. 1987, 77, 667-679.

Miron, Jeffrey A., and Stephen P. Zeldes, "Seasonality, Cost Shocks, and the Production Smoothing Model of Inventories," Econometrica, July 1988, 56, 877-908.

Ramey, Valerie A., "Non-Convex Costs and the Behavior of Inventories," unpublished, June 1988.

Ramey, Valerie A., "Inventories as Factors of Production and Economic Fluctuations," American Economic Review, June 1989, 79, 338-354.

Telser, Lester G., "Futures Trading and the Storage of Cotton and Wheat," Journal of Political Economy, 1958, 66, 233-55.

Thurman, Walter N., "Speculative Carryover: An Empirical Examination of the U.S. Refined Copper Market," Rand Journal of Economics, Autumn 1988, 19, 420-437.

West, Kenneth D., "A Variance Bounds Test of the Linear Quadratic Inventory Model," Journal of Political Economy, April 1986. 91, 374-401.

Williams, Jeffrey, "Futures Markets: A Consequence of Risk Aversion or Transactions Costs?" Journal of Political Economy, October 1987, 95, 1000-1023.

Working, Holbrook, "Theory of the Inverse Carrying Charge in Futures Markets," Journal of Farm Economics, February 1948, 30, 1-28.

Working, Holbrook, "The Theory of the Price of Storage," American Economic Review, 1949, 39, 1254-62.

FOOTNOTES

1. The standard deviations of monthly percentage changes in the spot prices of copper, lumber, and heating oil, for example, have all averaged more than 10 percent over the past two decades, and in some years have been three or four times higher.

2. See, for example, Blanchard (1983), Blinder (1986), and West (1986). But also see Fair (1989), who shows that the use of disaggregated (three- and four-digit SIC) data, for which units sold is measured directly rather than inferred from dollar sales, supports the production smoothing model.

3. See Blanchard (1983), Miron and Zeldes (1988), and Eichenbaum (1989). All of their models include a cost of deviating from a target inventory level, where the target is proportional to sales. As Kahn (1987) has shown, this is consistent with the use of inventories to avoid stockouts.

4. As these authors show, the solution to the firm's optimization problem can be stated as a set of parameter restrictions in a vector autoregression. For analytical studies of the linear-quadratic inventory model, see Blinder (1982) and Eichenbaum (1983).

5. See, for example, Brennan (1958) and Telser (1958).

6. Studies of manufactured inventories generally use Department of Commerce data in which production is computed from dollar sales, a deflator, and inventories. Fair (1989) shows that the resulting measurement errors add spurious volatility to the production series.

7. Two other studies of commodity inventories and prices should be mentioned. Bresnahan and Suslow (1985) show that with inventory stockouts, price can take a perfectly anticipated fall, i.e., the spot price can exceed the futures price. Hence capital gains are limited (by arbitrage through inventory holdings), but capital losses are unlimited. However, they ignore convenience yield, which, as we will see, is an important component of the total return to holding inventory. Also, Thurman (1988) develops a rational expectations model of inventory holding which he estimates for copper, but the model is linear and takes production as fixed.

8. This is supported by earlier studies of commodities (see Footnote 5), and by my data for copper, lumber, and heating oil. As for manufactured goods, Ramey (1989) models inventories as a factor of production, and her results imply that production cost can rise sharply as inventories become small. This view of inventories as an essential factor of production is consistent with my findings.

9. Thus ψ is the net flow of benefits that accrues from the marginal unit of inventory, a notion first introduced by Working (1948, 1949). Williams (1987) shows how convenience yield can arise from non-constant costs of processing.

10. Note that the <u>expected</u> future spot price, and thus the risk premium on a forward contract, will depend on the "beta" of the commodity. But because $\psi_{t,T}$ is the capitalized convenience yield, expected spot prices or risk premia do not appear in eqn. (1'). Indeed, eqn. (1') depends in no way on the stochastic structure of price evolution or on any particular model of asset pricing, and one need not know the "beta" of the commodity.

11. If the interest rate is non-stochastic, the present value of the expected daily cash flows over the life of the futures contract will equal 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 (which we would expect to be the case 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. For a rigorous proof of this result, see Cox, Ingersoll, and Ross (1981).

12. French (1983) compares the futures prices for silver and copper on the Comex with their forward prices on the London Metals Exchange, and shows that the differences are very small (about 0.1% for 3-month contracts).

13. During 1988, the net convenience yield for copper reached 40 cents per pound, which was nearly 30 percent of the price.

14. Ideally, an expression should be derived for ψ_t from a dynamic optimizing model of the firm in which there are stockout costs, costs of scheduling and managing production and shipments, etc., but that is well beyond the scope of this paper. However, Brennan (1986) shows that a functional form close to (2) can be derived from a simple transactions cost model.

15. If ψ_t is a convex function of N_t, the spot price should be more volatile than the futures or forward prices, especially when stocks are low. Fama and French (1988) show that this is indeed the case for a number of metals.

16. More general specifications could have been used for both direct cost and the cost of adjustment, but at the expense of adding parameters.

17. Note that if futures market data were unavailable, one could instead use the following equation:

$$R_{1t}E_tP_{t+1} - P_t = a_0 + \Sigma_j a_j D_j + \Phi_N(N_t, Q_{t+1}, P_t) \qquad \text{(i)}$$

i.e., firms hold inventory up to the point where the expected capital gain in excess of interest costs just equals the full marginal cost of storage, where the latter is the cost of physical storage less the gross benefit (marginal convenience yield) that the unit provides. This equation can be derived by using (4) to eliminate Q_t instead of y_t and then maximizing with respect to N_t. (If the only errors are expectational, the covariance matrix of (6), (7), and (i) would be singular, but this problem does not arise if there are also random shocks to current production and storage costs.) But expectational errors in (i) are likely to be large, so it is preferable to use the information in futures prices and estimate (8).

18. The model as specified above ignores the demand side of the market. Assuming a quadratic cost of adjusting consumption, the corresponding intertemporal optimization problem of buyers is: where $U(Q_t)$ is the utility (e.g., gross revenue product in the case of an industrial buyer) from consuming at a rate Q_t, given an index of aggregate economic activity X_t.

$$\max_{\{Q_t\}} E_t \sum_{\tau=t}^{\infty} R_{t,\tau}[U(Q_\tau) - P_\tau Q_\tau - (1/2)\beta_2(\Delta Q_\tau)^2] \quad \text{(i)}$$

The corresponding first-order condition is:

$$P_t = U_Q(Q_t,X_t) - \beta_1(\Delta Q_t - E_t R_{1t}\Delta Q_{t+1}) \quad \text{(ii)}$$

i.e., marginal utility must equal full marginal cost, where the latter equals the price of a unit plus the expected change in adjustment cost from consuming one more unit now. One can also estimate the expanded system, i.e., eqns. (6), (7), (8), and (ii).

19. I make the assumption of conditional homoscedasticity for simplicity. If the assumption is incorrect, the parameter estimates will still be consistent, but the standard errors and test of the overidentifying restrictions will not be valid.

20. Source: Metal Statistics (American Metal Market), various years. Note that only finished product stocks are included. Excluded are "in process" stocks, such as stocks of ore at mines and smelters, and stocks of unrefined copper at smelters and refineries.

21. Source: National Forest Products Association, Fingertip Facts and Figures. Most of the lumber consumed in the U.S. is softwood (e.g., pine and fir). Futures contracts for softwood lumber are traded on the Chicago Mercantile Exchange.

22. Source: U.S. Department of Energy, Monthly Energy Review, various issues.

23. The use of an average cost of capital for firms in the industry is also incorrect; we want a beta for a project that produces a marginal unit of the commodity, not a beta for equity or debt of the firm.

24. There are often additional thinly traded contracts, but the number of transactions may not suffice to measure the end-of-month spot price.

25. For a model that accounts for these sunk costs, see Brennan and Schwartz (1985).

FIGURE 1A

COPPER – SPOT PRICE AND NET CONVENIENCE YIELD

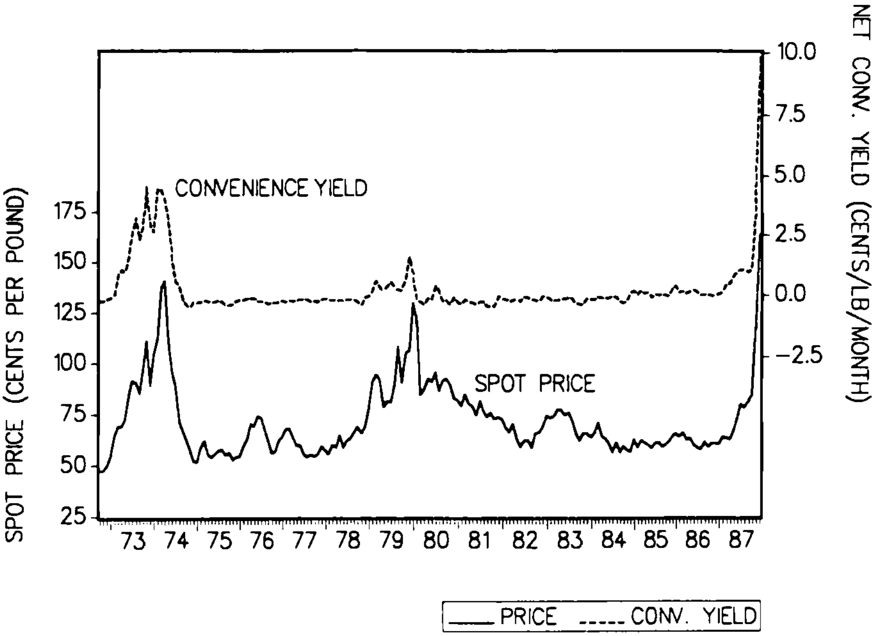

FIGURE 1B

LUMBER – SPOT PRICE AND NET CONVENIENCE YIELD

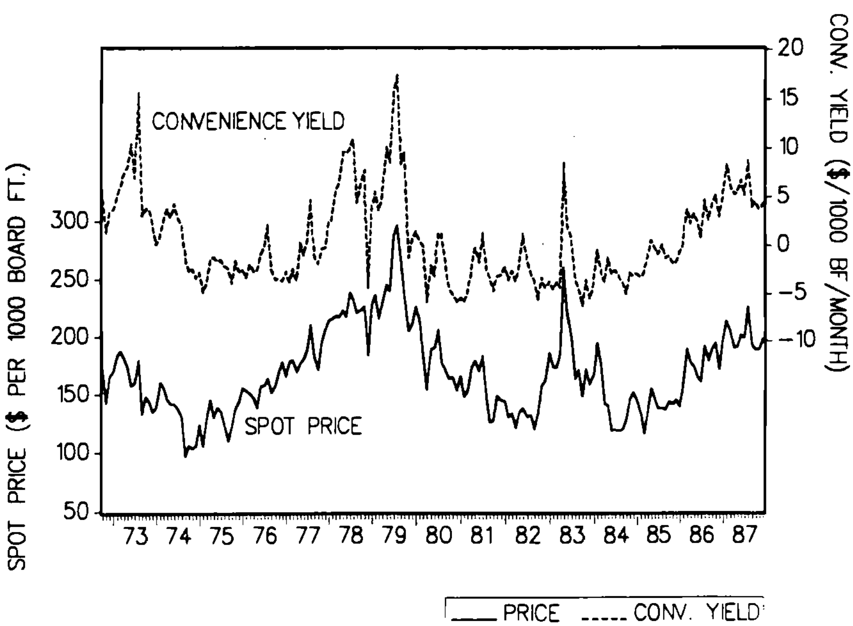

FIGURE 1C

HEATING OIL – SPOT PRICE AND NET CONVENIENCE YIELD

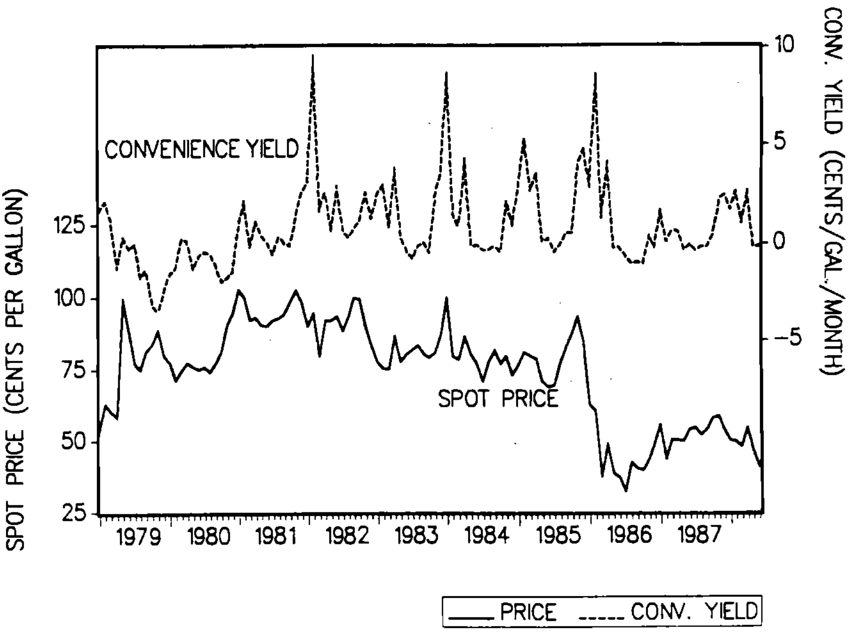

FIGURE 2A

COPPER – NET CONVENIENCE YIELD VS. N/Q

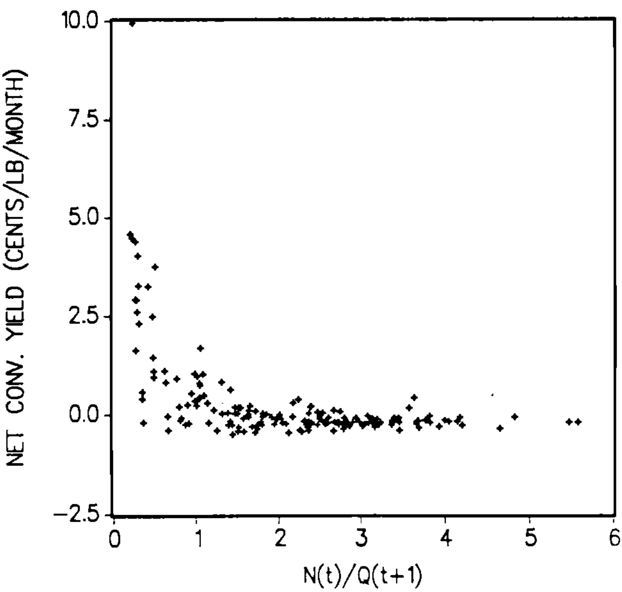

FIGURE 2B

LUMBER – NET CONVENIENCE YIELD VS. N/Q

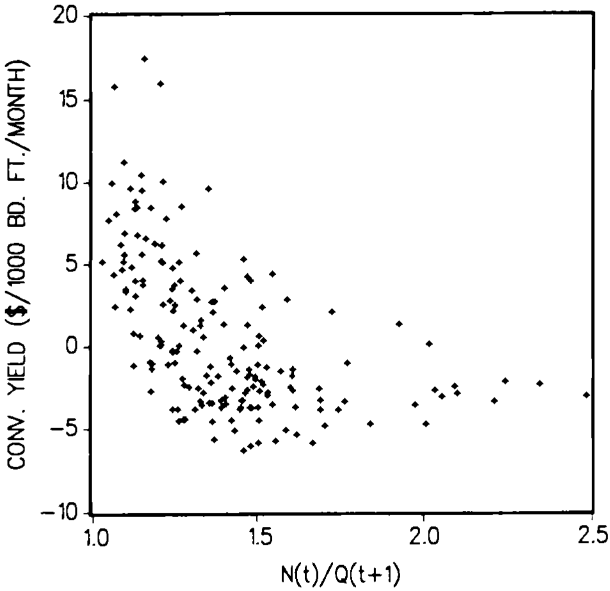

FIGURE 2C

HEATING OIL – NET CONVENIENCE YIELD VS. N/Q