RAND Journal of Economics
Vol. 25, No. 1, Spring 1994

Inventories and the short-run dynamics of commodity prices

Robert S. Pindyck*

Competitive producers hold inventories to reduce costs of adjusting production and to reduce marketing costs by facilitating scheduling and avoiding stockouts. Using data for copper, heating oil, and lumber, I estimate these costs within a structural model of production, sales, and storage, and I study their implications for inventory and price behavior. Unlike earlier studies, this work focuses on homogeneous and fungible commodities. This avoids aggregation problems, and it allows the use of direct measures of units produced, rather than inferences from dollar sales. Also, I estimate Euler equations and allow the marginal value of storage to be a convex function of the stock. This fits the data better, and helps explain the role of storage. Finally, I use futures prices to directly measure the marginal value of storage. I find a production-smoothing role for inventories only for heating oil, and during periods of low or normal prices. A more important role is to reduce marketing costs.

1. Introduction

■ The markets for many commodities are characterized by periods of sharp changes in prices and inventory levels. This article examines the role of inventories in the short-run dynamics of production and price, and it seeks to determine whether fluctuations in spot and futures prices can be explained in terms of rigidities in production and/or inventory demand.

In a competitive commodity market, producers hold inventories to reduce costs of adjusting production, but also to reduce marketing costs by facilitating production and delivery scheduling and avoiding stockouts. These latter factors make it costly for firms to reduce inventory holdings beyond some minimal level, even if marginal production cost is constant and adjustment costs are negligible. In general, the extent to which price fluctuates in the short run depends on costs of changing production as well as costs of drawing down inventories.

* Massachusetts Institute of Technology.

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, New Jersey, for providing futures market data. My thanks to Patricia Craig, John Simpson, and Yunyong Thaicharoen for their research assistance, and to Ben Bernanke, Michael Brennan, Zvi Eckstein, Martin Eichenbaum, Jeffrey Miron, Ariel Pakes, Julio Rotemberg, Menachem Sternberg, Lester Telser, seminar participants at MIT, Yale, Tel-Aviv University, and the NBER, and two anonymous referees for helpful comments and suggestions.

To determine these costs, I estimate a structural model of production, sales, and storage for three commodities: copper, heating oil, and lumber. 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 provided little support for the production-smoothing model of inventories; in fact, the variance of production generally exceeds the variance of sales in manufacturing.[1] There is more support for models of production-cost smoothing, in which inventories are used to shift production to periods of lower costs, and models in which inventories are used to avoid stockouts and reduce scheduling and other marketing costs.[2]

The data suggest that inventories play several roles in commodity markets. For two of the three commodities studied here, the variance of production is much less than the variance of sales, consistent with production smoothing. But the empirical results in this article show that for all three commodities, the cost of drawing down inventories rises rapidly as inventory levels fall, suggesting that inventories are needed to limit marketing costs. This would limit their use for production or production-cost smoothing, particularly during periods of high prices following shocks.

Besides their focus on manufactured goods, most earlier studies 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 the automobile industry 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.

Although convenient, the linear-quadratic specification is a major limitation of these models. First, marginal production cost might not be linear. But more important, a quadratic cost of deviating from a target inventory level implies that the cost of a marginal reduction in inventory is linear in the stock of inventory. Besides allowing negative inventories, this is a bad approximation. Early studies have demonstrated, and the data here confirm, that for commodities the marginal cost of drawing down inventories is highly convex in 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.[3] There is no reason to expect a linear approximation to be any better for manufactured goods.

The alternative approach is to abandon the linear-quadratic framework, adopt a more general specification, and estimate the Euler equations that follow from intertemporal optimization. This was done 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 accounts for unobservable cost shocks and seasonal fluctuations in sales, and by Ramey (1991), who uses a cubic cost function to show that declining marginal cost may help explain the excess volatility of production. However, in both of these studies

[1] See, e.g., Blanchard (1983), Blinder (1986), and West (1986). But Fair (1989) 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.

[2] 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. One of the earliest inventory studies is Holt et al., (1960), who estimate costs of inventory holdings and back orders using factory-level data. For a survey of recent research on inventories, see Blinder and Maccini (1991).

[3] Early studies include Brennan (1958) and Telser (1958). McCallum (1974) estimates a model of competitive price dynamics for the lumber industry, but he also restricts the marginal cost of drawing down inventory to be linear in the stock of inventory. Nonetheless, he shows that this cost plays an important role in price adjustment.

the cost of deviating from the target inventory level is quadratic, so that the marginal cost of drawing down inventories 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 it allows me to use direct measures of units produced, rather than inferences from dollar sales and inventories.[4] Second, as in Miron and Zeldes and Ramey, I estimate Euler equations but allow the marginal cost of drawing down inventory to be a convex function of the stock. This fits the data better, and it helps explain the value of storage and its role in the dynamics of price. Third, I use futures market data to directly measure the marginal value of storage, and thereby determine its dependence on the stock of inventory.[5]

The next section discusses the value of storage, presents basic data, 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

■ It is useful to separate a firm's costs of doing business into two components. The first is the direct cost of production, which depends on the prices and quantities of factor inputs. The marginal cost of production might or might not be rising, and there may also be costs of adjusting production. The second component of cost relates to the marketing of the firm's output, and it includes costs of scheduling production and deliveries and avoiding stockouts.

Both components of cost can create a value to holding inventory. If marginal inventory is sharply rising in the short run and/or there are substantial costs of adjusting production, inventories can be used to smooth production when demand is fluctuating, and thereby reduce cost. But even if marginal production cost is constant and there are no adjustment costs, inventories are needed as a lubricant to facilitate scheduling and thereby reduce marketing costs. The marginal value of storage is the savings in marketing costs resulting from one additional unit of inventory. This marginal value is likely to be small when the total stock of inventory is large, but it can rise sharply when the stock becomes very small.[6]

Letting N_t denote the end-of-period inventory level, P_t the price, and E_tQ_{t+1} the expected next-period sales, we can represent the total per-period marketing cost by a function $\Phi(N_t, E_tQ_{t+1}, P_t)$, with $\Phi_N < 0$, $\Phi_{NN} > 0$, $\Phi(0, Q, P) = \infty$, $\Phi \to 0$ for N large, and Φ_Q, $\Phi_P > 0$. Hence the benefit (in terms of reduced marketing costs) of an extra unit of inventory is $-\Phi_N$. This is commonly referred to as the *marginal convenience yield* from storage.

I will assume that there is a (constant) cost of physical storage of a dollars per unit per period. Thus total per-period marketing and storage costs are given by $\Phi(N_t, E_tQ_{t+1}, P_t) + aN_t$, and the *net* benefit of an extra unit of inventory is $-\Phi_N - a$. We

[4] 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.

[5] Two other related studies should be mentioned. Bresnahan and Suslow (1985) show that with 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 the nonspeculative value of inventory. Also, Thurman (1988) estimates a log-linear rational expectations model of inventory holding for copper in which production follows an AR(1) process and the marginal value of storage is a convex function of the stock of inventory.

[6] This is supported by earlier studies (see footote 3), and by the results of this article. As for manufactured goods, Ramey (1989) models inventories as an essential factor of production, and her results imply that production cost can rise sharply as inventories fall, which is consistent with my findings.

will let $\psi \equiv -\Phi_N - a$ denote this net benefit, i.e., ψ is the net marginal convenience yield.[7]

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

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

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 nonstochastic, so it must equal $r_T P_t$, from which (1) follows.

In keeping with the literature on inventories (see the references in footnotes 1 and 2), I 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[8]

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

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 *and* is positively (negatively) correlated with the spot price.[9] 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.[10] I therefore use the futures price, $F_{T,t}$, in place of the forward price in (1a).

Figures 1, 2, and 3 show spot prices for copper, lumber, and heating oil, together with the one-month net marginal convenience yield, $\psi_t \equiv \psi_{t,1}$. (My data for copper and lumber run from October 1972 through December 1987. Heating oil futures began trading only in late 1978, so data for this commodity cover November 1978 to June 1988. The data and construction of ψ_t are discussed 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, 1979–1980, and the end of 1987. On each occasion (and especially the first and third), the convenience yield also rose sharply. Likewise, when lumber prices

[7] This notion of marginal convenience yield was introduced by Working (1949). Williams (1987) shows how convenience yield can arise from nonconstant processing costs. If storage is always positive, price is the present value of the expected future flow of convenience yield. Pindyck (1993) tests this present-value model of commodity pricing.

[8] Note that the *expected* future spot price, and thus the risk premium on a forward contract, depends on the "beta" of the commodity. But expected spot prices or risk premia do not appear in (1a). Indeed, (1a) does not depend on the stochastic structure of price or on any model of asset pricing.

[9] If the interest rate is nonstochastic, 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).

[10] 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 .1% for three-month contracts).

FIGURE 1

COPPER: SPOT PRICE AND NET CONVENIENCE YIELD

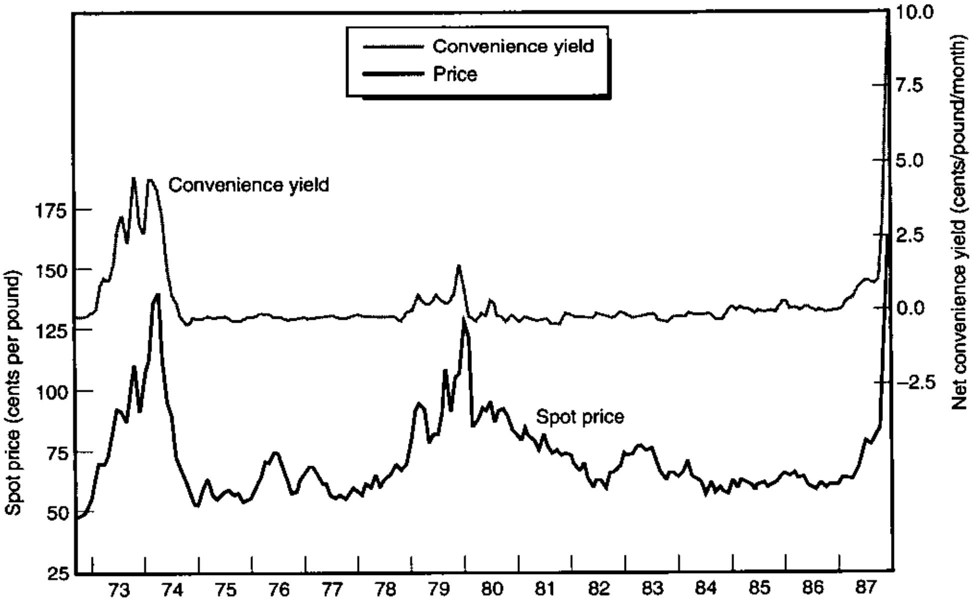

FIGURE 2

LUMBER: SPOT PRICE AND NET CONVENIENCE YIELD

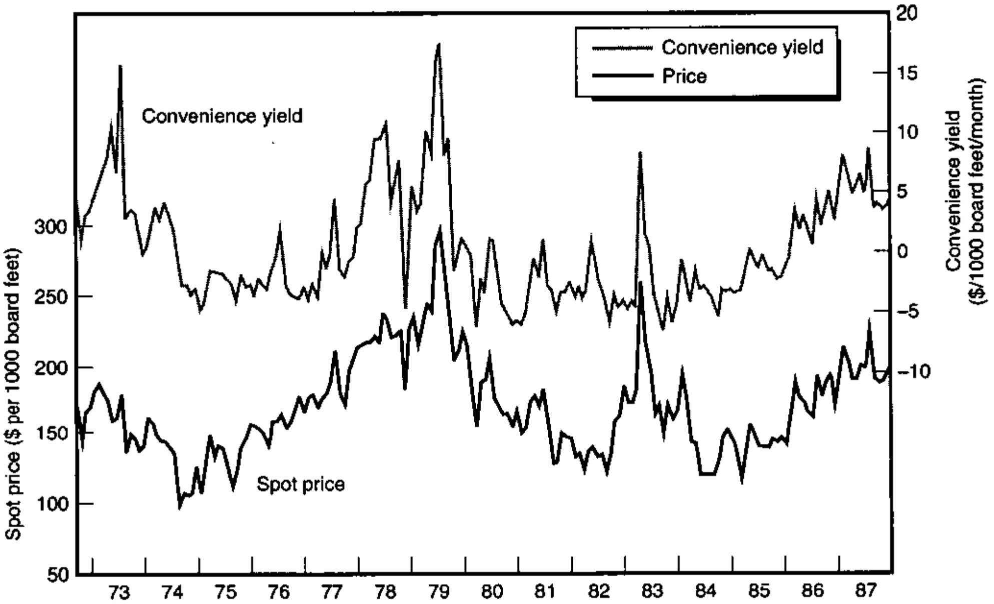

FIGURE 3

HEATING OIL: SPOT PRICE AND NET CONVENIENCE YIELD

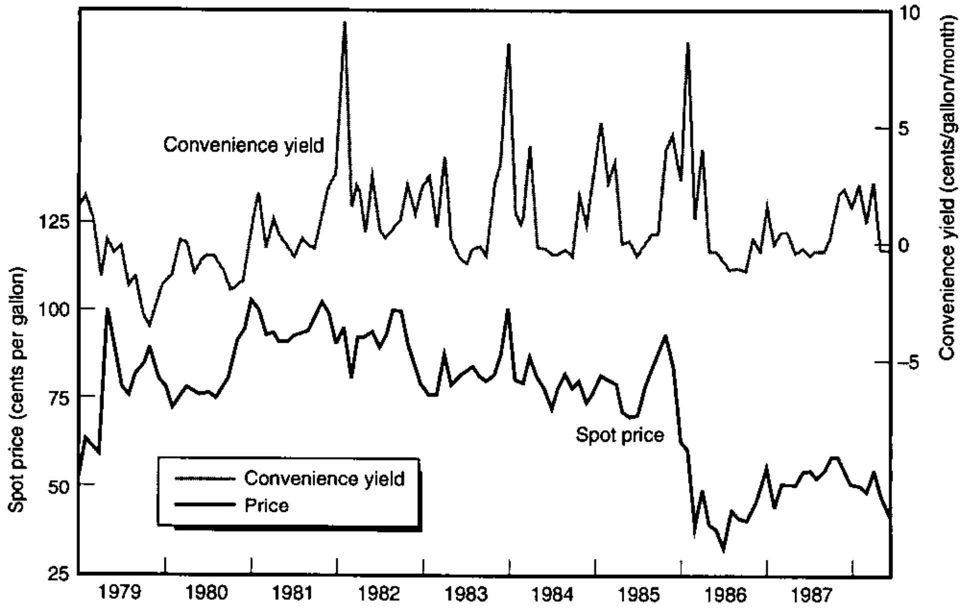

rose in early 1973, 1977–1979, 1983, and 1986–1987, the convenience yield also rose. For heating oil the comovement is smaller (and much of what there is is seasonal), but there has still been a tendency for price and convenience yield to move together.

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% of the price. Thus firms were paying 8% *per month*—plus interest and direct storage costs—to maintain stocks. (By 1988, the net convenience yield reached 40 cents per pound, which was nearly 30% of the price.) The net convenience yield for lumber and heating oil also reached peaks of 8% to 10% of price. During these periods of high prices and high convenience yields, inventory levels were below 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 role of inventories is to avoid stockouts and facilitate the scheduling of production and sales. This role probably dominates when prices are high and inventory levels are low.

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 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 for 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, but for copper and heating oil the variance of sales still exceeds that of production. However, as West (1986) and Kahn (1990, 1992) show, this need not imply that inventories are used to smooth production. Also, Kahn (1990), using a longer time series (1947–1987), finds the variances of production and sales to be approximately the same for copper.

TABLE 1 **Variance Ratios**

	Copper	Lumber	Heating Oil
$\mathrm{Var}(y)/\mathrm{Var}(Q)$	.701	.976	.380
$\mathrm{Var}(y)^*/\mathrm{Var}(Q^*)$	.680	1.011	.744
$(\bar{N}/\bar{y})^2\mathrm{Var}(y)/\mathrm{Var}(N)$	.191	3.187	.263
$(\bar{N}/\bar{y})^2\mathrm{Var}(y^*)/\mathrm{Var}(N^*)$	.149	9.035	.391
$\mathrm{Var}(y)/\mathrm{Var}(y^*)$	1.287	1.333	1.530
$\mathrm{Var}(N)/\mathrm{Var}(N^*)$	1.005	3.793	2.277
$\mathrm{Correl}(y, Q)$	.728	.964	.198
$\mathrm{Correl}(y^*, Q^*)$	.698	.962	.399

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

Table 1 also shows the ratio of the variance of production to that of inventories, normalized by the squared means. For copper and heating oil, inventories vary much more than production, whether or not the variables have been deseasonalized, suggesting that inventories are used to smooth production. But for lumber, the variances of production and sales are about the same, and production varies much more than inventories, especially after deseasonalizing. Also, production and sales track each other very closely. This suggests that production smoothing is not important for lumber, and instead inventories are needed to facilitate scheduling and avoid stockouts.

Finally, what do the data tell us about the dependence of the marginal convenience yield on the level of inventories? Marketing costs, and hence the marginal value of storage, should be roughly proportional to the price of the commodity, and they should also 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:

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

Ideally, ψ_t should be derived from a dynamic optimizing model of the firm in which there are stockout costs and costs of scheduling and managing production and shipments, etc., but that is beyond the scope of this article. However, Brennan (1991) shows that a functional form close to (2) can be derived from a simple transactions cost model.

Figures 4, 5, and 6 show ψ_t plotted against the inventory-sales ratio, N_t/Q_{t+1}, for each commodity. These figures suggest that ψ_t is well described by (2), with β, $\phi > 0$, and that the linear relationship used in most studies of inventories may be a poor approximation of what is in fact a highly convex function. Also, note that 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 several metals.

Table 2 shows nonlinear least squares estimates of (2), with monthly dummy variables included for a. (These dummies capture seasonal shifts in both the cost of storage and 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. Also, the monthly dummy variables are groupwise significant for every commodity. As expected, there are strong seasonal fluctuations in the net benefit from holding inventory.

3. The model

■ Intertemporal optimization by producers balances three costs: the cost of producing, 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, i.e, the increase in marketing costs less the savings in storage costs resulting from less in-

FIGURE 4

COPPER: NET CONVENIENCE YIELD VS. *N/Q*

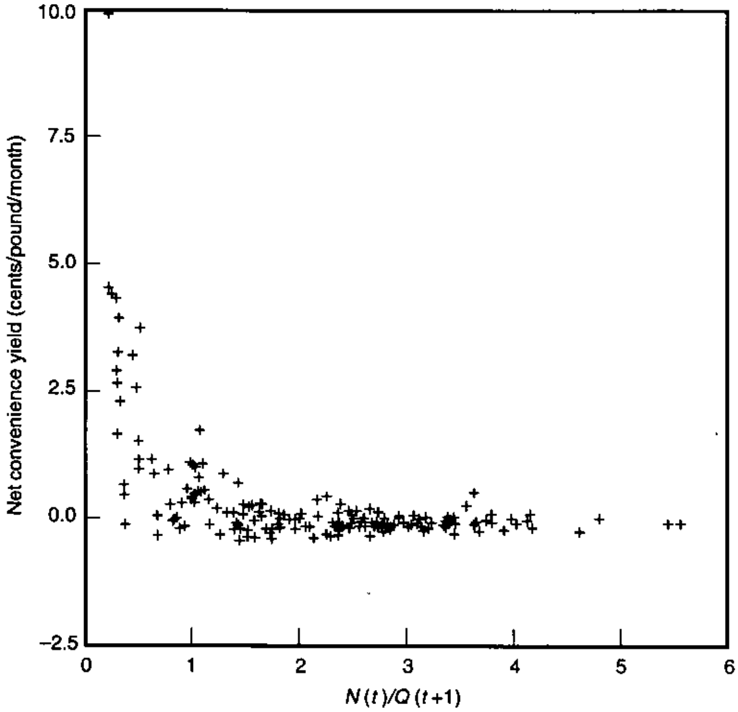

FIGURE 5

LUMBER: NET CONVENIENCE YIELD VS. *N/Q*

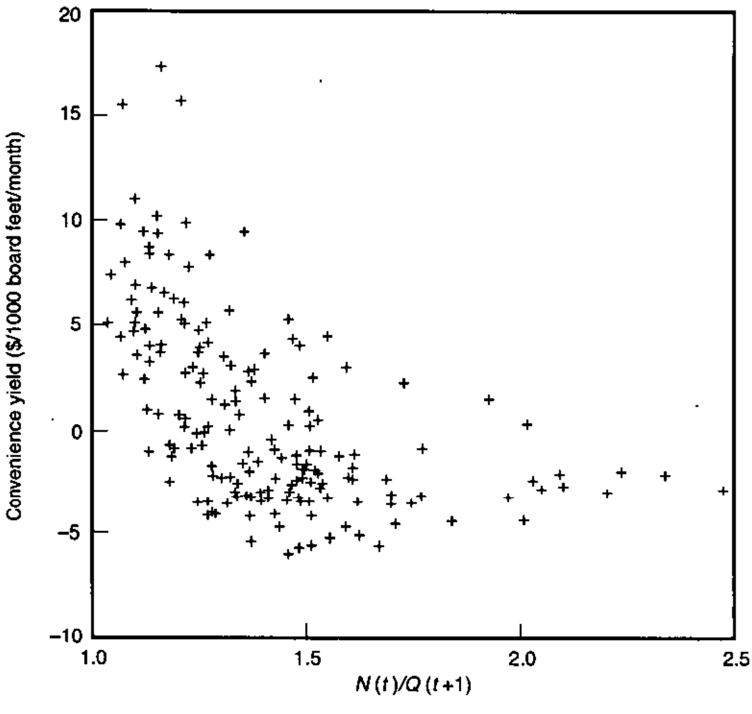

FIGURE 6

HEATING OIL: NET CONVENIENCE YIELD VS. *N/Q*

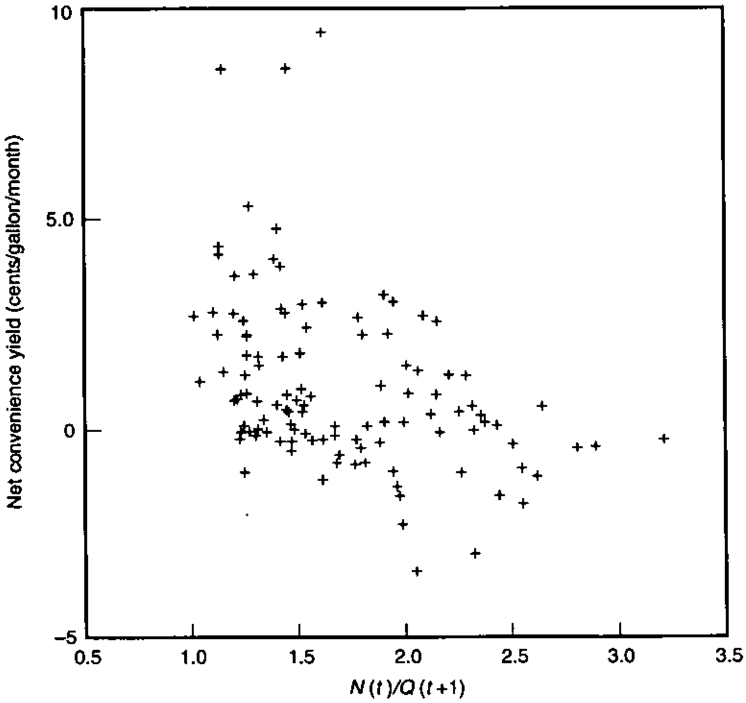

ventory. My objective is to estimate all three of these costs and determine their dependence on output, sales, and inventory levels. To do this, I make use of the fact that in the U.S. markets for copper, heating oil, and lumber, producers can be viewed as price takers. This, together with the fact that 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 (e.g., Blanchard (1983), Miron and Zeldes (1988), and Ramey (1991)).

I model the direct cost of production as quadratic in output, I assume that there is a quadratic cost of adjusting output, and I use (2) to represent the net savings in marketing costs from a marginal unit of inventory. Direct production cost, marketing costs, and storage costs are likely to fluctuate seasonally, so I introduce monthly dummy variables.

TABLE 2 **Nonlinear Least Squares Estimates of Equation (2)**

	$\hat{\beta}$	$\hat{\phi}$	$F(a_j)$	$\hat{\rho}$	R^2	DW
Copper	.0120	.9050	2.79*	0.75	.922	1.72
	(.0021)	(.1088)				
Lumber	.0934	.3029	2.75*	0.50	.814	1.62
	(.0073)	(.1540)				
Heating oil	.1107	.8110	3.66*	0.54	.609	2.21
	(.0223)	(.2788)				

Note: Asymptotic standard errors are in parentheses. $F(a_j)$ is the F statistic for significance of monthly dummy variables; a * indicates significance at the 5% level. ρ is the coefficient for AR(1) correction. See Section 4 for a discussion of the spot price series.

Allowing for unobservable shocks, total per-period cost can be written as

$$C_t = (c_0 + \sum_{j=1}^{11} c_j D_{jt} + \sum_{j=1}^{m} \gamma_j w_{jt} + \eta_t)y_t + (1/2)by_t^2 + (1/2)\beta_1(\Delta y_t)^2 + \Phi(N_t, E_t Q_{t+1}, P_t) + (a_0 + \sum_{j=1}^{11} a_j D_{jt} + \nu_t)N_t. \tag{3}$$

Here, the w_{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 w_{jt}s and the error terms η_t and ν_t allow for both observable and unobservable cost shocks.

Inventories must satisfy the following accounting identity:

$$N_t = N_{t-1} + y_t - Q_t. \tag{4}$$

Taking price as given, firms choose production and sales levels to maximize the present value of the flow of expected profits, subject to (4):

$$\max E_t \sum_{\tau=0}^{\infty} R_{\tau,t}(P_{t+\tau}Q_{t+\tau} - C_{t+\tau}), \tag{5}$$

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

To obtain first-order conditions, use (4) to eliminate y_t, then maximize with respect to Q_t and N_t. First, maximize with respect to Q_t, holding N_t and Q_{t+1} fixed. This yields

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

Second, maximize with respect to N_t, holding Q_t and future Ns and Qs fixed (so that $\Delta y_t = \Delta N_t$ and $\Delta y_{t+1} = -\Delta N_t$). Using $\Phi_N = -\beta P_t(N_t/Q_{t+1})^{-\phi}$, we have

$$0 = c_0(1 - R_{1t}) + \sum_{j=1}^{11} c_j(D_{jt} - R_{1t}D_{j,t+1}) + E_t\left[\sum_{j=1}^{m} \gamma_j(w_{jt} - R_{1t}w_{j,t+1}) + b(y_t - R_{1t}y_{t+1}) + \beta_1(\Delta y_t - 2R_{1t}\Delta y_{t+1} + R_{2t}\Delta y_{t+2}) + a_0 + \sum_{j=1}^{11} a_j D_{jt} - \beta P_t(N_t/Q_{t+1})^{-\phi}\right] + \eta_t - R_{1t}E_t\eta_{t+1} + \nu_t. \tag{7}$$

Equation (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_t fixed so that sales also increase by one unit, and keeping y_{t+1} and Q_{t+1} fixed) increases the current cost of adjustment (by $\beta_1\Delta y_t$), but it 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.

Equation (7) describes the tradeoff between selling out of inventory versus producing, holding Q fixed. To see this, move $a_0 + \sum_j a_j D_{jt} - \beta P_t(N_t/Q_{t+1})^{-\phi}$ to the left-hand side. The equation then says that net marginal convenience yield (the savings in marketing cost over the coming period from having another unit of inventory) must equal the expected change in production cost (the increase 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. This expected change in cost may be due to expected changes in factor prices ($R_{1t}E_t w_{j,t+1}$ may differ from w_{jt}), expected increases in cost due to convexity of the cost function, and changes in expected adjustment costs. Again, the error terms in (7) represent the unexplained parts of marginal production and marketing costs.

Equation (7) includes marginal convenience yield, so estimation of that equation can provide estimates of the parameters β and ϕ of the convenience yield function. Miron and Zeldes and Ramey estimate the parameters of this function (which they constrain to be linear) just this way. However, we can use the fact that net marginal convenience yield, $\psi_t = -\Phi_N - a_0 - \sum_j a_j D_{jt} - \nu_t$, can be inferred from futures prices. Using (1a) with a one-month futures price replacing the forward price gives the additional equation

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

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

Unlike the models of Bresnahan and Suslow (1985) and Deaton and Laroque (1992), in this model inventories are always positive, because $\psi_t \rightarrow \infty$ as $N_t \rightarrow 0$. This can be viewed as an approximation of the model, but a reasonably good one. One might ask whether stockouts in fact occur, even though we never observe zero inventories in the data. For the homogeneous and clearly defined commodities studied here, extremely liquid futures (and forward) markets make this unlikely. Any firm can easily buy or sell inventory through these markets. If some firms find their inventories low relative to their delivery commitments, they can (and in practice do) buy spot (bidding up the marginal convenience yield in the process). Also, as Kahn (1992) points out, the use of inventory to avoid stockouts is compatible with a very low probability of a stockout actually occurring (but a very high cost if it does occur).

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 (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:

$$\begin{aligned} -R_{1t}F_{1,t} + P_t = c_0(1 - R_{1t}) + \sum_{j=1}^{11} c_j(D_{jt} - R_{1t}D_{j,t+1}) + E_t\Bigg[\sum_{j=1}^{m} \gamma_j(w_{jt} - R_{1t}w_{j,t+1}) \\ + b(y_t - R_{1t}y_{t+1}) + \beta_1(\Delta y_t - 2R_{1t}\Delta y_{t+1} + R_{2t}\Delta y_{t+2}) \Bigg] + \eta_t - R_{1t}E_t\eta_{t+1}. \end{aligned} \tag{7a}$$

Note that this also eliminates inventories, N_t, as a variable in the model. Estimation of (6) and (7a) 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 (equations (6), (7), and (8) and equations (6) and (7a)) and the dataset.

□ **Estimation.** A natural estimator for a 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 equations (6), (7), and (8) using Hansen's (1982) generalized method of moments (GMM) procedure.

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, marketing, and storage costs. When estimating the model, actual values for variables at time $t + 1$ and $t + 2$ are used in place of expectations, which introduces expectational errors. For example, (6) becomes

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

Similarly, (7) will have a composite error term $\eta_t - R_{1t}\eta_{t+1} + \nu_t + \epsilon_{2,t+1} + \epsilon_{2,t+2}$. As with most Euler equation models, the expectational errors are needed for identification. Without them, the full model, i.e., (6), (7), and (8), would be underidentified.

Under rational expectations, the errors $\epsilon_{1,t+1}$ and $\epsilon_{2,t+2}$ (and the corresponding errors for (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. Also, errors may be serially correlated. Hence, I use as instruments only variables that 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 producer price index (PPI), the rate of growth of the Standard and Poors 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 for heating oil.

As Hansen and Singleton (1982) show, the minimized value of the objective function times the number of observations provides a 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. This statistic is used to test the model's overidentifying restrictions, and hence the hypothesis that agents are optimizing with rational expectations.

□ **Data.** 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. For copper, 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.[11] For lumber, y_t is monthly production and N_t is end-of-month inventories of softwood lumber. Units are millions of board feet.[12] For heating oil, y_t is monthly production and N_t is end-of-month inventories of distillate (no. 2) fuel oil. Units are millions of barrels.[13]

Unit sales for each commodity is calculated from unit production and end-of-month inventories using (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 1%.

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

[11] 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.

[12] Source: National Forest Products Association, *Fingertip Facts and Figures*. Most lumber consumed in the United States is softwood (pine and fir). Futures contracts are traded on the Chicago Mercantile Exchange.

[13] Source: U.S. Deparment of Energy, *Monthly Energy Review*, various issues.

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_{1t} should reflect the rate used to discount nominal cash flows at time t. For (8), which is an arbitrage relationship, this should clearly be the risk-free rate, e.g., the nominal Treasury bill rate. For (6) and (7), however, the rate should include a premium reflecting the systematic risk associated with production cost. Unfortunately, this risk is likely to vary across the components of cost (in the context of the Capital Asset Pricing Model, it depends 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. (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 the equity or debt of the firm.) I therefore ignore systematic risk and use the nominal Treasury bill rate, measured at the end of each month, to calculate R_{1t} and R_{2t}.

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 it 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 expiring in month t. This also has problems. First, the spot contract sometimes expires before the end of the month. Second, open interest in the spot contract (the number of contracts outstanding) falls sharply as expiration approaches and longs and shorts close out their positions, so by the end of the month there may be few spot transactions. Finally, for many commodities, active contracts do not exist for each month.

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:

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

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_{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. Equation (10) is used to construct a series for P_t. Finally, the commodity term structure is also used to infer the 30-day net marginal convenience yield by replacing F_{1t} on the left-hand side of (8) with $P_t(F_{1t}/P_t)^{(30/n_{01})}$.

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. To check that such errors are small, I compared these prices to actual spot contract prices for copper (available for 200 of my original 224 observations) and for lumber (available for 114 observations), and found the series to be very close.[14] (No spot contract prices were available for heating oil.) Finally, I constructed new price series for copper and lumber using the spot contract price when available and the imputed price otherwise, and used them to reestimate the model. The results were very close to those reported below.

[14] The root-mean-square and mean percent errors for the two series are, respectively, 1.21% and −.12% for copper and 3.99% and .39% for lumber. The simple correlations are .998 for copper and .983 for lumber.

TABLE 3 **Estimation of Equations (6), (7), and (8)**

Parameter	Copper	Lumber	Heating Oil
γ_1	−.3043 (.0650)	−1.741 (.6273)	.7406 (.4182)
γ_2	−.5916 (.6165)	−16.953 (7.878)	4.108 (2.441)
γ_3			−.0113 (.0531)
b	−.0000021 (.0000028)	−.00038 (.0027)	.1551 (.0343)
β_1	−.0000011 (.0000012)	−.00004 (.00093)	−.0689 (.0124)
β	.0133 (.0011)	.2378 (.0407)	.0854 (.0096)
ϕ	.7551 (.0710)	3.092 (.5230)	1.420 (.1555)
a_0	.3795 (.0588)	7.510 (.9113)	.8913 (.4526)
a_1	.0916 (.0427)	2.107 (.5476)	.0513 (.2645)
a_2	.0134 (.0361)	.9106 (.5252)	1.796 (.2569)
a_3	.0348 (.0391)	.4467 (.6557)	1.552 (.3633)
a_4	−.0266 (.0450)	−.0220 (.6037)	2.309 (.3123)
a_5	−.0506 (.0343)	−1.585 (.5807)	1.620 (.2998)
a_6	−.1186 (.0500)	−.4772 (.7055)	1.750 (.2922)
a_7	−.0400 (.0594)	−1.726 (.7782)	1.991 (.3052)
a_8	.0740 (.0435)	.8569 (.6483)	1.985 (.2737)
a_9	.1150 (.0394)	−1.107 (.6580)	2.045 (.3258)
a_{10}	.0642 (.0399)	−.8144 (.6183)	1.299 (.2957)
a_{11}	−.0622 (.0396)	−.3134 (.5459)	1.075 (.2917)
c_0	100.494 (9.953)	515.926 (165.070)	−152.506 (104.839)
c_1	2.528 (1.005)	−1.325 (2.985)	−3.428 (1.307)
c_2	3.598 (1.515)	−4.837 (3.722)	−4.473 (1.873)
c_3	2.655 (1.586)	−11.275 (4.338)	−1.185 (2.150)
c_4	−1.501 (1.724)	−6.183 (4.772)	−2.240 (2.438)

TABLE 3 Continued

Parameter	Copper	Lumber	Heating Oil
c_5	−2.353 (1.708)	−10.283 (5.046)	−2.714 (2.475)
c_6	−1.008 (1.711)	−10.173 (4.941)	−4.188 (2.499)
c_7	−3.138 (1.732)	−7.784 (4.602)	−1.781 (2.222)
c_8	−4.890 (1.766)	−18.421 (4.617)	.2928 (2.221)
c_9	−4.762 (1.682)	−16.350 (4.257)	.9881 (2.167)
c_{10}	−2.874 (1.473)	−6.732 (3.590)	3.489 (2.077)
c_{11}	−2.324 (1.311)	−8.223 (2.816)	.1761 (1.819)
ρ_1	.9443 (.0309)	.9886 (.0095)	.9962 (.0035)
ρ_2	.9895 (.0152)	.9550 (.0263)	.8962 (.0502)
J	77.43	93.38*	60.38

Note: ρ_1 and ρ_2 are AR(1) coefficients for (6) and (7). J is the minimized value of the objective function times the number of observations, distributed as χ^2 (59) for copper and lumber and χ^2 (52) for heating oil. A * indicates significance at the 5% level. Asymptotic heteroskedastic-consistent standard errors are in parentheses.

5. Results

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

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 reject these restrictions at the 5% 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, none of the commodities exhibit statistically significant costs of adjusting production. For both the full model and for equations (6) and (7a), β_1 is either insignificantly different from zero (for heating oil) or negative. 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). The results for heating oil, however, do provide evidence of rising marginal costs, and the estimates are economically meaningful. For example, $\hat{b}\bar{y}/\bar{P} \approx .15$, so this component of cost accounts on average for some 15% 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 result,

TABLE 4 **Estimation of Equations (6) and (7a)**

Parameter	Copper	Lumber	Heating Oil
γ_1	−.3534 (.0807)	.0882 (.8668)	1.551 (.6878)
γ_2	−.8654 (.4455)	−2.237 (6.100)	.5809 (2.955)
γ_3			−.0484 (.0854)
b	.0000018 (.0000024)	−.0057 (.0023)	.2357 (.0621)
β_1	−.000000052 (.00000083)	.0019 (.00069)	−.1048 (.0208)
c_0	107.779 (14.530)	1373.62 (295.529)	−169.038 (116.862)
c_1	1.918 (.9647)	2.691 (2.913)	−2.957 (1.206)
c_2	5.114 (1.588)	−.9123 (3.475)	−5.860 (1.709)
c_3	4.294 (1.698)	−1.341 (4.027)	−2.835 (2.030)
c_4	1.303 (1.713)	−2.410 (4.446)	−1.111 (2.466)
c_5	.6371 (1.829)	−7.935 (4.682)	−2.472 (2.552)
c_6	1.415 (1.857)	−11.697 (4.610)	−4.581 (2.400)
c_7	−.0618 (1.847)	−13.308 (4.649)	−1.716 (2.378)
c_8	−1.038 (1.769)	−22.296 (4.342)	1.501 (2.424)
c_9	−1.685 (1.495)	−20.210 (4.020)	2.689 (2.350)
c_{10}	−1.617 (1.254)	−11.744 (3.476)	5.224 (2.206)
c_{11}	−1.816 (.9689)	−8.023 (2.956)	2.527 (1.766)
ρ_1	.9694 (.0299)	1.000 (.00082)	.9899 (.0109)
ρ_2	.9757 (.0203)	.9878 (.0155)	.8444 (.0813)
J	43.32	65.26*	34.99

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

together with the numbers in Table 1, suggests that one role of heating oil inventories is to smooth production.

Second, the factor cost variables w_{jt} do little to explain price, and some of the γ_js are negative. Of course these are partial correlations, but in the case of copper, even the simple correlations between the spot price and the two factor cost indices are negative (−.04 and −.07). These indices may simply be too aggregated to capture costs in these industries.

Third, 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 be a role of inventories (at least for heating oil), that role is limited to periods when inventories are at normal-to-high levels. The use of inventories to smooth prices is likewise limited. As Figures 1, 2, and 3 show, sharp price increases are usually accompanied by sharp increases in convenience yield.

These estimates of β and ϕ are a rejection of the quadratic inventory model (and the linear marginal convenience yield it implies) that is central to most studies of manufacturing inventories. While the role of inventories may be different for manufacturing than for commodities, these results at least throw into question the validity of this standard assumption.

I also estimated several alternative versions of the model. First, one might argue that the value of inventory depends not only on expected sales, but also on the next-period expected price (rather than the current price). I reestimated the model substituting P_{t+1} for P_t in the expression for ψ_t in (7) and (8). The results were nearly the same as for the original model.

Second, equations (6), (7), and (8) and equations (6) and (7a) were estimated using quarterly data, on the grounds that intertemporal optimization may be feasible only over time horizons longer than one month. The results were not very different. Third, cubic terms were added to the production cost function, on the grounds that there may be 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 equations (6), (7), and (7a), but estimates of this parameter were insignificant and/or quantitatively unimportant.

6. Conclusions

■ Unlike models of manufacturing inventories, this article has 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 it may explain the rejection of the overidentifying restrictions for lumber as well as the failure to find any evidence of adjustment costs or, for copper and lumber, increasing marginal costs.

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. Then 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.

These caveats aside, the results suggest a production-smoothing role for inventories only in the case of heating oil. And even for this commodity, 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 the more important role for inventories is to reduce marketing costs by facilitating production and delivery schedules and avoiding stockouts. The importance of this role is clear from the fact that producers keep inventories on hand at an effective cost that is sometimes very high.

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 (1a) with negligible measurement error. Ignoring systematic risk, the difference between the futures price, $F_{T,t}$, and the forward price, $f_{T,t}$, is

$$F_{T,t} - f_{T,t} = -\int_t^T F_{T,w}^{\cdot}\text{cov}[(dF_{T,w}/F_{T,w}),(dB_{T,w}/B_{T,w})]dw, \tag{A1}$$

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

$$\% \text{ bias} \approx \bar{r}[\widehat{\text{cov}}(\Delta r/r, \Delta F/F)], \tag{A2}$$

where $\bar{r}$ is the mean monthly bond yield and $\widehat{\text{cov}}$ 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 hundredth of a cent for a one-month contract.

References

BLANCHARD, O.J. "The Production and Inventory Behavior of the American Automobile Industry." *Journal of Political Economy*, Vol. 91 (1983), pp. 365–400.

——— AND MELINO, A. "The Cyclical Behavior of Prices and Quantities: The Case of the Automobile Industry." *Journal of Monetary Economics*, Vol. 17 (1986), pp. 379–407.

BLINDER, A.S. "Can the Production Smoothing Model of Inventory Behavior Be Saved?" *Quarterly Journal of Economics*, Vol. 101 (1986), pp. 431–453.

——— AND MACCINI, L.J. "The Resurgence of Inventory Research: What Have We Learned?" *Journal of Economic Surveys*, Vol. 5 (1991), pp. 291–328.

BRENNAN, M.J. "The Supply of Storage." *American Economic Review*, Vol. 48 (1958), pp. 50–72.

BRENNAN, M.J. "The Cost of Convenience and the Pricing of Commodity Contingent Claims." In D. Lund and B. Øksendal, eds., *Stochastic Models and Option Values*. Amsterdam: North-Holland, 1991.

BRESNAHAN, T.F. AND SUSLOW, V.Y. "Inventories as an Asset: The Volatility of Copper Prices." *International Economic Review*, Vol. 26 (1985), pp. 409–424.

COX, J.C., INGERSOLL, J.E., JR., AND ROSS, S.A. "The Relation Between Forward Prices and Futures Prices." *Journal of Financial Economics*, Vol. 9 (1981), pp. 321–346.

DEATON, A. AND LAROQUE, G. "On the Behavior of Commodity Prices." *Review of Economic Studies*, Vol. 59 (1992), pp. 1–23.

ECKSTEIN, Z. AND EICHENBAUM, M.S. "Inventories and Quantity-Constrained Equilibria in Regulated Markets: The U.S. Petroleum Industry, 1947–1972." In T. Sargent, ed., *Energy, Foresight, and Strategy*. Washington, D.C.: Resources for the Future, 1985.

EICHENBAUM, M.S. "Rational Expectations and the Smoothing Properties of Inventories of Finished Goods." *Journal of Monetary Economics*, Vol. 14 (1984), pp. 71–96.

———. "Some Empirical Evidence on the Production Level and Production Cost Smoothing Models of Inventory Investment." *American Economic Review*, Vol. 79 (1989), pp. 853–864.

FAIR, R.C. "The Production-Smoothing Model is Alive and Well." *Journal of Monetary Economics*, Vol. 24 (1989), pp. 353–370.

FAMA, E.F. AND FRENCH, K.R. "Business Cycles and the Behavior of Metals Prices." *Journal of Finance*, Vol 43 (1988), pp. 1075–1093.

FRENCH, K.R. "A Comparison of Futures and Forward Prices." *Journal of Financial Economics*, Vol. 12 (1983), pp. 311–342.

HANSEN, L.P. "Large Sample Properties of Generalized Method of Moments Estimators." *Econometrica*, Vol. 50 (1982), pp. 1029–1054.

——— AND SINGLETON, K.J. "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models." *Econometrica*, Vol. 50, (1982), pp. 1269–1286.

HOLT, C.C., MODIGLIANI, F., MUTH, J.F., AND SIMON, H.A. *Planning Production, Inventories, and Work Force*. Englewood Cliffs, N.J.: Prentice-Hall, 1960.

KAHN, J.A. "Inventories and the Volatility of Production." *American Economic Review*, Vol. 77 (1987), pp. 667–679.

———. "The Seasonal and Cyclical Behavior of Inventories." Mimeo, University of Rochester, 1990.

———. "Why is Production More Volatile than Sales? Theory and Evidence on the Stockout-Avoidance Motive for Inventory Holding." *Quarterly Journal of Economics*, Vol. 107 (1992), pp. 481–510.

McCALLUM, B.T. "Competitive Price Adjustments: An Empirical Study." *American Economic Review*, Vol. 64 (1974), pp. 56–65.

MIRON, J.A. AND ZELDES, S.P. "Seasonality, Cost Shocks, and the Production Smoothing Model of Inventories." *Econometrica*, Vol. 56 (1988), pp. 877–908.

PINDYCK, R.S. "The Present Value Model of Rational Commodity Pricing." *Economic Journal*, Vol. 103 (1993), pp. 511–530.

RAMEY, V.A. "Inventories as Factors of Production and Economic Fluctuations." *American Economic Review*, Vol. 79 (1989), pp. 338–354.

———. "Nonconvex Costs and the Behavior of Inventories." *Journal of Political Economy*, Vol. 99 (1991), pp. 306–334.

TELSER, L.G. "Futures Trading and the Storage of Cotton and Wheat." *Journal of Political Economy*, Vol. 66 (1958), pp. 233–255.

THURMAN, W.N. "Speculative Carryover: An Empirical Examination of the U.S. Refined Copper Market." *RAND Journal of Economics*, Vol. 19 (1988), pp. 420–437.

WEST, K.D. "A Variance Bounds Test of the Linear Quadratic Inventory Model." *Journal of Political Economy*, Vol. 94 (1986), pp. 374–401.

WILLIAMS, J. "Futures Markets: A Consequence of Risk Aversion or Transactions Costs?" *Journal of Political Economy*, Vol. 95 (1987), pp. 1000–1023.

WORKING, H. "The Theory of Price of Storage." *American Economic Review*, Vol. 39 (1949), pp. 1254–1262.