Rochester Center for
Economic Research

The Seasonal and Cyclical Behavior of Inventories

Kahn, James A.

Working Paper No. 223
February 1990

University of
Rochester

The Seasonal and Cyclical Behavior of Inventories

James A. Kahn

Rochester Center for Economic Research
Working Paper No. 223

The Seasonal and Cyclical Behavior of Inventories

James A. Kahn

February 1990

Working Paper No. 223

*Department of Economics, University of Rochester, Rochester, NY 14627. I would like to thank Spencer Krane, Jeffrey Miron, Kenneth West and seminar participants at Harvard, Wharton, Rochester, and the 1989 ASSA winter meetings for their comments. I would also like to thank Ray Fair for providing me with his data.

Abstract

Recent research has suggested that the aggregate economy behaves similarly at seasonal and business cycle frequencies. This paper argues that inventory behavior—which accounts for some thirty percent of peak–to–trough changes in GNP—does not fit this pattern. Whereas seasonal production–smoothing appears important for industries with large seasonal fluctuations in sales, there is no evidence of smoothing behavior in those or other industries at lower frequencies.

The paper also offers other evidence against production–smoothing as a significant component of inventory movements, and offers evidence in favor of stockout–avoidance as an alternative theory. The paper derives a number of contrasting implications for the two approaches and then examines data from six 3–digit level industries. The paper concludes that stockout–avoidance accounts for the stylized facts about inventory behavior except that firms may smooth production in response to large seasonal disturbances. Production–smoothing is irrelevant for most seasonal and all business cycle fluctuations.

Recent research has suggested that the aggregate economy behaves similarly at seasonal and business cycle frequencies. This paper argues that inventory behavior—which accounts for some thirty percent of peak–to–trough changes in GNP—does not fit this pattern. Seasonal production–smoothing appears important for industries with large seasonal fluctuations in sales, but there is no evidence of smoothing behavior in those or other industries at lower frequencies.

The paper also contrasts the predictions of the production–smoothing model and the stockout–avoidance model. Despite the large fraction of the peak–to–trough decline in GNP in the U.S. accounted for by changes in inventory investment, there is no consensus even on what the primary motivation is for the holding of inventories by manufacturers and retailers. In fact, the debate within the inventory literature regarding the source of production volatility resembles the broader debate regarding the sources of aggregate fluctuations: The production–smoothing models and stockout–avoidance models emphasize demand uncertainty, while the cost–smoothing and increasing–returns–to–scale models tend to emphasize supply–side or technological factors. Since changes in inventory investment are such a significant aspect of cyclical fluctuations in GNP, one can argue that a better understanding of inventory behavior will lead to a better understanding of aggregate fluctuations.

Using high–quality data on physical quantities at the three–digit industry level, the paper shows that with the exception of those industries that smooth large seasonal fluctuations, production–smoothing simply cannot account for basic facts about inventory behavior, while stockout–avoidance can. Thus any significant role for production–smoothing is clearly limited to high frequency activity and is of little consequence for understanding business cycles. Moreover it would be a mistake to draw general inferences about inventories on the basis of their seasonal behavior.

The paper proceeds as follows: Section 1 provides two simple models—the polar cases of "pure" production–smoothing and "pure" stockout–avoidance—around which the discussions of the data can be organized. Section 2 describes the data. Section 3 analyzes various features of the data in light of the empirical implicationns of the models from Section 1. Section 4 concludes.

1. Background

Models of inventory behavior must confront a number of basic facts. First, the level of inventories in the U.S. economy is approximately three months' worth of sales. A good theory ought to account for the level of inventories as well as the dynamic behavior. Second, there is considerable evidence that inventories do not smooth production. This fact has been documented in terms of the relative variances of production and sales, and in terms of the covariance of inventories with sales. On the other hand, there is room for controversy: Recent research has called into question the evidence about "excess" production volatility (see Fair, 1989, Krane and Braun, 1989, Miron and Zeldes, 1989). Some of the results below will clarify the sense in which inventory behavior does violate the production–smoothing hypothesis.

Several explanations of the primary role of inventories have emerged in recent years as alternatives to the production–smoothing hypothesis. Some (e.g. Blinder, Eichenbaum, et al.) have argued that inventories are used to intertemporally substitute production in the face of cost or productivity shocks. A second approach invokes the idea that inventories play a direct role in production and puts inventories as an argument in the production function (e.g. Christiano, 1988). A third approach is the hypothesis of production–bunching in response to increasing returns to scale (e.g. Ramey, 1988, Cooper and Haltiwanger, 1989). Finally, the speculative, or stockout–avoidance approach assumes that firms hold inventories in response to demand uncertainty and a constraint that sales in a given time period cannot exceed accumulated inventories (see Kahn, 1987, 1989, Christiano and Fitzgerald, 1988.)

Each of these approaches encounters difficulties of its own. The cost shock hypothesis suffers from the problem that attempts to explain inventory movements by measurable changes in costs have not been very successful (see, e.g. Ramey, 1988, Kahn, 1989). While to some extent this poses a problem for *any* rational model of inventory behavior, it must be comparatively worse news for the approach that says cost shocks are the primary cause of fluctuations in inventory investment. It is also bad news for the increasing returns hypothesis, since this explanation also implies a systematic relationship between productivity and inventory accumulation.

The inventories–in–the–production–function approach is difficult to justify on theoretical grounds for most types of inventory–holding. It has arisen largely because of the difficulty of generating significant inventories in calibrated representative–agent real business cycle models (see Christiano and Fitzgerald, 1988). As I will argue below, a more conventional approach in a model with many goods has little trouble generating levels of inventory–holding of the right order of magnitude.

The difficulties with the stockout–avoidance model are primarily technical. It is hard to come up with a tractable model that makes powerful predictions about the data without a considerable set of auxiliary assumptions. In particular it is difficult to nest the stockout–avoidance model within the other traditional approaches, or vice–versa. Nonetheless I will argue that even a relatively simple version of the stockout–avoidance model can explain away a large number of the supposed anomalies associated with inventory data.

Two common implementations of the production–smoothing hypothesis are the class of linear–quadratic models (e.g. Blanchard, 1983, West, 1986) and the partial adjustment model (e.g. Lovell, 1961). The theoretical difficulty with these models is that the behavior predicted by them depends heavily on the arbitrary specification of an inventory "target". The empirical problem with them is that while they fit the data well (see Blanchard, 1983), they do poorly in explaining several very basic facts in the data, among them:

- the high level of inventories
- the procyclical behavior of inventory investment (and the related "excess" production volatility)
- the closely timed pattern of production and sales at seasonal frequencies

This section describes the stockout–avoidance model and contrasts it with the two standard versions of the production–smoothing hypothesis. It turns out that the stockout–avoidance model can rationalize these anomalies at the same time as fitting the data reasonably well.

1.1. The Production–Smoothing Hypothesis

The production–smoothing hypothesis holds that the chief motive for the holding of inventories is to smooth production in the face of variable demand. It does not rule out other motives, but (to have any content) does say they are secondary. Since it is impossible to devise a metric for how much of inventory behavior is described by production–smoothing relative to other explanations, the approach taken here is to describe the "pure" production–smoothing motive and then see how well it does in accounting for inventory behavior. A poor performance does not rule out production–smoothing as *one* motive, but it does rule it out as being of primary importance.

As far as macroeconomics is concerned, the focus of inventory models should be to explain data at "business cycle" frequencies. Notwithstanding Barsky and Miron's (1988) evidence for the similarity of the behavior of many macroeconomic variables at seasonal and cyclical frequencies, the *sine qua non* of any hypothesis about inventory behavior is its ability to describe behavior at business cycle frequencies. One should be cautious of drawing strong inferences about a model's implications for cyclical phenomena on the basis of seasonal observations. This does not necessarily work for or against production–smoothing: It simply says that both production–smoothing and its alternatives should be evaluated primarily by their ability to account for cyclical phenomena.

This section presents a very simple version of a linear–quadratic model of inventory behavior. It represents a special case of models considered by Blanchard (1983) and West (1986) among others. Firms solve an infinite–horizon cost–minimization problem treating sales as an exogenous stochastic process. In order to focus exclusively on the production–smoothing hypothesis, the firm's cost in period t is assumed to be a function only of the level of production and inventories in period t, i.e.

$$c(Y_t,N_t) = \frac{h}{2} Y_t^2 + \frac{d}{2}N_t^2, \tag{1.1}$$

where the variables are defined as before, and h and d are both positive. The problem the firm faces is to choose a sequence of production to minimize:

$$E_{t-1}[\sum_{s=t}^{\infty} \beta^{s-t}c(Y_t,N_t)] \tag{1.2}$$

subject again to the identity

$$N_t = N_{t-1} + Y_t - Z_t, \tag{1.3}$$

where Z_t is a stochastic process that can be arbitrary except that it cannot be expected to grow at a rate that exceeds $1/\beta$–1. Standard techniques yield an Euler equation condition

(1.4) $$(1+\gamma)Y_t = E_{t-1}(Z_t) - N_{t-1} + \gamma\beta E_{t-1}(Y_{t+1}).$$

where $\gamma = h/d$. With a little algebra this can be expressed as a stochastic difference equation for inventories:

(1.5) $$[1+\gamma(1+\beta)]N_t = \gamma N_{t-1} + \gamma\beta E_t(N_{t+1}) + -\gamma Z_t + \gamma\beta E_t(Z_{t+1}).$$

This equation can be solved by standard techniques to get feedback rules for inventory investment. The equation can be expressed in the form

(1.6) $$(1+\theta L + (1/\beta)L^{-1})E_t N_{t+1} = Z_t/\beta - E_t(Z_{t+1}).$$

where L is the lag operator (e.g. $LE_tZ_{t+s} = E_tZ_{t+s-1}$), L^{-1} the inverse lag or forward operator, and $\theta = [1+\gamma(1+\beta)]/\gamma$. This yields a first–order difference equation of the form

(1.7) $$N_t = \lambda_1 N_{t-1} - \lambda_2^{-1} E_t\{\sum_{s=t}^{\infty} \lambda_2^{-(s-t)}[Z_s/\beta - Z_{s+1}]$$

where λ_1 and λ_2 are roots of $(1+\theta L + \delta L^{-1})$ and satisfy $\beta\lambda_1 = \lambda_2^{-1}$. This can be simplified to get

(1.8) $$N_t = \lambda N_{t-1} - \lambda Z_t + (1-\lambda)E_t\{\sum_{s=t}^{\infty}(\beta\lambda)^{s+1-t} Z_{s+1}\}$$

(where $\lambda \equiv \lambda_1$). This in turn implies the following rule for production:

$$(1.9) \qquad Y_t = -(1-\lambda)N_{t-1} + (1-\lambda)E_t\{\sum_{s=t}^{\infty}(\beta\lambda)^{s-t}Z_s\}.$$

Equation (1.9) makes intuitive sense: As $\lambda \to 1$ (which occurs as $\gamma \to \infty$), $Var(Y) \to 0$; as $\lambda \to 0$ (i.e. as $\gamma \to 0$), $Var(N) \to 0$ and $Var(Y) = Var(Z)$.[1] Equation (1.9) also shows why some authors (e.g. Abel, 1985, Blinder, 1986) have defined production–smoothing to occur if $|dY_t/dN_t| < 1$.

This model has several important implications that will be discussed in more detail below. These include:

1. $Var(Y) + (1/\gamma)Var(N) \leq Var(Z)$ for any stationary Z process.
2. Seasonals in Z => Seasonals in Y are a moving average of seasonals in Z.

Fair (1989) also emphasizes the effects of future expected sales on current production as an implication of production–smoothing (though in the context of a stock–adjustment model like that described below). It will be clear, though, that such a finding is also consistent with stockout–avoidance. It should also be pointed out that there is no clear implication for inventory levels, nor is there a clear structural interpretation of error terms in estimated equations.

1.2. The Stockout–Avoidance Model

This section describes a simple version of the stockout–avoidance model. It is closely based on Kahn (1989), and again is a polar case: Costs are assumed to be linear, so there is no inherent desire smooth production. The key assumptions are, first, that when the firm makes production decisions it is uncertain about the state of

[1]The implications are expressed in terms of variances rather than levels because in practice the variables corresponding to Y, Z, and N are prefiltered (demeaned, detrended, seasonally adjusted, etc.), so the interpretation of the levels is less standard than that of the variances.

demand; and, second, there is a non–negativity constraint on inventories. Consequently sales in any one period are constrained by the firm's inventory stock. In deciding how much to make available for sale the firm must therefore trade off the cost of stocking out (lost profitable sales) and the cost of carrying inventory stock over to the subsequent period.

For simplicity the output price p_t is assumed for now to be predetermined, and *potential* sales X_t ("demand") conditional on price are assumed to be stochastic and exogenous. While these assumptions are somewhat restrictive, their purpose is primarily just to facilitate the exposition; the qualitative theoretical results are robust to variations in the economic environment. Another simplifying assumption is that the firm cannot backorder its excess demand. Again the predictions of the model are not sensitive to allowing backorders (see Kahn, 1987) so long as there is some cost—even if the cost is only delayed revenue.

Upon observing the inventory stock at the end of period t–1, denoted N_{t-1}, and taking account of any other information regarding future demand, the firm chooses production for period t, Y_t. Sales Z_t satisfy

$$Z_t = \min\{N_{t-1} + Y_t, X_t\}. \qquad (1.10)$$

If demand turns out to be less than the total stock made available by the firm ($N_{t-1} + Y_t$), the firm carries the remainder into the next period as inventory $N_t = N_{t-1} + Y_t - Z_t$. If, on the other hand, the firm "stocks out", the result is lost sales, and $N_t = 0$.[2] It will be convenient to assume that demand can be decomposed into a deterministic component plus a stochastic component with multiplicative uncertainty, i.e.

$$X_t = X_t^d + \tilde{X}_t, \qquad (1.11)$$

[2]Kahn (1987) allows for backorders, with similar results.

where $\tilde{X}_t = E_{t-1}(\tilde{X}_t)U_t$, and where U_t is an i.i.d. random variable with $E(U_t) = 1$. The deterministic part of X can represent deterministic seasonality, backorders, or any other part of demand that is known for certain. The quantity X_t should be thought of as potential sales conditional on price at time t. It differs from actual sales to the extent there are stockouts.

With each decision the firm makes its objective is to maximize expected profits over an infinite horizon from that moment forward. The stock of goods the firm has available for sale each period is equal to its inventory as of the end of the previous period plus the current period's production. Actual sales equal either that period's total demand or the amount available for sale, whichever is smaller. The firm discounts profits in period t+1 as of period t by a discount factor denoted β_t.

The production costs $c(Y_t;w_t)$ facing the firm are assumed to be linear in production, i.e.

$$c(Y_t;w_t) = c_tY_t. \tag{1.12}$$

where c_t is a scalar that varies stochastically over time and w_t represents a vector of factor prices. This cost function would be valid if the firm had a constant returns to scale technology and were a price taker in factor markets.

The firm stocks out if X_t exceeds the stock available for sale $N_{t-1}+ Y_t$, whereas if $N_{t-1}+ Y_t > X_t$, the leftover stock is carried over into the next period. It will be convenient for the exposition of the model (though not necessary for the results) to assume that at the time the firm chooses Y_t it knows p_t and c_t and has a forecast of c_{t+1} denoted E_tc_{t+1}, and to assume that the innovation in c_{t+1} is independent of the demand shock U_{t+1}.

To get an explicit characterization of the target, given the assumption of multiplicative uncertainty in demand, we can solve the firm's profit maximization problem

$$\max E_{t-1}\{\sum_{s=t}^{\infty}\beta^{s-t}[p_t Z_t - c_t Y_t]\}, \tag{1.13}$$

subject to (1.10), (1.11), and (1.12), using standard techniques. Let G denote the cumulative distribution function (assumed to be continuous and strictly monotonic) for the forecast error U_t, and let $A_t \equiv N_{t-1} + Y_t$, the stock available for sale. The first–order condition for the maximization problem is

$$(p_t - c_t)(1 - G[(A_t - X_t^d)/E_{t-1}(\tilde{X}_t)]) - (c_t - \beta_t E_t c_{t+1})G[(A_t - X_t^d)/E_{t-1}(\tilde{X}_t)] = 0. \tag{1.14}$$

We can then solve (1.14) for the optimal choice of A_t:

$$A_t^* = X_t^d + K(p_t/c_t, \beta_t E_t c_{t+1}/c_t)E_{t-1}(\tilde{X}_t), \tag{1.15}$$

where $K() = G^{-1}[(p_t - c_t)/(p_t - \beta_t E_t c_{t+1})]$. Equation (1.15) says that the amount the firm makes available for sale in period t equals a multiple K of its expected new demand. Since G^{-1} is non–negative and monotonically increasing, K (which is closely related to the inventory–sales ratio, as will be seen below) is non–negative and depends positively in its two arguments, as one would expect. The magnitude of K reflects the opportunity cost of stocking out relative to the cost of carrying inventory over into the next period. A higher markup means a higher cost of stocking out, while a higher value $\beta_t E_t c_{t+1}/c_t$ means production costs are lower in t relative to t+1, and therefore implies a smaller cost of carrying inventories into the next period. K also depends in a more complicated way on the dispersion of U. The same intuition applies to the behavior of production, given expected demand and initial inventories.

As an aside, it is interesting to compare the target specification in (1.15) with that used by, for example, Blanchard (1983) and West (1986) in the linear–quadratic framework:

$$N^*_{t-1} = KE_{t-1}(Z_t) \tag{1.16}$$

where N^*_{t-1} is desired inventories, which in their specifications can differ from actual N_{t-1} because of adjustment costs. There are minor differences: Equation (1.4) includes production in period t as part of the stock, and has expected demand rather than expected sales. But the main difference is that in (1.15) K is a function of cost variables, while in (1.16) it is assumed to be constant. The assumption in those models that K is a fixed parameter necessarily implies that serially correlated deviations of N from its target either fall into the residuals or are attributed to adjustment costs. Moreover, while it is difficult to say for sure what the effect of misspecifying the target is on econometric estimates of linear–quadratic models, it is at least possible that it is responsible for rejections such as those found by West (1986).

We can solve for sales in terms of demand expectations and innovations. Letting $K_t \equiv K(p_t/c_t, \beta_t E_t c_{t+1}/c_t)$, and recalling that $Z_t = \min\{A_t, X_t\}$, we have:

$$Z_t = X^d_t + E_{t-1}(\tilde{X}_t)V_t \tag{1.17}$$

where $V_t \equiv \min\{K_t, U_t\}$. Straight–forward substitutions yield

$$Y_t = X^d_t + E_{t-1}(\tilde{X}_t)K_t - E_{t-2}(\tilde{X}_{t-1})(K_{t-1} - V_{t-1}) \tag{1.18}$$

and

$$N_t = E_{t-1}(\tilde{X}_t)(K_t - V_t). \tag{1.19}$$

Equation (1.17) says that actual sales equal total demand or the stock available, whichever is smaller. A stockout occurs if $U_t > K_t$. Equation (1.18) can be rearranged to show that production in t is the sum of three components: Certain demand in t, the stochastic component of sales in t–1 (to replenish inventories), and the change in the desired stock available for sale.[3] Kahn (1987) shows that that positive serial correlation in demand implies that Var(Y) > Var(Z) even if costs are constant. Equation (1.19) shows that inventories do not reflect deterministic components of demand. Thus, for example, inventory/sales ratios will tend to be lower when a greater proportion of sales are attributable to backorders, seasonal demand, and the like.

Note that (1.17) and (1.18) imply that seasonals or other certain components of demand should be matched in both Y and Z, regardless of whether there are accompanying movements in output prices. On the other hand, seasonals in K are likely to result in movements in Y that are magnified relative to those in Y. To see this, suppose that $X_t^d = 0$, that p_t is constant, and that there are deterministic seasonals in K. Consider two polar cases: If $V_t = U_t \ \forall t$, then no stockouts ever occur, and there are no seasonals at all in Z. If $V_t = K_t$, then seasonals in Z match those in Y. In the intermediate case stockouts occur occasionally, as the theory predicts, and seasonals in Y are larger than seasonals in Z. If there are corresponding price movements, then these will induce some seasonality in Z, but presumably no more

[3]This derivation assumes that the solution is always at an interior point. Kahn (1988) takes explicit account of the non–negativity constraint on production.

than in Y. (Otherwise we would have the perverse result that firms would accumulate inventories when costs are high because the price goes up more than proportionately.)[4]

Overall the results contained in equations (1.15) and (1.17) are simple and intuitive. They say that the quantity A_t^*, the firm's planned stock available for sale in period t, is the product of two factors: costs in period t (absolutely and relative to period t+1), and expected demand in period t. Using small letters to denote the log of the corresponding capital letter (e.g. $z_t \equiv \log(Z_t)$), we have from (1.15) and (1.17)

$$
\text{(1.20)} \qquad z_t = a_t - k_t + v_t
$$

$$
= \begin{cases} a_t - k_t + u_t & \text{no stockout} \\ a_t & \text{stockout} \end{cases}
$$

where, again, $v_t \equiv \log(\min\{K_t,U_t\})$, and U_t is orthogonal to anything known prior to time t—which includes the other right–hand–side variables. Kahn (1989) estimates a generalized form of (1.20) to test this model on the automobile industry data and finds that while there are specific statistical rejections, for the most part the results are supportive. In regressions of sales on a_t and determinants of k_t, the coefficient on a_t is very close to one, there is little evidence of serial correlation in the residual (which the model implies is a truncated forecast error), and cost variables (except for the interest rate) enter as predicted.

Note another implication of the model: That $v_t = z_t - a_t$ reflects truncation due to stockouts.[5] If the underlying disturbances are symmetric (in logarithms), then this

---

[4]Another case to consider is exogenous changes in output prices, which might arise when the firm or industry being considered is small relative to the market, and has costs or demand that is not very correlated with other participants. This would look like seasonality in K.

[5]Stockouts do not necessarily correspond in the data to zero inventories. Apart from issues of timing, the quantities we observe are inevitably aggregates, so that there generally will be stockouts of some goods and not of others at any point in time.

observable quantity should exhibit negative skewness. For a given forecast error distribution we would expect the skewness to be more negative the lower the inventory–sales ratio.

We can summarize the basic implications of the stockout–avoidance model as follows:

1. Positive serial correlation in X => Var(Y) > Var(Z).
2. a. Seasonals in X => matched seasonals in Y and Z.
   b. Seasonals in K => seasonals in Z are damped relative to seasonals in Y.
3. Cross–sectional implications for inventory levels.
4. Negatively skewed structural forecast errors due to truncation.

These will be discussed in more detail below.

2. The Data

This section describes the data and discusses several stylized facts that a model of inventory behavior should be expected to address. The data come from several sources. Monthly auto industry data from 1966 to 1979 were provided by Olivier Blanchard. I updated this data set through 1983 as described in Kahn (1988).[6] Ray Fair has kindly provided me with his monthly data from seven three–digit industries: cigarettes, cigars, cement, tires, slab zinc, copper, and lead. Of these I focus only on the last five, because of some problems with the tobacco industry data. Cigarette inventories and sales appeared to have different trends (which is difficult to reconcile with any model), while the data on cigars showed evidence of the effects of aggregation over different sizes that made interpretation somewhat difficult. Except for the cement industry (for

[6]Kahn (1989) examines these data disaggregated to the level of divisions. The characteristics of the industry aggregates are similar to those of the individual divisions.

which the data are available only through 1964) the Fair data are available from 1947 to 1987.

One important difference between the automobile industry data and the Fair data is that the automobile sales data represent final sales to consumers, while the other sales data are really shipments to retailers or other manufacturers. Since the latter may themselves hold inventories (either as retail inventories or materials stocks), we actually only see one horizontal slice of the whole picture. Economically it makes no difference whether slab zinc inventories are held by producers as finished goods inventories or by manufactureres as materials inventories, yet the data consist only of the former.

Another difference is that the automobile industry has an annual product cycle that makes it difficult to smooth production at business cycle frequencies. The annual cycle also contributes to production volatility (and on a lesser scale to sales volatility) because of the annual mid–summer retooling period. But while it is therefore not surprising to find that seasonally adjusted production is not smoother than seasonally adjusted sales, the product cycle cannot explain a failure to find smoothing at seasonal frequencies (except for the contribution of the retooling period), nor can it explain why seasonally adjusted production actually varies much more than seasonally adjusted sales.

What distinguishes this entire collection of data from the standard Commerce Department data used in many studies is that they are data on physical quantities, and thus are not subject to errors from the use of price deflators. (See Miron and Zeldes (1988) for a discussion of problems with the Commerce Department data). For most of the analysis the data have been pre–filtered as follows: The automobile industry data (which show little evidence of trends) are either raw (except for an adjustment to control for the number of days in the month) or additively seasonally adjusted with month dummies; the other data have a common log–linear trend removed, and are analyzed with and without seasonal adjustment.[7] The data are adjusted for days in the month as in Fair (1989). The results are not sensitive to this adjustment.

3. Stylized Facts

This section describes some general features of the data that a model of inventory behavior should be expected to address. This will lead to some natural tests or calibration exercises for these models that will be discussed in the remainder of the section. Table 1 provides summary statistics for a number of variables that will be discussed below. Note that the variance of production exceeds the variance of sales (in some cases by a considerable margin) in the seasonally adjusted data for all six industries, while in some cases the unadjusted data indicate production–smoothing. This divergence between behavior at seasonal and nonseasonal frequencies is of crucial importance for understanding the nature of inventory behavior, as will be discussed below. The remainder of this section will discuss these and other facts in light of the models introduced in Section 1.

3.1 The Level of Inventories

One difficulty with the L–Q model emerges immediately: The steady state level of inventory holdings in this model is easily seen to be negative for any permissible Z_t process. As indicated in Table 1, however, finished goods inventories are typically between 1 and 4 months worth of sales. This shortcoming can be remedied with a variant of the cost function (1.19), e.g. $c(Y_t,N_t) = \frac{c}{2} Y_t^2 + \frac{d}{2}(N_t - N_t^*)^2$, where $N_t^* >$ 0 and could be constant or time–varying, possibly a function of the sales process. This is not, however, a very satisfactory fixup of the production–smoothing model. First, it implies that the marginal cost of holding inventories is negative at low levels of N_t.

[7] West (1987) shows that the presence of a common log–linear deterministic trend in inventories, sales, and production is consistent with the linear–quadratic model.

Second, it is not clear where N_t^* comes from, and if it is time–varying what its determinants are. In practice the only argument for such a specification is that there are stockout costs which make N_t^* a function of expected sales as in the partial–adjustment model. But then we no longer have a pure production–smoothing motive for inventories. The best one can say about the production–smoothing model regarding inventory levels is that it is silent.

It could be argued that inventories are positive simply because they have to be positive, and that the correct way to work out the production–smoothing model would be to incorporate a non–negativity constraint into the problem. This argument fails on two counts: First, positive inventories clearly are *not* a technological necessity; we observe many industries in which backorders are present and inventories are negligible or zero. Second, as we have seen, such a non–negativity constraint generates the stockout–avoidance motive for holding inventories even without any desire to smooth production (see also Kahn, 1987), provided the firm does not know the state of demand when it makes its production decisions.

It would perhaps be useful to work out a model in which a firm does know the state of demand when it makes its production decisions, has strictly convex costs, and must hold non–negative inventories. It seems unlikely, however, that such a model could rationalize the levels of inventories we observe with the very slight degree production–smoothing (in the sense of $Var(Y) < Var(Z)$) that takes place. (See Section 3.3 for an elaboration of this argument.)

Given the explicit non–negativity constraint on inventories, the stockout–avoidance model can obviously rationalize positive inventory levels without resort to an *ad hoc* target level, though the predictions of the model can be sensitive to assumptions about the length of the decision period. Suppose for simplicity we take the case in which costs, markups, and interest rates are fixed, and assume a lognormal multiplicative demand process. Equations (1.4) and (1.6) then imply that the formula for the average inventory sales ratio in this model is

$$(3.1) \qquad E(N/Z) = G^{-1}\left[\frac{p/c-1}{p/c-\beta}\right]e^{\sigma^2/2} - 1,$$

where G^{-1} is the inverse of the c.d.f. of the lognormal distribution, σ^2 is the log forecast error variance, and p, β, and c are as defined earlier. With a quarterly decision period, a real interest rate of 2 percent annually (β=0.995), and a 10 percent markup (p/c=1.1), the producer will have stock available to satisfy the 95th percentile of demand. This would generate an inventory–sales ratio of approximately one month with a 20 percent standard deviation of demand. With a monthly decision period (so that $\beta \cong 0.9984$) the producer will have stock available to satisfy approximately the 98.4th percentile, which corresponds to only about half a month of sales. Greater demand uncertainty, a lower discount rate, or a higher markup would generate larger inventory sales ratios. Also, regardless of the standard deviation, positive kurtosis (i.e. fatter tails) would imply greater inventory/sales ratios as well.

It might appear that the simple stockout–avoidance model has some difficulty generating observed inventory/sales ratios. In fact, however, once one recognizes that even 3–digit level industry data are highly aggregated across goods, and that demand uncertainty with respect to specific goods is potentially much greater than the industry aggregates would suggest, it is no longer so clear that stockout–avoidance alone does not account for inventory levels.[8] Goods are distinguished not only by the usual physical characteristics (size, color, etc., which might not be very important for primary metals and other raw materials) but by location as well.

There are also cross–sectional implications of the stockout–avoidance model that find some support from the inventory/sales ratio data. The stockout–avoidance model

[8]This is why Christiano and Fitzgerald's (1989) inability to generate significant inventory holdings in an aggregate one–good economy is not particularly informative about the stockout–avoidance motive.

would predict that, *ceteris paribus*, inventory/sales ratios for an aggregate of goods would be positively associated with the heterogeneity of the goods. This is because E(N/Z) is driven in part by the uncertainty about demand for the representative homogeneous good (e.g. a tire with specific attributes such as size and color), which will generally be larger than the uncertainty about demand for aggregates.[9] The model also predicts that inventory/sales ratios would be negatively associated with the proportion of sales that are deterministic or known ahead of time (e.g. seasonals, backorders, etc.).

Looking back at Table 1, we see that while the forecast standard errors are roughly comparable in order of magnitude, the inventory/sales ratios in the automobile and tire industries are larger generally larger than in the Cement, Lead, Zinc, and Copper industries. The former group consists of aggregates over relatively heterogeneous goods, while the latter are relatively homogeneous commodities. While the data one would need to control for other determinants of the ratio (the true forecast error distribution, the markup of price over marginal cost, and the length of the decision period) are not readily available, the numbers in Table 1 arguably provide some support for stockout–avoidance. At least one can say that the stockout–avoidance model offers some guidance about inventory/sales ratios, whereas the production–smoothing model is at best silent, and taken literally offers no reason why inventory levels should even be positive.

3.2. Seasonality

The distinguishing feature of seasonality is that in most cases it reflects movements in demand or supply that are more predictable, and usually much larger, than those at lower or so–called "business cycle" frequencies. Focusing on the seasonal

[9]This effect would be mitigated to the extent the aggregates are over goods that consumers consider close substitutes. In the limiting case of perfect substitutes the aggregate would be equivalent to a single homogenous good.

frequencies may therefore provide valuable information about firms' responses to forecastable disturbances. At the same time, although Barsky and Miron (1988) have argued otherwise, it is certainly possible that the sources of disturbances at seasonal frequencies differ from those at lower frequencies, or that the response to the same types of disturbances differs, so that it would be incorrect to extrapolate too much on the basis of results obtained from seasonal movements.

The structural model itself may offer guidance as to the correct handling of seasonality. Suppose we modify the Z_t process to include deterministic seasonality, e.g. $Z_t - s_t = \rho(Z_{t-1} - s_{t-1}) + \epsilon_t$, $s_t = s_{t-k}$, where k is the frequency of observations per year. Either version of the LQ model (with or without Z_t observable when choosing Y_t) would imply a pattern of seasonality for production very different from that of sales. It is easy to see from (1.9) that the seasonals in production would be a moving average of the seasonals in sales if the firm were behaving according to the LQ model —though of course as γ approaches zero the seasonals in production would approach the seasonals in sales.

The stockout–avoidance model has very different predictions about seasonality. As shown above, if the only source of seasonality is deterministic demand fluctuations, then production and sales will have identical seasonal patterns. If demand seasonals are stochastic, then production seasonals will match the forecastable part of sales seasonals. On the other hand, if seasonality occurs on the cost side, then the seasonality in sales may be damped or smoothed relative to the seasonality in production. In short, the pure stockout–avoidance model predicts that seasonals in production should be at least as large as those of sales; larger when caused by seasonals in costs, equal when caused by seasonals in demand.

Miron and Zeldes (1988) show that the Commerce Department data exhibit the property that seasonal patterns in production and sales are very similar. While they argue such findings suggest that production–smoothing may not be very important, one cannot firmly conclude that without looking at inventory behavior as well. After all, if the parameter γ in the linear–quadratic model described earlier is very small, production and sales could have very similar seasonal patterns, and yet inventories would still be serving to smooth production. As we will see, what rules out production–smoothing is the combination of this fact with the fact that inventories themselves exhibit high variance.

Figures 1A and 1B plot estimated seasonals for the six industries.[10] In the auto industry, the seasonals in production and sales match each other very closely, with some of the production seasonals larger than the corresponding seasonsals in sales. These large production seasonals occur in the July–August model year changeover. It is perhaps not too surprising that production falls by less than sales. The same may be said for the declines in December–January (due to holiday vacations), though there may be demand effects here as well. On the whole, the seasonal patterns (except for July and August) match each other very closely. Where they do not match, the production seasonal is larger, arguably due to supply–side factors. This is precisely the prediction of the stockout–avoidance model.

In the other three–digit industries there is also a high correlation between the seasonal patterns, but with some evidence of production–smoothing in Tires and Cement (324). With these two exceptions, the analysis of seasonals corroborates the Miron and Zeldes finding that there is not much production–smoothing going on at seasonal frequencies. The behavior of the seasonals (again with the possible exceptions of Tires and Cement) is, however, perfectly consistent with stockout–avoidance. It is interesting to note that of the five industries, tires and cement are the two with the largest seasonal variation in sales. In the cement industry in particular, seasonality dominates the data (see Figure 3A). Possibly the large seasonals lead to smoothing at seasonal

---

[10]Multiplicative seasonals were estimated using dummy variables, normalized so that they sum to zero. The scale in the graphs in Figures 1A and 1B is logarithmic—for example, the seasonal in tire sales ranges from +14 to –18 percent.

frequencies because of capacity constraints, while such constraints do not lead to smoothing at lower frequencies because of the relatively small magnitude of those fluctuations. This conjecture will be examined further when business cycle frequencies are examined.

3.3. The Variance of Production and the Variance of Sales

West (1986) has shown that the LQ model from Section 1 implies that the variance of production should be no greater than the variance of sales. While Blinder (1986) and others have argued that this is violated in the data, more recent work has questioned this (e.g. Miron and Zeldes (1989), Fair (1989), Krane and Braun (1989)). What needs to be emphasized, however, is that the simple variance inequality $\mathrm{Var}(Y) \leq \mathrm{Var}(Z)$ is a very weak implication of production–smoothing—much stronger statements can be made. In particular, it is a statement that is true for the entire parameter space (e.g. regardless of the value of γ). One can, however, also use the information contained in the variance of inventories to rule out large portions of the parameter space and thereby put a much tighter bound on the ratio of the variance of production to the variance of sales.[11]

The parametric inequality derived by West (under the assumption of trend stationarity of inventories, production, and sales) is

$$\mathrm{Var}(Y) + (1/\gamma)\mathrm{Var}(N) \leq \mathrm{Var}(Z). \tag{3.2}$$

This comes from the argument that it is feasible for the firm simply to set Y_t equal to Z_t and N_t equal to zero. This would result in costs proportional to the right–hand side of (3.2). What the firm actually does should be less costly.

---

[11]This is essentially the point of West's procedure in his 1986 paper.

This inequality can be rearranged to yield a minimum value for γ consistent with the three variances (assuming they are known):

$$\gamma \geq \text{Var(N)}/[\text{Var(Z)} - \text{Var(Y)}]. \tag{3.3}$$

Note that γ can be interpreted as the cost of production variability relative to inventory variability. Thus, for example, for a given degree of production–smoothing (as measured by [Var(Z) – Var(Y)]), a larger variance of inventories indicates a larger value of γ.

We can use this lower bound on γ to derive a more precise variance inequality for production and sales. This will be done on actual data below, but for the sake of illustration suppose $Z_t = \overline{Z}(1-\rho) + \rho Z_{t-1} + \epsilon_t$, with mean $\overline{Z}$ and forecast error variance σ^2_ϵ. Table 2 gives values for γ and the implied ratio Var(Z)/Var(Y), assuming $\beta = .99$:

[Table 2 here]

Thus to get a ratio above 0.9, for example, requires a combination of high serial correlation and a relatively small value for γ. By such calculations we can see that even those industries in which Var(Y)/Var(Z) < 1 there is still production variability far in excess of what the model predicts. Applying this analysis to Fair's data in those cases in which the seasonally adjusted data exhibits production–smoothing (that is, Lead, Zinc, and Copper) we get the following results:

[Table 3 here]

The point is that the variance of inventories is far too large to be consistent with the little bit of "production–smoothing" (i.e. the reduction of the variance of production below the variance of sales) that is taking place at business cycle frequencies. Simply looking at the variances of production and sales alone is not an appropriate test of the production–smoothing hypothesis. This is illustrated in Figures 2A and 2B with data from the Zinc industry. The variance of production is less than the variance of sales, but only very slightly, while inventories vary tremendously. Also note (and see Table 1) that inventories show no tendency to be negatively correlated with sales, as the production–smoothing model would predict.

This analysis can be extended to the slightly more complicated case in which the firm cannot observe Z_t when it chooses Y_t. Analagous to West's procedure in his 1986 paper, we can compare the value of the objective under a feasible suboptimal policy to the value of the objective using actual data. Since the policy of setting $N_t = 0 \; \forall t$ is infeasible when Z_t is not known, however, the simple comparison of variances as in (3.2) is no longer possible. A feasible policy for the firm now is to set $Y_t = E_{t-1}(Z_t) + Z_{t-1} - E_{t-2}(Z_{t-1})$.[12] This implies that $N_t = E_{t-1}(Z_t) - Z_t$. Maintaining the assumption that Z_t is an AR(1) process, we have $\mathrm{Var}(N) = \sigma_\epsilon^2$, and

$$(3.4) \qquad \gamma \geq [\mathrm{Var}(N) - \sigma_\epsilon^2]/[\mathrm{Var}(Z) - 2\rho\sigma_\epsilon^2 - \mathrm{Var}(Y)].$$

This clearly has ambiguous effects on the lower bound on γ, and therefore will not necessarily mitigate a finding of excess production volatility. In fact, in the three cases above this version of the test resulted in an even larger lower bound on γ.

---

[12]Note that simply setting $Y_t = E_{t-1}Z_t$, while feasible, implies that inventories are nonstationary. The proposed policy has the firm make up for the previous period's forecast error.

3.4. Frequency Domain Analysis

There is an easy way to examine the relative variances of production and sales by frequency, to see more directly the properties of production and sales at seasonal and business cycle frequencies. The power spectrum decomposes the variance of a series by frequency. Simply comparing the power of production versus sales at different frequencies indicates whether "smoothing" is going on at those frequencies. While the evidence from Table 1 suggests that production variability is close to or exceeds that of sales even after controlling for seasonality, there are problems with those statistics. In many of the divisions and industries, the data are dominated by seasonal fluctuations. Unless the seasonals are deterministic (and additive or multiplicative), it is very difficult by time domain methods alone to filter out seasonal variation from these series. Time–varying seasonals in the auto industry are documented in Kahn (1989). Figures 3A–C illustrate another case in point: The cement industry. Even small stochastic (or time–varying) elements in the seasonals mean that dummy variable or other parsimonious methods leave much of the variance at seasonal frequencies remaining. Figure 3A depicts the raw production and sales data, which indicate obvious seasonal production–smoothing. Figure 3B shows the detrended and seasonally adjusted (with dummy variables) data. It is clear that much of remaining variance still reflects stochastic seasonals. Nonetheless it would be incorrect to conclude that there is not a lot of cyclical variation in the industry. Figure 3C plots the 12th difference of log production against the NBER reference cycles (1 = peak to trough). This series has mean of 3.86 percent, a standard deviation of 8.65 percent, and a correlation of –0.31 with the reference cycle dummy. Thus simple time domain methods do not adequately control for seasonality, and it is useful, therefore, to move to the frequency domain.

Figure 4 plots the power spectra of the seasonally adjusted production and sales series. (Unadjusted data are, strictly speaking, inappropriate if the seasonals have a deterministic componenent, as they almost certainly do). In all of these figures the plots are from 0 to $\pi/3$, so that the midpoint is the seasonal frequency. For the Auto industry, the figures demonstrate that production–smoothing is taking place at neither seasonal nor business cycle frequencies. For the Tire and Cement industries, the pictures dramatically confirm the results from Table 1: Whatever production–smoothing is taking place is entirely at seasonal (or higher) frequencies. The variance of production at business cycle frequencies is at least as great as that of sales. Though less dramatically, the same can be said for the Lead and Zinc industries. The Copper industry shows no obvious difference between low and high frequencies. The spectra for the cement data (324) also illustrate the point made above that considerable seasonal variance remains after dummy variable adjustment.

3.5. Nonstationarity

There is a problem with looking at variances if we are not sure the data are stationary around a deterministic trend. If in fact the series are nonstationary, looking at sample variances may be misleading. West (1987) suggests an alternative test (for a slightly different hypothesis) that is robust to non–stationarity provided the series are stationary in first–differences. This test works as follows: Note that $\mathrm{Var}(Y) < \mathrm{Var}(Z)$ $\iff \mathrm{Cov}(\Delta N,Z) < -(1/2)\mathrm{Var}(\Delta N)$. (This follows from the identity (1.3).) While both $\mathrm{Cov}(\Delta N,Z)$ and $\mathrm{Var}(\Delta N)$ generally exist when N and Z have unit roots, one cannot use the convential sample analog of $\mathrm{Cov}(\Delta N,Z)$. But we also have $Z_t = \Delta Z_t + \Delta Z_{t-1} + \Delta Z_{t-2} + \ldots$, so that

$$\mathrm{Cov}(\Delta N,Z) = \sum_{s=0}^{\infty} \mathrm{Cov}(\Delta N,\Delta Z_{-s}) \qquad (3.5)$$

Let $\hat{C}_s$ be a consistent estimate of $Cov(\Delta N,\Delta Z_{-s})$. West shows that $\sum_{s=0}^{m}\hat{C}_s$, where $m \to \infty$ as the sample size $T \to \infty$, is a consistent estimator of $Cov(\Delta N,Z)$, provided $m/T^{1/2} \to 0$. One statistic analogous to Var(Y)/Var(Z) would then be

$$-0.5 - \sum_{s=0}^{m}\hat{C}_s/[T^{-1}\sum_{s=0}^{T}(\Delta N_s - E(\Delta N))^2],$$

which gives an estimate of the extent to which $Cov(\Delta N,Z)/Var(\Delta N)$ is less than $-1/2$. This was calculated for the five Fair industries and yielded results qualitatively similar to those in Table 1. The statistics were –0.170, –0.184, 0.156, 0.245, and 0.259 for seasonally adjusted data from Tires, Cement, Lead, Zinc, and Copper. (The value of m for these tests was 15 for the Cement and 20 for the others.) The similarity of the results may be due to the fact that the only series that look like they might be non–stationary are those from the Zinc industry.

Of course it is not clear what implications the production–smoothing model has for numbers such as $Cov(\Delta N,Z)/Var(\Delta N)$ when Z is non–stationary. While West shows that bounds analogous to $Var(Y) < Var(Z)$ can be calculated when the data have unit roots, the derivation of the test in Section 2.3.1 is clearly not valid. This is certainly an area that would benefit from more research, though it is not likely to alter the conclusions in this paper.

3.6 Forecast Error Distributions

The stockout–avoidance model does not have strong predictions about stockouts *per se*. Depending on the values of the determinants of k, stockouts could be so infrequent as to be indetectable in the data, or could be extremely common. Nonetheless the natural place to look for evidence of stockouts in the data on the basis of the model is in the shape of the error distribution. The simplest version of the model (with k constant) predicts that the quantity $z_t - a_t = v_t - k$ should be skewed to the left, since v_t reflects truncation due to stockouts. The idea is that large negative demand disturbances have a greater impact on sales than large positive demand disturbances, because the latter run up against stockouts.

More generally, if one maintains the assumption that the underlying stochastic terms k_t and u_t have symmetric distributions, then a finding that $z_t - a_t$ is skewed to the left would be evidence of stockouts in the data. Table 4 presents skewness statistics for this variable for all six industries, both seasonally adjusted and unadjusted.[13] (Note that the model suggests that seasonality in $z_t - a_t$ largely reflects seasonals on the cost side. The reason for looking at adjusted data is to control for the possibility that cost seasonals might be skewed.) Negative skewness of $z_t - a_t$ shows up in five of the six industries. Moreover, the exceptional behavior of the Tire industry data is entirely attributable to a single stretch or positive outliers in 1950–51.[14] If one starts the sample at a an arbitrary later date the positive skewness disappears. For example, starting in 1954 the skewness statistic for tires is –0.229, significant at 10 percent. While it is difficult to know whether the underlying disturbances are symmetric, absent any compelling reason to think that they would be skewed in such a way as to produce negative skewness in $z_t - a_t$, this finding provides additional support for the view that stockouts are a relevant consideration for these industries.

4. Conclusions

The results in the paper demonstrate that the production–smoothing model cannot account for the stylized facts about the cyclical behavior of inventories, while the stockout–avoidance approach shows considerable promise. Data from 3–digit industries show some evidence of smoothing at seasonal frequencies (only in the two industries

[13]Kahn (1989) also finds evidence of skewness with disaggregated automobile industry data.

[14]I have not been able to determine whether the unusual behavior of the data is the result of a strike, or because of the Korean war, or some other factor.

with the largest seasonals in sales), but none whatsoever at cyclical frequencies. Large seasonal variations may lead to smoothing, due perhaps to capacity constraints or other strict convexities that only kick in for very large deviations from the mean. Since business cycle fluctuations are much smaller than seasonal fluctuations, we would then see some smoothing at seasonal frequencies but none at business cycle frequencies.

The results suggest important modifications to the analyses of Barsky and Miron (1989) and of Miron and Zeldes (1988). Barsky and Miron argue that seasonal cycles look a lot like business cycles. Miron and Zeldes look at production and sales (2–digit Commerce department data) and argue that the seasonals of the two series are very similar. We have seen that when seasonals are very large there may be production–smoothing at seasonal frequencies that is not present at business cycle frequencies. This suggests that if we want to understand the role of inventories in business cycles it may be important to filter out seasonal patterns rather than extrapolate from them. While it is true that to explain non–seasonally–adjusted data the stockout–avoidance model may need to be modified to have some sort of capacity constraint that binds for industries with very large seasonal variations (or a convexity in costs that only kicks in with large variations), the need for such a modification does not alter the view that production–smoothing is not important at all at business cycle frequencies.

References

Abel, A. B., "Inventories, Stock–Outs, and Production Smoothing," *Review of Economic Studies 52* (1985), 283–293.

Barsky, R., and J. Miron, "The Seasonal Cycle and the Business Cycle," *Journal of Political Economy 97* (1989), 503–534.

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

Blinder, A., "Can the Production Smoothing Model of Inventory Behavior be Saved?" *Quarterly Journal of Economics 101* (1986), 431–454.

Christiano, "Why Does Inventory Investment Fluctuate So Much?" *Journal of Monetary Economics 21* (1988), 247–280.

Christiano, L. and T. Fitzgerald, "The Magnitude of the Speculative Motive for Holding Inventories in a Real Business Cycle Model," F.R.B. of Minneapolis Working Paper #419, 1988.

Cooper, R., and J. Haltiwanger, "Macroeconomic Implications of Production Bunching: Factor Demand Linkages," mimeo, 1989.

Eichenbaum, M., "Rational Expectations and the Smoothing Properties of Inventories of Finished Goods," *Journal of Monetary Economics 14* (July 1984), 55–70.

Eichenbaum, M., "Some Empirical Evidence on the Production Level and Production Cost Smoothing Models of Inventory Investment", NBER Working Paper No. 2523, February 1988.

Fair, R., "The Production–Smoothing Model is Alive and Well," NBER Working Paper No. 2877, March 1989.

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

Kahn, J. "Inventories and Production in a Competitive Equilibrium", mimeo, 1988.

Kahn, J., "Why is Production more Volatile than Sales? Theory and Evidence on the Stockout–Avoidance Motive for Inventory–Holding", Rochester Center for Economic Research Working Paper No. 168, revised 1989.

Krane, S., and S. Braun, "Production Smoothing Evidence from Physical–Product Data," Board of Governors of the Federal Reserve System, mimeo.

Lovell, M., "Manufacturer's Inventories, Sales Expectations, and the Acceleration Principle," *Econometrica 29* (1961), 293–314.

Miron, J. and S. Zeldes, "Seasonality, Cost Shocks and the Production Smoothing Model of Inventories," *Econometrica 56* (1988), 877–908.

Miron, J. and S. Zeldes, "Production, Sales, and the Change in Inventories: An Identity that Doesn't Add Up," *Journal of Monetary Economics*, 1989.

Ramey, V., "Non–Convex Costs and the Behavior of Inventories", University of California, San Diego, mimeo, 1988.

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

West, K., "Order Backlogs and Production Smoothing," NBER Working Paper No. 2385, 1987.

Table 1: Summary Statistics

Name	Var(Y)/Var(Z)		E(N/Z)[1]	Cov(N,Z)		S.E. of Z
	NSA	SA		NSA	SA	Forecast[2]
Autos	1.294	1.191	2.519	0.545	0.595	0.092
Tires	0.950	1.114	2.349	0.462	0.485	0.081
Cement	0.373	1.130	1.179	–0.485	–0.234	0.161
Lead	0.934	0.936	1.716	–0.230	–0.204	0.140
Zinc	0.909	0.916	1.258	0.038	–0.001	0.198
Copper	0.913	0.946	1.495	–0.479	–0.474	0.146

Notes:
1. N refers to the end–of–period inventory stock.
2. The forecast standard error is the s.e. from univariate ARMA equation in log(Z).

Table 2: Implied Variance Ratios

γ	Var(Y)/Var(Z) ρ=0	ρ=.5	ρ=.9
0.5	0.577	0.755	0.944
1	0.446	0.657	0.915
5	0.216	0.422	0.813
10	0.154	0.332	0.751
50	0.068	0.174	0.567
100	0.048	0.127	0.476

Table 3: Implied Versus Actual Production–Smoothing

Ind.	γ_{1b}	ρ	Var(Y)/Var(Z) implied	Var(Y)/Var(Z) actual
Lead	298.1	0.75	0.158	0.936
Zinc	79.2	0.90	0.507	0.916
Copper	172.8	0.81	0.255	0.946

Table 4: Skewness Statistics

Name	Skewness	
	NSA	SA
Autos	–0.423	–0.724**
Tires	1.078**	1.433**
Cement	–0.289*	–0.074
Lead	–0.651**	–0.647**
Zinc	–0.400**	–0.376**
Copper	–1.075**	–1.049**

Notes: Auto data (SA only) were also strike–adjusted.
* Significant at 10 percent.
** Significant at 1 percent.

Figure 1: SEASONALS

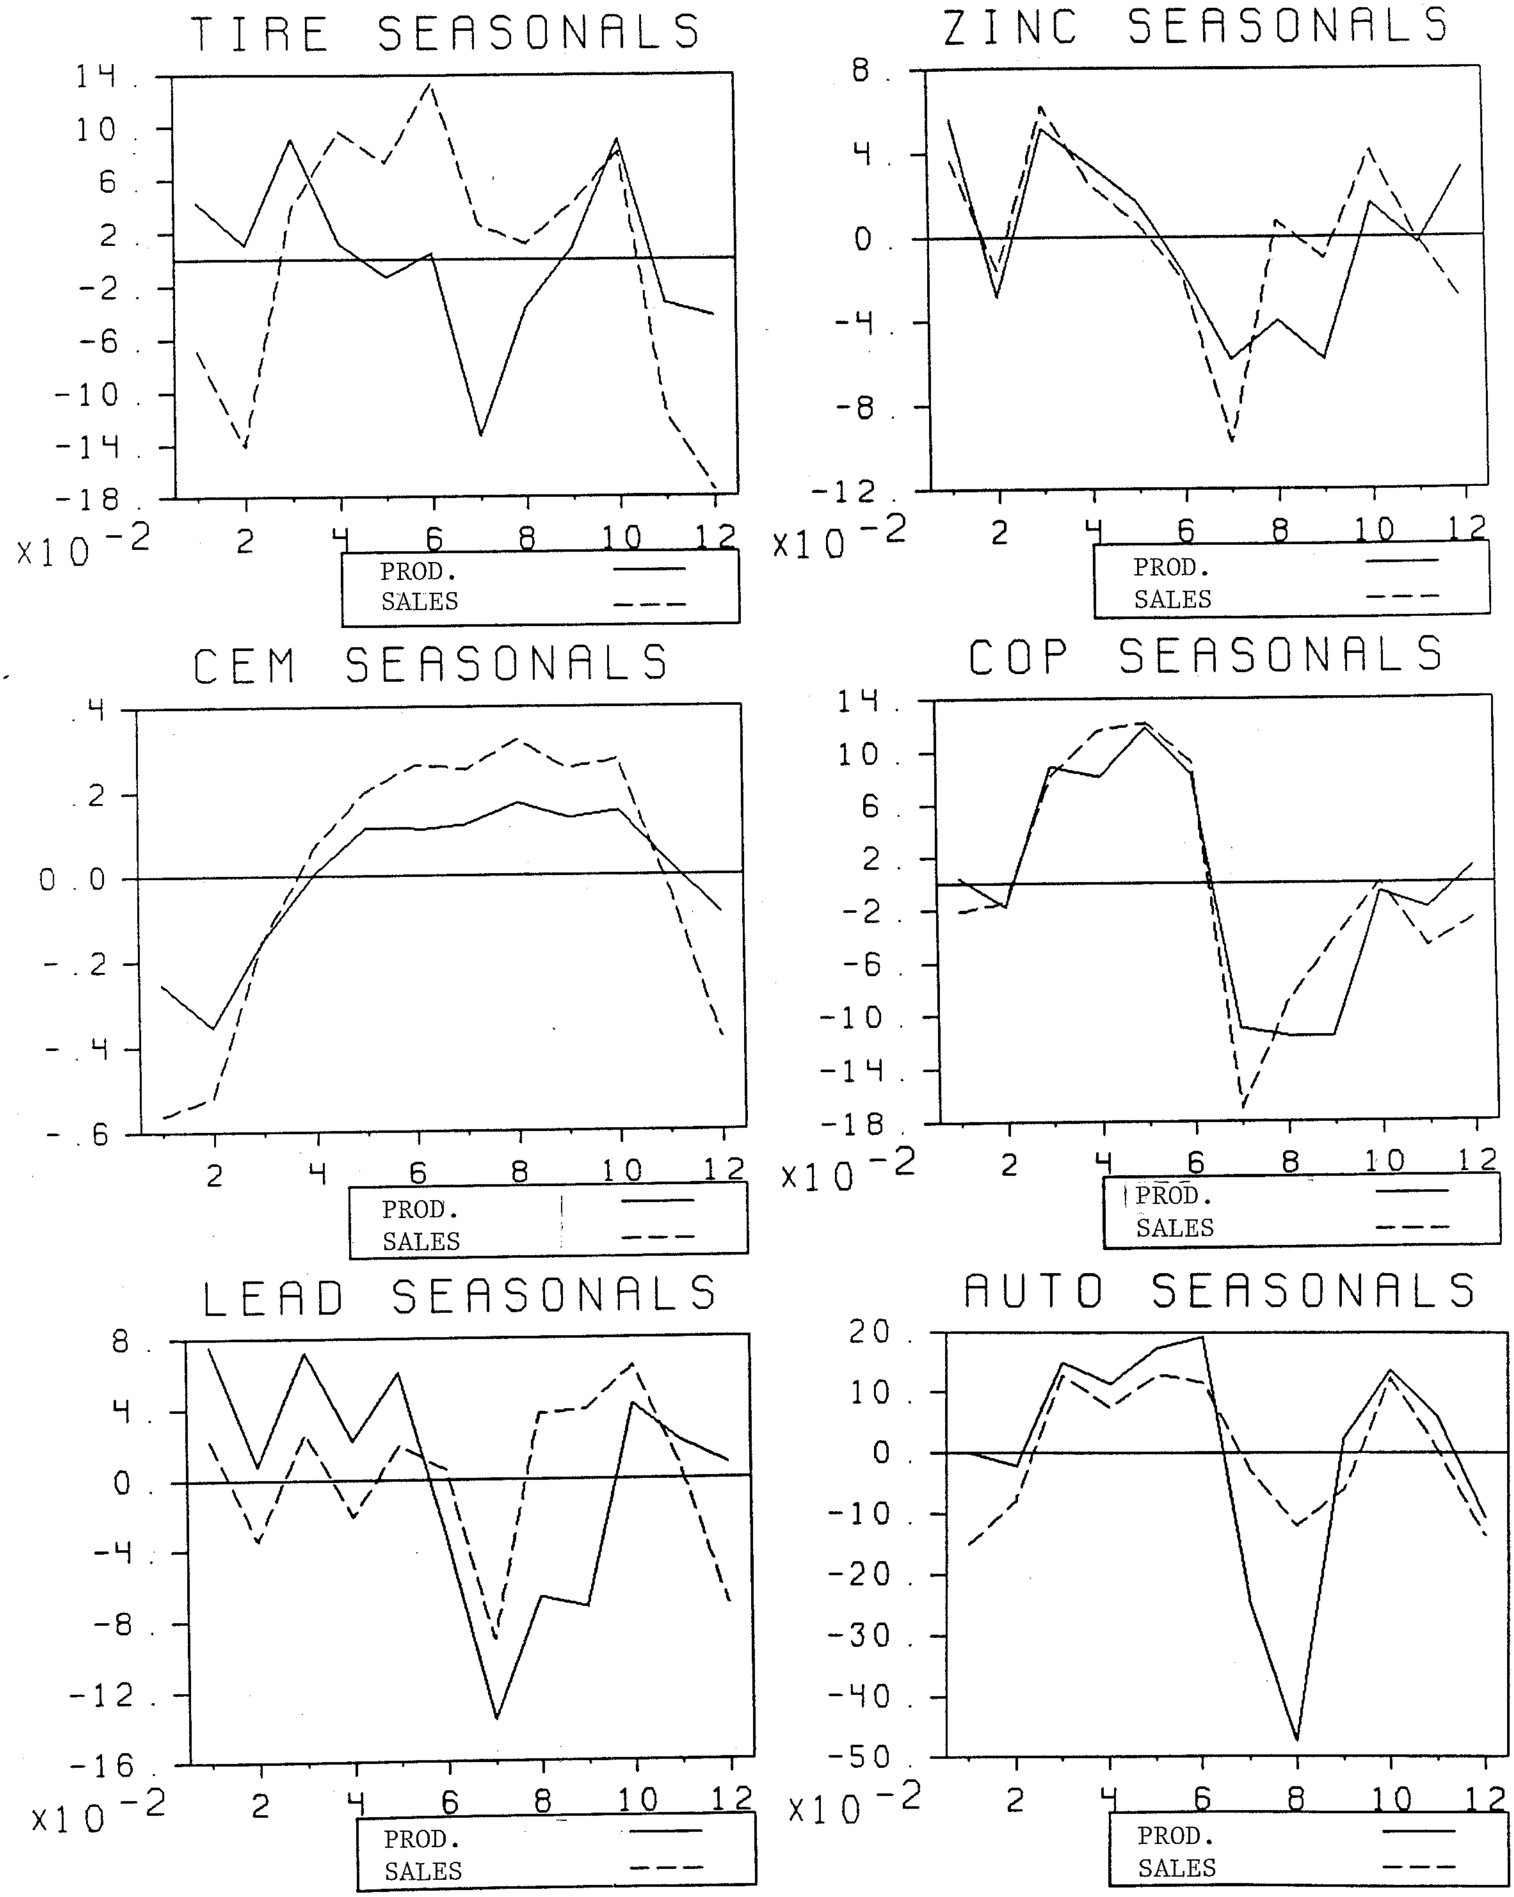

FIGURE 2B

Inventories and Sales, Slab Zinc

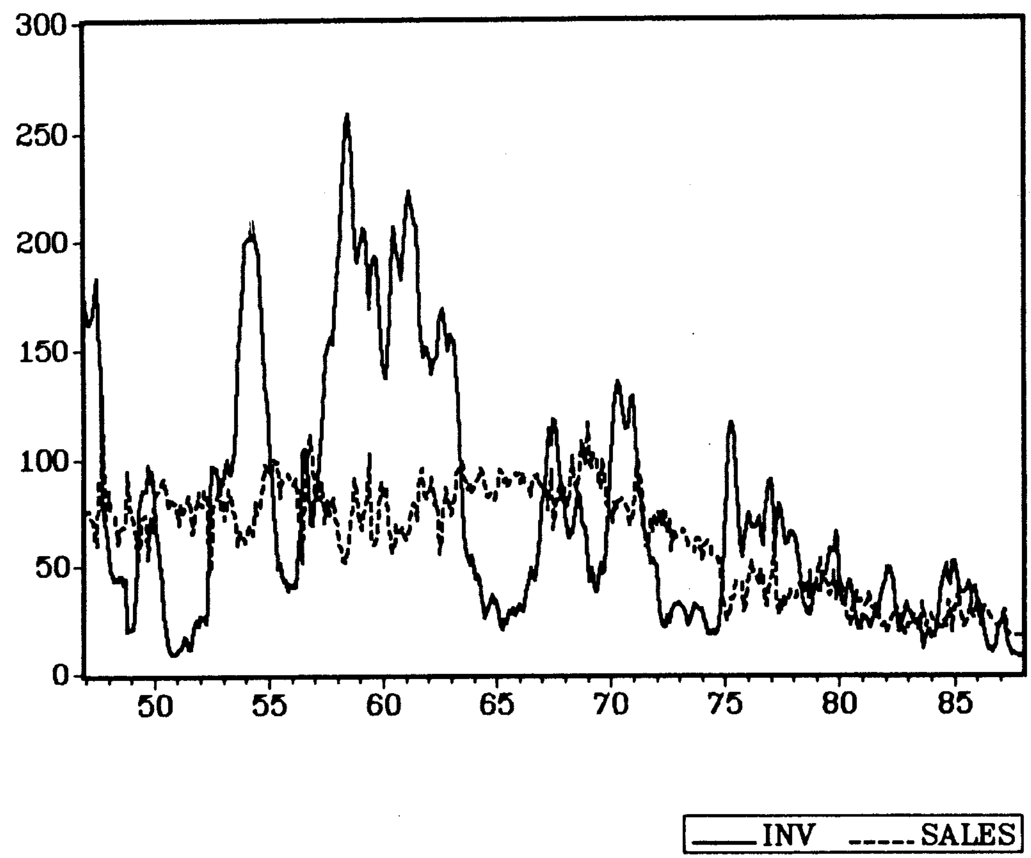

FIGURE 2A

Production and Sales, Slab Zinc

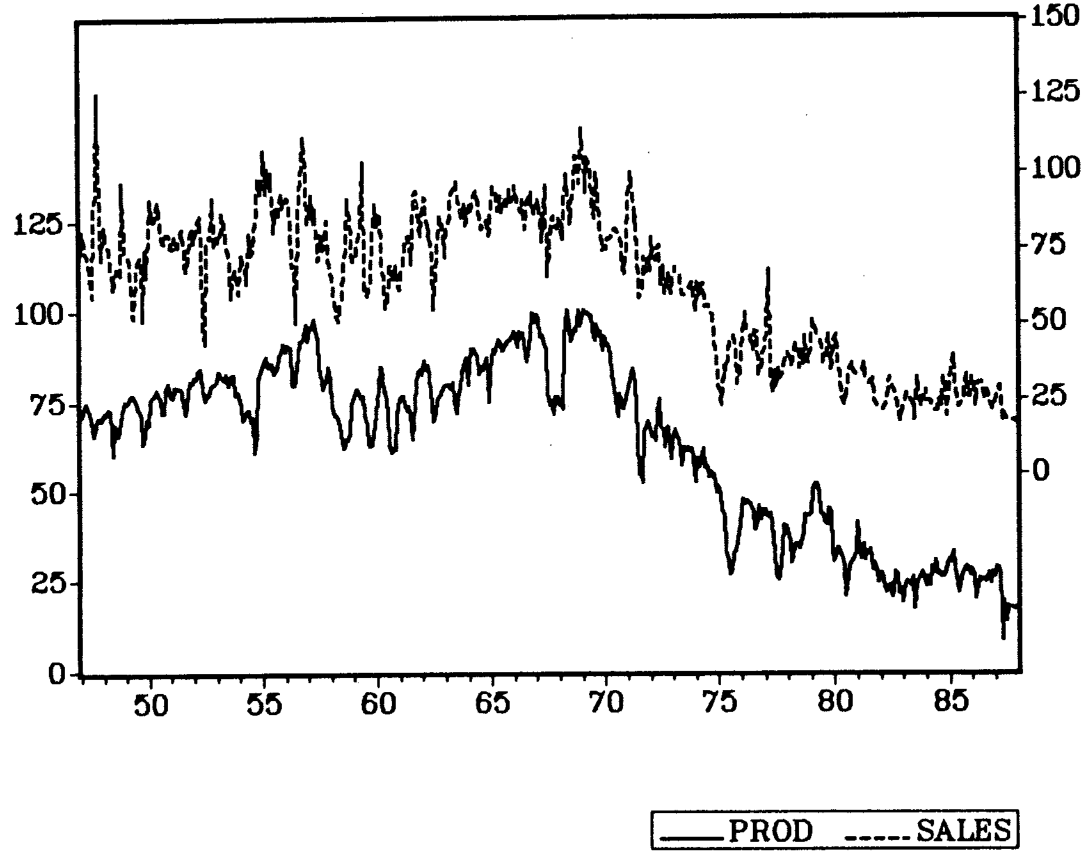

FIGURE 3A

Production and Sales, Cement

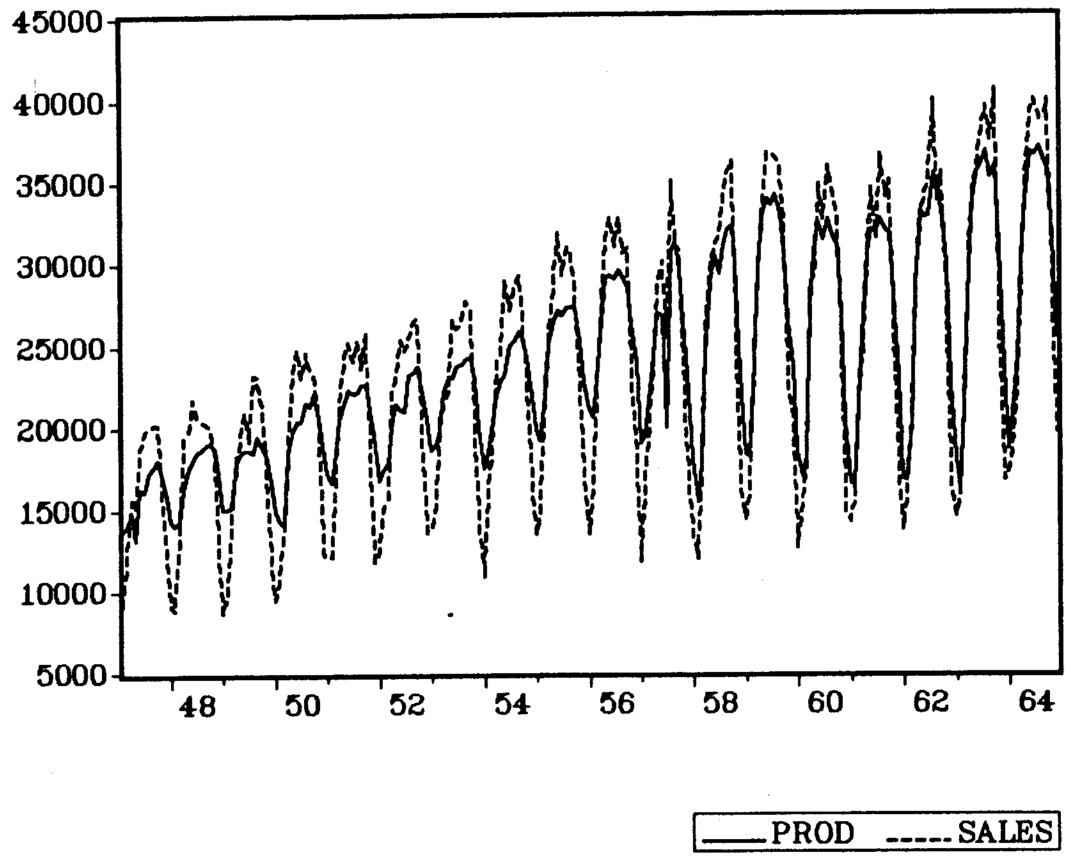

FIGURE 3B

Log Production, "Seasonally Adjusted", Cement Industry

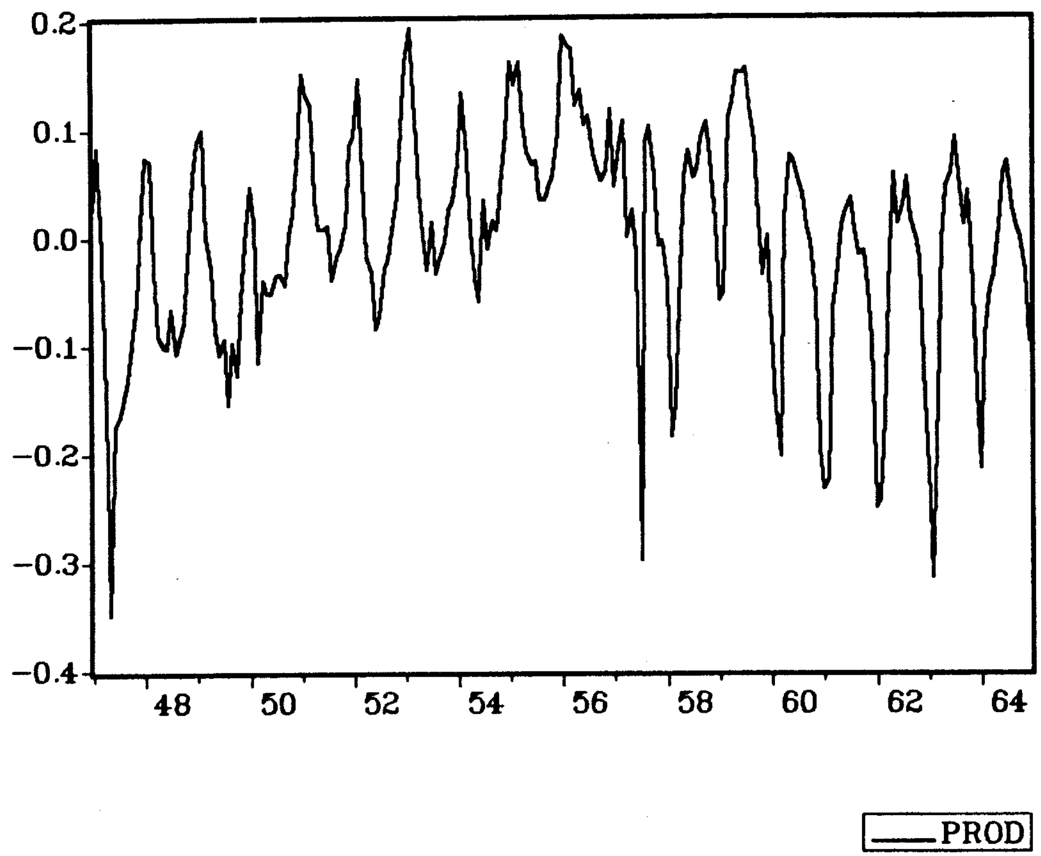

FIGURE 3C

Annual changes in log Production, NBER Reference Cycles
Cement Industry

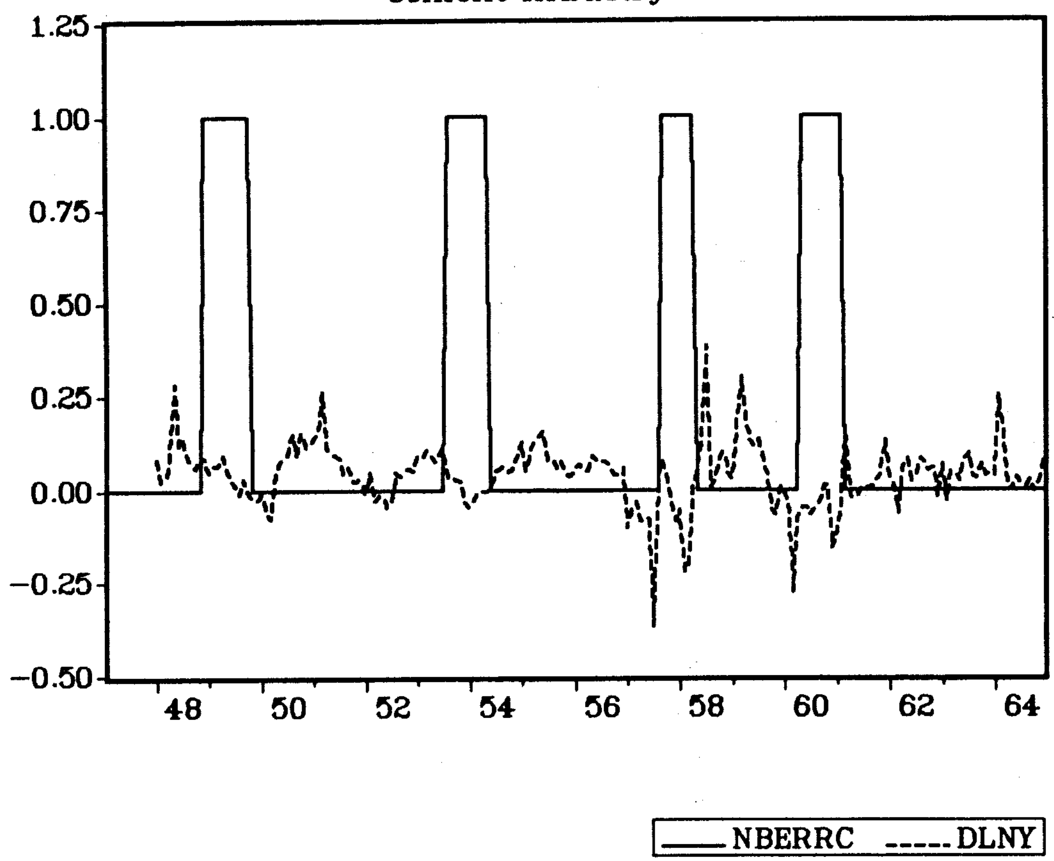

FIGURE 4: POWER SPECTRA