Econometrica, Vol. 56, No. 4 (July, 1988), 877–908

SEASONALITY, COST SHOCKS, AND THE PRODUCTION SMOOTHING MODEL OF INVENTORIES

BY JEFFREY A. MIRON AND STEPHEN P. ZELDES[1]

A great deal of research on the empirical behavior of inventories examines some variant of the production smoothing model of finished goods inventories. The overall assessment of this model that exists in the literature is quite negative: there is little evidence that manufacturers hold inventories of finished goods in order to smooth production patterns.

This paper examines whether this negative assessment of the model is due to one or both of two features: costs shocks and seasonal fluctuations. The reason for considering costs shocks is that, if firms are buffeted more by cost shocks than demand shocks, production should optimally be more variable than sales. The reasons for considering seasonal fluctuations are that seasonal fluctuations account for a major portion of the variance in production and sales, that seasonal fluctuations are precisely the kinds of fluctuations that producers should most easily smooth, and that seasonally adjusted data are likely to produce spurious rejections of the production smoothing model even when it is correct.

We integrate cost shocks and seasonal fluctuations into the analysis of the production smoothing model in three steps. First, we present a general production smoothing model of inventory investment that is consistent with both seasonal and nonseasonal fluctuations in production, sales, and inventories. The model allows for both observable and unobservable changes in marginal costs. Second, we estimate this model using both seasonally adjusted and seasonally unadjusted data plus seasonal dummies. The goal here is to determine whether the incorrect use of seasonally adjusted data has been responsible for the rejections of the production smoothing model reported in previous studies. The third part of our approach is to explicitly examine the seasonal movements in the data. We test whether the residual from an Euler equation is uncorrelated with the seasonal component of contemporaneous sales. Even if unobservable seasonal cost shocks make the seasonal variation in output greater than that in sales, the timing of the resulting seasonal movements in output should not necessarily match that of sales.

The results of our empirical work provide a strong negative report on the production smoothing model, even when it includes cost shocks and seasonal fluctuations. At both seasonal and nonseasonal frequencies, there appears to be little evidence that firms hold inventories in order to smooth production. A striking piece of evidence is that in most industries the seasonal in production closely matches the seasonal in shipments, even after accounting for the movements in interest rates, input prices, and the weather.

KEYWORDS: Inventories, seasonality, production, cost shocks

1. INTRODUCTION

A GREAT DEAL OF RESEARCH on the empirical behavior of inventories examines some variant of the production smoothing model of finished goods inventories. Blinder (1986a) emphasizes that, in the absence of cost shocks, the model implies that the variance of production should be less than the variance of sales, an inequality that is violated for manufacturing as a whole and most 2-digit

[1] The authors are grateful for helpful comments by Andrew Abel, Olivier Blanchard, Alan Blinder, John Campbell, Larry Christiano, Angus Deaton, Martin Eichenbaum, Stanley Fischer, Marvin Goodfriend, Tryphon Kollintzas, Louis Maccini, Trish Mosser, Chris Sims, Larry Summers, Ken West, members of the Macro Lunch Group at the University of Pennsylvania, two referees, and seminar participants at workshops at Michigan, Princeton, Johns Hopkins, Rochester, MacMaster, Northwestern, Econometric Society/NORC, and FRB Minneapolis and Philadelphia; and for excellent research assistance by Edward Gold.

industries. West (1986) derives a variance bounds test that extends this inequality in a number of ways and also finds that the data reject the model. Both Blinder and West conclude that there is strong evidence against the production smoothing model. Other authors, such as Blanchard (1983), Eichenbaum (1984), and Christiano and Eichenbaum (1986), present evidence that is less unfavorable to the model, but they reject it as well.

This paper examines the extent to which the negative assessment of the model is due to two features: cost shocks and seasonal fluctuations. Blinder (1986a) and West (1986) both note that the presence of cost shocks could explain the rejections that they report, and Blinder (1986b), Maccini and Rossana (1984), Eichenbaum (1984), and Christiano and Eichenbaum (1986) test the model in the presence of cost shocks, with partial success. The reason for considering these shocks is simply that if firms are buffeted more by cost shocks than demand shocks, production should optimally be more variable than sales.

Most of the empirical work on the production smoothing model uses data adjusted by the X-11 seasonal adjustment routine. This includes studies by Blinder (1986a, 1986b), Eichenbaum (1984), and Maccini and Rossana (1984). Blanchard (1983), Reagan and Sheehan (1985), and West (1986) begin with the seasonally unadjusted data and then adjust the data with seasonal dummies. Few studies examine whether the seasonal fluctuations themselves are consistent with the model of inventories. Exceptions are Ward (1978), who finds evidence that firms alter production rates differently in response to seasonal versus nonseasonal variations in demand; West (1986), who includes a version of his variance bounds test based on both the seasonal and nonseasonal variations in the data; and Ghali (1987), who uses data from the Portland Cement industry and finds that seasonal adjustment of the data is an important factor in the rejection of the production smoothing model.[2]

There are several reasons to think that using seasonally adjusted data to test inventory models is problematic. To begin with, seasonal fluctuations account for a major portion of the variation in production, shipments, and inventories. Table I shows the seasonal, nonseasonal, and total variance of the logarithmic rate of growth of production and shipments, for six 2-digit manufacturing industries.[3,4] For both variables, seasonal variation accounts for more than half of the total variance in most industries. Any analysis of production/inventory behavior that excludes seasonality at best explains only part of the story and fails to exploit much of the variation in the data.

Seasonal fluctuations are likely to be particularly useful in examining the production smoothing model because they are anticipated. Any test of the

[2] Irvine (1981) uses seasonally unadjusted data, with no seasonal dummies, to examine retail inventory behavior and the cost of capital.

[3] Table I includes results based on two different measures of production. See Section 4 for details.

[4] This table is similar to Table II in Blanchard (1983). As he points out, since the seasonal component is deterministic, it has no variance in the statistical sense. The numbers reported here for the seasonal variances are the average squared deviations of the twelve seasonal dummy coefficients from the sample mean of these coefficients.

TABLE I

SUMMARY STATISTICS, LOG GROWTH RATES OF PRODUCTION AND SHIPMENTS

		Food	Tobacco	Apparel	Chemicals	Petroleum	Rubber
Shipments	Mean:	0.16%	−0.09%	−0.06%	0.20%	0.16%	0.15%
	Variance:						
	Nonseasonal	5.60E-04	5.08E-03	1.74E-03	9.30E-04	5.10E-04	1.08E-03
	Seasonal	1.63E-03	1.61E-03	1.36E-02	2.86E-03	3.10E-04	4.33E-03
	Total	2.19E-03	6.69E-03	1.54E-02	3.79E-03	8.20E-04	5.41E-03
	Seasonal/Total	75%	24%	89%	75%	38%	80%
Production, $Y4$	Mean:	0.13%	−0.17%	−0.04%	0.19%	0.13%	0.15%
	Variance:						
	Nonseasonal	6.10E-04	6.03E-03	3.10E-03	8.00E-04	9.80E-04	1.60E-03
	Seasonal	1.63E-03	5.26E-03	1.16E-02	2.20E-03	2.20E-04	5.13E-03
	Total	2.24E-03	1.13E-02	1.47E-02	3.00E-03	1.20E-03	6.73E-03
	Seasonal/Total	73%	47%	79%	73%	18%	76%
Production, IP	Mean:	0.22%	−0.17%	0.01%	0.36%	0.11%	0.42%
	Variance:						
	Nonseasonal	1.50E-04	2.47E-03	1.80E-03	2.70E-04	4.50E-04	1.09E-03
	Seasonal	8.50E-04	1.70E-02	4.28E-03	4.30E-04	5.10E-04	3.06E-03
	Total	1.00E-03	1.95E-02	6.08E-03	6.90E-04	9.60E-04	4.14E-03
	Seasonal/Total	85%	87%	70%	62%	53%	74%

Notes: The sample period is 1967:5–1982:12. The log growth rates are defined as $\ln x_t - \ln x_{t-1}$. They have not been annualized.

production smoothing model involves a set of maintained hypotheses, one of which is the rationality hypothesis. Rejections of the model, therefore, are not usually informative as to which aspect of the joint hypothesis has been rejected. When a rational expectations model is applied to seasonal fluctuations, however, it seems reasonable to take the rationality hypothesis as correct, since if anything is correctly anticipated by agents seasonal fluctuations ought to be. Applying the production smoothing model to seasonal fluctuations may help determine which aspects of the model, if any, fail.

A final reason to avoid the use of seasonally adjusted data is that, since the true model must apply to the seasonally unadjusted data, the use of adjusted data is likely to lead to rejection of the model even when it is correct.[5] This is especially the case with data adjusted by the Census X-11 method because this technique makes the adjusted data a two-sided moving average of the underlying unadjusted data.[6,7] Therefore, the key implication of most rational expectations models, that the error term should be uncorrelated with lagged information, need not hold in the adjusted data even if it does hold in the unadjusted data.[8] If the data are adjusted by some other method, such as seasonal dummies, then the

[5] Summers (1981) emphasizes this point.

[6] X-11 is not literally a two-sided moving average filter. Rather, it can be well approximated by such filters. For more on this point, see Cleveland and Tiao (1976) and Wallis (1974).

[7] For example, Miron (1986) finds that the use of X-11 adjusted data is partially responsible for rejections of consumption Euler equations.

[8] See Sargent (1978).

time series properties of the adjusted data are not altered as radically as they are with X-11.

We integrate cost shocks and seasonal fluctuations into the analysis of the production smoothing model in three steps. First, we present a general production smoothing model of inventory investment that is consistent with both seasonal and nonseasonal fluctuations in production, sales, and inventories. The model allows for both observable and unobservable changes in marginal costs (cost shocks). The observables include wages, energy prices, raw materials prices, and interest rates, as well as weather variables (temperature and precipitation).[9] We examine a firm's cost minimization problem, so our analysis is robust to various assumptions about the competitiveness of the firm's output market. A key implication of the model is that, for any firm that can hold finished goods inventories at finite cost, the marginal cost of producing an additional unit of output today and holding it in inventories until next period must equal the expected marginal cost of producing that unit next period. With standard types of auxiliary assumptions about functional forms and identification, the model leads to an estimable Euler equation relating the rate of growth of production to the rate of growth of input prices, the level of inventories, and the interest rate. We estimate this Euler equation and test the overidentifying restrictions implied by the model, using data on six 2-digit manufacturing industries.

The second part of our approach is to perform the exact same estimations and tests of the above model using both seasonally adjusted and unadjusted data.[10] The goal here is to determine whether the incorrect use of seasonally adjusted data is responsible for the rejections of the production smoothing model reported in previous studies.

The third part of our approach is to explicitly examine the seasonal movements in the data. Since the predictable seasonal movement in demand is exactly the variation that should be most easily smoothed by firms, tests of the model at seasonal frequencies are particularly powerful. We therefore test whether the residual from the Euler equation is uncorrelated with the seasonal component of contemporaneous sales. Even if unobservable seasonal cost shocks make the seasonal variation in output greater than that of sales, the *timing* of the resulting seasonal movements in output should not match that of sales.

The estimation strategy that we employ involves a number of important identifying restrictions about the shifts over time in the firm's production function. We include a number of observable variables that account for the shifts in technology. There may, however, be additional shifts that are not accounted for by the measured variables, and these unobserved productivity shifts will appear in the error term of the equation that we estimate. In order to consistently

[9] Maccini and Rossana (1984) estimate a different style model of inventory accumulation (a general flexible accelerator model) using data on aggregate durables and nondurables inventory accumulation in which they include wages, energy costs, interest rates, and raw materials prices. They found that only raw materials prices had significant effects in their model.

[10] Constant dollar, seasonally unadjusted inventory data are not available and are therefore constructed. This is discussed further in Section 4.

estimate the Euler equation, therefore, we need to assume that this term is uncorrelated with the variables we use as instruments. Specifically, we assume that the unobserved productivity shifter is uncorrelated with lagged values of sales and with the part of current sales that is predictable based on lagged information, and that the growth of the unobserved productivity shifter is uncorrelated with lagged growth rates of input prices, lagged growth rates of output, and lagged interest rates. In addition, when we examine the seasonal fluctuations in production, we assume that any seasonal in unobserved productivity is uncorrelated with the seasonal in sales.

The remainder of the paper is organized as follows. Section 2 presents the basic production smoothing model that we employ throughout the paper and derives the first order condition that we estimate. In Section 3 we describe the identifying assumptions, the resulting testable implications, and the econometric techniques used to test those implications. In Section 4, we discuss the data used. Section 5 presents the basic results with seasonally adjusted and unadjusted data. In Section 6, we examine the seasonal-specific results. Section 7 concludes the paper.

2. THE MODEL

Consider a profit maximizing firm. Sales by the firm, the price of the firm's output, and the firm's capital stock may be exogenously or endogenously determined. The firm may be a monopolist, a perfect competitor, or something in between. The firm is, however, assumed to be a competitor in the markets for inputs. For any pattern of prices, sales, and the capital stock, the firm chooses its inputs over time so as to minimize costs.

The firm's intertemporal cost minimization problem is

$$(1) \qquad \min_{\{y_{t+j}\}} E_t \sum_{j=0}^{T} \Gamma_{t,t+j} C_{t+j}(y_{t+j})$$

subject to

$$n_{t+j} = n_{t+j-1}(1 - s'_{t+j-1}) + y_{t+j} - x_{t+j},$$
$$n_{t+j} \geq 0 \quad \forall_j,$$

where y_t is production in period t, x_t is sales in period t, and n_t is the stock of inventories at the end of period t, all measured in terms of the output good. The end of the firm's horizon is period T. C_t is the one period nominal cost function of the firm, to be derived shortly. $\Gamma_{t,t+j}$ is the nominal discount factor, defined as the present value at time t of one dollar at $t+j$. Thus,

$$\Gamma_{t,t+j} \equiv \left[\prod_{s=0}^{j-1}\left(\frac{1}{1+\tilde{R}_{t+s}}\right)\right], \quad \Gamma_{t,t} \equiv 1, \quad \text{and}$$
$$\tilde{R}_{t+s} \equiv (1 - m_{t+s+1}) R_{t+s}.$$

R_t is the pretax cost of capital for the firm, and m_t is the marginal tax rate. E_t indicates expectations conditional on information available at time t.

The term s_t' is the fraction of inventories lost due to storage costs. In the case of linear storage costs, s_t' is equal to a constant (call this s_1). Some researchers have modeled storage costs as being convex in the level of inventories. For example, convex inventory costs are the key factor driving Blinder and Fischer's (1981) model of the real business cycle. We capture these types of costs here by writing $s_t' = s_1 + (s_2/2) \cdot n_t$.[11]

For any cost minimizing firm that carries inventories between two periods, the marginal cost of producing an extra unit of output this period and holding it in inventories until next period must equal the expected marginal cost of producing an extra unit of output next period. This first-order condition can be written as

$$(2) \qquad MC_t = E_t\left[\frac{MC_{t+1}(1-s_t)}{1+\tilde{R}_t}\right]$$

or

$$(3) \qquad E_t\left[\frac{MC_{t+1}}{MC_t} \cdot \frac{(1-s_t)}{1+\tilde{R}_t}\right] = 1.$$

Rational expectations implies

$$(4) \qquad \frac{MC_{t+1}}{MC_t} \cdot \frac{(1-s_t)}{1+\tilde{R}_t} = 1 + \varepsilon_{t+1}$$

where $E_t[\varepsilon_{t+1}] = 0$, i.e. ε_{t+1} is orthogonal to all information available at time t. The marginal storage cost, s_t, is equal to $s_1 + s_2 \cdot n_t$.[12] The Euler equation (4) will not be satisfied if desired inventories are zero. We discuss this possibility below.

At this point it is worth pointing out the parallel between the production/storage problem of a cost minimizing firm and the consumption/saving problem of a utility maximizing consumer. The firm's problem is to minimize the expected discounted value of a convex cost function, subject to an expected pattern of sales and costs of holding inventories. The consumer's problem is to maximize the expected discounted value of a concave utility function, subject to an expected pattern of income and return to holding wealth. Not surprisingly, then, the solution to cost minimization yields a first-order condition analogous to the first-order condition implied by the stochastic version of the permanent income hypothesis (Hall (1978), Mankiw (1981), Hansen and Singleton (1983)), and we can apply the methods of that literature to testing the production smoothing model of inventories and output. Production, sales, inventories, the interest rate,

[11] If storage costs come in the form of depreciating inventories, then the accounting identity definition of output would be $y_t = x_t + (n_t - n_{t-1}(1 - s_{t-1}'))$. In this paper, we construct output in the standard way: $y_t = x_t + (n_t - n_{t-1})$. If, rather than coming in the form of depreciated stocks, storage costs are actually paid out and these costs are proportional to the replacement cost of the goods, then our model and our constructed output measure are consistent with one another. In either of these cases the equations are correct when the IP measure of output is used. If the costs are paid out in dollars, in an amount related to the goods stored, our equation is approximately correct.

[12] If average storage costs (s_t') are equal to $s_1 + (s_2/2)n_t$, this implies that marginal storage costs (s_t) are equal to $s_1 + s_2 n_t$.

and storage costs are analogous to consumption, income, wealth, the rate of time preference, and the return on wealth, respectively.[13] In the simplest version of this model, the real interest rate, the growth in the capital stock, and productivity growth are all constant over time. Given the production function that we employ, these assumptions imply that the expected growth in output is constant over time —i.e., real output follows a geometric random walk with drift. This is analogous to Hall's (1978) condition that consumption follow a random walk with drift.

To implement the model described above we need to specify the form of the cost function. We assume a standard Cobb-Douglas production function with m inputs (q_i, $i=1,\dots,m$). Let the last input (q_m) be the capital stock. In each period, the firm thus solves the following (constrained) problem:

$$(5)\qquad \operatorname*{Min}_{\{q_1,q_2,\dots,q_{m-1}\}} \sum_{i=1}^{m} w_i\cdot q_i \quad \text{subject to}$$

$$f(q_1,\dots,q_m)=\mu\prod_{i=1}^{m} q_i^{a_i}=\bar{y},$$

$$q_m=\bar{q}_m,$$

where w_i and q_i are the price and quantity, respectively, of input i, and f is the production function. Note that the production function includes a productivity measure μ that may shift over time in deterministic and/or stochastic ways.[14] Define $A\equiv\sum_{i=1}^{m-1}a_i$. The one period (constrained) cost function from this problem is:

$$(6)\qquad C(y)=w_m\cdot q_m+A\cdot q_m^{(A-1)/A}\left[\prod_{i=1}^{m-1}\left(\frac{w_i}{a_i}\right)^{a_i/A}\right]\cdot\mu^{-1/A}\cdot y^{1/A}$$

and the marginal cost function MC is:

$$(7)\qquad MC(y)=q_m^{A-1/A}\left[\prod_{i=1}^{m-1}\left(\frac{w_i}{a_i}\right)^{a_i/A}\right]\cdot\mu^{-(1/A)}\cdot y^{(1-A)/A}.$$

Equation (7) can be used to calculate the ratio of marginal costs in t and $t+1$:

$$(8)\qquad \ln\left(\frac{MC_{t+1}}{MC_t}\right)=\left[\sum_{i=1}^{m-1}(a_i/A)\cdot\ln\left(\frac{w_{it+1}}{w_{it}}\right)\right]-\left(\frac{1-A}{A}\right)\ln\left(\frac{q_{mt+1}}{q_{mt}}\right)$$
$$+\left(\frac{1-A}{A}\right)\ln\left(\frac{y_{t+1}}{y_t}\right)-\frac{1}{A}\ln\frac{\mu_{t+1}}{\mu_t}.$$

[13] The nonnegativity condition on inventories mentioned above is analogous to a borrowing constraint in the consumption literature. If time series/cross section data on firms were available, an approach similar to that of Zeldes (1985) could be applied here to test for the importance of this nonnegativity constraint on inventories.

[14] Unlike some previous studies, we do not include costs of adjusting the level of output. As Maccini and Rossana (1984) point out, the costs of adjusting output presumably arise because of the costs of changing one or more factors of production. These costs may be important, but we do not attempt to model them here.

The next step is to derive an expression for the growth rate of output. We do so by taking logs of the Euler equation (4), taking a first order Taylor expansion of $\ln(1-s_t)$ around $n_t = k$ for an arbitrary value of $k \geq 0$ and a second order Taylor expansion of $\ln(1+\varepsilon_{t+1})$ around $\varepsilon_{t+1} = 0$, substituting in equation (8), and rearranging. This gives:

$$(9) \qquad Gy_{t+1} = \big[(A/(1-A)) \times \big(-\ln(1-s_1-s_2k) - ks_2/(1-s_1-s_2k) - \tfrac{1}{2}\sigma_\varepsilon^2\big)\big] + \left(\frac{A}{1-A}\right)\left[\ln\big(1+\tilde{R}_t\big) - \sum_{i=1}^{m-1}(a_i/A)Gw_{it+1}\right] + Gq_{mt+1} + \frac{s_2}{1-s_1-s_2k}\left(\frac{A}{1-A}\right)n_t + \left(\frac{1}{1-A}\right)G\mu_{t+1} + \left(\frac{A}{1-A}\right)\left[\big(\tfrac{1}{2}\sigma_\varepsilon^2 - \tfrac{1}{2}\varepsilon_{t+1}^2\big) + \varepsilon_{t+1}\right]$$

where for any variable Z, $GZ_{t+1} \equiv \ln(Z_{t+1}/Z_t)$. We have added and subtracted $(A/(1-A))\frac{1}{2}\sigma_\varepsilon^2$ from the equation, so that the last term in brackets in equation (9) has mean zero.

Discussion of the Model

Equation (9) is the basis of all the estimations performed in this paper. It says that the growth rate of output is a function of the real interest rate (where the inflation rate used to calculate the real rate is a weighted average of the rates of inflation of factor prices), the growth in the capital stock, the level of inventories, productivity growth, and a surprise term. The key implication that we test in this paper is that no other information known at time t should help predict output growth.

As is well known, an advantage of estimating this Euler equation is that we avoid solving for firms' closed form decision rule for production. This allows us to step outside the linear-quadratic framework, and it allows us to estimate our model that includes stochastic input prices and interest rates. In addition, the Euler equation procedure yields testable implications for the growth rate rather than the level of output, so we do not need to assume that output is stationary around a deterministic trend. Our procedure is valid even if there is a unit root in the level of output, a condition that Nelson and Plosser (1982) find, for example, characterizes aggregate output series.[15]

In setting up the model we have imposed the constraint that inventories are nonnegative, and we have indicated that the Euler equation is valid in a given period only if the nonnegativity constraint is not binding in that period. We

[15] See Ghysels (1987) for an analysis of trends versus unit roots in manufacturing inventory and production data.

should point out here a potential problem related to this nonnegativity constraint. Consider a certainty version of our model without the nonnegativity constraint imposed. Assuming that s_1 and s_2 are nonnegative, equation (3) implies that when inventories are positive firms want the level of marginal costs to rise over time. If the marginal cost *function* is constant or falling (due to growth in the capital stock), this implies that output rises over time, i.e., that firms push production towards the future and run down inventory stocks today. In fact, only if inventories are negative could there be a steady state with constant marginal costs. This indicates that in a model in which the nonnegativity condition *is* imposed, it will at times bind, and therefore the Euler equation will not be satisfied in some periods.[16]

To partially avoid this problem, we follow Blinder (1982) and allow s_1 to be negative. This captures the fact that at low but positive levels of inventories, increases in inventories may lower total costs, i.e. there may be a convenience yield to holding inventories. With s_1 sufficiently negative, there is a steady state in the certainty version of the model that has a positive level of inventories.[17] Of course, this still does not imply, in the certainty or uncertainty version of the model, that inventories *never* hit the constraint.[18,19]

The model that we use allows for seasonal fluctuations in output growth in several ways. First, there may be seasonal movements in the observable or unobservable component of the productivity shifter. Second, there may be seasonals in the relevant input prices. Of course, it is not entirely accurate to describe these as determining the seasonal fluctuations in output growth, since in general equilibrium the seasonals in output growth and input prices are determined simultaneously. For an individual firm, however, and even for a 2-digit industry, the degree of simultaneity is likely to be small.

Rather than assuming that the productivity shifter μ is totally unobservable to the econometrician, we allow it to be a function of some observable seasonal variables and some unobservables. The observable variables are weather related: functions of current temperature and precipitation. It seems reasonable a priori that productivity would be affected by the current local weather. We write $\mu = e^{Z\gamma+\eta}$, where Z is a matrix of observable weather variables and η is the unobservable productivity shifter.

In the absence of shifts in the cost function (i.e., changes in μ), the model presented is a simple production smoothing model. For a given time path for the

[16] Even under these assumptions, firms will in general choose to use inventories in some periods to smooth production, i.e., to build up positive inventories in the periods in which sales are especially high and run them down in periods in which sales are low.

[17] This can be seen by using (3), assuming no uncertainty, letting $R_t = R$, $n_t = n$, and setting $MC_{t+1} = MC_t$. Rearranging gives $n = (-R - s_1)/s_2$, which will be greater than zero if $s_1 < -R$.

[18] For a further discussion of this issue see Schutte (1983). Another factor in our model that tends to push inventories positive is sales growth, although this will be reversed to the extent that it is accompanied by growth in the capital stock.

[19] In the industries that we use to estimate the equation, industrywide inventories are always positive. This does not of course imply that inventories are always positive for every firm in these industries.

capital stock, the derived cost function is convex, inducing firms to try to spread production evenly over time.[20] When productivity is allowed to vary over time, the result is no longer a pure production smoothing model. Although our model is consistent with the variance of production exceeding the variance of sales, the convexity of the cost function remains and we continue to refer to the model as a type of production smoothing model.

Blinder (1986a) states that introducing (unobservable) cost shocks into the analysis makes his variance bounds inequality untestable, because one could explain an arbitrarily large variance of production relative to sales by assuming unobservable cost shocks with appropriately large variance. The approach that we adopt in this paper avoids this problem in two ways. First, we include *measurements* of a number of factors that might influence the marginal cost of production, and account for these in the analysis. Second, we show that under reasonable identifying assumptions, the model described above has testable implications even in the presence of unobservable cost shocks. The most important assumption is that the unobserved component of productivity is uncorrelated with the component of sales that is predictable on the basis of information known at the time the firm makes its production decision. The testable implication is that once a number of cost variables are accounted for, the remaining movements in output should be uncorrelated with predictable movements in sales. In other words, even if production moves around a lot due to cost shocks, these movements should not be related to predictable movements in sales. This will be an especially useful test when applied to predictable *seasonal* movements in sales.

3. IDENTIFICATION AND TESTING

A. The General Approach

Equation (9), augmented to include the weather variables, can be written as:

$$
\begin{aligned}
(10)\qquad Gy_{t+1} = {} & (A/(1-A)) \\
& \times\left(-\ln(1-s_1-s_2k)-\left(ks_2/(1-s_1-s_2k)\right)-\tfrac{1}{2}\sigma_\varepsilon^2\right) \\
& +\left(\frac{A}{1-A}\right)\left[\ln\left(1+\tilde{R}_t\right)-\sum_{i=1}^{m-1}(a_i/A)Gw_{it+1}\right] \\
& +Gq_{mt+1}+\frac{s_2}{1-s_1-s_2k}\left(\frac{A}{1-A}\right)n_t+\left(\frac{1}{1-A}\right)(Z_{t+1}-Z_t)\gamma \\
& +\left(\frac{1}{1-A}\right)G\eta_{t+1}+\left(\frac{A}{1-A}\right)\left[\left(\tfrac{1}{2}\sigma_\varepsilon^2-\tfrac{1}{2}\varepsilon_{t+1}^2\right)+\varepsilon_{t+1}\right].
\end{aligned}
$$

[20] This smoothing that arises from a convex cost function is different than the smoothing induced by introducing costs of adjusting output (as in, for example, Eichenbaum (1984)). For further discussion, see Blanchard (1983).

We cannot estimate this equation by OLS because the right-hand side variables are in general correlated with the expectations error. We therefore use an instrumental variables procedure to estimate the equation. To do so, we must choose instruments that are correlated with the included variables but not with the error term. Recall that the error term includes two components: the expectations error and the growth in the unobserved productivity shifter. Any variable that is known at time t will, by rational expectations, be orthogonal to ε_{t+1}.[21] However, rationality of expectations does not imply that $G\eta_{t+1}$ is orthogonal to time t information—it is possible that there are predictable movements in productivity growth. Note that a reasonable possibility is that the productivity measure follows a geometric random walk, in which case the growth in productivity is i.i.d. and therefore orthogonal to lagged information.[22]

Garber and King (1984) point out that a number of studies that estimate Euler equations assume that there are no shocks in the sector that they are estimating —effectively ignoring the identification issue. In this paper, we allow some measurable shocks to this sector, and we make the following identifying assumptions about the relationship between the unobserved cost shifter and the included instruments. (i) The unobserved productivity shifter (η) is uncorrelated with lagged values of sales and with the part of current sales that was predictable based on lagged information. (ii) The growth of the productivity shifter is uncorrelated with lagged growth rates of input prices, lagged growth rates of output, and lagged interest rates.

We thus consider the following variables to be orthogonal to the error term in the regression: lagged growth in sales, lagged growth in output, lagged interest rate, lagged growth in factor prices, and lagged inventories. In some sets of results we relax the assumption that the lagged growth rate of output is uncorrelated with the growth in the productivity shifter. To test the model, we first estimate equation (10) with instrumental variables, including as instruments the variables in the above list. Since there are more instruments than right-hand side variables, the equation is overidentified. We then test the overidentifying restrictions by regressing the estimated residuals on all of the included instruments (including the predetermined right-hand side variables). The quantity T times the R^2 from this regression is distributed χ^2_j, where T is the number of observations and j is the number of overidentifying restrictions. One possible alternative hypothesis to our null is that firms simply set current output in line with current sales. In this case, we would expect the lagged growth rate of sales to enter significantly in our test of the overidentifying restrictions.

[21] We assume that production decisions for the month are made after information about demand and other economic variables is revealed, i.e., period t output decisions are made contingent on period t economic variables. An alternative assumption would be that production decisions are made before demand for the month is known. This creates a stockout motive for holding inventories (see Kahn, 1986). In Section 5, we also present results based on the alternative assumption that output must be chosen before demand for the period is known. For a further discussion of these timing issues, see Blinder (1986a).

[22] This is the assumption made by Prescott (1986).

B. Seasonality and Identification

It is possible that there are seasonal movements in productivity that are not captured by the weather variables. One possible way to capture these would be to allow the productivity measure to be an arbitrary function of seasonal dummies. We do this in our first set of results by including seasonal dummies in the estimation of equation (10).[23] This gives the same results as first regressing all of the variables on seasonal dummies, and then using the residuals from these regressions for estimation purposes.

In order to examine whether the use of X-11 adjusted data has been responsible for the rejections by others of the production smoothing model, we compare the tests of the model using seasonally unadjusted data and seasonal dummies to the tests using X-11 seasonally adjusted data.

When we include seasonal dummies in equation (10), we lose all power to test the model at seasonal frequencies, i.e., we cannot test whether the seasonal movements in the data are consistent with the model. In the latter part of the paper, therefore, we make the stronger identifying assumption that seasonal shifts in productivity not captured by weather variables are uncorrelated with the instruments used to estimate equation (10).[24] Under this assumption, we can exclude seasonal dummies and perform two further tests that directly use seasonal fluctuations in the data. We test the implication that once the other factors in the cost function are taken into account, the remaining movements in output should be uncorrelated with the seasonal movements in sales. This is a strong implication of the production smoothing model that has not been tested to date. In addition, we examine whether the model fits at purely seasonal frequencies. We describe these latter two tests in Section 6.

4. THE DATA

This section describes the data set that we employ. There are a number of technical issues to be considered with respect to both the adjusted and unadjusted data on inventories and production; we discuss these in detail below. Readers who are not interested in these details can skip to Section 5.

The equations are estimated using monthly data from May 1967 through December 1982.[25,26] Data on inventories and shipments at the 2-digit SIC level

[23] The fact that we include seasonal dummies does not mean that we assume purely deterministic seasonality. Since the right-hand side variables may exhibit stochastic seasonality, our model allows for both stochastic and deterministic seasonality in output growth. We should also note that because we are working in log first differences, using additive seasonal dummies allows for multiplicative seasonality in output.

[24] In the section below on seasonal results, we discuss the circumstances under which this assumption might not hold.

[25] Most of our data run through December 1984, but we only have weather data through December 1982.

[26] A month seems like a reasonable planning horizon for a firm, but there is no obvious reason why it need be so. For a discussion of time aggregation issues in inventory models, see Christiano and Eichenbaum (1986).

for 20 industries were obtained from the Department of Commerce. We estimate the equations only on the six industries identified by Belsley (1969) as being production to stock industries. The inventory data are end of month inventories of finished goods, adjusted by the Bureau of Economic Analysis (BEA) from the book value reported by firms into constant dollars.[27,28] We follow West (1983) and adjust the BEA series from "cost" to "market," so that shipments and inventories are in comparable units. Shipments data are total monthly figures in constant dollars.

Two different measures of production are used. The first comes from the identity that production of finished goods equals sales plus the change in inventories of finished goods. Commerce Department data for sales and the change in inventories are used to compute this production measure (which we call "Y4"). The second measure of production used is the Federal Reserve Board's index of industrial production (IP), also available at the 2-digit SIC level.

In principle, the two production series measure the same variable and should therefore behave similarly over time. As documented in Miron and Zeldes (1987), however, the two series are in fact quite different. For the six industries studied here, the correlations between growth rates of the two series range from .8 to .4 for the seasonally unadjusted data, and from .4 to less than .1 for the seasonally adjusted data. The serial correlation properties and seasonal movements of the two series are also different. Since we have not resolved this discrepancy, we present results based on both output measures.

The nominal interest rate is the yield to maturity on Treasury Bills with one month to maturity as reported on the CRSP tapes. The marginal corporate tax rate series is the one calculated by Feldstein and Summers (1979). The input price series are wages, the price of crude materials for further processing, and energy prices, representing the three largest variable inputs in the production process. Wages (average hourly earnings) and industrial production at the 2-digit SIC level, and aggregate measures of energy prices (the PPI for petroleum and coal products) and raw materials prices are available from the Citibank Economic Database.

The capital stock enters our equations as the number of machine days used per month. Since we did not have access to industry capital stock data, we model the growth in the capital stock as a constant plus a function of the growth in the number of nonholiday weekdays in the month. Any remaining month to month variation in the growth in the capital stock is included in the error term.

The weather data include estimates of total monthly precipitation and average monthly temperature. We construct a different temperature and precipitation

[27] This adjustment attempts to take into account whether firms used LIFO or FIFO accounting. See Hinrichs and Eckman (1981) for a description of how the constant dollar inventory series are constructed. See Reagan and Sheehan (1985) for a presentation of the stylized facts of these series at an aggregate (durables and nondurables) level.

[28] There is some disagreement over whether it is appropriate to use finished goods inventories only (West (1986)) or finished goods plus work in progress inventories (Blinder (1986a)). We estimate the equations separately for each definition. See footnote 32 in Section 5.

measure for each industry, equal to weighted averages of the corresponding measures in the different states. The weights are equal to the historical share of the total shipments of the industry that originated in each state.[29] To capture nonlinearities, we also include the weighted average of squared temperature, squared precipitation, and the cross-product of temperature and precipitation. Given our functional form assumptions, the first differences of these variables enter equation (10).

Seasonality

Whenever possible, we obtained both seasonally adjusted (*SA*) and seasonally unadjusted (*NSA*) data. The BEA reports real shipments and inventories data, but these constant dollar series are only available on a *SA* basis.[30] The Bureau of the Census reports *NSA* and *SA* current dollar shipments series and book value inventories series. As in Reagan and Sheehan (1985) and West (1986), we estimate the real *NSA* inventory series by multiplying the real *SA* series by the ratio of book value *NSA* to book value *SA*, thus putting back in an estimate of the seasonal. (Another way of thinking of this is that we deflate the book value *NSA* series by the ratio of the book value *SA* to real *SA* series.) We estimate real *NSA* shipments by multiplying the real *SA* series by the ratio of nominal *NSA* shipments to nominal *SA* shipments. These procedures assume that there are no seasonal movements in the factors that convert from book value to current dollar value or in the deflators used to convert the series from current dollar to constant dollar. An additional adjustment that we considered was to multiply the above series by the ratio of the *SA* to *NSA PPI* series for the finished goods, in order to adjust for the seasonal in the deflators. We found statistically significant evidence of seasonality in the price indexes in three out of six industries. However, the magnitudes of the seasonal movements in these prices are much smaller than in the corresponding quantities. We estimated the specifications in Tables II and V both with and without this adjustment and the results were virtually identical to each other.[31] We report only the results without this last adjustment.

The *IP* data are available both *NSA* and *SA*, and the energy price series, wage rates, raw materials prices, and interest rates are all unadjusted.

5. BASIC RESULTS

In this section, we examine the basic results from estimating equation (10) and testing the implied overidentifying restrictions. In order to determine whether the use of X-11 adjusted data has been responsible for previous rejections of the

[29] The weights change every five years but always correspond to averages of previous (never future) years.

[30] The reason for this has to do with the technique used to construct the constant dollar figures. The disaggregated nominal series are first seasonally adjusted, then deflated and then aggregated.

[31] We estimated the equation over a shorter sample period for the food, chemicals, and petroleum industries because seasonally adjusted *PPI*'s were unavailable for part of the sample period.

production smoothing model, we run the same set of tests with (i) the standard X-11 seasonally adjusted data and (ii) seasonally unadjusted data plus seasonal dummies.

A summary of the results is presented in Table II.[32] There are four sets of results, since we carry out the estimation with both unadjusted and adjusted data, and we do this for both the *Y4* and *IP* measures of output. In the first line of each set of results, we list the variables that entered equation (10) at a significance level of 5%. In the second line of each set, we present the R^2 from the regression of the residuals on all the instruments. Recall that $T \cdot R^2$ is distributed χ_j^2, where j is the number of overidentifying restrictions. On the same line, we report the marginal significance level of the test statistic $T \cdot R^2$. In the last line of each set, we list the variables that entered this auxiliary test significantly.

We make the following observations about the results. First, in no case does the interest rate or the growth rate in energy prices enter equation (10) significantly. In about one third of the cases, the growth in raw materials prices enters significantly, but usually with the wrong sign. Wage growth enters significantly only four times, twice with the wrong sign. Thus, the signs and statistical significance of the coefficient estimates are not supportive of the model.

The second observation we make is that the data reject the overidentifying restrictions on the model in all cases using the *Y4* data, and in two-thirds using the *IP* data. For the *Y4* data, the rejections are about as strong using seasonally adjusted as seasonally unadjusted data. For the *IP* data, the rejections are not quite as strong overall using the seasonally unadjusted data. On the whole, there is little evidence that the use of unadjusted data with seasonal dummies provides better results than using seasonally adjusted data.

Finally, note that in approximately half of the cases, at least one of the five weather variables enters the equation significantly. Even after including seasonal dummies, the weather has a significant influence on production in certain industries (tobacco, chemicals, and petroleum).

Thus far, we arrive at a negative assessment of the model for two reasons. First, the overidentifying restrictions are typically rejected. Second, the signs of the coefficient estimates are not sensible and rarely significant. Proponents of the model might make the following argument against these two reasons, respectively. First, the instrument list may include variables that are correlated with the error term even under the null hypothesis, thus invalidating the tests of the overidentifying restrictions. Second, the instruments may not do a very good job of explaining the right-hand side variables. If this is the case, one should not expect the parameter estimates to be statistically significant, even under the null. We discuss each of these arguments in turn.

There are two circumstances in which the instrument list employed, consisting of lagged values of production, sales, input prices, and inventories, may be

[32] Most of these estimations were also done using the sum of finished goods inventories and work in progress inventories as the definition of inventories. The results were almost identical to those reported in the text.

TABLE II

REGRESSION RESULTS, EQUATION (10) SEASONAL DUMMIES IN EQUATION AND INSTRUMENT LIST

		Food	Tobacco	Apparel	Chemicals	Petroleum	Rubber
Y4, NSA	What enters (10) significantly?	$sd, -rm$	$tem, -tem2$	sd	$-sd, tem$ $-tem2, -pre$	$-w, -day$ $-tem2, -pre2$	sd
	R^2, Significance level	.17, .000	.23, .000	.09, .050	.15, .001	.13, .004	.16, .000
	What is significant in test of OIR's?	$we_{-1}, -we_{-2}$	tem_{-1}	—	$-y_{-1}, x_{-1}$, $-n_{-2}, n_{-3}, sd$	we_{-2}	w, sd
Y4, SA	What enters (10) significantly?	—	rm, $tem, -tem2$	$-n_{-1}, rm$	rm	$-w, -day$ $-pre2, tpr$	w
	R^2, Significance Level	.10, .027	.14, .002	.11, .014	.14, .002	.14, .002	.11, .014
	What is significant in test of OIR's?	—	—	—	—	—	—
IP, NSA	What enters (10) significantly?	sd	$-sd, w, -pre2$	sd, day_{-1}	sd, rm	$sd, -pre2$	rm
	R^2, Significance Level	.16, .000	.16, .000	.09, .050	.04, .583	.08, .090	.02, .926
	What is significant in test of OIR's?	$-y_{-1}, -w_{-2}$	$-y_{-1}, sd$	$-y_{-1}$	—	—	—
IP, SA	What enters (10) significantly?	—	$-pre2$	$-pre2, day_{-1}$	$tem, -tem2$	$-day_{-1}$	rm, $tem, -tem2$
	R^2, Significance Level	.15, .001	.28, .000	.09, .050	.13, .004	.09, .050	.06, .257
	What is significant in test of OIR's?	$-y_{-1}, -we_{-2}$	$-y_{-1}, -y_{-2}$ $-tem_{-1}, tpr_{-1}$	$-y_{-1}$	x_{-1}, x_{-2}	—	$-n_{-1}$

Notes:

1. The sample period is 1967:5–1982:12.
2. The first line of each set of results lists the variables that entered equation (10) at the 5% significance level. We list seasonal dummies if one or more of the eleven dummies entered significantly.
3. The second line gives the R^2 from the regression of the residuals on the instruments, as well as the marginal significance level of this statistic. The quantity $T \times R^2$ is distributed χ_j^2, where j is the number of overidentifying restrictions and T is the number of observations. In the results presented here, there are 9 such restrictions.
4. The third line lists the variables that entered the regression of the residuals on the instruments at the 5% significance level.
5. w = wage growth, sd = seasonal dummies, y = output growth, x = sales growth, day = number of production days, we = energy price growth, rm = raw materials price growth, n = inventories, r = interest rate, pre = change in precipitation, $pre2$ = change in precipitation squared, tem = change in temperature, $tem2$ = change in temperature squared, tpr = change in temperature * precipitation.
6. A (−) before variable indicates that the sign of the coefficient was negative.
7. A subscript of $-l$ on a variable means that it is dated l periods earlier than the dependent variable.

correlated with the error term. First, lagged output growth may not be a valid instrument, even if other lagged variables are, because productivity growth might be serially correlated. Since productivity growth is correlated with output growth, this implies that lagged output growth will also be correlated with contemporaneous productivity growth (a component of the error term), making it an invalid instrument.

Second, if firms do not have complete current period information when they make their output decisions for period t, then variables dated time t may not be valid instruments. This could arise because firms do not know the demand for their own products for the period before choosing output (as in Kahn (1986) or Christiano (1986)). Alternatively, firms may know the total demand for their product, but not the aggregate component of demand. Since we are using data on firms aggregated to the industry level, this too might invalidate the use of time t instruments (see Goodfriend (1986)).

In order to take account of these possibilities, we have estimated equation (10) using two alternative instrument lists. The first excludes production from the instrument list and includes extra lags of sales. The second list excludes all variables dated time t and includes extra lags of the variables at earlier dates.

When we employ the first alternative instrument list we reject the overidentifying restrictions significantly less often than with the list used in our basic results. In this case, the restrictions are rejected in a majority of cases for the *Y4* data, but never for the *IP* data. When we employ the second alternative instrument list, we never reject the overidentifying restrictions. In both cases, however, we almost never find that the input price variables, the interest rate, or the level of inventories enter statistically significantly with the correct sign.

This brings us to the second issue. It is possible that we are not finding that expected changes in input prices affect the timing of production because there are no expected changes in input prices. That is, the instruments that we employ, either in our basic results or alternative results, may be of such poor quality that they have no explanatory power for the right-hand side variables in equation (10). If this is the case, the failure of these input prices to explain the pattern of production is not evidence against our model.

It is easy to check this possibility by examining directly the explanatory power of the instruments. For all three instrument lists, we find the following: there is statistically significant explanatory power in the instruments about half the time for wages; all the time for interest rates, energy prices, and all five weather variables; and almost never for raw materials prices. Thus, with the exception of raw materials prices, the failure of input prices to explain production in any of our results is valid evidence against the model.

To summarize, with our basic instrument list the results provide evidence against the production smoothing model, even when it is expanded to incorporate a stochastic interest rate, measurable and unmeasurable cost shocks, and non-quadratic technology. When two weaker sets of identifying restrictions are used, there is substantially less statistical evidence against the model, but there is still no evidence that it describes an important aspect of firm behavior. Using seasonally unadjusted data and seasonal dummies does little better than using X-11 adjusted data.

6. SEASONAL-SPECIFIC RESULTS

In this section we examine the extent to which the seasonal fluctuations in production, shipments and inventories are consistent with the production smoothing model. The results presented above incorporate seasonal fluctuations into the analysis by using seasonally unadjusted data and including seasonal dummies and weather variables in the equations. This approach does not tell us to what extent the seasonal movements in interest rates or input prices determine the seasonal movements in output growth, nor does it answer the question of whether the seasonal movements in the data themselves satisfy the production smoothing model. In order to answer these questions, we cannot include seasonal dummies in equation (10) and must therefore assume that any fluctuations (seasonal and nonseasonal) in the productivity measure not captured by the weather variables are orthogonal to the instruments used.

Before describing our formal tests, it is useful to consider a set of stylized facts about the seasonality in production, inventories, and sales. We saw in Table I that the seasonal variation in the data is large relative to the nonseasonal variation. In Table III, we present estimates of the ratio of the variance of production to the variance of sales, and we include estimates based separately on the seasonal and nonseasonal variation. Following Blinder (1986a), these numbers are based on detrended levels rather than growth rates.[33,34] As we have discussed above, if cost shocks are assumed to be "small," the production smoothing model restricts these ratios to be less than one. We focus here on the ratio of the seasonal variances. For three of the six industries, we estimate this ratio to be greater than one.[35]

While one could interpret a ratio greater than one as a rejection of the production smoothing model, there is no reason to expect the above ratio to be less than one if there are seasonal shifts in the cost function. Even in this case, however, there is information to be learned from examining the seasonal movements. Whether or not seasonal shifts in productivity affect the seasonal pattern

[33]Along the lines of Blinder, we use the following procedure to obtain detrended levels of the data. The log of each series is regressed on a constant, time, and a dummy variable that is one beginning in October 1973. The coefficients are estimated by GLS, assuming a second order autoregressive process for the error term. The antilogs of the fitted values of this regression are then subtracted from the levels of the raw data to define the detrended data. We convert the *IP* measure from an index into a constant dollar figure by multiplying it by the ratio of average *Y4* to average *IP* (i.e., we set the average of the two series equal to each other). We apply the detrending procedure to the resulting *IP*, as well as *Y4* and shipments. We then regress the detrended series on a constant and eleven seasonal dummies. The seasonal and nonseasonal variances are estimated using the fitted and residual values of this regression, respectively.

[34]In the last section of his paper, West (1986) describes a variance bounds test that includes deterministic seasonal variations in the data. He found that the variance bounds were rejected for each of the three industries that he examined.

[35]We examine these ratios for seasonally adjusted data in Miron and Zeldes (1987) and find significant differences between the ratios based on *IP* and *Y4* data.

TABLE III

VARIANCE OF PRODUCTION DIVIDED BY VARIANCE OF SALES

		Food	Tobacco	Apparel	Chemicals	Petroleum	Rubber
Y4	Nonseasonal	1.22	1.84	1.32	1.01	0.91	1.13
	Seasonal	1.71	4.71	0.58	0.72	2.73	0.99
	Total	1.50	2.53	0.80	0.89	0.94	1.09
IP	Nonseasonal	0.48	0.58	1.20	0.83	0.48	0.95
	Seasonal	1.64	6.28	0.21	0.18	7.91	0.61
	Total	1.14	1.95	0.50	0.55	0.59	0.86

Notes:
1. The sample period is 1967:5–1982:12.
2. The estimation procedure is the following. For both shipments and production, the log level is regressed on a constant, time and a time trend that is one beginning in October, 1973. The coefficients are estimated by GLS, assuming a second order autoregressive process for the error term. The antilogs of the fitted values of this regression are then subtracted from the actual data, in levels, to define the detrended data. The seasonal and nonseasonal variances are calculated as the variance of the fitted values and residuals, respectively, of a regression of the detrended series on seasonal dummies.
3. We convert the *IP* measure from an index to a constant dollar figure by multiplying it by the ratio of average *Y4* to average *IP*.

of production, there is no reason to expect that seasonal pattern to match the seasonal pattern of sales. Figures 1–6 show the seasonal patterns in output and shipments for the six industries we examine and document behavior potentially problematic for the production smoothing model.[36] *The seasonal movements in output and sales are in fact very similar*. The implication of these graphs is that inventories do not appear to be playing the role of smoothing seasonal fluctuations in sales.

In the tests we present in this section, we formalize this observation. First, we test whether the contemporaneous seasonal movement in sales growth helps predict residual output growth, once the movements in factor prices, the weather, and lagged inventories are taken into account. To do this, we use the same procedure as in Section 5, except that seasonal dummies are excluded from the regression and the instrument list, and the seasonal component of contemporaneous sales growth is added to the instrument list. It is unusual when running this type of orthogonality test to include as an instrument a contemporaneous variable, but since this series is deterministic, it is part of the lagged information set. Since it is also assumed orthogonal to the unobservable productivity shifter, it is a valid instrument.[37]

The interpretation of this procedure is the following. By excluding seasonal dummies from the equation, we force the seasonal and nonseasonal movements in the right-hand side variables to affect output growth via the same coefficients. Given this restriction, we are then testing whether the part of output growth not explained by these variables is correlated with the seasonal component of sales growth. This allows us to compare the seasonals in sales and output, after taking

[36] The seasonal coefficient plotted for each month is the average percentage difference in that month from a logarithmic time trend.

[37] The series we actually use is, of course, the estimated rather than the true seasonal in sales growth.

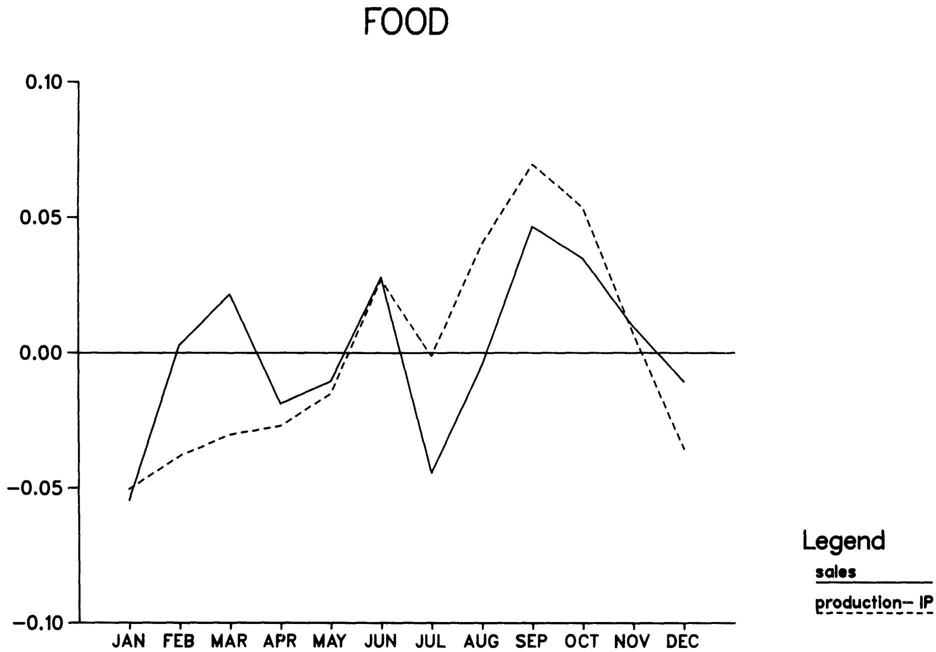
FOOD
0.10
0.05
0.00
−0.05
−0.10
JAN FEB MAR APR MAY JUN JUL AUG SEP OCT NOV DEC
Legend
sales
production– IP

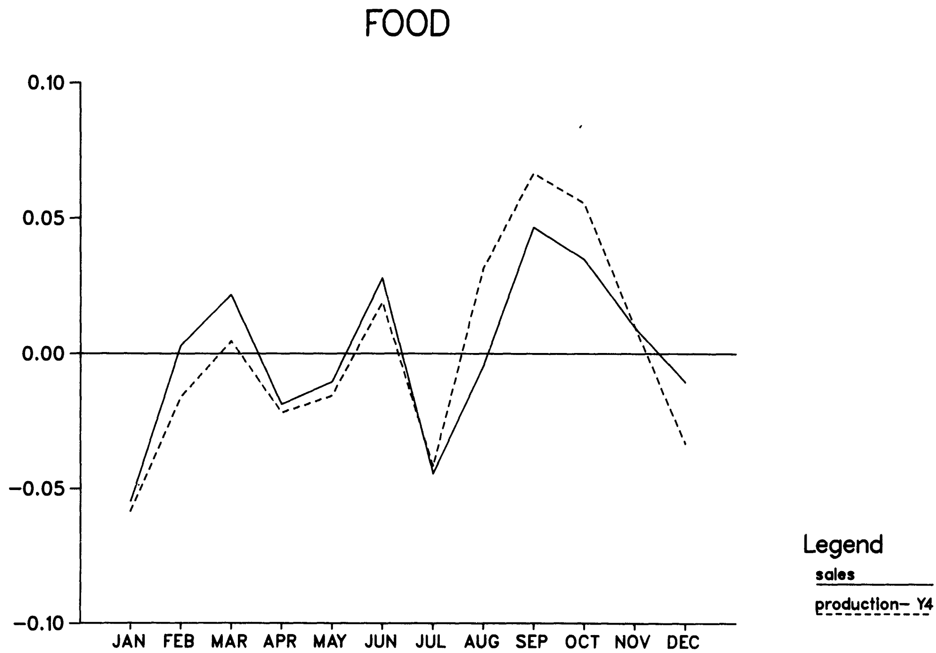
FOOD
0.10
0.05
0.00
−0.05
−0.10
JAN FEB MAR APR MAY JUN JUL AUG SEP OCT NOV DEC
Legend
sales
production– Y4

FIGURE 1

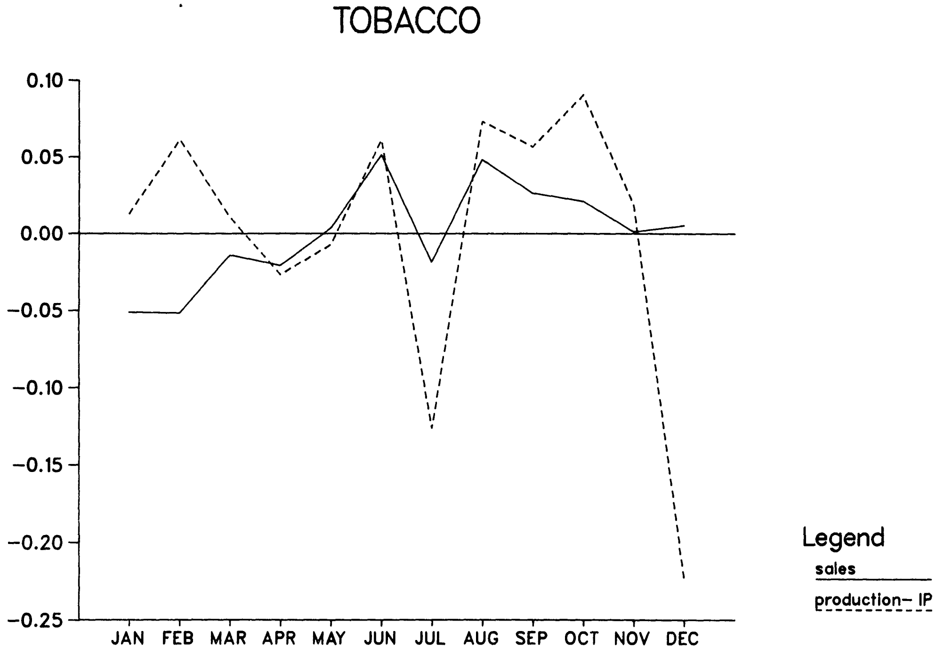
TOBACCO
0.10
0.05
0.00
−0.05
−0.10
−0.15
−0.20
−0.25
JAN FEB MAR APR MAY JUN JUL AUG SEP OCT NOV DEC
Legend
sales
production– IP

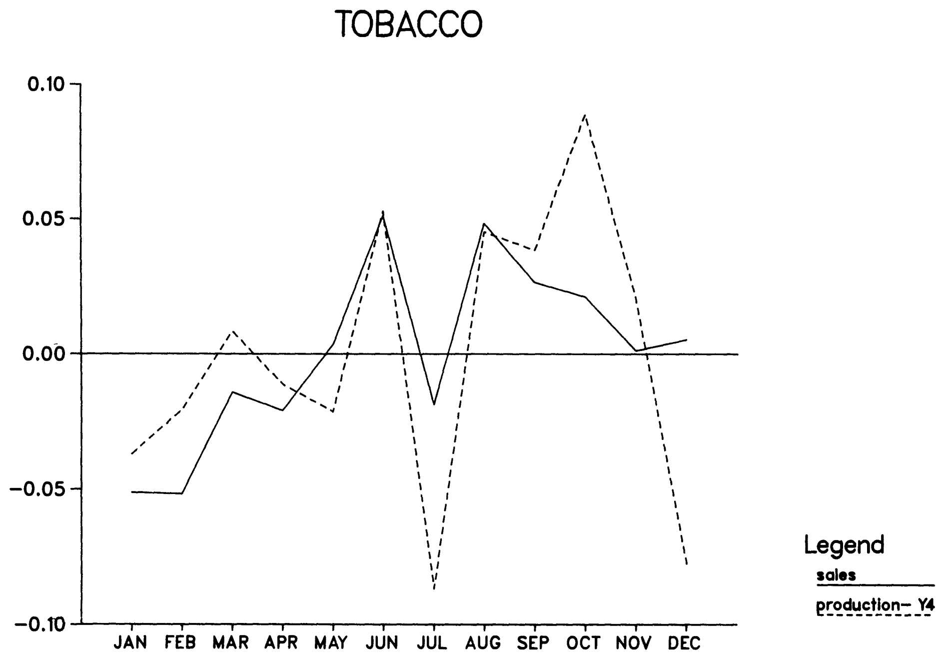
TOBACCO
0.10
0.05
0.00
−0.05
−0.10
JAN FEB MAR APR MAY JUN JUL AUG SEP OCT NOV DEC
Legend
sales
production– Y4

FIGURE 2

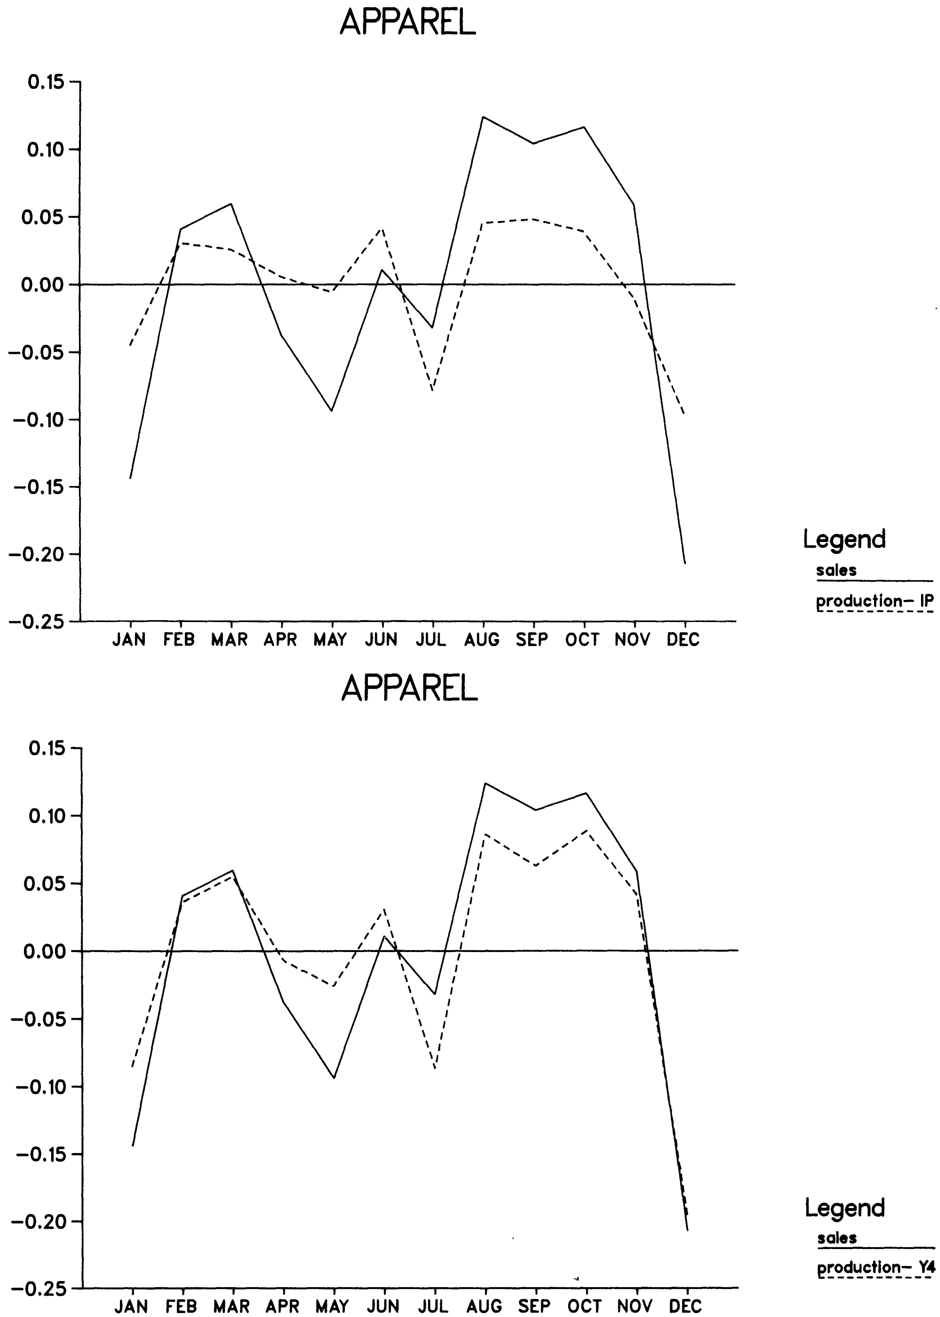
APPAREL
0.15
0.10
0.05
0.00
−0.05
−0.10
−0.15
−0.20
−0.25
JAN FEB MAR APR MAY JUN JUL AUG SEP OCT NOV DEC
Legend
sales
production– IP
APPAREL
0.15
0.10
0.05
0.00
−0.05
−0.10
−0.15
−0.20
−0.25
JAN FEB MAR APR MAY JUN JUL AUG SEP OCT NOV DEC
Legend
sales
production– Y4

FIGURE 3

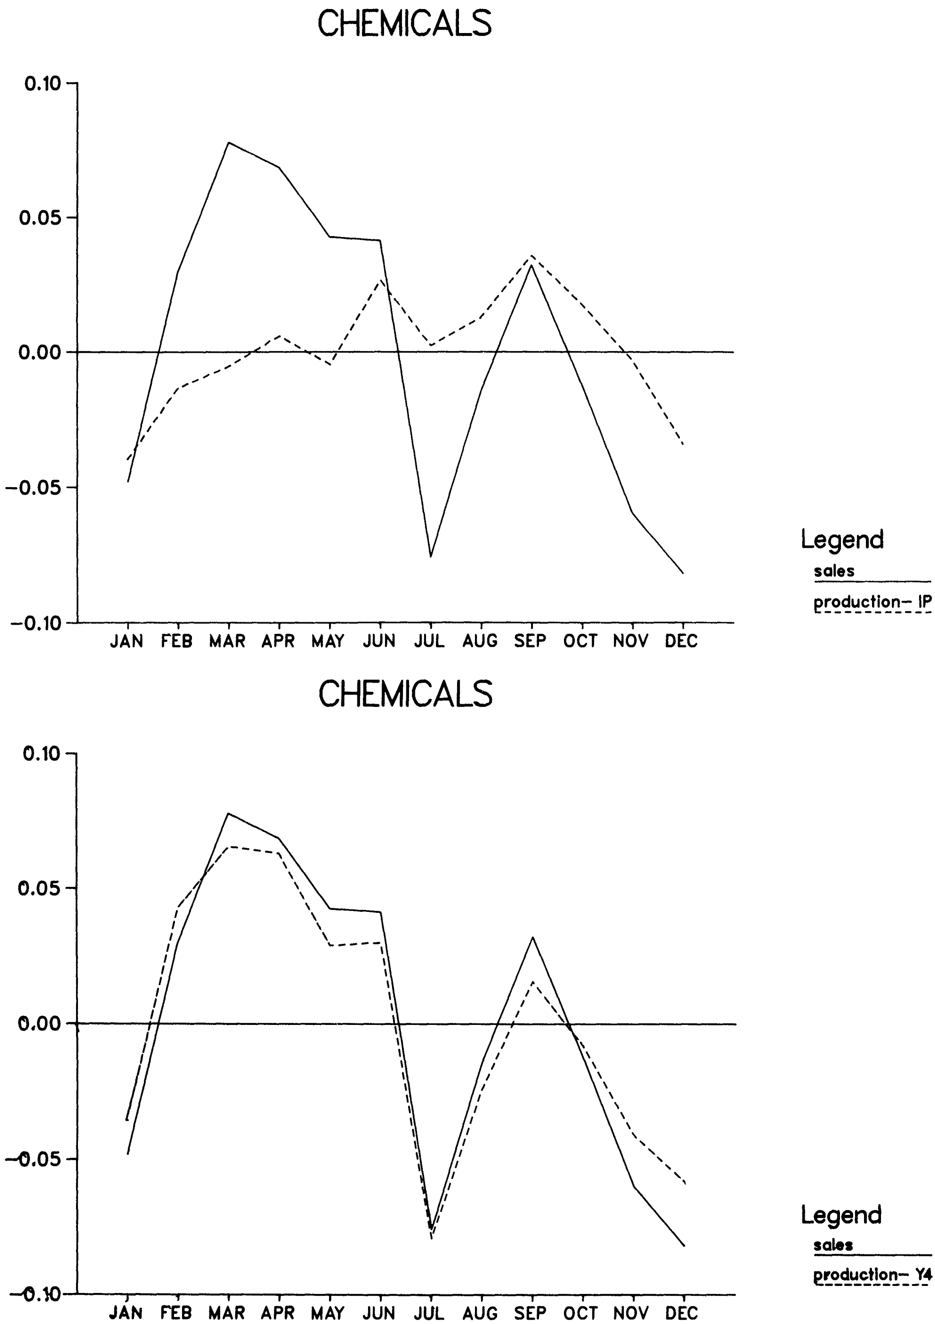
CHEMICALS
0.10
0.05
0.00
−0.05
−0.10
JAN FEB MAR APR MAY JUN JUL AUG SEP OCT NOV DEC
Legend
sales
production− IP
CHEMICALS
0.10
0.05
0.00
−0.05
−0.10
JAN FEB MAR APR MAY JUN JUL AUG SEP OCT NOV DEC
Legend
sales
production− Y4

FIGURE 4

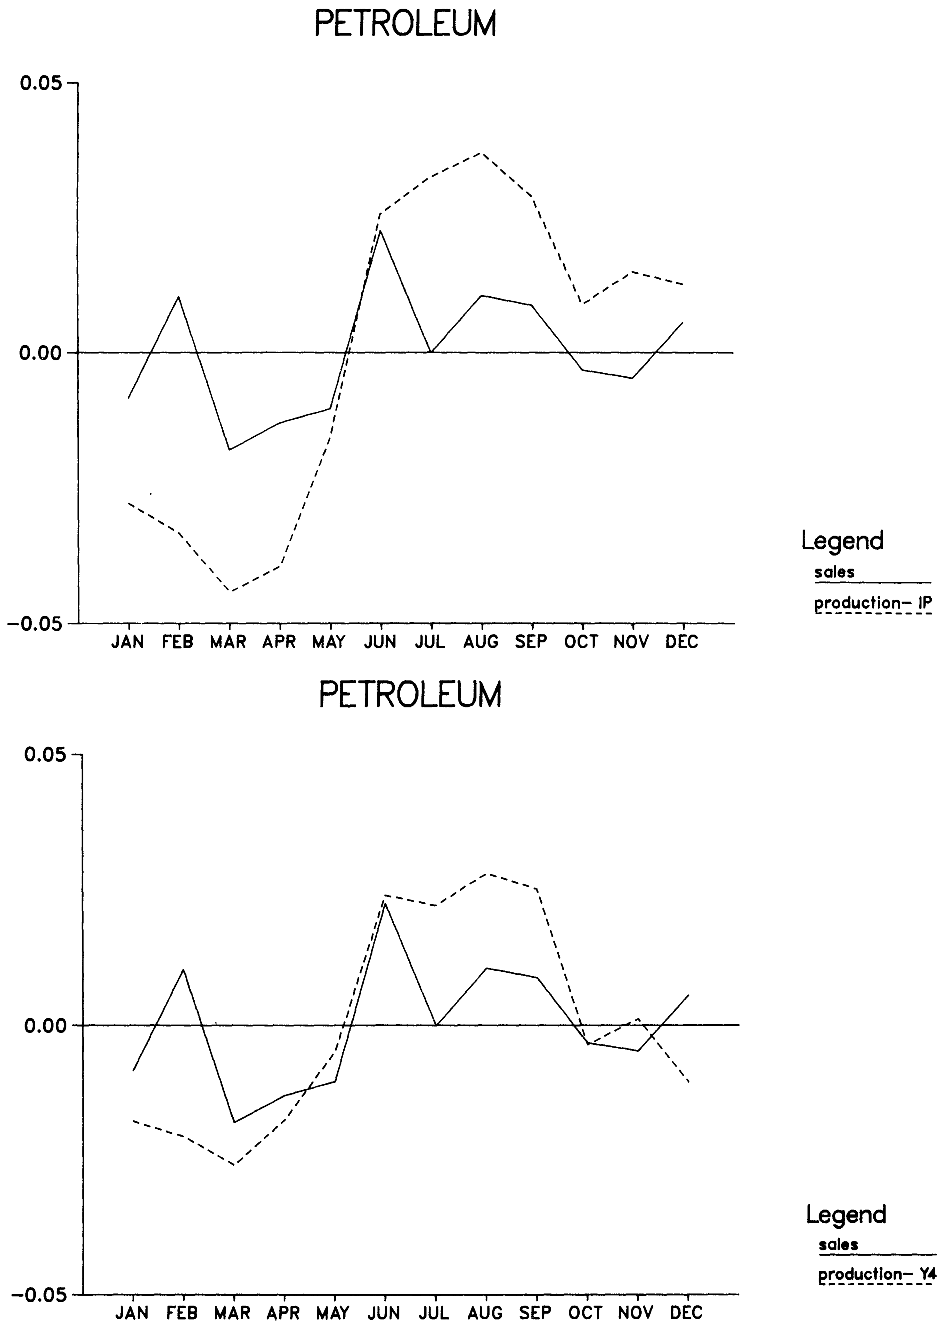
PETROLEUM
0.05
0.00
−0.05
JAN FEB MAR APR MAY JUN JUL AUG SEP OCT NOV DEC
Legend
sales
production– IP
PETROLEUM
0.05
0.00
−0.05
JAN FEB MAR APR MAY JUN JUL AUG SEP OCT NOV DEC
Legend
sales
production– Y4

FIGURE 5

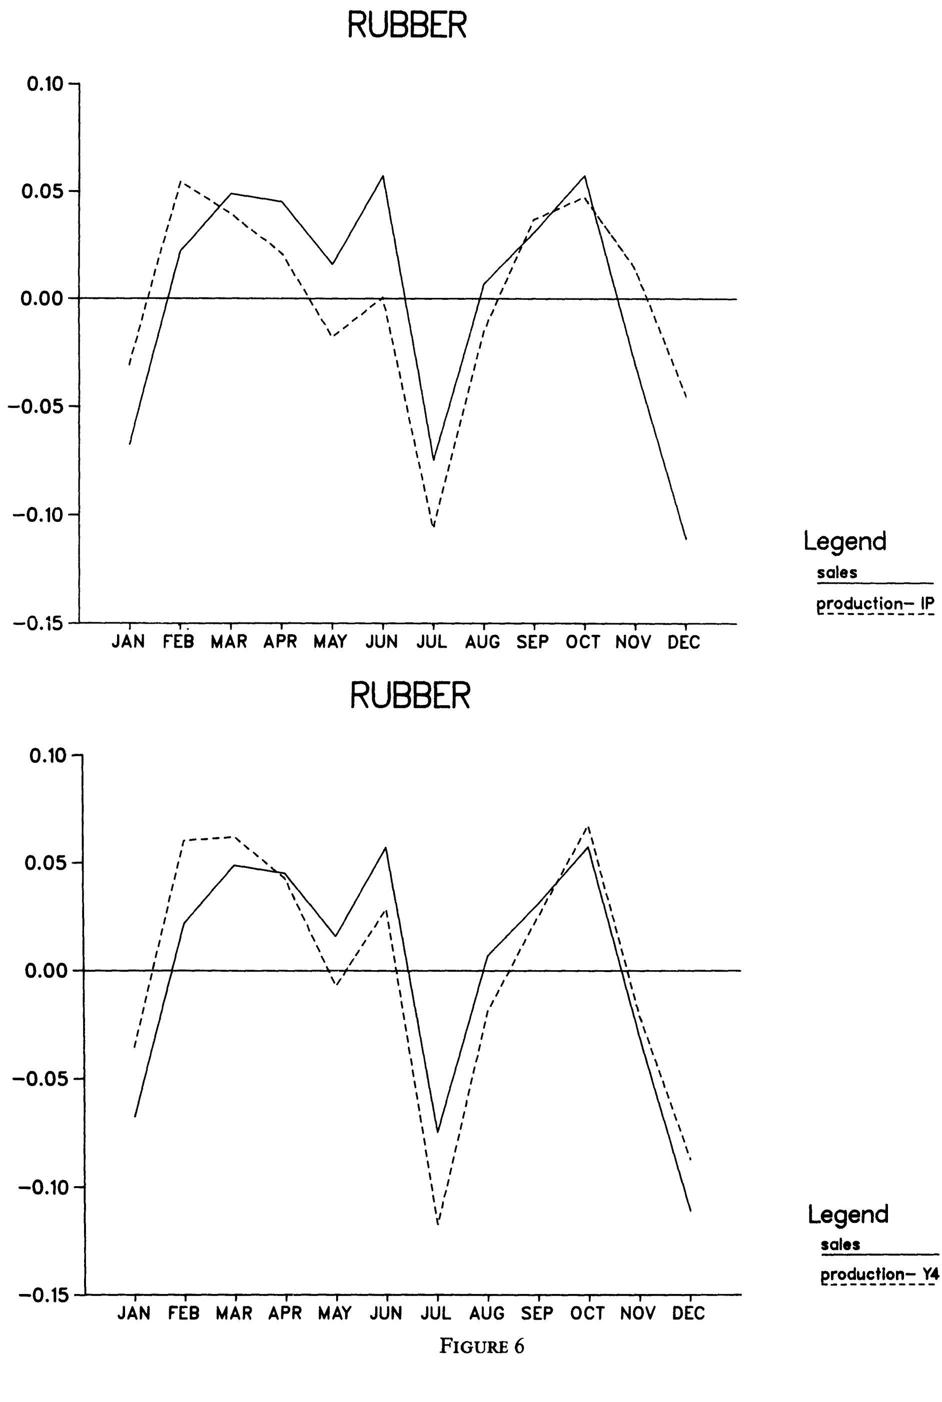
RUBBER
0.10
0.05
0.00
−0.05
−0.10
−0.15
JAN FEB MAR APR MAY JUN JUL AUG SEP OCT NOV DEC
Legend
sales
production– IP
RUBBER
0.10
0.05
0.00
−0.05
−0.10
−0.15
JAN FEB MAR APR MAY JUN JUL AUG SEP OCT NOV DEC
Legend
sales
production– Y4

FIGURE 6

into account the measured seasonality in factor costs, the weather, and the level of inventories.

This test of the production smoothing model using the seasonal fluctuations does involve one important maintained hypothesis, namely that the coefficients on the seasonal and nonseasonal components of input prices and the weather are the same. In our last set of tests, we relax this assumption and test whether the seasonal movements in the data, taken by themselves, are consistent with the model. This is accomplished as follows. The first step is to construct the seasonal component of each of the relevant variables (output growth, input prices, weather variables, etc.) by regressing them on seasonal dummies and calculating the fitted values of these regressions. We then regress the seasonal component of output growth on the seasonal in input prices, weather, the level of inventories, and the contemporaneous seasonal component of sales, and we test the restriction implied by the model that this last coefficient should be zero.[38,39]

The results are summarized in Tables IV and V. Table IV presents the same type of information as Table II, but it includes the t statistic on the seasonal component of contemporaneous sales in the test of the overidentifying restrictions. Table V is also set up similarly to Table II, but it simply reports whether the seasonal in sales significantly affects the seasonal in output growth, after controlling for the seasonal movements in input prices, the weather, and the level of inventories.

In both tables, there is striking evidence against the production smoothing model. In Table IV, we reject the overidentifying restrictions in every instance. In most cases, the seasonal component of sales is significantly correlated with the movement in output, even after taking account of any seasonals in input prices, lagged inventories, and the weather. This is true for five out of six industries using at least one of the output measures, and for three of the six industries using both output measures.

When we redo the estimates in Table IV using the alternative instrument lists discussed above (leaving out time t variables or lagged output growth) we again reject the overidentifying restrictions and find that the seasonal in sales growth is significantly correlated with the residual output growth in most cases.

In Table V (the seasonal-only results) the seasonal in sales growth is statistically significant in five out of six cases for the $Y4$ measure of output, and in four out of six cases for the IP measure. Variables other than sales almost never enter significantly.

These results on the behavior of production and sales at seasonal frequencies are perhaps the most problematic yet presented for the production smoothing model. To a large extent, firms appear to be choosing their seasonal production

[38] Since we found there to be essentially no seasonality in energy prices, raw materials prices, or interest rates, we excluded these variables.

[39] We implement the procedure above by estimating equation (10), with sales growth included, using seasonal dummies as the *only* instruments. The coefficient estimates are numerically identical to those produced by the procedure described in the text, but this instrumental variables procedure produces correct standard errors. The resulting t statistic on sales growth is reported in Table V.

TABLE IV

REGRESSION RESULTS, EQUATION (10) SEASONAL DUMMIES EXCLUDED, SEASONAL COMPONENT OF SALES IN INSTRUMENT LIST

		Food	Tobacco	Apparel	Chemicals	Petroleum	Rubber
Y4, NSA	What enters (10) significantly?	$-w, tem, -tem2$	$-w$	$w, tem, -tem2$	$-w, tem, -pre, -tem2$	$-w, -pre2, -day,$	$-w, tem, -tem2, -pre$
	R^2, Significance Level	.52, .000	.16, .001	.14, .003	.31, .000	.18, .000	.26, .000
	What is significant in test of OIR's?	$tem_{-1}, -tem2_{-1}, -x_{-2}, -rm_{-1}, xseas$	tem_{-1}	$xseas$	$tem_{-1}, xseas$	we_{-2}	$xseas$
	t-stat on *xseas*	3.51	0.51	2.84	4.14	1.87	3.41
IP, NSA	What enters (10) significantly?	$-w, tem, -tem2$	$-w$	$w, rm, tem, -tem2$	$tem, -tem2, -pre$	$-w, -pre2$	$tem, -tem2$
	R^2, Significance Level	.30, .000	.14, .003	.12, .012	.15, .002	.13, .007	.24, .000
	What is significant in test of OIR's?	$-y_{-1}, -y_{-2}, -w_{-2}, xseas$	$-y_{-1}, w_{-2}$	$-y_{-1}, xseas$	—	$-y_{-1}, x_{-1}, xseas$	$xseas$
	t-stat on *xseas*	2.86	0.99	2.62	1.82	3.39	3.41

Notes:

1. The sample period is 1967:5–1982:12.
2. The first line of each set of results lists the variables that entered equation (10) at the 5% significance level.
3. The second line gives the R^2 from the regression of the residuals on the instruments, as well as the marginal significance level of this statistic. The quantity $T \times R^2$ is distributed χ^2_j, where j is the number of overidentifying restrictions and T is the number of observations. In the results presented here, there are 10 such restrictions.
4. The third line lists the variables that entered the regression of the residuals on the instruments at the 5% significance level.
5. w = wage growth, y = output growth, x = sales growth, *day* = number of production days, *we* = energy price growth, *rm* = raw materials price growth, n = inventories, r = interest rate, *pre* = change in precipitation, *pre2* = change in precipitation squared, *tem* = change in temperature, *tem2* = change in temperature squared, *tpr* = change in temperature* precipitation, *xseas* = seasonal component of sales growth (defined as the fitted values from a regression on seasonal dummies).
6. A (−) before variable indicates that the sign of the coefficient was negative.
7. A subscript of $-l$ on a variable means that it is dated l periods earlier than the dependent variable.

TABLE V

REGRESSION RESULTS, EQUATION (10) SEASONAL COMPONENTS ONLY

		Food	Tobacco	Apparel	Chemicals	Petroleum	Rubber
Y4	What enters (10) significantly?	$-n, x$	x	x	x	—	x
	t stat on x	9.21	4.91	1.98	11.69	1.10	19.99
IP	What enters (10) significantly?	x	$-pre, x$	—	x	—	x
	t stat on x	3.19	3.39	1.14	2.31	.62	3.97

Notes:
1. The sample period is 1967:5–1982:12.
2. The coefficients were obtained by estimating equation (10), with sales growth included, using seasonal dummies only as instruments. This gives coefficient estimates numerically identical to the procedure described in the text, but it produces correct standard errors. The t statistic on sales growth reported in the table is from this instrumental variables regression.
3. The first line of each set of results lists the variables that entered equation (10) at the 5% significance level.
4. w = wage growth, y = output growth, x = sales growth, *day* = change in number of production days, n = inventories, *pre* = change in precipitation, *pre*2 = change in precipitation squared, *tem* = change in temperature, *tem*2 = change in temperature squared, *tpr* = change in temperature * precipitation.
5. A (−) before variable indicates that the sign of the coefficient was negative.

patterns to match their seasonal sales patterns, rather than using inventories to smooth production over the year. Moreover, since the seasonal variation in production and sales growth generally accounts for more than 50 percent of the total variation in these variables, this problematic behavior is a quantitatively important feature of the data.

A key assumption that we have made here is that the seasonal in the productivity shifter is uncorrelated with the seasonal in demand. Are there circumstances under which this assumption would not hold? An example that comes to mind is the case of an economy-wide seasonal in labor supply, namely that individuals, all else equal, would rather take vacations in certain months. This would induce a corresponding seasonal in output. If each industry's output is an input into another industry, then we might expect to see a corresponding seasonal in shipments, leading optimally to the same seasonal patterns in output and shipments.

Theoretically, our approach accounts for this by including the wage as a determinant of desired production. However, if the measured wage differs from the true shadow cost of utilizing labor, then the residual will include the seasonal in labor supply and therefore still be correlated with the seasonal in shipments. This explanation suggests that we should see the same seasonal movements in output in all industries. In Figures 1–6, we do see common seasonal patterns in output across industries, but we also see a fair amount of seasonal movement that is different across industries.

It is not clear what conclusion to draw from this discussion. It is possible that the hypothesis proposed above is the explanation for the seasonal results. If so, we should ask whether the same type of arguments could be made about nonseasonal movements, i.e., whether we believe that the failure of the production smoothing model at nonseasonal frequencies is due to economy-wide changes in desired labor supply that are not captured by measured wages.

7. CONCLUSIONS

The results presented above show a strong rejection of the production smoothing model. This is despite the fact that we have extended the standard model considerably, by allowing for nonquadratic technology, a stochastic interest rate, convex costs of holding inventories, and measurable and nonmeasurable cost shocks, and by including seasonal fluctuations explicitly. Although previous work has examined many of these features, none has simultaneously allowed for all of them.

The rejections of the basic production smoothing model that we report are robust with respect to the treatment of seasonal fluctuations. To begin with, we reject the model about as strongly when we treat seasonality in the standard way, by using adjusted data, as when we treat it more explicitly by specifying the economic sources of the seasonal movements in production and inventories. Even more surprisingly, our results show that the seasonal movements in production, inventories, and shipments are inconsistent with the basic model. Specifically, the seasonal component of output growth, even after adjusting for the seasonality in interest rates, wages, energy prices, raw materials prices, and the weather, is still highly correlated with the seasonal component of sales growth, contrary to the prediction of the model.

We conclude the paper by discussing what we believe to be the implications of our results for a number of hypotheses that have been offered for the failure of the production smoothing model. We first discuss those hypotheses on which our results provide direct evidence and then turn to more indirect implications.

Our results provide direct evidence that the limited role given to cost shocks in previous papers is not the major reason for the rejections of the model. In this paper we have included a more general set of cost shocks than in earlier work, and we still find that the data reject production smoothing. Moreover, we find relatively little evidence that cost shocks play any role in determining the optimal timing of production. It is possible, of course, that we have omitted the "key" cost shock, or that one of our identifying assumptions is invalid. We believe, however, that the set of costs we have included covers all of the major ones, and we think that the identifying assumptions we make are minimally restrictive. It seems to us unlikely, therefore, that the treatment of cost shocks is a major factor in explaining the poor performance of the model.

The second area in which our results provide direct evidence is on whether the inappropriate use of data seasonally adjusted by X-11 has been responsible for the failure of the model. As we discussed above, X-11 data are (approximately) a two-sided moving average of the underlying seasonally unadjusted data. This means that such data likely violate the crucial orthogonality conditions that are tested in the kinds of models considered above, even if the unadjusted data satisfy them. Although it seemed likely to us on a priori grounds that the use of

X-11 adjusted data was a major problem, our results indicate otherwise. The particular method of treating seasonal fluctuations does not appear crucial to an evaluation of the model.

So much for direct implications. We now turn to more indirect implications, specifically, the implications of our tests using the seasonal movements in the data. These implications are subject to the critique that the production smoothing model may fit differently at different frequencies, in which case we may not be able to learn about the validity of the model at nonseasonal frequencies from its performance at seasonal frequencies. However, to the extent that the same model is underlying the different movements in the data, we can draw the following conclusions.

To begin with, since seasonal fluctuations are anticipated, it seems unlikely that the failure of the model at seasonal frequencies could be due to any kind of irrationality or disequilibrium. If so, this rules out a large class of possible explanations of the failure of the model.

A second issue that is illuminated by our seasonal specific results is that of costs of changing production. We have omitted costs of changing production (or, more generally, costs of changing inputs) from our specifications above; the addition of these costs might "help the data fit the model." We regard this tactic as unsatisfactory, however. The fact that there are extremely large seasonal changes in the rate of production makes it seem quite unlikely that there are large costs of adjustment, although it is true that costs may be lower when they are anticipated.

Blinder (1986a) suggests that the production smoothing model could be saved by including persistent demand shocks and small cost shocks. Even if the nonseasonal movements in sales are very persistent, however, the same is not true of the seasonal movements. Therefore, our seasonal results suggest that Blinder's explanation will not suffice to "save" the production smoothing model.

Finally, our seasonal specific results allow us to rule out a concern regarding the choice of appropriate instruments. In our estimation, we assume that firms know current demand, and therefore time t sales is a valid instrument. In contrast, others, such as Kahn (1986), assume that firms do not know the level of current period demand when they choose the current period level of production. If this assumption is a more appropriate abstraction, then our general results are inconsistent. When we correct for this by using only variables dated $t-1$ and earlier as instruments, we can no longer reject the model. However, it is still valid to include the seasonal component of contemporaneous sales growth, since the seasonal component of demand would be known even if the overall level were not. Since the results from this test show a strong rejection of the model, this suggests that the assumption that firms observe demand before choosing output is not, by itself, to blame.

What remains, then, as a possible explanation for the failure of the production smoothing model? There are two main possibilities: nonconvexity of the cost function, and stockout costs. Giving up convexity of costs is unappealing because it requires also giving up much of neoclassical theory. This does not mean it is not the correct explanation; it simply suggests that we should turn to it only as a last resort. We end, therefore, by discussing the role of stockout costs.

An important maintained assumption above is that firms always hold positive inventories, which implies that firms do not stock out. Total inventories for each industry are always positive in our data, but this may not be the case for each individual firm or product. Kahn (1986) presents a model in which, because stockouts are costly, firms may not smooth production.[40] However, there is as yet relatively little direct evidence that stockout costs are high, or that firms cannot simply hold unfilled orders as a type of negative inventories. This line of research deserves further attention, in particular direct empirical testing.

Department of Economics, University of Michigan, Ann Arbor, MI 48109, U.S.A.
and
Finance Department, The Wharton School, University of Pennsylvania, Philadelphia, PA 19104-6367, U.S.A.

Manuscript received December, 1986; final revision received August, 1987.

REFERENCES

ABEL, ANDREW B. (1985): "Inventories, Stock-Outs and Production Smoothing," *Review of Economic Studies*, 52, 283–293.

BELSLEY, DAVID A. (1969): *Industrial Production Behavior: The Order-Stock Distinction*. Amsterdam: North-Holland.

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

BLINDER, ALAN S. (1986a): "Can the Production Smoothing Model of Inventory Behavior Be Saved?" *Quarterly Journal of Economics*, 101, 431–454.

——— (1986b): "More on the Speed of Adjustment in Inventory Models," *Journal of Money, Credit, and Banking*, 18, 355–365.

——— (1982): "Inventories and Sticky Prices: More on the Microfoundations of Macroeconomics," *American Economic Review*, 72, 334–348.

BLINDER, ALAN S., AND STANLEY FISCHER (1981): "Inventories, Rational Expectations, and the Business Cycle," *Journal of Monetary Economics*, 8, 277–304.

CHRISTIANO, LAWRENCE J. (1986): "Why Does Inventory Investment Fluctuate So Much?" manuscript, Federal Reserve Bank of Minneapolis.

CHRISTIANO, LAWRENCE J., AND MARTIN EICHENBAUM (1986): "Temporal Aggregation and Structural Inference in Macroeconomics," manuscript, Carnegie-Mellon University.

CLEVELAND, W., AND G. TIAO (1976): "Decomposition of Seasonal Time Series: A Model for the Census X-11 Program," *Journal of the American Statistical Association*, 71, 581–587.

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

FELDSTEIN, MARTIN, AND LAWRENCE SUMMERS (1979): "Inflation and the Taxation of Capital Income in the Corporate Sector," *National Tax Journal*, 32, 445–470.

GARBER, PETER, AND ROBERT KING (1984): "Deep Structural Excavation? A Critique of Euler Equation Methods," NBER Technical Working Paper #83-14.

GHALI, MOHEB (1987): "Seasonality, Aggregation, and Testing of the Production Smoothing Hypothesis," *American Economic Review*, 77, 464–469.

GHYSELS, ERIC (1987): "Cycles and Seasonals in Inventories: Another Look at Non-Stationarity and Induced Seasonality," manuscript, University of Montreal.

GOODFRIEND, MARVIN (1986): "Information-Aggregation Bias: The Case of Consumption," manuscript, Federal Reserve Bank of Richmond.

[40]Abel (1985) also examines production smoothing in a model with stockouts.

HALL, ROBERT E. (1978): "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, 86, 971–987.

HANSEN, LARS, AND KENNETH J. SINGLETON (1983): "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns," *Journal of Political Economy*, 91, 249–266.

HINRICHS, JOHN C., AND ANTHONY D. ECKMAN (1981): "Constant Dollar Manufacturing Inventories," *Survey of Current Business*, 61, 16–23.

IRVINE, F. OWEN, JR. (1981): "Retail Inventory Investment and the Cost of Capital," *American Economic Review*, 71, 633–648.

KAHN, JAMES (1986): "Inventories and the Volatility of Production," manuscript, University of Rochester.

MACCINI, LOUIS J., AND ROBERT J. ROSSANA (1984): "Joint Production, Quasi-Fixed Factors of Production, and Investment in Finished Goods Inventories," *Journal of Money, Credit, and Banking*, 16, 218–236.

MANKIW, N. GREGORY (1981): "The Permanent Income Hypothesis and the Real Interest Rate," *Economics Letters*, 7, 307–311.

MIRON, JEFFREY A. (1986): "Seasonality and the Life Cycle-Permanent Income Model of Consumption," *Journal of Political Economy*, 94, 1258–1279.

MIRON, JEFFREY A., AND STEPHEN P. ZELDES (1987): "Production, Sales, and the Change in Inventories: An Identity that Doesn't Add Up," Rodney L. White Working Paper #20-87, The Wharton School, University of Pennsylvania.

NELSON, CHARLES R., AND CHARLES I. PLOSSER (1982): "Trends and Random Walks in Macroeconomic Time Series: Some Evidence and Implications," *Journal of Monetary Economics*, 10, 139–162.

PRESCOTT, EDWARD C. (1986): "Theory Ahead of Business Cycle Measurement," Carnegie-Rochester Conference Series, 25, 11–44.

REAGAN, PATRICIA, AND DENNIS P. SHEEHAN (1985): "The Stylized Facts about the Behavior of Manufacturers' Inventories and Backorders over the Business Cycle: 1959–1980," *Journal of Monetary Economics*, 15, 217–246.

SARGENT, THOMAS J. (1978): "Rational Expectations, Econometric Exogeneity, and Consumption," *Journal of Political Economy*, 86, 673–700.

SCHUTTE, DAVID P. (1983): "Inventories and Sticky Prices: Note," *American Economic Review*, 73, 815–816.

SUMMERS, LAWRENCE (1981): "Comment on 'Retail Inventory Behavior and Business Fluctuations' by Alan Blinder," *Brookings Papers on Economic Activity*, 2, 513–517.

WALLIS, KENNETH F. (1974): "Seasonal Adjustment and the Relation between Variables," *Journal of the American Statistical Association*, 69, 13–32.

WARD, MICHAEL P. (1978): "Optimal Production and Inventory Decisions: An Analysis of Firm and Industry Behavior," unpublished Ph.D. thesis, University of Chicago.

WEST, KENNETH, D. (1983): "A Note on the Econometric Use of Constant Dollar Inventory Series," *Economics Letters*, 13, 337–341.

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

ZELDES, STEPHEN P. (1985): "Consumption and Liquidity Constraints: An Empirical Investigation," Rodney L. White Center Working Paper #24-85, The Wharton School, University of Pennsylvania.