Rochester Center for
Economic Research

What Inventory Behavior Tells Us About Business Cycles

Bills, Mark and James A. Kahn

Working Paper No. 428
September 1996

University of
Rochester

What Inventory Behavior Tells Us About Business Cycles

Mark Bils and James A. Kahn

Rochester Center for Economic Research
Working Paper No. 428

September 1996

SB August 1996

What Inventory Behavior Tells Us about Business Cycles

Mark Bils

University of Rochester and NBER

and

James A. Kahn

University of Rochester

August 1996*

Abstract

We argue that the behavior of manufacturing inventories provides evidence against models of business cycle fluctuations based on productivity shocks, increasing returns to scale, or favorable externalities, whereas it is consistent with models with short-run diminishing returns and procyclical work effort. Both finished goods and work-in-process inventories move proportionally much less than sales or production over the business cycle, facts which we show imply procyclical marginal cost. Obvious measures for marginal cost do not show temporarily high marginal cost near peaks, as required to rationalize the inventory behavior, because measured factor productivity rises and then gradually declines during the peak phase of the cycle. We can explain the cyclical behavior of inventory holdings by allowing for procyclical work effort, the cost of which is internalized by firms but is not contemporaneously reflected in measured wage rates.

* We gratefully acknowledge support from the National Science Foundation. We thank Anil Kashyap, Peter Klenow, Valerie Ramey, and participants at numerous seminars for their helpful comments. We also thank Leonard Loebach of the Commerce Department for providing the inventory data.

I. Introduction

There are two distinct views on the nature of business cycle fluctuations. In one view business cycle peaks represent expansion of production possibilities, lower production costs, and increases in productivity--much like a bountiful harvest. This would include models based on increasing returns such as Farmer and Guo (1994), as well as standard real business cycle models (e.g., Kydland and Prescott, 1982). According to the second view, at peaks of the cycle capacity constraints and diminishing returns kick in, driving up the costs of production and partially stabilizing fluctuations. Many researchers have viewed the procyclical behavior of inventory investment as evidence favoring countercyclical marginal cost because it suggests that firms bunch production more than is necessary to match the fluctuations in sales. If short-run marginal cost curves are fixed and upward sloping (the argument goes), firms will smooth production relative to sales, making inventory investment countercyclical.[1]

This paper argues that this reasoning is false: We should expect to observe procyclical inventory investment even with increasing marginal cost. What the above reasoning overlooks are changes in the shadow value of inventories, which we argue increase with the level of production and expected sales.[2] We propose a model in which

[1]Hall (1991, pages 34-37) presents this view. See West (1985), Blinder (1986), and Fair (1989) for evidence on production volatility and the cyclical behavior of inventory investment. Eichenbaum (1989) introduces unobserved cost shocks that generate simultaneous expansions in production and inventory investment. Ramey (1991) estimates a downward sloping short-run marginal cost function, which of course reverses the production-smoothing prediction. Cooper and Haltiwanger (1992) adopt a nonconvex technology on the basis of observations about inventory behavior. Others (e.g., Gertler and Gilchrist, 1993, Kashyap, Lamont, and Stein, 1994) argue that credit market imperfections--essentially countercyclical inventory holding costs for some firms--are responsible for what is termed "excess volatility" in inventory investment.

[2]Pindyck (1994) makes a related point regarding what he calls the "convenience yield" of inventories. A number of papers in the inventory literature do include a target inventory-sales ratio as part of a more general cost function to similarly generate a procyclical inventory demand. Many of these papers, for example Blanchard (1983), West (1986), Krane and Braun (1991), Kashyap and Wilcox (1993), and Durlauf and Maccini (1995), estimate upward-sloping marginal cost in the presence of procyclical inventory investment. This appears consistent with our evidence that marginal cost is procyclical.

finished inventories facilitate sales, with the marginal value of inventories proportional to the level of sales. This generates a desired inventory to expected sales ratio.

The well-known finding that inventory investment is procyclical implies that inventory stock is procyclical, but not that it keeps up with sales. In fact, inventory-sales ratios are extremely countercyclical. Figure 1 plots monthly the ratios of finished goods inventory to sales for 1947 to September 1995 along with the detrended logarithm of constant-dollar production for manufacturing. NBER defined recessions are shaded. The inventory-sales ratio increases dramatically in each recession, typically by 10 to 20 percent. Note that this rise does not simply reflect a transitory response to an unexpected fall in sales, but rather continues throughout recessions. Replacing sales with forecasted sales generates a virtually identical figure. The correlation between detrended production and the ratio of finished goods inventories to sales in manufacturing is -0.54. The upshot is that although inventory investment is procyclical, it is not nearly as procyclical as sales.[3]

What really needs explaining is why inventory investment is not *more* procyclical. The marginal unit of finished goods is predictably associated with more sales at a peak than at a trough. We provide evidence from trends and cross-sectional data of the absence of any scale effects, so that absent higher costs or a lower markup the firm should want to add more inventories in those periods with low stock-to-flow ratios so as to equate the ratios (and hence the "returns") over time. That they fail to do so implies that marginal cost must be high relative to the future in those periods with relatively high production and sales.

The next section presents a model in which finished goods inventories facilitate sales and the holding of work in process facilitates production. We then examine data for six two-digit manufacturing industries that produce primarily to stock. These data

[3]We find similar results for finished goods inventories and works-in-process for new housing construction and for finished goods inventories in wholesale and retail trade.

reinforce the picture from aggregate data in Figure 1--finished and work-in-process inventories are highly countercyclical relative to sales or production. Yet obvious measures for marginal cost based on measured inputs and factor prices do not show temporarily high marginal cost in booms, as required to explain this inventory behavior. This is because input prices are not very procyclical and productivity does not fall in booms. We can rationalize the cyclical patterns of inventory holdings, however, by allowing for procyclical work effort that is not contemporaneously reflected in wage rates. This explanation is consistent with our finding that inventories lag production and sales during a boom as the capital stock is in the process of catching up with labor input and output.

We find the joint behavior of inventories, costs, and quantities consistent with the following view of business cycles: In a boom the capital stock lags behind labor and output. This temporarily raises marginal cost, inducing firms to squeeze inventory-sales ratios. Short-run fixity of factors also leads to higher factor utilization, generating a transitory rise in Solow residuals. which hides the true short-run increase in costs.

II. Inventory Fundamentals

A. *The Derived Demand for Finished Goods Inventories*

In this section we present the production and inventory decision for a representative producer. We show that producers choose an optimal ratio of stock (in excess of overhead) available to expected sales that varies in relation to anticipated interest rates and growth in marginal cost.

We define the stock of goods available for sale during period t, a_t, to equal the inventory, i_t, of unsold goods carried forward from the previous period plus the goods completed in t, z_t. Deviations of z_t from output y_t imply variations in a firm's inventory of work in process, which we will treat separately below. Producers maximize the expected present-discounted value of profits.

$$\max_{z_t,\, y_t} \quad E_t \sum_{i=0}^{\infty} \beta_{t,t+i}[\, p_{t+i}s_{t+i} - C_{t+i}(y_{t+i})]$$

subject to:

$$a_t = i_t + z_t = a_{t-1} - s_{t-1} + z_t, \text{ and}$$

$$s_t = d_t(p_t)[a_t - \hat{a}]^{\phi} .$$

Here s_t and p_t denote respectively sales and price (relative to a numeraire) in period t. $C_t(y_t)$ is the cost of producing output y_t during t, also in terms of the numeraire, based on a production technology which we will describe in more detail below. $\beta_{t,t+i}$ denotes the real rate of market discount at time t for i periods ahead in terms of the numeraire.[4]

The key feature of the producer's problem is that sales are a function of the producer's available stock. For a given price, a producer views its sales as increasing with an elasticity of ϕ with respect to its available stock above some threshold value $\hat{a}$. Our approach is consistent, for example, with a competitive market that allows for the possibility of stockouts (e.g., Kahn, 1986, Thurlow, 1995). In this case $\hat{a}$ would equal zero and ϕ would equal one. Alternatively, one could think of the stock as an aggregate of similar goods of different sizes, colors, locations, and the like. This relationship could approximate a matching function in which a random draw of purchasers arrives with a demand for a specific type of good. More generally, we simply want to capture the idea

[4]In the empirical work we incorporate a storage cost for both finished inventories and work in process. This effectively lowers β_{t+1} as it now reflects both a rate of storage cost, δ, as well as a real interest rate r_{t+1}: $\beta_{t+1} = \frac{1-\delta}{1 + r_{t+1}}$.

We also explored whether the relative cost of storing to producing goods varies cyclically. We examined a case in which storage is completely capital intensive, so that storing each unit of inventories requires a certain number of units of capital. If the output-capital ratio is procyclical, as it is, then the storage cost will be procyclical. When we introduced this procyclical storage cost into the estimation, however, it failed to enter consistently significantly or correctly signed, and had little effect on the estimates of other parameters. For these reasons, below we present results for a constant δ.

that a producer's demand increases at least slightly with its available inventory; if it did not, there would be no reason to hold positive inventories on average, much less the one to three months' worth of sales that we typically observe.

We also allow the demand for the producer to move proportionately with a stochastic function $d_t(p_t)$. Again, this is consistent with a perfectly competitive market in which charging a price below the market price yields sales equal to a_t and charging a price above market clearing implies zero sales. The function $d_t(p_t)$ will more generally depend on total market demand and available supply. In fact, it can depend on any number of factors. All we require is that the impact of the firm's stock a_t be captured by the separate multiplicative term $(a_t - \hat{a})^{\phi}$.

If z_t is positive (which we assume), then the impact on expected discounted profits of producing one more unit of z_t (holding $z_t = y_t$) must equal zero. This condition is

$$E_t \{ - c_t + \phi d_t(p_t)[a_t - \hat{a}]^{\phi\text{-}1} p_t + [1 - \phi d_t(p_t)[a_t - \hat{a}]^{\phi\text{-}1}]\beta_{t+1} c_{t+1} \} = 0 .$$

(For convenience we write $\beta_{t,t+1}$ as simply β_{t+1}.) The producer incurs marginal cost $c_t \equiv C_t'(y_t)$. By increasing the available stock, sales are increased by $\phi d_t(p_t)[a_t - \hat{a}]^{\phi\text{-}1}$. These sales are at price p_t. To the extent the increase in stock available does not increase sales, it does increase the inventory carried forward to t+1. This inventory can displace a comparable amount of production in t+1, saving its marginal cost c_{t+1}.

Note that the marginal impact on sales, $\phi d_t(p_t)[a_t - \hat{a}]^{\phi\text{-}1}$, is proportional to the ratio of sales to stock available above some threshold $\hat{a}$, equalling $\frac{\phi s_t}{a_t - \hat{a}}$. Making this substitution and rearranging gives

(1) $$E_t \{ [\frac{\phi s_t}{a_t - \hat{a}} m_t + 1] v_{t+1} \} = 1 ,$$

$$\text{where } v_{t+1} = \frac{\beta_{t+1} c_{t+1}}{c_t} ,$$

$$\text{and } m_t = \frac{p_t - \beta_{t+1}c_{t+1}}{\beta_{t+1}c_{t+1}} .$$

v_{t+1} equals the discounted gross rate of growth in marginal cost. In a pure production smoothing model, ϕ equal to zero, its expectation is always one. m_t is the percent markup of price above the present value of marginal cost in t+1. We denote this by m_t for markup because, with production to stock, $\beta_{t+1}c_{t+1}$ is the opportunity cost of selling a unit during t. Therefore, it is the relevant measure for marginal cost determining pricing decisions.

It is useful to consider the very special case in which the growth in costs, the real interest rate, and price-cost markups are all constant through time. This implies the expectation of $\frac{s_t}{a_t - \hat{a}}$ is constant. If $\hat{a}$ is close to zero (which we discuss momentarily), then all predictable movements in sales are matched by proportional movements in the stock available.[5] To generate persistent procyclical movements in the ratio of sales to inventory, which we see very strongly in the data, requires at least one of the following: A countercyclical markup, a procyclical real interest rate, or procyclical costs. Thus the behavior of inventories actually points *against* increasing returns or countercyclical costs. It also points against the idea that credit constraints bind in contractions, causing some firms to shed inventories. To account for the data, credit constraints would need to bind in expansions, which is precisely opposite the scenario emphasized by Gertler and Gilcrest (1993), Kashyap, Lamont, and Stein (1994), and others.

We view the case of $\hat{a}$ equal to zero as a useful benchmark. Inventory quantities are often stated in terms of an inventory-sales ratio. The model produces a desired stock available relative to sales, and therefore a desired inventory-sales ratio, only if the threshold value $\hat{a}$ equals zero. We can examine the behavior of $\frac{s_t}{a_t}$ to see whether assuming $\hat{a}$ equal to zero is reasonable, as it implies that the steady-state sales to inventory ratio should be independent of the size of the industry or firm.

[5] In a steady state with constant marginal cost the ratio $\frac{a - \hat{a}}{s}$, equals $\frac{\phi m(1-\delta)}{r + \delta}$; r here is the real interest rate and δ is the rate of storage cost.

Some evidence can be gleaned from seeing whether the ratio changes very much over time in industries where there is substantial growth. Below we examine six manufacturing industries in detail (tobacco, apparel, lumber, chemicals, petroleum and rubber). With the exception of tobacco, all these industries display substantial upward trends in production and sales. For all but tobacco and petroleum, sales increased by 50 percent or more from 1967 to 1995. In Figure 2 we present the behavior of $\frac{s_t}{a_t}$ for each of the industries for that period. For all six industries any long-run movement in the ratio is fairly small even when the level of s_t changes considerably. The largest trend movement is for lumber, where the ratio rises by about 15 percent, whereas in apparel it actually declines by about 5 percent. There is also little trend in most of the ratios of production to work-in-process inventories. Exceptions are chemicals, where it rises by about 25 percent, and petroleum, where it falls by about 30 percent.

An assumption that the threshold $\hat{a}$ is close to zero is also supported by cross-sectional evidence. Kahn (1992) reports average inventory-sales ratios and sales across divisions of U.S. automobile firms. These are reproduced here.

Sales and Inventory-Sales Ratio by Auto Division, 1966-1983

	Sales (1000 cars)	Inventory/Sales
General Motors		
Chevrolet	545	0.79
Pontiac	199	0.68
Oldsmobile	204	0.67
Buick	178	0.69
Cadillac	67	0.54
Ford Motors		
Ford	448	0.79
Lincoln-Mercury	136	0.95
Chrysler		
Chrysler-Plymouth	210	1.30
Dodge	131	1.24
American Motors	67	1.03

These data show no particular tendency for the ratio to be related to the size of the division, either within or across firms.

Gertler and Gilchrest (1993) present inventory-sales ratios for manufacturing by firm size, with size defined by firm assets. We reproduce their Table 4 here.

Inventory-Sales Ratio for Manufacturing by Size Class

	Cumulative Asset Size Class (in Millions of Dollars)				
Year	≤ 25	≤ 50	≤ 250	≤ 1000	All Manuf.
1960	0.58	0.61	0.66	0.66	0.72
1970	0.51	0.54	0.63	0.70	0.74
1980	0.47	0.49	0.53	0.55	0.53
1990	0.47	0.49	0.52	0.54	0.52

Again there is little relation between size and inventory-sales ratio. If anything, larger firms hold a higher inventory-sales ratio.

We conclude that scale effects do not appear a promising explanation for the fall in inventory-sales ratios in booms. When we do estimate the size of the threshold term $\hat{a}$ below, it is typically less that 20 percent of the average size of a_t, and its introduction does not significantly affect our results. Furthermore, allowing for a threshold $\hat{a}$ that is half as large as the average stock available does not qualitatively affect our results that marginal cost must be quite procyclical to explain inventory behavior.

B. *The Role of Work-in-Process Inventories and Cyclical Time to Build*

Finished-goods inventories appear to earn a higher return in booms. We consider procyclical marginal cost the most likely explanation, though a countercyclical output price markup would also work. We now examine the behavior of materials and work-in-process (WIP) inventories. Because decisions regarding these stocks involve only technological tradeoffs on the production side and are separate from the selling technology (and specifically from the output price), their behavior provides a simpler test for the cyclical behavior of marginal cost. The empirical work will focus on WIP inventories only, but inventories of materials tell a very similar story.

We specify a production technology that includes a role for work in process following Kydland and Prescott (1982), Christiano (1988), and Ramey (1989).[6] We assume that holding larger WIP stocks, like capital, increases labor's productivity. Increasing WIP could correspond to having more parallel processes going for each worker, thereby facilitating specialization and increased productivity. Reducing the number of assembly lines while increasing their speed would correspond to decreasing WIP while maintaining output. Presumably this would require more labor input (otherwise it would represent unexploited profits).

By holding WIP inventories firms lower their production costs, justifying the holding cost. Of course there is a tradeoff: WIP inventories will be high *relative to production* if interest rates are low or if marginal cost is transitorily low. Intuitively, if marginal cost is temporarily low it is a good time to accumulate work in process in order to lower next period's production cost. In Section 3 we show that the ratio of work in process to production is highly countercyclical, suggesting either a procyclical real interest rate or procyclical marginal cost. Let x_t be the beginning of period stock of work in process. Let y_t denote the value of production in t, whereas z_t denotes the value of completed output. By definition, x_t increases if y_t is greater than z_t; that is, Δx_t equals $(y_t - z_t)$.

We assume a Cobb-Douglas production function in terms of production labor n_t, nonproduction labor l_t, capital k_t, and work in process x_t.

$$(2) \qquad y_t = \theta_t n_t^{\alpha} l_t^{\nu} k_t^{\eta} [x_t - \hat{x}]^{\rho} .$$

$\hat{x}$, similarly to $\hat{a}$, is a parameter included for generality. Technology shocks, or omitted

[6]Kydland and Prescott (1982) and Christiano (1988) do not distinguish empirically between work in process and finished inventories. Ramey (1989) estimates cost functions involving each of the three stages of production separately.

factors, are allowed through the term θ_t. Note that we have not imposed constant returns to scale. We calibrate the power on x_t, ρ, to be quite small. on the order of 0.01. This value is consistent with observed ratios of x to y given reasonable values for interest rates and storage costs.

As both production, y_t, and completed output, z_t, are choice variables, this yields an additional first-order condition. Cost minimization requires that increasing the value of production in t then decreasing it in t+1, holding constant the value of completions, has zero impact on costs. This condition yields

(3) $$E_t \{ [\frac{\rho y_{t+1}}{x_{t+1} - \hat{x}} + 1] v_{t+1} \} = 1 .$$

Just as finished goods inventories should be judged relative to the size of expected sales, it is important to judge work in process relative to the flow of production. Work-in-process will be high *relative to production* if interest rates are low or if marginal cost is transitorily low.[7] Empirically, work in process is typically much less procyclical than production. This requires either a procyclical real interest rate. procyclical marginal cost, or important scale affects through the parameter $\hat{x}$. We do not find an explanation based on scale effects promising. As discussed above. differences in size cross-sectionally and across time do not reveal important scale effects. Furthermore. our estimates below for $\hat{x}$ suggest it is not an important factor.

The first-order conditions for finished inventories and work in process can be combined to arrive at

(4) $$E_t \{ [\frac{\phi s_t}{a_t - \hat{a}} m_t - \frac{\rho y_{t+1}}{x_{t+1} - \hat{x}}] v_{t+1} \} = 0 .$$

Suppose $\hat{a}$ and $\hat{x}$ are each relatively, which estimation below would suggest. Then. absent important cyclical fluctuations in markup of price over marginal cost, the ratio of

[7] In a steady state with constant marginal cost the ratio $\frac{x - \hat{x}}{y}$, equals $\frac{\rho(1-\delta)}{r + \delta}$. r is the real interest rate and δ is the rate of storage cost.

sales to stock available and the ratio of production to work in process should move together. In fact, both ratios are highly procyclical.

C. *Digression: Relation to the Linear-Quadratic Model*

Much of the inventory literature estimates linear-quadratic cost-function parameters (e.g. West, 1986, Eichenbaum, 1989, or Ramey, 1991). A typical specification of the single-period cost function (ignoring additive disturbance terms) is[8]

$$C(y_t, i_t) = \psi y_t^2 + \zeta(i_t - \alpha s_t)^2 .$$

The slope of marginal cost is governed by the parameter ψ. Note that $\zeta > 0$ allows for a target inventory-sales ratio. While many researchers (Blinder, 1985, Fair, 1989, among others) have focused on the relative volatility of production and sales, it is easy to show (see West, 1986) that for $\alpha > 0$, having $\psi > 0$ does not imply that the variance of sales exceeds the variance of production, or that inventory investment is countercyclical.

We would argue that use of the linear-quadratic specification misses the more revealing information coming from the behavior of the inventory-sales ratio. As a consequence, much effort has gone into accounting for why inventory investment is procyclical and why production is more volatile than sales, whereas the questions should be: Why is inventory investment not *more* procyclical? Why is production not even more volatile?

To see this, consider the following specification

$$C(y_t, i_t) = \psi y_t^2 + \zeta(a_t - \alpha s_t)^2 .$$

where either ψ or ζ, or both, are positive. We replace i_t in the previous expression by $a_t = i_t + y_t$, as in our model. It is straightforward to prove by a variance bounds

[8]A number of papers in this literature include a cost term in the change in output. Its exclusion in our discussion is simply for convenience.

argument, similar to that of West (1986), that if $\frac{a_t}{s_t}$ is countercyclical then in the absence of procyclical cost shocks γ must be positive. The proof also goes through if we replace sales by expected sales, and similar bounds can be found for other specifications. Thus it could not be optimal for a firm to systematically have a low ratio when sales are high if its marginal cost is decreasing in production. Such a firm could reduce costs by bunching production in periods when sales are high, thereby having a procyclical ratio. This supports the more general intuition that countercyclical inventory-sales ratios indicate procyclical marginal cost.

So how it is that some researchers (in particular Ramey, 1991) who include an inventory-sales target still find downward-sloping marginal cost with data that have countercyclical stock-sales ratios? One possibility is that the linear quadratic specification is not a good approximation to the true model. This is suggested by West's (1986) rejections of the model based on variance bounds tests and by Pindyck's (1994) results for the shape of his convenience yield function. Alternatively, those results may be sensitive to the details of the specification, such as whether a_t or i_t enters the targets, how parameters are normalized (see Krane and Braun's, 1991, discussion), the choice of instruments, and whether cost shocks are allowed. Furthermore, our reading of the literature is that most authors do typically estimate marginal cost to be upward sloping (e.g., Blanchard, 1983, West, 1986, Krane and Braun, 1991, Kashyap and Wilcox, 1993, and Durlauf and Maccini, 1995).

In contrast to the linear-quadratic literature, our model explicitly considers the revenue side of the firm's problem. This allows us to account for significant features of inventory data caused by price variations. According to our model, the return on finished inventory is proportional to the price markup, thus sales relative to stock available should move inversely with the markup.

The tobacco industry provides an excellent experiment in this regard. The price of tobacco products rose very dramatically from 1984 to 1993. Figure 2 shows the behavior of the producer price for tobacco relative to the general PPI as well as the ratio of sales to stock available. The relative price doubled. Although material costs in tobacco rose during this period, the relative price change appears to have largely reflected a rise in price markup (Howell et al., 1994). Consistent with the model, the ratio of sales to finished goods available fell by about 15 percent, whereas the ratio of production to work in process was essentially unchanged. More striking is what occured in 1993. During one month, August 1993, the price of tobacco products fell by 25 percent, presumably reflecting a breakdown in collusion (see Figure 2). Within 3 months the ratio of sales to finished goods available rose dramatically, as predicted by the model, by at least 25 percent. In contrast, but also predicted by our model, the ratio of production to work in process showed no noticeable effect. In sum, the linear-quadratic model is silent on the large movements in inventory-sales ratios that accompanied these developments, whereas the model in this paper contains a ready explanation.

III. Empirical Implementation

A. *Measuring Marginal Cost of Production*

To estimate first-order conditions (1) and (3) requires estimating the behavior of marginal cost, c_t. In turn, this requires specifying a cost function or production technology. We assume that production is Leontief with respect to material input. The cost of output can be written as

$$c_t = \lambda_t \omega_t + \tilde{c}_t .$$

ω_t denotes the price of materials; and λ_t denotes the real material content of a unit of output. Note that λ_t is allowed to vary through time, but is independent of the choice of output or other inputs. $\tilde{c}_t$ is the marginal cost of producing a unit of real value added. We are assuming a Cobb-Douglass production technology as described in equation (2).

Let w_t denote the real wage cost, relative to a numeraire price deflator, of a marginal increase in labor input. Then the marginal cost of value added is

$$\tilde{c}_t = (\frac{1}{\alpha})\frac{w_t n_t}{y_t} ,$$

which is proportional to the wage divided by production workers' labor productivity.[9] This result for marginal cost does not depend on our treatment of work in process. For example, it continues to be true if the power on work-in-process inventory, ρ, equals zero. It also holds even if work-in-process enters in some more general fashion. The measure also allows for technology shocks, the impact of which appear through output.

With data on output, materials cost, production hours, and the effective wage rate. marginal cost can be calculated given a value for the production labor parameter α.

$$c_t = \lambda_t \omega_t + (\frac{1}{\alpha})\frac{w_t n_t}{y_t} . \tag{5}$$

A value for α equal to labor's share roughly corresponds to perfect competition. Because higher values for α imply lower values of marginal cost, they also imply a higher measure of the markup given data on real output prices.

Our approach differs substantially from the inventory literature. That literature, cited above, typically estimates a structural quadratic cost function including quadratic terms in output and typically in the change in output. Measures for cost shocks, such as wage changes, are also sometimes included (e.g., Ramey, 1991, or Durlauf and Maccini, 1995). By contrast we are exploiting the production function to work with the reduced form for marginal cost. This measure allows not only for changes in output, wages, the

[9]Bils (1987) takes this approach to measuring marginal cost.

cost of capital, or other inputs, but also for shocks to productivity. Thus we do not need instruments that are orthorgonal to technology or other shocks.

B. *Measuring the Marginal Price of Labor Input*

The standard practice in the literature is to measure the price of production labor by a measure of average hourly earnings for production earnings. We deviate from the literature by considering a competing measure that allows for the possibility that worker utilization, that is effort, is procyclical, and that this procyclical effort is largely neglected in contemporaneous wage rates.[10]

Total factor productivity is markedly procyclical for most manufacturing industries. One interpretation for this finding is that factors, including labor, are utilized more intensively in booms, but these movements in utilization are not captured in the measured cyclicality of inputs (e.g., Solow, 1973). This does not necessarily invalidate our measure of marginal cost. If hourly wage rates vary to capture the cyclicality of workers' effort then it continues to give the correct answer. Our concern is that wages may not reflect the spot price of labor, but rather are smoothed relative to labor's effective price for convenience or to smooth workers' incomes. (See Hall, 1980.)

Let μ_{nt} denote the effort per hour of production labor; so total production labor input in t equals $\mu_{nt}n_t$. Then the marginal cost of value added is, similarly to before,

$$\tilde{c}_t = \left(\frac{1}{\alpha}\right)\frac{w_t(\mu_{nt}n_t)}{y_t} .$$

[10]We also considered a wage measure that captures the fact that actual compensation packages are convex with respect to the workweek, so that the marginal wage rate is higher and more cyclical than average hourly earnings. Bils (1987) shows that a marginal wage is much more procyclical than an average wage rate for manufacturing because in booms a marginal expansion in labor is much more likely to incur overtime pay. We estimated marginal wage rates for our six manufacturing industries. In all six the estimated marginal wage was substantially more procyclical than the industry's average wage rate. Nonetheless, the implications for inventory investment were not significantly altered when we employed this wage measure in our empirical work in place of average hourly earnings. Although the workweek is strongly procyclical, it apparently fails to decline throughout an expansion as required to explain the persistent fall in inventory-sales ratios during booms.

w should be interpreted as the marginal price of a unit of effort, or *effective* labor. If average hourly earnings vary with $w_t\mu_{nt}$ then there is no need to adjust for procyclical effort. Otherwise it is necessary to gauge the size of variations in μ_{nt}. For our second wage measure we assume that average hourly earnings are set on the basis of an average anticipated effort level, say $\overline{\mu}_n$, but employers treat the true labor cost as $w_t\mu_{nt}$. Thus it is necessary to "add back" the variations in μ_{nt} onto average hourly earnings.

To construct a measure of movements in effort we make several assumptions. The first is that the production of value added exibits constant-returns-to-scale

$$y_t = \theta(\mu_{nt}n_t)^{\alpha}(\mu_{lt}l_t)^{\nu}(\mu_{kt}k_t)^{1-\alpha-\nu} .$$

Note that we have allowed for unmeasured variations in nonproduction labor and capital (μ_{lt} and μ_{kt}) as well as for production labor. We do not allow for transitory high-frequency technological shocks in the parameter θ.[11] For this exercise we ignore work in process as a factor of production by setting the parameter ρ equal to zero. In practice this is unimportant as the values for ρ we consider below are very small (less than 0.01).

Marginal cost of value added is then

$$(6) \qquad \tilde{c}_t = \theta^{\frac{-1}{\alpha}} \frac{w_t}{\alpha}\left(\frac{y_t}{\mu_{kt}k_t}\right)^{\frac{1-\alpha}{\alpha}}\left(\frac{\mu_{kt}k_t}{\mu_{lt}l_t}\right)^{\nu} = \theta^{\frac{-1}{\alpha+\nu}} \frac{w_t}{\alpha}\left(\frac{y_t}{\mu_{kt}k_t}\right)^{\frac{1-\alpha-\nu}{\alpha+\nu}} .$$

The second equality reflects an additional assumption that there are no relative price changes between production and nonproduction workers, which implies that utilized nonproduction labor moves proportionately with utilized production labor.[12] It is necessary to take some stand here because there are no data on effort levels--or for that

[11]For estimating the WIP first-order condition (3) we can allow for random-walk variations in θ_t, as this has no impact on expected growth of marginal cost, in our estimated equations.

[12]It is sometimes argued that skill premia fall in booms. If true, this implies effort that is even more procyclical.

matter on hourly workweeks--for nonproduction workers. Effort is judged to move in relation to movements in the ratio of output to utilized capital, with the factor of proportionality being capital's share $(1-\alpha-\nu)$ relative to one minus capital's share. In the empirical work we gauge the cyclicality of μ_k on the basis of the cyclicality of electricity consumption per unit of physical capital.

IV. Data and Results

A. *The Behavior of Inventories*

We begin by examining the cyclical behavior of the ratio of sales to stock available for sale $\frac{s_t}{a_t}$, and the ratio of production to WIP inventories $\frac{y_t}{x_t}$. We obtained manufacturing sales and inventories by stage of construction, all in constant dollars and seasonally adjusted, from the Department of Commerce. These series are monthly and disaggregated to the SIC two-digit level. They are available back to 1967. We construct monthly production from the identity for inventory accumulation, with production equal to sales plus inventory investment.

We examine six manufacturing industries. The six are tobacco, apparel, lumber, chemicals, petroleum, and rubber. These are roughly the six industries commonly identified as production for stock industries (Belsley, 1969). We have deleted food and added lumber. We are concerned that some large food industries, such as meat and dairy, hold relatively little inventories. Thus any compositional shift during cycles could generate sharp shifts in inventory ratios. Our impression of the lumber industry is that it is for all practical purposes production to stock, though there are very small orders numbers collected. This view was reinforced by discussions with Census.

Figure 3 presents the ratios $\frac{s_t}{a_t}$ and $\frac{y_{t+1}}{x_{t+1}}$ for each of the six industries. The ratio of production to work in process is for t+1 to be consistent with first-order condition (3). The figure also includes linearly detrended production for that industry as well shading for NBER defined aggregate recessions. The period is for 1967 through 1993.

In almost every case the ratios of sales to stock available and production to work in process are highly procyclical. An industry boom is associated with a much larger percentage increase in sales than the available stock in each of the six industries. Furthermore, the decline in stock available relative to sales persists through the cycle. Thus it makes little difference if we replace sales with expected sales. This is reinforced by Table 1 which presents correlations between the ratios of sales to stock available and detrended output. The correlations are all significantly positive. Again, replacing sales with a forecast of sales would generate similar results. The ratios for production relative to WIP are similarly procyclical in four of the six industries. These correlations appear in Table 1 as well. In tobacco the correlation is essentially zero and for petroleum it is small and negative.[13]

We want to stress the tendency for these ratios to be procyclical is not peculiar to these six industries. Figure 1 depicted a similar picture for aggregate manufacturing, and we also observe it in home construction, the automobile industry, and in wholesale and retail trade. Furthermore, for most of these six industries production is more volatile than sales, as it is for aggregate manufacturing.[14]

---

[13]We also first differenced the series, looking at the correlation of the changes in the ratios $\frac{s_t}{a_t}$ and $\frac{y_t}{x_t}$ with the rate of growth in output. For the change in $\frac{s_t}{a_t}$ the correlations are very positive, ranging across industries from 0.35 to 0.54, and averaging 0.46. (Using forecasted growth in $\frac{s_t}{a_t}$ yields even higher correlations.) For the changes in $\frac{y_t}{x_t}$ the correlations are sensitive to the timing. The change in $\frac{y_t}{x_t}$ is very positively correlated with output growth, but the change in $\frac{y_{t+1}}{x_{t+1}}$ is somewhat negatively correlated with output growth in t. High output growth in t predicts negative output growth in t+1, which is then associated with a fall in $\frac{y_{t+1}}{x_{t+1}}$. We introduced timing in the model in such a way that $\frac{y_{t+1}}{x_{t+1}}$ enters into the works in process investment decision in t. But this is very arbitrary. If we had introduced end-of-period WIP into costs (i.e., it is expensive to complete every single unit of work in process in a period) then y_t, rather than y_{t+1}, would enter into period t's WIP investment decision.

[14]Production and sales are measured relative to a linear trend. Production is more volatile in apparel, lumber, petroleum, and rubber. Sales are more volatile in tobacco. For chemicals the two series are equally volatile.

B. *Implications for the Behavior of Marginal Cost*

We have argued that the failure of inventory stocks to keep pace with sales and production suggests that marginal cost is temporarily high in booms and low in recessions. One way to depict this is to ask what behavior of interest rates and costs would be exactly consistent with the observed behavior of inventories. Assuming the variables in first-order condition (1) are conditionally distributed jointly lognormal, the equation can be manipulated to yield

$$(1') \qquad E_t\left(\frac{c_{t+1}}{c_t}\right) \approx E_t\left[r_{t+1} - \frac{\phi m s_t}{a_t}\right] + \text{constant} .$$

The constant term reflects covariances between the random variables as well as storage costs.[15] For now we set the "threshold level" for a_t, $\hat{a}$, equal to zero. We reintroduce $\hat{a}$ in the estimation below, but it has limited impact. For purposes here we assume a constant markup. We set ϕm equal to 0.038, this is consistent with a steady state ratio $\frac{s}{a}$ equal to 0.6 for a monthly storage plus interest rate cost of 2.25 percent: $\frac{s}{a}$ equal to 0.6 is roughly consistent with what we typically observe for our six industries. To get a conditional expectation of the right-hand side of (1′) we project onto a set of variables Γ that contains $\frac{s_{t-1}}{a_{t-1}}$, $\frac{y_t}{x_t}$, c_t, p_{t-1}, $\tilde{y}_t$, $\frac{c_t}{c_{t-1}}$, r_{t-1}, R_t, and $\frac{y_t \tilde{c}_t}{x_t c_t}$. R_t denotes the nominal commercial paper rate, r_t the ex post real rate based on PPI inflation.

Results for this exercise appear in the first half of Table 2 for two separate cases. First we set the real interest rate equal to a constant that equals the mean of the six-month commercial paper rate for 1967 to 1993. (This could also be interpreted as computing the implied expected growth in marginal cost relative to the real interest rate.) Secondly we allow the real interest to vary with the actual movements in the

[15]This approximation is arbitrarily good for small values for the real interest rate r and for the ratio $\frac{m\phi s}{a}$. For monthly data r should be on the order of 0.01 or less. In steady-state the ratio $\frac{m\phi s}{a}$ equals r plus the monthly storage rate. So we would argue this is a small fraction on the order of 0.02.

commercial paper rate. Looking first at the constant interest rate case, the behavior of finished goods inventories requires that marginal cost is temporarily high in booms in every one of the six industries. This can also be seen from Table 2, Column 1, which displays the correlation between expected growth and costs and detrended output. This correlation is $-.35$ or more negative in each case. Taking account of movements in the commercial paper rate does not alter this basic picture. Figure 4 presents only the second measure plotted against detrended output, because the two measures are sufficiently close that they would be hard to distinguish in the pictures.

A comparable exercise can be conducted based on explaining the predictable variations in the ratio of production to work-in-process inventories. The comparable equation for works in process, again assuming conditional joint log normality, is

$$(3') \qquad E_t\left(\frac{c_{t+1}}{c_t}\right) \approx E_t\left[\frac{\rho y_{t+1}}{x_{t+1}}\left(\frac{\tilde{c}_{t+1}}{c_{t+1}}\right) - r_{t+1}\right] + \text{constant} .$$

Here we set the threshold $\hat{x}$ equal to zero. When we introduce $\hat{x}$ in the estimation below it has fairly minor impact. For this simulation we set $\rho(\frac{\tilde{c}_{t+1}}{c_{t+1}})$ equal to 0.006. This is consistent with a steady-state interest plus storage cost of 2.25 percent monthly and a production to work-in-process ratio of four. This ratio is typical of our industries, though it varies considerably.

The results for this exercise appear in Figure 5 and Columns 3 and 4 of Table 2. The results are very similar to those based on finished goods inventories, except for tobacco and petroleum. For these two industries there is no requirement that marginal cost be cyclical regardless of whether we assume a constant or varying interest rate.

C. *The Behavior of Marginal Cost Measures*

We now compare the movements in marginal cost required by the inventory fluctuations, as depicted in Figures 4 and 5, with actual movements based on our empirical measures derived above. Repeating from above, our basis for measuring marginal cost is the equation

$$(5) \qquad c_t = \lambda_t \omega_t + \left(\tfrac{1}{\alpha}\right)\frac{w_t n_t}{y_t} .$$

Above we described alternative measures for the price of labor, w. Here we need to describe how me measure the cost of materials $\lambda_t \omega_t$ and the production parameter α.

We know of no monthly data on material inputs or material price deflators. We construct our own monthly price of materials index, ω_t, for each industry as follows. Based on the 1977 input output matrix, we note every 4-digit industry whose input constituted at least 2 percent of gross output for one of our six industries. This adds up to 13 industries. We then construct a monthly index for each industry weighting the price movements for those 13 goods by their relative importance. For most of the industries one or two inputs constitute a large fraction of material input; for example, crude petroleum for petroleum refining or leaf tobacco for tobacco manufacture. For the residual material share we use the general producer price index. Although we assume that materials are a fixed input per unit of output, we do not impose that this input, λ, be constant through time. We allow for low frequency movements in λ by imposing that our series $\lambda\omega_t$ exhibit the same Hodrick-Prescott filter as that industry's material input as measured from the annual survey of manufacturing. (These data are from the NBER Productivity Database.)

To measure α we proceed as follows. We assume a constant share for labor broadly defined to include production and nonproduction labor. This implies a constant sum of the parameters $\alpha + \nu$. We allow α to vary at low frequencies, however, based on how production workers' earnings have varied as a share of total labor compensation. Production workers' earnings are based on BLS Establishment data; total compensation is based on National Income Accounts. Low frequency movements are again defined by Hodrick-Prescott filters. It remains necessary to tie down the absolute level of α. We do so by choosing an average price markup over marginal cost of 10 percent.[16]

The results measuring the price of labor by average hourly earnings appear in Table 3, Column 1 and in Figure 6. The correlations and the figure relate expected growth in marginal cost (based on the instruments discussed above) to detrended output. Marginal cost is very countercyclical for each industry--temporarily low and rising in a boom and temporarily high and falling in a recession. This reflects the behavior of labor productivity. As is well known, labor productivity is typically high in the beginning of an expansion, but then declines. This implies mirror image behavior for this measure of marginal cost, which the model argues is not compatible with the behavior of inventories.

Our alternative measure of marginal cost assumes constant returns to scale production and disallows high-frequency variations in technology, but allows for cyclical variations in the effort exerted by labor. As discussed at length above, mechanically this involves replacing labor productivity in equation (5) with a measure of the ratio of utilized capital to output, raised to a power equal to capital's share in value added relative to one minus capital share. Thus it is not surprising that it could dramatically reverse the implications for marginal cost.

Our measure of industry capital stocks comes from the Commerce Department and is annual for 1967 to 1993. We interpolate to obtain a monthly series. To obtain a measure of utilization we project annual electricity consumption per unit of capital on the industry annual output to capital stock ratio. The electricity consumption is from the Annual Surveys of Manufacturing. The elasticity of electricity use with respect to output averages 0.56 across the six industries, varying from 0.42 to about 0.67. We then multiply this elasticity by the monthly output–capital stock ratio to generate a monthly utilization rate for capital.

[16]Marginal cost here is measured by $\beta_{t+1}c_{t+1}$, where β_{t+1} reflects both the rate of storage cost and the real interest rate; so even a perfectly competitive industry would exhibit a small average price markup by this measure.

The results for the third measure appear in Table 3, Column 3 and in Figure 7. With this measure, the expected change in marginal cost now becomes dramatically procyclical for tobacco, lumber, chemicals, and rubber. The negative correlations in Table 3 for these four industries are all highly statistically significant.[17] For apparel marginal cost is now acyclic. Petroleum is largely unaffected, reflecting the small share of value added in gross output; it remains quite countercyclical. The lack of procyclical cost for petroleum is very noteworthy because (recall Table 2 and Figures 4 and 5) its inventory behavior does not require procyclical marginal cost.

Thus allowing for unobserved movements in effort dramatically alters the cyclical behavior of marginal cost. The implied movements in effort are not particularly large: The standard deviation of our implicit effort measure (relative to an H-P trend) is substantially less than the standard deviation of measured hours for every industry except tobacco. The median cases are lumber and petroleum, for which the standard deviations in effort are respectively 30 and 40 percent the sizes of the standard deviations in measured hours.

D. *Composition Effects*

There is one remaining feature of the data that we address prior to estimation. For some industries there are low frequency movements in the series $\frac{s}{a}$ and $\frac{y}{x}$ that the theory suggests should be stationary. A natural explanation for these movements, given the relatively high level of aggregation of the data, is that there are composition effects due to shifts in production and sales in the direction of industries with higher or lower inventory-sales ratios. To gauge this we examine 4-digit annual inventory, sales, and output data from the *Annual Survey of Manufacturers.* We look at the ratios $\frac{s}{a}$ and $\frac{y}{x}$ across the 4-digit industries near the midpoint of the sample (end of 1979) and then fit

[17] Projecting one variable on the other generates t-statistics ranging from 8 to 32.

trend lines to the composition effect created by changes in the composition of sales or output across these 4-digit industries over time. Thus for a 2-digit industry in which sales shifted towards high $\frac{s}{a}$ 4-digit industries, we would make a composition adjustment by removing what would be a positive compositional trend in $\frac{s}{a}$.

For some industries (notably rubber and tobacco) these composition effects are substantial. The results presented below incorporate this adjustment in $\frac{s}{a}$ and $\frac{y}{x}$ for these industries, but it turns out to have little effect on the results. More naive filtering out of low-frequency movements, such as linear or Hodrick-Prescott detrending, also do not substantially alter the results.

E. *Estimation of the First-Order Conditions*

We now examine the model more formally by estimating the first-order conditions by GMM. The statistics presented thus far suggest that the effort wage measure is more consistent with inventory behavior than the AHE measure. We will examine this more closely by estimating the first-order conditions for finished inventories and for WIP separately. Prior to that, however, we can do some structural estimation without taking a stand on costs by estimating the combined equation (4). We do that first, and then run a "horse race" between the two marginal cost measures before finally estimating the first-order conditions.

The combined equation (4) does not depend on marginal cost or the interest rate except through the markup and $\frac{\tilde{c}_{t+1}}{c_{t+1}}$. Its loglinearized version is

$$E_t[\, \frac{\rho y_t \tilde{c}_{t+1}}{x_t c_{t+1}} - \frac{\phi m_t s_t}{a_t} \,] + \kappa \;=\; 0 \,.$$

κ denotes a constant. For this exercise we set the threshold parameters $\hat{a}$ and $\hat{x}$ equal to zero. If, for the moment, we assume that m_t and $\frac{\tilde{c}_{t+1}}{c_{t+1}}$ are constants, then we can estimate $\frac{\phi}{\rho}$ independently of how we measure cost and interest rates.

In order to easily evaluate the estimates across industries, we multiply the ratios $\frac{y}{x}$ and $\frac{s}{a}$ by the values for ρ and ϕ that are consistent with that industry's steady state values for $\frac{y}{x}$ and $\frac{s}{a}$. These values, ρ^* and ϕ^*, reflect assumptions about average markups, interest rates and storage costs as discussed above. The values by industry are:

Steady-State Parameter Values

	ϕ^*	ρ^*
Tobacco	.094	.0034
Apparel	.386	.0140
Lumber	.354	.0144
Chemicals	.398	.0100
Petroleum	.659	.0200
Rubber	.369	.0094

We estimate the equation in the form

$$\frac{\rho^* y_t \tilde{c}_{t+1}}{x_t c_{t+1}} = \frac{\phi/\rho}{\phi^*/\rho^*} E_t[\frac{\phi^* ms_t}{a_t}] - \frac{\rho\kappa}{\rho^*} .$$

According to the model, the estimated parameter is $\frac{\phi/\rho}{\phi^*/\rho^*}$, which should be about one.

The results are in Table 4. All the estimates are significantly positive and the right order of magnitude, but typically greater than one, suggesting either that $\phi > \phi^*$ and/or $\rho < \rho^*$. Thus the two ratios do covary with each other as the model predicts, but with work-in-process inventories somewhat more variable relative to finished goods inventories than predicted.

Next we let the two cost measures face off against each other. To highlight the their behavior we write the finished-inventory equation (1′) as

$$E_t[\, r_{t+1} - \frac{\phi^*(\hat{a})ms_t}{a_t - \hat{a}}] \approx E_t(\frac{c_{t+1}}{c_t}) + \kappa$$

$$= \gamma_1 E_t(\frac{c_{t+1}}{c_t})_{AHE} + \gamma_2 E_t(\frac{c_{t+1}}{c_t})_{Effort} + \kappa .$$

κ, again, denotes a constant. $(\frac{c_{t+1}}{c_t})_{AHE}$ and $(\frac{c_{t+1}}{c_t})_{Effort}$ denote respective measures of growth in costs using average hourly earnings or our effort wage. We have multiplied s_t by a parameter $\phi^*(\hat{a})$ corresponding to the value for ϕ consistent with that industry's average $\frac{s}{a}$ ratio given an estimate for the parameter $\hat{a}$. Values for ϕ^* for $\hat{a}$ equal to zero appear on the page above. Also we impose a constant markup as otherwise m_t would depend on the marginal cost measure. (We relax this below.) Conditional on the inventory model being correct, the parameter γ_1 should equal one and γ_2 should equal zero if average hourly earnings is the correct measure of the price of labor. Conversely, if the effort wage is the correct measure, then γ_1 should equal zero and γ_2 should equal one.

We estimate a comparable horserace from the behavior of WIP inventory, using equation (3′)

$$E_t[\, r_{t+1} - \frac{\rho^*(\hat{x}) y_{t+1} \tilde{c}/c}{x_t - \hat{x}}] \approx E_t(\frac{c_{t+1}}{c_t}) + \kappa$$

$$= \gamma_1 E_t(\frac{c_{t+1}}{c_t})_{AHE} + \gamma_2 E_t(\frac{c_{t+1}}{c_t})_{Effort} + \kappa\, .$$

κ again denotes a constant. We have imposed a constant for the ratio $\frac{\tilde{c}}{c}$, as otherwise the ratio depends on the particular cost measure. (We relax this momentarily.) We have multiplied y_{t+1} by a parameter $\rho^*(\hat{a})$ corresponding to the value for ρ consistent with the industry's average ratio $\frac{y}{x}$ given an estimate for the parameter $\hat{x}$. Values for ρ^* for $\hat{x}$ equal to zero appear above.

Results for finished goods appear in Table 5. For each industry we first estimate γ_1 and γ_2 imposing that the parameter $\hat{a}$ be equal to zero. We then estimate the parameter $\hat{a}$ as well. In Table 5 we report the estimate of $\hat{a}$ relative to the average value of stock available during the sample period, $\overline{a}$. The results strongly indicate that the cost series based on the effort wage is the one consistent with inventory behavior, except for petroleum, where both measures fare poorly. The coefficient is significantly positive for the effort wage in the remaining five industries, and in several cases with magnitude on the order of one. For the cost measure based on average hourly earnings, by contrast, the coefficient is close to zero or of the wrong sign with the sole exception of tobacco. And even for tobacco, the effort wage has the stronger impact.[18] These statements are essentially unaffected by allowing a nonzero threshold parameter $\hat{a}$, with the possible exception of lumber. The estimates for $\hat{a}$ are significantly positive only for lumber and chemicals. In these industries the threshold value $\hat{a}$ equals respectively 0.27 and 0.19 of the average stock available. For the remaining four industries the estimate of $\hat{a}$ is insignificantly different from zero, except in apparel where it is significantly negative.[19]

Results for WIP are given in Table 6. The results again strongly indicate that the cost series based on the effort wage is the one consistent with inventory behavior. The sole exception is petroleum, where both measures are insignificant. These results parallel those for finished goods very closely. The main difference is that the coefficient for the effort-based cost is of the right magnitude in almost all cases. The result that inventory behavior is consistent with the procyclical effort-based cost holds true after allowing for a nonzero parameter $\hat{x}$. The estimate for $\hat{x}$ is significantly positive in chemicals and rubber. It is significantly negative for apparel. For two industries. tobacco and petroleum, estimation of $\hat{x}$ does not converge. To check for robustness, we estimate γ_1 and γ_2 setting $\hat{x}$ equal to one half the average industry value for x. A value for $\frac{\hat{x}}{\bar{x}}$ of 0.5 is larger than any that we estimate. This has a dramatic effect on the estimates only in petroleum; where previously they had been insignicant at any rate.

Lastly, we reestimate the Euler equations relaxing the assumption that the markup is a constant and relaxing the assumption that $\frac{\tilde{c}_t}{c_t}$ (the marginal cost of value added relative to all marginal cost) is a constant. Given the relative success of the effort-

[18]We also estimated comparable equations multiplying the ratio $\frac{s}{a}$ by zero rather than by ϕ^*. This corresponds to a pure production smoothing model. The results continue to strongly favor the cost measure that incorporates some procyclical effort.

[19]We also considered specifications in which $\hat{a}$ could have an exponential trend, with broadly similar results.

wage cost measure in Table 5, we do the remaining estimation solely based on this cost measure. The estimated equations are

$$E_t[\, r_{t+1} - \frac{\phi^*(\hat{a})ms_t}{a_t - \hat{a}}] = \gamma E_t(\frac{c_{t+1}}{c_t}) + \kappa \ , \text{ and}$$

$$E_t[\, r_{t+1} - \frac{\rho^*(\hat{x})y_{t+1}\tilde{c}/c}{x_t - \hat{x}}] = \gamma E_t(\frac{c_{t+1}}{c_t}) + \kappa \quad .$$

The parameter γ is, of course, predicted to be one in all cases, but we will estimate it freely for both equations and all industries. We first present results imposing that the threshold parameters $\hat{a}$ and $\hat{x}$ equal zero. We then attempt to estimate these parameters as well.[20]

The results for finished goods appear in Tables 7 and 8. To illustrate the impact of the varying markups, we first present results in Table 7 continuing to impose constant markups. These results are well anticipated by Table 5. The estimate of γ is significantly positive in all but petroleum. The estimate is also of the right magnitude in apparel, chemicals, and rubber. Furthermore, the results are largely unaffected by allowing for a nonzero estimate for $\hat{a}$.

Two points stand out from the results in Table 8 adding varying markups. The parameter estimates increase dramatically in magnitude. Apparel, lumber, chemicals. and rubber all have significantly positive estimates for γ, but for all but lumber the estimate is well above the predicted value of one. Secondly, the parameter estimate for tobacco becomes significantly negative. This is puzzling in light of the fact that we have already argued that the qualitative behavior of $\frac{s_t}{a_t}$ in response to the major price fluctuations in the tobacco industry supported the model. It would appear, however,

[20] In calculating variations in markups, we set the mean markups at 10 percent. Tobacco, however, because of the large runup of prices in the 1983-93 period, required a larger average markup (about 30 percent) to avoid having persistent negative markups prior to 1983. Also, petroleum, which has a materials share of 90 percent, requires a smaller average markup (we set it at 5 percent). Note that the interest rate enters into the calculation of the markups. We also use the 6-month commercial paper rate in calculating the markup.

that our constructed markup movements are proportionally larger than the contrary movements in $\frac{s_t}{a_t}$, with the result that the movements in $\frac{m_t s_t}{a_t}$ are negatively correlated with $\frac{s_t}{a_t}$. In other words, m_t "overcorrects" for the the movements in $\frac{s_t}{a_t}$. The most plausible explanation is that the true markup did not move as much as the price and cost data suggest, perhaps because some of the markup rents are dissipated in ways not accounted for in the input-output tables.

Estimates for $\hat{a}$ in the presence of time varying markups do not converge. To check for robustness to the parameter $\hat{a}$, we estimate γ setting $\hat{a}$ equal to one half the average industry value for a. This is a larger value for $\frac{\hat{a}}{a}$ than were found in any of the constant-markup estimates from Table 7. Changing the value of $\frac{\hat{a}}{a}$ from zero to 0.5 has very little impact on the results.

Finally, Table 9 presents the estimation results for WIP. The estimate of γ is significantly positive in tobacco, lumber, chemicals, and rubber. This estimate is of the correct order of magnitude in each of these industries except lumber. Again these results are not at all sensitive to estimating a nonzero value for $\hat{x}$ (or setting $\hat{x}$ to half of the average value of x, when estimates for $\hat{x}$ do not converge).

In sum, the evidence strongly favors our measure of marginal cost that interprets movements in labor productivity as reflecting small variations in work effort rather than as shifts in technology. There are some limitations to our findings: The model's precise quantitative predictions are not consistently borne out by the data, and the model's overidentifying restrictions are consistently rejected. The model is also entirely (though not surprisingly) unsuccessful as applied to the petroleum industry. It should be noted, though, that our approach places very strong restrictions on the data. In particular, we measure marginal cost rather than parameterize and estimate it, and we do not freely estimate trend terms or filter out other low-frequency movements of our key variables.

V. Conclusions

Evidence from cross-sectional and low frequency data indicate that firms' demands for finished goods and work-in-process inventories are derived proportionally from firms' sales and production. Yet during business cycles these inventories are highly countercyclical relative to sales and production. Obvious measures for marginal cost do not show temporarily high marginal cost in booms as required to justify this inventory behavior, because factor productivity rises then gradually declines during booms. Though we consider alternative explanations—-including overhead inventory and convex storage costs—-we find little evidence for anything but procyclical marginal cost (broadly defined to include the interest rate) as an explanation of the cyclical behavior of inventories. We show that the cyclical patterns of inventory holdings are most consistent with the interpretation of fluctuations in labor productivity as primarily mismeasured work effort, the cost of which is internalized by firms but not contemporaneously reflected in measured wage rates.

Our results challenge the empirical basis for models that generate large fluctuations from procyclical productivity due to technology shocks, increasing returns, or favorable externalities. Any such model should have to explain why inventory-holding firms fail to increase production more during booms when their inventory stocks are more productive, and less during recessions when the return on these stocks is down. Our view that procyclical work effort accounts for this puzzle is consistent with other evidence that workers work more intensively in booms (Fay and Medoff, 1985, Bernanke and Parkinson, 1991, and Bils and Cho, 1994), but leaves open the question of what is the ultimate source of business cycle fluctuations.

References

Belsley, D.A., *Industry Production Behavior: The Order-Stock Distinction.* Amsterdam: North-Holland, 1969.

Bernanke, Ben S., and Martin L. Parkinson, "Procyclical Labor Productivity and Competing Theories of the Business Cycle: Some Evidence from Interwar U.S. Manufacturing," *Journal of Political Economy* 99 (June 1991), 439-59.

Bils, Mark, "The Cyclical Behavior of Marginal Cost and Price," *American Economic Review* (1987).

Bils, Mark and Jang-Ok Cho, "Cyclical Factor Utilization," *Journal of Monetary Economics* 33 (March 1994).

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

Blinder, Alan, "Can the Production Smoothing Model of Inventory Behavior be Saved?" *Quarterly Journal of Economics* 101 (1986), 421-54.

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

Cooper, R. and J. Haltiwanger, "Macroeconomic Implications of Production Bunching: Factor Demand Linkages," *Journal of Monetary Economics* 30 (1992), 107-128.

Durlauf, Steven N. and Louis J. Maccini, "Measuring Noise in Inventory Models," *Journal of Monetary Economics* 36 (August 1995), 65-89.

Eichenbaum, Martin, "Some Empirical Evidence on the Production Level and Production Cost Smoothing Models of Inventory Investment," *American Economic Review* 79 (Sept. 1989), 853-64.

Fair, Ray C., "The Production-Smoothing Model is Alive and Well," *Journal of Monetary Economics* 24 (Nov. 1989), 353-70.

Fay, John A., and James L. Medoff, "Labor and Output over the Business Cycle: Some Direct Evidence," *American Economic Review* 75 (September 1985), 638-55.

Farmer, Roger, and Jang Ting Guo, "Real Business Cycles and the Animal Spirits Hypothesis," *Journal of Economic Theory* (1994).

Hall, Robert E., "Employment Fluctuations and Wage Rigidity," *Brookings Papers on Economic Activity* (1980), 93-123.

Hall, Robert, E., "Labor Demand, Labor Supply, and Employment Volatility," *NBER*

Macroannual (1991): 17-54.

Howell, C., F. Congelio, and R. Yatsko, "Pricing Practices for Tobacco Products, 1980-1994," *Monthly Labor Review* (December 1994), 3-16.

Gertler, Mark, and Simon Gilchrist, "Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms," Board of Governors of the Federal Reserve System Discussion Paper 93-4, 1993.

Kahn, James A., "Inventories and the Volatility of Production." *American Economic Review* 77 (1987), 667-79.

Kahn, James A., "Why is Production more Volatile than Sales? Theory and Evidence on the Stockout-Avoidance Motive for Inventory Holding." *Quarterly Journal of Economics* 107 (1992).

Kashyap, Anil, Owen Lamont, and Jeremy Stein, "Credit Conditions and the Cyclical Behavior of Inventories: A Case Study of the 1981-82 Recession," NBER Working Paper 4211. 1992.

Kashyap. Anil. and David W. Wilcox. "Production and Inventory Control at the General Motors Corporation in the 1920's and 1930's," *American Economic Review* 83 (June 1993). 383-401.

Kydland. Fynn E. and Edward C. Prescott. "Time to Build and Aggregate Fluctuations." *Econometrica* 50 (1982), 1345-1370.

Pindyck, Robert S.. "Inventories and the Short-Run Dynamics of Commodity Prices," *The Rand Journal of Economics* 25 (1994), 141-159.

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

Ramey, Valerie. "Non-Convex Costs and the Behavior of Inventories." *Journal of Political Economy* 99 (1991).

Solow, Robert M., "Some Evidence on the Short-Run Productivity Puzzle," in *Development and Planning*, edited by J. Baghwati and R.S. Eckaus. New York: St. Martin's, 1968.

Thurlow, Peter H., *Essays on the Macroeconomic Implications of Inventory Behavior.* Ph. D. thesis, University of Toronto, 1993.

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

Table 1--Correlations Between $\frac{s_t}{a_t}$, $\frac{y_{t+1}}{x_{t+1}}$, and detrended log Output, $\tilde{y}_t$.

	Correlation $\frac{s_t}{a_t}$ and $\tilde{y}_t$	Correlation $\frac{y_{t+1}}{x_{t+1}}$ and $\tilde{y}_t$
Tobacco	.49*	−.03
Apparel	.37*	.19*
Lumber	.46*	.50*
Chemicals	.65*	.49*
Petroleum	.68*	−.15*
Rubber	.34*	.48*

*Correlation is significantly different from zero at .01 level.

Table 2—Correlations Between the *Required* Expected Growth in Marginal Cost and detrended log Output

	Correlation Between Required $E(\frac{c_{t+1}}{c_t})$ and $\tilde{y}_t$			
	Based on $\frac{s_t}{a_t}$		Based on $\frac{y_{t+1}}{x_{t+1}}$	
	For constant r	For varying r	For constant r	For varying r
Tobacco	−.49*	−.36*	−.03	−.11
Apparel	−.52*	−.49*	−.21*	−.43*
Lumber	−.51*	−.70*	−.58*	−.77*
Chemicals	−.70*	−.79*	−.56*	−.80*
Petroleum	−.75*	−.33*	.13	−.02
Rubber	−.35*	−.84*	−.52*	−.83*

*Correlation is significantly different from zero at .01 level.

Table 3—Correlations Between Measures of Expected Growth in Marginal Cost and detrended log Output

Correlation Between $E(\frac{c_{t+1}}{c_t})$ and $\tilde{y}_t$

	Wage Measure	
	Average Hourly Earnings	Effort Wage
Tobacco	.75*	−.88*
Apparel	.73*	−.07
Lumber	.61*	−.40*
Chemicals	.35*	−.75*
Petroleum	.41*	.40*
Rubber	.34*	−.52*

*Correlation is significantly different from zero at .01 level.

Table 4--GMM Estimates of Combined Inventory Investment Equation

	$\frac{\phi/\rho}{\phi^*/\rho^*}$	χ^2 Statistic
Tobacco	2.908 (0.138)	86.9*
Apparel	1.099 (0.107)	135.7*
Lumber	1.519 (0.116)	122.3*
Chemicals	2.591 (0.100)	120.8*
Petroleum	2.421 (0.201)	129.8*
Rubber	2.694 (0.119)	137.0*

---

*Significant model rejection at .01 level. Standard errors are in parentheses.

Táble 5——Horse Race Results between AHE and Effort Wage Cost Measures for Finished Inventory

	γ_1	γ_2	$\frac{\hat{a}}{\bar{a}}$	χ^2 statistics
Tobacco	0.212 (0.047)	0.306 (0.058)		37.4*
	0.209 (0.053)	0.308 (0.060)	−0.067 (0.521)	37.3*
Apparel	0.040 (0.032)	0.653 (0.103)		86.2*
	−0.044 (0.036)	0.304 (0.117)	−0.864 (0.328)	57.1*
Lumber	−0.130 (0.037)	0.195 (0.086)		72.3*
	−0.110 (0.037)	0.097 (0.088)	0.274 (0.064)	72.7*
Chemicals	−0.021 (0.077)	1.294 (0.171)		64.5*
	0.087 (0.085)	1.514 (0.188)	0.193 (0.042)	52.1*
Petroleum	−0.016 (1.083)	−0.207 (1.103)		36.4*
	0.071 (1.061)	−0.293 (1.081)	−0.164 (0.344)	37.3*
Rubber	−0.094 (0.050)	0.887 (0.127)		43.8*
	−0.075 (0.053)	0.963 (0.132)	0.056 (0.040)	40.2*

*Significant model rejection at .01 level. Standard errors are in parentheses.

Table 6--Horse Race Results between AHE and Effort Wage Cost Measures for WIP

	γ_1	γ_2	$\frac{\hat{x}}{\bar{x}}$	χ^2 statistics
Tobacco	0.459 (0.100)	0.693 (0.122)		47.4*
	0.626 (0.188)	1.055 (0.212)	0.5†	33.6*
Apparel	0.066 (0.027)	0.798 (0.113)		77.5*
	0.005 (0.030)	0.197 (0.122)	−0.794 (0.218)	52.0*
Lumber	−0.129 (0.041)	0.265 (0.104)		52.9*
	−0.127 (0.041)	0.247 (0.107)	0.081 (0.101)	52.5*
Chemicals	−0.104 (0.085)	1.219 (0.100)		78.0*
	0.051 (0.092)	1.633 (0.211)	0.326 (0.030)	49.7*
Petroleum	−1.869 (1.557)	1.622 (1.585)		60.1*
	6.741 (2.660)	−7.131 (2.703)	0.5†	84.9*
Rubber	−0.078 (0.055)	0.785 (0.146)		74.9*
	−0.039 (0.058)	0.938 (0.153)	0.129 (0.036)	68.4*

*Significant model rejection at .01 level. Standard errors are in parentheses.

†Set at 0.5. Estimation failed to converge.

Table 7—–GMM Estimates of Finished Inventory Investment Equation, Effort Wage, with Constant Markup

	γ	$\frac{\hat{a}}{\bar{a}}$	χ^2 statistic
Tobacco	0.129 (0.031)		79.6*
	0.157 (0.043)	−1.001 (1.374)	74.8*
Apparel	0.622 (0.103)		84.0*
	0.328 (0.115)	−0.811 (0.287)	63.6*
Lumber	0.225 (0.073)		83.5*
	0.120 (0.077)	0.249 (0.057)	84.3*
Chemicals	1.311 (0.167)		64.9*
	1.458 (0.182)	0.186 (0.041)	51.2*
Petroleum	−0.222 (0.036)		36.7*
	−0.221 (0.036)	−0.165 (0.342)	37.6*
Rubber	0.962 (0.111)		47.4*
	1.040 (0.119)	0.080 (0.036)	42.0*

*Significant model rejection at .01 level. Standard errors are in parentheses.

Table 8--GMM Estimates of Finished Inventory Investment Equation, Effort Wage, with Variable Markup

	γ	$\frac{\hat{a}}{a}$	χ^2 statistic
Tobacco	−0.402 (0.063)		74.0*
	−0.467 (0.068)	0.5†	67.9*
Apparel	3.798 (0.602)		96.3*
	4.520 (0.768)	0.5†	86.8*
Lumber	1.047 (0.220)		146.2*
	1.163 (0.234)	0.5†	152.9*
Chemicals	4.637 (0.659)		121.3*
	5.757 (0.937)	0.5†	108..2*
Petroleum	−0.954 (0.234)		58.4*
	−0.832 (0.036)	0.5†	60.7*
Rubber	2.975 (0.330)		104.0*
	6.278 (0.911)	0.5†	75.5*

---

*Significant model rejection at .01 level. Standard errors are in parentheses.

†Set at 0.5. Estimation failed to converge.

Table 9--GMM Estimates of Work In Process Investment Equation, Effort Wage

	γ	$\frac{\hat{x}}{\bar{x}}$	χ^2 statistic
Tobacco	0.615 (0.096)		94.6*
	0.759 (0.146)	0.5†	75.9*
Apparel	0.090 (0.115)		38.7*
	0.053 (0.124)	−0.56 (0.069)	37.5*
Lumber	0.198 (0.084)		66.6*
	0.184 (0.086)	0.051 (0.077)	66.1*
Chemicals	1.451 (0.185)		91.4*
	1.710 (0.219)	0.383 (0.024)	50.0*
Petroleum	−0.245 (0.053)		122.3*
	−0.308 (0.062)	0.5†	123.0*
Rubber	0.830 (0.144)		105.8*
	0.840 (0.152)	0.196 (0.028)	98.7*

---

*Significant model rejection at .01 level. Standard errors are in parentheses.

†Set at 0.5. Estimation failed to converge.

Figure 1: The Cyclical Behavior of the Inventory-Sales Ratio

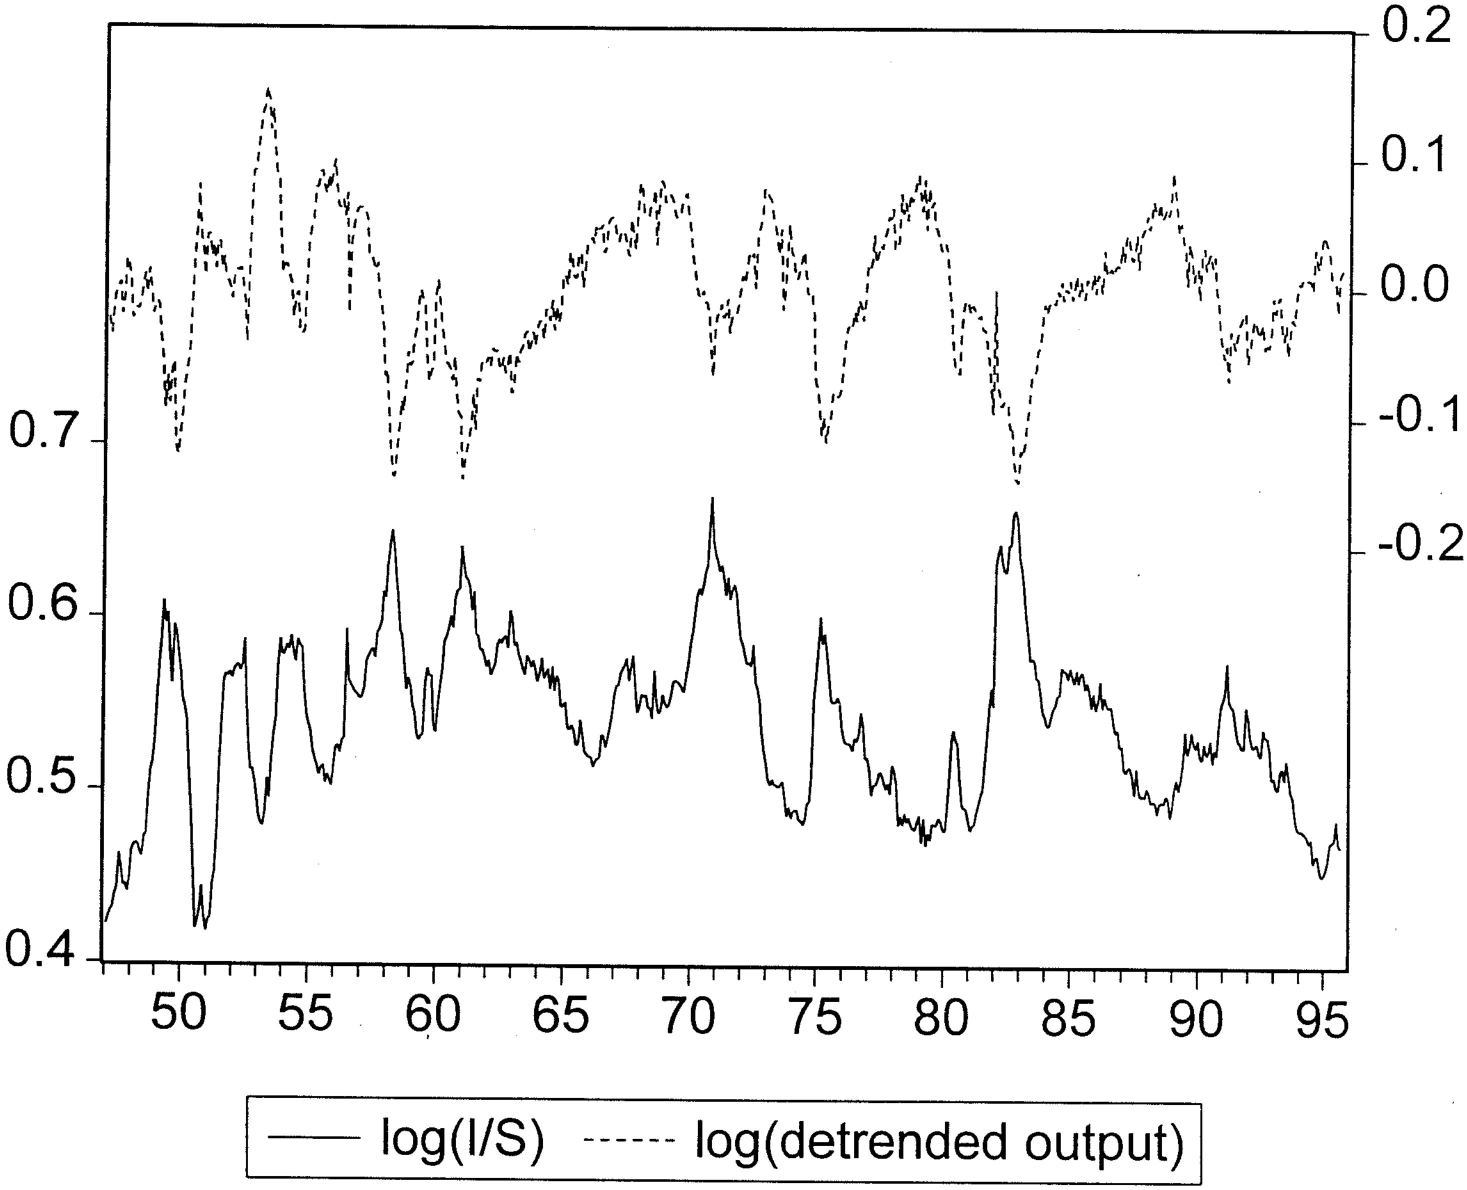

Figure 2: S/A and Price in the Tobacco Industry

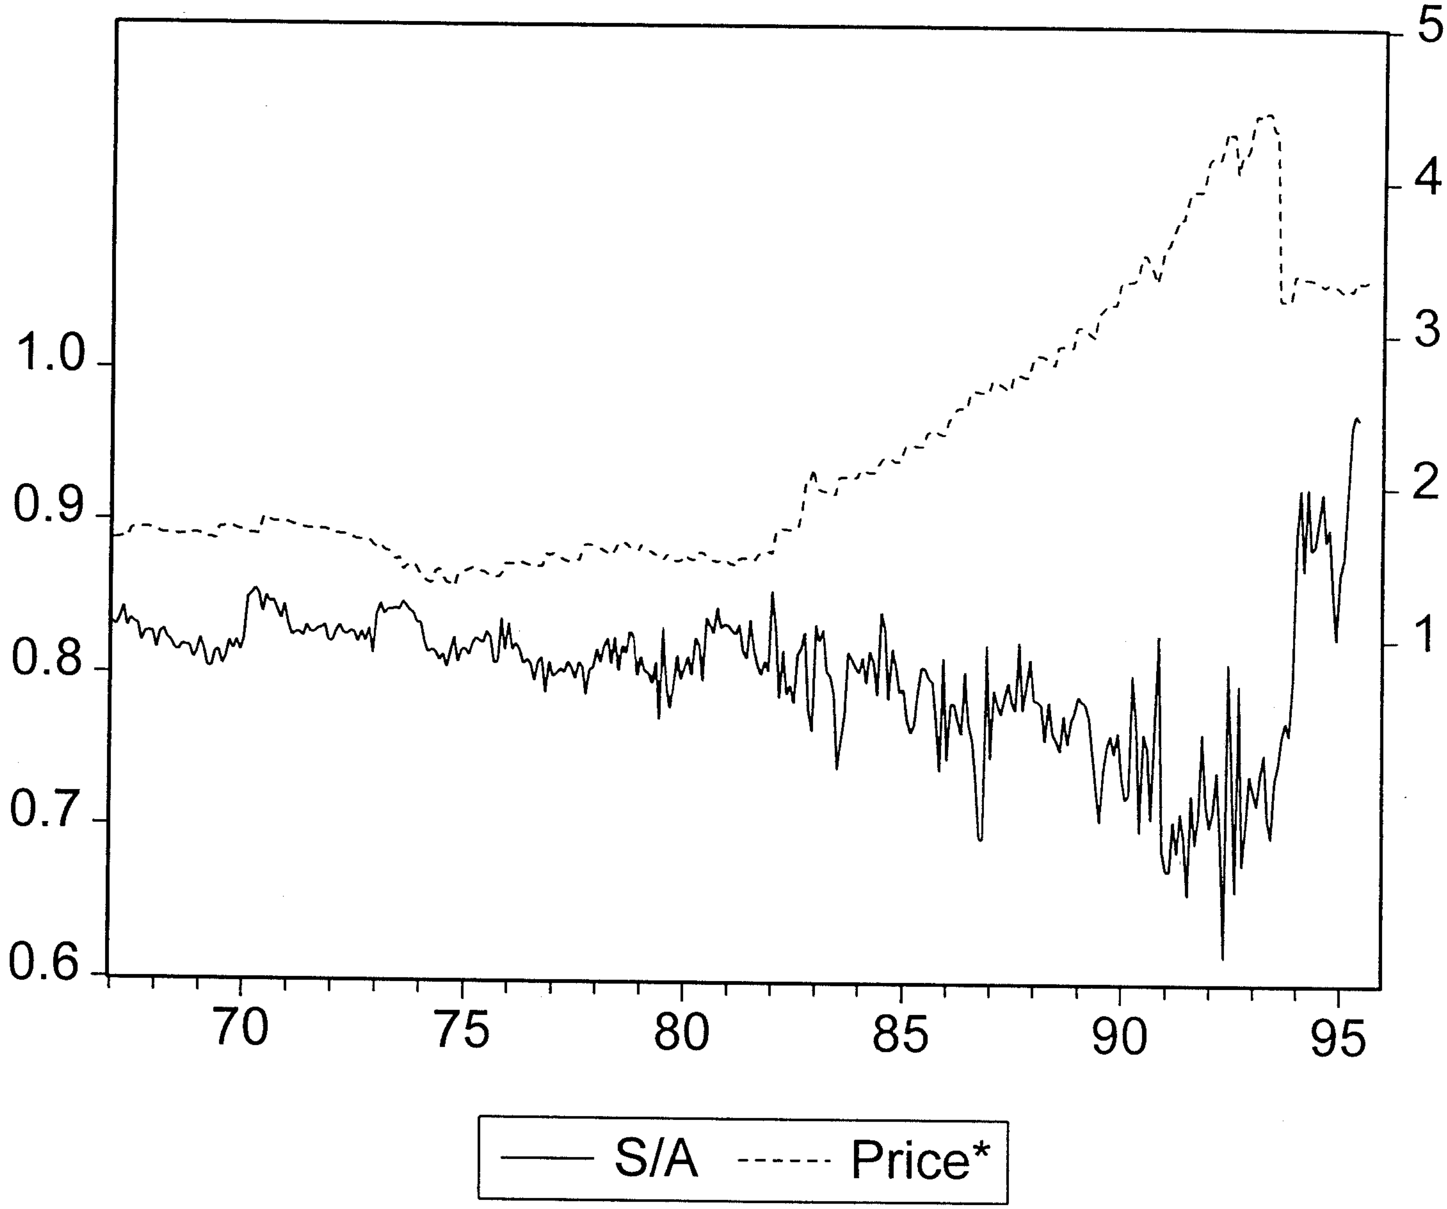

*Deflated by the Producer Price Index

Figure 3: The Cyclical Behavior of Flow-Stock Ratios

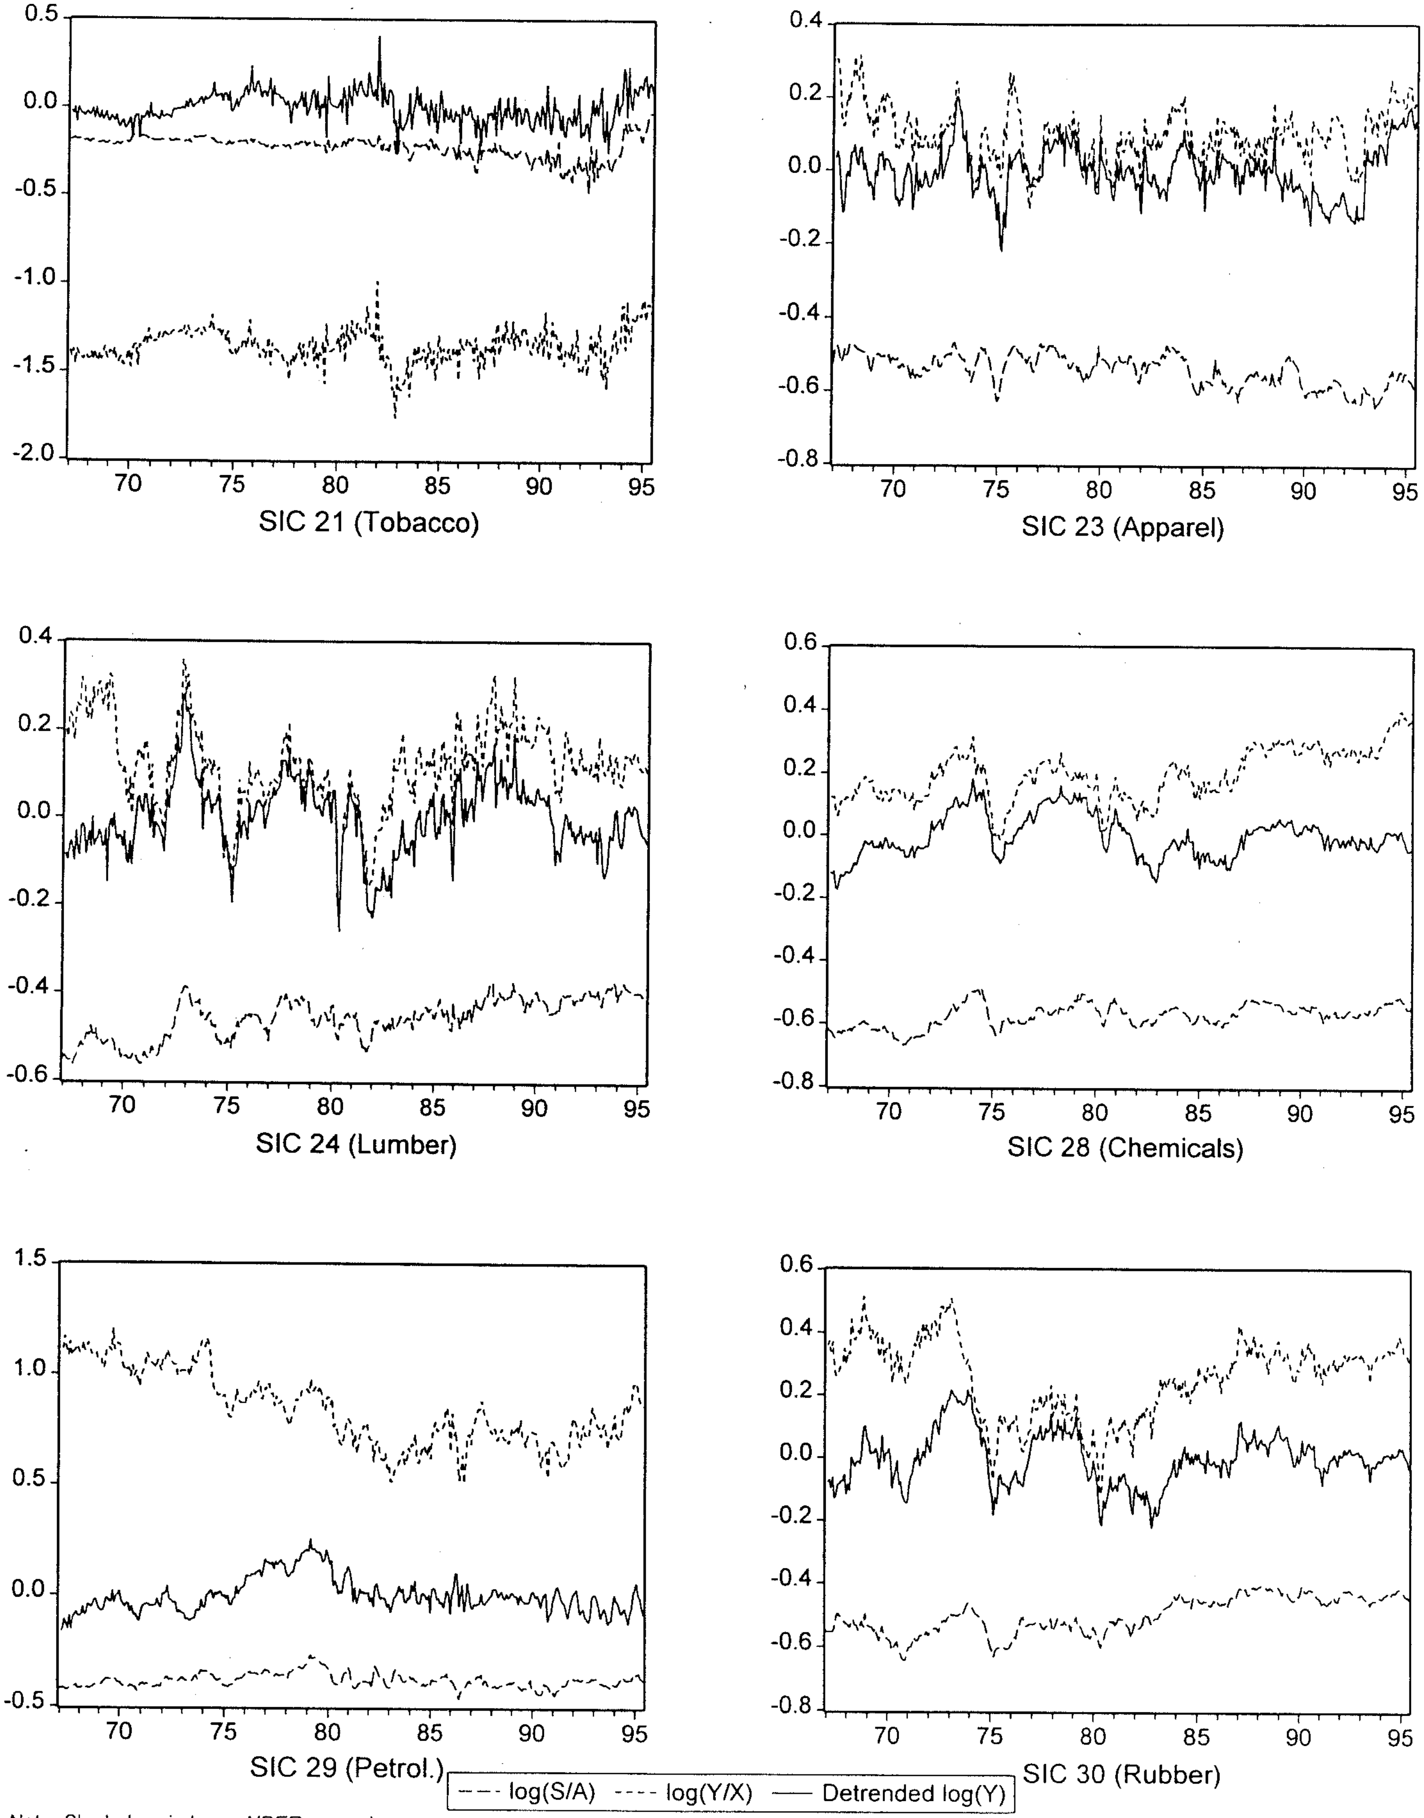

Note: Shaded periods are NBER recessions.

Figure 4: Implied Expected Marginal Cost Growth (S/A Basis)

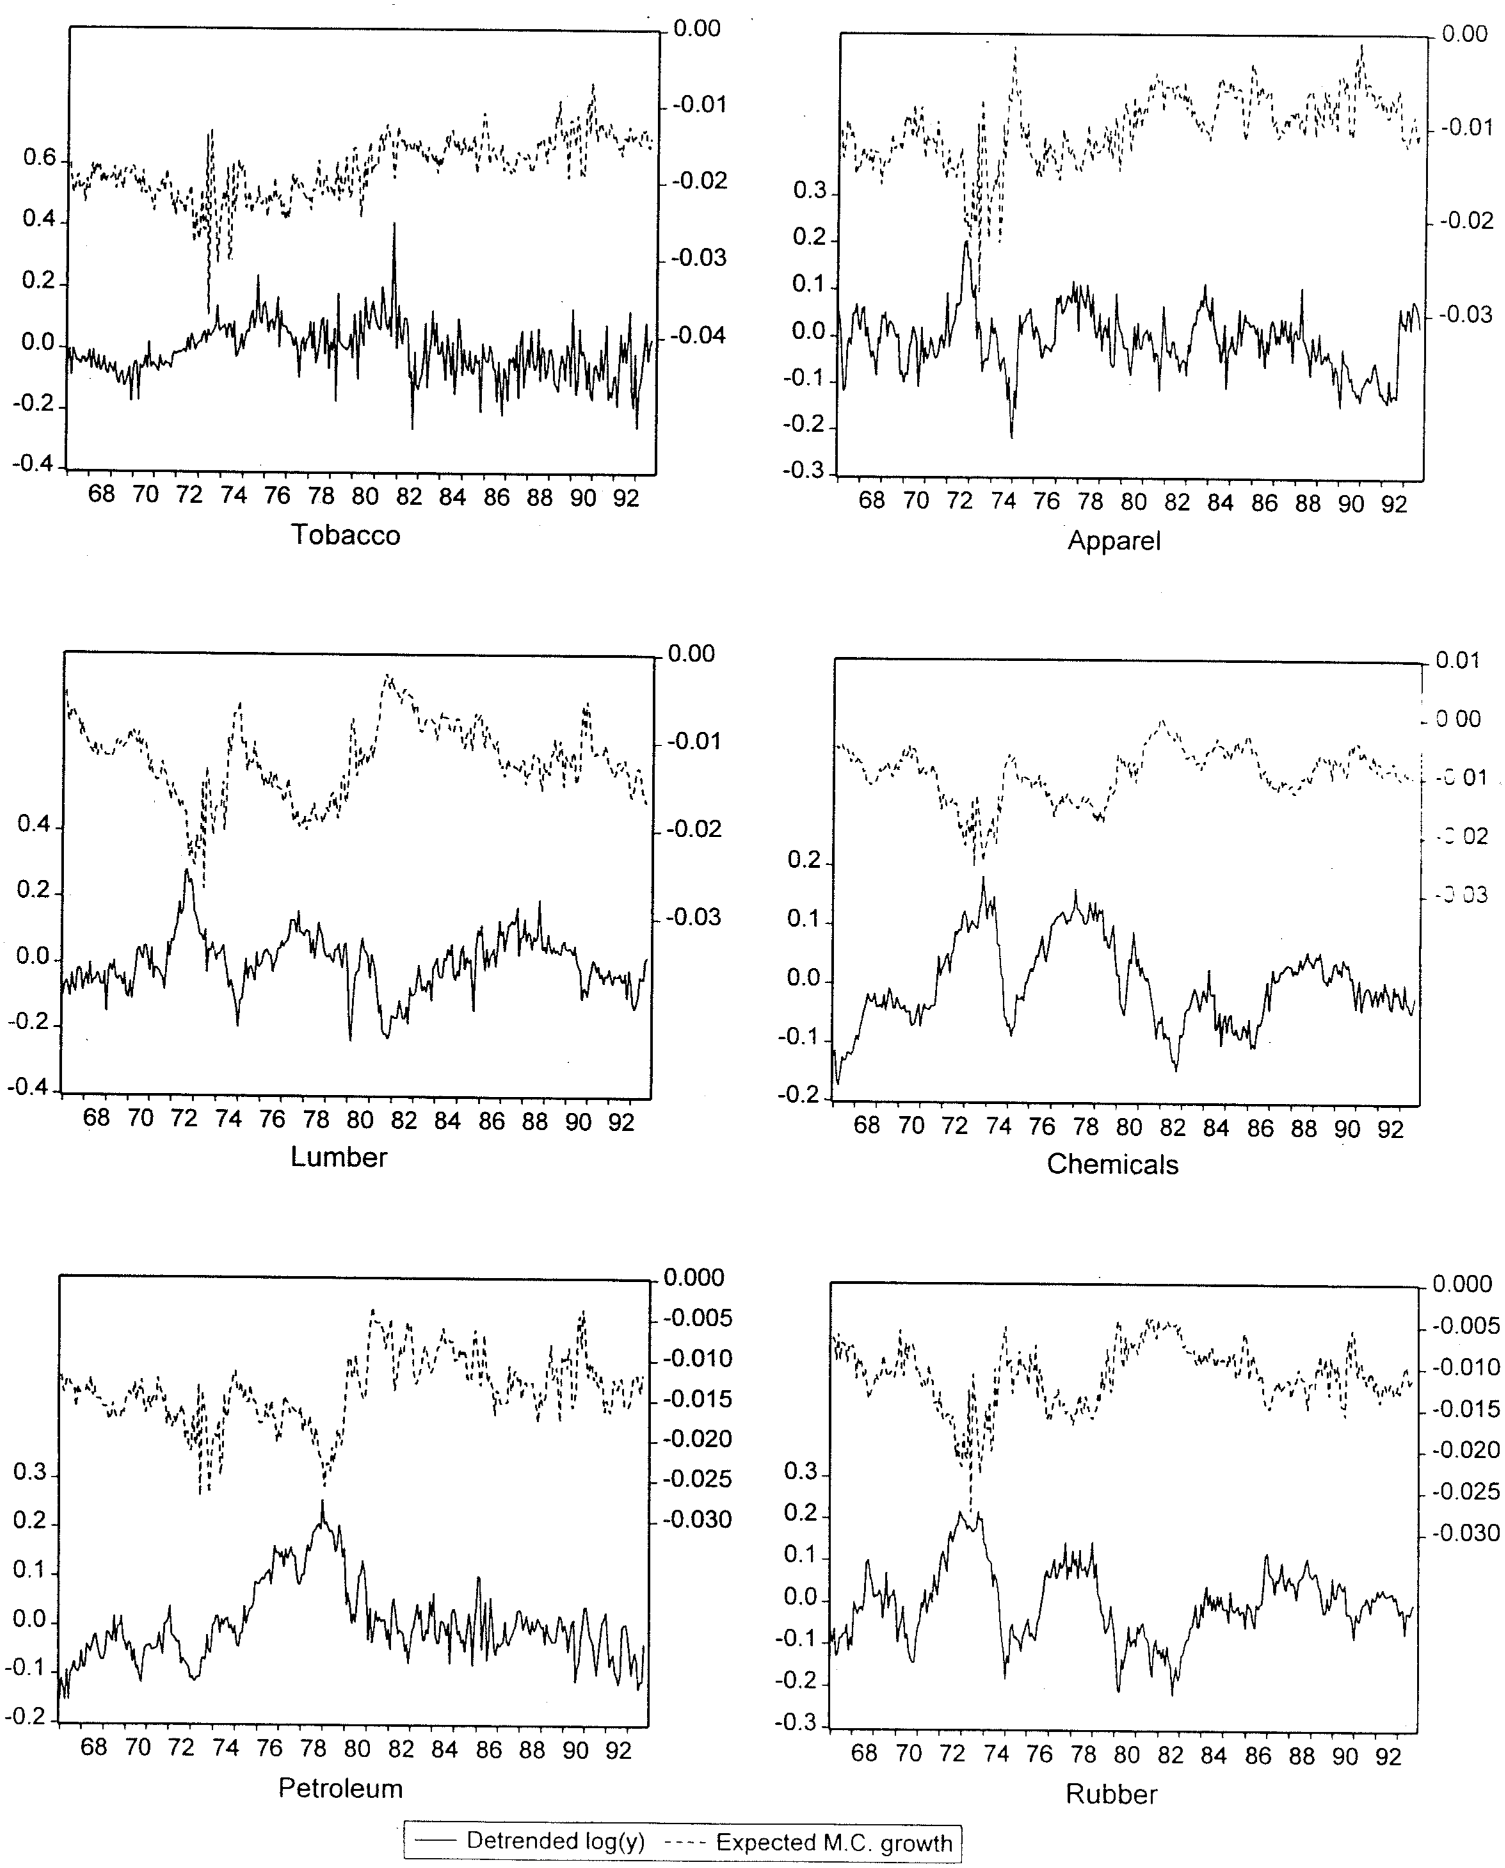

Figure 5: Implied Expected Marginal Cost Growth (Y/X Basis)

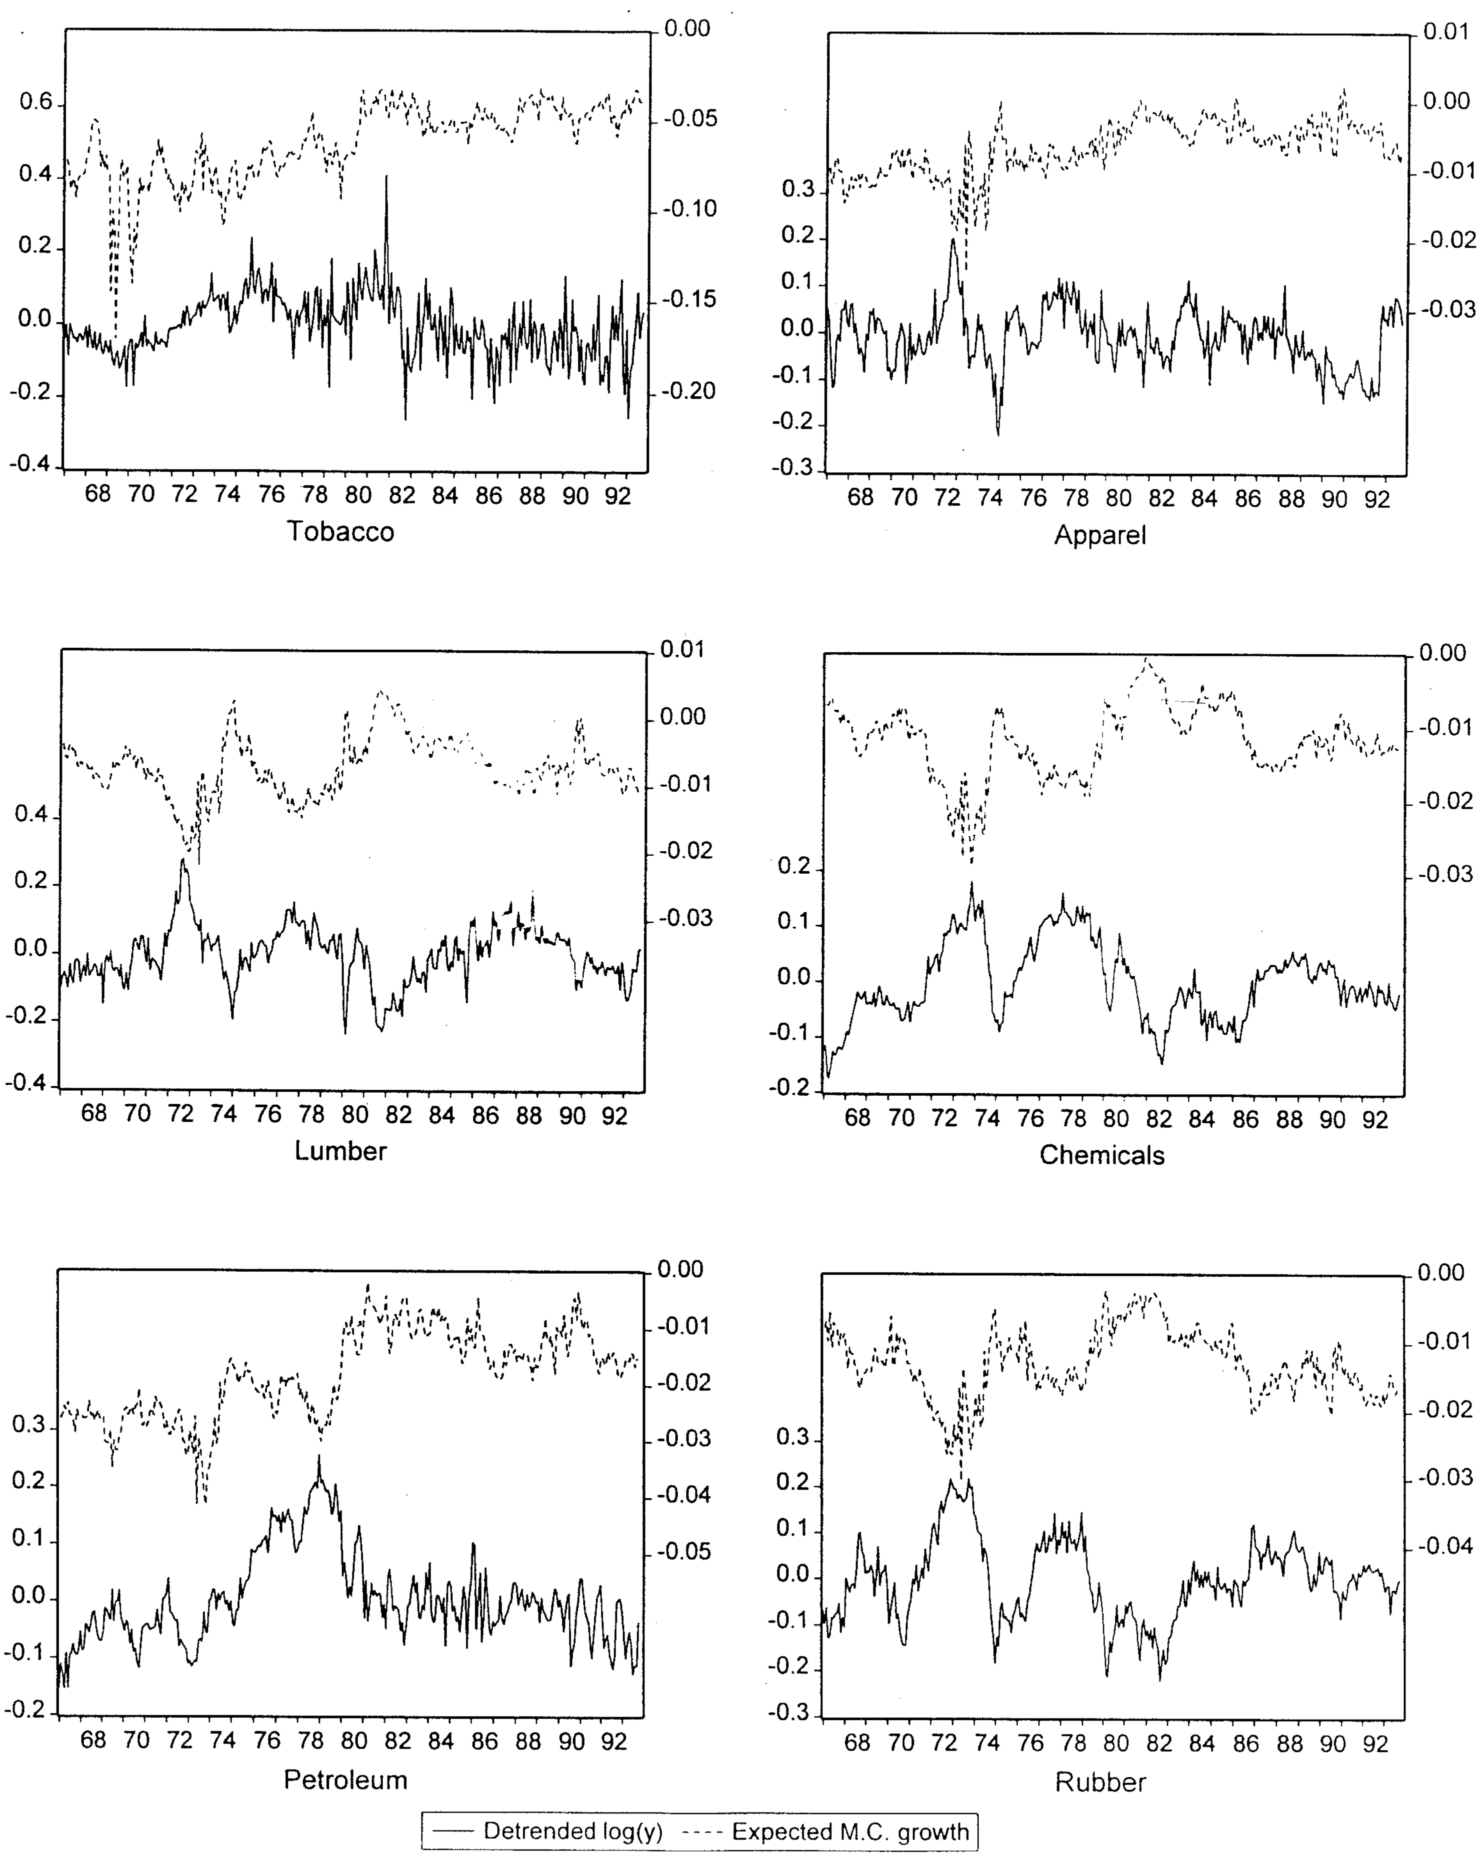

Figure 6: The Cyclical Behavior of Marginal Cost (Avg. Hourly Earnings Basis)

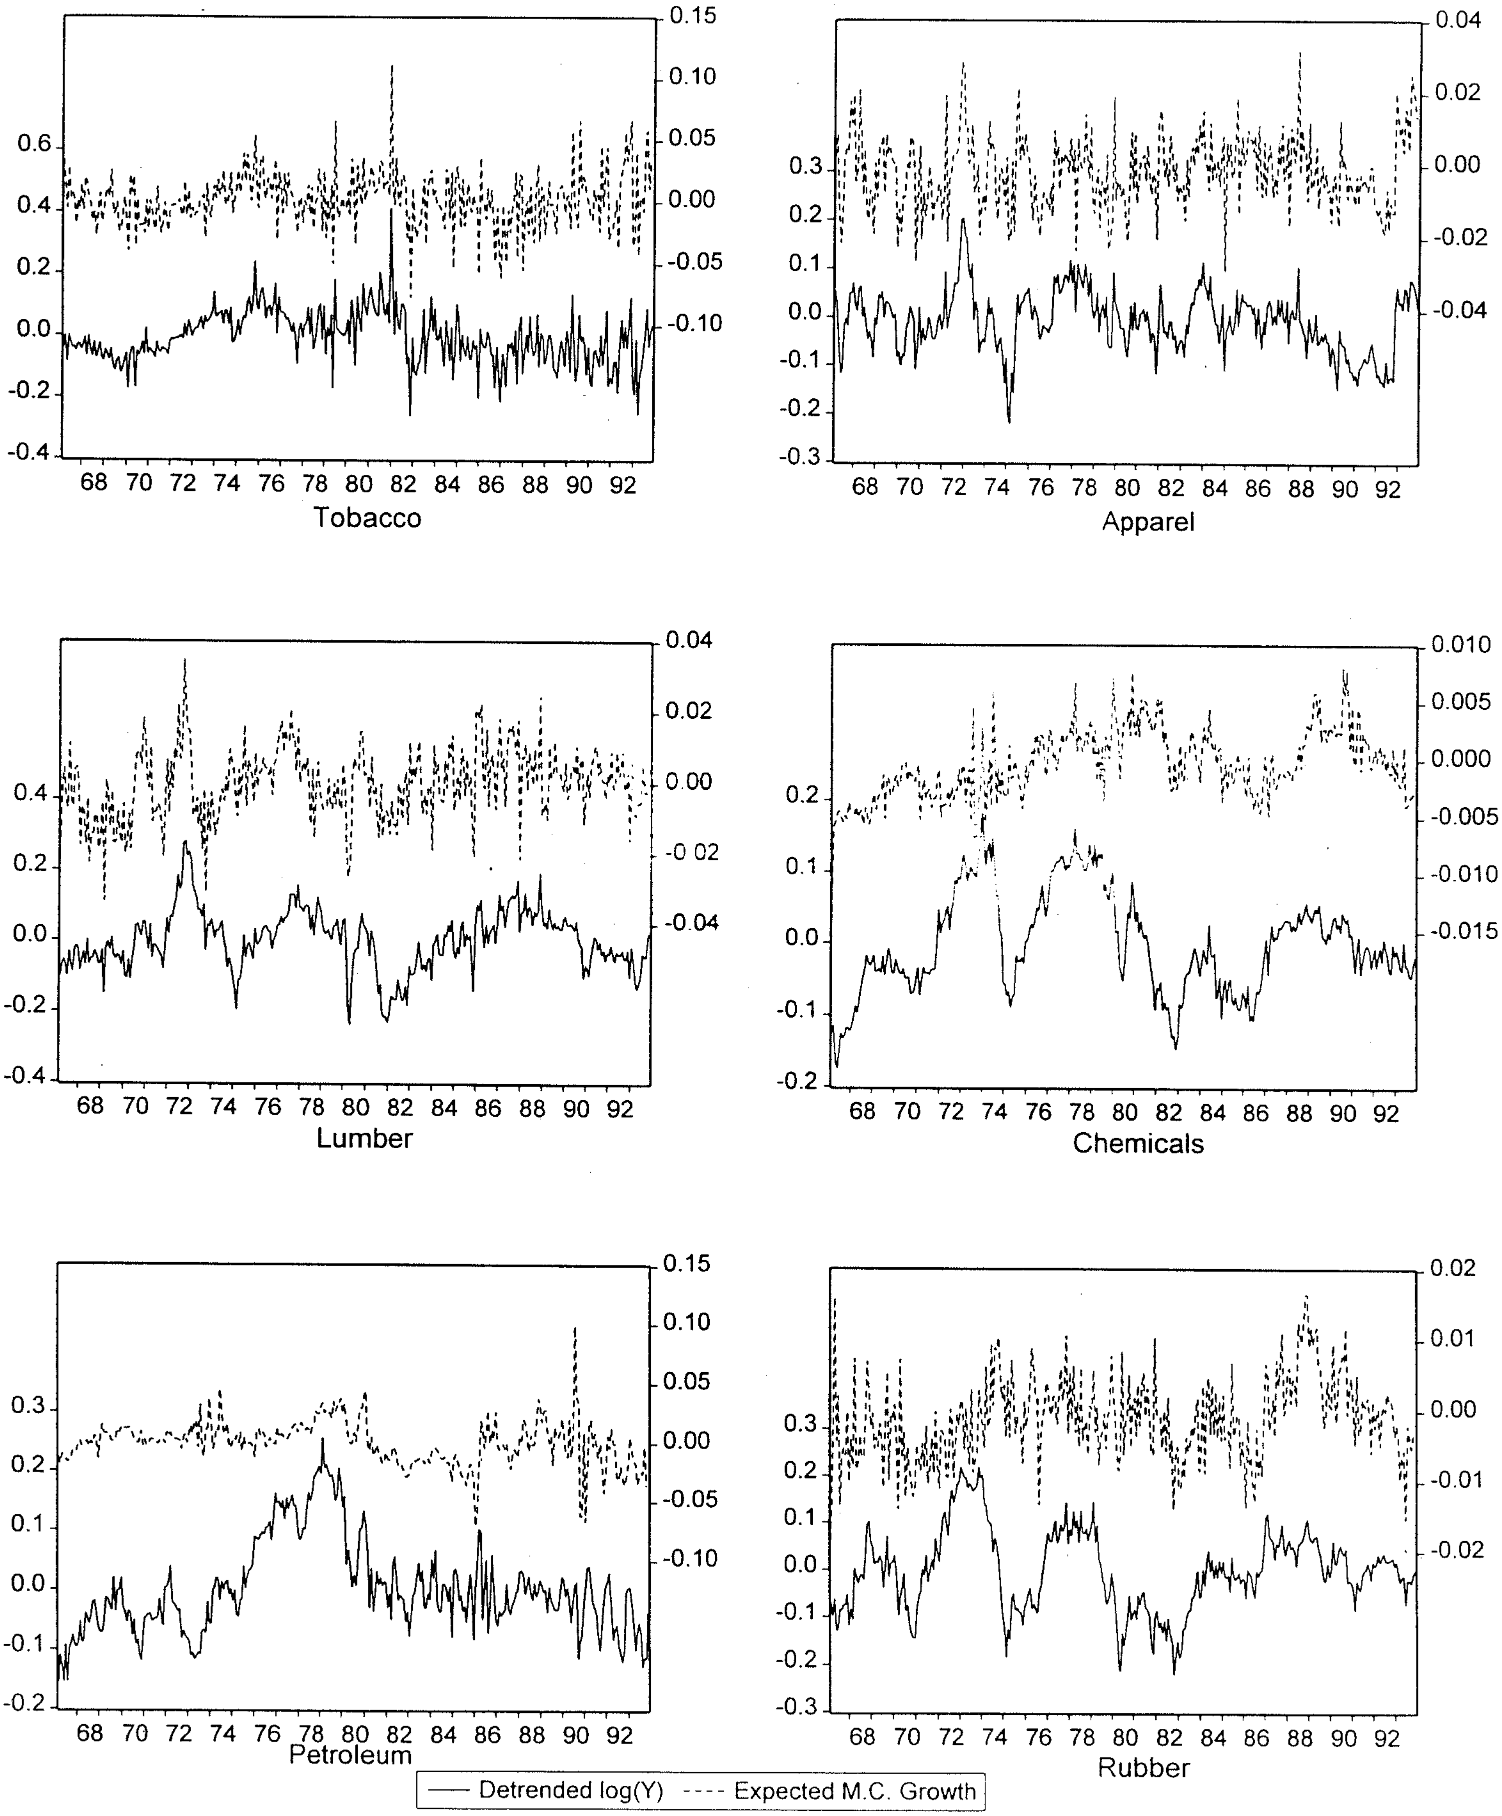

Figure 7: The Cyclical Behavior of Marginal Cost (Effort Wage Basis)