WHAT INVENTORY BEHAVIOR TELLS US ABOUT BUSINESS CYCLES

Mark Bils and James A. Kahn

Federal Reserve Bank of New York
Research Paper No. 9817

July 1998

What Inventory Behavior Tells Us about Business Cycles

Mark Bils

University of Rochester and NBER

and

James A. Kahn

Federal Reserve Bank of New York

May 1998*

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. Finished goods inventories move proportionally much less than sales or production over the business cycle, which we show implies procyclical marginal cost and countercyclical price markups. Obvious measures for marginal cost do not show high marginal cost near peaks, as required to rationalize the inventory behavior, because measured factor productivity rises during the peak phase of the cycle. We can better explain the cyclical behavior of inventory holdings by allowing for procyclical factor utilization, the cost of which is internalized by firms but is not contemporaneously reflected in measured wage rates.

* Support from the National Science Foundation while both authors were at the University of Rochester is gratefully acknowledged. We also thank Anil Kashyap, Peter Klenow, Valerie Ramey, Julio Rotemberg, Ken West, Michael Woodford, two referees, and participants at a number of seminars for helpful comments. The views expressed are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of New York or of the Federal Reserve System.

I. Introduction

There are two distinct views on the nature of business cycle fluctuations. In one view business cycle peaks represent times of increased productivity and lower production costs. This includes models based on increasing returns such as Farmer and Guo (1994), as well as real business cycle models (e.g., Kydland and Prescott, 1982) driven by fluctuations in technology. According to the second view, at cyclical peaks capacity constraints and diminishing returns kick in, driving up the costs of production relative to input prices and relative to other periods. Many researchers have viewed the procyclical behavior of inventory investment as evidence for the first view because it suggests that firms bunch production more than is necessary to match the fluctuations in sales. If short-run marginal cost curves were fixed and upward sloping (the argument goes), firms would smooth production relative to sales, making inventory investment countercyclical.[1] Countercyclical marginal cost in turn is viewed as evidence for procyclical technology shocks, increasing returns, or positive externalities.[2]

We argue that this reasoning is false: Inventory investment should be procyclical even with increasing marginal cost. The argument outlined above overlooks changes in the shadow value of inventories--which we argue increases with the level of production and expected sales.[3] We propose a model in which finished inventories facilitate sales,

[1]See West (1985), Blinder (1986), and Fair (1989) for evidence on production volatility and the cyclical behavior of inventory investment.

[2]West (1991) explicitly uses inventory behavior to decompose the sources of cyclical fluctuations into cost and demand shocks. 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. See also Hall (1991). 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.

with the marginal value of inventories proportional to the level of sales. This implies that (holding price and marginal cost fixed) inventories should rise proportionately with anticipated increases in sales.

Procyclical inventory investment implies that inventory stocks are also procyclical, but does not imply that stocks rise as much as sales. In fact, inventory-sales ratios are extremely countercyclical. Figure 1 plots the monthly ratios of finished goods inventory to sales (shipments) in aggregate manufacturing for 1959 through 1997 along with output, where output is detrended according to a Hodrick-Prescott (H-P) filter. The inventory-sales ratio increases dramatically in each recession, typically by 5 to 10 percent. Note that these increases do not simply reflect a transitory response to an unexpected fall in sales, but are highly persistent for the duration of each recession. Replacing sales with forecasted sales generates a very similar picture. The correlation between the two series in the figure is 0.675. In the empirical work below we examine data for six two-digit manufacturing industries that produce primarily to stock. These data reinforce the picture from aggregate data in Figure 1−−inventories fail to keep up with sales over the business cycle.[4]

We find the puzzle then to be why inventory investment is not *more* procyclical. Inventories sell with predictably higher probability at peaks, suggesting that firms should add more inventories in booms so as to equate the ratios (and hence the "returns") over

[3]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. West (1991) demonstrates that the estimated importance of cost versus demand shocks in output fluctuations is very dependent on the size of the target inventory-sales ratio.

[4]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.

time. Our model shows that this striking fact implies that in booms either 1) marginal cost must be high relative to *discounted* future marginal cost, or 2) the markup of price over marginal cost must be low.

We initially examine the case of a constant markup. This allows us to measure expected movements in marginal cost by expected movements in price. For sales to increase relative to inventories then requires that the rate of price increase be less than the interest rate. This is sharply rejected for the six industries we study--in fact the opposite is true. We turn then to allowing for both cyclical movements in marginal cost and price-cost markup. We ask what behavior of marginal cost is consistent with marginal cost being temporarily high and/or price being low relative to marginal cost in booms, as required to justify the cyclical behavior of inventories.

When we measure marginal cost based on inputs and factor prices, however, we do not find high marginal cost in booms, or countercyclical markups, because input prices are less procyclical than productivity. But we can explain the cyclicality of inventory holdings by allowing for procyclical factor utilization that is not reflected contemporaneously in wage rates.

We find the joint behavior of inventories, prices, and productivity consistent with the following view of business cycles: In a boom the capital stock lags behind labor and output. This raises marginal cost and is associated with cuts in price markups. In turn, this induces 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. Inventory behavior thus points to the second view of business cycles that stresses short-run diminishing returns.

II. The Demand for Inventories

A. *A Firm's Problem*

We examine the production to inventory decision for a representative producer, relying on little more than the following simple elements: Profit maximization, a production function, and an inventory technology that is specified to reflect the fact that inventory-sales ratios appear to be stationary and independent of scale. To achieve the latter, we assume that finished inventories are productive in generating greater sales at a given price (see Kahn, 1987, 1992). Related approaches in the literature include Kydland and Prescott (1982), Christiano (1988), and Ramey (1989), who introduce inventories as a factor of production. Inventory models that incorporate a target inventory-sales ratio, or that recognize stockouts, create a demand for inventories in addition to any value for production smoothing.

A producer maximizes expected present-discounted profits according to

$$(1) \qquad \max_{y_t} \quad E_t \sum_{i=0}^{\infty} \beta_{t,t+i}[\, p_{t+i}s_{t+i} - C_{t+i}(y_{t+i};\lambda_t,\theta_t,z_t)]$$

subject to:

$$\text{i)} \qquad a_t = i_t + y_t = a_{t-1} - s_{t-1} + y_t,$$

$$\text{ii)} \qquad s_t = d_t(p_t)a_t^{\phi},$$

$$\text{iii)} \qquad y_t = \min\,[\frac{q_t}{\lambda_t}, (\theta_t n_t^{\alpha} l_t^{\nu} k_t^{1-\alpha-\nu})^{\gamma}] \; .$$

In the objective function, s_t and p_t denote sales and price in period t, λ_t is a vector of material input requirements, θ_t a technology shock, and z_t is a vector of input prices. $C_t(y_t)$ is the cost of producing output y_t during t. $\beta_{t,t+i}$ denotes the nominal rate of discount at time t for i periods ahead. For example $\beta_{t,t+1}$, which for convenience we write β_{t+1}, equals $\frac{1}{1+R_{t+1}}$ where R_{t+1} is the nominal interest rate between t and t+1.[5] Constraint (i) is just a standard stock-flow identity, taking the stock of goods available for sale during period t, a_t, as consisting of the inventory i_t of unsold goods carried forward from the previous period plus the y_t goods produced in t.

Constraint (ii) depicts the dependence of sales on finished inventories. For a given price, a producer views its sales as increasing with an elasticity of ϕ with respect to its available stock. This approach is consistent, for example, with a competitive market that allows for the possibility of stockouts (e.g., Kahn, 1987, Thurlow, 1995). Alternatively, one can view the stock as an aggregate of similar goods of different sizes, colors, locations, and the like. A larger stock in turn facilitates matching with potential purchasers, who arrive with preferences for a specific type of good. Empirically it is important that a producer's demand increases at least somewhat with its available inventory; otherwise it becomes difficult to rationalize even systematically positive finished goods inventory holdings, 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. All we require is that the impact of the firm's stock a_t be captured by the separate multiplicative term a_t^{ϕ}.

From constraint (iii), output is produced using both a vector of material inputs, q_t, and value added produced by a Cobb-Douglas function of production labor, n_t, nonproduction labor, l_t, and capital, k_t. Material inputs are proportional to output as

[5] In the empirical work we incorporate a storage cost for inventories. We let the cost of storing a unit from period t to t+1 equals δ times the cost of production in t. (This follows, for instance, if storing goods requires the use of capital and labor in proportions similar to producing goods.) The storage cost then effectively lowers β_{t+1} as it now reflects both a rate of storage cost, δ, as well as an interest rate R_{t+1}: $\beta_{t+1} = \frac{1-\delta}{(1 + R_{t+1})}$.

dictated by a vector of per unit material requirements, λ_t. (We do not treat elements of λ_t as choice variables, but λ_t can vary over time.) The value-added production function has returns to scale γ, potentially greater than one.

B. *The First-Order Condition*

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

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

The expectations operator conditions on variables known when choosing period t's output. The producer incurs marginal cost $c_t \equiv C_t'(y_t)$. By increasing the available stock, sales are increased by $\phi d_t(p_t)a_t^{\phi-1}$. 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, $\phi d_t(p_t)a_t^{\phi-1}$, is proportional to the ratio of sales to stock available, equalling $\frac{\phi s_t}{a_t}$. Making this substitution and rearranging gives

$$(2) \qquad E_t \{ [\frac{\phi s_t}{a_t} m_t + 1]v_{t+1} \} = 1 ,$$

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

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

v_{t+1} equals the discounted gross rate of growth in marginal cost. In a pure production smoothing model, with 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 $\beta_{t+1}c_{t+1}$ is the opportunity cost of selling a unit during t.

Consider a scenario in which the growth in discounted costs, v_{t+1}, and the price-cost markup are both constant through time. This implies that $\frac{E_t(s_t)}{a_t}$ is constant, i.e. all predictable movements in sales are matched by proportional movements in the stock available.[6] To generate persistent procyclical movements in the ratio of sales to inventory such as we see in the data requires either a countercyclical markup or that marginal cost is transitory high, relative to the expected discounted value of next periods marginal cost. Thus the behavior of inventories actually points *against* increasing returns or countercyclical costs. It also points against a role for credit market imperfections in accounting for the cyclical behavior of inventories. To account for the data, credit constraints would need to bind in expansions, thereby driving up current marginal cost relative to discounted future marginal cost. This is opposite the scenario emphasized by Gertler and Gilchrist (1993), Kashyap, Lamont, and Stein (1994), and others.[7]

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, because sales (conditional on price) are a power function of the available stock. We can examine the behavior of $\frac{s_t}{a_t}$ to see whether this functional form assumption is reasonable, as it implies that the steady-state sales to inventory ratio should be independent of the size of the industry or firm.

Some evidence can be gleaned from observing how the ratio $\frac{s_t}{a_t}$ changes over time in industries with substantial growth. Below we examine in detail the six manufacturing industries tobacco, apparel, lumber, chemicals, petroleum, and rubber. For all but

[6] In a steady state with a constant rate of growth in marginal cost the ratio $\frac{a}{s}$ equals $\frac{\phi m(1-\delta)}{r+\delta}$; r here is a real interest rate equalling R minus the inflation rate in marginal cost and δ is the rate of storage cost.

[7] Credit crunches during downturns could still account for the differences in behavior between firms labelled constrained and unconstrained by these authors.

tobacco, sales increased by 50 percent or more from 1959 to 1997. Figure 2 presents the behavior of $\frac{s_t}{a_t}$ for each industry for that period. None of the six industries display large long-run movements in the ratio, even when the level of s_t changes considerably. The largest trend movements are for apparel, where the ratio declines by about 25 percent, and in rubber, where it rises by about 20 percent.

The model's implication that stock available is proportional to expected sales is also supported by cross-sectional evidence. Kahn (1992) reports average inventory-sales ratios and sales across divisions of U.S. automobile firms. These data show no tendency for the ratio to be related to the size of the division, either within or across firms. Gertler and Gilchrist (1993) present inventory-sales ratios for manufacturing by firm size, with size defined by firm assets. Their data similarly show 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.[8]

C. *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 is[9]

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

[8] In a previous version (available as Rochester Center for Economic Research Working Paper #428, September 1996) we allow for the more general functional form s_t equal to $d_t(p_t)[a_t - \hat{a}]^{\phi}$, implying s_t increases with a_t only after the available stock reaches a threshold value $\hat{a}$. This generates a scale effect in inventory holdings, providing another possible explanation for the failure of inventories to keep pace with sales over the business cycle. Our estimates for the threshold term $\hat{a}$ are typically less that 20 percent of the average size of a_t; and its introduction did not significantly affect other estimated results.

[9] A number of papers include a cost term in the change in output. Its exclusion here is simply for convenience. Measures for cost shocks, such as wage changes, are also sometimes included (e.g., Ramey, 1991, or Durlauf and Maccini, 1995).

where, as before, y_t, a_t, and s_t are output, stock available, and sales during t. The slope of marginal cost is governed by the parameter ψ. Note that $\alpha > 0$ allows for a target inventory-sales ratio. Many researchers (e.g., Blinder, 1985, Fair, 1989) have focused on the relative volatility of production and sales or, relatedly, on explaining why inventory investment is procyclical. But if $\alpha > 0$, having $\psi > 0$ does not imply that the variance of sales exceeds the variance of production or that inventory investment is countercyclical (West, 1986). On the other hand, under this linear-quadratic model, it can only be optimal for a firm to systematically have a low $\frac{a_t}{s_t}$ ratio when sales are high if its marginal cost is relatively high in those periods. Otherwise costs could be reduced by bunching production in high-sales periods, thereby generating a procyclical ratio.[10] Given inventories are so countercyclical relative to sales, we therefore believe the question should be: Why is inventory investment not *more* procyclical?

Our approach differs substantially from the linear-quadratic literature in at least two ways. First, rather than estimate parameters of a quadratic cost function, we exploit the production function to measure marginal cost directly in terms of observables and parameters of the underlying production technology. This measure allows not only for variation in wages, the cost of capital, and other inputs, but also potentially for shocks to productivity. Second, our model explicitly considers the firm's revenue side. This allows us to account for variations in inventory holdings caused by price variations. In our

[10]For example, in the absence of cyclical cost shocks, one can prove by a variance bounds argument similar to that of West (1986) that if a_t/s_t is countercyclical then ψ must be positive. So how do researchers (in particular, Ramey, 1991) who include an inventory-sales target find downward-sloping marginal cost with data exhibiting countercyclical stock-sales ratios? One possibility is that the linear-quadratic specification poorly approximates the true model, as suggested by results in West (1986) and Pindyck's (1994). Also, the findings of downward sloping marginal cost may be sensitive to how parameters are normalized (see Krane and Braun's, 1991, discussion), choice of instruments, and whether cost shocks are allowed. Furthermore, our reading of the literature is that most authors do 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).

model the return on finished inventory is proportional to the price markup; so sales relative to stock available should move inversely with the markup.

The tobacco industry provides an excellent experiment. The price of tobacco products rose very dramatically from 1984 to 1993. Figure 3 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 largely reflected a rise in price markup (Howell et al., 1994). Consistent with the model, the ratio $\frac{s_t}{a_t}$ fell by about 15 percent. More striking is what occured in 1993. During one month, August 1993, the price of tobacco products fell by 25 percent, apparently reflecting a breakdown in collusion (see Figure 3). Within 3 months the ratio $\frac{s_t}{a_t}$ rose dramatically, as predicted by the model, by at least 25 percent. Whereas the linear-quadratic model is silent on these large movements in inventory-sales ratios, the model in this paper contains a ready explanation.

III. Empirical Implementation

A. *The Case of a Constant Markup*

Inventory investment is closely related to variations in marginal cost. A transitory decrease in marginal cost motivates firms to produce now, accumulating inventory. A higher markup of price over marginal cost also motivates firms to accumulate inventory. For this reason, much of the empirical work is directed at the behavior of marginal cost. But first we consider the case of a constant markup. This not only eliminates markup changes as a factor, but also implies that intertemporal cost variations can be measured simply by variations in price. This clearly holds regardless of how we specify the production function or costs of production in equation (1).

Under production to stock, the expected opportunity cost of selling a unit of inventory is equal to $\tilde{E}_t[\beta_{t+1}c_{t+1}]$. $\tilde{E}_t$ denotes the expectations operator conditioned on information available at the time of sales during t. In addition to variables incorporated in E_t, we assume it also reflects s_t and p_t. Assuming a constant markup m therefore implies that p_t equals $(1+m)\tilde{E}_t[\beta_{t+1}c_{t+1}]$. Substituting p_t appropriately for discounted future cost in the firms first-order condition (2), taking expections, and rearranging yields

$$(3) \qquad \tilde{E}_{t-1}\{ \frac{\beta_t p_t}{p_{t-1}}[1 + \frac{\phi m s_t}{a_t}] \} = 1 .$$

(3) predicts strong procyclical movements in the ratio $\frac{s_t}{a_t}$ only if there are opposite cyclical movements in $\frac{\beta_t p_t}{p_{t-1}}$. $\frac{\beta_t p_t}{p_{t-1}}$ will be countercyclical if interest rates are procyclical relative to the expected inflation in the firm's price. We demonstrate below that $\frac{\beta_t p_t}{p_{t-1}}$ exhibits no such cyclical behavior. Consequently, we drop the assumption of a constant markup and proceed to measure movements in marginal cost and markups.

B. *Measuring Marginal Cost of Production*

From the firm's problem (1), marginal cost c_t equals $(\lambda_t\omega_t + c_t^v)$. ω_t is the price of materials, with $\lambda_t\omega_t$ being the cost of materials per unit of output. c_t^v is the marginal cost of labor and capital required to produce a unit of output from those materials.

Let w_t denote the wage for marginally increasing production labor. Given that production labor enters as a power function in technology in (1), the marginal cost of value added is $(\frac{1}{\gamma\alpha})\frac{w_t n_t}{y_t}$, which is proportional to the wage divided by production workers' labor productivity. (See Bils, 1987). This result allows for technology shocks, the impact of which appear through output. A value for $\gamma\alpha$ equal to labor's share roughly corresponds to perfect competition. Higher values for $\gamma\alpha$ reduce marginal cost.

With data on output, materials cost, production hours, and the production labor wage, marginal cost can be calculated given a value for the parameter combination $\gamma\alpha$.

$$(4) \qquad c_t = \lambda_t \omega_t + \left(\frac{1}{\gamma\alpha}\right)\frac{w_t n_t}{y_t}.$$

Part 1 of the appendix describes how we construct monthly indices of materials cost, $\lambda_t \omega_t$, for our six industries. Part 2 of the appendix shows how the parameter combination $\gamma\alpha$ can be related to observables, such as labor's share, and to the returns to scale parameter γ. γ, in turn, is estimated in Section IV. Here we focus attention on the key question of measuring the effective price of labor, w_t.

C. *Measuring the Marginal Price of Labor Input*

It is standard practice to measure the price of production labor by average hourly earnings for production workers. We also consider a competing measure that allows for the possibility that average hourly earnings do not reflect true variations in the price of labor, but rather are smoothed relative to labor's effective price for convenience or to smooth workers' incomes. (See Hall, 1980.) Specifically, we allow for procyclical factor utilization that drives a cyclical wedge between the effective or true cost of labor and average hourly earnings. This may be because in booms workers transitorily boost efforts without *contemporaneous* increases in wages or because capital is transitorily worked longer hours (more shifts) with workers being payed shift premia that do not fully compensate them for the disutility of late work (Shapiro, 1995).

Total factor productivity is markedly procyclical for most manufacturing industries. One interpretation for this finding is that factors are utilized more intensively in booms, with these movements in utilization not captured in the measured cyclicality of inputs (e.g., Solow, 1973). Let x_t denote the effort or exertion per hour of labor for both production and nonproduction workers, and let u_t denote a utilization rate for capital. The production function then becomes

$$y_t = [\theta_t(x_t n_t)^{\alpha}(x_t l_t)^{\nu}(u_t k_t)^{1-\alpha-\nu}]^{\gamma} .$$

We assume firms choose x_t and u_t subject to the constraint that working labor harder requires higher wages as a compensating differential, and working capital longer requires working labor at later, odder hours thereby also requiring higher wages. (Shapiro's estimates suggest that the effective premia for night work represents the primary cost of working capital longer hours.) Therefore the *effective* hourly production worker wage is a function of x_t and s_t, $w_t(x_t, u_t)$, and similarly for the wages of nonproduction workers. Cost minimization requires that firms choose value for x_t and u_t such that the elasticities of these wage functions equal one with respect to x_t and equal $\frac{1-\alpha-\nu}{\alpha+\nu}$ with respect to u_t.[11]

If data on wages capture the contemporaneous impact of x_t and u_t on required wages then the measure for marginal cost in equation (4) remains correct. Higher factor utilization increases labor productivity but at the same times increases the price of labor, $w_t(x_t, u_t)$. Our concern is that hourly wages may reflect a typical level of effort and capital utilization, say $w_t(\overline{x}, \overline{u})$, with employers bearing the cost of their choices for x_t and u_t only gradually over time. Shapiro, for instance, estimates that the effective shift premia is on the order of 25 percent, but the observed premium paid in response to shifting a worker to a late shift is often 5 percent or less. Thus the cost of working capital longer hours may appear as a higher wage rate on day shifts in the future, rather than the contemporaneous payment of the effective shift premium.

If data on hourly earnings reflect average levels of utilization, $\overline{x}$ and $\overline{u}$, then

[11]Implicit here is that the wage functions for production and nonproduction workers exhibit the same elasticities with respect to x_t and u_t at particular values for x_t and u_t.

movements in the effective cost of labor can be approximately related to the observed movements in hourly earnings and unobserved movements in x_t and u_t according to

$$\tilde{w}_t(x_t, u_t) \approx \tilde{w}_t(\overline{x}, \overline{u}) + \tilde{x}_t + \frac{1-\alpha-\nu}{\alpha+\nu}\,\tilde{u}_t\,,$$

where the tilde over a variable denotes the deviation of the natural log of that variable from its longer-run path (defined below by an H-P filter). Here we use the result that the elasticities of wage movements with respect to x_t and s_t must be respectively approximately one and $\frac{1-\alpha-\nu}{\alpha}$. But note that from growth accounting

$$\tilde{x}_t + \frac{1-\alpha-\nu}{\alpha+\nu}\,\tilde{u}_t = \frac{1}{\alpha+\nu}[\frac{1}{\gamma}\tilde{y}_t - \tilde{\theta}_t - \alpha\tilde{n}_t - \nu\tilde{I}_t - (1-\alpha-\nu)\tilde{k}_t]\,.$$

If we assume that high-frequency fluctuations in θ are negligible, then combining these two equations yields our alternative measure of movements in the effective wages

$$(5)\quad \tilde{w}_t(x_t, u_t) \approx \tilde{w}_t(\overline{x}, \overline{u}) + \frac{1}{\alpha+\nu}[\frac{1}{\gamma}\tilde{y}_t - \alpha\tilde{n}_t - \nu\tilde{I}_t - (1-\alpha-\nu)\tilde{k}_t]\,.$$

Thus our measure of the wage augmented to account for varying factor utilization essentially adds high-frequency movements in TFP (scaled by $\frac{1}{\alpha+\nu}$) to average hourly earnings. It interprets those movements as variations in utilization for which labor is compensated, though the compensation does not show up contemporaneously in average hourly earnings. The constructed wage movements in (5), and the implied movements in marginal cost, depend on the returns to scale γ. From the cyclical behavior of inventories we can estimate γ, and thereby judge the extent to which the procyclical behavior of factor productivity reflects increasing returns or procyclical factor utilization.

IV. Results

A. *The Behavior of Inventories*

We begin by examining the behavior of the ratio of sales to stock available for sale, $\frac{s_t}{a_t}$, for the six manufacturing industries: Tobacco, apparel, lumber, chemicals, petroleum, and rubber. These are roughly the six industries commonly identified as production for stock industries (Belsley, 1969).[12] We obtained monthly data on sales and finished inventories, both in constant dollars and seasonally adjusted, from the Department of Commerce. The series are available back to 1959. We construct monthly production from the identity for inventory accumulation, with production equal to sales plus inventory investment.[13]

Figure 2 presents the ratio $\frac{s_t}{a_t}$ for each of the six industries along with industry sales. The period is for 1959.1 to 1997.9. For every industry the ratio of sales to stock available is 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. Table 1, Column 1 presents industry correlations between the ratio $\frac{s_t}{a_t}$ and output with both series H-P filtered. The correlations are all large and positive, ranging from .46 to .84. To show that these correlations do not merely reflect mistakes, e.g. sales forecast errors,

[12]In comparison to Belsley, 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. On the other hand, our understanding 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.

[13]West (1983) discusses that the relative size of inventories is somewhat understated relative to sales because inventories are valued on the basis of unit costs whereas sales are valued at price. We recalculated output adjusting upward the relative size of inventory investment to reflect the ratio of costs to revenue in each of our 6 industries as given in West. This had very little effect. The correlation in detrended log of output with and without this adjustment is greater than 0.99 for each of the industries. It also has very little impact on the estimates of the Euler equation for inventory investment presented below. Therefore we focus here solely on results from simply adding the series for inventory investment to sales.

Column 2 of Table 1 presents correlations between a conditional expectation of $\frac{s_t}{a_t}$ and output. The expectation is conditioned on a set of variables Γ_t and Γ_{t-1}, where Γ_t includes $\ln(a_t)$, $\ln(\frac{s_{t-1}}{a_{t-1}})$, $\ln(y_t)$, $\ln(\frac{p_{t-1}}{p_{t-2}})$, R_t, $\ln(\frac{w_t}{w_{t-1}})$, $\ln(\frac{\lambda_t \omega_t}{\lambda_{t-1}\omega_{t-1}})$, $\ln(\frac{n_t}{y_t})$, $\ln(\frac{n_{t-1}}{y_{t-1}})$, $\ln(TFP_t)$, and $\ln(TFP_{t-1})$. Price, p_t, is measured by the industry's monthly Producer Price Index, and R_t refers to the nominal interest rate measured by the 90-day bankers' acceptance rate. Replacing sales with forecasted sales yields even larger correlations, ranging from .52 to .88.[14]

We want to stress that the strong tendency for $\frac{s_t}{a_t}$ to be procyclical is not peculiar to these six industries. Figure 1 depicted a similar finding for aggregate manufacturing. We also observe this pattern 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.

B. *The Behavior of Marginal Cost and Markups*

Our model suggests that the procyclicality of $\frac{s_t}{a_t}$ requires that marginal cost is temporarily high in booms or that the price markup be countercyclical. We next ask whether costs and markups in fact behave in that manner. We start with the case of a constant markup, so that expected discounted cost can be measured by expected price. We then drop the assumption of a constant markup, and see how well we can explain inventory behavior under our two competing measures of the cost of labor.

With a constant markup the first-order condition for inventory investment reduces to equation (3). If we assume the two variables in this equation are conditionally

[14]Data sources for hours, wages, and TFP are described in part 3 of the appendix. All variables are H-P filtered, also as detailed in the appendix. We also first differenced the series, looking at the correlation of the changes in the ratios $\frac{s_t}{a_t}$ with the rate of growth in output. The correlations are very positive, ranging across industries from 0.18 to 0.70, and averaging 0.47. (Using forecasted growth in $\frac{s_t}{a_t}$ yields even higher correlations, ranging from 0.57 to 0.86.)

distributed jointly lognormal, then it can be written[15]

$$(3') \qquad \tilde{E}_{t-1}[\phi m \frac{s_t}{a_t} + \ln(\frac{\beta_t p_t}{p_{t-1}})] + \kappa \approx 0.$$

β_t reflects the nominal interest rate from t-1 to t as well as a rate of storage cost. The constant term κ reflects covariances between the random variables as well as m. (3′) implies we should see a strong negative relation between expectations of the two variables $\frac{s_t}{a_t}$ and $\ln(\frac{\beta_t p_t}{p_{t-1}})$.

We first report, by industry, the correlation of $\tilde{E}_{t-1}[\ln(\frac{\beta_t p_t}{p_{t-1}})]$ with output. $\tilde{E}_{t-1}$ is based on the set of variables Γ_{t-1} and Γ_{t-2}, where Γ_t is as defined above, plus the variables $\ln(\frac{s_{t-1}}{a_{t-1}})$ and $\ln(\frac{p_{t-1}}{p_{t-2}})$. All variables are H-P filtered. Results are in the first column of Table 2. The correlation is significantly positive for every one of the six industries. This is precisely the opposite of what is necessary to explain the procyclicality of the ratio $\frac{s_t}{a_t}$. The correlation of $\tilde{E}_{t-1}[\ln(\frac{\beta_t p_t}{p_{t-1}})]$ with $\tilde{E}_{t-1}[\frac{s_t}{a_t}]$ appears in Table 2, Column 2. Again the correlation is positive, significant, and large for every industry, ranging from .34 to .72. For equation (3′) to hold these variables need to be negatively correlated. Also, estimating (3′) by GMM yields a statistically significant, negative coefficient estimate for ϕ for every one of the six industries.

We interpret the evidence in Table 2 as strongly rejecting the constant-markup assumption. Indeed it leaves us with even more to explain: Absent changes in markups, we would expect $\frac{s_t}{a_t}$ to be not merely acyclical, but actually countercyclical. This requires a bigger role for the markup (together with intertemporal marginal cost) in accounting for inventory behavior. Therefore we proceed by allowing the markup to vary, as in first-order condition (2). Again assuming variables in the first-order condition

[15] This approximation is arbitrarily good for small values for the real interest rate r and for the ratio $\frac{m\phi s}{a}$. 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.

are conditionally distributed jointly lognormal, the equation can be written

$$(2') \qquad E_t[\frac{\phi m_t s_t}{a_t} + \ln(\frac{\beta_{t+1} c_{t+1}}{c_t})] + \kappa \approx 0 .$$

Where κ reflects covariances between the random variables.

Before estimating (2′) we report correlations of discounted growth in marginal cost, $E_t[\frac{\beta_{t+1} c_{t+1}}{c_t}]$, and the markup, m_t, with detrended output and with $E_t[\frac{s_t}{a_t}]$. Approximating (2′) around average values of m_t and $\frac{s_t}{a_t}$ yields

$$(2'') \qquad E_t[\phi \overline{m} \frac{s_t}{a_t} + \phi \frac{\overline{s}}{\overline{a}} m_t + \ln(\frac{\beta_{t+1} c_{t+1}}{c_t})] - \phi \overline{m} \frac{\overline{s}}{\overline{a}} + \kappa \approx 0 .$$

Thus the procyclicality of $E_t[\frac{s_t}{a_t}]$ requires countercyclical movements in the expectations of $\frac{\beta_{t+1} c_{t+1}}{c_t}$ and/or m_t. To obtain conditional expectations of the variables we again project onto the set of variables Γ_t and Γ_{t-1} described above. For this exercise we assume constant returns to scale in calibrating the size of markups and the parameter α. (See part 2 of the appendix.)

The results, by industry and for each of the two measures of the price of labor, appear in Tables 3 and 4. Considering first the average hourly earnings measure (the first two columns of each table), note that for every industry the growth in marginal cost is very significantly positively correlated with both output and $E_t[\frac{s_t}{a_t}]$. The correlations with output range from .45 to .82. The correlations with $E_t[\frac{s_t}{a_t}]$ range from .30 to .72. Markups do not display a consistent pattern across industries. They are procyclical, and vary positively with $E_t[\frac{s_t}{a_t}]$, in apparel, lumber, and chemicals, whereas they are countercyclical, and vary negatively with $E_t[\frac{s_t}{a_t}]$, in tobacco, petroleum, and rubber. Taken together, these correlations do not bode well for the average hourly earnings-based measure of marginal cost: $E_t[\frac{s_t}{a_t}]$ fails to be consistently negatively related to expected growth in marginal cost or markups, as required by (2″).

The remaining correlations use the wage augmented to reflect uncompensated variations in factor utilization. These correlations appear in the last set of columns in Tables 3 and 4. In Table 3 we see that the cyclical behavior of marginal cost changes completely, with expected growth in marginal cost negatively correlated with output except in the petroleum industry. (Value added is very small in petroleum. So adjustments to the cost of value added have very little impact.) But in Table 4 we see that, even though both $E_t[\frac{s_t}{a_t}]$ and expected growth in marginal cost are strongly procyclical, the two are not systematically correlated with each other. Expected growth in marginal cost is actually positively correlated with $E_t[\frac{s_t}{a_t}]$ in five of the six industries, though significantly so only for petroleum. For tobacco the two variables are significantly negatively related. Using the augmented wage rate does dramatically decrease the magnitude of the correlation between expected growth in marginal cost and $E_t[\frac{s_t}{a_t}]$, except in petroleum.

The expected markups based on our alternate wage and cost measure are much more consistently and dramatically countercyclical. Looking at the far right columns of Tables 3 and 4, the markup is highly countercyclical in all but the lumber industry. Excluding lumber, the correlations of expected markup with output vary from $-.43$ to $-.90$. For lumber the correlation is slightly positive. The correlations of expected markup with $E_t[\frac{s_t}{a_t}]$ varies from $-.49$ to $-.79$, again excluding lumber where it is significantly positive.

C. *Estimation of the First-Order Condition*

The statistics presented thus far suggest that the wage measure augmented to reflect procyclical factor utilization is *qualitatively* more consistent with inventory behavior. We now evaluate the alternative cost measures more formally by estimating first-order condition (2′) by GMM. Bearing in mind that the two measures reflect polar assumptions regarding the interpretation of short-run productivity movements, we do not necessarily expect either measure to rationalize inventory behavior completely; nonetheless we will evaluate which one does so more successfully.

Looking at (2′), the equation contains explicitly the parameter ϕ and implicitly the returns to scale parameter γ through both c_t and m_t. We first estimate (2′) by GMM to obtain separate estimates of γ and ϕ. Note, however, that given values for the real interest rate, storage costs, and returns to scale γ, equation (2′) requires a particular value of ϕ in order for the implied steady-state value of $\frac{s_t}{a_t}$ to be consistent with the average observed value of $\frac{s_t}{a_t}$ for each industry. (See the appendix, part 2.) We therefore secondly estimate equation (2′) imposing this constraint on ϕ as a function of $\hat{\gamma}$.

We estimate equation (2′) separately for each of the two cost measures. Table 5 contains results using the average hourly earnings-based wage, while Table 6 contains results based on our alternative wage. The results in Table 5 using average hourly earnings are nonsensical, overwhelmingly indicating misspecification. Returns to scale are estimated at a very large positive or very large negative number (greater than 16 in absolute value) for all industries but petroleum. To interpret this, note that marginal cost of value added reflects a weight of $\frac{1}{\gamma}$. So by estimating an absurdly high absolute value for γ, the estimating is essentially negating the marginal cost of value added.

The results in Table 6 using the augmented wage are much more reasonable. The constraint that ϕ take the value implied by the steady-state level of $\frac{s_t}{a_t}$ is rejected only for the lumber and rubber industries. Turning first to the constrained estimates, the estimate for returns to scale is implausibly large for tobacco (about 2.9), but varies between 1.09 and 1.42 for the other five industries. For the unconstrained estimates, ϕ is not always estimated very precisely. The estimate of ϕ is positive for four of the industries, but significantly so for only three: apparel, chemicals, and petroleum. The estimates of returns to scale are more robust: They are very similar for the constrained and unconstrained estimates, with the exception of petroleum, which shows mildly decreasing returns to scale ($\hat{\gamma}$=0.82, though not significantly different from the 1.09 constrained estimate) when ϕ is estimated separately.[16]

The estimates in Table 6 suggest that augmenting marginal cost for uncompensated fluctuations in factor utilization goes quite far in explaining the behavior of inventory investment. Nevertheless, we would not argue that Table 6 reflects an exact or "true" measure of marginal cost. The estimate of ϕ comes primarily from the relationship between $\frac{m_t s_t}{a_t}$ and discounted growth of marginal cost. To the extent we have an imperfect measure of marginal cost, the signal-to-noise ratio in the growth rate of marginal cost might be rather low (especially if c_t is close to a random walk). This may suggest focusing on the constrained estimates of ϕ. The fact that we presumably have an imperfect measure of marginal cost could also be reflected in the tendency to reject the overidentifying restrictions of the model according to the J-statistic. (The restrictions are rejected in four of the six industries). On the other hand, the model is fairly successful in accounting for most of the persistence of $\frac{s_t}{a_t}$ without resorting to *ad hoc* adjustment costs. With the exception of the lumber industry, the Durbin-Watson statistics do not suggest the presence of a large amount of unexplained serial correlation.

Our approach of adding back short-run TFP movements to construct an effective wage succeeds in explaining the very procyclical behavior of $\frac{s_t}{a_t}$ by generating a procyclical time-series for marginal cost and a countercyclical one for markups. Table 4 reported correlations of $E_t[\frac{s_t}{a_t}]$ with the expected growth in (discounted) marginal cost

[16]These results are for data with low frequency movements in the variables removed by an H-P filter. Parameter estimates based on unfiltered data are very similar to those in Table 6. The primary difference is that the test statistics for overidentifying restrictions and for the constraint on ϕ more typically reject.

and with the expected markup assuming constant returns to scale. Those correlations suggest that much of the impact of augmenting marginal cost for procyclical factor utilization acts through making the markup very countercyclical. This remains true allowing for the modest increasing returns we estimate in Table 6. Using the estimates of γ from Table 6, Figure 4 presents the implied markup together with the ratio $\frac{s_t}{a_t}$ for each industry. Not only are the markups highly countercyclical, with the exception of lumber, but their movements are quantitatively important. Figure 3 showed that the large shifts in price markups in tobacco in the 1980's and 1990's were accompanied by opposite movements in the ratio $\frac{s_t}{a_t}$ as predicted by the model. Figure 4 shows that, more generally, most of the striking shifts in $\frac{s_t}{a_t}$ that occured in these six industries reflect large opposite movements in the markup.[17]

Our alternate measure for marginal cost movements is much more successful in explaining inventory investment. But it may be excessive, at least for certain industries. In the case of rubber, for instance, augmenting the wage for utilization makes the markup sufficiently countercyclical that $E_t[\frac{m_t s_t}{a_t}]$ is very countercyclical, despite $E_t[\frac{s_t}{a_t}]$ being very procyclical. (This is true, though to a lesser extent, for every industry except lumber.) At the same time, it causes the expected growth in marginal cost, $E_t[\frac{\beta_{t+1} c_{t+1}}{c_t}]$, to become countercyclical. (See the last row of Table 3.) Thus both variables in the equation become countercyclical, resulting in a negative (though insignificant) estimate of ϕ. Presumably a somewhat smaller adjustment would yield a positive estimate for ϕ.

[17] Several empirical papers have examined the cyclicality of markups. (Rotemberg and Woodford, 1995, survey some of these.) Our definition of the markup is slightly different, as it compares price to discounted next period's marginal cost. The markup of price relative to contemporaneous marginal cost, however, behaves extremely similarly to the markups pictured in Figure 4.

V. Conclusions

Evidence from cross-sectional and low frequency data indicates that firms' demands for finished goods inventories are proportional to their expected sales. Yet during business cycles these inventories are highly countercyclical relative to sales. We can explain this behavior if firms exhibit procyclical marginal cost and countercyclical price markups. Obvious measures for marginal cost do not show high marginal cost in booms because factor productivity rises during expansions. We show that the cyclical patterns of inventory holdings can be rationalized by interpreting fluctuations in labor productivity as arising primarily from mismeasured cyclical factor utilization, 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 factor utilization accounts for this puzzle is consistent with other evidence that factors are worked more intensively in booms (for example, Bernanke and Parkinson, 1991, Shapiro, 1993, Bils and Cho, 1994, Burnside, Eichenbaum, and Rebelo, 1995, and Gali, 1997). The results further suggest that countercyclical markups (either purposeful or reflecting price rigidities) may contribute to business cycles by muting the role of diminishing returns in partially stabilizing fluctuations.

References

Basu, Susantu and John G. Fernald, "Returns to Scale in U.S. Production: Estimates and Implications," Journal of Political Economy 105 (1997): 249-283.

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.

Burnside, Craig, Eichenbaum, Martin, and Sergio Rebelo, "Capital Utilization and Returns to Scale," in B.S. Bernanke and J.J. Rotemberg, eds., Macroeconomics Annual. Cambridge, MA: MIT Press, 1995.

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.

Farmer, Roger, and Jang Ting Guo, "Real Business Cycles and the Animal Spirits

Hypothesis," *Journal of Economic Theory* (1994).

Gali, Jordi, "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations," manuscript, New York, University, 1997.

Gertler, Mark, and Simon Gilchrist, "Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms," *Quarterly Journal of Economics* 109 (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.

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.

Krane, Spencer D. and Steven N Braun, "Production Smoothing Evidence from Physical Product Data," *Journal of Political Economy* 99 (1991), 558-581.

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).

Rotemberg, Julio J. and Michael Woodford, "Dynamic General Equilibrium Models with Imperfectly Competitive Markets," in *Frontiers of Business Cycle Research*, edited by T. Cooley. Princeton: Princeton University Press, 1995.

Shapiro, Matthew D., "Cyclical Productivity and the Workweek of Capital," American Economic Review: Papers and Proceedings 83 (May 1993): 229-233.

Shapiro, Matthew D., "Capital Utilization and the Marginal Premium for Work at Night," manuscript, University of Michigan, 1995.

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 Note on the Econometric Use of Constant Dollar Inventory Series," *Economics Letters* 13 (1983), 337-341.

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

West, Kenneth, "The Sources of Fluctuations in Aggregate Inventories and GNP," *Quarterly Journal of Economics* 105 (1991), 939-972.

Appendix

1. *The Cost of Materials*

We know of no monthly data on 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. This contrasts with Durlauf and Maccini (1995), who scale up the shares for those inputs they consider so that they sum to one, which results in more volatile input price indices than ours.

Although we assume that materials are a fixed input per unit of output, we do not impose that this input be constant through time. We allow low frequency movements in the input λ_t by imposing that our series $\lambda_t \omega_t$ exhibit the same H-P filter as does the industry's material input measured by the annual survey of manufacturing (from the NBER Productivity Database).

2. *Production Function Parameters*

From equation (4) the relative contributions to marginal cost of materials versus factors in value added depends on the parameter combination $\gamma\alpha$. $\gamma\alpha$ differs from labor's share not only due to markups, but also because production is to stock. Evaluated at steady-state, the first-order condition equating the wage to the present value of marginal revenue product yields

$$\text{(A1)} \qquad \gamma\alpha = \tilde{\beta}(1+\mathrm{m})\frac{\mathrm{wn}}{\mathrm{py}}\left[\frac{1}{1-\tilde{\beta}(1+\mathrm{m})\frac{\lambda\omega}{\mathrm{p}}}\right] .$$

The term $\tilde{\beta}$ equals $\frac{1-\delta}{1+\mathrm{r}}$, reflecting discounting for a real rate of interest r and the storage cost δ.

Conditional on the size of economic profits, the size of the markup is closely related to the degree of increasing returns to scale. Define π as the rate of profit per unit cost of value added. The steady-state relation between the markup, returns to scale, and the profit rate is

$$\text{(A2)} \qquad 1 + \mathrm{m} = \frac{(\gamma+\pi)[1 - (1-\frac{\mathrm{s}}{\mathrm{a}})\tilde{\beta}]}{\tilde{\beta}[\frac{\mathrm{s}}{\mathrm{a}} + (\gamma-1+\pi)[1 - (1-\frac{\mathrm{s}}{\mathrm{a}})\tilde{\beta}]]}$$

Under production to stock the markup is positive even under constant returns (γ equal one) and zero profits. A small markup is needed to cover costs of holding inventories.

Substituting the markup from (A2) into (A1) yields

$$\text{(A3)} \qquad \gamma\alpha = \frac{(\gamma + \pi)\frac{\mathrm{wn}}{\mathrm{py}}}{\dfrac{\frac{\mathrm{s}}{\mathrm{a}}}{1 - (1-\frac{\mathrm{s}}{\mathrm{a}})\tilde{\beta}} - \frac{\lambda\omega}{\mathrm{p}}} .$$

This expresses $\gamma\alpha$ in terms of the parameter combination $(\gamma+\pi)$, which we estimate, and observables. We measure the term $\frac{\frac{\mathrm{s}}{\mathrm{a}}}{1 - (1-\frac{\mathrm{s}}{\mathrm{a}})\tilde{\beta}}$ by its sample average, where $\tilde{\beta}$ reflects a monthly storage cost of one percent and the real interest rate implied by the interest rate as measured by the 90-day bankers' acceptance rate. We measure $\frac{\mathrm{wn}}{\mathrm{py}}$ and $\frac{\lambda\omega}{\mathrm{p}}$ respectively by H-P filters fit to series for production labor and material shares of gross output. Therefore $\gamma\alpha$ varies at low frequencies as well.

In the text we refer to estimating returns to scale γ. From (A3) we see that the model actually identifies γ plus the profit rate π. A number of studies have suggested that profit rates in manufacturing are fairly close to zero. For example, Basu and Fernald (1997) experiment with several different industry cost of capital series and always find very low profit rates, on the order of three percent, for manufacturing industries. With this rationale, we implicitly treat the profit rate as equalling zero for convenience in the text discussion. More generally, our estimates of $\gamma-1$ in Tables 5 and 6 can be interpreted approximately as estimates of $(\gamma-1+\pi)$, that is, returns to scale plus the profit rate.[18]

Finally, the first-order condition in the steady state implies a relationship between ϕ, γ, $\tilde{\beta}$, and the long-run value of $\frac{\mathrm{s_t}}{\mathrm{a_t}}$. This relationship is

$$\text{(A4)} \qquad \phi = \frac{1 - \tilde{\beta}}{\mathrm{m}\tilde{\beta}\frac{\mathrm{s}}{\mathrm{a}}}$$

where m is taken from (A2) above as a function of γ (or more generally $\gamma+\pi$). This is the constraint imposed for the second set of estimation results in Tables 5 and 6.

[18] Technically, the augmented wage in equation (4) has γ entering separately from π. But in practice this makes no difference: One can still interpret the estimates of $\gamma-1$ in Table 6 more generally as estimates of $\gamma-1+\pi$, even if profit rates are as high as 10 percent.

3. *Data sources for Hours, Wages, and TFP*

Monthly data for hours and wages for production workers are from the the Bureau of Labor Statistics (BLS) Establishment Survey. We employ TFP as an instrument and in adjusting the wage according to equation (4). Output in TFP is measured by sales plus inventory accumulation, as described in the text. In addition to output and production labor, TFP reflects movements in nonproduction labor and capital. Employment for nonproduction workers is based on the BLS Establishment Survey. There are no monthly data on workweeks for nonproduction workers. We assume workweeks for nonproduction workers vary according to variations in workweeks for production workers. We have annual measures of industry capital stocks from the Commerce Department for 1959 to 1996, which we interpolate to get monthly stocks.

4. *Detrending Procedures*

Although the first-order condition (2′) suggests that quantities such as $\frac{m_t s_t}{a_t}$ and $\ln(\frac{\beta_{t+1} c_{t+1}}{c_t})$ ought to be stationary (or at least cointegrated), this may not necessarily hold over the nearly 40-year period covered by the sample. Changes in product composition or inventory technology, for example, could produce low frequency movements in these variables that are really outside the scope of this paper. We therefore remove low frequency shifts in these variables with a Hodrick-Prescott filter, using a parameter of 86,400. (The conventional choice of 14,400 for monthly data is only appropriate for series with significant trends--for the above variables it would take out too much high frequency variation.) Moreover, because these variables depend on an estimated parameter (γ), it is necessary to detrend subcomponents that do not involve γ. We use the same H-P filter on the instruments.

We also use the same filter on log(TFP) in constructing the augmented wage, though here the purpose is different. Our assumption is that low-frequency movements in log(TFP), the part removed by the filter, reflect technical change, so we remove that component before using the residual (which we assume reflects varying utilization) to augment average hourly earnings.

Table 1 – – Correlations Between $\ln y_t$ and $\frac{s_t}{a_t}$ †

	Correlation $\frac{s_t}{a_t}$ and $\ln y_t$	Correlation $E_t(\frac{s_t}{a_t})$ and $\ln y_t$
Tobacco	.663	.854
Apparel	.484	.567
Lumber	.644	.723
Chemicals	.837	.880
Petroleum	.455	.516
Rubber	.791	.853

† The sample is 1959.1 to 1997.9. All correlations have p-values < 0.01. y_t is output; $\frac{s_t}{a_t}$ is the ratio of sales to the stock available for sale.

Table 2--Constant Markup Case: Correlations of $\tilde{E}_{t-1}[\ln(\frac{\beta_t p_t}{p_{t-1}})]$ with $\ln y_t$ and with $\tilde{E}_{t-1}[\frac{s_t}{a_t}]$[†]

	Correlation with $\ln y_t$	Correlation with $\tilde{E}_{t-1}(\frac{s_t}{a_t})$
Tobacco	.256	.415
Apparel	.367	.351
Lumber	.190	.335
Chemicals	.570	.722
Petroleum	.270	.579
Rubber	.233	.449

[†] The sample is 1959.1 to 1997.9. All correlations have p-values < 0.01.

$\frac{\beta_t p_t}{p_{t-1}}$ is the discounted growth in output price from t-1 to t; y_t is output; $\frac{s_t}{a_t}$ is the ratio of sales to the stock available for sale.

Table 3--Correlations of ln y_t with $E_t[\frac{\beta_{t+1}c_{t+1}}{c_t}]$ and with $E_t[m_t]$ †

	Using Average Hourly Earnings		Using the Augmented Wage Measure	
	$E_t[\frac{\beta_{t+1}c_{t+1}}{c_t}]$	$E_t[m_t]$	$E_t[\frac{\beta_{t+1}c_{t+1}}{c_t}]$	$E_t[m_t]$
Tobacco	.821	−.336	−.882	−.433
Apparel	.750	.210	−.266	−.615
Lumber	.612	.533	−.228	.076*
Chemicals	.450	.192	−.186	−.836
Petroleum	.562	−.331	.467	−.526
Rubber	.458	−.360	−.238	−.897

† The sample is 1959.1 to 1997.9.

y_t is output; $\frac{\beta_{t+1}c_{t+1}}{c_t}$ is the discounted growth in marginal cost from t to t+1; m_t is the net markup of price over discounted next period marginal cost.

* All p-values < 0.05, except for this correlation.

Table 4—Correlations of $E_t[\frac{s_t}{a_t}]$ with $E_t[\frac{\beta_{t+1}c_{t+1}}{c_t}]$ and $E_t[m_t]$ †

	Using Average Hourly Earnings		Using the Augmented Wage Measure	
	$E_t[\frac{\beta_{t+1}c_{t+1}}{c_t}]$	$E_t[m_t]$	$E_t[\frac{\beta_{t+1}c_{t+1}}{c_t}]$	$E_t[m_t]$
Tobacco	.720	−.439	−.721	−.506
Apparel	.300	.157	.010*	−.486
Lumber	.396	.609	.044*	.219
Chemicals	.328	.299	.109*	−.741
Petroleum	.356	−.211	.366	−.311
Rubber	.428	−.322	.067*	−.790

† The sample is 1959.1 to 1997.9.
$\frac{s_t}{a_t}$ is the ratio of sales to the stock available for sale; $\frac{\beta_{t+1}c_{t+1}}{c_t}$ is the discounted growth in marginal cost from t to t+1; m_t is the net markup of price over discounted next period's marginal cost.

* All p-values < 0.05, except for these correlations.

Table 5--GMM Estimates of Inventory Demand Using Average Hourly Earnings[†]

	ϕ	γ-1	J-statistic	D-W statistic
Tobacco	−.070 (.042)	74.67 (34.30)	27.0	2.30
	.011*	101.3 (55.91)	27.2	2.35
Apparel	.026 (.021)	−59.66 (33.11)	31.3	1.30
	.020*	−69.80 (45.00)	31.2	1.31
Lumber	.005 (.025)	179.7 (959.0)	39.1	0.90
	.024*	23.49 (14.72)	39.5	0.87
Chemicals	−.068 (.028)	15.23 (6.53)	27.8	1.17
	.018*	37.46 (31.64)	31.5	1.09
Petroleum	.524 (.061)	1.491 (.778)	24.4	1.40
	.494*	.091 (.028)	22.2	1.33
Rubber	−.043 (.016)	−330.7 (1057.7)	25.6	1.00
	.020*	−112.9 (104.7)	20.9	0.94

[†]The sample is 1959.1 to 1997.9. Standard errors are in parentheses. The 0.05 critical value for the J-statistic is 28.87.

*Constrained based on estimates of γ, according to steady state. The constraint on ϕ and γ is rejected at a 0.05 critical value for all industries except apparel and lumber.

Table 6--GMM Estimates of Inventory Demand Using the Wage Augmented for Variations in Factor Utilization†

	ϕ	γ-1	J-statistic	D-W statistic
Tobacco	−.021 (.043)	1.932 (.134)	21.3	2.34
	.023*	1.988 (.136)	23.3	2.36
Apparel	.180 (.056)	.188 (.025)	45.4	1.85
	.205*	.175 (.023)	45.5	1.84
Lumber**	.033 (.027)	.395 (.051)	38.8	0.92
	.106*	.408 (.051)	36.4	0.84
Chemicals	.076 (.029)	.395 (.047)	35.7	1.41
	.094*	.416 (.043)	37.0	1.38
Petroleum	.325 (.101)	−.181 (.174)	21.1	1.45
	.486*	.094 (.021)	23.6	1.32
Rubber**	−.047 (.054)	.186 (.039)	37.6	1.80
	.160*	.248 (.030)	42.5	1.73

---

†The sample is 1959.1 to 1997.9. Standard errors are in parentheses. The 0.05 critical value for the J-statistic is 28.87.

*Constrained based on estimates of γ, according to steady state.

**Constraint on ϕ and γ is rejected with a 0.05 critical value.

Figure 1: The Cyclical Behavior of the Sales-Stock Ratio in Aggregate Manufacturing

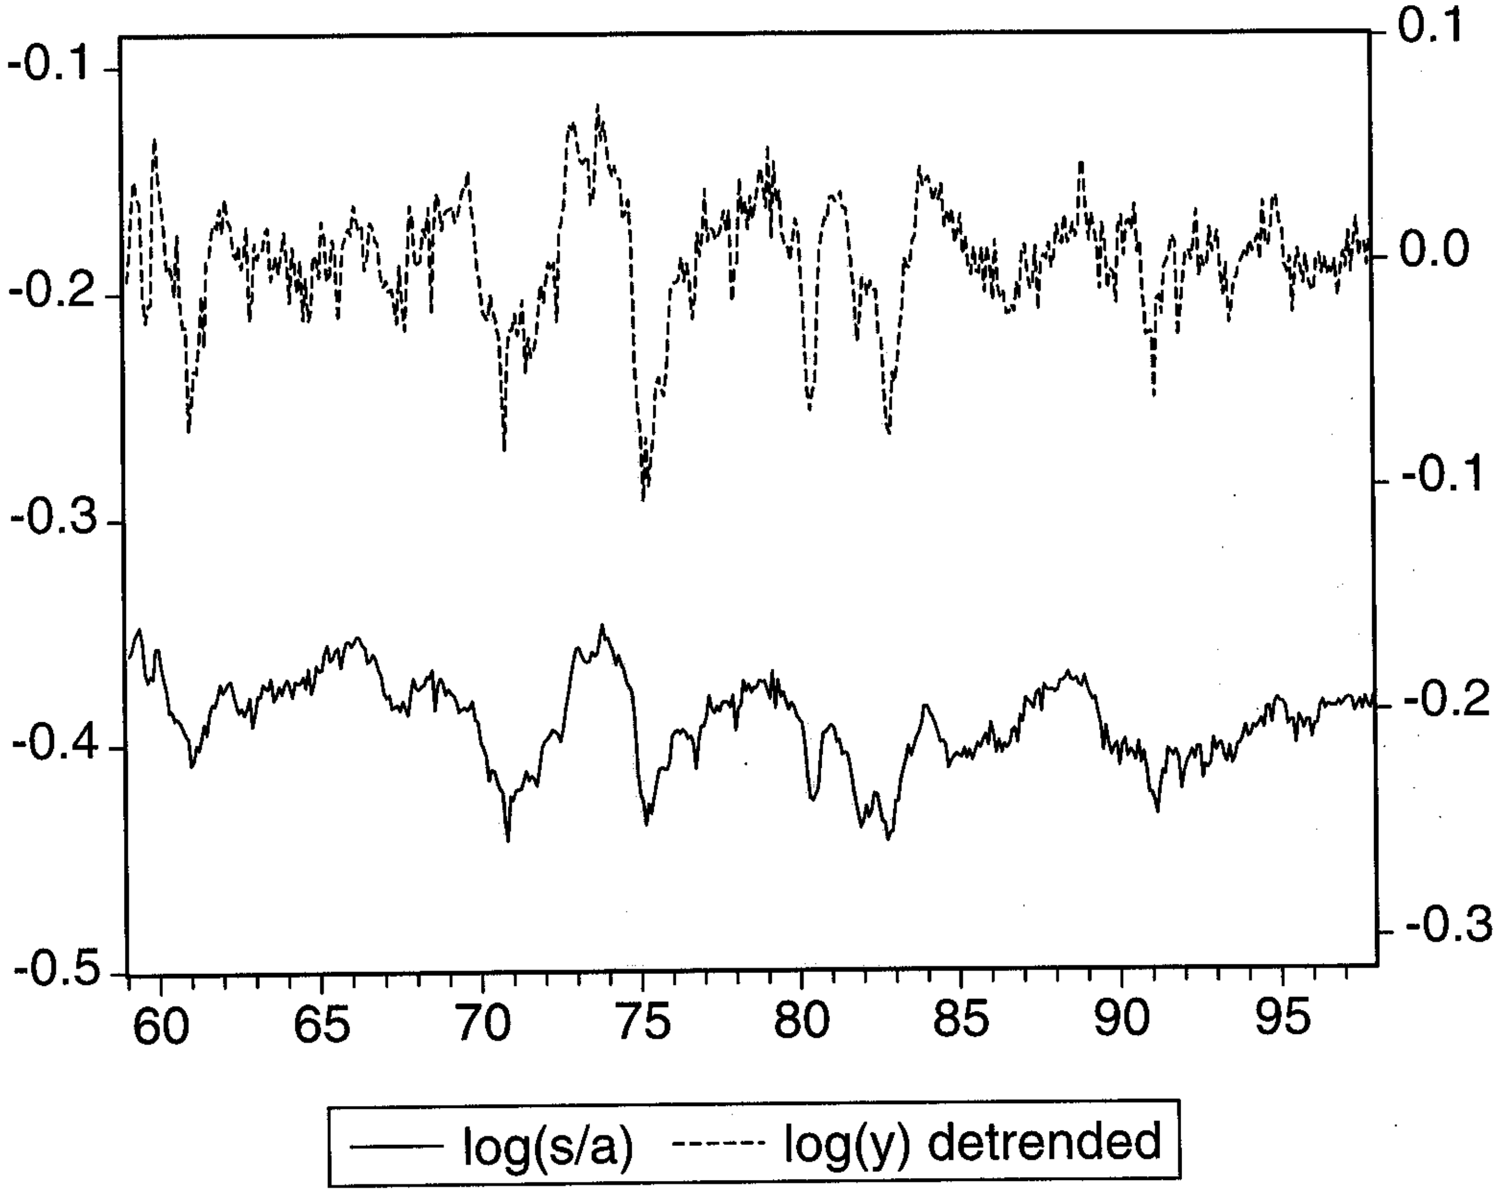

Note: Shaded areas indicate recessions. log(y) detrended with H-P filter.
y=output, s=sales, a=i+y, i=beginning finished goods inventory stock.

Figure 2: Cycles and Trends in s/a and s

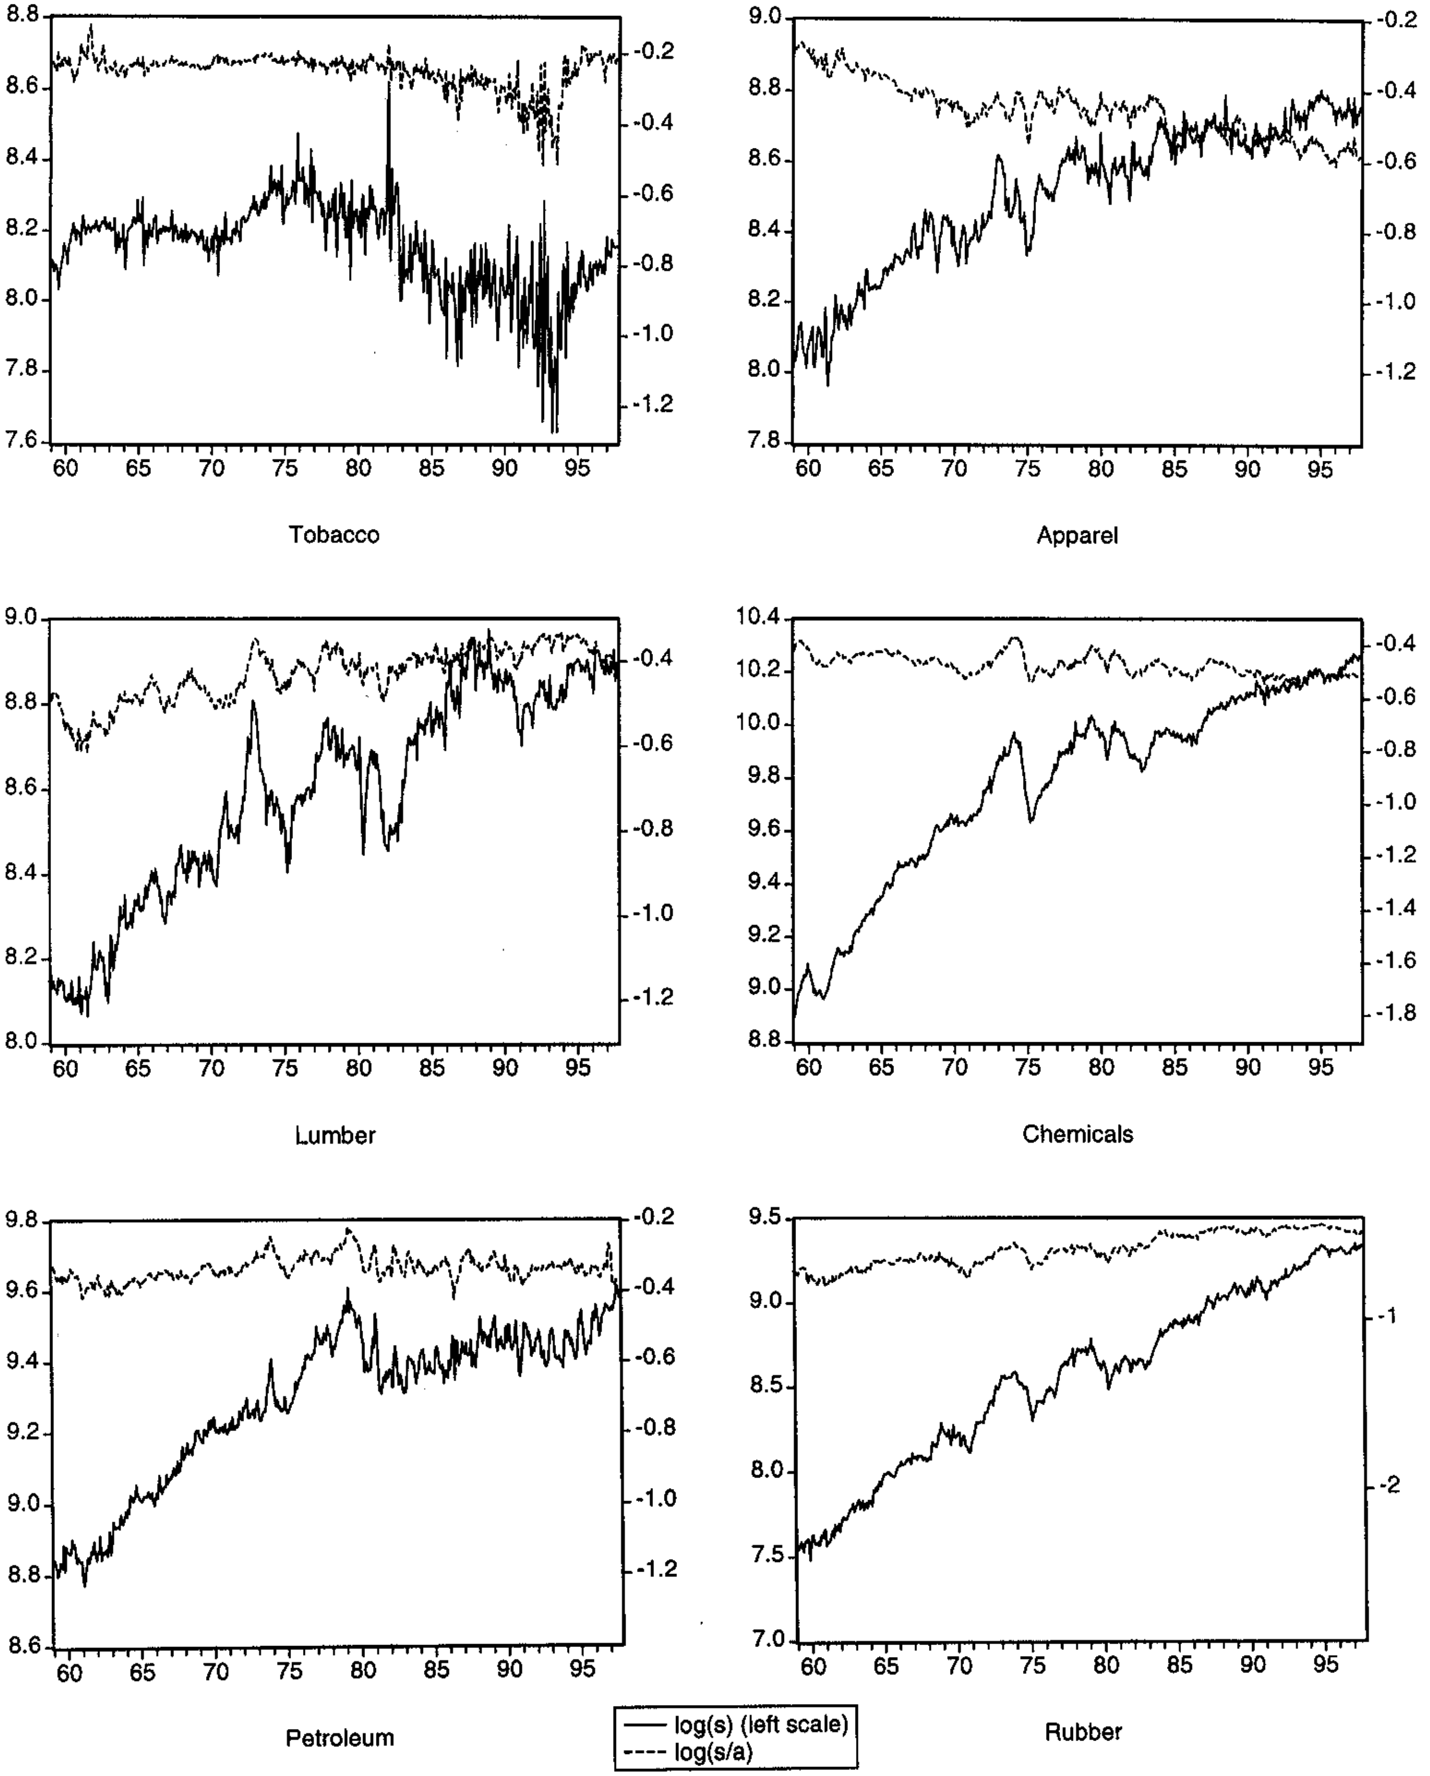

Shaded areas indicate recessions
s=sales, a=i+y, i=beginning finished goods inventory stock.

Figure 3: Price and s/a in the Tobacco Industry

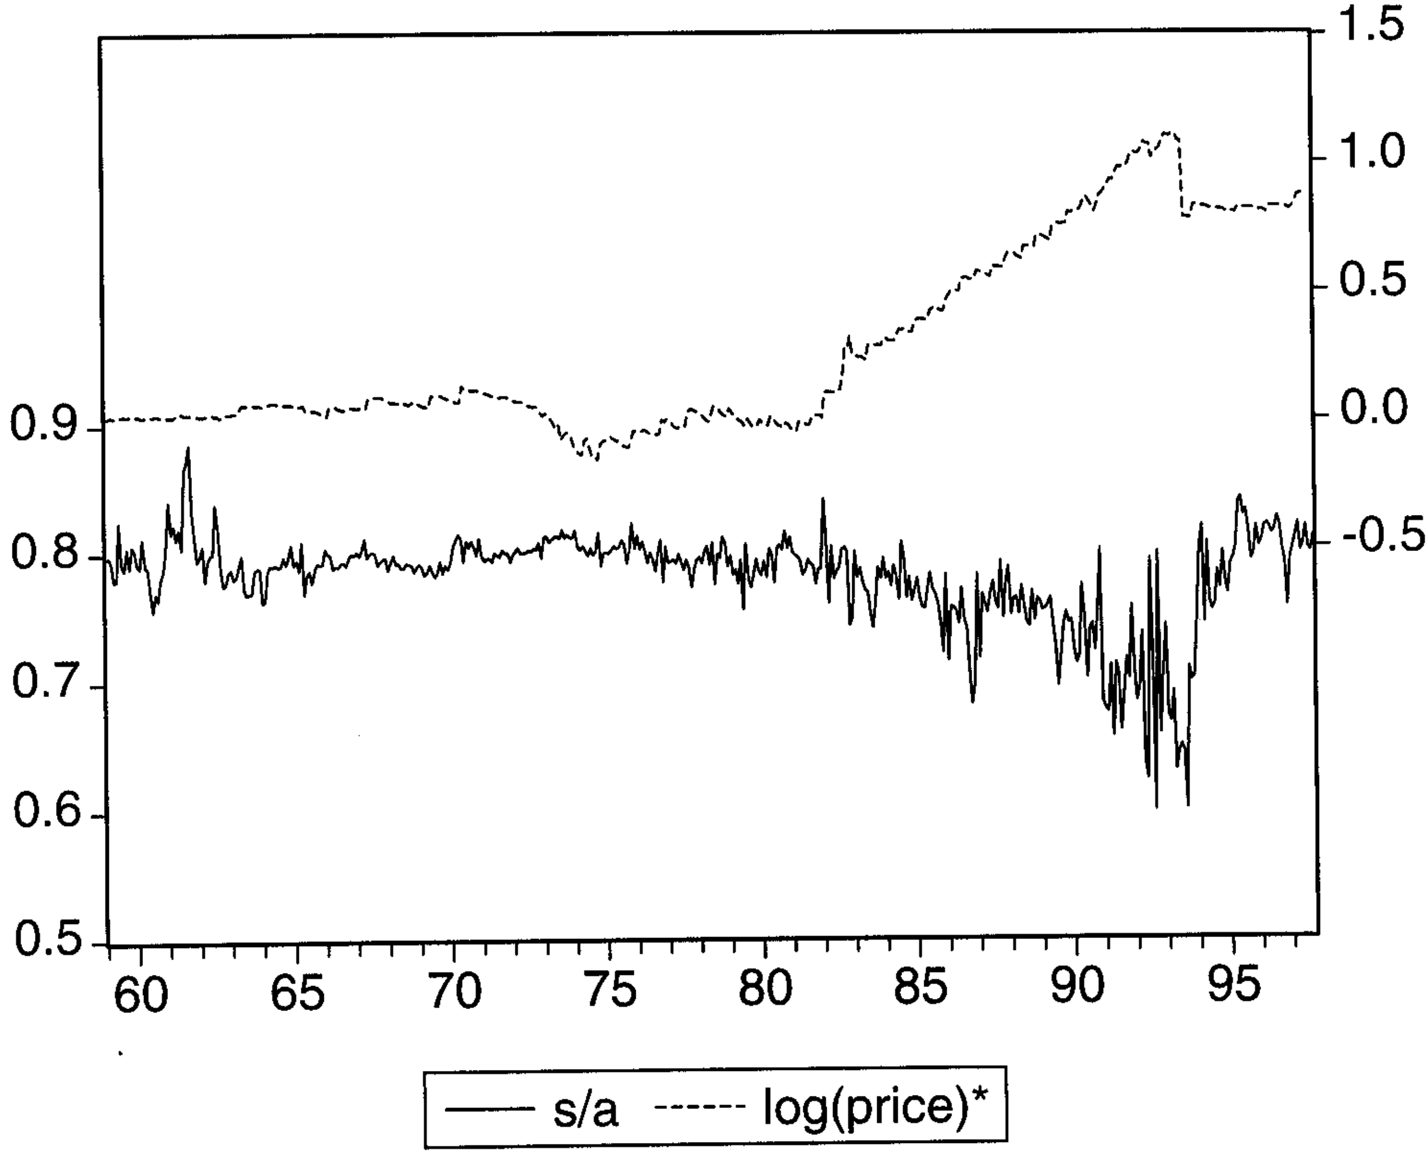

*Tobacco products price deflated by the general Producer Price Index

Figure 4
Markups and s/a Ratio with Estimated Returns to Scale

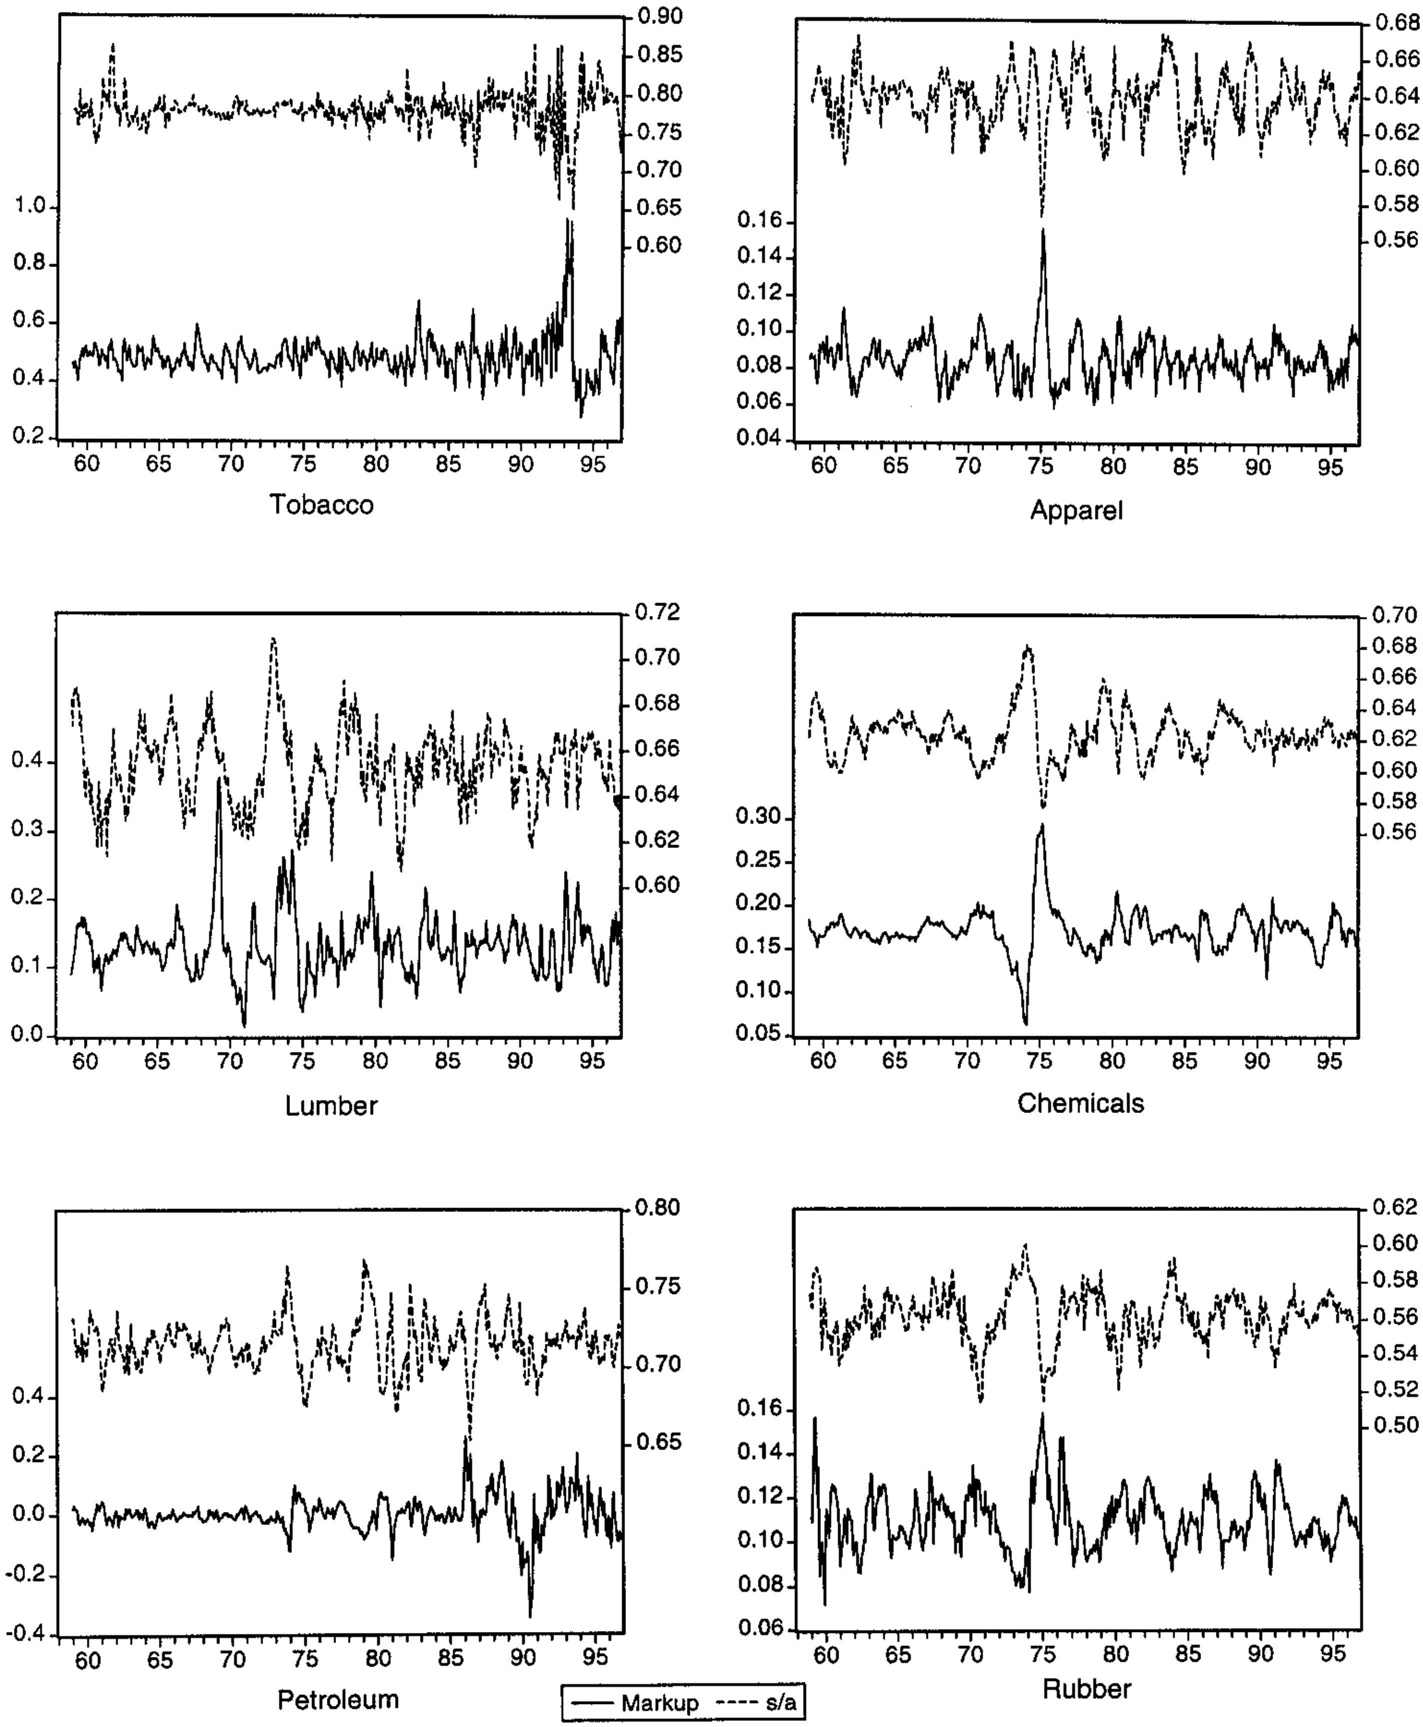

s=sales, a=i+y, i=beginning finished goods inventory stock.
The markup is price in t relative to discounted marginal cost in t+1, minus 1