What Inventory Behavior Tells Us about Business Cycles

By Mark Bils and James A. Kahn*

The countercyclical pattern of inventory-sales ratios is a striking feature of inventory behavior. In a model where inventories are productive for sales, both the markup of price over marginal cost and expected changes in marginal cost are key determinants of that ratio. This paper argues that costly variation in factor utilization gives rise to countercyclical markups in production-to-stock manufacturing industries. The markup turns out to be more important than intertemporal substitution in explaining the behavior of inventory-sales ratios. (*JEL* E22, E32)

Researchers have studied inventory behavior because it provides clues to the nature of business cycles. Many have viewed the procyclical behavior of inventory investment as evidence that costs of producing are lower in an expansion 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

---

*Bils: Department of Economics, University of Rochester, Rochester, NY 14627. Kahn: Research Department, Federal Reserve Bank of New York, 33 Liberty Street, New York, NY 10045. We acknowledge support from the National Science Foundation. We also thank Anil Kashyap, Peter Klenow, Valerie Ramey, Julio Rotemberg, Ken West, Michael Woodford, three referees, and participants at a number of seminars for helpful comments. A data diskette is available upon request that contains both the raw and filtered data series that are used in our empirical work. The views expressed are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of New York or the Federal Reserve System.

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

in turn is viewed as evidence for procyclical technology shocks, increasing returns, or positive externalities.[2]

We claim that this reasoning is false. The argument outlined above overlooks changes in the shadow value of inventories, which we argue increases with expected sales.[3] We propose a model in which finished goods inventories facilitate sales. The model implies that, holding prices fixed, inventories should vary in proportion to anticipated sales, as in fact they do in the long run. Over the business cycle, however, the ratio of sales to stocks is highly persistent and procyclical, which suggests that inventory stocks behave sluggishly in the short run. This seemingly paradoxical feature of inventory behavior—the sluggish adjustment of stocks even to relatively small changes in targets—has been noted by researchers going back at least to Martin S. Feldstein and Alan J. Auerbach (1976).

Figure 1 plots the monthly ratio of sales (shipments) to the sum of beginning-of-period finished goods inventories plus production (what we define as the "stock available for sale") in aggregate manufacturing for 1959 through 1997, together with production. Production is detrended using a Hodrick-Prescott (H-P) filter. The sales-stock ratio decreases dramatically in each recession, typically by 5 to 10 percent. Note that these decreases do not simply reflect transitory sales surprises, 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

---

[2] West (1991) explicitly uses inventory behavior to decompose the sources of cyclical fluctuations into cost and demand shocks. Martin S. Eichenbaum (1989) introduces unobserved cost shocks that generate simultaneous expansions in production and inventory investment. Valerie A. Ramey (1991) estimates a downward sloping short-run marginal cost function, which of course reverses the production-smoothing prediction. See also Robert E. Hall (1991). Russell Cooper and John C. Haltiwanger (1992) adopt a nonconvex technology on the basis of observations about inventory behavior. Others (e.g., Mark L. Gertler and Simon Gilchrist, 1993, Anil K. Kashyap et al., 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.

[3] Robert S. 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 Olivier J. Blanchard (1983), West (1986), Spencer D. Krane and Steven N. Braun (1991), Kashyap and David W. Wilcox (1993), and Steven N. Durlauf and Louis J. 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.

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]

Because in the long run inventories do track sales one for one, we find the real puzzle to be why inventory investment is not *more* procyclical. Inventories sell with predictably higher probability at peaks, suggesting that—*ceteris paribus*—firms should add more inventories in booms so as to equate the ratios (and hence the "returns") over time. Our model shows that this striking fact implies that in booms marginal cost must be high relative to either (1) discounted future marginal cost; or (2) the price of output. The former chokes off intertemporal substitution of production, while the latter implies a relatively small payoff from additional sales.

We start by assuming that markups are constant, which allows us to measure expected movements in marginal cost by expected movements in price. For sales to increase relative to inventories in an expansion then requires that the rate of expected 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 the task of measuring marginal cost separately from output prices. 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. Therefore, we cannot rationalize why inventories fail to keep up with sales over the business cycle.

Finally, we ask: What if the procyclicality of productivity reflects workers providing a greater effort in boom periods, with this expanded effort not contemporaneously reflected in measured average hourly earnings? By assuming that short-run movements in total factor productivity reflect either (estimated) increasing returns to scale or procyclical worker effort, we

[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.

construct an alternative measure of the price of labor and marginal cost. Our alternative measure of marginal cost takes an admittedly extreme position on the source of short-run productivity movements, i.e. that technology shocks are relatively unimportant. But we find that it is much more successful in accounting for the behavior of inventories relative to sales.

Under our alternative measure, an output expansion is associated with an increase in real marginal cost or, equivalently, a countercyclical decrease in the markup of price over marginal cost. This increase in real marginal cost arises, not from diminishing returns to labor in the production function, but from a higher shadow cost of labor. The countercyclical markup is reflected in a fall in desired inventories relative to expected sales during expansions, justifying why inventories fail to keep up with sales over the business cycle. We find little reason for firms to engage in the standard production or cost-smoothing envisioned in conventional inventory models. Such intertemporal substitution requires forecastable changes in marginal cost relative to interest rates that we cannot find in the data. The last finding is important given that the linear-quadratic inventory model—by far the most commonly employed model of inventory behavior—imposes a constant target sales-stock relationship and requires that persistent deviations from that target be the result of intertemporal substitution.

We find the joint behavior of inventories, prices, and productivity consistent with the following view of business cycles: Real marginal cost is procyclical, but changes are not sufficiently predictable relative to real interest rates to give rise to intertemporal substitution. The rise in real marginal cost during an expansion is equivalent to a decline in the markup; it damps production by reducing optimal inventory holdings relative to expected sales. Thus the salient features of inventory behavior are not the result of persistent deviations from a fixed target sales-stock ratio; rather, the target ratio itself varies systematically over the cycle due to countercyclical markups.

I. 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 elements: Profit maximization, a production function, and an inventory technology that is specified to reflect the fact that inventory-sales ratios appear to be independent of scale (which we document below). 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 Finn E. Kydland and Edward C. Prescott (1982), Lawrence J. 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

$$(1) \qquad \max_{y_t} E\left(\sum_{i=0}^{\infty} \beta_{t,t+i}[p_{t+i}s_{t+i} - C_{t+i}(y_{t+i};\theta_t,\boldsymbol{z_t})] \,\middle|\, \mathcal{I}_t \right)$$

subject to

$$\begin{aligned} i) \quad & a_t = i_t + y_t = a_{t-1} - s_{t-1} + y_t \\ ii) \quad & y_t = \min\left\{ \frac{\boldsymbol{q_t}}{\boldsymbol{\lambda}}, (\theta_t n_t^{\alpha} l_t^{\nu} k_t^{1-\alpha-\nu})^{\gamma} \right\} \\ iii) \quad & s_t = d_t(p_t)a_t^{\phi}. \end{aligned}$$

The expectation is conditioned on a set of variables $\mathcal{I}_t$ known when production is chosen in t. In the objective function, s_t and p_t denote sales and price in period t, θ_t a technology shock, and $\boldsymbol{z_t}$ is a vector of input prices. $C_t(y_t)$ is the cost of producing period t's output y_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 $(1+R_{t+1})^{-1}$, where R_{t+1} is the nominal interest rate between t and $t+1$.[5] We assume that when firms choose production for t they know

[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 the proportions used in 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}: $\beta_{t+1} = (1-\delta)/(1+R_{t+1})$.

realizations of the variables θ_t and $\boldsymbol{z_t}$ that determine the costs of producing (as well as the nominal interest rate), but not the realizations of price p_t or sales s_t.

Constraint (i) is just a 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. The presumption here is that any output completed during period t can be sold during period t as well.

Constraint (ii) specifies that output is produced using both a vector of material inputs, $\boldsymbol{q_t}$, and value added produced by a Cobb-Douglas function of production labor, n_t, nonproduction labor, l_t, and capital, k_t. The value-added production function has returns to scale γ, potentially greater than one. Material inputs are proportional to output as dictated by a vector of per unit material requirements, $\boldsymbol{\lambda}$.[6] Although we do not treat elements of $\boldsymbol{\lambda}$ as choice variables, in the empirical work we allow for low-frequency movements in $\boldsymbol{\lambda}$. (For convenience we write $\boldsymbol{\lambda}$ without a time subscript.)

Constraint (iii) 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, where $0 \leq \phi \leq 1$. This approach is consistent, for example, with a competitive market that allows for the possibility of stockouts (e.g., Kahn, 1987, Peter H. Thurlow, 1995). This corresponds to the case $\phi = 1$, because a competitive firm can sell as much as it wants up to a_t. At the other extreme, $\phi = 0$ represents a pure cost-smoothing model, where the firm decouples the timing of production from sales. More generally, 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, but the marginal benefit of this diminishes in a relative to expected sales. This corresponds to the intermediate case of ϕ between 0 and 1. Pindyck (1994) provides evidence for a similar functional form.

[6]Here and elsewhere in the paper, notation such as $\boldsymbol{q}/\boldsymbol{\lambda}$ refers to element-by-element division.

The data strongly suggest that firms do value inventories beyond their role in varying production relative to sales. We typically observe that firms hold stocks of finished inventories that are the equivalent of one to three months' worth of sales. But under a pure production smoothing model it is difficult to even rationalize systematically positive holdings.

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}. In the absence of perfect competition, firms maximize the objective in (1) with respect to a choice of price as well as output. We focus, however, on the choice of output given that price. From constraint (*i*) expanding production translates directly into a higher stock available. Given price, constraint (*iii*) then dictates how that extra stock available translates into greater sales versus greater inventory for the following period.

B. The First-Order Condition for Inventory Investment

In a pure production smoothing model of inventories a firm's expected discounted costs, at an optimum, are not affected by *marginally* increasing current production in conjunction with decreasing subsequent production. Our firms face a similar dynamic first-order condition, but with the additional consideration of the marginal impact of the stock on expected sales. For our firms the appropriate perturbation is producing one more unit during t, adding that unit to the stock available for sale, and then producing less at $t+1$ to the extent that the extra unit for sale during t fails to generate an additional sale. This yields the first-order condition

$$(2) \qquad E\Big(-c_t + \phi d_t(p_t) a_t^{\phi-1} p_t + [1 - \phi d_t(p_t) a_t^{\phi-1}]\beta_{t+1} c_{t+1} \ \big|\ \mathcal{I}_t\Big) = 0$$

The producer incurs marginal cost $c_t \equiv C'(y_t)$. 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 $t+1$. If production is positive at $t+1$ (which we assume), then this inventory can displace a unit 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 equal to $\phi s_t/a_t$, i.e. is proportional to the ratio of sales to stock available. Making this substitution and rearranging gives

$$E\left(\left[\phi\frac{s_t}{a_t}m_t+1\right]\frac{\beta_{t+1}c_{t+1}}{c_t}\ \Big|\ \mathcal{I}_t\right)=1 \tag{3}$$

where $m_t=(p_t-\beta_{t+1}c_{t+1})/(\beta_{t+1}c_{t+1})$. Here m_t is the percent markup of price in t over discounted marginal cost in $t+1$. We refer to this as the markup because $\beta_{t+1}c_{t+1}$ is the opportunity cost of selling a unit at date t. For $\phi>0$, $E_t(m_t)>0$,[7] even under competition and zero profits, as firms require an expected markup to rationalize the costs of inventory holdings. Suppose we denote an aggregate output price deflator by P_t. Note that the term $\beta_{t+1}c_{t+1}/c_t$, the growth rate of nominal marginal cost relative to a nominal interest rate, is equivalent to the growth rate of *real* marginal cost, $c_t/P_{t,}$ relative to the real interest rate defined net of the rate of inflation in P_t.

In a pure production smoothing model ($\phi=0$), the discounted expected growth of marginal cost would always equal 1. That is, nominal marginal cost would always be expected to grow at the nominal interest rate; otherwise it would be profitable to shift production intertemporally. But with $\phi>0$ the desire to smooth costs is balanced against the desire to have a_t track expected sales multiplied by the markup.

If $E_t(\beta_{t+1}c_{t+1})/c_t$ and m_t were both constant through time, then $E_t(s_t)/a_t$ would be constant, i.e. all predictable movements in sales would be matched by proportional movements in the stock available.[8] To generate persistent and systematic procyclical movements in the ratio of sales to inventory such as we see in the data therefore requires either:

[7] In the text we write $E_t(x_t)$ as shorthand for $E(x_t \mid \mathcal{I}_t)$. For example, here $E_t(m_t)$ refers to $E(m_t \mid \mathcal{I}_t)$.

1. Procyclical marginal cost, judged relative to $E_t[\beta_{t+1}c_{t+1}]$
2. A countercyclical markup.

Suppose for the moment that the markup were constant. Then we would observe high expected sales relative to a_t only if marginal cost is high relative to expected next period's marginal cost, i.e. the firm would let a_t fall short of its target only to the extent that marginal cost is temporarily high. Thus, although firms systematically accumulate inventories during expansions, the strong procyclicality of s_t/a_t requires, under a constant markup, that marginal cost be temporarily high in expansions—that is, high relative to next period's discounted cost. The impetus for marginal cost to be temporarily high could be internal (i.e. from a movement along an upward sloping marginal cost curve), or external through input prices.[9] We will refer to this motive for procyclical s_t/a_t as "intertemporal substitution."

Alternatively, suppose that $E_t(\beta_{t+1}c_{t+1})/c_t$ does not vary, i.e., discounted marginal cost is a random walk (possibly with drift). Then we would observe high expected sales relative to a_t only if the markup is low, or, equivalently, real marginal cost (marginal cost relative to output price) is high. The return on holding inventories is largely their ability to generate sales; so a lower markup requires a higher $E_t(s_t)/a_t$ to yield the same return. The strong procyclicality we document for s_t/a_t would, in the absence of an intertemporal substitution motive, thus require countercyclical markups.

We can justify the functional form in constraint (iii) above by examining the low frequency behavior of s_t/a_t, where we can arguably neglect movements in $\beta_{t+1}c_{t+1}/c_t$ and m_t. The model then yields a constant desired ratio of expected sales to stock available (akin to the inverse of the usual inventory-sales ratio) because sales, conditional on price, are a power

---

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

[9] This also suggests little 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 (for example, by increasing the effective interest rate, thereby reducing β_{t+1}). This is opposite the scenario emphasized by Gertler and Gilchrist (1993), Kashyap et al. (1994), and others.

function of the available stock. This implies an absence of scale effects; in other words, the steady-state s/a ratio should be independent of the size of the industry or firm.

Some evidence can be gleaned from observing how the ratio 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 tobacco, sales increased by 50 percent or more from 1959 to 1997. Figure 2 presents the behavior of 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 such movements are for apparel, where the ratio declines by about 25 percent, and in rubber, where it rises by about 20 percent. Clearly there are no *systematic* scale effects on s_t/a_t, though there are some secular changes in some industries.

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 have higher inventory-sales ratios. We conclude that scale effects do not appear to be a promising explanation for the cyclical behavior of s_t/a_t.[10]

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

---

[10] 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 $d_t(p_t)[a_t - \overline{a}\,]^{\phi}$, implying s_t increases with s_t only after the available stock reaches a threshold value $\overline{a}$. 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 $\overline{a}$ were typically less that 20 percent of the average size of a_t; and its introduction did not significantly affect other estimated results.

cost function is[11]

$$C(y_t, a_t) = \frac{\psi}{2} y_t^2 + \frac{\rho}{2}(a_t - \mu s_t)^2 + (\boldsymbol{\lambda' \omega_t} + \xi w_t + \epsilon_t) y_t. \tag{4}$$

where, as before, y_t, a_t, and s_t are output, stock available, and sales during t, and where ρ and μ are assumed to be positive. The last term multiplying y_t represents input costs, including materials costs $\boldsymbol{\lambda' \omega_t}$ (where $\boldsymbol{\omega}$ is a vector of materials prices), labor input ξ with wage w, and a general cost shock ϵ (which could be correlated with output), all expressed in real terms. The slope of marginal cost arising from convexity of the production function is governed by the parameter ψ, whereas the cyclical behavior of marginal cost depends in addition on the behavior of $\boldsymbol{\lambda' \omega_t} + \xi w_t + \epsilon_t$. Note that $\mu > 0$ allows for a target s_t/a_t ratio.

The first-order condition for minimizing the present discounted value of costs based on this cost function is

$$E\Big(\rho(a_t - \mu s_t) - (\beta c_{t+1} - c_t) \;\Big|\, \mathcal{I}_t\Big) = 0. \tag{5}$$

where β is the discount factor and marginal cost $c_t \equiv \psi y_t + \boldsymbol{\lambda' \omega_t} + \xi w_t + \epsilon_t$. Thus a_t deviates from μs_t only to the extent that c_t is expected to deviate from βc_{t+1}. Functional form aside, this condition is very similar to our condition (3). The crucial difference is that here μ is just a parameter, whereas the term in (3) that corresponds to μ is proportional to a time-varying markup.

Many researchers (e.g., Blinder, 1986, Fair, 1989) have focused on the relative volatility of production and sales, or on the related question of why inventory investment is procyclical. Yet procyclical inventory investment and production varying mroe than sales are both perfectly consistent with marginal cost being either procyclical or countercyclical (see West, 1986, or Kahn, 1987). On the other hand, this linear-quadratic model implies that a procyclical s_t/a_t ratio requires procyclical marginal cost. Otherwise the firm could reduce its costs by bunching production in periods with high sales, thereby generating a countercyclical ratio.[12] Thus the

[11] A number of papers include a cost of changing output. Its exclusion here is simply for convenience. Also, note that the specification in terms of a_t is observationally equivalent to the more typical specification with i_t.

cyclical behavior of s_t/a_t is more revealing than the cyclical behavior of the stock alone. What needs explaining, therefore, is not why inventory investment is so procyclical, but rather why it is not *more* procyclical, i.e. why a_t fails to keep up with s_t over the cycle.

Our approach differs substantially from the linear-quadratic literature in at least two ways. First, 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 measurable shocks to productivity. Second, and more important, our model explicitly considers the revenue side of the firm's maximization problem. This allows us to account for variation in target inventory holdings caused by variation in markups. In our model the return on finished inventory is proportional to the markup; so sales relative to stock available should move inversely with the markup. The standard specification of the linear-quadratic model does not permit variation in μ, and therefore requires that all persistent deviations of inventories from their target be the result of intertemporal substitution.[13]

In fact, we find that movements in the markup (and, hence, movements in the desired s/a ratio) are the dominant explanation for the procyclical behavior of s_t/a_t in five of the six industries we examine. Failure to allow for a cyclical markup represents a potentially serious misspecification in the linear-quadratic model, as its effects will be confounded with other cyclical variables in the model, biasing the parameter estimates. This could account for the rather mixed success of the linear-quadratic model, and for why estimates of the slope of marginal cost in the linear-quadratic model have varied so much in the literature.[14]

---

[12] 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 s_t/a_t is procyclical then ψ must be positive.

[13] This distinction between a persistently varying target and persistent deviations from a fixed target dates back to Feldstein and Auerbach (1976). They argued that persistent deviations from a fixed target were inconsistent with the apparent ease with which firms could and did adjust inventory stocks to sales surprises.

[14] For example, Ramey (1991) estimates downward-sloping marginal cost, while others find it upward sloping (e.g., Blanchard, 1983, West, 1986, Krane and Braun, 1991, Kashyap and Wilcox, 1993, and Steven N. Durlauf and Louis J. Maccini, 1995). The linear-quadratic model may also be misspecified in functional form. This is suggested by Pindyck (1994), who finds evidence of a convex marginal convenience yield of inventories consistent with our specification with $\phi < 1$.

The tobacco industry provides an excellent case study to illustrate the importance of the markup. 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 s_t/a_t fell over the same period by about 15 percent. More striking is what occurred 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 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.

II. 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 our 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 (1).

The expected opportunity cost of selling a unit of inventory as of the moment p_t is determined is equal to $E\left(\beta_{t+1} c_{t+1} \mid \mathcal{I}_t,\, s_t,\, p_t\right)$. Note that this expected cost is conditioned not only on the set of variables, $\mathcal{I}_t$, known at the time of choosing production for t, but also includes sales and price for t. Assuming a constant markup m therefore implies that p_t equals $(1+m)E\left(\beta_{t+1}c_{t+1} \mid \mathcal{I}_t,\, s_t,\, p_t\right)$. Substituting p_t appropriately for discounted future cost in the firm's first-order condition (3), taking expectations, and rearranging yields

$$E\left(\frac{\beta_t p_t}{p_{t-1}}\left[1+\frac{\phi m s_t}{a_t}\right] \,\middle|\, \mathcal{I}_{t-1},\, s_{t-1},\, p_{t-1}\right) = 1. \tag{6}$$

Equation (6) predicts strong procyclical movements in the ratio s_t/a_t only if there are opposite cyclical movements in $\beta_t p_t/p_{t-1}$. $\beta_t p_t/p_{t-1}$ will be countercyclical if interest rates are procyclical relative to the expected inflation in the firm's price. We demonstrate below that $\beta_t p_t/p_{t-1}$ exhibits no such cyclical behavior; in fact, movements in $\beta_t p_t/p_{t-1}$ are strongly *positively* correlated with movements in s_t/a_t. Based on this striking result, we consequently 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 $\boldsymbol{\lambda'\omega_t} + v_t$; where $\boldsymbol{\omega_t}$ is a vector of materials prices (hence $\boldsymbol{\lambda'\omega_t}$ is the cost of materials per unit of output), and v_t 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 $(\gamma\alpha)^{-1}w_t n_t/y_t$, which is proportional to the wage divided by production workers' labor productivity. This result allows for technology shocks, the impact of which appear through output. A value for $\gamma\alpha$ equal to labor's share in revenue corresponds to marginal cost equal to price. Higher values for $\gamma\alpha$ are associated with lower marginal cost.

Marginal cost then depends on the observables: output y_t, materials cost $\boldsymbol{\lambda'\omega_t}$, production hours n_t, and the production labor wage w_t; and it depends on the parameter combination $\gamma\alpha$. We have

$$c_t = \boldsymbol{\lambda}'\boldsymbol{\omega_t} + \left(\frac{1}{\gamma\alpha}\right)\frac{w_t n_t}{y_t} \tag{7}$$

Part 2 of the appendix constructs a measure of α based on observable variables (conditional on a profit rate), which turns out to be

$$\alpha = \frac{\frac{wn}{py}}{\zeta - \frac{\boldsymbol{\lambda}'\boldsymbol{\omega}}{p}}. \tag{8}$$

The ratios $wn/(py)$ and $\boldsymbol{\lambda}'\boldsymbol{\omega}/p$ are measured by smooth H-P filters fit to each industry's time series for production labor's and materials' shares in revenue. $\zeta \leq 1$ denotes the sample average of $(s/a)/(1-(1-s/a)\tilde{\beta})$, where the term $\tilde{\beta} = (1-\delta)/(1+r)$ reflects discounting for a real rate of interest r and rate of storage cost δ. As explained in the appendix, ζ adjusts the price of output for the average cost of holding inventories. (Note that $s/a = 1$ implies $\zeta = 1$.) Equation (8) implicitly assumes that firms do not earn pure economic profits. The appendix treats the more general case with pure profits. (It also discusses evidence for a small profit rate.) In the empirical work we consider the robustness of results to profit rates as high as 10 percent of costs.

In estimating first-order condition (3) we proceed as follows. Together, equations (7) and (8) express marginal cost in terms of observables and the parameter γ. We substitute this expression for marginal cost into (3), yielding an equation that depends on observables and the two parameters we estimate, γ and ϕ.

We will describe the data in greater detail below in Section III. Part 1 of the appendix describes how we construct monthly indices of materials cost, $\boldsymbol{\lambda}'\boldsymbol{\omega_t}$, for our six industries. We now consider how to measure the price of labor.

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 depart from this practice by considering 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. (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 because in booms workers transitorily boost efforts without *contemporaneous* increases in measured average hourly earnings.

Total factor productivity (TFP) 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., Robert M. Solow, 1973). We now generalize the production function to allow for variations in worker effort.

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

where x_t denotes the effort or exertion per hour of labor. We treat the choice of x_t as common for production and nonproduction workers.

We assume firms choose x_t subject to the constraint that working labor more intensively requires higher wages as a compensating differential (as in Gary S. Becker, 1985). Therefore the *effective* hourly production worker wage is a function of x_t, $w_t(x_t)$, and similarly for the wages of nonproduction workers. If data on wages capture the contemporaneous impact of x_t on required wages then the measure of marginal cost in equation (7) remains correct. Higher factor utilization increases labor productivity, but at the same time increases the price of labor.

Our concern is that hourly wages may reflect a typical level of effort, say $w_t(\overline{x})$. Employers bear the cost of their choice for x_t, but perhaps in bonuses or promotions that are not reflected, at least concurrently, in data on average hourly earnings. More exactly, suppose we break the marginal price of labor $w_t(x_t)$ into average hourly earnings $w_t(\overline{x})$, reflecting a typical effort level, plus a "bonus payment" $B(x_t - \overline{x})$ that (for convenience) is zero for $x_t = \overline{x}$, and increases with x_t.

(10) $$w_t(x_t) \;=\; w_t(\overline{x}) \;+\; B(x_t - \overline{x}) \;\approx\; w_t(\overline{x}) \;+\; B'(0)[x_t - \overline{x}]\,.$$

The approximately equals in the second equation refers to a first-order Taylor approximation near $x_t = \overline{x}$.

Cost minimization requires that firms choose x_t to minimize the price of labor *per efficiency unit*, $w_t(x_t)/x_t$. This, in turn, requires a choice for x_t that yields an elasticity of one for $w_t(x_t)$ with respect to x_t. In our notation, this requires that $B'(0) = w_t(\overline{x})/\overline{x}$. Making this substitution in the equation above yields that $w_t(x_t)$ approximately equals $w_t(\overline{x})[x_t/\overline{x}]$, or

(11) $$\hat{w}_t(x_t) \approx \hat{w}_t(\overline{x}) + \hat{x}_t,$$

where a circumflex over a variable denotes the deviation of the natural log of that variable from its longer-run path. (We define this longer-run path empirically by an H-P filter--see the appendix, part 5).

But applying productivity accounting to equation (9), note

(12) $$\hat{x}_t = \frac{1}{\alpha+\nu}[\frac{1}{\gamma}\hat{y}_t - \hat{\theta}_t - \alpha\hat{h}_t - \nu\hat{l}_t - (1-\alpha-\nu)\hat{k}_t].$$

If we assume that high-frequency fluctuations in θ are negligible, then combining these two equations yields our alternative wage measure:

(13) $$\hat{w}_t(x_t) \approx \hat{w}_t(\overline{x}) + \frac{1}{\alpha+\nu}\left[\frac{1}{\gamma}\hat{y}_t - \alpha\hat{h}_t - \nu\hat{l}_t - (1-\alpha-\nu)\hat{k}_t\right].$$

Cyclical (H-P filtered) movements in TFP are interpreted as reflecting either increasing returns to scale or varying effort. Therefore, we augment average hourly earnings to capture varying effort simply by adding TFP movements, to the extent those movements are not attributable to increasing returns, scaled by $1/(\alpha+\nu)$. This equation can be written alternatively as

(14) $$\hat{w}_t(x_t) \approx \hat{w}_t(\overline{x}) + \frac{1}{\alpha+\nu}\left[\widehat{\mathrm{TFP}}_t - \frac{\gamma-1}{\gamma}\hat{y}_t\right],$$

where $\widehat{\mathrm{TFP}}_t = \hat{y}_t - \alpha\hat{h}_t - \nu\hat{l}_t - (1-\alpha-\nu)\hat{k}_t$.

We estimate a value for γ based on explaining the time-series behavior of inventories. Given that estimate for γ, we can then judge the extent to which the procyclical behavior of factor productivity reflects increasing returns or procyclical factor utilization.

III. Results

A. The Behavior of Inventories

We begin by examining the behavior of the ratio of sales to stock available for sale 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 (David A. Belsley, 1969).[15] 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.[16]

Figure 2 presents the ratio 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 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, Column 2 of Table 1 presents correlations between a conditional expectation of s_t/a_t and output. The expectation is conditioned on a set of variables $\mathcal{Z}_t$ and $\mathcal{Z}_{t-1}$, where $\mathcal{Z}_t = \{\ln(a_t)$, s_{t-1}/a_{t-1}, $\ln(y_t)$, $\ln(p_{t-1}/p_{t-2})$, R_t, $\ln(w_t/w_{t-1})$, $\ln(\boldsymbol{\omega_t/\omega_{t-1}})$, $\ln(n_t/y_t)$, $\ln(n_{t-1}/y_{t-1})$, $\ln(\mathrm{TFP}_t)$, $\ln(\mathrm{TFP}_{t-1})\}$. Price p_t is measured by the industry's monthly Producer Price Index, and R_t refers to the nominal interest

---

[15] 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.

[16] 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.

rate measured by the 90-day bankers' acceptance rate. Replacing sales with forecasted sales yields even larger correlations, ranging from .52 to .88.[17]

We want to stress that the strong tendency for 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 s_t/a_t requires that marginal cost is temporarily high in booms or that the price-marginal cost 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 (6). If we assume the two variables in this equation are conditionally distributed jointly lognormal, then (6) can be written[18]

$$E\left(\phi m \frac{s_t}{a_t} + \ln\left(\frac{\beta_t p_t}{p_{t-1}}\right) \,\middle|\, \mathcal{I}_{t-1},\ s_{t-1},\ p_{t-1}\right) + \kappa \approx 0, \tag{15}$$

where β_t reflects the nominal interest rate from $t-1$ to t measured by the 90-day bankers'

---

[17] Data sources for hours, wages, and TFP are described in part 4 of the appendix. All variables are H-P filtered as described in part 5 of the appendix. We also first differenced the series, looking at the correlation of the changes in the ratios 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 s_t/a_t yields even higher correlations, ranging from 0.57 to 0.86.)

[18] This approximation is arbitrarily good for small values for the real interest rate r and for the ratio $m\phi s/a$. In steady-state the ratio $m\phi s/a$ equals r plus the monthly storage rate. So we would argue this is a small fraction on the order of 0.02.

acceptance rate as well as a one percent monthly storage cost. The constant term κ reflects covariances between the random variables. Equation (15) implies we should see a strong negative relation between expectations of the two variables s_t/a_t and $\ln(\beta_t p_t/p_{t-1})$.

We first report, by industry, the correlation of $E(\ln(\beta_t p_t/p_{t-1})\,|\,\mathcal{I}_{t-1},\ s_{t-1},\ p_{t-1})$ with output. This expectation of $\ln(\beta_t p_t/p_{t-1})$ is constructed based on the variable sets $\mathcal{Z}_{t-1}$ and $\mathcal{Z}_{t-2}$, defined directly above in Section C, plus the variables s_{t-1}/a_{t-1} and $\ln(p_{t-1}/p_{t-2})$, which are part of $\mathcal{Z}_t$. 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 s_t/a_t. The correlation of $E(\ln(\beta_t p_t/p_{t-1})\,|\,\mathcal{I}_{t-1},\ s_{t-1},\ p_{t-1})$ with $E(\ln(s_t/a_t)\,|\,\mathcal{I}_{t-1},\ s_{t-1},\ p_{t-1})$ appears in Table 2, Column 2. Again the correlation is positive, significant, and large for every industry, ranging from .34 to .72. For equation (15) to hold these variables need to be negatively correlated. Also, estimating (15) 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 s_t/a_t to be not merely acyclical, but actually countercyclical. Therefore we proceed by allowing the markup to vary, as in first-order condition (3). Again assuming variables in the first-order condition are conditionally distributed jointly lognormal, the equation can be written

$$(16)\qquad E\left(\frac{\phi m_t s_t}{a_t} + \ln\left(\frac{\beta_{t+1} c_{t+1}}{c_t}\right) \,\middle|\, \mathcal{I}_t\right) + \kappa \approx 0,$$

where κ reflects covariances between the random variables.

Before estimating (16), we report correlations of discounted growth in marginal cost, $E_t(\beta_{t+1}c_{t+1}/c_t)$, the markup, $E_t(m_t)$, and $E_t(m_t s_t/a_t)$, with detrended output and with $E_t(s_t/a_t)$. Approximating (16) around average values of m_t and s_t/a_t (denoted with bars) yields

$$(17) \qquad E\left(\phi\overline{m}\frac{s_t}{a_t} + \phi\left(\frac{\overline{s}}{a}\right)m_t + \ln\left(\frac{\beta_{t+1}c_{t+1}}{c_t}\right) \,\Big|\, \mathcal{I}_t\right) - \phi\overline{m}\left(\frac{\overline{s}}{a}\right) + \kappa \approx 0,$$

Thus the procyclicality of $E_t(s_t/a_t)$ requires countercyclical movements in the expectations of $\beta_{t+1}c_{t+1}/c_t$ or m_t. Marginal cost is given by equation (7), with α as defined as in (8). (This assumes zero pure profits. See part 2 of the appendix.) For this exercise we impose constant returns to scale ($\gamma = 1$). To obtain conditional expectations of the variables we again project onto the set of variables $\mathcal{Z}_t$ and $\mathcal{Z}_{t-1}$ described above.

The results, by industry and for each of the two measures of the price of labor, appear in Tables 3 and 4. Consider first the measure based simply on average hourly earnings, represented by the first three columns of each table. For every industry the growth in marginal cost is very significantly positively correlated with both output and $E_t(s_t/a_t)$. The correlations with output range from .45 to .82. The correlations with $E_t(s_t/a_t)$ range from .30 to .72. Markups, on the other hand, do not display a consistent pattern across industries. They are procyclical, and vary positively with $E_t(s_t/a_t)$, in apparel, lumber, and chemicals, whereas they are countercyclical, and vary negatively with $E_t(s_t/a_t)$, in tobacco, petroleum, and rubber. Taken together, these correlations do not bode well for the average hourly earnings-based measure of marginal cost: $E_t(s_t/a_t)$ fails to be consistently negatively related to expected growth in marginal cost or markups, as required by (17).

Next consider correlations that use the wage augmented for variations in worker effort as described by equation (14), assuming for now that $\gamma = 1$. These 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 dramatically, 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 despite the fact that $E_t(s_t/a_t)$ is strongly procyclical (Table 1) and expected growth in marginal cost is strongly countercyclical (Table 3, Column 3), the two variables are not systematically correlated with each other. Expected growth in marginal cost is actually positively correlated with $E_t(s_t/a_t)$ 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 $E_t(s_t/a_t)$, 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 $E_t(s_t/a_t)$ varies from $-.49$ to $-.79$, again excluding lumber where it is significantly positive.

Tables 3 and 4 additionally report the correlations of the composite term $E_t(m_t s_t/a_t)$, with detrended output and with $E_t(s_t/a_t)$. Focusing on the augmented wage measure, we see from Table 3 that $E_t(m_t s_t/a_t)$ is clearly countercyclical for every industry but lumber. Thus the markup is sufficiently countercyclical to *more than offset* the strong procyclical movements in s_t/a_t. In fact, we can see from Table 4 that, again with the exception of lumber, the composite $E_t(m_t s_t/a_t)$ is even negatively correlated with $E_t(s_t/a_t)$. The implication is that countercyclical movements in the markup are more than sufficient to explain the procyclicality of s_t/a_t.

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 the parameters ϕ and γ from the first-order condition (16). Bearing in mind that the two wage measures reflect polar assumptions regarding the interpretation of short-run productivity movements, we do not necessarily expect either measure to rationalize inventory behavior completely; but we can evaluate which one does so more successfully. The parameter estimate for γ also provides information on the slope of marginal cost, that is, the response of marginal cost to an increase in output holding input prices fixed. This is distinct from our discussion to this point, which has focused on the reduced form cyclical behavior of marginal cost.

Equation (16) contains explicitly the parameter ϕ and implicitly the returns to scale parameter γ through both c_t and m_t. Using (7) to substitute for c_t and c_{t+1} in (16), and using the definition of m_t, we get

$$(18) \qquad E\left(\phi\left(\frac{s_t}{a_t}\right)\left(\frac{p_t}{\beta_{t+1}[\boldsymbol{\lambda'\omega_{t+1}} + \frac{1}{\gamma}\frac{w_{t+1}n_{t+1}}{\alpha y_{t+1}}]} - 1\right) + \ln\left(\frac{\beta_{t+1}[\boldsymbol{\lambda'\omega_{t+1}} + \frac{1}{\gamma}\frac{w_{t+1}n_{t+1}}{\alpha y_{t+1}}]}{\boldsymbol{\lambda'\omega_t} + \frac{1}{\gamma}\frac{w_t n_t}{\alpha y_t}}\right) \,\middle|\, \mathcal{I}_t\right) + \kappa = 0$$

where α is measured as in equation (8). To facilitate detrending, we approximate the second part of this equation to obtain

$$(19) \quad E_t\left(\phi\left(\frac{s_t}{a_t}\right)\left(\frac{p_t}{\beta_{t+1}[\boldsymbol{\lambda'\omega_{t+1}} + \frac{1}{\gamma}\frac{w_{t+1}n_{t+1}}{\alpha y_{t+1}}]} - 1\right) + \ln\beta_{t+1} + \psi(\gamma)\ln\left(\frac{\boldsymbol{\lambda'\omega_{t+1}}}{\boldsymbol{\lambda'\omega_t}}\right) + (1-\psi(\gamma))\ln\left(\frac{w_{t+1}n_{t+1}/y_{t+1}}{w_t n_t/y_t}\right) \,\middle|\, \mathcal{I}_t\right) + \kappa = 0\,.$$

$\psi(\gamma) = \gamma(\boldsymbol{\lambda'\omega}/p)/(\zeta + (\gamma-1)\boldsymbol{\lambda'\omega}/p)$. Recall that $\boldsymbol{\lambda'\omega}/p$ is measured by an H-P trend. ζ denotes the sample average of $(s/a)/(1-(1-s/a)\tilde{\beta})$, where $\tilde{\beta} = (1-\delta)/(1+r)$ reflects discounting for a real rate of interest r and for a storage cost δ. (Again, see parts 1 and 2 of the Appendix for more details). We remove low-frequency movements from the variables as described in part 4 of the appendix. Note that with the alternative wage measure, γ also enters the estimated equation as part of the wage through the term $-[(\gamma-1)/\gamma]\hat{y}_{t+1}$ in equation (14). The expectation is again constructed by conditioning on the set of variables $\mathcal{Z}_t \cup \mathcal{Z}_{t-1}$, where $\mathcal{Z}_t = \{\ln(a_t), s_{t-1}/a_{t-1}, \ln(y_t), \ln(p_{t-1}/p_{t-2}), R_t, \ln(w_t/w_{t-1}), \ln(\boldsymbol{\lambda'\omega_t/\lambda'\omega_{t-1}}), \ln(n_t/y_t), \ln(\mathrm{TFP}_t)\}$. Thus including a constant term there are nineteen instrumental variables.

We first estimate (19) by nonlinear GMM[19] to obtain unconstrained estimates of γ and ϕ for each wage measure. But given values for the real interest rate, storage costs, and returns to scale γ, the first-order condition implies a particular value of ϕ in order for the implied steady-state value of s_t/a_t to be consistent with the average observed value of s_t/a_t for each industry. This constraint is

$$\phi = \frac{1-\tilde{\beta}}{(\frac{p}{c}-\tilde{\beta})(s/a)}\,. \tag{20}$$

This is described in detail in part 3 of the appendix, including how the static markup p/c can be related to the returns to scale γ and an industry profit rate. We therefore also estimate equation (19) imposing this constraint on ϕ as a function of γ.

Table 5 contains results using the wage measured as average hourly earnings, 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 $1/\gamma$. So by estimating an absurdly high absolute value for γ, the estimation is essentially zeroing out this measure of 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 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 very 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, and significantly so for three: apparel, chemicals, and petroleum. The estimates of returns to scale are more robust: Even where the

[19]We used the Model procedure in SAS, using the GMM option with the quadratic spectral kernel applied to the weighting matrix. The procedure first linearizes the equation around an initial set of parameter values and uses Two-Stage Least Squares to obtain an initial weighting matrix. It then iterates, updating both the parameter estimates and the weighting matrix at each iteration, until convergence.

constraint is rejected the two estimates of γ are very similar, and the one case in which the point estimates differ substantially (petroleum), the difference is not statistically significant.[20]

As we discuss momentarily, for many of the industries, the exceptions being chemicals and petroleum, the intertemporal substitution term $E_t(\ln(\beta_{t+1}c_{t+1}/c_t))$ is largely acyclical. If discounted marginal cost literally follows a random walk, then the parameter ϕ is not identified. This suggests focusing largely on the constrained estimates of ϕ. Furthermore, although we find that allowing for uncompensated fluctuations in factor utilization goes quite far in explaining the behavior of inventory investment, we would not argue that Table 6 reflects an exact or “true” measure 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 will be rather low if c_t is close to, though not literally, a random walk. Evidence that we have an imperfect measure of marginal cost may 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 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.

We have focused on implications for the cyclical behavior of marginal cost—both relative to price and relative to expected future marginal cost—that come from inventory behavior. Much of the inventory literature, however, has focused more narrowly on estimating cost function parameters and, in particular, the relationship between output and marginal cost holding input prices constant, which we refer to as the “slope of marginal cost.” We would argue that the broader cyclicality measure is more relevant both for inventory behavior and many broader questions about the nature of business cycles. The slope of marginal cost does not enter separately from overall marginal cost in the Euler equation; and for many questions

[20] 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, and are available on request. The primary difference is that the test statistics for overidentifying restrictions and for the constraint on ϕ more typically reject.

about the nature of cyclical fluctuations the distinction between internal and external convexity or diminishing returns is not germane.

Nonetheless, for the sake of comparison, we can look at the implications of our estimates for the slope of marginal cost arising from convexity of the production function. If we assume that capital is fixed in the short run, then marginal cost is upward sloping if and only if $\gamma(\alpha+\nu)<1$. The estimates of γ in Table 6 bear a close inverse relationship to each industry's total labor exponent $\alpha+\nu$, which is provided in the first column of the table. By this criterion only chemicals ($\gamma(\alpha+\nu)=$ 0.88) and petroleum (0.69) exhibit significantly upward sloping marginal cost. The other four industries have very close to flat marginal cost.[21] Given a relation between short-run marginal cost and output that is relatively flat, then the extent to which overall marginal cost (allowing for changes in input prices) is procyclical rests largely on the behavior of input prices, and in particular the shadow price of labor.

D. Cyclical Markups versus Intertemporal Substitution

Our approach of adding back short-run TFP movements to construct an effective wage explains the procyclical behavior of s_t/a_t by some combination of procyclical marginal cost (relative to discounted future marginal cost—i.e. intertemporal substitution) and countercyclical markups. Can we say which factor is more important? Recall that Table 4 reported correlations of $E_t(s_t/a_t)$ with the expected growth in marginal cost, with the expected markup, and with $E_t(m_t s_t/a_t)$, 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

[21] The tobacco and rubber industries display very slightly downward sloping marginal cost, with $\gamma(\alpha+\nu)$ estimated at 1.02 in tobacco and 1.04 in rubber. In a model where inventories are held only to minimize costs a lack of short-run diminishing returns to labor can lead to failure of the second-order condition that accompanies first-order condition (3) for optimizing. This is not the case for our model. For $\phi<1$, there are diminishing returns to the available stock, a_t, in generating sales. This provides an incentive to smooth the stock available, and therefore production as well, even if there is no direct cost motive for smoothing production. In fact, our estimate for ϕ is less than 0.5 for each of the six industries, implying considerable diminishing returns in increasing the stock a_t. Related to this, the second-order condition for an optimum is satisfied for each of our industries based on the estimates in Table 6.

markup very countercyclical (except for the lumber industry), and not through the intertemporal cost term.

This conclusion is strengthened when we incorporate estimates for returns to scale. Using the estimates of γ from Table 6, Figure 4 presents the implied markup together with the ratio s_t/a_t for each industry. The movements in the markups are highly countercyclical (except for lumber) and quantitatively important. Several empirical papers have examined the cyclicality of markups. (Julio J. Rotemberg and Michael Woodford, 1999, 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. Figure 3 showed that the large shifts in price markups in tobacco in the 1980s and 1990s were accompanied by opposite movements in the ratio s_t/a_t as predicted by the model. Figure 4 shows that, more generally, most of the striking shifts in s_t/a_t that occurred in these six industries are associated with large opposite movements in the markup. It also highlights by comparison the relative unimportance of intertemporal movements in marginal cost in explaining inventory holdings.

In fact, we can say more. Table 7 presents correlations for the terms: $E_t(\beta_{t+1}\text{c}_{t+1}/\text{c}_t)$, $E_t(m_t)$, and $E_t(m_t s_t/a_t)$ with detrended output and with $E_t(s_t/a_t)$. This parallels Tables 3 and 4, except now the cost and markup terms are constructed using returns to scale as estimated in Table 6. In contrast to results under constant returns, the intertemporal substitution factor is now significantly positively correlated with output and with $E_t(s_t/a_t)$ in each of the six industries. By itself this would push s_t/a_t in the direction of being *countercyclical*; the movements in the markup have to *more than offset* the behavior of intertemporal cost in order to generate procyclical s_t/a_t. As under constant returns, with the exception of lumber, the anticipated markup is very countercyclical and significantly negatively correlated with $E_t(s_t/a_t)$, though the magnitudes are now somewhat smaller.[22] Table 7 also shows that

---

[22] In the case of lumber, even though both the expected markup and the intertemporal cost terms are slightly procyclical, they are negatively related to each other, as required by the model.

$E_t(m_t s_t/a_t)$ is negatively related to both detrended output and with $E_t(s_t/a_t)$ in all industries but lumber. This indicates that forecastable movements in m_t do indeed more than offset the cyclical behavior of s_t/a_t. Intertemporal substitution, due to temporarily high marginal cost during expansions, does not explain why a_t fails to keep pace with fluctuations in $E_t(s_t)$.

Although intertemporal substitution fails to play a key role, the alternative marginal cost measure that allows for cyclical work effort is still crucial in explaining inventory behavior. Allowing for the impact of cyclical work effort on the shadow cost of labor sufficiently alters our measure of marginal cost to enable the model to account for the procyclical behavior of s_t/a_t. While it does not make marginal cost procyclical, in the sense of being transitorily high at business cycle peaks, it does moves it in that direction. More importantly, it makes marginal cost procyclical relative to the price of output.

IV. Conclusions

Evidence from cross-sectional and low-frequency time-series 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. This behavior requires that during booms firms exhibit either high marginal cost relative to discounted future marginal cost (prompting intertemporal substitution) or low price markups.

Measures of marginal cost based on measured prices and productivity fail to explain this behavior because factor productivity rises during expansions relative to input prices. We show that the cyclical patterns of inventory holdings can be rationalized by interpreting fluctuations in labor productivity as arising primarily from mismeasured cyclical utilization of labor, the cost of which is internalized by firms but not contemporaneously reflected in measured average hourly earnings. Our view that procyclical factor utilization accounts for the inventory puzzle is consistent with other evidence that factors are worked more intensively in booms (for example, Ben S. Bernanke and Martin L. Parkinson, 1991, Matthew D. Shapiro, 1993, Bils and Jang-Ok Cho, 1994, Craig Burnside et al. 1995, and Jordi Galí, 1999).

It turns out, however, that it is not intertemporal substitution that accounts for the cyclical behavior of inventories. The standard story that firms deviate from a fixed target inventory-sales ratio because of transitory changes in marginal cost is not borne out by our analysis. Instead, what drives inventory behavior is primarily countercyclical markups, which have the effect of changing the target ratio. Thus the failure of inventories to keep pace with shipments is mirrored by the failure of price to keep pace with marginal cost.

In aggregate, observing a countercyclical markup is equivalent to observing procyclical *real marginal cost*, that is marginal cost that is procyclical relative to a general price deflator. What we see in the industry-level data is consistent with the following picture of the aggregate economy: An aggregate expansion in output is associated with an increase in real marginal cost. This rise in real marginal cost emanates not from diminishing returns to labor in the production function but from a higher shadow cost of labor. For a persistent increase in output, however, this does not justify predicting a negative growth rate for real marginal cost (relative to real interest rates) as needed to give rise to intertemporal substitution. For our model, a rise in real marginal cost, or equivalently a drop in the markup, directly reduces the value of inventory holdings by reducing the valuation of sales generated by those inventories. Therefore, a persistent rise in real marginal cost, absent intertemporal substitution, creates a persistent reduction in inventory holdings relative to expected shipments, as in Figure 1.

In recent years a number of papers have attempted to explain why firms might cut price markups during expansions. Rotemberg and Woodford (1995) survey many of these. As outlined by Rotemberg and Woodford, among others, such pricing can dramatically exacerbate cyclical fluctuations by reducing the distortionary impact of price markups on employment and output during booms. Our results clearly support these efforts.

References

Basu, Susantu and Fernald, John G., "Returns to Scale in U.S. Production: Estimates and Implications." *Journal of Political Economy*, April 1997, *105*(2), pp. 249-283.

Becker, Gary S., "Human Capital, Effort, and the Sexual Division of Labor." *Journal of Labor Economics*, January 1985, *3*(1, Part 2), S33-S58.

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

Bernanke, Ben S., and Parkinson, Martin L., "Procyclical Labor Productivity and Competing Theories of the Business Cycle: Some Evidence from Interwar U.S. Manufacturing." *Journal of Political Economy*, June 1991, *99*(3), 439-459.

Bils, Mark, "The Cyclical Behavior of Marginal Cost and Price." *American Economic Review*, December 1987, *77*(5) pp. 838-855.

Bils, Mark and Jang-Ok Cho, "Cyclical Factor Utilization." *Journal of Monetary Economics*, April 1994, *33*(2), pp. 319-354.

Blanchard, Olivier J., "The Production and Inventory Behavior of the American Automobile Industry." *Journal of Political Economy*, June 1983, *91*(3), pp. 365-400.

Blinder, Alan S., "Can the Production Smoothing Model of Inventory Behavior be Saved?" *Quarterly Journal of Economics*, August 1986, *101*(3), pp. 421-454.

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

Christiano, Lawrence J., "Why Does Inventory Investment Fluctuate So Much?" *Journal of Monetary Economics*, March/May 1988, *21*(2/3), pp. 247-280.

Cooper, Russell. and Haltiwanger, John C., "Macroeconomic Implications of Production Bunching: Factor Demand Linkages." *Journal of Monetary Economics*, *30*(1), October 1992, pp. 107-128.

Durlauf, Steven N. and Maccini, Louis J., "Measuring Noise in Inventory Models," *Journal of Monetary Economics,* August 1995, *36*(1), pp. 65-89.

Eichenbaum, Martin S., "Some Empirical Evidence on the Production Level and Production Cost Smoothing Models of Inventory Investment." *American Economic Review*, September 1989, *79*(4), pp. 853-864.

Fair, Ray C., "The Production-Smoothing Model is Alive and Well." *Journal of Monetary Economics*, November 1989, *24*(3), pp. 353-370.

Feldstein, Martin S., and Auerbach, Alan J., "Inventory Behavior in Durable-Goods Manufacturing: The Target-Adjustment Model." *Brookings Papers on Economic Activity*, 1976, (2), pp. 351-404.

Galí, Jordi, "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations." *American Economic Review*, March 1999, *89*(1), pp. 249-271.

Gertler, Mark L., and Gilchrist, Simon, "Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms." *Quarterly Journal of Economics*, May 1994, *109*(2), pp. 309-340.

Hall, Robert E., "Employment Fluctuations and Wage Rigidity." *Brookings Papers on Economic Activity*, 1980, (1), pp. 93-123.

Hall, Robert, E., "Labor Demand, Labor Supply, and Employment Volatility," in Stanley Fischer, ed., *NBER Macroeconomics Annual 1991*. Cambridge, MA: MIT Press, 1991, pp. 17-54.

Howell, Craig, Congelio, Frank, and Yatsko, Ralph, "Pricing Practices for Tobacco Products, 1980-1994." *Monthly Labor Review*, December 1994, *117*(12), pp. 3-16.

Kahn, James A., “Inventories and the Volatility of Production.” *American Economic Review*, September 1987, *77*(4), pp. 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*, May 1992, *107*(2), pp. 481-510.

Kashyap, Anil K., Lamont, Owen, and Stein, Jeremy, “Credit Conditions and the Cyclical Behavior of Inventories.” *Quarterly Journal of Economics*, August 1994, *109*(3), pp. 565-592.

Kashyap, Anil K., and Wilcox, David W., “Production and Inventory Control at the General Motors Corporation in the 1920's and 1930's.” *American Economic Review*, June 1993, *83*(3), pp. 383-401.

Krane, Spencer D. and Braun, Steven N., “Production Smoothing Evidence from Physical Product Data.” *Journal of Political Economy*, June 1991, *99*(3), pp. 558-581.

Kydland, Fynn E. and Prescott, Edward C., “Time to Build and Aggregate Fluctuations.” *Econometrica*, November 1982, *50*(6), pp. 1345-1370.

Pindyck, Robert S., “Inventories and the Short-Run Dynamics of Commodity Prices.” *Rand Journal of Economics*, Spring 1994, *25*(1), pp. 141-159.

Ramey, Valerie A., “Inventories as Factors of Production and Economic Fluctuations.” *American Economic Review*, June 1989, *79*(3), 338-354.

Ramey, Valerie A., “Non-Convex Costs and the Behavior of Inventories.” *Journal of Political Economy*, April 1991, *99*(2), pp. 306-334..

Rotemberg, Julio J. and Woodford, Michael, “Dynamic General Equilibrium Models with Imperfectly Competitive Markets,” in Thomas Cooley, ed., *Frontiers of Business Cycle Research*. Princeton, NJ: Princeton University Press, 1995, pp. 243-293.

Rotemberg, Julio J. and Woodford, Michael, "The Cyclical Behavior of Prices and Costs." National Bureau of Economic Research (Cambridge, MA) Working Paper No. 6909, January 1999.

Shapiro, Matthew D., "Cyclical Productivity and the Workweek of Capital." *American Economic Review: Papers and Proceedings*, May 1993, *83*(2), pp. 229-233.

Solow, Robert M., "Some Evidence on the Short-Run Productivity Puzzle," in Jagdish N. Baghwati and Richard S. Eckaus, eds., *Development and Planning.* New York: St. Martin's, 1968, pp. 316-325.

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

West, Kenneth D., "A Note on the Econometric Use of Constant Dollar Inventory Series." *Economics Letters*, 1983, *13*(4), pp. 337-341.

West, Kenneth D., "A Variance Bounds Test of the Linear Quadratic Inventory Model." *Journal of Political Economy*, April 1986, *94*(2), pp. 374-401.

West, Kenneth D., "The Sources of Fluctuations in Aggregate Inventories and GNP." *Quarterly Journal of Economics*, November 1990, *105*(4), pp. 939-971.

Appendix

1. *The Cost of Materials*

We know of no monthly data on material price deflators by industry. We construct our own monthly price of materials index for each industry (ω_t) 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 per unit material input by imposing that our series $\lambda'\omega_t$ 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*

We choose to calibrate the production labor exponent α. Because we do not impose that price equals marginal cost, we cannot calibrate the parameter based simply on production labor's share of value added. Even if firms do not earn profits, price must exceed production's marginal cost to cover the holdings costs of inventories. Secondly, if there are increasing returns, this implies average cost exceeds marginal cost; so zero profits implies that price exceeds marginal costs. Thirdly, firms in principle may earn profits.

Average cost per unit of production, call it $\overline{c}$, equals

(A1)
$$\overline{c} = \boldsymbol{\lambda}'\boldsymbol{\omega} + \left(\frac{1}{\alpha}\right)\frac{wn}{y}.$$

Let Φ denote the present-discounted flow of revenue generated by each unit of production. Evaluated under a constant average probability of selling s/a, and for a constant rate of nominal price inflation and nominal interest rate, Φ is given by

(A2)
$$\Phi \;=\; \frac{s}{a}p \;+\; (1-\frac{s}{a})\tilde{\beta}\frac{s}{a}p \;+ (1-\frac{s}{a})^2\tilde{\beta}^2\frac{s}{a}p \;+ \ldots.. \;=\; \frac{(s/a)p}{1-(1-s/a)\tilde{\beta}}\,.$$

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

Let the present-discounted value of profits be equal to a fraction π of costs. This requires that Φ be equal to (1+π)$\overline{c}$, or substituting from equations (A1) and (A2)

(A3)
$$\frac{(s/a)p}{1-(1-s/a)\tilde{\beta}} = (1+\pi)\Big[\boldsymbol{\lambda}'\boldsymbol{\omega} + \Big(\frac{1}{\alpha}\Big)\frac{wn}{y}\Big].$$

Rearranging for α yields

(A4)
$$\alpha = \frac{\frac{wn}{py}}{\frac{\zeta}{(1+\pi)} - \frac{\boldsymbol{\lambda}'\boldsymbol{\omega}}{p}}.$$

where $\zeta \equiv (s/a)/(1-(1-s/a)\tilde{\beta}) \leq 1$.

Note that in the absence of production to stock (i.e. $s/a = 1$) and with π equal to zero, α simply equals production labor's share of value added. More generally, labor's share understates α, due to the larger average markup necessary to make up for the cost of holding inventories (as well as any profits). We measure ζ by its sample average, where $\tilde{\beta}$ is the average value of $(\beta_{t+1}p_{t+1})/p_t$ assuming a monthly storage cost of one percent and a nominal interest rate measured by the 90-day bankers' acceptance rate. Thus α is directly related to observables except for the profit rate π. For the bulk of our estimation we assume that the steady-state level of economic profits is zero. A number of studies have suggested that profit rates in manufacturing are fairly close to zero. For example, Susanto Basu and John G. 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. We also explore robustness to profit rates as high as ten percent.

To allow for secular changes in factor and material cost shares we measure "steady state" $wn/(py)$ and $\boldsymbol{\lambda'\omega}/p$ respectively by H-P filters fit to series for production labor and material shares of gross output. Consequently, α varies at low frequencies as well. We do impose a constant industry value for $\alpha + \nu$, which reflects the sum of production and nonproduction labor shares, adjusted as in (A4).

3. *Constraining the value of* ϕ

In the estimation we consider the impact of constraining parameter ϕ to take a value consistent with an industry's long-run ratio of sales to stock available. We constrain ϕ as follows. Evaluating first-order condition (3) at a steady-state yields

$$\tilde{\beta}\left(1 + \frac{\phi m s}{a}\right) = 1.$$

Using the definition of the markup from (3) and rearranging

(A5) $$\phi = \frac{1 - \tilde{\beta}}{(\frac{p}{c} - \tilde{\beta})(s/a)} \,.$$

Substituting for *p* from equation (A3) and substituting $\boldsymbol{\lambda'\omega} + (\gamma\alpha)^{-1} wn/y$ for marginal cost yields

(A6) $$\frac{p}{c} = \frac{\gamma(1+\pi)}{\zeta \;+\; (\gamma - 1)(1+\pi)\frac{\boldsymbol{\lambda'\omega}}{p}}.$$

where, again, $\zeta \equiv (s/a)/(1-(1-s/a)\tilde{\beta}) \leq 1$. Substituting for p/c in (A5) from (A6) relates ϕ to the parameters γ, π, and the long-run values of s/a, $\tilde{\beta}$, and $\boldsymbol{\lambda'\omega}/p$.

In estimating we proceed as follows: (1) Choose a value for π; (2) set s/a and $\boldsymbol{\lambda'\omega}/p$ to industry sample averages (the latter allowed to drift according to an H-P filter) and $\tilde{\beta}$ to be consistent with an interest rate measured by the 90-day bankers' acceptance rate and a one percent monthly storage cost; (3) estimate γ, based partly on its influence on the constrained parameter ϕ. In the estimation reported in Tables 5 and 6 we impose a zero profit rate π for reasons discussed directly above. We did examine profit rates of 0.05 and 0.1 and found that the results were robust.

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

Monthly data for hours and wages for production workers are from the Bureau of Labor Statistics (BLS) Establishment Survey. For the augmented wage we compute TFP and adjust the wage according to equation (14) (except for the term involving γ, which is estimated). Output, for the purpose of measuring 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.

5. *Detrending Procedures*

Although the first-order condition (16) suggests that quantities such as s_t/a_t and $\ln(\beta_{t+1}c_{t+1}/c_t)$ 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 (H-P) 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 business cycle variation.) Because the Euler equation is nonlinear, it is necessary to detrend certain combinations of variables that are linear in the parameters. Specifically, we detrend the quantities $\beta_{t+1}\boldsymbol{\lambda'\omega_{t+1}}/(p_t s_t/a_t)$, $\beta_{t+1}w_{t+1}n_{t+1}/(\alpha y_{t+1}p_t s_t/a_t)$, $\ln(\beta_{t+1}w_{t+1}n_{t+1}/(w_t n_t))$, $\ln(\beta_{t+1}\boldsymbol{\omega_{t+1}/\omega_t})$, and $\ln(y_t)$, where w_t here refers either to average hourly earnings or to the augmented wage under the assumption $\gamma = 1$. Equation (19) can be expressed in terms of these variables multiplied by parameters or by functions of parameters.

We also use the same filter on ln(TFP) in constructing the augmented wage, though here the purpose is different. Our assumption is that low-frequency movements in ln(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: The Cyclicality of s_t/a_t in Manufacturing

Industry	Correlation of $\ln(y_t)$ with s_t/a_t	Correlation of $\ln(y_t)$ with $E\left(\frac{s_t}{a_t} \mid \mathcal{I}_t\right)$
Tobacco	.663	.854
Apparel	.484	.567
Lumber	.644	.723
Chemicals	.837	.880
Petroleum	.455	.516
Rubber	.791	.853

Note: The sample is 1959.1-1997.9. s_t/a_t is the ratio of sales to the stock available for sale; y_t is output. All correlations have p-values < 0.01.

Table 2: The Constant Markup Assumption

Industry	Correlation of $E\left(\ln\left(\frac{\beta_t p_t}{p_{t\text{-}1}}\right) \,\middle|\, \mathcal{I}_{t-1}, s_{t-1}, p_{t-1}\right)$ with	
	$\ln(y_t)$	$E\left(\frac{s_t}{a_t} \,\middle|\, \mathcal{I}_{t-1}, s_t, p_t\right)$
Tobacco	0.256	0.415
Apparel	0.367	0.351
Lumber	0.190	0.335
Chemicals	0.570	0.722
Petroleum	0.270	0.579
Rubber	0.233	0.449

†Note: The sample is 1959.1-1996.12. s_t/a_t is the ratio of sales to the stock available for sale; $\beta_t p_t/p_{t-1}$ is the discounted growth in output price from t-1 to t.
All correlations have p-values < 0.01.

Table 3: Cyclicality of Key Variables Relative to ln(y_t)[†]

	Correlation of ln(y_t) with					
	Average hourly earnings-based			Augmented wage-based		
	$E_t\left(\ln\left(\frac{\beta_{t+1}c_{t+1}}{c_t}\right)\right)$	$E_t(m_t)$	$E_t(\frac{s_t}{a_t}m_t)$	$E_t\left(\ln\left(\frac{\beta_{t+1}c_{t+1}}{c_t}\right)\right)$	$E_t(m_t)$	$E_t(\frac{s_t}{a_t}m_t)$
Tobacco	0.826	-0.257	-0.251	-0.880	-0.403	-0.434
Apparel	0.756	0.164	0.227	-0.348	-0.648	-0.679
Lumber	0.620	0.505	0.532	-0.270	0.030^*	-0.032^*
Chemicals	0.471	0.132	0.188	-0.106	-0.861	-0.875
Petroleum	0.563	-0.299	-0.333	0.479	-0.500	-0.548
Rubber	0.464	-0.202	-0.238	-0.284	-0.834	-0.865

[†]Note: All correlations are of H-P detrended series assuming γ=1. The sample is 1959.1-1996.12. s_t/a_t is the ratio of sales to the stock available for sale; $\beta_{t+1}c_{t+1}/c_t$ is the discounted growth in marginal cost from t to $t+1$; m_t is the markup as defined in the text. $E_t(x_t)$ is short for $E(x_t \mid \mathcal{I}_t)$.

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

Table 4: Cyclicality of Key Variables Relative to $E_t(s_t/a_t)^{\dagger}$

	Correlation of $E_t(s_t/a_t)$ with					
	Average hourly earnings-based			Augmented wage-based		
	$E_t\left(\ln\left(\frac{\beta_{t+1}c_{t+1}}{c_t}\right)\right)$	$E_t(m_t)$	$E_t(\frac{s_t}{a_t}m_t)$	$E_t\left(\ln\left(\frac{\beta_{t+1}c_{t+1}}{c_t}\right)\right)$	$E_t(m_t)$	$E_t(\frac{s_t}{a_t}m_t)$
Tobacco	0.725	-0.353	-0.358	-0.717	-0.483	-0.521
Apparel	0.312	0.152	0.220	-0.009^{*}	-0.423	-0.416
Lumber	0.399	0.596	0.632	0.020^{*}	0.127	0.135
Chemicals	0.343	0.230	0.302	0.195	-0.773	-0.788
Petroleum	0.357	-0.126	-0.148	0.379	-0.238	-0.269
Rubber	0.433	-0.162	-0.184	0.066^{*}	-0.724	-0.743

†Note: All correlations are of H-P detrended series assuming γ=1. The sample is 1959.1-1996.12. s_t/a_t is the ratio of sales to the stock available for sale; $\beta_{t+1}c_{t+1}/c_t$ is the discounted growth in marginal cost from t to $t+1$; m_t is the markup as defined in the text. $E_t(x_t)$ is short for $E(x_t \mid \mathcal{I}_t)$.

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

Table 5: GMM Estimates of Model Parameters with Wage equal to Average Hourly Earnings†

Industry	ϕ	$\gamma - 1$	*J*-statistic	*D-W* statistic
Tobacco	-0.070	74.7	27.0	2.30
	(0.042)	(34.3)		
	0.011*	101.3	27.2	2.35
		(55.9)		
Apparel	0.026	-59.7	31.3	1.30
	(0.021)	(33.1)		
	0.020*	-69.8	31.2	1.31
		(45.0)		
Lumber**	0.005	179.7	39.1	0.90
	(0.025)	(959.0)		
	0.024*	23.5	39.5	0.87
		(14.7)		
Chemicals	-0.068	15.2	27.8	1.17
	(0.028)	(6.5)		
	0.018*	37.5	31.5	1.09
		(31.6)		
Petroleum	0.524	1.49	24.4	1.40
	(0.061)	(0.78)		
	0.494*	0.091	22.2	1.33
		(0.028)		
Rubber**	-0.043	-326.6	25.6	1.00
	(0.016)	(1051.0)		
	0.020*	-111.0	20.9	0.93
		(102.9)		

†The sample is 1959.1 to 1996.12. Standard errors are in parentheses. The 0.05 critical value for the J-statistic is 26.3 when ϕ is estimated $(\chi^2(16))$, 27.6 in the restricted case $(\chi^2(17))$.
*Constrained based on estimate of γ.
**Constraint on ϕ and γ is rejected with a 0.05 critical value.

Table 6: GMM Estimates of Model Parameters with Augmented Wage†

Industry $(\alpha+\nu)$††	ϕ	$\gamma - 1$	*J*-stat.	*D-W* stat.
Tobacco	-0.021	1.932	21.3	2.34
	(0.043)	(0.134)		
(0.349)				
	0.023*	1.988	23.3	2.36
		(0.136)		
Apparel	0.180	0.188	45.4	1.85
	(0.056)	(0.025)		
(0.832)				
	0.205*	0.175	45.5	1.84
		(0.023)		
Lumber**	0.033	0.395	38.8	0.92
	(0.027)	(0.051)		
(0.660)				
	0.106*	0.408	36.4	0.84
		(0.051)		
Chemicals	0.076	0.395	35.7	1.41
	(0.029)	(0.047)		
(0.621)				
	0.094*	0.416	37.0	1.38
		(0.043)		
Petroleum	0.325	-0.181	21.1	1.45
	(0.101)	(0.174)		
(0.635)				
	0.486*	0.094	23.6	1.32
		(0.021)		
Rubber**	-0.049	0.186	37.4	1.80
	(0.055)	(0.039)		
(0.820)				
	0.161*	0.246	41.5	1.72
		(0.030)		

†The sample is 1959.1 to 1996.12. Standard errors are in parentheses. The 0.05 critical value for the J-statistic is 26.3 when ϕ is estimated $(\chi^2(16))$, 27.6 in the restricted case $(\chi^2(17))$.
††$\gamma(\alpha+\nu) < 1$ indicates short-run diminishing returns to scale. We measure $\alpha+\nu$ directly for each industry (see part 2 of the Appendix).
*Constrained based on estimate of γ
**Constraint on ϕ and γ is rejected with a 0.05 critical value.

Table 7: The Relative Importance of the Markup and Intertemporal Substitution in Accounting for Inventory Behavior*

	Correlation of					
	$\ln(y_t)$ with			$E_t(s_t/a_t)$ with		
	$E_t\left(\ln\left(\frac{\beta_{t+1}c_{t+1}}{c_t}\right)\right)$	$E_t(m_t)$	$E_t(\frac{s_t}{a_t}m_t)$	$E_t\left(\ln\left(\frac{\beta_{t+1}c_{t+1}}{c_t}\right)\right)$	$E_t(m_t)$	$E_t(\frac{s_t}{a_t}m_t)$
Tobacco	0.110	− 0.212	− 0.137	0.172	− 0.276	− 0.239
Apparel	0.154	− 0.371	− 0.406	0.197	− 0.234	− 0.092
Lumber	0.118	0.186	0.277	0.217	0.274	0.408
Chemicals	0.386	− 0.659	− 0.711	0.653	− 0.595	− 0.601
Petroleum	0.497	− 0.372	− 0.521	0.381	− 0.186	− 0.263
Rubber	0.172	− 0.574	− 0.649	0.430	− 0.486	− 0.456

*Note: All correlations are of H-P detrended series. $E_t(x_t)$ is short for $E(x_t \mid \mathcal{I}_t)$.

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

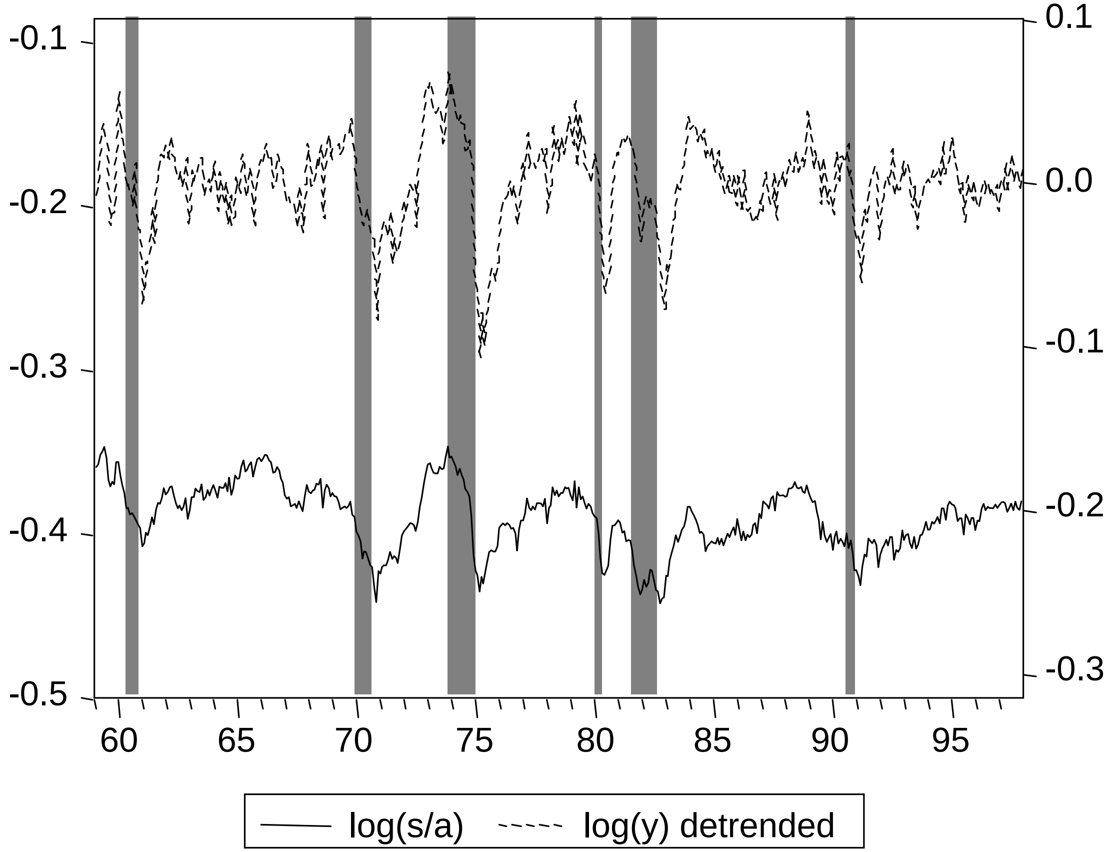

Note: Shaded areas indicate recessions. log(y) detrended with H-P filter.
y=output, s=sales, a=i+y, i=beginning finished goods inventory stock.

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

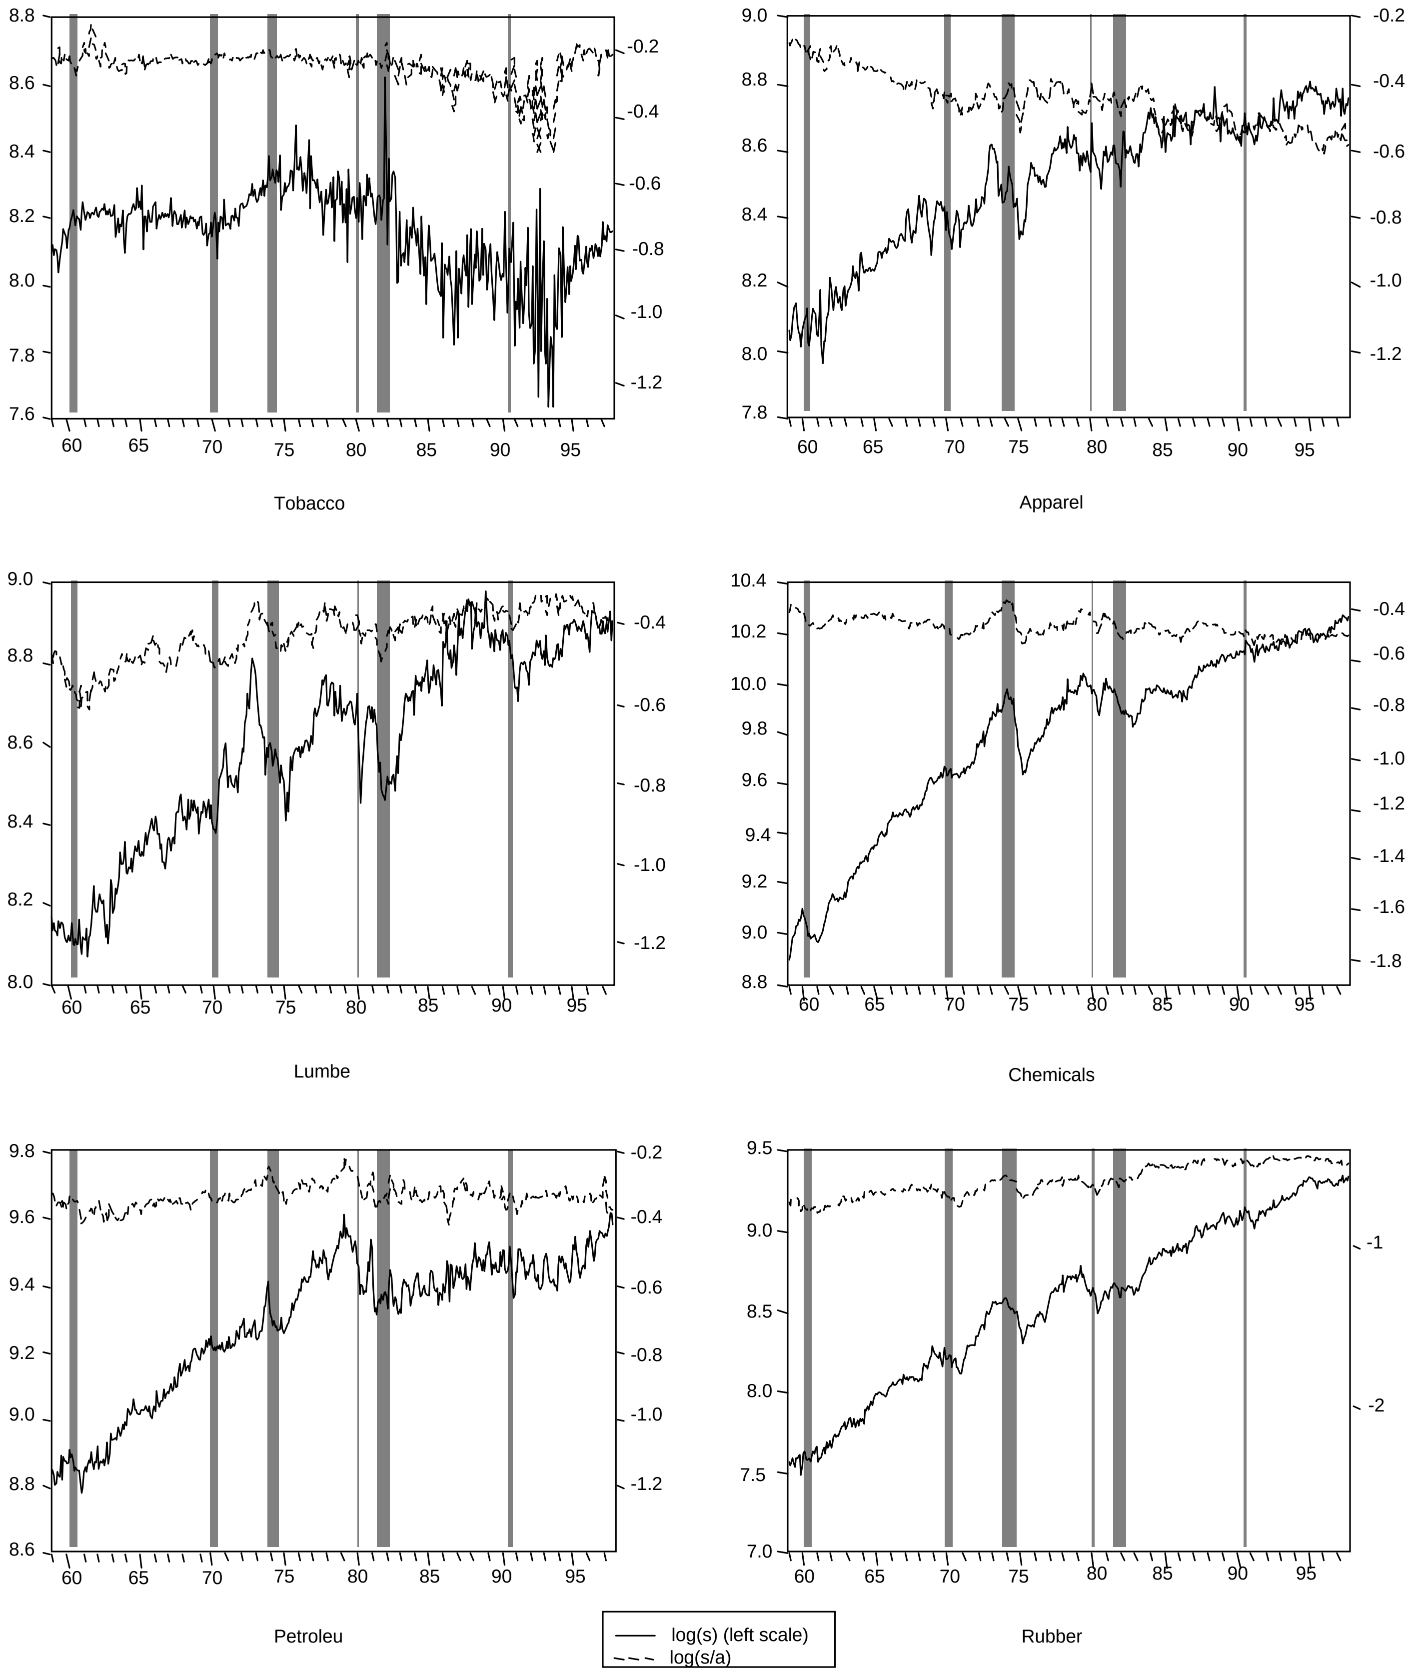

Shaded areas indicate recessions

s=sales, a=i+y, i=beginning finished goods inventory stock.

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

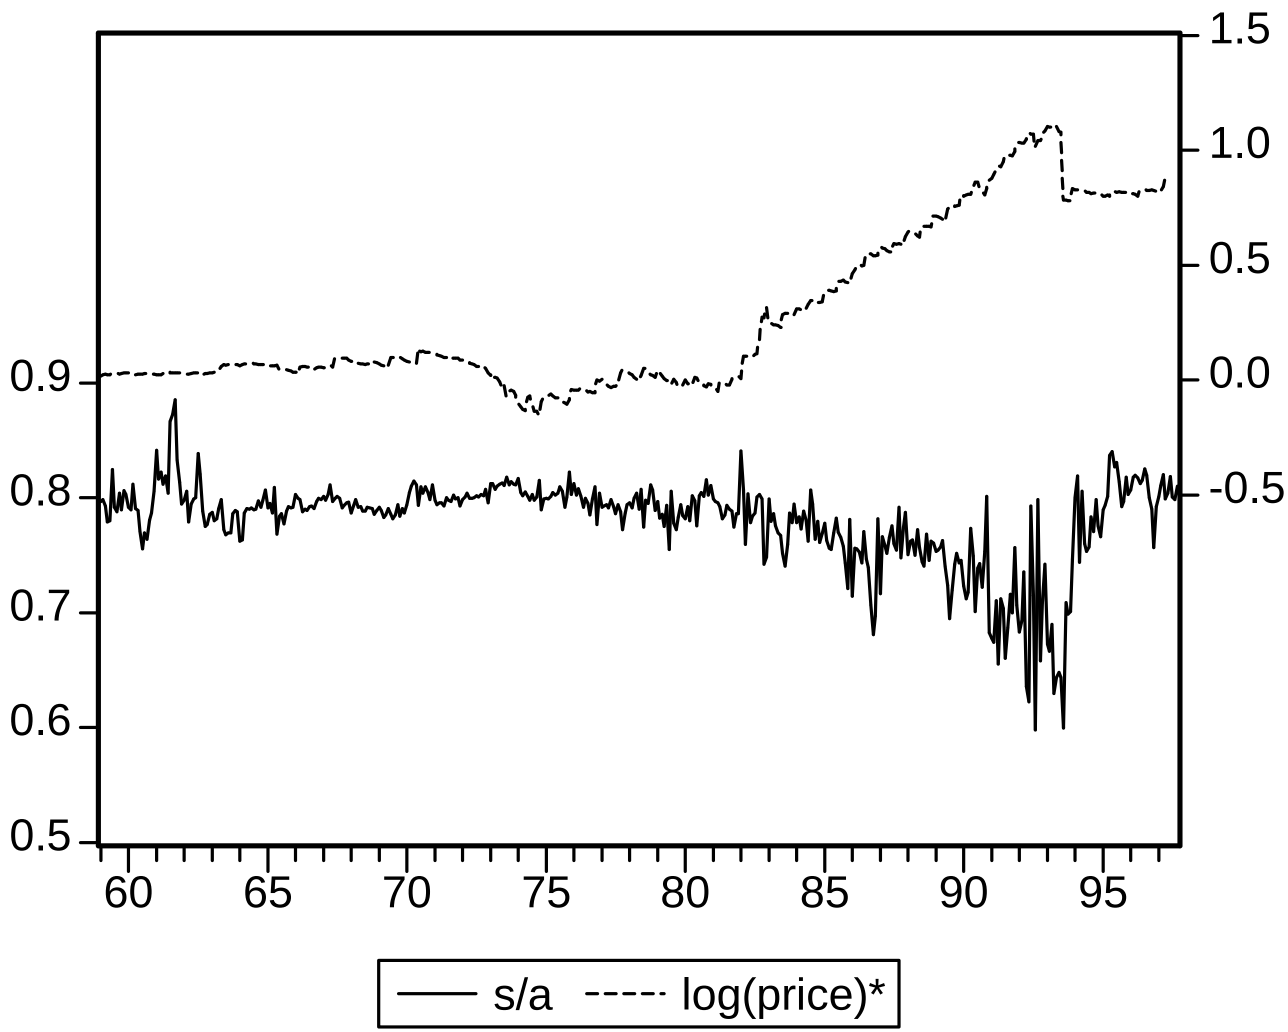

*Tobacco products price deflated by the general Producer Price Index

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

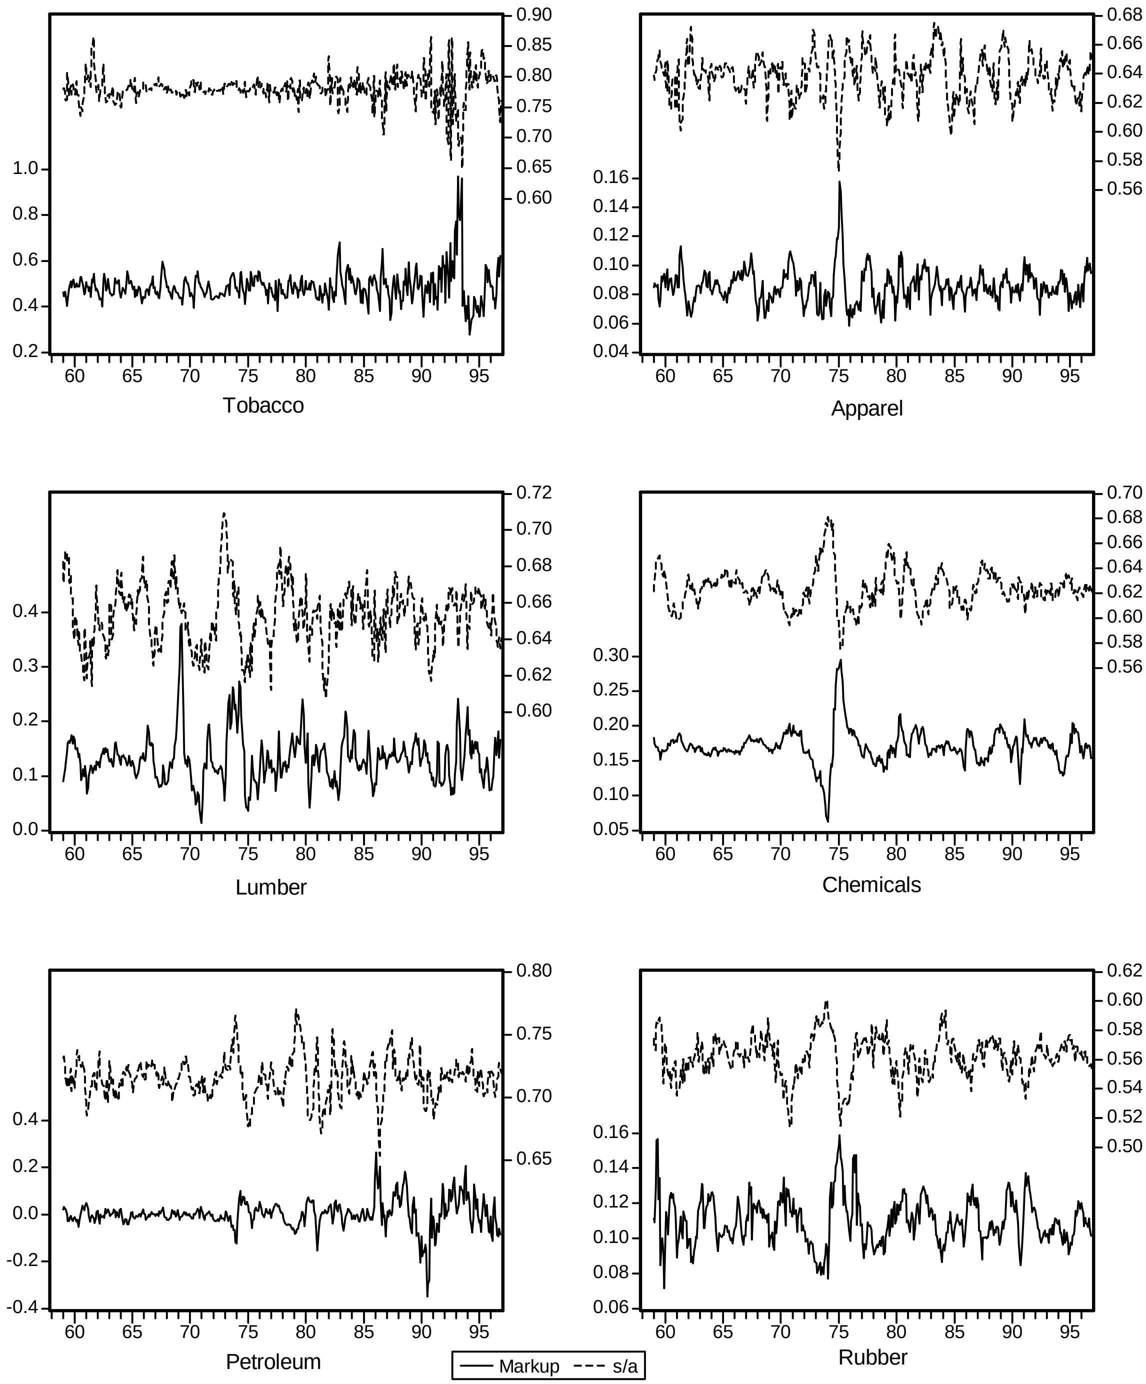

s=sales, a=i+y, i=beginning finished goods inventory stock.
The markup is price in t relative to discounted marginal cost in t+1, minus 1