Center for Research on Economic and Social Theory

and

Department of Economics

Working Paper Series

Intermediate Goods and Business Cycles: Implications for Productivity and Welfare

Susanto Basu

September, 1993
Number 93-23

DEPARTMENT OF ECONOMICS
University of Michigan
Ann Arbor, Michigan 48109-1220

Revised September 1993

INTERMEDIATE GOODS AND BUSINESS CYCLES: IMPLICATIONS FOR PRODUCTIVITY AND WELFARE

Susanto Basu

Department of Economics
University of Michigan
611 Tappan Street
Ann Arbor, MI 48109-1220

313-764-5359

internet: sbasu@econ.lsa.umich.edu

This is a revised version of Chapter 1 of my unpublished 1992 Harvard dissertation. I would like to thank, without implicating, Larry Ball, Bob Barsky, Steve Cecchetti, Brad DeLong, John Fernald, John Leahy, Valerie Ramey, David Romer, the members of the Harvard Macro Lunch Group, and seminar participants at various universities for helpful comments and advice. I would especially like to thank my advisor Greg Mankiw for his unstinting help, advice, and support. Financial support from a National Science Foundation Graduate Fellowship and fellowships from the Chiles Foundation, the Social Sciences Research Council, and the National Bureau of Economic Research are gratefully acknowledged.

ABSTRACT

This paper studies a business-cycle model with imperfect competition where intermediate goods are used in production. It is an example of a class of models in which markups are countercyclical. One major result is that in this setting, demand-driven output movements cause productivity to be procyclical. The paper studies a number of theoretical and empirical implications of this source of productivity fluctuations. In a subset of models, countercyclical markups result from assuming that there are fixed costs of changing nominal prices. The paper shows that modeling the use of intermediate goods in this type of model greatly expands the extent of price rigidity, leading to larger welfare losses from business cycles.

This paper studies a business-cycle model with imperfect competition where intermediate goods are used in production. It is an example of a class of models in which markups are countercyclical. One major result is that in this setting, demand-driven output movements cause productivity to be procyclical. The paper studies a number of theoretical and empirical implications of this source of productivity fluctuations. In a subset of models, countercyclical markups result from assuming that there are fixed costs of changing nominal prices. The paper shows that modeling the use of intermediate goods in this type of model greatly expands the extent of price rigidity, leading to larger welfare losses from business cycles.

It is an old idea that an industrialized economy, with its greater interdependence and more roundabout production, is more subject to cyclical output fluctuations. The idea has been present at least since the work of Gardiner C. Means (1935), who presented evidence that different industries had very different patterns of price changes versus quantity changes in the Great Depression. Means showed that simple goods, such as agricultural products, declined heavily in price, while their quantity was almost unchanged. Complex manufactured goods, on the other hand, showed the opposite pattern, with small price changes and consequently huge declines in the quantity of sales. Crude manufactured goods fell somewhere in between. Means's suggestive evidence has led many to speculate on the relationship between output fluctuations and roundabout production; see, for example, Robert J. Gordon (1990).

Means was concerned with the comovement of output and prices during the Depression, but in recent years another stylized fact of business cycles — the procyclicality of productivity — has attracted greater interest. Early work on real business cycles ascribed measured fluctuations in productivity to actual changes in production technology. However, Robert E. Hall (1988) and Charles L. Evans (1992) have shown that productivity is correlated with variables that are exogenous with respect to technology. Hall (1988, 1990) explains cyclical productivity as a consequence of imperfect competition and increasing returns. Hall's explanations imply that sectoral changes in cyclical productivity should be a function only of changes in sectoral output. Ricardo Caballero and Richard K. Lyons (1990a, 1992) document, however, that changes in sectoral productivity are also correlated with aggregate output fluctuations. They interpret this finding as evidence for technological spillovers between sectors. Ben S. Bernanke and Martin L. Parkinson (1991), on the other hand, interpret a similar set of results for sectoral productivity during the Depression as evidence of labor hoarding.

Here I suggest a new mechanism that explains changes in productivity over the cycle. As in Hall (1988), the explanation relies on imperfect competition. The model differs from Hall's in that it takes account of intermediate goods in production, and generates countercyclical markups. Therefore the explanation for cyclical productivity is quite different from Hall's; in particular, even with constant returns in production, the "cost-based Solow residual" — which Hall (1990) shows is the right measure of total factor productivity under imperfect competition — is also procyclical. As a result of these differences, the model predicts that sectoral productivity should in fact be correlated with aggregate activity. I show that this result implies potentially substantial biases in Hall's estimates of the markup and the degree of returns to scale. The explanation I propose can be distinguished empirically from those of Caballero and Lyons and Bernanke and Parkinson. If the component of procyclical productivity that appears as an external effect were caused by technological externalities or cyclical factor utilization, this effect should be evident in both gross-output and value-added data. On the other hand, if the explanation I propose here is correct, external effects should appear in value-added data, but not in gross-output data. Performing this test with data from U.S. manufacturing industries supports the model I present here. Other empirical tests also support the predictions of the model.

These are the main results of the paper when it is broadly construed as an example of models of countercyclical markups. The particular model I present, however, generates countercyclical markups by assuming that firms face small costs of changing prices (menu costs) and therefore have rigid prices over some range of shocks. In this setting, modeling firms' use of intermediate goods in production also validates Means's conjecture. For parameter values taken from U.S. manufacturing, the model shows that the roundabout nature of production allows sticky-price models to explain much larger output fluctuations and more severe welfare losses than heretofore thought plausible.[1]

[1] See the discussions by Laurence Ball and David Romer (1989, 1990) of the early menu-cost models of N. Gregory Mankiw (1985) and George Akerlof and Janet Yellen (1985).

I model the use of intermediate goods in an input-output production structure, so all firms use intermediate inputs in production. If price changes are costly, they are presumably costly for all firms — including those producing intermediate goods — so intermediate goods should also have rigid prices. Intermediate goods, however, act as a multiplier for price stickiness: a little price rigidity at the level of an individual firm leads to a large degree of economy-wide price inflexibility.

The reason is straightforward. The representative firm is connected by a complex input-output relationship to many other firms. But each firm cares only about the ratio of its price to its marginal cost of production; an increase in aggregate demand induces a firm to raise its price only to the extent that its profits are squeezed between a fixed output price and rising input costs. With intermediate goods in production, the increase in firms' costs depends on whether other firms raise prices. So in response to a demand shock, each firm simply "waits by the mailbox" to see if other firms have raised their prices. If other prices go up, then the firm will also be obliged to raise its own price. But if all firms follow this reasonable strategy, no input prices — and hence no output prices — will increase. In the limit as intermediate goods become the only variable input to production, firms never change prices and output is determined solely by aggregate demand.[2]

The assumption of sticky intermediate goods prices is supported by the evidence. The most detailed analysis of nominal rigidities studied intermediate goods. George J. Stigler and James K. Kindahl (1970) collected data on actual transaction prices for a large number of such products. Dennis W. Carlton's (1986) analysis of this data set showed that for some substances, particularly steel, paper, chemicals, stone, and glass products, prices can be rigid for long periods of time — in some cases years.[3]

---

[2] Therefore intermediate goods are a "real rigidity" in the sense in which the term is used by Ball and Romer (1990). However the real rigidities that they provide as examples are all auxiliary assumptions about the behavior of labor or product markets — e.g. efficiency wages and kinked demand curves — whose existence and extent are controversial. By contrast, usage of intermediate goods is a widespread and easily documented feature of the production process in any modern economy.

[3] In fact, Carlton's results imply so much price rigidity that some (perhaps including Carlton) have refused to believe that these prices are allocative. Instead, they hold that long-term relationships between buyers and suppliers make the observed spot prices a bad indicator of the true shadow cost of intermediate inputs. If this were true, however, then there should be a strong positive correlation between the observed rigidity of input prices and the length of buyer-seller association. But in fact, Carlton finds a *negative* relationship between price rigidity and the length of association, making it unlikely that the "installment payment" interpretation of price rigidity is correct.

In the paper, I present evidence that intermediate goods prices are less procyclical than labor costs. This also is consistent with my hypothesis that prices of intermediate goods are relatively rigid.

There is a similarity between the conclusion of this part of the paper and that of Olivier J. Blanchard (1983), but any similarity in the models is more apparent than real. I model price stickiness as state-dependent and the production structure as following an input-output relationship, while Blanchard has time-dependent pricing and in-line production. In Blanchard's model, the degree of price stickiness is a function of the number of stages of processing only because price setting is assumed to be staggered along the chain of production. If the pricing decision in Blanchard's model were made state-dependent then, since the "first good" is made without intermediate goods, there would be no increase in price rigidity regardless of the number of stages of production. In my state-dependent model, price stickiness depends upon the use of intermediate inputs because the input-output structure of production ensures that all firms use intermediate goods.

Given this difference between the two models, one might ask whether production should in fact be modeled as an input-output process, or as an irreversible chain where goods move in only one direction down the stages of processing. While the "chain of production" seems plausible *prima facie*, it naturally leads one to ask whether in the real world there are empirically relevant "first goods," the ones produced without any intermediate inputs. Input-output studies certainly do not support the chain-of-production view; even the most detailed input-output tables show surprisingly few zeros.[4] Empirically, the biggest source of any industry's inputs is usually itself: that is, the diagonal entries of input-output matrices are almost always the largest elements of each column (see BEA, 1984). This seems to lend credence to the view of "roundabout" rather than "in-line" production.

The paper is organized into six sections. Section I presents a simple menu-cost model and uses it to demonstrate that, with intermediate goods, sticky-price models can explain larger output fluctuations. Section II investigates the issue of reasonable parameter values for the model. Section III shows how the model broadly construed — that is, independent of whether prices are sticky — can generate procyclical

---

[4] In its discussion of the 1977 input-output table, the BEA notes that the table "shows heavy interdependence among industries. Seventy-six of the [eighty-five] industries shown in the table required inputs of at least 40 commodities, and 52 industries required inputs of at least 50 commodities (1984, p. 50)."

productivity movements. Section IV shows how the results derived in the previous section affect Hall's estimates of the markup and the degree of returns to scale. Section V examines empirical evidence from U.S. manufacturing and asks whether the evidence is consistent with the predictions of the model. Section VI concludes.

I. The Model Narrowly Construed: Menu Costs

The model is based on Mankiw (1991). There is a continuum of goods, indexed on [0,1]. The representative consumer maximizes an utility function that is additively separable in goods, real balances, and leisure.

$$U = \frac{1}{1-\phi}\int_0^1 Q_{i,F}^{\,1-\phi}\,di + \log\left[\frac{M}{P}\right] - L \qquad (1)$$

where

$Q_{i,F}$ is the quantity of product i used for final consumption,

ϕ is the reciprocal of the elasticity of substitution between different products ($0 < \phi < 1$),

M is money demand (assumed equal to money supply by money market equilibrium),

P is the general price level, and

L is labor supply.

Money is put in the utility function as a shortcut for generating money demand. [5] I have assumed a constant disutility of labor. Note that the assumption of additive separability makes the quantity of any product consumed independent of the prices of all other products.

The price level, P is defined by

$$P = \left(\int_0^1 P_i^{(\phi-1)/\phi}\,di\right)^{\phi/(\phi-1)}$$

[5] The alternative is to simply write down Y=M/P, as in Ball and Romer (1989, 1990). Putting money in the utility function generates a similar aggregate demand relationship, except that the elasticity of aggregate demand with respect to real balances becomes $1/\phi$ rather than one.

The price level is, of course, homogeneous of degree one in all prices. The consumer maximizes (1) subject to a standard budget constraint; the first order conditions are derived in the Appendix.

The production side of the economy is composed of a continuum of monopolistic firms, each producing one variety of product. Each firm maximizes profits, given the production function

$$Q_i = L_i^{\alpha} I_i^{1-\alpha} \qquad (2)$$

where

$$I_i = \left(\int_0^1 I_{ki}^{\,1-\phi}\,dk\right)^{\frac{1}{(1-\phi)}},$$

L_i is the labor input of firm i, and I_{ki} is the quantity of the kth intermediate input used by firm i.

I assume that all goods can serve either as final outputs or as inputs for the production of other goods. There is, therefore, no distinction between firms producing manufactured inputs and those producing final goods — all firms produce for both markets. For simplicity, I also assume that firms' elasticity of substitution between manufactured inputs in production is the same as consumers' elasticity of substitution between goods in consumption.[6] Finally, I take the aggregate production function to be constant returns to scale and Cobb-Douglas, with the share of nonproduced inputs (here only labor) being α.

Under these conditions, each firm's profit-maximizing nominal price, P_i^*, is

$$P_i^* = \frac{1}{1-\phi}\left(kW^{\alpha}P^{1-\alpha}\right)$$
$$= \mu\left(kW^{\alpha}P^{1-\alpha}\right) \qquad (3)$$

where μ is the markup, W is the nominal wage, and k is an unimportant constant.[7] The intuition for (3) comes from the fact that the output price is set as a markup on marginal cost; since the production function is Cobb-Douglas, marginal cost is a geometric average of the wage and the overall price of intermediate goods, where the weights are the shares of the inputs in production.

[6] This simplification is inessential for any of the results below. Its only purpose is to ensure that firms face a constant elasticity of substitution demand curve for their output.
[7] See the Appendix for details.

It follows from the equality of the wage and the nominal money supply shown in the Appendix[8] that the optimal relative price for each firm i, p_i^*, is

$$p_i^* = \frac{P_i^*}{P} = k\mu \left(\frac{M}{P}\right)^{\alpha} \tag{4}$$

The important point to note about (4) is that the optimal relative price depends upon real balances raised to the power α, rather than to the power 1 as in the simple menu cost model of Mankiw (1991). Therefore the change in the optimal price following a monetary shock is α times the percentage change in money. Depending on the value of α, this can substantially reduce the loss to a firm that does not adjust its price after a monetary shock, relative to the case where the use of intermediate goods is not modeled. The intuition is straightforward. Aggregate demand is proportional to real money balances, so with fixed prices an increase in money raises output and the demand for labor. Workers are always on their labor supply curves, so the increased use of labor raises the real wage. Since firms set their optimal relative price as a markup on real marginal cost, the increase in the real wage raises the optimal price. But if intermediate goods are used in production, firms' marginal costs rise only in proportion to labor's share, α, since intermediate goods prices are fixed.

As Mankiw (1985) and Akerlof and Yellen (1985) point out, the loss to a monopolistic firm of not changing its price in response to a money shock is of second order in the change in the optimal price. For price stickiness to be a Nash equilibrium, it must be the case that the loss to each firm of not adjusting prices, *assuming that no other firm adjusts*, is less than the menu cost of changing prices. Each firm calculates the private cost of leaving prices unchanged under the Nash assumption that all other prices will remain at their current levels.[9] The calculations performed in this section show how this private cost of price rigidity varies as a function of the importance of intermediate goods in production.

---

[8] The relation W=M follows from the assumption of a constant disutility of labor. A constant disutility of labor makes the real wage less procyclical than it otherwise would be, and therefore understates the reduction in cyclicality of marginal cost from introducing intermediate goods. In Section III, it also makes markups less countercyclical than they would otherwise be. Thus, the assumption actually works against the results of the paper. It does, however, simplify the model considerably.

[9] There are actually many conjectural variations that firms could make. In general, each would conjecture that k% of all other firms in the economy would change their prices in response to the money shock. The standard assumption in the menu cost literature is implicitly that firms take k=0. So each firm contemplates a change in the price of its own product, assuming that no other firm will change its price.

To a second order approximation, the change in profit to a firm from not adjusting its price in response to a monetary shock is given by

$$\pi(p^*_{new}) - \pi(p^*_{old}) \cong \frac{1}{2}\pi''(p^*)(p^*_{new} - p^*_{old})^2 \tag{5}$$

where π'' is the second derivative of the profit function with respect to prices. The difference between p^*_{new} and p^*_{old} is of course proportional to the change in real balances.

We wish to see how the profit loss (5) varies with respect to α, the share of labor in total cost. As far as the change in the firm's optimal relative price is concerned, this is not difficult. From the discussion above, it is clear that the change is proportional to α. If the loss were based solely on the square of the change in the optimal price, it would be the case that the loss would diminish in proportion to α^2. However, there are also the changes in the π'' term to be considered. Unfortunately, the expression for the derivative of π'' with respect to α is positive. The two effects work against each other, so it is not possible to sign the derivative unambiguously. Therefore I present numerical evaluations of (5) for different values of α.

The numerical results are reported in Table 1. Table 1 gives the loss to a firm of not adjusting its price in response to a 1 percent money shock as a function of the two main parameters, α and ϕ. The loss is expressed as a fraction of firm profits. To facilitate comparison, the numbers in each column are normalized so that the profit loss in the base case ($\alpha = 1$) is 1. So each entry in Table 1 reports the quantity

$$\frac{(\pi^*-\pi)/\pi^*}{[(\pi^*-\pi)/\pi^*]_{|\alpha=1}} \tag{6}$$

where starred variables reflect optimal values and the denominator is always evaluated at $\alpha = 1$. To give some idea of the absolute size of the losses, the numbers in parentheses below each entry gives the loss expressed as a percentage of the firm's profits before the monetary shock — that is, the numerator of (6).

In either case, firms' losses diminish significantly as α decreases. The results of Table 1 conform to expectation. As normalized by firm output, profit losses fall by approximately a factor of α^2. The deviation from the loss being precisely α^2 comes from the fact that the π'' term also changes with α. This second effect is insignificant, however, showing up in the fourth decimal place if at all. The effect of the elasticity of demand, ϕ, is to change the levels of output and profits. Once the loss is normalized by profits, there is no further effect of the elasticity of demand on the relative profit loss (although it does matter for the absolute loss).

If we assume that the firm is indifferent between changing its price and leaving it constant for a one percent money shock at $\alpha = 1$, what is the percent change in money required to leave it similarly indifferent for lower α values? Since losses are basically proportional to the square of the money shock, the required money shock is given by the square root of the inverse of each entry in Table 1. Looking at the line for $\alpha = 0.9$, it is interesting to note that even with low use of intermediate goods in production, the allowable money shock jumps by more than 10 percent.[10]

The results of the next two sections suggest reasonable ranges for α and μ, given the share of intermediate inputs in revenue and econometric estimates of the markup. Computations similar to those in Table 1 are reported in the second and third lines of Table 2. Each column is computed for a different value of α which, as the next section shows, defines an implicit value of μ. The markups corresponding to these values of α are reported in the first line of the table. The second line of Table 2 gives the profit loss for a one percent money shock, normalized as before so that it is a fraction of the profit loss for $\alpha = 1$. The third line reports these losses in absolute terms, as percentages of profit.

The results indicate that the introduction of intermediate goods with sticky prices into a menu cost model is quantitatively important. I show later that the relevant columns of Table 2 are the last two on the right, corresponding to $\alpha = 0.2$ and $\alpha = 0.1$. These show profit losses to firms declining by a factor of 25 to 100. Thus, the size of the maximum shock to money for which non-adjustment is a Nash equilibrium jumps five- to ten-fold if one models the use of intermediate goods in production: menu-cost models can explain larger business cycles.

The welfare consequences are also immediate. Ball and Romer (1989, 1990) summarize welfare by examining the ratio, R, of the social cost of output fluctuations to the private cost. Since the social cost in this model comes solely from the disutility of fluctuations in consumption, it is unaffected by introducing

[10] The percentage change in the quantity of final goods produced is $1/\phi$ times the percent change in money, so the size of output fluctuations that can be explained by a menu cost of a given size can also be computed from Table 1.

intermediate goods. But as I have shown, for plausible parameter values the private cost of business cycles falls considerably. Therefore R increases by a factor of 25 to 100, so menu-cost models can also explain inefficient business cycles.

These results follow logically from the assumption of sticky intermediate (and final) goods prices. One might ask, however, whether this assumption is a reasonable description of the world: just how cyclical are the costs of intermediate inputs, particularly relative to the cost of labor? Before answering the question it is necessary to dispel a common misconception. One typically thinks of intermediate inputs as "materials" — raw commodities whose prices are known to be volatile and procyclical.[11] But raw materials and energy are actually only a small fraction of total intermediate input. In a modern economy, by far the largest share of these inputs is devoted to purchases of goods manufactured by other firms. This paper takes the same view as the National Income Accounts: intermediate goods are properly distinguished by use, not by type of good. Once one takes the correct input-output view of intermediate inputs, it is easy to believe that intermediate goods have relatively rigid prices. At the least, the assumption of rigid materials prices is no less reasonable than the assumption of rigid final goods prices: in many cases, materials *are* final goods.[12] Understanding the correct definition of intermediate inputs helps understand why Stigler and Kindahl (1970) found that intermediate goods prices were rigid for long periods of time. In Section V, I present evidence showing that intermediate goods prices are less procyclical than labor costs. This is further evidence in support of my hypothesis that prices of intermediate goods are relatively rigid.

II. Choosing Reasonable Parameter Values.

[11] This confusion results in large part from terminology. In the production function literature, inputs are classified as KLEM — capital, labor, energy, and materials. That literature takes the correct view of "materials" as all intermediate goods and services, but the word is confusing in this context because it leads one to naturally — but incorrectly — identify "materials" with unprocessed commodities.

[12] One commonly hears the claim that Blanchard (1983) and Kevin Murphy, Andrei Shleifer, and Robert Vishny (1989) have shown that intermediate input prices are more procyclical than final goods prices. Their conclusions are also a consequence of incorrectly equating intermediate goods with unprocessed goods.

How can the model be calibrated to judge what are reasonable values of α? In their examination of gross output in U.S. manufacturing, Dale W. Jorgenson, Frank Gollop, and Barbara Fraumeni (1987) find that the share of intermediate inputs in total manufacturing output is 50 percent or greater over the period 1947-1979. (Interestingly, the share of materials is even higher in the agriculture and service industries.) So a value of 0.50 seems conservative.

To use the figure cited above, we must derive a relationship between the revenue share of materials and α. The share of intermediate goods in total revenue is $(1 - Q_F/Q)$, where Q_F is final production (or value added), and Q is total (gross) output. Using the equations in the Appendix, we obtain

$$(1 - \alpha) = \mu \left(1 - \frac{Q_F}{Q}\right). \qquad (7)$$

This equation defines a negative relationship between α and the markup. That, combined with the restriction that α must lie strictly between zero and one, defines a range of possible values. So the upper bound for α is 0.5 — this corresponds to the case where there are no markups — but with markups the true α may be much smaller. Table 2 gives a range of values for α, with corresponding implied values of μ.

The intuition for the appearance of the markup in the expression for the share of intermediate goods comes from the fact that the economy is imperfectly competitive. Since the production function is Cobb-Douglas, the share of intermediate goods in total *cost* is $(1-\alpha)$. But in an economy with monopolistic competition, the cost share equals the revenue share multiplied by the markup. To calculate the cost share from the observable revenue share, we must take a position on the size of the markup.

Clearly if the markup is high, (i.e. the elasticity of demand is close to one), the share of profits in output becomes implausibly large. The most natural way to reconcile high markups with low observed profits is to suppose that there are large fixed costs of production. In this view, output in excess of variable cost is largely consumed by fixed costs.[13] The production function of equation (2) can easily be amended to allow for fixed costs, without any change in the preceding analysis.[14] With fixed costs,

13 This is the view of Robert E. Hall (1986).

14 The production function would then be $Q_i = (L_i - L^*)^{\alpha} I_i^{1-\alpha}$, where L^* is the fixed cost. This alteration would not affect any of the analysis, which depends only on changes in variable cost.

however, the interpretation of α changes: α is no longer the share of labor in total cost, but the share of *variable* labor in total *variable* cost. This is important to keep in mind, because in later sections we will see that plausible values of the markup imply that α is small — on the order of 0.2 to 0.1. But the share of non-produced inputs (capital and labor) in total cost is about 0.5. With overhead labor and capital, however, these different figures are not contradictory: it is perfectly possible for the share of variable labor and capital to be small, while their total share (inclusive of fixed costs) is large.[15]

III. The Model Broadly Construed: Countercyclical Markups

This section shows that a business cycle model with constant returns in production, no technological externalities, and no technology shocks can account for one of the major stylized facts of business cycles, procyclical labor and total factor productivity. As noted in the introduction, the result of procyclical productivity and its consequences, which are explored in the following sections, would obtain in any model with intermediate goods and countercyclical markups. Price rigidity is one way of generating countercyclical markups, but other explanations — such as the customer-market model of Edmund S. Phelps and Sidney G. Winter (1970), or the supergame-theoretic model of Julio J. Rotemberg and Garth Saloner (1986) — would serve equally well. Thus, the results of this section and the following ones apply to a much broader class of models than the one considered in Section I, and constitute an interesting mechanism for the transmission of shocks in purely real models of countercyclical markups such as Rotemberg and Michael Woodford (1991).

The result of procyclical productivity is driven by three properties of the model. First, since firms are imperfectly competitive — prices are above marginal costs — the equilibrium is inefficient. Second, the inefficiency is lower at higher levels of output: the model has the feature that markups are countercyclical. Third, intermediate goods are used in production. Blanchard and Nobuhiro Kiyotaki

15 In this discussion, I have assumed that fixed costs take the form of overhead capital and labor, not intermediate goods. While it is analytically possible that fixed costs include "overhead goods," there do not seem to be interesting examples of such costs in the way that factory buildings and non-production workers constitute examples of overhead capital and labor costs. (In certain industries, one exception might be advertising.)

(1987) present a model that has the first and second features, but not the third. In their model, the presence of markup pricing distorts the labor-leisure choice but not producers' decisions about the mix of inputs to employ: since they assume that labor is the only input to production, there is no "input mix" to distort. Thus, even though higher levels of output raise welfare in their model, this increase in welfare does not raise productivity. But when manufactured inputs are used in production, markup pricing also makes firms' production decisions socially suboptimal. In particular, at the initial equilibrium firms use too much of the primary input, labor, and too few manufactured inputs. In this model, the decrease in markups that accompanies output movements causes the ratio of input prices to be closer to the marginal rate of transformation between goods and labor. Thus, firms making their input choices make decisions about the quantity of labor to employ versus the quantity of manufactured inputs to use that are closer to being socially optimal. This increase in social efficiency not only increases welfare, it also raises productivity.

So far the discussion of the paper has used gross output as the concept of production. However, gross output is the total output of a firm, including the output used by other firms as intermediate goods. But the efficiency of an economy is judged by its ability to produce final goods from a given quantity of non-produced inputs (here only labor). So the correct statistic to examine is value added relative to labor, rather than gross output relative to labor.

To examine the issue of cyclical productivity, we first derive the economy-wide (and sectoral) value-added production function. Since labor is the only input to production and there are constant returns to scale, we can represent the net output or value-added production function as

$$Q_F = A\,L \qquad (8)$$

where A is this economy's "total factor productivity."[16]

We wish to examine the change in productivity as a function of the change in value added, assuming that menu costs are large enough to prevent prices from changing in response to the money shock that causes the output expansion. Alternatively, one can view the output movement as an expansion

[16] Since the gross-output production function (2) is separable in labor and materials, a value-added function exists. In fact, since the production function is Cobb-Douglas, a value-added function would exist even if there were more than one primary input (capital and labor, for instance).

of real aggregate demand in a flexible-price model where, for any of the reasons given in the models cited above, a one percent increase in final output is accompanied by an $\alpha\phi$ percent reduction in the markup.

The percent change in A is derived from this experiment. Taking the appropriate derivatives and evaluating the resulting expression at the real wage that prevails at the initial equilibrium yields:

$$\frac{dA}{A} = \left(\frac{\phi(1-\alpha)(\mu-1)}{(\mu+\alpha-1)}\right)\frac{dQ_F}{Q_F}\,. \qquad (9)$$

This expression is positive and rises monotonically as α falls.[17] It shows why imperfect competition is necessary for the result: if $\mu = 1$, so that there are no distortions in production, productivity is not procyclical. Of course, if $\alpha = 1$ there are no intermediate goods and productivity is again acyclical: this is the special case of Blanchard and Kiyotaki (1987). Note that the result is *not* being driven by Hall's (1988) argument that the Solow residual is procyclical if there are markups. In this model economy, dA/A is the correct measure of the change in productive efficiency; in Hall's terminology it is the "cost-based Solow residual." Hall (1990) claims that this cost-based residual should be invariant even with markup pricing, unless firms have increasing returns to scale. This model provides a counterexample: all firms produce with constant returns to scale, but the cost-based residual is procyclical. The intuition for this result, and its consequences for Hall's tests, are discussed below.

Values for the coefficient in (9) are reported in the fourth line of Table 2 for settings of α and μ that satisfy calibration. The table shows that, for a given value of the markup, productivity changes are higher for a lower α. Thus, as one would expect, changes in total factor productivity are larger as intermediate goods become more important in production. For the U.S. over the period 1962-84, a one percent growth in output resulting from a demand shock is associated with a 0.59 percent growth in total factor productivity.[18] The model as calibrated would predict a growth in total factor productivity of between 0.02 percent and 0.33 percent in response to a one percent growth in output resulting from a

[17] Note that the size of α actually has two opposing effects. First, a low α means that there is high usage of intermediate goods in production, so a reduction in the relative price of intermediate goods is very important. On the other hand, a low α means that the markup does not fall as much in an output expansion. In this model, the first effect always outweighs the second.

[18] Rotemberg and Lawrence Summers (1990, Table II).

demand shock. Obviously, for some empirically reasonable parameter values, the model is capable of accounting for a significant proportion of demand-induced changes in total factor productivity.

A striking feature of this model is that it predicts a positive correlation between sectoral productivity and aggregate activity. The increase in productivity comes from the fact that each industry is making more efficient choices about the mix of factors to employ in production.[19] Given the assumption that each industry uses as intermediate goods mostly the outputs of other industries, the increase in efficiency depends on an across-the-board reduction in markups. An uniform reduction in the markup for all varieties of intermediate goods will be correlated with changes in aggregate activity, but not with industry-specific changes in output. Caballero and Lyons (1990a, 1992) present empirical evidence that sectoral productivity is in fact correlated with aggregate output in both European and U.S. manufacturing. The model provides an economic explanation of their finding.

This explanation of the Caballero-Lyons stylized fact does not assume that there are true technological spillovers operating at business-cycle frequencies. Not only is it difficult to model such externalities, it is difficult even to tell an intuitive story for what form they might take. The advantage of this model is that it does not rely on high-frequency shifts in the production function, but rather on cyclical changes in the relative price of inputs to production.

Consequently the model delivers a sharp prediction about how we can distinguish between these two explanations for the Caballero-Lyons findings. It predicts that such an effect should be found when a production function is estimated with value-added data, but not with gross-output data. Estimating the gross-output production function amounts to estimating equation (2). By assumption, there are no technological externalities in the production of gross output, so a correct estimation of (2) will reveal none. But each productive unit becomes more efficient at creating value-added because the markup is

[19] Here is another point where taking the correct, broad, view of intermediate inputs provides better intuition. By using a Cobb-Douglas production function, I have assumed that the elasticity of substitution between labor and intermediate goods is one. Rotemberg and Woodford (1992) use econometric estimates and an imperfectly competitive general equilibrium model to calibrate this elasticity equal to 1.2 — *greater* than the Cobb-Douglas case. These relatively large elasticities of substitution are not surprising when one recalls that intermediate inputs include inputs of services, which are an increasingly large share of total inputs. The elasticity of substitution between using, for example, an in-house computer technician and an outside repairman is surely very high.

smaller. This increase in efficiency is correlated with increases in aggregate output. This distinction between the two explanations is testable; it is examined in Section V.[20]

IV. Implications for Hall's Tests

In a series of papers, Hall (1986, 1988, 1990) has explored various ingenious methods of determining the markup of price over marginal cost and the degree of returns to scale from time series data. One implication of the result of Section III is that if intermediate goods are used in production and markups are countercyclical, many of Hall's estimates are likely to be biased upward.

In his 1986 and 1988 papers, Hall attempts to estimate the markup in various industries and at an aggregate level. He notes that if output is sold with a markup, the standard "revenue-based" Solow residual should be procyclical in response to demand shocks. He isolates demand shocks by using instruments and shows that the Solow residual is in fact strongly procyclical when output is driven by exogenous changes in demand. In his 1990 paper, he argues that if production takes place with constant returns to scale then the true measure of total factor productivity — the cost-based Solow residual — should be unaffected by shocks to demand, regardless of markup pricing. The cost-based residual also fails his "invariance" test, so Hall concludes that there are increasing returns to scale.

Hall errs in both cases by assuming that true productivity, as measured by the cost-based Solow residual, is procyclical only if there are increasing returns. The previous section has shown that even with constant returns to scale, the cost-based Solow residual may rise in an expansion because of an external effect working through relative factor prices. In the world described by this model, Hall's (1990) tests would incorrectly conclude that there were increasing returns to scale. The cyclicality of the revenue-based Solow residual results from a combination of the cyclicality of the cost-based residual and the presence of markup pricing. By attributing the entire cyclicality of the revenue-based residual to markups, Hall's 1986 and 1988 tests would overstate the size of the markup.

[20] This predicted difference will also hold if the Caballero-Lyons results in fact stem from cyclical factor utilization. A change in utilization is a shift of the gross-output production function, and will show up as such in gross-output data as well as in value-added data.

Econometrically, the problem is that the error terms in Hall's regressions are correlated with the right-hand-side variables, even if they are instrumented. This bias is not eliminated — indeed, it is exacerbated — by his use of instruments. Hall assumes that technology shocks are the only source of error in his regressions. In the model I have presented here there are (by assumption) no true technology shocks, but if they were modeled there would be two parts to the error term in Hall's regressions: technology shocks and changes in sectoral productivity from countercyclical markups. Hall's instruments are uncorrelated with the first part of the error term but not the second. Indeed, when Hall uses aggregate demand instruments, he is studying only the changes in sectoral productivity that come from aggregate demand. Thus, his instrumented estimates correspond to exactly the case studied here, where by assumption sectoral and aggregate output changes are perfectly correlated and all fluctuations are driven by aggregate demand shocks. Hence, all of my results below should be understood as applying to Hall's IV estimates. In order to solve the problem I identify here, Hall would have to use instruments that are correlated with sectoral demand but uncorrelated with aggregate demand. It is difficult to conceive of such instruments.

Hall also allows for changes in the capital stock over the business cycle. Here I have not modeled capital as a factor of production,[21] but if one does so assuming the capital stock is fixed (or, more realistically, that movements in the capital stock are not correlated with high-frequency output movements[22]), all the results below go through.[23] If there is capital, the interpretation of each variable changes: one should think of output and labor input as being measured in "per capital" magnitudes.

It is easy to show that in this model economy characterized by imperfect competition, constant returns, and markup pricing — just what Hall (1988) assumes — Hall's methodology leads to a systematic upward bias in the estimate of the markup. Hall's estimating equation is

$$(\Delta \ln Q_F) = \mu^* (\sigma \Delta \ln L) \tag{10}$$

[21] Actually, of course, there is some concept of "capital" already in the model, since it is assumed that profit is paid to the owners of the firm rather than to workers.
[22] This assumption is not far off the mark: in two-digit U.S. manufacturing over 1959 to 1984, the average correlation between the annual growth of the capital stock and the annual growth of output is -0.027.
[23] This claim depends of course on how one models the productive contribution of capital. I assume that capital and labor enter multiplicatively in the production of value added.

where σ is the share of labor in value added

$$\sigma = \frac{WL}{PAL}$$

and μ^* is claimed to be an unbiased estimate of the markup.[24]

In the model developed here, however,

$$\Delta \ln Q_F = \Delta \ln A + \Delta \ln L$$

If $\Delta\ln A$ did not comove with output or labor input — or were uncorrelated with instrumented changes in labor input, which would happen if $\Delta\ln A$ were a pure technology shock — the expectation of the estimated markup would be given by

$$E(\mu^* \,|\, \Delta\ln A = 0) = \frac{\operatorname{cov}(\Delta\ln Q_F, \sigma\Delta\ln L)}{\operatorname{var}(\sigma\Delta\ln L)} = \frac{1}{\sigma} = \frac{P}{W/A} = \frac{\mu - (1-\alpha)}{1-(1-\alpha)}. \tag{11}$$

Even without considering the problems posed by the correlation of $\Delta\ln A$ with changes in labor input, we see that the expectation of μ^* is not equal to the true μ. This is because Hall uses value-added data for estimating the markup. As Ian Domowitz, R. Glenn Hubbard, and Bruce C. Petersen (1988) note, if materials are used in production then the markup estimated from value added exceeds the true markup since it divides profits by a smaller denominator (value added rather than gross output).[25] This is easier to see if we use (7) to rewrite (11) as

$$E(\mu^* \,|\, \Delta\ln A = 0) = \frac{\mu(Q_F/Q)}{1 - \mu(1 - Q_F/Q)} \equiv \mu^{VA}, \tag{12}$$

recalling that $(1 - Q_F/Q)$ is the share of intermediate inputs in revenue. But given an unbiased estimate of μ^{VA}, it is relatively easy to back out the true μ, since the relation between the two depends only on the observable share of intermediate goods in revenue. So if one ignores changes in sectoral productivity that are correlated with aggregate demand instruments, Hall's methodology at least gives an unbiased estimate of μ^{VA}, from which we can calculate the true markup, μ.

But as shown in (9), changes in A, the productivity of labor, covary systematically and positively with changes in Q_F and L. Therefore, if one takes into account the fact that $\Delta\ln A$ changes predictably in

[24] Actually, since Hall uses instruments, he runs the regression in reverse and estimates $1/\mu^*$. For the reasons discussed above, I have not explicitly modeled this procedure, but the results are unaffected by this change.
[25] Hall is clearly aware of this aspect of the value-added/gross-output distinction; see Hall (1986).

response to demand shocks, Hall's methodology creates an upward bias even in the estimate of μ^{VA} in this economy. The size of the bias can be computed from the equation for the true expectation of μ^*:

$$E(\mu^*) = \frac{\mathrm{cov}\,(\Delta \ln L + \Delta \ln A\,, \sigma \Delta \ln L)}{\mathrm{var}\,(\sigma \Delta \ln L)} = \mu^{VA}\left(1 + \frac{\mathrm{cov}(\Delta \ln A\,, \Delta \ln L)}{\mathrm{var}(\Delta \ln L)}\right). \quad (13)$$

It is evident that the bias is a positive one (since productivity comoves positively with labor input). Calculations of the size of this bias are given in Table 3. Again, computations for the calibrated parameter values are reported in the last line of Table 2. These show that for some admissible values of the relevant parameters, the bias can be large.

This result provides one way to pin down the calibration of the model. The bias of μ^{VA} relative to μ found in (12) and the bias of μ^* relative to μ^{VA} found in (13) are both consequences of Hall's use of value-added data. In that case, the estimation of Domowitz, Hubbard, and Petersen (1988) using gross-output data should be unbiased. They find an average markup of about 1.6. This figure should be compared to those labeled "Implied Markup" in the second line of Table 2. So, as previously claimed, the last two columns, with implied true markups of 1.5 and 1.7, seem to be the most germane. It is clear that the parameter values that seem most relevant — high values of μ and low values of α — are also the ones which imply that introducing intermediate goods into a model with countercyclical markups has strong economic implications.

Hall's 1990 paper tested the invariance of the cost-based Solow residual at both the economy-wide and two-digit SIC levels. In this model, Hall's estimating equation would be:

$$\Delta \ln Q_F = \gamma(\sigma_C \Delta \ln L), \quad (14)$$

where σ_C is the share of labor in the total cost of producing value added (always one in this case) and γ is Hall's estimate of the degree of returns to scale.[26] For the reasons given above and in the previous section, this methodology would wrongly conclude that firms had increasing returns to scale ($\gamma > 1$), when in fact the true γ is identically equal to 1. The percentage biases would be the same as those given in Tables 2 and 3. An easy way to test the hypothesis that Hall's estimate of γ is biased up is to estimate (14) using gross-output data — which should not be subject to this bias — and contrast the results with Hall's

26 Again, Hall's use of instruments allowed him to run this equation "backwards" and estimate $1/\gamma$.

value-added estimates. I present evidence below which indicates that Hall's estimates are in fact subject to precisely this bias. The gross-output estimates imply returns to scale that are about constant (or slightly decreasing) — a sharp contrast with Hall's finding of strongly increasing returns, and evidence in favor of this model.

The reason for the bias is the confusion of external effects with internal increasing returns to scale. In the model considered here, all firms produce with constant returns, but an aggregate demand shock increases both output and productivity. Thus, an increase in the output of every firm is correlated with an increase in the economy-wide efficiency of production — an external effect which Hall's procedure would mistake for increasing returns to scale at the firm level.

V. Some Empirical Evidence

The model has two types of empirical implications. First, it predicts cyclical movements of some ratios that are not often studied in business cycle theory. For example, it implies that the prices of intermediate goods should be countercyclical relative to the price of labor. Also, the quantities of intermediate goods used should be procyclical, again relative to labor input.

More direct tests of the model examine its predictions regarding the substitution of materials for labor and the behavior of total factor productivity.[27] First, the model predicts that there should be a positive correlation between changes in the materials-output ratio and changes in the ratio of wages to the price of intermediate inputs (since the increased use of materials results from a change in the relative price of inputs). Second, if the change in the relative price has the effect claimed in the paper, then there should also be a positive correlation between procyclical total factor productivity and procyclical usage of materials relative to labor. The most novel prediction of the paper is the third effect: as discussed above, estimates of external effects using value-added data should be significant, but similar estimates using gross-output data should be insignificant.

27 I would like to thank an anonymous referee for suggesting the first two tests.

In this section, I check these predictions against U.S. time series data. All of the predictions are empirically verified.

The data I use were compiled by Dale Jorgenson and a number of coworkers, and are thoroughly described in Jorgenson, Gollop, and Fraumeni (1987).[28] A major improvements in this data set relative to standard NIPA data is that all the inputs are quality-adjusted. The labor input series, for example, is a quality-weighted index of hours worked by different categories of workers, rather than the usual measure of the sum of hours worked by all workers which implicitly assigns equal weight to all workers. The data set employs gross production as the relevant concept of output, and therefore reports quantities used of capital, labor, energy, and materials.[29] Since some of the predictions to be tested involve materials usage, this makes the Jorgenson data set especially suitable.

The data used for the tests are a panel of annual observations on 21 manufacturing industries in the U.S. from 1959 to 1984. The definitions of the industries are standard two-digit S.I.C., with the exception that the Jorgenson data set separates Motor Vehicles (S.I.C. 371) from Other Transportation Equipment (S.I.C. 372-79). Thus, there are 21 industries rather than the usual 20.

All of the regressions involve testing for cyclical effects. To avoid the possibility that the cyclicality of both output and, say, the price of materials relative to labor are driven by technology shocks that make materials usage more attractive in a boom, I typically instrument the right hand side variables. The instruments are those suggested by Valerie A. Ramey (1989): the change in the world price of oil, changes in military expenditures, and the political party of the president.

A. Cyclical Regularities

[28] I thank Dale Jorgenson for making this data set available.

[29] The materials series has been modified to reflect usage per year rather than purchases. However, although work in process and intermediate goods inventories are strongly procyclical, as documented by Ramey (1989), the results are not sensitive to this modification (inventories are quite small relative to total intermediate input use).

As noted above, one of the predictions of the predictions of the model is that materials prices are countercyclical, relative to the prices of substitutes such as labor and capital.[30] This prediction is tested using the Jorgenson data. I estimate the following equation (where all variables are in logs):

$$(\Delta P_{m,it} - \Delta P_{l,it}) = \text{constant}_i + \beta_1 \Delta Y_{it} \qquad (15)$$

$P_{m,i}$ and $P_{l,i}$ are the prices of intermediate goods and labor inputs to a given industry i; Y_i is sectoral output.

The first line of Table 4 reports the estimate of β_1 (with the elasticity constrained to be equal across industries). The elasticity is negative and significant, as the model predicts.[31]

This result, that one measure of the real wage is significantly procyclical, may seem at odds with the conventional wisdom that the real wage is acyclical or only slightly procyclical. The cost of labor input is more procyclical in the Jorgenson data set because labor quality is significantly countercyclical.[32] A slightly higher real wage paid to lower quality workers implies that the cost of an efficiency unit of labor is much higher in booms.[33] The basic prediction of the model is confirmed: in an expansion, labor become more expensive relative to intermediate inputs, which must lead producers to economize on labor and use intermediate goods more intensively.

The next prediction I check is the claim that materials usage is procyclical relative to labor. The basic equation I estimate is:

$$(\Delta M_{it} - \Delta L_{it}) = \text{constant}_i + \beta_2 \Delta Y_{it} \qquad (16)$$

[30] The usual claim that the observed wage is less procyclical than the shadow wage works against the results of this section. If the true wage is more procyclical than it appears from the data, the result that materials prices are countercyclical is even stronger than it appears.

[31] For purposes of empirical work, there are clearly different concepts of "labor" and "intermediate inputs" that might be used. We can think of "labor," as used in this model, as comprising both capital and labor inputs; materials may or may not include energy. In the empirical work I have used the standard definition of labor and a concept of materials that excludes energy (which anyway is only about 5 percent of a typical sector's materials input). Excluding energy shows that the results are not being driven by the oil price shocks of the 70s. However, the findings are robust to using all combinations of these different concepts; in many cases the results are strengthened by using different definitions of "labor" and "materials."

[32] The importance of this composition effect for analyzing real wage cyclicality has been pointed out by Finn E. Kydland and Edward C. Prescott (1988) and Gary Solon, Robert Barsky, and Jonathan Parker (1992).

[33] The Jorgenson data show that the *average* wage is procyclical. We are concerned with the marginal cost of labor, so we want to know how the *marginal* wage changes. Mark Bils (1987) shows that the marginal wage increases about eight times faster than the average wage (because of overtime payments). Making Bils's correction would therefore greatly strengthen the finding that the relative price of intermediate goods is countercyclical.

where M_i and L_i are (the log of) intermediate goods and labor inputs. The second line of Table 4 gives the estimate of β_2 for U.S. manufacturing. It is apparent that the elasticity is positive and significant. This is just what is predicted by the model.

One might believe, however, that the result of procyclical intermediate goods usage is being driven by labor hoarding.[34] Estimates of production functions and the degree of returns to scale are often thought to be subject to cyclical measurement error. In this view, the apparent acyclicality of labor hours may stem from unmeasured procyclical work effort. If so, true labor input will be procyclical even if measured hours are not. If this effect is not taken into account, one might wrongly conclude that the ratio of intermediate goods to labor is procyclical.

One way to control for labor hoarding is to include right-hand-side variables that are plausible proxies for cyclical labor utilization. One such variable, average hours worked per employee (AGH), has been proposed by Thomas Abbott, Zvi Griliches, and Jerry Hausman (1989). Following Caballero and Lyons (1992), I also use two other variables to control for labor hoarding: the ratio of production to non-production workers (PNP), and the average number of overtime hours worked (OVT).[35]

The results are found in the third line of Table 4. They confirm the hypothesis of labor hoarding: the variables that control for changes in effective labor input always have the correct sign and are usually significant. As expected, taking labor hoarding into account reduces β_2. But even accounting for labor hoarding, intermediate input usage remains strongly procyclical relative to labor.

Another way to see if labor hoarding is responsible for the results is to examine the cyclicality of the ratio of intermediate inputs to industry output.[36] In fact, this test is biased against finding procyclicality. Changes in both industry and aggregate output are likely to be driven by common productivity shocks and oil price shocks, imparting a negative bias to the results. For this reason, I do not instrument the explanatory variable, changes in aggregate output. Almost all of the explanatory power of the instruments used previously comes from oil prices. But using oil prices as an instrument would only exacerbate the bias, by isolating those changes in aggregate output that are most strongly correlated with the error term.

---

[34] I am indebted to Albert Ando for this observation.
[35] I thank Ricardo Caballero for providing these data.
[36] I thank Valerie Ramey for suggesting this test.

The results are reported in the fourth line of Table 4. The ratio of intermediate inputs to output is procyclical and statistically significant. The numerical magnitudes are smaller, but this is not surprising, given the negative bias in the results noted above. These results should be thought of as a lower bound on the procyclicality of intermediate goods usage; it is apparent that even the lower bound is positive, as the model predicts. The existence of labor hoarding does not alter this basic finding.

B. Specific Predictions

In this section I test the more specific predictions of the model.

First, I test whether changes in the intermediate goods-output ratio are consistently related to changes in the relative price of these inputs. To test the prediction, I run the regression

$$(\Delta M_{it} - \Delta Y_{it}) = \text{constant}_i + \beta_3 (\Delta P_{l,it} - \Delta P_{m,it}) \tag{17}$$

The result is found in the last line of Table 4; β_3 is positive and significant. Note that if the production function is in fact Cobb-Douglas, the estimated coefficient in this regression should be α. The coefficient is 0.12, which is in line with the calibrated value of α: based on evidence regarding the size of the markup, α was predicted to be between 0.2 and 0.1.

Next I test the prediction that changes in the input mix are responsible for changes in total factor productivity. Here, however, there is the problem with labor hoarding discussed above. If there is a significant degree of unmeasured factor utilization that applies to labor but not to intermediate goods, the measured procyclicality of total factor productivity and of the intermediate goods-labor ratio may both be driven by cyclical measurement error. There is a way to distinguish these two hypotheses, however. To the extent that cyclical measurement error is driving the finding of procyclical productivity, this effect should be apparent in both gross-output and value-added data. However, as argued above, the procyclicality resulting from countercyclical markups should be found only in value added. This suggests that we should regress two different measures of the Solow residual on changes in the materials-labor ratio — one measure calculated from gross output and the other from value added. If the value-added estimate is significantly larger, this will imply that the mechanism identified in the paper is at work. So I estimate

$$\Delta TFP_{it} = \text{constant}_i + \beta_4 (\Delta M_{it} - \Delta L_{it}) \qquad (18)$$

In calculating the growth rate of total factor productivity, I use cost shares[37] rather than revenue shares to avoid the problem pointed out by Hall (1988): productivity calculated using revenue shares appears spuriously procyclical if firms price their product above marginal cost. I do not allow for increasing returns in production, which Hall (1990) argues is responsible for the failure of invariance in the cost-based Solow residual. Recent empirical work by Martin N. Baily, Charles Hulten, and David Campbell (1992) applies Hall's procedure to plant-level gross-output data from the Longitudinal Research Database and finds essentially constant returns to scale. In the next series of tests, I come to the same conclusion using industry-level gross-output data. Hence, it appears that Hall's (1990) finding of increasing returns is a figment of the value-added bias identified in the previous section; this is another piece of evidence in favor of the model presented here.

The results of the test using total factor productivity are reported in Table 5. Note first that there is evidence of significant procyclicality of the gross-output residual (ΔTFP^{GO}) in response to changes in the materials-labor ratio. Therefore, as indicated above, part of the movement in this ratio most likely reflects changes in unmeasured labor utilization. However, it is also clear that the value-added estimate significantly exceeds — by almost a factor of three — the gross-output estimate. Therefore, although it appears that some of the correlation between total factor productivity and changes in the materials-labor ratio reflect labor hoarding, the data support the contention that some other mechanism like the one identified here is also at work.

Finally, I test the prediction that if the mechanism proposed by the model is responsible for the finding that procyclical productivity is an external effect, we should be able to detect the effect in value-

[37] The cost shares are calculated as in Hall (1990). However, I use capital-specific depreciation rates and tax parameters (the investment tax credit and the present value of depreciation allowances) that vary by industry.

added data but not in gross-output data.[38] The empirical procedure follows Caballero and Lyons (1990b). I estimate the equation:

$$\Delta Y_{it} = \text{constant}_i + \gamma \Delta X_{it} + \kappa \Delta X_t, \qquad (19)$$

where ΔX_i is the cost share-weighted sum of sectoral input growths, and ΔX is growth of aggregate (manufacturing) inputs, similarly cost-weighted. ΔY_i is either the growth of value added or of gross output. In the value-added regressions, the inputs are capital and labor; in the gross-output regressions the inputs are capital, labor, energy, and materials. Value added is constructed using the cost share of materials (see notes 37 and 38). γ is the degree of internal returns to scale; κ captures external effects from aggregate activity. (See Caballero and Lyons (1990b) for a fuller description of the procedure.)

The results are found in Table 5. It is clear that the predictions of the model are exactly verified. In the gross-output regressions, the point estimate of κ is 0.01 — very close to zero, and not significantly different from zero. On the other hand, the value-added estimate is 0.80, with a t-statistic that exceeds 9. Therefore it is apparent that the predictions of the model are borne out. This finding is also significant for interpreting the Caballero-Lyons stylized fact. The difference between the two sets of results indicates that their finding of a large positive κ is not evidence for a true technological externality, but rather an indication that a more subtle effect, having to do with cyclical changes in markups, is at work.

This simple model therefore has a number of quite detailed predictions about cyclical patterns of input use, input prices, and the behavior of productivity over the business cycle. Data from U.S. manufacturing industries confirms these predictions.

V. Conclusion

[38] Susanto Basu and John G. Fernald (1993a, 1993b) propose a different explanation for why value-added data should give incorrect estimates of returns to scale and external effects. Their explanation basically rests on the correct claim that with imperfect competition, value added should be calculated using the cost share of materials rather than the revenue share. (For the definition of value added, see Kenneth Arrow (1974).) To meet this objection, I used cost shares to construct the value-added data that I use in equation (19). As the results show, even with this correction there are external effects in value added and none in gross output, as the model predicts.

A wide variety of evidence indicates that modern economies are characterized by imperfectly competitive behavior, and many business-cycle models of imperfect competition imply countercyclical markups. Intermediate goods are widely used in production. This paper has explored the implications of these two sets of stylized facts, and finds that in conjunction they lead to a number of strong results. In purely real models they imply that productivity, even correctly measured, is procyclical. In models where countercyclical markups are a consequence of output price rigidity, modeling the use of intermediate goods in production implies that business cycles are both larger and more costly.

The model makes a number of predictions about cyclical productivity that accord with the facts. Among other things, it predicts that sectoral productivity should appear in the data as an external effect: the productivity of one sector should be correlated with aggregate rather than sectoral activity. The model also implies that Hall's estimates of markups and returns to scale are biased up. Other authors have interpreted the finding of external effects in productivity as evidence for technological spillovers or for labor hoarding. I show that there is a sharp empirical test that can discriminate among these various hypotheses. If the type of model presented here is responsible for the finding of external effects, these effects should be present in value-added data but not in gross-output data. If the spillover or labor hoarding hypotheses are at work, on the other hand, then they should be present in both gross output and value added. It turns out that the spillovers are found only in value added, which confirms the predictions of the model. The paper predicts that the biases in Hall's work should also be a function of his use of value-added data. Using Hall's procedure to estimate returns to scale from gross-output data shows that there are constant returns, not the strongly increasing returns that Hall finds. This finding also supports a model of the kind presented here, with countercyclical markups and intermediate goods in production.

It is important to stress that although price rigidity is not necessary to derive the results on cyclical productivity, these results do follow naturally from a sticky-price model. Thus, a setting in which the menu-cost assumption easily explains large welfare losses — a model with imperfect competition and heavy usage of intermediate goods — also enables these models to explain many of the stylized facts on cyclical productivity. This paper, then, provides a link between the purely real and purely nominal literatures within the New Keynesian economics. This link should be a subject of future research.

REFERENCES

Abbott, Thomas, Zvi Griliches, and Jerry Hausman (1989). "Short Run Movements in Productivity: Market Power versus Capacity Utilization." Mimeo.

Akerlof, George, and Janet Yellen (1985). "A Near-Rational Model of the Business Cycle, with Wage and Price Inertia." *Quarterly Journal of Economics* 100 (Aug.) 823-38.

Arrow, Kenneth (1974). "The Measurement of Real Value Added." In P. David and M. Reder (eds.) *Nations and Households in Economic Growth* (New York: Academic Press).

Baily, Martin N., Charles Hulten, and David Campbell (1992). "Productivity Dynamics in Manufacturing Plants," *Brookings Papers on Economic Activity* (Microeconomics), 187-267.

Ball, Laurence, and David Romer (1989). "Are Prices too Sticky?" *Quarterly Journal of Economics* 104 (Aug.) 507-524.

______ and ______ (1990). "Real Rigidities and the Non-Neutrality of Money." *Review of Economic Studies.*

Basu, Susanto and John G. Fernald (1993a). "Constant Returns and Small Markups in U.S. Manufacturing" Mimeo, University of Michigan.

______ and ______ (1993b). "Are Apparent Productive Spillovers a Figment of Specification Error?" Mimeo, University of Michigan.

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

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

Blanchard, Olivier J. (1983). "Price Asynchronization and Price-Level Inertia. In Rudiger Dornbusch and Maria Henrique Simonson (eds.) *Inflation, Debt, and Indexation* (Cambridge: MIT Press).

______ and Nobuhiro Kiyotaki (1987). "Monopolistic Competition and the Effects of Aggregate Demand." *American Economic Review* 77 (Sept.) 647-666.

Bureau of Economic Analysis (1984). "The Input-Output Structure of the U.S. Economy, 1977." U.S. Department of Commerce, Survey of Current Business, May 1984, 42-84.

Caballero, Ricardo J., and Richard K. Lyons (1990a). "Internal Versus External Economies in European Industry." *European Economic Review* 34, 805-30.

______, and ______ (1990b). "The Role of External Economies in U.S. Manufacturing." National Bureau of Economic Research Working Paper 3033.

______, and ______ (1992). "External Effects in U.S. Procyclical Productivity." Mimeo, Columbia University. Forthcoming in *Journal of Monetary Economics.*

Carlton, Dennis W. (1986). "The Rigidity of Prices." *American Economic Review* 76 (Sept.) 637-658.

Domowitz, Ian, R. Glenn Hubbard, and Bruce C. Petersen (1988). "Market Structure and Cyclical Fluctuations in U.S. Manufacturing." *Review of Economics and Statistics* 55-66..

Evans, Charles L. (1992). "Productivity Shocks and Real Business Cycles." *Journal of Monetary Economics* 29 (April) 191-208.

Gordon, Robert J. "What is New-Keynesian Economics?" *Journal of Economic Literature* 28 (Sept.) 1115-1171.

Hall, Robert E. (1986). "Market Structure and Macroeconomic Fluctuations." *Brookings Papers on Economic Activity*, 2, 285-322.

______ (1988). "The Relation Between Price and Marginal Cost in U.S. Industry." *Journal of Political Economy* 96 (Oct.) 921-947.

______ (1990). "Invariance Properties of Solow's Productivity Residual." In Peter Diamond (ed.) *Growth, Productivity, Employment* (Cambridge: MIT Press).

Jorgenson, Dale W., Frank Gollop, and Barbara Fraumeni (1987). *Productivity and US Economic Growth*. Cambridge: Harvard University Press.

Kydland, Finn E. and Edward C. Prescott (1988). "Cyclical Movements of the Labor Input and Its Real Wage." Federal Reserve Bank of Minneapolis Working Paper no. 413.

Mankiw, N. Gregory (1985). "Small Menu Costs and Large Business Cycles: A Macroeconomic Model of Monopoly." *Quarterly Journal of Economics* 100 (May) 529-539.

______ (1991). "Appendix: A General Equilibrium Example." In N. Gregory Mankiw and David Romer (eds.) New Keynesian Economics (Cambridge: MIT Press) pp. 36-40.

Means, Gardiner C. (1935). "Industrial Prices and Their Relative Inflexibility." U.S. Senate Document 13, 74th Congress, 1st Session, Washington D.C.

Murphy, Kevin M., Andrei Shleifer, and Robert W. Vishny (1989). "Building Blocks of Market Clearing Business Cycle Models." In Olivier J. Blanchard and Stanley Fischer (eds.) NBER Macroeconomics Annual (Cambridge: MIT Press).

Phelps, Edmund S. and Sidney G. Winter (1970). "Optimal Price Policy Under Atomistic Competition." In Edmund S. Phelps (ed.) *Microeconomic Foundations of Employment and Inflation Theory* (New York: W. W. Norton).

Ramey, Valerie A. (1989). "Inventories as Factors of Production and Economic Fluctuations." *American Economic Review* 79 (June) 338-54.

Rotemberg, Julio J., and Garth Saloner (1986). "A Super Game Theoretic Model of Price Wars during Booms." *American Economic Review* 77 (Dec.) 917-926.

______, and Lawrence H. Summers (1990). "Inflexible Prices and Procyclical Productivity." *Quarterly Journal of Economics* 105 (Nov.) 851-874.

______, and Michael Woodford (1991). "Markups and the Business Cycle." In Olivier J. Blanchard and Stanley Fischer (eds.) NBER Macroeconomics Annual (Cambridge: MIT Press).

______, and ______ (1992). "The Effects of Energy Price Increases on Economic Activity." Mimeo.

Solon, Gary, Robert Barsky, and Jonathan Parker (1992). "Measuring the Cyclicality of Real Wages: How Important is Composition Bias." Mimeo, University of Michigan.

Stigler, George J. and James K. Kindahl (1970). *The Behavior of Industrial Prices*. New York: NBER.

Appendix: Derivations

Maximizing (1) subject to a standard budget constraint gives the consumer's first-order conditions:

$$W = M, \tag{A1}$$

and

$$Q_{iF} = W^{1/\phi} P_i^{-1/\phi}. \tag{A2}$$

Minimizing costs subject to the production function (2) gives the input demands for each firm:

$$L_{ki} = \left(\frac{P_k}{P}\right)^{-1/\phi} \left(\frac{1-\alpha}{\alpha}\frac{W}{P}\right) L_i, \tag{A3}$$

and

$$L_i = \left(\frac{\alpha}{1-\alpha}\frac{P}{W}\right)^{1-\alpha} Q_i. \tag{A4}$$

The cost function is then given by substituting (A3) and (A4) into the expression for the cost of production:

$$C_i = k\, W\, Q_i \left(\frac{P}{W}\right)^{1-\alpha} \tag{A5}$$

where

$$k = \left(\frac{\alpha}{1-\alpha}\right)^{1-\alpha} + \left(\frac{\alpha}{1-\alpha}\right)^{-\alpha}.$$

The total output of each firm, Q_i, is given by the sum of demands for its output as final goods (equation A2) and as intermediate inputs (the integral of (A3) over i). This gives

$$Q_i = \left[\left(\frac{\alpha}{1-\alpha}\right)^{-\alpha}\left(\frac{W}{P}\right)^{\alpha} P^{1/\phi}\, Q + W^{1/\phi}\right] P_i^{-1/\phi} \tag{A6}$$

where $Q = \int_0^1 Q_i \, di$ is aggregate (gross) output. One can solve for the equilibrium aggregates Q_F and Q from (A2) and (A6), using the fact that all firms' quantities and prices are equal in this symmetric equilibrium.

Table 1. Losses Resulting from a One Percent Money Shock
(as percent of profit)

	$\phi = 0.9$ ($\mu = 1.1$)	$\phi = 0.5$ ($\mu = 2$)	$\phi = 0.1$ ($\mu = 10$)
$\alpha = 1.0$	1 (6.17E-6)	1 (1.00E-4)	1 (0.0045)
$\alpha = 0.9$	0.8102 (5.00E-6)	0.8102 (8.10E-5)	0.8102 (0.0036)
$\alpha = 0.7$	0.4901 (3.02E-6)	0.4901 (4.90E-5)	0.4901 (0.0022)
$\alpha = 0.5$	0.2500 (1.54E-6)	0.2500 (2.50E-5)	0.2500 (0.0011)
$\alpha = 0.3$	0.0900 (5.56E-7)	0.0900 (9.00E-6)	0.0900 (4.00E-4)
$\alpha = 0.1$	0.0100 (6.17E-8)	0.0100 (1.00E-6)	0.0100 (4.50E-5)

Table 2. Computations for Parameter Values Satisfying Calibration

	$\alpha = 0.45$	$\alpha = 0.4$	$\alpha = 0.3$	$\alpha = 0.2$	$\alpha = 0.1$
Implied Markup	1.042	1.137	1.327	1.516	1.706
Loss from Fixed Prices for a 1 % Money Shock[1]	0.202	0.160	0.090	0.040	0.010
Loss from Fixed Prices as Fraction of Profit[2]	5.92E-3	4.84E-4	5.60E-5	1.14E-5	1.71E-6
Percent Change in Productivity for1% Output Change	0.002	0.018	0.090	0.196	0.326
Percent Bias in Estimated μ^{VA}	0.165	1.88	9.87	24.39	48.37

[1] Loss is normalized by loss in the base case ($\alpha = 1$) as in Table 1.
[2] Also computed for a one percent shock to money.

Table 3. Percent Bias in Hall's Estimate of the Markup

	$\mu = 10$	$\mu = 2$	$\mu = 1.1$
$\alpha = 1.0$	0	0	0
$\alpha = 0.9$	8.82	2.67	0.11
$\alpha = 0.7$	33.73	9.89	0.41
$\alpha = 0.5$	73.44	19.76	0.91
$\alpha = 0.3$	158.14	36.99	1.94
$\alpha = 0.1$	404.55	69.49	4.93

Table 4. Empirical Regularities

Dependent Variable	Independent Variables: ΔY_i	ΔY	PNP_i	OUT_i	AGH_i	$\Delta P_{l,i} - \Delta P_{m,i}$
$\Delta P_{m,i} - \Delta P_{l,i}$	-0.20 (0.02)					
$\Delta M_i - \Delta L_i$	0.56 (0.08)					
	0.41 (0.07)		0.004 (0.002)	0.006 (0.004)	0.026 (0.008)	
$\Delta M_i - \Delta Y_i$		0.18 (0.02)				
						0.12 (0.01)

Standard errors in parentheses.

Table 5. Value-Added and Gross-Output Results

Dependent Variable	Independent Variables: $\Delta M_i - \Delta L_i$	ΔX_i^{VA}	ΔX^{VA}	ΔX_i^{GO}	ΔX^{GO}
ΔTFP_i^{VA}	0.33 (0.02)				
ΔTFP_i^{GO}	0.12 (0.01)				
ΔY_i^{VA}		0.63 (0.04)	0.79 (0.08)		
ΔY_i^{GO}				0.96 (0.01)	0.01 (0.02)

Standard errors in parentheses.