MEASURING THE CYCLICALITY OF REAL WAGES: HOW IMPORTANT IS THE FIRM'S POINT OF VIEW?

Eric T. Swanson*

Abstract—There is a growing consensus among economists that real wages in the postwar United States have been procyclical, greatly bolstering technology-driven theories of business cycles at the expense of more classical models. This paper makes the point that technological movements in firm's labor demand curves should be tested with a wage that is deflated by the firm's own price of output, with appropriate controls for intermediate inputs, and with respect to the cyclical state of the firm's own industry, as opposed to the state of the aggregate economy. Failure to control for these factors is found to lead to substantial overrejection of the classical model. In detailed industry data, with controls for changes in worker composition, I find that a vast majority of sectors have paid real product wages that vary inversely (that is, countercyclically) with the state of their industry.

I. Introduction

WHAT causes business cycles? This is one of the most fundamental questions faced by macroeconomists, yet answers remain controversial.

One notable explanation, proposed by Kydland and Prescott (1982), is that economic fluctuations are the result of exogenous shocks to the real productive potential of the economy. Their basic model, and the standard variations of it, imply a positive correlation between real wages and the quantity of labor employed, due to these technology shocks:

> If utility is separable over time and all goods are superior, then we can generate an increase in today's consumption and work effort—hence a decline in today's leisure—only if there is a change in the current technological parameter α_t that generates an upward shift in today's schedule for the marginal product of labor. . . . But notice that the real wage rate, which equals the marginal product of labor, must rise along with the increases in output and work effort. In other words, a procyclical pattern for the real wage rate is central to our theoretical analysis. (Barro & King, 1984, pp. 832–833)

In contrast, the classical and Keynesian explanations for economic fluctuations were typically based upon a stable labor demand curve, with workers' labor supply curves shifting due to nominally rigid union contracts or money illusion (in the classical framework), or with worker's labor supply curves roughly fixed but (dis)equilibrium being attained off of the labor supply curve (in Keynes's framework):

> In emphasizing our point of departure from the classical system, we must not overlook an important point of agreement. For we shall maintain the first postulate [I. The wage is equal to the marginal product of labor (p. 5)] as heretofore, subject only to the same qualifications as in the classical theory, and we must pause, for a moment, to consider what this involves.
>
> It means that for a given organisation, equipment, and technique, real wages and the volume of output (and hence of employment) are uniquely correlated, so that, in general, an increase in employment can only occur to the accompaniment of a decline in the rate of real wages. Thus, I am not disputing this vital fact which the classical economists have (rightly) asserted as indefeasible. (Keynes, 1936, p. 17).

Thus, these more traditional theories imply an inverse relationship between real wages and employment, exactly opposite the predictions of the technology-driven models.

One can thus try to discriminate between these two sets of theories by using the correlation between real wages and employment at cyclical frequencies as a guide. Specifically, the null hypothesis that firms' labor demand curves are stable at business cycle frequencies (the classical or Keynesian theory) can be tested against the alternative that firms' labor demand curves shift at business cycle frequencies. It should be noted that a number of theories are consistent with rejection of this null, including the technology-driven class of real business cycle (RBC) models described above, models of countercyclical markups (for example, Rotemberg & Woodford, 1992),[1] and models of external increasing returns to scale (for example, Bartelsman, Caballero, & Lyons, 1994). I have chosen to focus on the technology-driven class of RBC models as the alternative for ease of exposition, and because this class of models is the most widespread in the literature.[2]

Received for publication January 19, 2001. Revision accepted for publication February 11, 2003.

* Board of Governors of the Federal Reserve System.

This paper is based on chapter 3 of my doctoral dissertation at Stanford; an earlier version circulated under the title "Measuring the Cyclicality of Real Wages: How Important is Aggregation Across Industries?" I thank Michael Horvath, Bob Hall, John Taylor, John Pencavel, Mun Ho, Kevin Stiroh, Charles Fleischman, Jonathan Parker, Mark Schweitzer, and Bob Gordon; seminar participants at Stanford, the Federal Reserve Board, the 1999 Federal Reserve System Conference in Cleveland, and the 2001 Meetings of the American Economic Association; and two anonymous referees for helpful discussions, comments, and suggestions. The views expressed in this paper, and all errors and omissions, should be regarded as those solely of the author, and are not necessarily those of the individuals listed above, the Federal Reserve System, or its Board of Governors.

[1] A number of textbook Keynesian models in which prices are fixed and firms supply whatever output is demanded at the given price (for example, Mankiw, 2002; Abel and Bernanke, 2001) fall into this category. Some of these models have sticky wages as well as prices, implying acyclical real wages; some assume that wages are less sticky than prices, implying procyclical real wages.

[2] It should also be noted that some recent papers (for example, Basu, Fernald, & Kimball, 1998) present dynamic general equilibrium (DGE) models in which technology shocks have *contractionary* implications for employment, and thus imply a countercyclical real wage. I will focus on the mainstream class of DGE models in this paper.

FIGURE 1.—THE LABOR MARKET FROM WORKER (LEFT) AND FIRM (RIGHT) POINTS OF VIEW

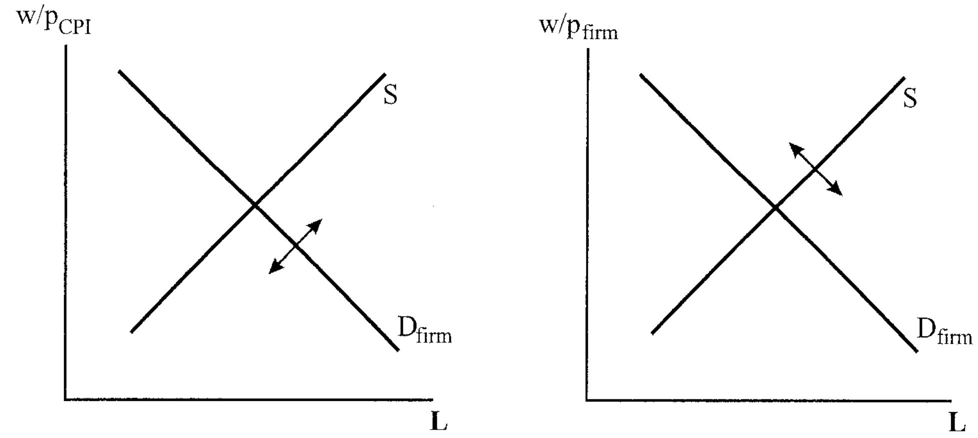

A simple diagram (figure 1) provides the intuition for the key difference between this paper and previous ones in the wage cyclicality literature [for examples of this literature, see the survey by Abraham and Haltiwanger (1995)]. Both the left- and right-hand panels of figure 1 graph the labor demand and labor supply curves faced by a given firm. In both panels, labor demand traces out the marginal product of labor, and labor supply the marginal disutility of work. The two charts differ in that the diagram on the left is graphed with respect to the real CPI-deflated wage, whereas the diagram on the right is graphed with respect to the real product wage (the nominal wage deflated by the firm's price of output).[3] With respect to the real CPI-deflated wage, labor supply is fixed (assuming no taste shocks), and labor demand shifts as the firm's product price varies relative to the CPI. With respect to the real product wage, labor demand is fixed (assuming no technology shocks and no changes in the price of nonlabor inputs), and labor supply shifts as the firm's product price varies relative to the overall CPI.

Previous studies of wage cyclicality have followed the approach in the left-hand panel of figure 1, deflating wages by an aggregate price index such as the CPI, and using measures of the aggregate quantity of labor, such as the national unemployment rate. In contrast, this paper looks at wages and employment from the firm's point of view, as in the right-hand panel of figure 1.[4] If the classical model is correct, then firms' labor demand curves should be stable at cyclical frequencies, and we should see an inverse (that is, countercyclical) relationship between real product wages and industry employment. Alternatively, if the technology-driven class of RBC models is correct, then firms' labor demand curves shift substantially at cyclical frequencies, and the real-wage–employment correlation should be positive (despite the ambiguity in the figure, a positive correlation is unequivocally predicted by the theory, as in the quotation from Barro and King above).[5] It is important to note that the left-hand panel of figure 1 is not appropriate for distinguishing between these two classes of theories, because firms' labor demand curves shift under both the null and alternative hypotheses, as described in the preceding paragraph.

Previous studies' findings of real-wage procyclicality, or at least an absence of countercyclicality, have led many economists to question the relevance of the classical and Keynesian models. However, these studies, including the panel studies by Bils (1985), Keane, Moffitt, and Runkle (1988), and Solon, Barsky, and Parker (1994),[6] have consistently adopted a non-firm-oriented point of view, as in the left-hand panel of figure 1, using an aggregate price index such as the GNP deflator to deflate wages, instead of the price of the firm's own output, and measuring cyclicality with respect to the state of the aggregate economy, instead of conditions in the industry in which the worker and firm are employed. Although these choices of aggregate deflator and cycle indicator may be appropriate for investigating the effects of worker composition change on the cyclicality of

[3] It is assumed in this paper that economic shocks often lead to variation in relative prices (and employment). For example, an energy price shock disproportionately affects firms in energy-intensive industries; an interest-rate shock disproportionately affects firms in financing-intensive industries (housing, automobiles); and a government purchases shock disproportionately affects firms in government-purchases-intensive industries (defense). Swanson (1999) models these types of effects in a DGE framework.

[4] Note that, to the extent that relative prices and economic conditions vary across firms and industries over time, the proper approach is a firm- or sector-level measure of both product prices and employment, as depicted in the right-hand panel of figure 1. In particular, deflating aggregate wages by the (aggregate) PPI and regressing them on aggregate employment would not be taking seriously the issues raised here and in the rest of the paper, although it would be a small step in the right direction.

[5] Barro and King's analysis uses a representative-agent, representative-firm framework, and one might thus object that their model might have different implications for real-product-wage–employment correlations when carried out within a heterogeneous, multisector framework. I will argue in section VII that my results call for exactly such a heterogeneous, multisector approach.

[6] One of the primary advantages of these panel studies is that they can control for changes in workforce composition over the business cycle. I will discuss and take great care to allow for these composition effects in my analysis below.

aggregate wage statistics, they are not appropriate for testing the null hypothesis that firms' labor demand curves are stable at cyclical frequencies (the classical or the Keynesian theory), as discussed above and in more detail below. Performing this test is the focus of the present paper, and requires the more firm-oriented approach in the right-hand panel of figure 1.

One might question whether industry employment is the most appropriate measure for testing a theory of business cycles, which are aggregate phenomena. The business cycle theories described above generally do not draw a distinction between industry and aggregate cyclicality, assuming instead that all industries comove perfectly, as in a representative-agent, representative-firm framework. In contrast, the present paper does not assume perfect comovement across industries, and recognizes that to test for movement in firms' labor demand curves, the industry measure of employment is the more relevant one. I will discuss to what extent my results have implications for aggregate economic fluctuations in section VI.

The present paper is divided into seven sections. Section II presents a standard, classical model of wage and employment determination. Section III describes the two data sets used to test the theory: the 458-sector NBER Productivity Database and Jorgenson's 34-sector KLEM data set. Section IV presents the main results. Section V compares and contrasts the results with previous studies in the literature, and section VI discusses the broader implications of the results. Section VII concludes.

II. A Simple Model of Labor Demand

Consider the case of a classical, cost-minimizing firm with no monopsony power. The firm faces a production function $y_t = F(k_t, l_t, e_t, m_t, t)$, where y denotes output, k capital, l labor, e energy, m nonenergy intermediate inputs from other sectors, and t time. Assume that k_t is freely chosen at the beginning of each period t and then fixed for the remainder of the period, while l_t, e_t, and m_t can be varied freely in response to shocks.[7] Assume also that F is increasing and has a well-defined interior profit-maximizing choice of l_t, e_t, and m_t for any given k_t and for any set of prices in the range of the data.[8]

[7] A period is 1 year in my data below. The assumption of quasi-fixity of capital is the norm in the DGE literature (for example, King, Plosser, & Rebelo, 1988); allowing for a more freely variable capital stock would require including capital stock data in the regressions (instead of a smooth trend), and capital stock data are generally of poor quality. Also, previous studies of wage cyclicality have avoided using capital stock data, so doing so here maintains comparability of this paper's results with the earlier literature.

[8] Assuming F is homogeneous of degree 1 in its first four arguments, is strictly concave, and satisfies gradient Inada-type conditions in its middle three arguments would be sufficient. In general, we need to rule out cases in which F has increasing returns to scale or its gradient is too flat as its inputs approach 0. Capital k_t is required to trend smoothly over time below, so the assumption of a unique profit-maximizing k_t is not necessary, allowing constant returns to scale.

If the function F is Cobb-Douglas, so that $y_t = f(t)k_t^{\theta}l_t^{\phi}e_t^{\beta}m_t^{\gamma}$, it is straightforward to show that

$$\log \frac{w_t}{p_t} = \alpha \log l_t + \beta \log e_t + \gamma \log m_t + g(t, k_t, \mu_t), \tag{1}$$

where w_t denotes the nominal wage, p_t the price of the firm's output, and μ_t the firm's markup (price divided by marginal cost), where $\alpha \equiv \phi - 1 < 0$, and where the function g encompasses a constant and terms relating to t, k_t, and μ_t. Equation (1) states that the equilibrium real product wage is log linear in the firm's inputs; it may also be regarded as a (log) first-order approximation to the firm's labor demand curve for a given capital stock and markup and a completely general production function F with factor-augmenting technical change.

The classical (or Keynesian) model assumes that the firm's technology and capital stock are stable at cyclical frequencies; I will also assume that the markup μ_t is stable at these frequencies (this is the case under constant markups, or under the standard assumptions of perfect competition and profit maximization, for example). With these assumptions as part of the null hypothesis, the function $g(t, k_t, \mu_t)$ trends smoothly over time, and will be well approximated by a polynomial or Hodrick-Prescott trend $h(t)$. The equation

$$\log \frac{w_t}{p_t} = \alpha \log l_t + \beta \log e_t + \gamma \log m_t + h(t) + \varepsilon_t \tag{2}$$

is then a stable relationship that will serve as the basis for my empirical work below.[9] It should be emphasized that under the null hypothesis, $h(t)$ captures all relevant shifts of the labor demand curve (2) due to changes in capital stocks, markups, and technology, so that the error terms ε_t (which represent deviations of the data from the Cobb-Douglas or log-linearity assumption) have no systematic correlation with the regressors. Finally, because of nonstationarity or near-nonstationarity in the dependent and independent variables, as described in section III below, I will typically estimate equation (2) in first differences rather than in levels.

Equation (2) is essentially the same specification as that of previous researchers, with the addition of energy and nonenergy intermediate input terms on the right-hand side. One can gain intuition for the inclusion of these additional terms in equation (2) as follows: The aim of this paper is to

[9] One could also consider the specification

$$\log l_t = a \log (w_t/p_t) + b \log (p_{e_t}/p_t) + c \log (p_{m_t}/p_t) + \tilde{g}(t, k_t, \mu_t) + \eta_t, \tag{2'}$$

which states that the quantity of labor demanded is a function of the prices of the inputs. This specification is less comparable with the usual specification (3) below, but its results are very similar and are discussed in footnotes to the main results below.

shed light on the contrasting implications of the business cycle theories mentioned above. Although exogenous shocks to the prices of inputs, such as oil, are interesting in their own right and are undoubtedly an important source of economic fluctuations, they are not at odds with any of the classical, Keynesian, or technology-driven models of business cycles. Put another way, if one is interested in testing for technological movements in labor demand, one must control for movements in labor demand that are due to nontechnological factors. Failing to do so could lead to rejection of the classical model when in fact it is completely consistent with the data.

If the assumptions underlying the model above do not hold, then equation (2) will be misspecified. For example, if technology does not trend smoothly over time, but is instead a major source of cyclical fluctuations in firm employment and real wages, then equation (2) will not satisfy the classical regression assumptions—in particular, the error term will generally be strongly and positively correlated with log l_t.[10] Thus, even though theory predicted that $\alpha < 0$, standard regression procedures applied to equation (2) may yield estimates of α that are insignificantly different from 0, or even positive.[11] Standard technology-driven models of RBCs in fact predict this last result.

Previous empirical work using regression specifications similar to equation (2) have in fact often estimated values of α that are positive. This has led many economists to question whether the classical assumption of cyclically stable technology and markups is appropriate, and greatly bolstered alternative theories of business cycles. However, none of these empirical studies has used data that are consistent with the firm's point of view, so conclusions about the shifting of firms' labor demand curves based on these results are inappropriate and premature.

In fact, estimates of α based on these non-firm-oriented methods are likely to lead to overrejection of the classical model. For example, it is common in the literature to estimate

$$\log \frac{w_t}{P_t} = \tilde{\alpha} \log L_t + \tilde{h}(t) + \tilde{\varepsilon}_t, \tag{3}$$

where P_t refers to some aggregate price index and L_t some aggregate measure of labor. We would expect several sources of error from running regression (3) instead of (2). First, to the extent that changes in the aggregate price index do not fully capture changes in the prices of intermediate inputs (and large changes in oil prices over the sample make this a valid concern), we would expect estimates of $\tilde{\alpha}$ in equation (3) to suffer from an omitted variables bias, relative to the true α in equation (2).[12] Second, if l_t and L_t are not perfectly correlated (so that $L_t = l_t + \eta_t$, where η_t is a stochastic error term), then estimates of $\tilde{\alpha}$ in equation (3) suffer from an errors-in-variables bias for α as well. Finally, to the extent that p_t and P_t are not perfectly correlated, we would expect an increase in the variance of the left-side variable in equation (3), leading to a further loss of precision in the estimates (though no bias).

The implications of estimating equation (2) with non-firm-oriented data are thus potentially serious. All three effects described above predict a deterioration in the quality of the estimates. More importantly, two of these effects are predicted to lead to an errors-in-variables or omitted variables bias in tests of the model's predictions under the null. That these issues are important in practice as well as in principle is demonstrated below.

III. Data and Methods

In estimating equation (2), one would ideally like to have comprehensive data on wages, prices, and hours worked for a large number of matched workers and firms within a variety of industries over a significant period of time. Unfortunately, such data are not presently available. The NBER Productivity Database [available from the NBER's Web site, and documented in Bartelsman and Gray (1996)], however, does contain sectoral data for all manufacturing industries at the four-digit (SIC) level between 1958 and 1994 at annual frequency. I will also use data at roughly the two-digit level compiled by Dale Jorgenson and his coworkers [available from Jorgenson's Web site, and documented in Jorgenson, Gollop, and Fraumeni (1987)], because its method of construction has some unique advantages and because having a second data set provides corroboration and an additional perspective on the results.

Detailed industry data may in fact be preferable to firm- or plant-level data in many respects, or at least not a significant drawback. For example, technology may be adopted not by existing plants or firms, but rather by new firms entering into the industry; similarly, industry productivity may improve as a result of obsolete firms exiting the sector. In these situations, cyclical fluctuations in patterns of technology adoption could be very important to the industry

[10] Some recent papers (Galí, 1999; Basu, Fernald, & Kimball, 1998) have focused on the possibility that technology shocks might be contractionary with respect to employment. As mentioned previously, I will focus on the mainstream view of technology shocks as the alternative hypothesis in this paper.

[11] The sign of the bias on α in equation (2) depends on the interaction between l_t, e_t, m_t, and ε_t. Letting x denote the price of a materials-energy composite input (which greatly simplifies the following), it is not difficult to show that this bias is a positive constant times $\rho_{l\varepsilon} - \rho_{xl}\rho_{x\varepsilon}$, where ρ denotes the correlation between the two corresponding variables in equation (2). Technology shocks in a given sector will tend to induce $\rho_{l\varepsilon} > 0$ and $\rho_{x\varepsilon} > 0$. Empirical studies such as that of Murphy, Shleifer, and Vishny (1989) document $\rho_{xl} > 0$. The sign of the bias on α will thus be positive if ρ_{xl} and $\rho_{x\varepsilon}$ are small enough relative to $\rho_{l\varepsilon}$. A priori there is no reason to think $\rho_{x\varepsilon} > \rho_{l\varepsilon}$; in fact, materials prices tend to be very volatile, so we might expect $\rho_{x\varepsilon} < \rho_{l\varepsilon}$ to hold—that is, the expected bias on α is positive.

[12] It is in this sense that I use the term "bias" in what follows. When interpreted as a raw sample correlation statistic, a regression coefficient can never be biased. As a test statistic for the implications of the classical model above, some estimators of α can be biased and lead to overrejection of the null.

as a whole, and yet not to its individual firms. By looking at detailed industry data instead of firms or plants, I allow for the possibility that technology may influence the economy in this way. In addition, to the extent that firms within a given industry face competitive input and output markets and maximize profits within this framework, theorems of indirect aggregation assure us that an industry-wide labor demand curve, derived from an industry-wide profit function, exists and satisfies all of the properties described in the previous section.[13]

Data on all variables in the NBER Productivity Database are at the four-digit (SIC) level. For labor input, I used total production worker hours (*PRODH*); for the nominal per-unit wage, I used total production worker wages (*PRODW*) divided by *PRODH*. I chose to focus on production worker hours and wages for two reasons: first, data on hours for nonproduction workers are generally unavailable (even at the two-digit level, they must be imputed from worker surveys); and second, production workers form a more homogeneous input than do all workers, so that calculation of "the" wage for a unit of labor and "the" quantity of labor employed is a more valid approximation. Product prices were measured as the price deflator for the value of shipments (*PISHIP*). There are two intermediate input price indices for each sector in the NBER Productivity Database: energy prices (*PIEN*), and the price of all other intermediate inputs to production in the sector, which must be derived from the raw data series.[14]

In addition to the product price deflators described above, I constructed value-added deflators for purposes of comparison. The method is analogous to that used for nonenergy materials prices above.[15] A value-added deflator for the aggregate manufacturing sector as a whole was also constructed, using the Törnqvist index method once again.

[13] Although the indirect aggregation theorems also apply to the U.S. economy taken as a whole, the assumptions underlying the theorems are more difficult to maintain for the overall economy. For example, labor is more heterogeneous across the aggregate economy than it is within a single four-digit manufacturing industry. Capital also varies greatly across sectors of the economy in its type, tax treatment, depreciation, and riskiness of return, which makes the assumption of a competitive capital market across sectors much less tenable than within a given four-digit manufacturing industry.

[14] See Bartelsman and Gray (1996). The nonenergy materials price deflator was constructed as follows: first, a real index of energy input was constructed using nominal energy expenditure (*ENERGY*) divided by *PIEN;* a real index of all intermediate inputs (energy plus nonenergy) was then constructed using nominal materials expenditure (*MATCOST*) divided by the price deflator for all materials (*PIMAT*); next, a real index of nonenergy materials was constructed using a theoretically ideal Törnqvist index as in Jorgenson, Gollop, and Fraumeni (1987); finally, the price for nonenergy materials was constructed as the nominal nonenergy materials expenditure (*MATCOST* − *ENERGY*) divided by this real index.

[15] First, nominal gross output was defined as nominal value added (*VADD*) plus the cost of intermediate inputs (*MATCOST*); real gross output was then constructed by dividing this number by the price deflator for shipments (*PISHIP*); next, real materials input was defined as *MATCOST* divided by the price deflator for all materials (*PIMAT*); finally, indices of real value added and value-added deflators were constructed from these gross output and materials numbers using a Törnqvist methodology.

In contrast to the NBER data, Jorgenson's KLEM data are at roughly the two-digit (SIC) level for the manufacturing and mining sectors, and at the one-digit level for other sectors of the economy. The data cover the years 1958–1996 at annual frequency. The KLEM data set complements the NBER Productivity Database in two key respects: first, although it is at a coarser level of detail, it covers nonmanufacturing sectors in addition to manufacturing; and second, Jorgenson and his associates have expended considerable effort constructing input and output indices that are adjusted for changes in composition. For example, labor input is divided into several hundred cells, corresponding to different levels of educational attainment, experience, sex, union status, managerial-production-clerical classification, and so on; the change in labor input is then calculated for each cell separately for each year (drawing on data from the CPS, the Census in benchmark years, and the BLS establishment surveys); and finally, these individual changes are aggregated into a single Törnqvist index for each sector based on the theory of ideal index number construction. Cyclical changes in labor force composition, which were found by Solon et al. (1994) to have important effects on standard aggregate measures of wages, will thus have little or no effect on the indices in the KLEM data set, because they have been constructed with quality adjustment taken into account. Finally, the KLEM data provide corroboration and additional perspective on the results obtained using the NBER productivity data. As with the NBER data, I also construct indices of the price of value added for each sector, and aggregate measures of wages, prices, labor input, and the price of value added for the whole KLEM economy, using a Törnqvist methodology.

Finally, in contrast to aggregate data, sectoral data at the two- and four-digit levels is often extremely variable and clearly nonstationary over the given sample period. For this reason, the regressions below are all estimated in first differences rather than in levels. Most previous studies of wage cyclicality have also used first-differenced data, so first-differencing preserves comparability with earlier work as well. Estimates using a low-parameter Hodrick-Prescott filter or a cubic or higher polynomial trend lead to qualitatively similar results. It should also be noted that the trend break in productivity and real wages that is present in the aggregate data around 1970 is essentially invisible in the detailed sectoral data; this is again because of the sectoral data's large, nonstationary movements or very pronounced trends.

IV. Results

A. NBER Productivity Data

Summary results for the 458 sectors in the NBER Productivity Database are presented in figures 2 and 3 and in table 1. Results for different time periods, such as 1970–1994, are very similar; in particular, all of the observations

FIGURE 2.—WAGE CYCLICALITY FROM FIRM AND NONFIRM POINTS OF VIEW

Firm-Oriented Specification (2) (a)

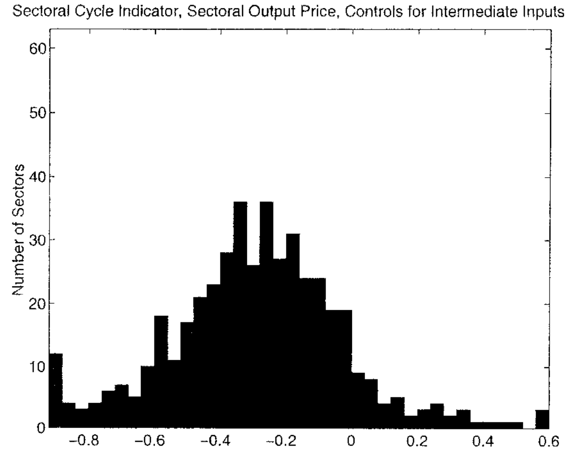

Countercyclical:	411
at 10% level:	270
at 5% level:	229
at 1% level:	176
Procyclical:	47
at 10% level:	8
at 5% level:	3
at 1% level:	2
Mean:	−.292
Weighted Mean:	−.310
Median:	−.286
Mean Abs. t-stat:	2.240

Industry nominal wages, deflated by industry prices, regressed on industry cycle indicator, with controls for intermediate inputs. The label (a) is for comparability with Figure 3.

Typical Nonfirm-Oriented Specification (3) (f)

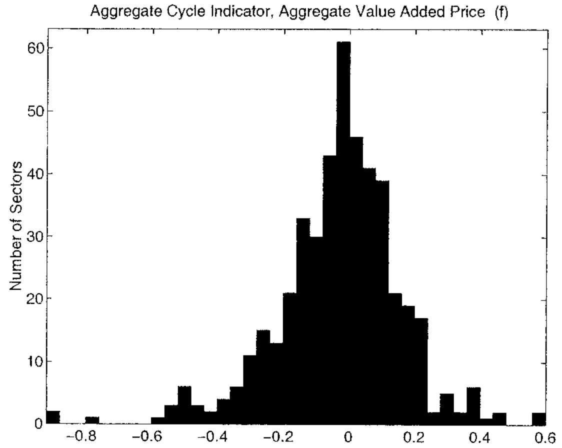

Countercyclical:	255
at 10% level:	34
at 5% level:	15
at 1% level:	8
Procyclical:	203
at 10% level:	21
at 5% level:	15
at 1% level:	6
Mean:	−.031
Weighted Mean:	−.004
Median:	−.016
Mean Abs. t-stat:	0.849

Industry nominal wages, deflated by aggregate manufacturing value added deflator, regressed on agg. cycle indicator, no controls for intermediate inputs. The label (f) is for comparability with Figure 3.

Completely Aggregate Specification (3) (g)

Coefficient:	.036
Std. Error:	.073

Aggregate manufacturing wage, deflated by agg. manufact. value added deflator, regressed on agg. cycle indicator, no controls for intermed. inputs. The label (g) is for comparability with Figure 3.

Regression coefficients for 458 NBER Productivity Database industries at the four-digit level. A. Firm-Oriented Specification (2) (a) A: Industry nominal wages, deflated by industry prices, regressed on industry cycle indicator, with controls for intermediate inputs. The label (a) is for comparability with figure 3. B. Typical Non-firm-Oriented Specification (3) (f) B: Industry nominal wages, deflated by aggregate manufacturing value-added deflator, regressed on aggregate cycle indicator, no controls for intermediate inputs. The label (f) is for comparability with figure 3. C. Completely Aggregate Specification (3) (g) Coefficient: .036, Std. Error: .073 C: Aggregate manufacturing wage, deflated by aggregate manufact. value-added deflator, regressed on aggregate cycle indicator, no controls for intermediate inputs. The label (g) is for comparability with figure 3.

that follow hold for this later time period as well. The robustness of the results across time periods should be regarded as evidence that my specification in equation (2) is a good one; in particular, it appears that I have done a reasonable job of controlling for the effects of changes in oil prices.

Each graph in figures 2 and 3 is a histogram of the 458 point estimates of the coefficient α for various specifications of regression equation (2) or (3); the panels differ in their choice of measures for p_t and l_t, and in the inclusion or exclusion of intermediate inputs in the regression. The nominal wage measure is the same in each panel, and is the

FIGURE 3.—INTERMEDIATE REGRESSION SPECIFICATIONS

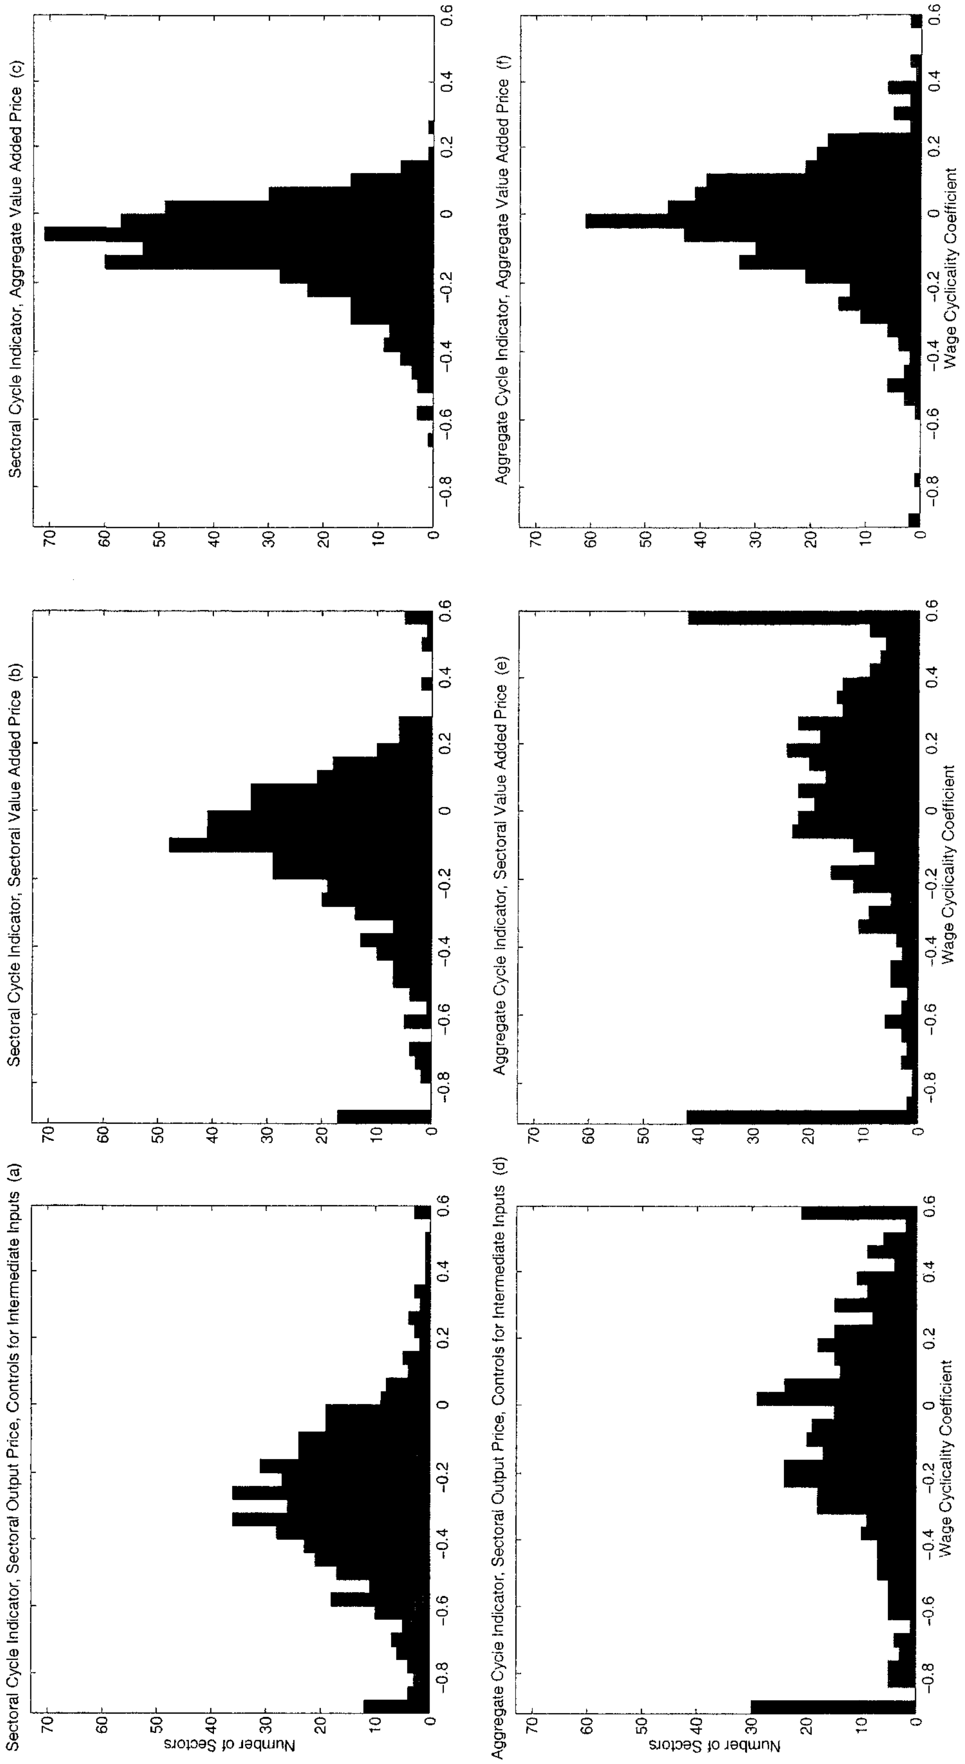

Regression coefficients for 458 NBER Productivity Database Industries.

TABLE 1.—SUMMARY STATISTICS FOR 458 NBER PRODUCTIVITY DATABASE REGRESSIONS

Firm-Oriented (a)		(b)		(c)	
Countercyclical:	411	Countercyclical:	321	Countercyclical:	356
at 10% level:	270	at 10% level:	113	at 10% level:	173
at 5% level:	229	at 5% level:	84	at 5% level:	135
at 1% level:	176	at 1% level:	56	at 1% level:	104
Procyclical:	47	Procyclical:	137	Procyclical:	102
at 10% level:	8	at 10% level:	14	at 10% level:	13
at 5% level:	3	at 5% level:	10	at 5% level:	9
at 1% level:	2	at 1% level:	5	at 1% level:	3
Mean:	−.292	Mean:	−.144	Mean:	−.102
Weighted mean:	−.310	Weighted mean:	−.122	Weighted mean:	−.083
Median:	−.286	Median:	−.089	Median:	−.079
Mean abs. t-stat.:	2.240	Mean abs. t-stat.:	1.306	Mean abs. t-stat.:	1.590
(d)		(e)		Typical Nonfirm (f)	
Countercyclical:	258	Countercyclical:	200	Countercyclical:	255
at 10% level:	85	at 10% level:	50	at 10% level:	34
at 5% level:	63	at 5% level:	42	at 5% level:	15
at 1% level:	44	at 1% level:	31	at 1% level:	8
Procyclical:	200	Procyclical:	258	Procyclical:	203
at 10% level:	51	at 10% level:	70	at 10% level:	21
at 5% level:	35	at 5% level:	45	at 5% level:	15
at 1% level:	21	at 1% level:	26	at 1% level:	6
Mean:	−.108	Mean:	−.067	Mean:	−.031
Weighted mean:	−.095	Weighted mean:	−.055	Weighted mean:	−.004
Median:	−.073	Median:	.063	Median:	−.016
Mean abs. t-stat.:	1.326	Mean abs. t-stat.:	1.263	Mean abs. t-stat.:	0.849
		Completely Aggregate (g)			
		Coefficient:	.036		
		Std. error:	.073		

Panels (a)–(g) correspond to figures 2 and 3.

detailed industry's average production worker wage, as defined in the previous section. Tables in figure 2 and table 1 present summary statistics for each regression specification. The term "countercyclical" in the tables refers to point estimates of α that are less than 0, whereas "procyclical" refers to estimates of α that are greater than 0. For each set of 458 regressions, the number of countercyclical and procyclical point estimates is given, as well as the number of each that are significant at the 10%, 5%, and 1% levels. The mean and median of the 458 coefficient estimates are presented, together with a weighted mean, with weights given by each sector's average share of total production worker hours in the first and last years of the sample. The mean absolute t-statistic provides a measure of the average precision of the coefficient estimates for each specification.

Note that I have chosen to estimate the wage cyclicality coefficients for each of the 458 industries separately, rather than impose the constraint that the coefficients are the same across industries. This allows for the possibility that technology-driven fluctuations in labor demand may be more important in some sectors of the economy than in others—for example, in newer industries in which the technology and production process are less well established than in older sectors of the economy. However, it should be kept in mind that these 458 coefficient estimates are not independent, for there is likely to be contemporaneous correlation in the residuals across sectors, as a result of macroeconomic disturbances.[16] Despite this limitation of the data, I am able to demonstrate significant differences between firm-oriented and non-firm-oriented specifications below.

Figure 2 summarizes the basic result of the paper. The first panel (a), labeled as firm-oriented, is the correct specification (2) under the assumptions of section II—the variable p_t is the detailed industry price of output, l_t is the detailed industry quantity of labor, and detailed industry energy and nonenergy intermediate inputs are included as regressors. Panel (f), the typical non-firm-oriented specification, corresponds to the regression specification (3), which has typically been run in the literature (for aggregate or panel studies of workers): the variable P_t is the value-added deflator for aggregate manufacturing, L_t is total production worker hours for aggregate manufacturing, and the regressions include no controls for energy or nonenergy intermediate inputs; specification (f) differs from most previous studies in the literature, however, in that it continues to use detailed industry nominal wages for w_t. Panel (g) of figure 2, the "Completely Aggregate Specification," aggre-

[16] Ideally, one would like to use a seemingly unrelated regression (SUR) framework to take the cross-correlation in the error terms into consideration, thereby improving the efficiency of the estimates. However, because the panel consists of 458 industries with only 36 observations in the time dimension, one cannot estimate a 458-by-458 variance-covariance matrix that has rank greater than 36, so the result is highly singular and an inversion for use in any kind of GLS procedure is impossible.

gates these wage data as well, and thus matches exactly the completely aggregate (wages, hours, and value-added prices) specifications considered in early studies of wage cyclicality.

The difference in results between the firm-oriented and non-firm-oriented specifications is striking. First, the firm-oriented specification (a) yields estimates that are much more often countercyclical ($\alpha < 0$) than procyclical ($\alpha > 0$). This is true both in terms of the raw numbers of point estimates, and at every level of statistical significance. In contrast, the typical non-firm-oriented specification (f) yields estimates that are substantially more procyclical and much less precise—in fact, the statistical significance of the estimates is not much different from pure sampling variation around 0. The mean, weighted mean, and median coefficients are virtually zero, and the mean absolute t-statistic for the regressions is much smaller than for specification (a). These observations are in agreement with the hypothesis of section II that there is a substantial omitted variables and errors-in-variables bias in the regression specification (3), the typical non-firm-oriented approach.[17] Finally, the point estimate and standard error for the completely aggregate specification (g) are very similar to those in the literature using completely aggregate methods (see Abraham & Haltiwanger, 1995), and are even more procyclical and less precise than specification (f).[18]

The effects of changes in worker composition are unable to explain these findings. First, estimates using the KLEM data (in figure 4 below), which control for composition change directly, are very similar. Second, because the data in each of these regressions is for production workers within a single detailed industry, changes in workforce composition are likely to be less important for this more homogeneous population than for the set of all (nonproduction and production) workers in the economy as a whole. Finally, to the extent that changes in workforce composition do affect the estimates, they will affect the estimates in each specification in figure 2 simultaneously, and hence will be differenced out when comparing the panels with each other.

In comparing panels (a) and (f) of figure 2, the question naturally arises what factors are the most important in leading to the disparate results. Figure 3 attempts to answer this question.

The format of Figure 3 and its accompanying table 1 are the same as in figure 2, with panels (b) through (e) of figure 3 and Table 1 considering various intermediate specifications between (a) and (f). The top panels [(a) through (c)] use the quantity of labor employed in each detailed industry as their cycle indicator l_t, whereas the bottom panels [(d) through (f)] take l_t to be the quantity of labor employed in aggregate manufacturing. The left-hand panels [(a) and (d)] use detailed industry prices for p_t and include energy and nonenergy intermediate inputs as regressors; the middle panels [(b) and (e)] use detailed industry prices of value added as their measure of p_t with no controls for intermediate inputs; and the right-hand panels [(c) and (f)] use the price of value added for aggregate manufacturing for their measure of p_t, again with no controls for intermediate inputs. Results for each of these will be discussed in turn below.

There are a number of conclusions to be drawn from figure 3. First, the use of an aggregate cycle indicator as a proxy for sectoral industry conditions appears to induce a large amount of noise and bias into the results, as was predicted in section II. A comparison of the top panels of figure 3 [(a) through (c)] with the bottom panels [(d) through (f)] reveals that the point estimates in the bottom panels are much more dispersed; the mean absolute t-statistics in the bottom panels are much smaller despite the greater dispersion of the point estimates; the number of significant point estimates is much less than for the upper panels at every level of significance; and the mean, weighted mean, and median coefficients for the bottom panels are all virtually zero, indicating a nontrivial upward or zero bias in the estimates, as compared to the specifications in the upper panels [(a) through (c)]. All of these observations are consistent with the hypothesis that the aggregate cycle indicator (aggregate number of production worker hours) is a poor proxy for the state of economic conditions in individual four-digit industries, leading to classic errors-in-variables bias.[19]

Second, it is apparent from figure 3 and table 1 that the coefficient estimates using sectoral value-added deflators

[17] The results for the analogous panels (a), (f), and (g) for the specification (2′) in footnote 9 are very similar: (a) countercyclical: 364, 161, 128, 102; procyclical: 94, 12, 6, 1; mean −.404, weighted mean −.365, median −.424, mean absolute t-statistic 1.574; (f) countercyclical: 255, 34, 15, 8; procyclical: 203, 21, 15, 6; mean −.016, weighted mean .019, median −.019, mean absolute t-statistic .849; (g) coefficient .232 (.433). Note that the coefficients in (f) and (g) for the specification (2′) are essentially reverse regression coefficients for (f) and (g) in figure 2; the point estimates are related by the ratio of the variances, as is well known, but the t-statistics are identical, a fact which is not often recognized.

[18] Note that this is despite the fact that the right-side variables in (g) are exactly the same as in (f). This is because the more cyclical industries, such as durable-goods manufacturing and construction, generally pay higher wages than do the less cyclical industries, such as services, so that, all else equal, a boom is a time when the aggregate economy has a greater proportion of high-wage workers. This yields a procyclical *industry composition effect* on the aggregate (Chirinko, 1980), in contrast to the countercyclical *worker composition effect* (low-wage workers at each firm tend to be fired first) emphasized by Solon et al. (1994).

[19] Although comovement between one-digit industries and the aggregate economy is fairly high (Murphy et al., 1989), comovement at the four-digit level is substantially lower, with an average correlation coefficient of approximately 0.4 (that is, an R^2 of 0.16). Although all sectors tend to move downward during a recession, idiosyncratic factors make up the bulk of each sector's variance. This is apparent when one looks at graphs of the sectoral data individually as well. Even in sectors which one could think of as highly cyclical (steel), some downturns are felt much more strongly than others (1982), because they tend to interact with other factors [energy prices, shifts of production overseas that led to large-scale plant closures and a permanent change in the nature of the industry (minimills, more specialized products)]. These types of stories are evident in the data for a large number of detailed industries.

[panels (b) and (e)] are extremely poor. The point estimates are very dispersed, much more so than for the correct specification (a), and yet the t-statistics for these regressions and the number of significant coefficients are relatively small despite the large point estimates. These findings corroborate those of Jorgenson et al. (1987) and Basu and Fernald (1997), who find that the assumptions required for a value-added production function to exist are often not met in sectoral data, and lead to poor estimation results when imposed. In addition, there appears to be a substantial positive bias in the results, relative to panels (a) and (d).[20] These results again suggest an omitted variables and errors-in-variables bias in the estimation of α using these specifications—omitted variables because of poor controls for changes in intermediate inputs, and errors-in-variables because the value-added deflators are very noisy proxies for the true price of output, p_t.[21]

Third, specifications using the aggregate value-added deflator [(c) and (f)] actually perform better than those using sectoral value-added deflators. This is again evidence that sectoral value-added deflators are not appropriate concepts. Turning to a comparison of specifications (c) and (f) with (a) and (d), it is apparent that the former provide significantly less precise estimates than the latter; in addition, the estimates using (c) and (f) are more tightly clustered around 0, indicating that an errors-in-variables bias toward 0 in these specifications is very plausible.

To summarize, the results are very supportive of the hypothesis that the firm-oriented specification (a) is the correct approach and the panels to the right and below in figure 3 are inferior proxies. In addition, the results in panel (a) are consistent with the classical model of section II. The deterioration in regression performance from the firm-oriented specification (a) to the typical non-firm-oriented specification (f) appears to be due to all three sources of error described in section II: omission of intermediate input terms, poor proxies for l_t, and poor proxies for p_t. No one factor appears to dominate the others: deterioration in regression performance is evident both as one moves to the right in figure 3 and as one moves down that figure. The next subsection demonstrates that these findings are robust across data sets and to controls for changes in worker composition.

B. KLEM Data

The results for the 34 nongovernmental sectors of Jorgenson's KLEM data set confirm those described above. Figure 4 and its accompanying table 2 present results in a format analogous to figure 3 and Table 1. In addition, table 3 presents point estimates and standard errors for each of the 34 KLEM sectors using firm-oriented specification (a).

The same patterns are evident in figure 4 that were apparent in figure 3.[22] First, the point estimates for the firm-oriented specification (a) are almost uniformly countercyclical rather than procyclical.[23] Second, as in figure 3, there is a substantial deterioration in regression performance from the firm-oriented specification (a) to the typical non-firm-oriented specification (f); moreover, the coefficients in (f) are again, statistically speaking, difficult to distinguish from pure random sampling around 0—evidence that the biases and sources of error discussed in section II are at work in these data as well. Third, deterioration in regression performance is evident both as one moves to the right in figure 4 and Table 2 and as one moves down, so that no one source of error in the typical nonfirm specification (f) appears to be dominant. However, it is noteworthy that specifications (b) and (e), which use value-added deflators, do not perform as disastrously as the corresponding specifications in figure 3 [though they still perform worse than specification (a)]. This may be because the value-added concept does better in two-digit- than in four-digit-level data, though the results still suggest that it is a poor proxy relative to using true gross output and intermediate input prices.

To summarize, the results are very much in line with those from the NBER Productivity Database and the predictions of the classical model of section II.

V. Comparison with the Literature

Many of the specifications (a) to (f) from the previous section have been considered in isolation elsewhere in the literature, with results that are consistent with those of the present paper.

First, as noted earlier, the completely aggregate specification (g), using NBER productivity data, replicates the findings of other studies in the literature that use similarly aggregate methods and composition-unadjusted data (Abraham & Haltiwanger, 1995).

[20] This is evident in the figures, but not in the mean and weighted mean statistics in (b) and (e), due to the fat negative tails of the distributions.

[21] Approximately half of the deterioration in regression performance is due to the omission of intermediate inputs, and approximately half to the use of value-added prices instead of gross output prices for p (not shown in figure 3 and table 1).

[22] Results for panels (a), (f), and (g) using the specification (2′) instead of (2) are very similar: (a) countercyclical: 29, 9, 5, 2; procyclical: 5, 1, 1, 0; mean −.124, weighted mean −.159, median −.118, mean absolute t-statistic 1.209; (f) countercyclical: 19, 1, 1, 1; procyclical: 15, 2, 1, 1; mean −.004, weighted mean .055, median −.016, mean absolute t-statistic .895; (g) coefficient .025 (.231).

[23] The mean, weighted mean, and median coefficients are around −.35, which is roughly consistent with the values from table 1. Note that the KLEM data use indices of total labor input and wages paid to all labor, as compared to production worker hours and production worker wages in the NBER Productivity Database. Thus, despite the presence of controls for changes in workforce composition in the KLEM data (which should lead to more procyclical point estimates ceteris paribus (Solon et al., 1994), it is not necessarily the case that the coefficients obtained using these data should be more procyclical than the NBER data. For example, production worker hours are approximately twice as variable (in terms of standard deviation) as are the KLEM indices of labor input, so a given change in wages will have a coefficient roughly twice as large when regressed on the less variable KLEM labor index.

FIGURE 4.—REGRESSION COEFFICIENTS FOR 34 KLEM DATABASE INDUSTRIES

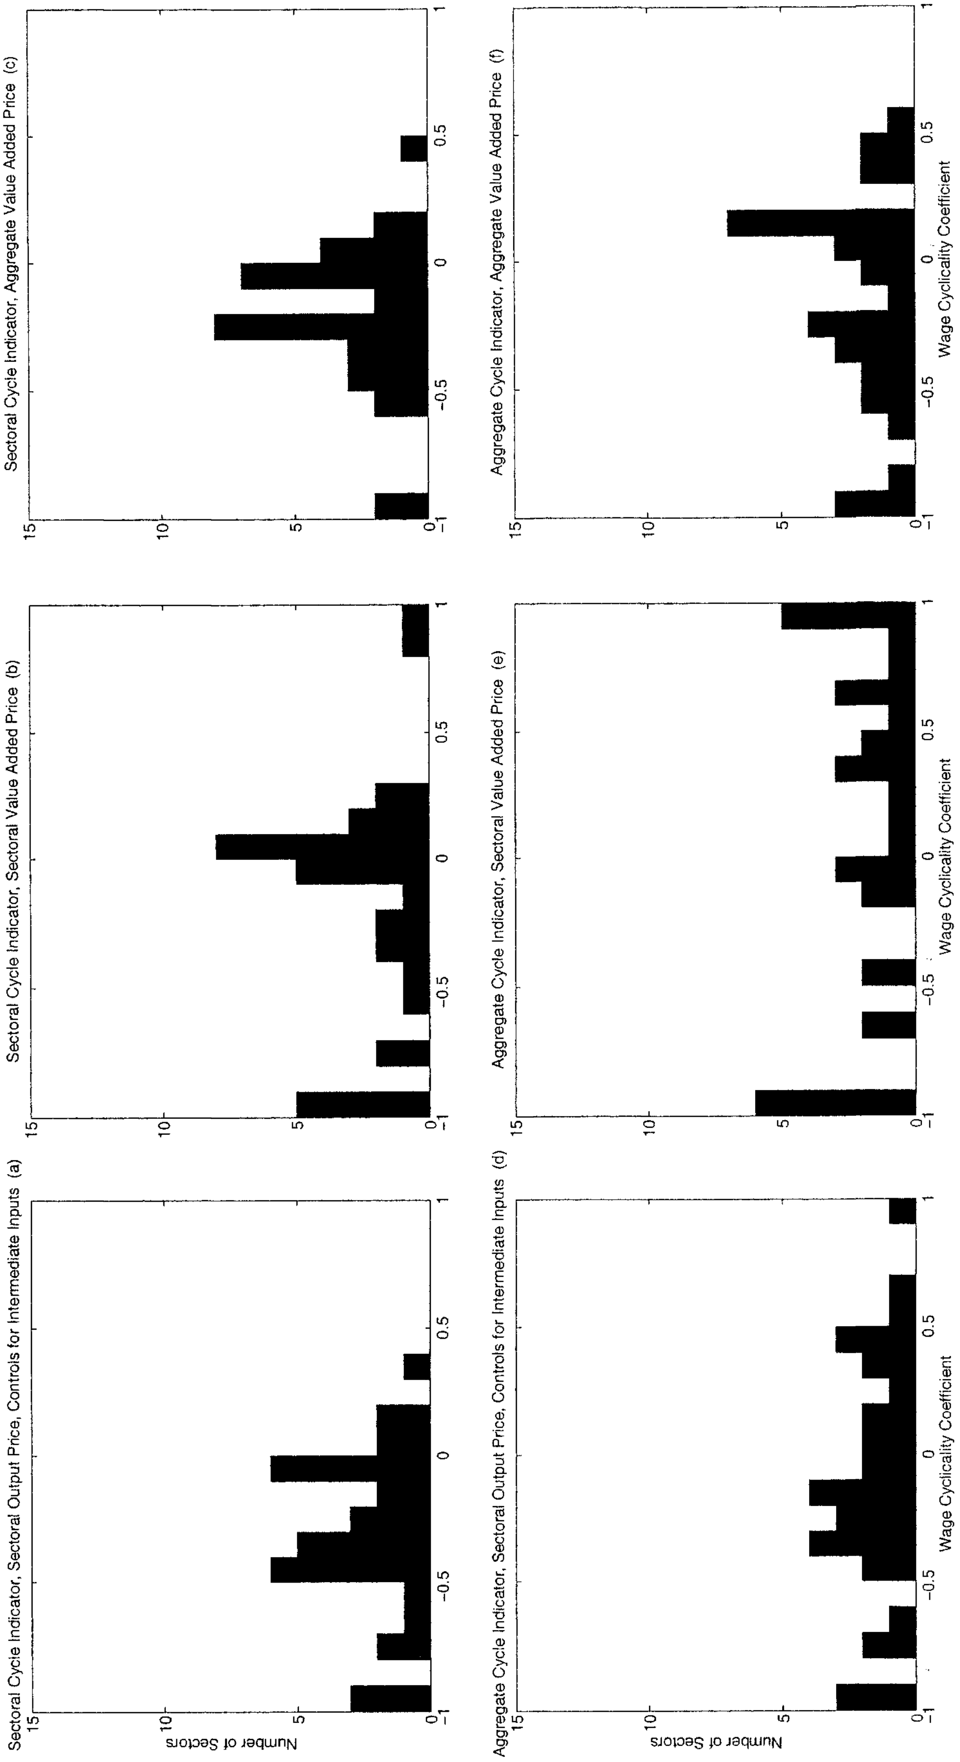

TABLE 2.—SUMMARY STATISTICS FOR 34 KLEM DATABASE INDUSTRIES

	Firm-Oriented (a)	(b)	(c)
Countercyclical:	29	19	27
at 10% level:	13	6	5
at 5% level:	7	5	4
at 1% level:	5	4	3
Procyclical:	5	15	7
at 10% level:	0	1	1
at 5% level:	0	0	1
at 1% level:	0	0	1
Mean:	−.369	−.270	−.248
Weighted mean:	−.330	−.101	−.227
Median:	−.332	−.073	−.235
Mean abs. t-stat.:	1.302	1.232	1.141

	(d)	(e)	Typical Nonfirm (f)
Countercyclical:	21	15	19
at 10% level:	4	3	1
at 5% level:	3	2	1
at 1% level:	1	2	1
Procyclical:	13	19	15
at 10% level:	1	3	2
at 5% level:	1	2	1
at 1% level:	0	1	1
Mean:	−.128	−.062	−.165
Weighted mean:	−.167	.041	.033
Median:	−.137	.192	−.112
Mean abs. t-stat.:	0.908	1.266	0.895

	Completely Aggregate (g)
Coefficient:	.013
Std. error:	.123

Panels (a)–(f) correspond to figure 4; panel (g) is for the aggregate economy

Second, this paper's finding of extremely poor estimates using sectoral value-added data [specifications (b) and (e)] is consistent with earlier work by Jorgenson et al. (1987) and Basu and Fernald (1997). These authors derive the necessary conditions for a well-defined concept of value added to exist, show that these conditions are often not met empirically in sectoral data, and find that imposing the value-added restrictions on the sectoral data yields poor empirical results.

Third, the present paper's findings for specification (c) are in line with earlier work by Waud (1968). Waud uses quarterly data on sectoral production worker hours and wages, deflated by the aggregate GNP deflator, to investigate wage-employment correlations in seventeen two-digit manufacturing industries.[24] Waud also stresses the importance of using a sectoral rather than an aggregate indicator of economic conditions: "The number of cycles n, of course, varies from industry to industry . . ." (p. 414). Using specification (c), Waud finds, as this paper does, very strong evidence of an inverse (that is, countercyclical) relationship between real wages and hours of employment, with the estimated coefficients being negative in 16 out of the 17 industries in his data set.

[24] Unlike the present study, Waud does not have data on industry output prices and input prices with which to run the more correct specification (a), nor does he compare his results with those that are obtained using the more typical non-firm-oriented specification (f).

Fourth, this paper's findings for specifications (e) and (f) are consistent with earlier work by Solon and Barsky (1989). Those authors use one-digit-level data on wages to examine real wage cyclicality with respect to the aggregate unemployment rate, and use both the aggregate GNP deflator and one-digit-level producer price indices (with no controls for intermediate input prices) as their measure of prices. Like the present paper, they find little evidence of cyclicality in either direction using those specifications.

Fifth, this paper's findings for the preferred, firm-oriented specification (a) are very much in line with the results of Pencavel and Craig (1994). Pencavel and Craig look at only a single manufacturing sector—plywood manufacturing in the Pacific Northwest—but they do so at the firm level for a very homogeneous set of workers and firms. They examine the relationship between employment, hours, wages, intermediate inputs (in their case, tree logs), and output prices for 35 firms for each of the years 1968 through 1986. Although they frame their study as a test of profit maximization under the maintained assumption that their firms' production functions are stable over the sample period, one can flip the maintained and tested hypotheses, and interpret their results as a test of whether their firms' production functions are stable over time, under the assumption that they maximize profits. Pencavel and Craig do not reject the hypothesis of stable labor demand curves over their sample period, and document a very strong inverse (that is, coun-

TABLE 3.—KLEM DATABASE COEFFICIENTS BY INDUSTRY

Industry	Coefficient	Std. Error	t-Stat.
1 Agriculture	−1.316	0.720	−1.829
2 Metal mining	−0.374	0.201	−1.862
3 Coal mining	−0.725	0.262	−2.769
4 Oil and gas extraction	−0.544	0.196	−2.771
5 Nonmetallic mining	−0.264	0.244	−1.081
6 Construction	−0.462	0.150	−3.073
7 Food & kindred products	−0.130	0.487	−0.267
8 Tobacco products	−0.046	0.524	−0.087
9 Textile mill products	0.193	0.433	0.445
10 Apparel	−0.019	0.195	−0.099
11 Lumber & wood products	−0.460	0.247	−1.858
12 Furniture & fixtures	0.030	0.197	0.154
13 Paper & allied products	0.188	0.295	0.638
14 Printing, publishing, & allied	−0.193	0.175	−1.103
15 Chemicals	−0.499	0.412	−1.211
16 Petroleum & coal products	−2.331	1.330	−1.752
17 Rubber & plastic products	−0.379	0.167	−2.272
18 Leather products	−0.461	0.256	−1.801
19 Stone, clay, glass products	−0.324	0.194	−1.667
20 Primary metal	−1.012	0.389	−2.602
21 Fabricated metal products	−0.052	0.168	−0.310
22 Machinery, nonelectrical	0.021	0.260	0.082
23 Machinery, electrical	−0.220	0.211	−1.039
24 Motor vehicles	−0.097	0.296	−0.329
25 Transport. equip. & ordnance	−0.341	0.158	−2.158
26 Instruments & electronics	−0.244	0.171	−1.421
27 Misc. manufacturing	0.337	0.404	0.833
28 Transportation	−0.496	0.262	−1.892
29 Communications	−0.033	0.180	−0.185
30 Utilities, electric	−0.473	0.194	−2.440
31 Utilities, gas	−0.775	0.522	−1.484
32 Wholesale & retail trade	−0.068	0.288	−0.238
33 Finance, insurance, & real estate	−0.647	0.400	−1.617
34 Services	−0.343	0.376	−0.912

Using firm-oriented specification (a)—product wages, industry conditions, and controls for intermediate input variation.

tercyclical) relationship between real wages (deflated by plywood output prices) and production worker hours, controlling for changes in intermediate inputs. Their findings are thus very supportive of the present paper's results—in fact, I am able to corroborate their findings both for the specific industry they consider (table 3) and for many two-digit sectors of the economy more generally.

A. *Panel Studies of Workers*

There is some discrepancy, however, between the findings of this paper and those of researchers using panel data on workers, such as Solon et al. (1994), Bils (1985), and Keane et al. (1988). Controlling for changes in worker composition using panel data, those authors find significantly procyclical real wages, whereas, using the composition-adjusted KLEM data set, I am unable to replicate their finding. There are a number of reasons for this discrepancy.

First and most importantly, the present paper uses detailed firm product prices [specifications (a), (b), (d), (e)] and detailed measures of industry conditions [specifications (a), (b), (c)], whereas the panel studies of workers use an aggregate price index and the aggregate unemployment rate, as in specification (f), which was found to be the most procyclical specification of all those considered in this paper.

Second, the worker panel studies' sample period is the 1970s and 1980s, due to data availability, and this period was one of substantially more procyclical wages than others in U.S. history, probably due to the large changes in oil prices over the period. For example, specification (g) with the KLEM data restricted to 1967–1988 (Solon et al.'s sample period) yields a coefficient of .143, compared to .013 for the full 1958–1996 period.[25]

Third, the strong wage procyclicality documented by the worker panel studies holds for men, but generally not for women,[26] whereas the KLEM data indices cover all labor input, including women. This difference in sample population leads to more procyclicality in the panel studies than will be found in the KLEM data.

Fourth, the KLEM data and the panel studies treat workers who change jobs very differently. A worker who changes jobs from a low-paying sector (such as services) to a high-paying sector (such as manufacturing) will contribute a large positive wage change to the worker panel data estimate, but the same worker changing between the same jobs has no effect on the KLEM wage indices—intuitively, the price of labor in services and the price of labor in manufacturing have not changed.[27] Because a large fraction of real-wage procyclicality in the panel data is attributable to workers who change jobs (Bils, 1985), the worker panel studies yield estimates that are more procyclical than the KLEM data.

Fifth and finally, the panel studies weight all workers equally, whereas the KLEM indices give more weight to high-income workers, who generally have less cyclical wages (Swanson, 1998).[28] This difference in sample population will also lead to estimates of greater wage cyclicality in worker panel data than in the KLEM data.

The last two points highlight the fact that Jorgenson's KLEM project and the panel studies of workers cited above

[25] Note that specifications (a) and (d) of this paper control for changes in oil prices—something the panel studies of workers do not do—and find coefficients and results that are stable across the two time periods. This finding provides further corroboration that the firm-oriented specification (a) of this paper successfully removes the effects of nontechnological shocks from the analysis.

[26] See Table II and pp. 14–15 of Solon et al. (1994). This difference in cyclicality may be due to the fact that men are more likely to hold jobs in sectors that use extra shifts and overtime (Swanson, 1998).

[27] More technically, as far as the KLEM data set is concerned, if a worker's wages change as a result of moving across cells, this counts as a change in the composition of labor input, rather than as a change in wages; only a change in wages within a given cell counts toward a change in wages for the sector.

[28] The KLEM (Törnqvist) index methodology weights workers' log wage changes by their shares in aggregate income (essentially their marginal product times hours worked). The latter approach is more appropriate from a production-theoretic point of view, because the reason a worker earns a higher wage is presumably that he or she contributes more to the production of output. Because the objective of this paper has been to test for movements in firm's labor demand curves, the production-theoretic KLEM definition is more appropriate than the equal-weighted panel approach.

have taken different approaches to composition control that have different effects on the measurement of wage cyclicality. The worker panel studies have naturally adopted a more worker-oriented approach, whereas the KLEM indices were developed for a production-theoretic analysis, which happens to coincide closely with the requirements for a hypothesis test of labor demand shifts, as conducted in the present paper.[29]

VI. Discussion

The results above lead to a number of observations about wages, markups, and labor demand shifts at cyclical frequencies within individual industrial sectors of the economy. For example, on the basis of the results above, there seems to be little reason to reject the classical (or Keynesian) hypothesis that firms' labor demand curves are stable at cyclical frequencies. In particular, technology shocks do not seem to be playing a major role in real wage and employment fluctuations at the detailed industry level.

Similarly, models of countercyclical markups, to the extent that they predict a positive correlation between real wages and employment within individual industrial sectors, also do not seem to be supported by the data.[30] Because these models are often formulated in terms of procyclical price wars by firms within detailed sectors of the economy (Rotemberg & Saloner, 1986; Rotemberg & Woodford, 1992), the findings of the present paper bring direct evidence to bear on these theories.[31]

A. *Aggregate Versus Industry Cycles*

One must take some care when relating the empirical findings of this paper to business cycles in the aggregate economy. For example, it is possible that idiosyncratic demand shocks (such as changes in consumer tastes or, as described in footnote 3, shocks to monetary policy, government spending, exchange rates, or oil prices that affect demand for some industries' products more than others) dominate most industries most of the time—indeed, the empirical findings of this paper seem to indicate that this is the case. However, the idiosyncratic components of these shocks will tend to average out in the aggregate, reducing their importance for explaining fluctuations in aggregate employment and aggregate output. Thus, despite our finding little or no role for technology shocks at the detailed industry level, it might still be the case that technology shocks are an important source of fluctuations for aggregate U.S. economic data. For example, it might be the case that aggregate technology shocks are the only shocks that do not cancel out in the aggregate, despite being relatively rare and of secondary importance at the detailed industry level.[32]

Figure 5 presents some evidence that this story of idiosyncratic demand shocks but aggregate technology shocks is not taking place in the data. If this type of story were true, one would expect industrial sectors that are more correlated with the aggregate business cycle to have, in general, more procyclical real wages, or at least less countercyclical real wages, than sectors of the economy that are highly idiosyncratic. Figure 5 plots each sector's estimated wage cyclicality coefficient as a function of each sector's correlation with the aggregate business cycle. The horizontal axes in the figure represent the R^2 from a regression of employment changes in each of the 458 NBER Productivity Database sectors on a constant and the employment change for the aggregate economy; sectors with high R^2 values are *cyclical* industries (such as automobile-related industries), whereas sectors with low R^2 values are highly *idiosyncratic* (such as sectors that manufacture food and kindred products or tobacco).[33]

The top panels in figure 5 plot the estimated wage cyclicality coefficient for each sector, using regression specification (a) from section IV. The left-hand panel presents the raw coefficient estimates, and the right-hand panel presents the corresponding t-statistics, with dashed reference lines at ± 2. I have also fitted a solid regression line to the data points in each panel, to serve as a summary reference line.

As can be seen in the figure, there does not appear to be any tendency for more cyclical industries to have more procyclical real wages. Aggregate technology shocks do not appear to be driving real-wage and employment movements even in the sectors that are most representative of the aggregate economy.

In contrast, there is some evidence that the measurement error and econometric specification error emphasized in section II are playing a role. The bottom panels of figure 5 plot sectoral wage cyclicality coefficients against industry cyclicality, this time using the less preferred specification (d) for the vertical coordinate of each point (recall that this specification used changes in *aggregate* employment as the

[29] It should also be emphasized that, even if composition change were insufficiently allowed for in the KLEM data, this paper's conclusions across specifications in figure 4 (and figure 3) would still be valid, because composition bias would be present to the same extent in each panel of the figures. The fact that the present paper's results are corroborated by the micro-level study of Pencavel and Craig (1994) lends additional weight to the claim that the findings here are robust.

[30] Only models with very strongly countercyclical markups generate a procyclical real wage. Mildly countercyclical markups are not sufficient to offset a procyclically varying marginal cost.

[31] A detailed comparison of the results of this paper with those in Bils (1987) is provided by Swanson (1998). In brief, I found that Bils's cubic detrending procedure did a poor job of extracting the cyclical components from the very nonstationary sectoral data; HP-filtering and first-differencing did much better. Also, though Bils found a number of significant coefficients for wages and prices separately, his methods yielded generally insignificant estimates for real wages (wages divided by prices). HP filtering and first-differencing yielded significantly countercyclical real wages in the large majority of his sectors.

[32] I thank an anonymous referee for suggesting discussion of this point.

[33] There are few sectors with a negative raw correlation with the aggregate economy, and no sectors with a large negative correlation. The results are very much the same using raw correlations along the horizontal axis instead of R^2 values.

FIGURE 5.—WAGE CYCLICALITY FOR IDIOSYNCRATIC VERSUS CYCLICAL SECTORS

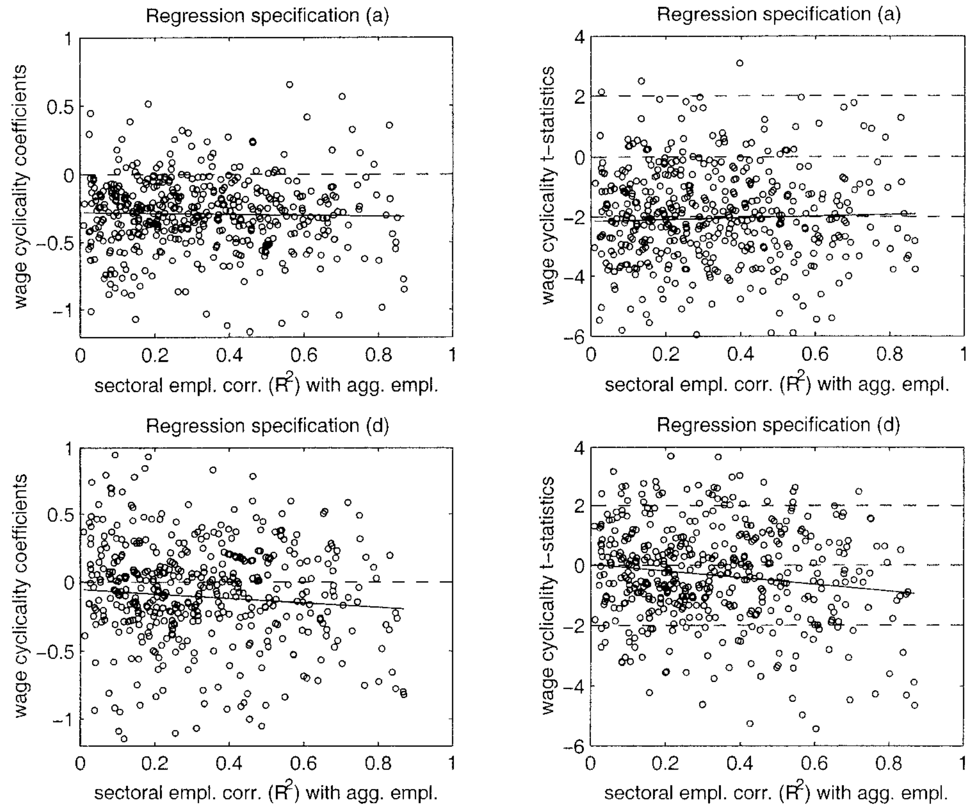

cycle indicator, rather than changes in sectoral employment).[34] The results across the two specifications are more consistent when the aggregate business cycle is a better indicator of conditions in the sector—that is, as we move to the right in the figure.

How can wage cyclicality from the firm point of view be consistently countercyclical when more aggregate measures [such as specification (g)], or results from the panel studies of workers cited earlier, are acyclical or even procyclical? There are a number of reasons, most of them compositional. Reasons why results for completely aggregate specification (g) were acyclical were discussed in section IV—in particular, there was evidence that changes in industry composition over the cycle (high-wage industries tend also to be more cyclical) were playing a role.[35] For the panel studies, many factors seemed to be at work, including their sample period, a specification that didn't control for oil price changes, and the fact that much of the cyclicality that they find is due to workers changing jobs (Bils, 1985), many across detailed sectors of the economy. The findings of this paper are thus consistent with those of Basu and Fernald (1997), who found that much of the procyclicality of aggregate productivity could be attributed to cyclical changes in the industrial composition of the U.S. economy over the business cycle, rather than a procyclical Solow residual in each sector separately.

VII. Conclusions

On the basis of these results, and in contrast to many previous studies of wage cyclicality, I find little reason to reject the classical or the Keynesian model of wage and employment determination. I have tried to improve on earlier studies by (1) using detailed industry data (on wages, output prices, intermediate inputs, and industry conditions) to capture the variables that theory says are relevant for testing the predictions of the model, (2) investigating the effects of using aggregate proxies for these sectoral variables, and (3) controlling for changes in workforce composition.

It should be emphasized that acceptance of the classical model of countercyclical real product wages does not require workers to suffer from money illusion or rigid nominal wage contracts (though these may be important), nor does it require rejecting the preponderance of empirical evidence that labor productivity is procyclical at the plant and industry level (see, for example, Foster, Haltiwanger, & Krizan, 1998; Bartelsman & Dhrymes, 1994). A multisector classical model can be consistent with all of these observations. For example, a positive fiscal spending shock that

[34] The horizontal coordinates (R^2 measures) are the same as in the top panels of the figure. Results using regression specification (f) are quite similar.

[35] See the discussion of the KLEM results in section IV.

impacts one sector of the economy more than others can be seen as leading to an increase in the price of that sector's good, a corresponding decrease in that sector's real wage deflated by its product price, and hence an increase in employment and the utilization of capital (and labor) in that sector. This change in capital and labor utilization is consistent with an increase in measured labor productivity in the sector, despite the fall in real wages, properly deflated. These effects can be demonstrated rigorously in a more fully specified general equilibrium framework (Swanson, 1999). The results of this paper suggest that further empirical and theoretical work in this direction might be illuminating.

REFERENCES

Abel, Andrew, and Ben Bernanke, *Macroeconomics* (Boston: Addison-Wesley, 2001).

Abraham, Katherine, and John Haltiwanger, "Real Wages and the Business Cycle," *Journal of Economic Literature* 33 (1995), 1215–1265.

Barro, Robert, and Robert King, "Time-Separable Preferences and Intertemporal-Substitution Models of Business Cycles," *Quarterly Journal of Economics* 99 (1984), 817–839.

Bartelsman, Eric, Ricardo Caballero, and Richard Lyons, "Customer- and Supplier-Driven Externalities," *American Economic Review* 84 (1994), 1075–1084.

Bartelsman, Eric, and Phoebus Dhrymes, "Productivity Dynamics: U.S. Manufacturing Plants, 1972–86," Federal Reserve Board finance and economics discussion series no. 94-1 (1994).

Bartelsman, Eric, and Wayne Gray, "The NBER Manufacturing Productivity Database," NBER technical working paper no. 205 (1996).

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

Basu, Susanto, John Fernald, and Miles Kimball, "Are Technology Improvements Contractionary?" Federal Reserve Board International Finance discussion paper no. 625 (1998).

Bils, Mark, "Real Wages over the Business Cycle: Evidence from Panel Data," *Journal of Political Economy* 93 (1985), 666–689.

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

Chirinko, Robert, "The Real Wage Rate over the Business Cycle," this REVIEW 62 (1980), 459–461.

Foster, Lucia, John Haltiwanger, and C. J. Krizan, "Aggregate Productivity Growth: Lessons from Microeconomic Evidence," NBER working paper no. 6803 (1998).

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

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

Keane, Michael, Robert Moffitt, and David Runkle, "Real Wages over the Business Cycle: Estimating the Impact of Heterogeneity with Micro Data," *Journal of Political Economy* 96 (1988), pp. 1232–1266.

Keynes, John Maynard, *The General Theory of Employment, Interest, and Money* (London: MacMillan, 1936).

King, Robert, Charles Plosser, and Sergio Rebelo, "Production, Growth, and Business Cycles," *Journal of Monetary Economics* 21 (1988), 195–232.

Kydland, Finn, and Edward Prescott, "Time to Build and Aggregate Fluctuations," *Econometrica* 50 (1982), 1345–1370.

Mankiw, N. Gregory, *Macroeconomics* (New York: Worth Publishers, 2002).

Murphy, Kevin, Andrei Shleifer, and Robert Vishny, "Building Blocks of Market Clearing Business Cycle Models," in *NBER Macro Annual* (1989).

Pencavel, John, and Ben Craig, "The Empirical Performance of Orthodox Models of the Firm: Conventional Firms and Worker Cooperatives," *Journal of Political Economy* 102 (1994), 718–744.

Rotemberg, Julio, and Garth Saloner, "A Supergame-Theoretic Model of Price Wars during Booms," *American Economic Review* 76 (1986), 390–407.

Rotemberg, Julio, and Michael Woodford, "Oligopolistic Pricing and the Effects of Aggregate Demand on Economic Activity," *Journal of Political Economy* 100 (1992), 1153–1207.

Solon, Gary, and Robert Barsky, "Real Wages over the Business Cycle," NBER working paper no. 2888 (1989).

Solon, Gary, Robert Barsky, and Jonathan Parker, "Measuring the Cyclicality of Real Wages: How Important is Composition Bias?" *Quarterly Journal of Economics* 109 (1994), 3–25.

Swanson, Eric, "Individual and sectoral heterogeneity, reallocation, and aggregate fluctuations," doctoral dissertation (Stanford, CA: Stanford University, 1998).

Swanson, Eric, "Models of Sectoral Reallocation," Federal Reserve Board finance and economics discussion series no. 99-03 (1999).

Waud, Roger N., "Man-Hour Behavior in U. S. Manufacturing: A Neoclassical Interpretation," *Journal of Political Economy* 76 (1968), 407–427.