

Measuring the Cyclicalities of Real Wages: How Important Is Aggregation Across Industries?

Eric T. Swanson
Board of Governors of the Federal Reserve System
eswanson@frb.gov

Abstract

There is a growing consensus among economists that real wages in the post-war U.S. have been moderately to strongly procyclical, particularly in panel data on workers. From the point of view of hiring decisions of firms, however, this conclusion may be premature or even erroneous. Whether a firm's labor demand curve is stable or shifting at business cycle frequencies should be tested with a wage that is deflated by the firm's own price of output, with appropriate controls for the prices of intermediate inputs, and with respect to the cyclical state of the firm's own industry, as opposed to the state of the aggregate economy. I find that failing to control for these factors has led to a substantial procyclical bias in previous estimates of wage cyclicalities. In two-digit and four-digit level (SIC) industry data on wages, with controls for changes in worker composition, I find that a substantial majority of sectors have paid real product wages that vary inversely (*i.e.*, countercyclically) with the state of their industry.

JEL Classification: E24, E32, J23

Version 0.2.1
October 19, 1999

This paper is based on Chapter 3 of my doctoral dissertation at Stanford University. I thank Michael Horvath, Bob Hall, John Taylor, John Pencavel, Mun Ho, Kevin Stroh, Charlie Fleischman, and seminar participants at Stanford and the Federal Reserve Board for helpful discussions. All errors and omissions are mine. The views expressed in this paper are solely those of the author and do not necessarily represent those of the Federal Reserve System or its Board of Governors.

1 Introduction

What causes business cycles? This is one of the most fundamental questions facing macroeconomists, yet potential 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: “Economic theory implies that, given the nature of the shocks to technology and people’s willingness and ability to intertemporally and intratemporally substitute, the economy will display fluctuations like those the U.S. economy displays,” (Prescott (1986), p. 21). Their basic model, and the standard variations of it, imply a positive correlation between real wages and the quantity of labor employed:¹

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 and King (1984), pp. 832–3)

In contrast, the Classical and Keynesian explanations for changes in employment over the business cycle 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’ 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

¹One can think of this correlation as the result of large, technology-driven shifts in firms’ labor demand curves, while the labor supply curves of workers remain relatively stable, as follows. Although workers’ labor supply curves at time t may shift markedly in response to technology shocks, as the future path of interest rates changes, this effect under standard parameterizations is not enough to offset the positive relationship between real wages and employment that typical RBC models imply. Thus, it is more useful to think about the model in terms of a *reduced-form labor supply curve*, which corresponds to the locus of equilibria between real wages and employment that obtain as we range through all possible values for the technology shock. We may then think of the technology shock as shifting firms’ labor demand curves along this stable reduced-form locus.

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 infeasible. (Keynes (1936), p. 17).

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

A potentially powerful test between these two sets of theories thus lies in the cyclicity of real wages, with the aim of determining whether firms' labor demand curves are shifting at business cycle frequencies, or simply trending outward smoothly over time. Because such a test fundamentally involves the labor demand of firms, to perform the test correctly, one should use detailed firm- or industry-level data on wages and employment, deflating wages by the firm- or industry-level price of output, and controlling for variation in the prices of intermediate inputs.²

Previous studies have not, however, controlled for these factors. Historically, the typical investigation of wage cyclicity has used aggregate data, such as the BLS's average hourly earnings statistic, regressed on an indicator for the cyclical state of the aggregate economy, such as the unemployment rate or aggregate employment. These studies find only a very weak correlation between wages and the state of the aggregate economy, with the correlation usually being slightly procyclical—see Abraham and Haltiwanger (1995) for a survey. Because of the failure to control for the factors mentioned above, however, one might expect these aggregate estimates to be imprecise, or even biased. I will demonstrate below that these fears are justified.

²One might object that Barro and King's analysis is very much a representative-agent, representative-firm argument, and thus might not be inconsistent with other predictions when carried out within a multi-sector framework. Indeed, based on my results below, I will argue for exactly that approach. In the meantime, let it suffice to note that the Classical and Keynesian theories are also very much representative-agent, representative-firm models, so that a comparison of the predictions of the two classes of models (Classical/Keynesian and Technology-driven) is not inappropriate.

These findings of real wage procyclicality, or, at least, an absence of countercyclicality, have led many economists to question the usefulness of the Classical or Keynesian models for explaining business cycles. Findings of even greater real wage procyclicality in panel studies of workers (such as Bils (1985) and Solon, Barsky, and Parker (1994)) have strengthened the case for alternative business cycle models even further.³ However, these panel studies continue to adopt a nonfirm-oriented point of view—in particular, they deflate wages by the aggregate GNP deflator, or a similar aggregate index, and measure cyclicity with respect to an indicator of aggregate economic conditions, rather than the conditions of the industry in which the worker and firm are employed. I find below that this introduces a significant positive bias into the results, if one’s question of interest is regarding the labor demand of firms.

The paper is divided into six sections. Section two derives a model of labor demand and discusses the expected effects of substituting various aggregate measures of wages, prices, and cyclical indicator into the model. Section three describes the two datasets used to investigate these sectoral real wage cyclicity relationships: the 458-sector NBER Productivity Database, and Jorgenson’s 34-sector KLEM dataset. Section four runs the regressions suggested by the model and discusses the results. Section five compares and contrasts the results to other studies in the literature. Section six concludes.

2 A Simple Model of Labor Demand

Consider the case of a profit-maximizing firm in an industry with perfectly competitive output and factor markets. Let the firm face a production function $y_t = F(k_t, l_t, m_t, e_t, t)$, where y denotes output, k capital, l labor, m materials input, e energy input, and t time. Assume that k is fixed at any point in time, and that $F(k_t, \cdot, \cdot, \cdot, t)$ is increasing, twice-differentiable, and concave in its middle three arguments. In addition, assume that the capital stock is chosen at the beginning of each period, before shocks to prices and wages

³One of the primary advantages of these panel studies is that they can control for changes in workforce composition over the business cycle—*i.e.*, recessions are times when workers with the lowest seniority and least amount of firm-specific human capital, and hence the lowest wages, tend to be laid off. This composition effect will tend to impart a countercyclical movement in average measures of wages over time, all else equal. I will take great care to account for these composition effects in the regressions below.

are realized; that the function F is homogeneous of degree one in its first four arguments at all times t ; and that firms may enter and exit the industry freely at the beginning of each period. This implies that the optimal dynamic profit-maximizing strategy of the firm is simply to maximize profits period by period. Then, at any time t , there is a well-defined mapping from the prices of output, labor, materials, and energy to the profit-maximizing quantities of these three variables.

If the function F is Cobb-Douglas, so that $y_t = f(t)k_t^\mu l_t^\nu m_t^\theta e_t^\varphi$, it is straightforward to show that

$$\log l_t = a \log \frac{w_t}{p_t} + b \log \frac{p_{m_t}}{p_t} + c \log \frac{p_{e_t}}{p_t} + g(t, k_t) \quad (1)$$

for some constants a , b , and c , where p_t denotes the firm's price of output, w_t the price of labor, p_{m_t} the price of materials, p_{e_t} the price of energy, and the function g encompasses a constant and terms relating to t and k_t . Equation (1) may also be regarded as a log-linear approximation to the labor demand curve resulting from a completely general production function F with factor-augmenting technical change.

Rearranging terms in (1) yields

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

This specification agrees more closely with others in the wage cyclicality literature.

Standard microeconomic theory implies that the coefficient a must be negative, and hence $\alpha < 0$.⁴ Under the Classical or Keynesian assumption that technology and capital stocks are stable at business cycle frequencies, the function $g(t, k_t)$ trends smoothly over time, and thus may be well approximated by a polynomial or Hodrick-Prescott trend $h(t)$. With this assumption, the equation

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

is a stable relationship that will serve as the basis for my empirical work below.⁵ This is essentially the same specification that has been used by previous researchers, with the

⁴In addition, assuming materials and labor are complements in production, we will also have $b < 0$, and hence $\beta > 0$, and assuming energy and labor are complements, $\gamma > 0$.

⁵Note that, under the assumption that $h(t)$ captures all relevant shifts of the labor demand curve due to changes in capital and technology, there is no need to instrument for any of the variables in equation (2).

addition of the materials and energy price terms on the right-hand side. Note that, for reasons given in the Data section below, I will typically estimate (2) in first differences rather than in levels.

If the assumptions underlying the model 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 fluctuations in firm employment and real wages, equation (2) will not satisfy the classical regression assumptions—in particular, the error term will be strongly and positively correlated with $\log l_t$. Thus, even though the theory predicted that the coefficient α should be less than 0, standard regression procedures applied to (2) may yield estimates of α that are positively biased, insignificantly different from 0, or even greater than 0.⁶ Models in which technology is the dominant source of fluctuations in employment would predict this last effect.

In fact, previous empirical work using regression specifications similar to (2) have often estimated values of α that are positive. This has led many to question whether the Classical and Keynesian assumptions regarding perfect competition and smoothly trending technology are appropriate. The case for models of real business cycles, and for models of countercyclical markups, have thus been significantly bolstered by these findings.

It is the point of this paper that these conclusions may be premature or even erroneous. The empirical studies referred to above have typically used highly aggregate data on prices, wages, and labor input, with no allowance for changes in the cost of intermediate materials and energy. This will be shown below to have important and positively biased effects on the estimated value of α . Even the panel data studies of Bils (1985) and Solon, Barsky, and Parker (1994) have continued to use aggregate rather than sectoral data on prices and industry conditions (*i.e.*, the cyclical indicator), with no controls for

⁶The sign of the bias on α in (2) depends on the interaction between $\log l_t$, $\log(p_{m_t}/p_t)$, $\log(p_{e_t}/p_t)$, 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 (2). Technology shocks in a given sector will tend to induce $\rho_{l\varepsilon} > 0$ and $\rho_{x\varepsilon} > 0$, the latter because price tends to move inversely to the technology shock. Empirical studies such as Murphy, Shleifer, and Vishny (1989) and the countercyclical markups literature 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 $\rho_{l\varepsilon}$. Since *a priori* there is no reason to think $\rho_{x\varepsilon} > \rho_{l\varepsilon}$ —in fact, since materials prices, particularly energy, tend to be more volatile than employment, we might expect $\rho_{x\varepsilon} < \rho_{l\varepsilon}$ to hold—the expected bias on α is positive.

variation in intermediate input prices. While these choices of deflator and right-hand side variables may be relatively ancillary when the main question of interest is the effect of cyclical changes in workforce composition on aggregate wages (or may even be appropriate choices for these variables when the question of interest involves the effects of recessions on workers' standards of living, or the labor supply decisions of individuals), when testing the Classical or Keynesian implications for the employment-real wage relationship, as described above, it is clear that we should use firm- or industry-level data on prices and quantities.

These are potentially serious sources of bias. Consider, for example, the effects of using an aggregate indicator L_t as a proxy for the state of a particular sectoral labor market, l_t . It is clear that, 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 estimation of equation (2) with L_t instead of l_t will lead to a classical errors-in-variables bias. We would thus expect the estimate of α obtained in this way to be both biased toward zero and imprecisely estimated.

Second, suppose we attempted to estimate a variant of equation (2) using sectoral value added deflators, rather than data on gross output and input prices. Our specification would have to be modified as follows:

$$\log \frac{w_t}{p_t} = \alpha \log l_t + h(t) + \varepsilon_t \quad (3)$$

where p_t here refers to the sectoral “price of value added” rather than that of gross output. We would expect two effects from running regression (3) instead of (2). First, to the extent that changes in the price of value added do not fully capture changes in the price of intermediate inputs (and with the large changes in energy prices over this period, this is a valid concern), we would expect estimates of α in (3) to suffer partially from an omitted-variables bias (the controls for intermediate input price variation that are present in equation (2) being effectively only partially implemented in the value-added specification of equation (3)). Second, to the extent that value-added deflators are noisier than gross output prices at the sectoral level (as is found in the data),⁷ we would expect the residuals

⁷In fact, sectoral value-added deflators are found to have higher variance than both gross output prices and intermediate input prices in both of the datasets below.

in equation (3) to have a higher variance, and estimates of α from that equation to be more imprecise.

Finally, we could go one step further and run regression specification (3) using an aggregate value added deflator P_t as our measure of prices. Assuming that the aggregate “price of value added” P_t is a poor proxy for the true, underlying sectoral price of value added p_t (so that $P_t = p_t + \eta_t$), we would expect this substitution to increase the variance of the left-hand-side variable in (3). We would thus expect a further increasing of the variance of the residuals in that equation, and a further deterioration in the accuracy of the estimates of α , from what would be observed using sectoral data on p_t in that equation.

The implications of estimating equation (2) with aggregate rather than sectoral data are thus potentially serious. Not surprisingly, all of the misspecifications considered above anticipate a deterioration in the quality of the estimates. More importantly, some of the stated misspecifications would also be expected to lead to a substantial errors-in-variables or omitted-variables bias in the results. The empirical and practical importance of these factors is considered below.

3 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 is not presently available. The NBER Productivity Database (available from the NBER’s website, and documented in Bartelsman and Gray (1996)), however, does contain sectoral data for all manufacturing industries at the 4-digit (SIC) level between 1958 and 1994 at annual frequency. I will also use data at roughly the 2-digit level compiled by Dale Jorgenson and his coworkers (available from Jorgenson’s website, and documented in Jorgenson, Gollop, and Fraumeni (1987)), because its method of construction has some unique advantages and because having a second dataset 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 ways, 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 as a whole, and yet not to its individual firms. By looking at detailed industry data instead of firms or plants, we 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.⁸

Data on all variables in the NBER Productivity Database is at the 4-digit (SIC) level unless otherwise noted. 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 is generally unavailable—even at the two-digit level, it 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 materials price indexes for each sector in the NBER Productivity Database: energy prices, and the price of all other materials. For the energy price, I use the price deflator for energy (PIEN); 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 materials (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 Törnqvist

⁸ While 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 quality and skills differ much more dramatically across broader sectors of the economy than they do within a detailed 4-digit manufacturing industry. Capital also varies greatly across sectors of the economy in terms of 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 4-digit manufacturing industry.

index (discrete approximation to Divisia) as in Jorgenson, Gollop, and Fraumeni (1987); finally, the price for nonenergy materials was then constructed as the nominal nonenergy materials expenditure ($\text{MATCOST} - \text{ENERGY}$) divided by this real index.

In addition to the product price deflators described above, I also constructed value-added deflators for purposes of comparison. The method is analogous to that used for nonenergy materials prices above: first, nominal gross output was defined as nominal value added (VADD) plus the cost of materials (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, indexes of real value added and value added deflators were constructed from these gross output and materials numbers using a Törnqvist methodology. A value-added deflator was also constructed for the aggregate manufacturing sector as a whole, using the Törnqvist index method once again.

In contrast to the NBER data, Jorgenson's KLEM data are at roughly the 2-digit (SIC) level for the manufacturing and mining sectors, and at the 1-digit level for other sectors of the economy. The data cover the years 1947–1991 at annual frequency. The KLEM dataset 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 indexes 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, etc.; 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); finally, these individual changes are aggregated into a single Törnqvist index number for each sector.⁹ The KLEM data thus control for cyclical changes in labor force composition, which was found by Solon, Barsky, and Parker (1994) to have important effects on aggregate measures of wages. Finally, the KLEM data provide us

⁹The details of construction for this and the other KLEM indexes are provided in Jorgenson, Gollop, and Fraumeni (1987).

with a corroboration and additional perspective on the results obtained using the NBER Productivity data. As with the NBER data, I also construct aggregate measures of wages, prices, and employment for the whole KLEM economy, using a Törnqvist index method.

In contrast to aggregate data, sectoral data at the two- or four-digit level 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.

4 Results

NBER Productivity Data

Summary results for the 458 sectors in the NBER Productivity Database are presented in Figure 1 and Table 1.¹⁰ Each graph in the Figure is a histogram of the 458 point estimates of the coefficient α for various specifications of regression equation (2); the panels differ in their choice of measures of p_t and l_t , and in the inclusion or exclusion of intermediate input prices (materials and energy) in the regression. The nominal wage measure is the same in each panel, and is the sector’s average production worker wage, as defined above. The first panel (1a) should be regarded as the “correct” specification under the hypotheses of section 2—the variable p_t is the sectoral price of gross output, l_t is the sector’s total production worker hours, and the sectoral prices of energy and of nonenergy materials are included as regressors. The last panel (1f) corresponds more closely to what has traditionally been run in the literature (in aggregate or panel studies of wage cyclicality,

¹⁰Results for different time periods, such as 1970–94, are very similar. In particular, all of the observations that follow hold for this later time period as well.

for example): the variable p_t is the value added deflator for aggregate manufacturing (as constructed above), l_t is total production worker hours for all of manufacturing, and the regressions include no controls for the prices of intermediate materials or energy. The panels (1b) through (1e) correspond to various intermediate specifications between these two extremes.

Table 1 presents summary statistics for the coefficient estimates for each panel in Figure 1, and are labeled accordingly. The term “countercyclical” in the Table refers to point estimates of α that are less than 0, while “procyclical” refers to estimates of α greater than 0. For each set of 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. To provide some idea of the magnitude of the coefficient estimates, in conjunction with Figure 1, the mean and median point estimates are presented, as well as a weighted mean, with weights given by the 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 precision of the estimates for each set of regressions.

The Figure and Table bring to light several important observations. First, the estimates for the “correct” specification in panel (1a) are much more often countercyclical ($\alpha < 0$) than procyclical ($\alpha > 0$). This is true of the raw point estimates, and at every level of statistical significance. The mean, weighted mean, and median coefficients are all roughly -0.1 , implying that a 1% increase in production worker hours is associated with a 0.1% decrease in real wages per hour paid by firms in the industry. Because production worker hours are relatively volatile (changing 10–15% over the course of a business cycle), the implied changes in real wages are roughly 1–1.5% over the course of a business cycle.

The effects of changes in worker composition are too small to explain the results in these regressions. Estimates using the KLEM data (in Table 2 below), which control for composition change, verify this directly. Alternatively, one can derive an estimate of the magnitude of composition change effects using panel data from the PSID, as in Solon, Barsky, and Parker (1994) and Chapter 2 of Swanson (1998). I find that the magnitude of these effects, regressed on total production worker hours (the cyclical indicator for the regressions in panels (1d), (1e), and (1f)), is about -0.025 , or only about one-fourth

the magnitude of the coefficients observed above.¹¹ Moreover, since the data in these regressions is for production workers only, changes in workforce composition may be less important for this relatively homogeneous population than for the economy as a whole. Finally, to the extent that changes in workforce composition do affect the estimates, they will affect the estimates in every panel of Figure 1, and hence will be differenced out when comparing any one panel of the Figure to the others.

The estimates in panel (1f) of Figure 1 correspond much more closely to regressions that have been run in previous studies of wage cyclicality. In particular, they use an aggregate manufacturing price index, do not control for variation in intermediate input prices, and use an indicator of aggregate rather than sectoral economic conditions; however, the regressions in (1f) differ from most previous studies in that they continue to use sectoral data on nominal wages. As predicted in Section 2, the estimates for this specification are substantially more procyclical and less accurate than for the correct specification in (1a). The mean, weighted mean, and median coefficients in (1f) are all virtually 0, lying in the range of $-.01$ to $-.03$; the number of procyclical point estimates is much larger than in (1a); the number of estimates that are statistically different from 0 is much smaller at every level of significance; and the mean absolute t-statistic for the regressions is dramatically smaller than for specification (1a). These results all suggest a very substantial omitted variables and errors-in-variables bias in estimation of equation (2) using the traditional approach.

Panel (1g) of Table 1 presents estimates of equation (2) using nothing but aggregate data (wages, hours, and value-added prices), constructed from the NBER Productivity Database as described in Section 3 above. No attempt is made to control for changes in intermediate input prices (this would be basically impossible to derive from the sectoral

¹¹This number is calculated as follows. An aggregate wage statistic for the PSID (household heads) is constructed as the sum of annual income for the sample divided by the sum of annual hours; this mirrors the construction of the wage statistics in the NBER Productivity Database and the construction of the aggregate hourly earnings measure by the BLS. This index is then first-differenced and regressed on a constant, time trend, and the change in log total production worker hours from the NBER Productivity Database, resulting in a coefficient on the last variable of .100. A comparable, composition-free index of wages is constructed as the median change in log wages that is observed in the sample of workers each year. When this index is regressed on a constant, time trend, and change in log total production worker hours, the coefficient on the last variable is .125. This implies that cyclical changes in workforce composition have an effect of roughly $-.025$ on the aggregate statistic, with respect to this cyclical indicator. Composition changes may have dramatically different effects in different sectors, however, so the KLEM estimates should be regarded as more authoritative.

Figure 1: Regression Coefficients for 458 NBER Productivity Database Industries

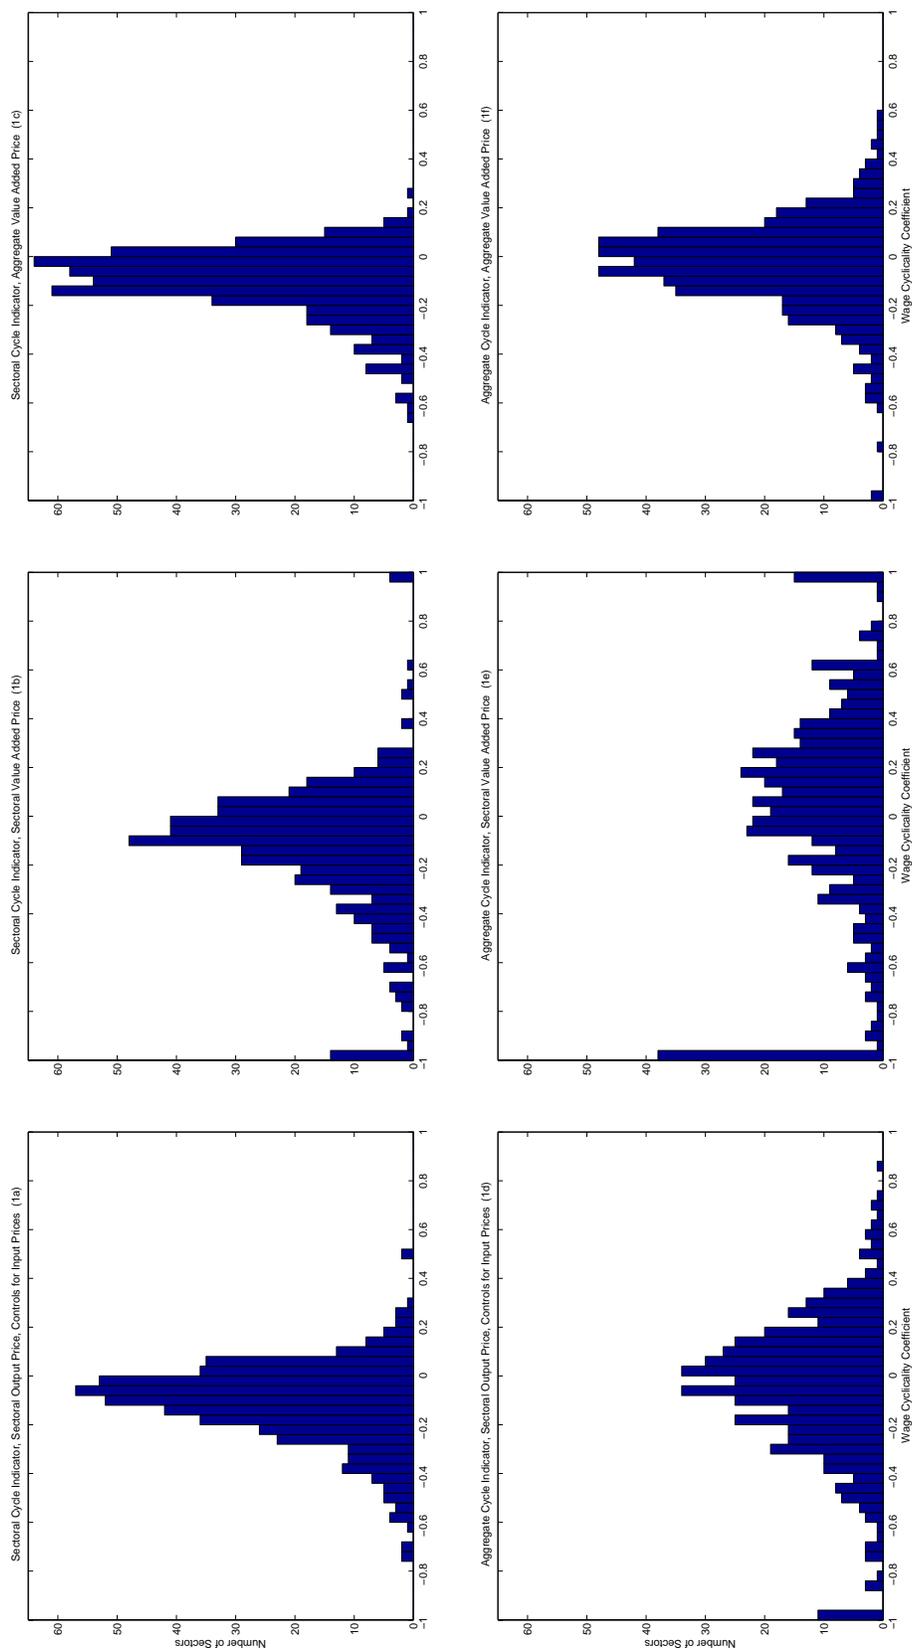


Table 1: Summary Statistics for 458 NBER Productivity Database Regressions
 (1a) through (1f) correspond to the panels in Figure 1; (1g) is for aggregate manufacturing

“Correct” (1a)		(1b)		(1c)	
Countercyclical:	352	Countercyclical:	321	Countercyclical:	355
at 10% level:	157	at 10% level:	113	at 10% level:	169
at 5% level:	127	at 5% level:	84	at 5% level:	132
at 1% level:	92	at 1% level:	56	at 1% level:	98
Procyclical:	106	Procyclical:	137	Procyclical:	103
at 10% level:	12	at 10% level:	14	at 10% level:	13
at 5% level:	9	at 5% level:	10	at 5% level:	8
at 1% level:	4	at 1% level:	5	at 1% level:	2
Mean:	−.114	Mean:	−.144	Mean:	−.103
Weighted Mean:	−.100	Weighted Mean:	−.122	Weighted Mean:	−.084
Median:	−.091	Median:	−.089	Median:	−.082
Mean Abs. t-stat:	1.528	Mean Abs. t-stat:	1.306	Mean Abs. t-stat:	1.563
(1d)		(1e)		“Proxied” (1f)	
Countercyclical:	246	Countercyclical:	200	Countercyclical:	250
at 10% level:	70	at 10% level:	50	at 10% level:	35
at 5% level:	49	at 5% level:	42	at 5% level:	15
at 1% level:	31	at 1% level:	31	at 1% level:	9
Procyclical:	212	Procyclical:	258	Procyclical:	208
at 10% level:	40	at 10% level:	35	at 10% level:	20
at 5% level:	15	at 5% level:	15	at 5% level:	15
at 1% level:	8	at 1% level:	9	at 1% level:	6
Mean:	−.062	Mean:	−.067	Mean:	−.035
Weighted Mean:	−.039	Weighted Mean:	−.055	Weighted Mean:	−.006
Median:	−.031	Median:	.063	Median:	−.021
Mean Abs. t-stat:	1.117	Mean Abs. t-stat:	1.263	Mean Abs. t-stat:	0.852
“Aggregate” (1g)					
Coefficient:		.036			
Std. Error:		.073			

data, anyway). Thus, this specification exactly matches the aggregate wage cyclical regressions run by previous authors, and surveyed by Abraham and Haltiwanger (1995). Perhaps not surprisingly, the results are even more procyclical than for the individual sectors regressed on the same variables (panel (1f)). The reason for this is that the more cyclical industries, such as durable goods, in general pay higher wages than do the less cyclical industries, such as nondurables, so that, all else equal, a recession will tend to be a time with a greater proportion of low-wage workers.¹² This industry composition effect induces additional procyclical into the average wage that is not present in the individual sectors separately. Note that the estimate for the wage cyclical coefficient in (1g) is very typical of other estimates in the literature using these methods—in particular, the coefficient α is slightly procyclical and statistically insignificant (Abraham and Haltiwanger (1995)).

The question naturally arises as to what factors in particular are the most important in leading to the biases observed in panel (1f)? Panels (1b) through (1e) in Figure 1 and Table 1 attempt to answer this question.

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 2. A comparison of the top panels of Figure 1 and Table 1 ((1a), (1b), and (1c)) to the bottom panels ((1d), (1e), and (1f)) reveals that the point estimates in the bottom panels vary much more widely; the mean absolute t-statistics in the bottom set of panels is much smaller; and the number of significant point estimates at every level of significance is much less than for the upper set of panels, which use the sectoral cycle indicator. Finally, the mean, weighted mean, and median coefficients for the bottom panels are all virtually zero or even slightly positive, and indicate a substantial upward or zero bias in the estimates, as compared to the correct specification (1a). All of these observations are consistent with the hypothesis that the aggregate cycle indicator (aggregate production worker hours) is a poor proxy for the state of economic conditions in individual four-digit

¹²This is an “industry composition effect”, as opposed to the “worker composition effect” emphasized by Solon, Barsky, and Parker (1994). Note that the industry composition effect induces a procyclical bias into the average wage statistic, while the worker composition effect (low-wage workers at each firm tend to be fired first) induces a countercyclical bias. These two biases are discussed in more detail in Swanson (1998).

industries, leading to classic errors-in-variables bias.^{13,14}

Second, it is apparent from Figure and Table 1 that the coefficient estimates using sectoral value-added deflators are quite poor. The point estimates are very disperse, much more so than for the correct specification (1a), yet the t-statistics for these regressions and the number of significant coefficients are small despite the large point estimates. These findings corroborate those of Jorgenson, Gollop, and Fraumeni (1987) and Basu and Fernald (1997), who find that the tenuous assumptions needed 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 slight positive bias in the coefficient estimates as well.¹⁵ These results are very much in line with the predictions of Section 2, which suggested that the use of value-added deflators would tend to create omitted-variables and errors-in-variables biases in the estimation of equation (2)—omitted-variables because of poor controls for intermediate input price variation, and errors-in-variables simply because the value-added deflators are very noisy at the four-digit level. Both of these factors appear to be at work.¹⁶

The final observation to be made from Figure and Table 1 is that, despite the poor performance of sectoral value-added deflators, the aggregate value-added deflator seems to perform quite well in panels (1c) and (1e), particularly in combination with sectoral

¹³ Although comovement between one-digit industries and the aggregate economy is fairly high (Murphy, Shleifer, and Vishny (1989)), comovement at the four-digit level is substantially lower, with an average correlation coefficient of about 0.4 (which would correspond to an R -squared of 0.16). Although all sectors tend to move downward during a recession, idiosyncratic factors make up the bulk of each sector's variance.

¹⁴ To the extent that estimates using the aggregate cycle indicator are slightly procyclical, rather than slightly countercyclical or zero, there is some indication that real wages may be correlated more positively with aggregate shocks than with idiosyncratic sectoral shocks. This would be the case if aggregate shocks have tended to shift firms' labor demand curves to a greater extent than have idiosyncratic sectoral shocks. While the regression results are consistent with this hypothesis, the evidence is not conclusive. Although a tendency toward procyclical coefficient estimates is present in Figures (1d) and (1e), it is not evident in Figure 2, and even in Figure 1 the mean and weighted mean coefficients are slightly negative, due to the presence of fat tails at the low end of the distribution. However, the hypothesis does offer a possible explanation for why many point estimates are slightly procyclical, when the errors-in-variables prediction would be simply for an attenuation bias toward zero.

¹⁵ This is not apparent in the mean and weighted mean statistics in panels (1b) and (1e), due to the fat tails at the lower end of the distributions.

¹⁶ In an additional set of regressions (not shown), I find that roughly half of the discrepancy between panels (1a) and (1b), or (1d) and (1e), is due to poor controlling for intermediate input price variation, and roughly half to the greater volatility of the sectoral value-added deflators (which are about 50% more variable than gross output prices, in terms of standard deviation).

measures of industry conditions, as in (1c). The point estimates are tightly clustered, to an extent comparable to the correct specification (1a), and the mean absolute t-statistic for panel (1c) is virtually identical to, and even slightly greater than, that for the correct specification (1a). Apparently, very little has been lost by moving from sectoral input and output prices to an aggregate value added deflator, in stark contrast to the predictions of Section 2. There it was argued that, because an aggregate value added deflator is a poor proxy for sectoral gross output and intermediate input prices, its use would introduce additional noise into the left-hand side variable, and yield estimates that are even more imprecise than those derived from sectoral value-added deflators. Clearly, this is not the case here. There are a few possible explanations. First, the aggregate deflator is much less noisy than the sectoral value-added deflators, and may still control for the most important variations in intermediate input prices, such as changes in energy prices; this might allow the aggregate deflator to yield estimates that are generally more precise than those using sectoral value-added prices, particularly if the signal-to-noise ratio in the sectoral deflators is very low, while that in the aggregate deflator is very high. Second, if the sectoral gross output and intermediate input price indexes are themselves poor proxies for the “true” values of these underlying variables, then it is possible that aggregate value added might perform just as well in the four-digit level regressions of panel (1c), as compared to (1a). Further insights into this empirical finding would be welcome.

KLEM Data

The results for the 34 non-governmental sectors of Jorgenson’s KLEM dataset are very similar to those above. Figure 2 presents histograms for six different variations of regression equation (2) that are analogous to those performed in Figure 1. Table 2 contains corresponding summary statistics for these regressions. In addition, Table 3 presents a list of the point estimates and standard errors for each of the 34 sectors using the most correct specification, (2a).¹⁷

¹⁷For later time periods, such as 1970–91, the number of significant coefficients is slightly less, and the number of positive point estimates is slightly greater, than those given in the text, but the main observations below all remain the same, and are very similar to those of Figure 1, above.

Figure 2: Regression Coefficients for 34 KLEM Industries

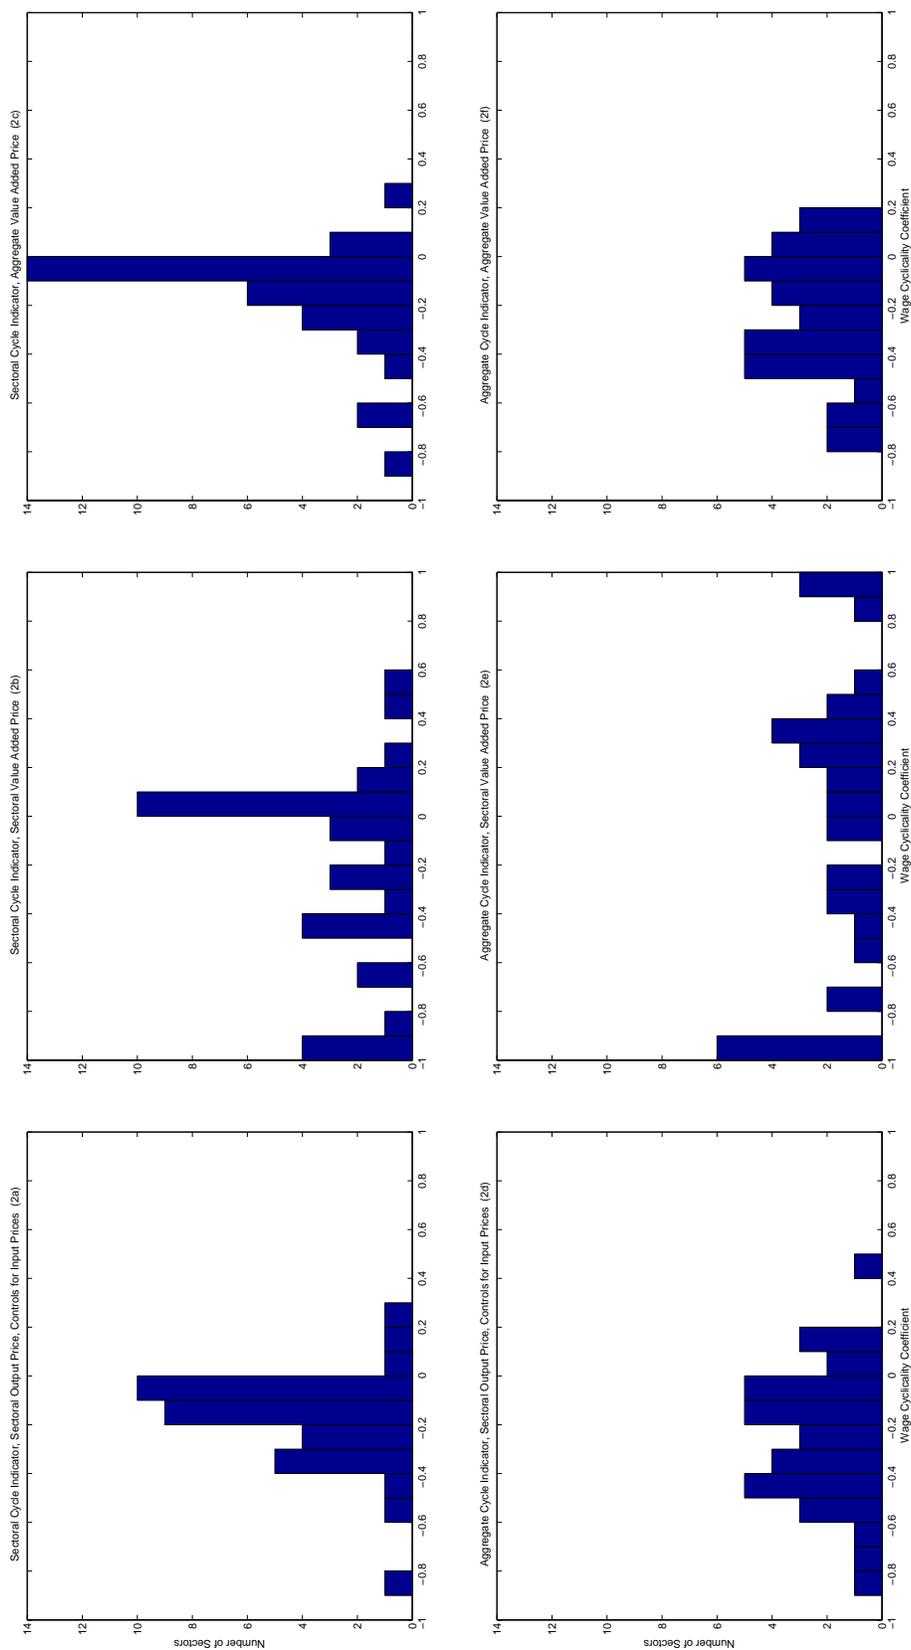


Table 2: Summary Statistics for 34 KLEM Data Regressions
 (2a) through (2f) correspond to the panels in Figure 2; (2g) is for the aggregate economy

“Correct” (2a)		(2b)		(2c)	
Countercyclical:	31	Countercyclical:	19	Countercyclical:	30
at 10% level:	21	at 10% level:	13	at 10% level:	17
at 5% level:	17	at 5% level:	9	at 5% level:	14
at 1% level:	13	at 1% level:	9	at 1% level:	11
Procyclical:	3	Procyclical:	15	Procyclical:	4
at 10% level:	2	at 10% level:	2	at 10% level:	0
at 5% level:	0	at 5% level:	1	at 5% level:	0
at 1% level:	0	at 1% level:	0	at 1% level:	0
Mean:	−.176	Mean:	−.232	Mean:	−.161
Weighted Mean:	−.118	Weighted Mean:	−.092	Weighted Mean:	−.131
Median:	−.147	Median:	−.055	Median:	−.097
Mean Abs. t-stat:	2.145	Mean Abs. t-stat:	1.809	Mean Abs. t-stat:	2.120
(2d)		(2e)		“Proxied” (2f)	
Countercyclical:	28	Countercyclical:	16	Countercyclical:	27
at 10% level:	15	at 10% level:	7	at 10% level:	18
at 5% level:	14	at 5% level:	5	at 5% level:	15
at 1% level:	13	at 1% level:	5	at 1% level:	12
Procyclical:	6	Procyclical:	18	Procyclical:	7
at 10% level:	0	at 10% level:	7	at 10% level:	0
at 5% level:	0	at 5% level:	3	at 5% level:	0
at 1% level:	0	at 1% level:	2	at 1% level:	0
Mean:	−.250	Mean:	−.349	Mean:	−.242
Weighted Mean:	−.172	Weighted Mean:	−.108	Weighted Mean:	−.183
Median:	−.241	Median:	.033	Median:	−.238
Mean Abs. t-stat:	1.723	Mean Abs. t-stat:	1.609	Mean Abs. t-stat:	1.873
“Aggregate” (2g)					
		Coefficient:	−.195		
		Std. Error:	.107		

The same general patterns are evident in Figure 2 that were apparent in Figure 1. First, the point estimates for the “correct” specification (2a) are almost uniformly countercyclical rather than procyclical. The mean, weighted mean, and median coefficients are in the range of -0.12 to -0.17 —almost double the magnitude of the estimates in Table 1, but in general agreement with them once one recognizes that the KLEM labor data are indexes of total labor input, as opposed to production worker hours, which are roughly twice as variable (in terms of standard deviation). The fact that the KLEM coefficients are somewhat less than twice as large may reflect the fact that these data control for changes in labor force composition, and thus largely eliminate the countercyclical worker composition bias that we might expect to be present in the NBER Productivity data.

The estimates in panel (2f) correspond more closely to regressions that have been run in previous studies of wage cyclicality, using an aggregate value added deflator, aggregate cycle indicator, and omitting controls for intermediate input prices. As was the case in Figure 1, these estimates are substantially more inaccurate than those for the correct specification, (2a). The point estimates are much more widely distributed; the mean absolute t-statistic is substantially smaller despite the larger point estimates; and the number of statistically significant coefficients is smaller at every level of significance. It appears at first, however, that in contrast to Figure (1f) and the predictions of Section 2, Figure (2f) does not display any upward or attenuation bias in the point estimates.¹⁸ This can be explained as follows: sectoral labor input is more variable than aggregate labor input (in fact, about twice as variable in this data, in terms of standard deviation); thus, regressions of wage changes on the less volatile aggregate labor index will tend to yield coefficients that are twice as large, all else equal. The fact that we do not observe coefficients of this magnitude is evidence of attenuation bias, as was predicted—aggregate labor input does not appear to be a perfect proxy for its sectoral counterparts.¹⁹ The results are thus very much in line with the predictions of Section 2.

¹⁸In fact, none of the panels using the aggregate cycle indicator ((2d), (2e), and (2f)) appears to be biased toward zero, in contrast to Figure 1. For later time periods (such as 1970–94), however, a positive shift in the distributions is apparent in moving from the top set of panels to the bottom panels.

¹⁹This also explains the larger values for the mean, median, and weighted mean point estimates in panel (2f) of Table 2. Note that this analysis applies to the coefficients in Figure 1 as well, so that the attenuation bias noted there is even stronger than might appear at first glance.

Table 3: KLEM Coefficients by Industry

	Industry	Coefficient	Std. Error	t-stat.
1	Agriculture	-0.417	(0.233)	-1.794
2	Metal mining	-0.158	(0.059)	-2.704
3	Coal mining	-0.213	(0.121)	-1.752
4	Oil and gas extraction	-0.334	(0.109)	-3.055
5	Non-metallic mining	-0.323	(0.068)	-4.791
6	Construction	-0.092	(0.058)	-1.583
7	Food & kindred products	-0.232	(0.128)	-1.818
8	Tobacco products	-0.397	(0.123)	-3.240
9	Textile mill products	-0.144	(0.112)	-1.280
10	Apparel	-0.081	(0.087)	-0.931
11	Lumber & wood products	-0.135	(0.063)	-2.166
12	Furniture & fixtures	-0.133	(0.062)	-2.160
13	Paper & allied products	-0.256	(0.077)	-3.339
14	Printing, publishing & allied	-0.377	(0.154)	-2.445
15	Chemicals	-0.260	(0.089)	-2.909
16	Petroleum & coal products	-0.002	(0.007)	-0.268
17	Rubber & plastic products	-0.080	(0.076)	-1.062
18	Leather products	-0.018	(0.090)	-0.195
19	Stone, clay, glass products	-0.055	(0.044)	-1.241
20	Primary metal	0.051	(0.069)	0.737
21	Fabricated metal products	-0.042	(0.072)	-0.581
22	Machinery, non-electrical	-0.022	(0.039)	-0.566
23	Machinery, electrical	-0.173	(0.048)	-3.627
24	Motor vehicles	0.106	(0.056)	1.903
25	Transport. equip. & ordnance	-0.089	(0.028)	-3.162
26	Instruments & electronics	-0.107	(0.054)	-1.997
27	Misc. manufacturing	-0.181	(0.080)	-2.258
28	Transportation	-0.157	(0.067)	-2.332
29	Communications	-0.304	(0.101)	-3.010
30	Utilities, electric	-0.594	(0.118)	-5.032
31	Utilities, gas	-0.151	(0.056)	-2.688
32	Wholesale & retail trade	-0.093	(0.112)	-0.831
33	Finance, insurance & real estate	-0.809	(0.216)	-3.743
34	Services	0.281	(0.162)	1.734

Panel (2g) of Table 2 presents estimates of equation (2) using nothing but aggregate data (wages, hours, and value-added prices) for the whole economy, constructed from the sectoral KLEM data as described in Section 3 above. Like the results in panel (1g), the coefficient is statistically insignificant; unlike the earlier result, however, the coefficient in (2g) is no more procyclical than those in (2f). The reason is the composition adjustment in the KLEM data—the aggregate wage statistic in (2g) is not aggregate wages divided by aggregate hours, but rather a composition-adjusted (Törnqvist) index of wages for the aggregate economy. The coefficient of this average is thus roughly equal to the average of the coefficients in panel (2f), which was not the case in panel (1g). Note that this provides an independent verification of the effectiveness of the KLEM composition adjustments.

Panels (2b) through (2e) of Figure and Table 2 tell a similar story to that of Figure and Table 1. The use of an aggregate cycle indicator as a proxy for sectoral industry conditions appears to induce a substantial amount of noise and attenuation bias into the estimates. Each of the bottom panels ((2d), (2e), and (2f)) exhibits more disperse point estimates and lower mean absolute t-statistics than its corresponding top panel ((2a), (2b), and (2c)), though the number of significant coefficients is sometimes slightly higher (though never higher than in (2a)). This finding agrees with the predictions of Section 2 and those in the NBER Productivity data.

The KLEM data also confirm the very poor results achieved using sectoral value-added deflators. In panels (2b) and (2e), the point estimates are extremely disperse, the mean absolute t-statistics are substantially smaller despite the larger point estimates, and the number of significant coefficients is dramatically lower for every test size. Even in two-digit level data, it thus appears that imposing value added relationships on the data leads to very poor estimation results.

Finally, just as in Figure 1 and Table 1, the estimates using an aggregate value-added deflator perform surprisingly well (panels (2c) and (2e)). The point estimates are tightly distributed, the mean absolute t-statistics are large (particularly in (2c)), and the number of significant coefficients is only slightly less than for the correct specification, (2a). Thus, the somewhat surprising finding in the previous section (that the aggregate value-added deflator is a very good proxy for sectoral input and output prices) appears to be robust.

5 Discussion

There are a number of conclusions to be drawn from the above analysis. First, real wages in sectoral data, deflated by sectoral product prices, controlling for both changes in intermediate input prices and changes in workforce composition, appear to be countercyclical—*i.e.*, inversely related to an indicator of the cyclical state of each sector. The noise surrounding the estimates tends to increase, and the magnitude of the estimated coefficients tends to decrease, as one moves toward more aggregate specifications of the regression equation. These findings are consistent with an omitted variables and errors-in-variables explanation.

Second, deflating wages by sectoral value-added deflators appears to yield extremely poor coefficient estimates. Standard errors using these data are very large, and the point estimates appear to be positively biased, as well as very disperse. The poor performance of the sectoral value-added deflators may be because the assumptions required for existence of a value-added production function are often not met in sectoral data, as found by Jorgenson, Gollop, and Fraumeni (1987). Sectoral value-added deflators also appear to be extremely noisy, substantially more so than sectoral gross output and input prices. I conclude, like Basu and Fernald (1997), that any results using sectoral value-added data should be regarded with a great deal of suspicion. Future studies should avoid using these data, in favor of gross output and input prices, if at all possible.

Finally, and surprisingly, the poor results using sectoral value-added data do not appear to generalize to results using an aggregate GDP deflator. Use of the latter to deflate sectoral data on wages yields estimates that appear to be virtually as precise as those using fully disaggregate gross output and intermediate input prices. There does not appear to be an obvious explanation for this finding.

These results can be checked against those of other studies in the literature. The specifications in this paper that use aggregated NBER Productivity data (panel (1g) of Figure and Table 1) are directly comparable to other studies of wage cyclicality using data on aggregate manufacturing, such as those surveyed in Abraham and Haltiwanger (1995). The coefficient estimates reported in those studies are very much in line with the results above, for the aggregated data—*i.e.*, a slightly procyclical point estimate, but one that is statistically insignificant.

The results using sectoral data are also very much in line with an earlier study by Waud (1968), who uses quarterly data on sectoral production-worker hours and wages, deflated by an aggregate (GNP) deflator, to investigate wage-employment correlations in 17 two-digit manufacturing industries.²⁰ His estimates thus correspond directly to those of panels (1c) and (2c), above. As in those regressions, he finds very strong evidence of an inverse (*i.e.*, countercyclical) relationship between wages and hours of employment: “the sign of the estimated coefficient of $\ln W_1$ is negative in all but one case...” (p. 415). Also, as I do above, he notes the importance of using a sectoral rather than an aggregate indicator of industrial conditions: “The number of cycles n , of course, varies from industry to industry...” (p. 414). He does not, however, have data on gross output and intermediate input prices with which to run the more correct specification (1a) or (2a), nor does he compare his results to those that are obtained using the more traditional specification (1f) or (2f). All in all, however, his results agree very closely with those of the present paper.^{21,22}

²⁰I thank John Pencavel for this reference.

²¹Another closely related study using sectoral wage data is that of Estevão and Wilson (1998). These authors look at wage cyclicality from the point of view of both aggregate and sectoral data, and interpret the use of aggregate data as helping to identify labor supply, and the use of detailed sectoral data as helping to identify labor demand. Using the NBER Productivity Database, they run a pooled regression, both instrumented and uninstrumented, and find strong evidence of a negative relationship between real wages (deflated by sectoral prices, with controls for intermediate input prices) and sectoral employment. This result thus broadly corroborates the findings of Section 4, above. The results are not directly comparable, however, because they are pooled into a single regression, and thus produce a coefficient that would be most comparable to the “weighted mean” statistic given in Tables 1 and 2 of this paper (Estevão and Wilson do allow the wage-employment coefficient to vary across sectors by interacting it with the industry concentration ratio, capital-labor ratio, and an indicator for whether the given sector produces durable goods, but do not allow the wage-employment correlation to vary more freely than this.) This approach has two shortcomings as far as the present analysis is concerned. First, pooling of the data prevents a detailed analysis of the sensitivity of regression specification (2) to the various substitutions of aggregate variables for their sectoral counterparts, and the omitted- and errors-in-variables biases that this induces. This has been a major focus of the present paper. Second, pooling the sectoral regressions in this manner leads to contemporaneous correlation of the residuals (due to comovement of sectors during recessions), and will bias reported standard errors downward; I have been able to avoid that problem with the more disaggregated approach taken here. Finally, Estevão and Wilson do not attempt to address the possible effects of cyclical changes in workforce composition on their results. Verifying the robustness of these findings to possible worker composition bias has been a major concern of the present paper.

²²An apparently contradictory study is that of Solon and Barsky (1989). They run a series of regressions using one-digit-level data on wages and prices, and fail to find any significant relationship in either direction with respect to the aggregate unemployment rate (see esp. their Table 5). Their findings hold generally whether they deflate by an aggregate GNP deflator, or by one-digit-level producer price indexes (with no controls for intermediate input prices). Thus, their specification corresponds almost exactly to that of panels (1e) and (1f), and (2e) and (2f), above. Solon and Barsky’s results are thus not particularly surprising when considered in light of the results of Section 4, above. There I showed that the use of an

The results using composition-adjusted data, however, are somewhat different from those of earlier studies using panel data on workers, such as Solon, Barsky, and Parker (1994). One would expect that the composition-adjusted KLEM series constructed by Jorgenson *et. al.*, plugged into aggregate specification (2g), would yield very similar results to the composition-adjusted series constructed by Solon, Barsky, and Parker (hereafter SBP), plugged into their regression framework, which matches (2g) directly. In fact, the KLEM data yields a negative coefficient estimate, while the SBP series results in a strongly positive coefficient. A number of details in the construction of the two series helps to explain much, if not all, of this discrepancy. First, the SBP series, which is derived from the PSID, covers only the years 1967 to 1987. This period is one of generally more procyclical wages than others in postwar U.S. history (see Abraham and Haltiwanger (1995)). Estimating the regression specification in panel (2g) using the KLEM data restricted to this time period yields a coefficient estimate that is noticeably more procyclical (.040 (with standard error .203) as compared to -0.195). Second, the SBP index is constructed entirely from data on male heads of household, while the Jorgenson KLEM data are indices of all labor input, including women and teenagers. This affects the relative cyclical behavior of the two series because the wages of men are much more procyclical than those of women; Solon, Barsky, and Parker (1994) note: “Whatever the reason for the gender difference in wage cyclicity, it poses a new difficulty for analyzing the discrepancy in wage cyclicity as measured in aggregate statistics versus longitudinal microdata. We now see that the measures may differ not only because of composition bias, but also because the aggregate measure combines the disparate wage cyclicality of men and women” (p. 15).²³ Third, the KLEM indexes of labor input give substantially more weight to high-income workers

aggregate cycle indicator, as opposed to an indicator of sectoral industry conditions, led to a substantial deterioration in the accuracy of the coefficient estimates, and even a substantial upward bias in the (non-composition-adjusted) NBER Productivity data (although there is greater comovement between one-digit industries and the aggregate economy, so that an aggregate cycle indicator ought to be a better proxy in this case, this may be more than offset by the highly aggregate nature of the one-digit industry wage data, which will contain more serious (and procyclical) industry composition effects (see footnote 12)). Thus, it seems reasonable to assume that their results would have been more significant (and countercyclical) had they used sectoral employment, or sectoral output, as their indicator of industry conditions.

²³Solon, Barsky, and Parker’s composition-adjusted index for women in fact exhibits *less* procyclicality than their aggregate wage statistic, in contrast to the composition-adjusted index for men, which exhibits substantially greater procyclicality than the aggregate measure. The latter, however, is the main focus of the SBP paper.

than does the SBP index. This is because the latter is constructed as a simple average of workers' log wage changes, with each worker receiving equal weight regardless of the number of hours worked or wages earned; in contrast, the KLEM (Törnqvist) methodology weights workers' log wage changes by their share in the total wage bill (*i.e.*, hours times wages, or nominal income). Neither method is correct or incorrect—they are simply different approaches, with each being more appropriate for its own class of applications.²⁴ The two approaches do, however, have an important effect on the end result, because the wages of low-income men are much more procyclical than those of high-income men (Swanson (1998)). Finally, there is a major difference between the two composition-adjusted indexes in their treatment of workers who change jobs. A worker who changes jobs from a low-paying sector (*e.g.*, Services) to a high-paying sector (*e.g.*, Construction or Manufacturing) will contribute a large positive wage change to the SBP index; in contrast, the same worker changing between the same jobs will have *no* effect on the KLEM index of wages. This is because the KLEM methodology equates higher wages in the high-paying sector with a greater efficiency-equivalent unit of labor input; the movement of labor from the low-paying sector to the high-paying sector thus shows up in the KLEM data as an increase in the quality-adjusted index of labor input, rather than as a change in wages. Because a large fraction of the procyclicality of wages is attributable to the effects of workers changing jobs (Bils (1985)),²⁵ the SBP index will be more procyclical than the KLEM index for this reason as well. Thus, although both the KLEM data and SBP series are composition-adjusted, their methods of construction, and thus their implicit underlying philosophy of adjustment for changes in workforce composition, are remarkably different. The SBP methodology, and time period, result in a substantially greater degree of real wage procyclicality. However, I have argued above (in footnote 24) that the KLEM index is more appropriate for testing the predictions of the Classical (or Keynesian) versus

²⁴The Jorgenson approach, however, is generally more appropriate when the question of interest is related to production theory (for example, labor demand). Assuming that workers are paid wages in proportion to their marginal products, Jorgenson's approach is the correct one for creating efficiency-equivalent indexes of labor input and wages, weighting each type of labor by its marginal product as well as its number of hours worked. This suggests that the KLEM methodology is more appropriate for testing the implications of the Classical/Keynesian versus Technology-driven models outlined in the Introduction.

²⁵Solon, Barsky, and Parker (1994) find a less important, though still substantial, effect of workers changing jobs on their results.

Technology-driven models of business cycles outlined in the Introduction.

Pencavel and Craig (1994) provide us with a study that is highly complementary to the above results. Their paper is unique in that it covers only a single sector (plywood manufacturing in the Pacific Northwest), but does so at the firm level for a very homogeneous set of workers and firms. The authors examine the relationship between employment, hours, wages, input prices, and output prices for their sample of 35 firms for each of the years 1968 to 1986. Although they frame their regressions as a test of profit maximization under the maintained assumption that their firms' input demand and output supply curves are stable over the sample period, one can flip the maintained and tested hypotheses, and interpret their results as a test of whether firms' input demand and output supply curves are stable over time, under the assumption that they maximize profits. This is exactly the type of test outlined in Section 2 and performed in Section 4 of this paper. Framed in this way, Pencavel and Craig do not reject the hypothesis of stable labor demand curves over their sample period. They document a very strong inverse (*i.e.*, countercyclical) relationship between real wages (deflated by plywood output prices) and production worker hours in their sample, controlling for changes in the price of inputs. Their findings are thus extremely supportive of those in the present paper. Moreover, they provide independent corroboration of the results for the Lumber and Wood Products industry in Table 3 using the more aggregate data of Section 4.

Finally, although Bils (1987) focuses on the cyclicity of markups rather than real wages, his reported results do provide estimates of real wage cyclicity by two-digit manufacturing industry (albeit without standard errors, as these are reported only for wages and prices separately, and for overall markup cyclicity). Somewhat surprisingly, Bils' point estimates of the real wage-employment relationship are almost uniformly procyclical. There are a number of possible reasons for this. First, Bils uses sectoral value-added deflators, which were shown above to yield very noisy, and even positively biased, results (panels (1b), (1e), (2b), and (2e) of Section 4). Second, he uses a levels specification detrended by a cubic polynomial. In sectoral data, however, fluctuations in wages, prices, and employment are often very large or nonstationary, and cubic detrending does not leave behind what one would think of as business-cycle-frequency fluctuations. Waud (1968)

makes a similar observation: “The usual practice of inserting time and time squared in the regression is certainly inappropriate in view of the above observations [that trends in wages and labor productivity can be complex],” (p. 413). A low-parameter HP filter, or annual first-differences specification, tends to do a significantly better job. Third, Bils’ labor detrending method is a relatively high-frequency filter (he uses deviation from a five-year centered moving average), which does not match the relatively low-frequency filter that is being used for real wages (cubic detrending). This mismatch of filters may lead to noisy or spurious results. Finally, as a check on the effects of Bils’ methods, I replicate his regressions using very comparable data (Hall’s (1988) data from the BEA), and find, like him, that a very large majority of point estimates are positive; however, they are very imprecisely estimated, with only three of the twenty being statistically significant at even the 10% level (and each of these three in fact being negative, or countercyclical). Using more standard methods of detrending (such as HP-filtering or first-differencing), and detrending both real wages and employment by the same filter, leads to many more countercyclical (and significantly countercyclical) point estimates. Details of these regressions and robustness checks are provided in Swanson (1998). I thus do not find a significant contradiction with the results of the present paper.

6 Conclusions

On the basis of these results, I find little reason to reject the Classical (or Keynesian) model of cyclical wage and employment determination. This is in contrast to many previous studies of wage cyclicality. I have tried to improve on earlier studies by 1) using more detailed, sectoral data on wages, prices, and industry conditions wherever possible, 2) investigating the effects of using aggregate proxies for these sectoral variables, and 3) controlling for changes in workforce composition.

The magnitudes of my estimates for real wage cyclicality, deflating by sectoral product prices and controlling for changes in intermediate input prices, are centered roughly around -0.1 , so that a 10% increase in production worker hours appears to be associated, on average, with a 1% decrease in production worker wages. It should be noted that these

are not estimates of the slope of a labor demand curve, as I have not attempted to use instrumental variables that would identify labor demand. Rather, the results should be regarded as an empirical test of the predictions of the Classical, or Keynesian, theory of cyclical real wage behavior, as compared to the predictions of standard technology-driven models. I find relatively little support for the predictions of the latter in the data.²⁶

Finally, it should be noted that an acceptance of the Classical model of counter-cyclical real wage behavior does not necessarily require the view that workers suffer from money illusion or are bound by rigid nominal wage contracts (though these may be important), nor does it require rejecting the preponderance of empirical evidence that labor productivity at the plant or detailed industry level is procyclical (see, for example, Foster, Haltiwanger, and Krizan (1998) and Bartelsman and Dhrymes (1994)). A multi-sector Classical model can be consistent with all of these views. For example, a positive fiscal spending shock that impacts one sector of the economy more than the 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 labor productivity, despite the fall in real wages (properly deflated). A general equilibrium version of this model, considering the effects of other shocks as well, is worked out in detail in Swanson (1999) and Chapters 4 and 5 of Swanson (1998). Further empirical and theoretical work in this direction thus seems promising.

²⁶ One can, and should, augment these models to include fiscal shocks, such as changes in government purchases and taxes, as in Christiano and Eichenbaum (1992), Braun (1994), and McGrattan (1994). However, the predictions of these models are still for real wages to be procyclical. Swanson (1999), and Chapters 4 and 5 of Swanson (1998), present a multisector dynamic general equilibrium model in which technology is not the main driving force of the business cycle; rather, the observed procyclicality of aggregate productivity (TFP) is primarily the result of changes in the composition of output, in response to, for example, fiscal shocks, cyclical changes in consumption patterns, or technological innovations in an investment goods-producing sector.

References

- ABRAHAM, KATHERINE AND JOHN HALTIWANGER. "Real Wages and the Business Cycle," *J. Econ. Lit.* 33 (1995), pp. 1215–65.
- BARRO, ROBERT AND ROBERT KING. "Time-Separable Preferences and Intertemporal-Substitution Models of Business Cycles," *Q. J. E.* 99 (1984), pp. 817–39.
- BARTELSMAN, ERIC AND PHOEBUS DHRYMES. "Productivity Dynamics: U.S. Manufacturing Plants, 1972–86," *Federal Reserve Board Finance and Economics Discussion Series* 94–1 (1994).
- BARTELSMAN, ERIC AND WAYNE GRAY. "The NBER Manufacturing Productivity Database," *NBER Technical Working Paper* 205 (1996).
- BASU, SUSANTO AND JOHN FERNALD. "Returns to Scale in U.S. Production: Estimates and Implications," *J. P. E.* 105 (1997), pp. 249–83.
- BILS, MARK. "Real Wages Over the Business Cycle: Evidence from Panel Data," *J. P. E.* 93 (1985), pp. 666–89.
- BILS, MARK. "The Cyclical Behavior of Marginal Cost and Price," *A. E. R.* 77 (1987), pp. 838–55.
- BRAUN, R. ANTON. "Tax Disturbances and Real Economic Activity in the Postwar United States," *J. Monetary Econ.* 33 (1994), pp. 441–62.
- CHRISTIANO, LAWRENCE AND MARTIN EICHENBAUM. "Current Real-Business-Cycle Theories and Aggregate Labor-Market Fluctuations," *A. E. R.* 82 (1992), pp. 430–50.
- ESTEVEÃO, MARCELLO, AND BETH ANNE WILSON. "Nominal Wage Rigidity and Real Wage Cyclicity," *Federal Reserve Board Finance and Economics Discussion Series* 98–21 (1998).
- FOSTER, LUCIA, JOHN HALTIWANGER, AND C.J. KRIZAN. "Aggregate Productivity Growth: Lessons from Microeconomic Evidence," *NBER Working Paper* 6803 (1998).
- HALL, ROBERT. "The Relation between Price and Marginal Cost in U.S. Industry," *J. P. E.* 96 (1988), pp. 921–47.
- JORGENSEN, DALE, FRANK GOLLOP, AND BARBARA FRAUMENI. *Productivity and U. S. Economic Growth* (Harvard Univ. Press: Cambridge, 1987).
- KEYNES, JOHN MAYNARD. *The General Theory of Employment, Interest, and Money* (London: MacMillan, 1936).
- MCGRATTAN, ELLEN. "A Progress Report on Business Cycle Models," *Fed. Res. Bank Minneapolis Quart. Rev.*, Fall (1994), pp. 2–16.
- MURPHY, KEVIN, ANDREI SHLEIFER, AND ROBERT VISHNY. "Building Blocks of Market Clearing Business Cycle Models," *NBER Macro Annual* (1989), pp. 247–87.
- PENCAVEL, JOHN AND BEN CRAIG. "The Empirical Performance of Orthodox Models of the Firm: Conventional Firms and Worker Cooperatives," *J. P. E.* 102 (1994), pp. 718–44.
- PRESCOTT, EDWARD. "Theory Ahead of Business Cycle Measurement," *Fed. Res. Bank Minneapolis Quart. Rev.*, Fall (1986), pp. 9–22.

- SOLON, GARY AND ROBERT BARSKY. "Real Wages Over the Business Cycle," *NBER Working Paper* 2888 (1989).
- SOLON, GARY, ROBERT BARSKY, AND JONATHAN PARKER. "Measuring the Cyclicity of Real Wages: How Important is Composition Bias?" *Q. J. E.* 109 (1994), pp. 3–25.
- STIGLER, GEORGE AND JAMES KINDAHL. *The Behavior of Industrial Prices* (New York: Columbia University Press, 1970).
- SWANSON, ERIC. *Individual and Sectoral Heterogeneity, Reallocation, and Aggregate Fluctuations* (Stanford, CA: Stanford University Doctoral Dissertation, 1998).
- SWANSON, ERIC. "Models of Sectoral Reallocation," *Federal Reserve Board Finance and Economics Discussion Series* 99–03 (1999).
- WAUD, ROGER N. "Man-Hour Behavior in U. S. Manufacturing: A Neoclassical Interpretation," *J. P. E.* 76 (1968), pp. 407–27.