

NBER WORKING PAPER SERIES

MEASURING THE CYCLICALITY
OF REAL WAGES:
HOW IMPORTANT IS COMPOSITION BIAS?

Gary Solon

Robert Barsky

Jonathan A. Parker

Working Paper No. 4202

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 1992

This research has been supported by National Science Foundation Grant No. SES-8709217 and by an Alfred P. Sloan Research Fellowship awarded to Solon. The authors are grateful for the excellent research assistance of Deborah Laren and Robert Wood and for comments from seminar participants at Michigan, Harvard, Kent, LSE, McMaster, NBER, Princeton, Rochester, Toronto, and Warwick. This paper is part of NBER's research programs in Monetary Economics and Economic Fluctuations. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

NBER Working Paper #4202
October 1992

MEASURING THE CYCLICALITY
OF REAL WAGES:
HOW IMPORTANT IS COMPOSITION BIAS?

ABSTRACT

In the period since the 1960's, as in other periods, aggregate time series on real wages have displayed only modest cyclicality. Macroeconomists therefore have described weak cyclicality of real wages as a salient feature of the business cycle. Contrary to this conventional wisdom, our analysis of longitudinal microdata indicates that real wages have been substantially procyclical since the 1960's. We also find that the substantial procyclicality of men's real wages pertains even to workers that stay with the same employer and that women's real wages are less procyclical than men's.

Numerous longitudinal studies besides ours have documented the substantial procyclicality of real wages, but none has adequately explained the discrepancy with the aggregate time series evidence. In accordance with a conjecture by Stockman (1983), we show that the true procyclicality of real wages is obscured in aggregate time series because of a composition bias: the aggregate statistics are constructed in a way that gives more weight to low-skill workers during expansions than during recessions. We conclude that, because real wages actually are much more procyclical than they appear in aggregate statistics, theories designed to explain the supposed weakness of real wage cyclicality may be unnecessary, and theories that predict substantially procyclical real wages become more credible.

Gary Solon
Department of Economics
University of Michigan
Ann Arbor, MI 48109

Jonathan A. Parker
Department of Economics
MIT
Cambridge, MA 02139

Robert Barsky
Department of Economics
University of Michigan
Ann Arbor, MI 48109
and NBER

Measuring the Cyclicalities of Real Wages:
How Important Is Composition Bias?

I. Introduction

Because aggregate time series on real wages display little cyclicalities, macroeconomists commonly have described weak cyclicalities of real wages as a salient feature of the business cycle. According to Lucas (1977, p. 17), for example, "Observed real wages are not constant over the cycle, but neither do they exhibit consistent pro- or countercyclical tendencies." Mankiw (1989, p.86) likewise has stated that, "over the typical business cycle, employment varies substantially while the determinants of labor supply -- the real wage and the real interest rate -- vary only slightly." In Blanchard and Fischer's (1989, p. 19) words, "The correlation between changes in real wages and changes in output or employment is usually slightly positive but often statistically insignificant." Very similar remarks have appeared recently in Abel and Bernanke (1992, pp. 338-40, 411-12, 448, and 454), Christiano and Eichenbaum (1992, pp. 430-31), Greenwald and Stiglitz (1988, pp. 223, 225, and 241), Hall and Taylor (1991, pp. 444-45), and Prescott (1986, p. 28) among others.

The widely shared view that real wages are at most weakly procyclical has profoundly affected the development of macroeconomic theory. Of course, the absence of countercyclical real wages has long been cited as a reason for dismissing theories that attribute cyclical labor market fluctuations to shifts in effective labor supply along a stable labor demand curve. Such theories include the model of nominal wage rigidity in Chapter 2 of Keynes' General Theory (1936) and the price misperceptions models of Friedman (1968) and Phelps (1970). More recently, macroeconomists generally have ascribed cyclical labor market fluctuations to shifts in labor demand along a highly elastic effective labor supply curve. As neatly summarized by Hall (1988, pp. 261-62), the appeal of this characterization is that it "accounts for...significant output and employment fluctuations and small real wage fluctuations."

Despite the agreement on this general paradigm, modern macroeconomists disagree quite vehemently about its details. Some macroeconomists attribute the shifts in labor demand to real productivity shocks, while others reserve a major role for aggregate demand disturbances.¹ Macroeconomists differ also with regard to the reasons for the high elasticity of effective labor supply. Some have followed Lucas and Rapping (1969) in explaining elastic labor supply as a reflection of intertemporal substitution behavior. Others have argued that the magnitude of the intertemporal substitution elasticity needed to reconcile observed employment fluctuations with small real wage fluctuations is implausibly large.² These macroeconomists have formulated alternative theories for why effective labor supply is so elastic that cyclical shifts in labor demand generate only small real wage movements. These theories -- including efficiency wage models, implicit contract models in which employers provide real wage insurance to workers, and insider-outsider models -- are surveyed in Chapter 9 of Blanchard and Fischer (1989). Despite differences in preferred explanations, however, macroeconomists generally have shared the premise that real wages vary only slightly over the business cycle.

The main conclusion of our paper is that the apparent weakness of real wage cyclicality in the United States has been substantially exaggerated by a statistical illusion. According to evidence from longitudinal surveys that have tracked individual workers since the 1960's, real wages have been highly procyclical in that period even though aggregate real wage data for the same period have not been nearly

¹As explained by Rotemberg and Woodford (1991), the latter approach generally involves imperfect competition with countercyclical price markups.

²In principle, the observed employment fluctuations could be labor supply responses to variation in real interest rates as well as real wages. But, as suggested in the above quotation from Mankiw, real interest rates also have been judged insufficiently procyclical to generate the observed employment movements via intertemporal substitution. Also, as explained by Barro and King (1984), a labor supply explanation centered on real interest rate variation leads (under time-separable preferences) to the counterfactual prediction that consumption and employment move in opposite directions over the business cycle.

so procyclical. Although our finding of substantial procyclicality in wage data from the Panel Study of Income Dynamics (PSID) was foreshadowed by numerous previous studies based on the PSID or the National Longitudinal Surveys of labor market experience, the discrepancy between the longitudinal evidence and the evidence from aggregate wage statistics has not been well-understood. We find that the two types of evidence differ within the same time period because the aggregate statistics are constructed in a way that gives more weight to low-skill workers during expansions than during recessions. This composition effect, first pointed out by Stockman (1983), biases the aggregate statistics in a countercyclical direction and obscures the true real wage procyclicality that workers typically experience. Unlike some previous researchers, we find Stockman's composition bias to be quantitatively important. When we purposefully distort the PSID data by imposing the same sort of weighting used in the aggregate statistics, we replicate the weak cyclicity displayed by the aggregate data. Because the composition bias in aggregate wage statistics seems likely to have been important in earlier periods (and perhaps other countries) as well, we conclude that near-noncyclicality of real wages should not be accepted as a salient feature of the business cycle. Accordingly, theories designed to explain the supposed weakness of real wage cyclicity may be unnecessary, and theories that predict substantially procyclical real wages become more credible.

Section II of this paper lays out our econometric framework, summarizes the aggregate time series evidence, and discusses the composition bias issue in detail. Section III presents our analysis of longitudinal data from the Panel Study of Income Dynamics and documents the empirical importance of composition bias. Section IV synthesizes our evidence with results from other longitudinal studies, and Section V discusses implications for macroeconomic theory.

II. Why Measuring Real Wage Cyclicity Is Trickier Than It Seems

A simple statistical model for characterizing the cyclicity in aggregate real wage data is

$$(1) \quad \ln W_t = \gamma_1 + \gamma_2 t + \gamma_3 t^2 + \gamma_4 (U_t - \delta_1 - \delta_2 t - \delta_3 t^2) + e_t$$

where W_t is some aggregate real wage measure in year t , U_t is the civilian unemployment rate (or some other indicator of the stage of the business cycle), and e_t is a random error term. A quadratic time trend is included in the wage equation, and the unemployment rate is entered as a deviation from its own quadratic trend, in order to focus on the cyclical components of wage and unemployment variation. With the unemployment rate as the cycle indicator, $\gamma_4 \gtrless 0$ as W_t is countercyclical, noncyclical, or procyclical.

Because e_t typically is highly serially correlated or even nonstationary, it is useful to first-difference equation (1) to obtain

$$(2) \quad \Delta \ln W_t = \beta_1 + \beta_2 t + \beta_3 \Delta U_t + v_t$$

where $v_t = \Delta e_t$, $\beta_1 = \gamma_2 - \gamma_3 + \gamma_4 (\delta_3 - \delta_2)$, $\beta_2 = 2(\gamma_3 - \gamma_4 \delta_3)$, and $\beta_3 = \gamma_4 \gtrless 0$ as the real wage variable is countercyclical, noncyclical, or procyclical. Equation (2) is precisely the same specification used by Bils (1985) in his aggregate time series analyses and, as will be seen later, it dovetails neatly with the specifications used by Bils, ourselves, and others for analyzing longitudinal microdata.³

Table I presents results from ordinary least squares (OLS) estimation of equation (2). Following Bils, we initially measure W_t with average hourly earnings of production or nonsupervisory workers in private nonagricultural employment, deflated by the implicit GNP deflator. The earnings data are

³It also may be worth noting that the t -ratio from ordinary least squares estimation of the relationship between $\Delta \ln W_t$ and ΔU_t is invariant to whether $\Delta \ln W_t$ is regressed on ΔU_t or vice versa.

generated by the Bureau of Labor Statistics (BLS) establishment survey. Details on data sources for these and other variables are provided in the Appendix.

The first column shows the results for 1947-48 to 1986-87 with the unemployment rate as the cycle indicator. The estimated coefficient of ΔU_t , $\hat{\beta}_3 = -.0030$, indicates mild, but statistically significant procyclicality of W_t . It implies that, when the unemployment rate increases by an additional percentage point, real wage growth declines by less than a third of a percentage point.⁴ The second column shows the results from using the natural logarithm of real GNP, instead of the unemployment rate, as the cycle indicator. Here we are regressing real wage growth on real GNP growth, and $\hat{\beta}_3$ estimates the elasticity of W_t with respect to real GNP. The elasticity estimate .153 again indicates mild procyclicality.

Next, for purposes of comparability with our later analyses of longitudinal microdata, we reestimate equation (2) for only the more recent subperiod 1967-68 to 1986-87. Doing so approximately doubles the estimated procyclicality, with the estimated coefficient of ΔU_t increasing from $-.0030$ to $-.0060$ and that of real GNP growth increasing from .153 to .293. For this later period, we are able to construct another cycle indicator, per capita hours of work, which is calculated as the product of the civilian employment/population ratio and average work hours of the employed. With this variable entered logarithmically, we obtain a .373 estimate of the elasticity of W_t with respect to aggregate hours of work.

⁴Bils' estimate based on the BLS time series is $\hat{\beta}_3 = -.0120$, four times larger than ours for 1947-48 to 1986-87. His estimate is based on only ten observations: 1966-67, 1967-68, 1968-69, 1969-70, 1970-71, 1971-73, 1973-75, 1975-76, 1976-78, and 1978-80. Some of his observations are two-year differences to maintain consistency with his microdata from the National Longitudinal Surveys of labor market experience. We have replicated his analysis with a more recent revision of the implicit GNP deflator series and obtained a slightly smaller $\hat{\beta}_3 = -.0104$. Filling in his missing years, i.e., using one-year differences for all years from 1966-67 to 1979-80, further reduces $\hat{\beta}_3$ to $-.0085$. Extending the sample period to 1986-87 produces $\hat{\beta}_3 = -.0062$. Then extending the sample period back to 1947-48 produces our estimate $-.0030$.

The tendency to estimate greater real wage procyclicality in more recent years has been noted previously by Coleman (1984) and Kniesner and Goldsmith (1987) among others. But even the estimates based on the more recent years are not so large. For example, suppose one assumes that cyclical labor market fluctuations arise from labor demand shifts along a stable short-run labor supply curve, which is positively sloped because of intertemporal substitution in labor supply. Then the regression of real wage growth on growth of work hours estimates the inverse elasticity of the short-run aggregate labor supply function. Our implied estimate of the labor supply elasticity itself is $1/.373 = 2.68$. If the short-run labor supply curve is not perfectly stable, however, this estimate is subject to simultaneity bias. As shown in Leamer (1981), if the supply equation's error term has positive variance and is uncorrelated with the demand error term, the inconsistency of our inverse estimator is in an upward direction. On the other hand, when the price and quantity variables are positively correlated (as they are here), OLS estimation of the direct regression of hours growth on real wage growth produces a downward-inconsistent estimator of the supply elasticity.⁵ Applying the direct estimator here generates an estimated supply elasticity of 1.20. Thus, if one assumes that cyclical hours fluctuations lie on a short-run notional labor supply function, the magnitudes of the cyclical hours and wage variation in aggregate data suggest a supply elasticity between 1.2 and 2.7. Many economists view such an elasticity as implausibly large. Consequently, numerous writers -- such as Abel and Bernanke (1992), Ashenfelter (1984), Fischer (1988), Greenwald and Stiglitz (1988), and Mankiw (1989) -- have concluded that the observed magnitude of real wage procyclicality is too small to generate the observed magnitude of cyclical hours variation via notional labor supply behavior.

Our finding of only modest cyclicity in measured aggregate real wages is not novel. It is the typical finding in a huge literature and is robust to variations in time unit, dynamic specification,

⁵Both inconsistency results hold a fortiori when the variables are subject to classical measurement error.

disaggregation by industry, treatment of overtime pay, and choice of deflator. (See Solon and Barsky [1989] for a series of robustness checks and references to the literature.) The point of the present paper, however, is to show that this typical finding is a spurious artifact of a composition bias in the aggregate wage statistics. The possibility of such a bias was first pointed out in an unpublished 1983 article by Stockman, but, as indicated by the citations in our first paragraph, the practical importance of this bias has not been taken seriously by the economics community. In the course of our empirical investigations, we have come to take it very seriously, and the rest of this paper tells why.

To understand Stockman's composition bias, one needs to understand that aggregate wage statistics like the BLS average hourly earnings measures are calculated as the ratio of the relevant sector's total wage bill B_t to its total work hours H_t . Now suppose the relevant worker population is divided into groups $j = 1, 2, \dots, J$ with B_j denoting the j^{th} group's wage bill, H_j its total work hours, $S_j = H_j/H_t$ its share of the population's work hours, and $W_j = B_j/H_j$ its average hourly earnings. Then the overall wage statistic W_t can be expressed as

$$\begin{aligned}
 (3) \quad W_t &= B_t/H_t \\
 &= \sum_{j=1}^J B_j/H_t \\
 &= \sum_{j=1}^J H_j W_j/H_t \\
 &= \sum_{j=1}^J S_j W_j .
 \end{aligned}$$

As the last version of equation (3) makes clear, the aggregate wage statistic is a weighted average of the group-specific wage statistics with the groups weighted by their hours shares.

The problem with this sort of wage statistic for measuring wage cyclicality is that the groups' hours shares vary with the business cycle. In particular, a long history of studies has shown that the work hours of low-wage groups tend to be more cyclically variable than those of high-wage groups. Kosters and Welch (1972), Okun (1973), Clark and Summers (1981), and Mitchell, Wallace, and Warner (1985) have documented the extreme cyclical hours sensitivity of youth and blacks, and Kydland (1984) has documented a similar pattern for less educated workers. Because the hours shares of low-wage groups tend to be procyclical, the aggregate wage statistics commonly used in time series studies give greater weight to low-skill workers during expansions than during recessions. This induces a countercyclical composition bias, which could obscure the degree of real wage procyclicality that the typical worker in any group really faces.

To see the point more formally, note that the derivative of the aggregate wage statistic with respect to a cycle indicator U_t is

$$(4) \quad dW_t/dU_t = \sum_{j=1}^J d(S_j W_j)/dU_t$$

$$= \sum_{j=1}^J S_j (dW_j/dU_t) + \sum_{j=1}^J W_j (dS_j/dU_t)$$

or, in logarithms,

$$(5) \quad d \ln W_t/dU_t = (1/W_t) (dW_t/dU_t)$$

$$= \sum_{j=1}^J S_j^* (d \ln W_j/dU_t) + \sum_{j=1}^J (W_j/W_t) (dS_j/dU_t)$$

where $S_j^* = S_j W_j/W_t$ is group j 's share of the wage bill. Equation (5) says that the cyclical variation in the aggregate wage statistic consists of a weighted average of the cyclical wage changes experienced

by the J groups plus a second term reflecting the cyclical change in the skill composition of total work hours. If groups with low relative wages W_j/W_i have procyclical hours shares, the second term contributes a countercyclical bias. For example, suppose that $J = 2$ with $W_1 > W_2$ (i.e., group 2 is less skilled than group 1) and that both groups experience the same real wage cyclicity $\beta_3 = d \ln W_1 / dU_t = d \ln W_2 / dU_t$. Then equation (5) simplifies to

$$(6) \quad d \ln W_t / dU_t = \beta_3 + [(W_2 - W_1) / W] (dS_2 / dU_t).$$

Thus, if dS_2 / dU_t is procyclical, the measured aggregate wage cyclicity $d \ln W / dU_t$ is systematically less procyclical than β_3 , the true wage cyclicity faced by each group of workers.

Since the source of the measurement problem in aggregate wage data is cyclically shifting weights, the most direct solution to the problem is obvious – construct a wage statistic without cyclically shifting weights. Doing so is straightforward if one has access to longitudinal microdata. Then one can hold composition constant by following the exact same workers over time with fixed weights. Numerous recent studies have done essentially this and, unlike most of the literature based on published aggregate data, they have estimated strong procyclicality in real wages. As will be discussed in Section IV, however, most of the authors of these studies have overlooked the crucial role of composition bias for explaining the seeming contradiction between the longitudinal and time series evidence. In the next section we will present our results on the importance of composition bias, and in Section IV we will show that other longitudinal researchers' results actually are fully consistent with our interpretation of the evidence.

III. Evidence from Longitudinal Microdata

Our main empirical analysis is based on the Panel Study of Income Dynamics (PSID), a national longitudinal survey that has collected data on members of the same families every year since 1968. Our

sample is drawn from the PSID's 1988 cross-year family-individual response-nonresponse file, which is documented in Survey Research Center (1991). Each year's PSID interviews collect information for the preceding calendar year on the annual labor income and hours of work of household heads and their spouses. Therefore, the data from the 1968-88 interviews include labor income and hours measures for 1967-87. Our measure of an individual's hourly wage rate in a given year is simply his/her ratio of annual labor income to annual hours of work.⁶ Like the BLS average hourly earnings statistic analyzed in the preceding section, this wage measure includes work on overtime and second jobs. Our complete data description is relegated to the Appendix, but we will highlight key data details in this section's text as they become relevant.

Simple Evidence from a "Balanced" Sample of Prime-Age Men

Our most straightforward method for avoiding composition bias is to construct a wage statistic that gives fixed weights to the exact same PSID workers over time. To do so, we must have a wage observation in each year from 1967 to 1987 for every worker in the sample. We therefore focus this subsection's analysis on prime-age men, the group most likely to have positive work hours in every year of the sample period. In particular, we restrict this subsection's sample to men born between 1928 and 1942 who were household heads every year from 1968 to 1988 and reported positive labor income and at least 100 hours of work for every year from 1967 to 1987.⁷ The birth year restriction assures that the sample members are between the ages of 25 and 59 throughout the sample period. The resulting sample is "balanced" in the sense that each year's wage information pertains to the exact same 355 men

⁶In contrast, studies based on the National Longitudinal Surveys of labor market experience have used wage reports pertaining to the time of interview. The results of these studies are summarized in Section IV.

⁷Also, for comparability with the analyses in the next subsection, we require information on years of education.

who meet all of the above criteria. The virtue of this sample is that it avoids composition bias in the most direct way imaginable. Its disadvantage is that the wage cyclicality observed for this highly restricted sample may not represent the wage cyclicality experienced by other groups in the labor force. We therefore will analyze broader samples in subsequent subsections.

In Section II, we estimated equation (2) for $\Delta \ln W_t$ with W_t measured by the BLS average hourly earnings statistic, which is contaminated by composition bias. In this subsection, we reestimate equation (2) with $\ln W_t$ measured instead by the sample mean of the log real wage in year t among our 355 prime-age men. As in Section II, we use the implicit GNP deflator to convert from nominal to real wages. Table II shows the resulting time series of $\Delta \ln W_t$. Although the sample means for real wage growth are quite volatile, some systematic tendencies are visible. As one would expect when following a cohort over the life cycle, average real wage growth is positive in most years, but it trends downward. The table also reveals a tendency for real wage growth to be unusually low in recession years.

The first column of Table III shows the results from OLS estimation of equation (2) with the new wage statistic as the dependent variable and time and change in the unemployment rate as the regressors. The estimated cycle coefficient $\hat{\beta}_3 = -.0135$ indicates that, when the unemployment rate increases by an additional percentage point, prime-age men's real wage growth tends to decline by more than a percentage point.* This estimate is more than double Table I's corresponding estimate of $-.0060$, which was based

*If time squared is included as an additional regressor, its estimated coefficient is insignificant, and $\hat{\beta}_3 = -.0136$ with estimated standard error $.0037$. Some readers also may wish to know what happens if we exclude the Survey of Economic Opportunity portion of the PSID sample, which overrepresents the low-income population. Our sample size then declines to 252, and $\hat{\beta}_3$ increases in magnitude to $-.0149$ with estimated standard error $.0038$.

on aggregate data subject to composition bias. The difference between the two estimates of β_3 is statistically significant at the .05 level.⁹

Another difference between the two wage statistics, besides in their susceptibility to composition bias, is that Table I's measure is the log of an average while our new measure is the average of a log. To check that this difference is not responsible for the disparity in results, we reestimate with $\ln W$, measured by the log of the sample mean wage (instead of the sample mean log wage) for our PSID men. As shown in the next-to-last column of Table III, the resulting $\hat{\beta}_3 = -.0168$ suggests even more procyclicality of real wages.

The last column of Table III shows the results from deflating by the CPI instead of by the implicit GNP deflator. Again, the estimated wage cyclicality becomes more procyclical. The difference between the estimates in the first and last columns, $-.0147 - (-.0135) = -.0012$, indicates how much more procyclical all of this paper's estimated coefficients of ΔU_t for 1967-87 would become if the CPI were used in place of the implicit GNP deflator.

The other columns of Table III report the results from using alternative cycle regressors. Using change in the unemployment rate imposes an assumption that the current and lagged unemployment rates have "equal and opposite" coefficients. The specification in the second column relaxes that restriction by entering the current and lagged values as separate regressors. The two coefficient estimates are opposite in sign, significantly different from zero, and again large in magnitude relative to the estimates in Table I. A *t*-test of the equal-and-opposite restriction accepts it at any conventional significance

⁹As Angelo Melino and Dwayne Benjamin pointed out to us, the significance of the difference is readily testable by estimating the regression of the difference between the two wage statistics on time and ΔU_t , and then checking the *t*-ratio for the coefficient of ΔU_t . That coefficient is estimated at .0075 (the difference between -.0060 and -.0135), and its *t*-ratio is 2.22.

level.¹⁰ The third and fourth columns of Table III, which can be compared to the last two columns of Table I, display the results from using growth in real GNP or in per capita hours as the cycle regressor. Like the analyses based on the unemployment rate, these analyses show that the procyclicality of the PSID-based wage statistic is about double that of the BLS statistic.

All these results suggest that, over the 1967-87 period, real wages of prime-age men who worked at least 100 hours every year were considerably more procyclical than one might have supposed from the sort of aggregate wage statistic discussed in Section II. But it certainly is reasonable to wonder whether these results might be an artifact of small sample size or specific to the sample's age range, attachment to the labor force, gender, or other characteristics. Accordingly, the next two subsections pursue more detailed analyses of broader samples.

Additional Evidence from Unbalanced Samples of Men

We now adapt our model for aggregate data to a form suitable for analyzing "unbalanced" longitudinal microdata. Our model for the log real wage of individual i in year t is

$$(7) \ln W_{it} = \gamma_1 + \gamma_2 t + \gamma_3 t^2 + \gamma_4 (U_t - \delta_1 - \delta_2 t - \delta_3 t^2) + \gamma_5' Z_i + \gamma_6 X_{it} + \gamma_7 X_{it}^2 + \epsilon_{it}$$

where U_t again is the civilian unemployment rate, Z_i is a vector of time-invariant worker characteristics such as race and years of education, X_{it} is the worker's years of work experience as of year t , and ϵ_{it} is a random error term. Equation (7) extends the standard log earnings function, popularized by Mincer (1974), to incorporate the general time trend and business cycle regressors from equation (1). If one averages equation (7) by year across all members of a balanced sample, one obtains aggregate equation (1) where $\ln W_t$ denotes the sample mean log wage in year t , the time-invariant sample mean of $\gamma_5' Z_i$

¹⁰We also have tried the lead of the unemployment rate as an additional regressor. Its estimated coefficient is insignificantly different from zero at any conventional level.

is impounded in the intercept, and the sample means of $\gamma_6 X_{it}$ and $\gamma_7 X_{it}^2$ are perfectly collinear over time with t and t^2 . Viewed from this perspective, the trouble with wage statistics that do not hold composition constant is that the yearly mean of $\gamma_5' Z_i + \gamma_6 X_{it} + \gamma_7 X_{it}^2$ changes over time and is positively correlated with the unemployment rate. Consequently, unless this temporal variation in average worker characteristics somehow is fully controlled for, $\hat{\beta}_3 = \hat{\gamma}_4$ is subject to a countercyclical omitted-variables bias.

First-differencing equation (7) yields

$$(8) \quad \Delta \ln W_{it} = \beta_1 + \beta_2 t + \beta_3 \Delta U_t + \beta_4 X_{it} + v_{it}$$

where $v_{it} = \Delta e_{it}$, $\beta_2 = 2(\gamma_3 - \gamma_4 \delta_3)$ as in equation (2), $\beta_1 = \gamma_2 - \gamma_3 + \gamma_4 (\delta_3 - \delta_2) + \gamma_6 - \gamma_7$ encompasses real wage growth due to the accumulation of individual experience as well as general time trends, $\beta_4 = 2 \gamma_7 < 0$ reflects the concavity of the log wage/experience profile, and again $\beta_3 = \gamma_4 \frac{z}{c} > 0$

as the real wage is countercyclical, noncyclical, or procyclical. It is important to note that equation (8) not only accounts for the wage growth effect of experience, but also controls implicitly for the wage effects of race, years of education, and the myriad of less readily measured elements of Z_i (such as motivation and ability) by "differencing out" their effects. This is crucial for treating the composition bias problem. In the previous subsection, we avoided composition bias by restricting our sample so that it contained the exact same workers every year. In this subsection's less restricted samples, who is in the sample will change somewhat from year to year, but sample composition changes with respect to both observable and unobservable elements of Z_i will be accounted for by the differencing approach. This is precisely the same way that Stockman (1983), Bils (1985), and many other previous longitudinal researchers have addressed the composition bias issue.

The first column of Table IV shows the results from OLS estimation of equation (8) using the microdata on the balanced sample of 355 men from the preceding subsection. With 20 year-to-year change observations for each individual, the total number of person-year observations is 7,100. The estimated coefficient of ΔU_t , $\hat{\beta}_3 = -.0135$, is identical to the corresponding Table III estimate based on the aggregated data from the same sample. The two estimates are necessarily equal because the only regressor in equation (8) that varies cross-sectionally, the work experience variable X_{it} , is perfectly correlated in the intertemporal dimension with t . Consequently, the only difference in estimated slope coefficients between the two analyses is that, in the aggregated analysis, the coefficient of time picks up both the general time effect and the effect of the aging of the sample cohort, while the microdata analysis separates these two effects.

Another minor difference between the analyses is that the standard error estimate in Table IV is .0033, as compared to .0035 in Table III. The estimate in Table IV is biased downward by its neglect of the cross-sectional and serial dependence of the error term v_{it} . As discussed by Coleman (1986), v_{it} is cross-sectionally correlated because different workers' error terms share common time effects. In addition, an analysis of serial correlation in the OLS residuals estimates that v_{it} , which is the first difference of ϵ_{it} , has a first-order autocorrelation of about -.4 and higher-order autocorrelations close to zero.¹¹ Accounting for both types of dependence in v_{it} would require a complicated generalized least squares procedure, which we have not undertaken. Instead, we merely emphasize that, based on the comparison of corresponding results in Tables III and IV, the standard error estimates in Table IV appear to be biased slightly downward.

¹¹Similar results are reported in MaCurdy (1982) and Topel (1990). The negligible higher-order autocorrelations support the omission of individual-specific intercepts from equation (8). If "fixed effects" in wage growth were empirically important for our purposes, they would contribute toward substantial positive autocorrelations at all lags.

We now apply our framework for analyzing the microdata to broaden our sample and then to test for various sorts of heterogeneity in real wage cyclicality. Our first modification of the sample is that we undo the "balancedness" restriction -- we no longer require sample members to report positive labor income and at least 100 work hours in every year. Instead, if an individual fails these requirements in some years, we still use his remaining observations for other years. For example, suppose worker i meets the requirements in every year except that, in 1983, he worked fewer than 100 hours. Whereas we previously excluded this worker altogether, he now contributes 18 observations, i.e., every year-to-year change from 1967-68 to 1986-87 except for 1982-83 and 1983-84.¹² As the second column of Table IV shows, this modification dramatically increases the sample size to 17,128 observations of 1,314 workers.¹³ The estimate $\hat{\beta}_3 = -.0145$ becomes a little more procyclical and, of course, its estimated standard error shrinks.

Then we broaden the sample further by relaxing the 1928-42 birth year restriction. Instead, we use all available person-year observations provided that the worker was at least age 16 in that year. As shown in the third column of Table IV, this increases the sample size even more dramatically to 64,847 observations of 7,225 workers. The estimate $\hat{\beta}_3 = -.0140$ is still about the same, and its estimated standard error is now only .0015. The robustness of the result that $\beta_3 = -.014$ to these large changes

¹²The endogeneity of annual hours of work raises the possibility of sample selection bias. As discussed in Bills (1985, pp. 676-77) and Solon and Barsky (1989, p. 20), such bias most likely would be in the direction of underestimating the procyclicity of real wages. If zero or very low work hours represent a labor supply response to poor wage opportunities, such as worsened opportunities in a recession, exclusion of such observations will tend to obscure some of the procyclicity of wage opportunities.

¹³Although the occurrence of years with fewer than 100 work hours is one source of the discrepancy between the numbers of workers in the balanced and unbalanced samples (and the most important source in the women's samples analyzed below), far and away the most important source for men is attrition from and entry into the survey. See Beckett et al. (1988) for a lucid discussion of changes in the PSID sample.

in sample membership suggests that men's real wage cyclicality does not vary greatly with respect to labor force attachment or age.

Next we check directly for various types of heterogeneity in men's real wage cyclicality. First, to see whether real wage cyclicality might vary with skill level, we augment the regressor vector with a variable for the individual's years of education and its interaction with ΔU_i . The results, in the fourth column of Table IV, indicate again that the log real wage response to a unit change in U_i is about -.014 for someone with typical education. The estimated coefficient of the interaction term is small and statistically insignificant.¹⁴ This suggests that the well-known tendency for low-skill workers to have more cyclically variable earnings¹⁵ arises mainly from their greater cyclicality in hours of work, not in hourly wages.¹⁶

Our investigations of heterogeneity with respect to union status and job attachment run into a data complication. The income and hours questions used to construct our wage variable pertain to the calendar year preceding the interview, but the questions on union status and tenure with employer pertain to the time of the interview. We juxtapose the variables as correctly as possible by performing the necessary cross-year manipulations. For example, our 1984 wage observation for worker i is based on his 1985 interview, and our 1984 observation of his union status is based on his 1984 interview. But to simplify these manipulations, we revert to the balanced sample of 355 men. Also, we lose the 1967-68 observations because there was no 1967 interview to elicit 1967 union and tenure status.

¹⁴Also, like Bils (1985), we obtain a minuscule and insignificant coefficient estimate for the uninteracted education variable. Our result for this coefficient, which represents the effect of education on wage growth when $\Delta U_i = 0$, runs counter to the common belief that more educated workers experience faster wage growth.

¹⁵See Vroman (1977), for example.

¹⁶Stockman (1983), Bils (1985), and Keane and Prasad (1991a) also reported that their estimates of real wage cyclicality varied insignificantly with education.

For comparison purposes, we reestimate equation (8) for the balanced sample with 1967-68 dropped from the sample period. The resulting $\hat{\beta}_3 = -.0138$, shown in the fifth column of Table IV, differs only slightly from the original estimate including 1967-68. Then, to check for a union-nonunion difference in real wage cyclicality, we add to equation (7) a dummy variable that equals 1 if the worker is a union member in the relevant year along with interactions of that variable with time, time squared, and the unemployment rate. We estimate the first-differenced form of this expanded equation. The coefficient of ΔU_t reflects the cyclicality of the real wage for nonunion workers, and the coefficient of the change in the interaction of the unemployment rate with the union dummy reflects the incremental cyclicality for union workers. As shown in the next-to-last column of Table IV, the estimated coefficient of ΔU_t is $-.0135$ for nonunion workers and $-.0142 = -.0135 - .0007$ for union workers. The estimated contrast between union and nonunion wage cyclicality is small, but is very imprecisely measured. Our failure to detect a union-nonunion difference in wage cyclicality, however, is altogether consistent with aggregate evidence presented in Pencavel and Hartsog (1984) and Solon and Barsky (1989).¹⁷

Finally, we investigate whether the procyclicality of real wages is due solely to the relatively favorable opportunities for switching employers that arise during an expansion or whether real wage procyclicality also is experienced by workers that stay with the same employer. The PSID is not suited for a definitive contrast between stayers and leavers because measurement error and the various definitions of job tenure used over the years make it difficult to pinpoint who changed employers when.¹⁸ Nevertheless, we are able at least to produce some reliable evidence on wage cyclicality for stayers. We borrow the complex algorithm described in Appendix 2 of Altonji and Shakotko (1985) that

¹⁷Stockman (1983) also reported an insignificant contrast between union and nonunion workers, as well as between white-collar and blue-collar workers. In addition, he and Bils (1985) both reported that their estimates of real wage cyclicality varied insignificantly with industry.

¹⁸See Altonji and Shakotko (1985) and Brown and Light (1992).

uses the various tenure variables in the PSID to impute tenure with employer for every year.¹⁹ Then we adopt a very conservative standard for classifying a worker as a stayer -- he is counted as having stayed with the same employer between years $t - 1$ and t if his employer tenure variable equals at least 1 in year $t - 1$, at least 1.5 in t , and at least 2 in $t + 1$. We consider this standard conservative in the sense that it excludes nearly all leavers at the expense of also excluding some stayers. Then we reestimate equation (8) with only the person-year observations that we are confident are stayer observations. As shown in the last column of Table IV, the resulting $\hat{\beta}_3$ is -.0124 with estimated standard error .0028. If we also exclude self-employed stayer observations, $\hat{\beta}_3$ falls slightly to -.0120 with estimated standard error .0028.

Though silent about how much more wage procyclicality is experienced by leavers, our results constitute strong evidence that prime-age men who stay with the same employer also receive substantially procyclical real wages. Our results are somewhat different from Bils' (1985) results for young men in the National Longitudinal Surveys of labor market experience. His estimates of β_3 for stayers are -.0064 (with estimated standard error .0040) for whites and -.0044 (with estimated standard error .0065) for blacks.

To summarize, our general finding that $\hat{\beta}_3 = -.014$ for men is remarkably robust to variations in sample selection criteria. The procyclicality of men's real wages does not appear to vary much with education level, and it is quite pronounced even for workers that stay with the same employer.

¹⁹We thank Joseph Altonji for his extraordinary helpfulness in sharing this algorithm.

Evidence for Women

Table V presents the results from some analyses of real wage cyclicality for women in the PSID. These analyses correspond to the men's analyses reported in the first three columns of Table IV. The first column of Table V pertains to a balanced sample of women – those born between 1928 and 1942 who were household heads or spouses every year from 1968 to 1988 and reported positive labor income and at least 100 hours of work every year from 1967 to 1987. For women, these restrictions lead to a sample of only 146 individuals. Estimation of equation (8) for this sample produces an estimated cyclicality coefficient of $\hat{\beta}_3 = -.0046$. This is noticeably less procyclical than the estimates for men, but it is imprecisely measured with estimated standard error .0059.

As shown in the second column, relaxing the balancedness restriction while maintaining the birth year restriction leads to a massive increase in sample size. The cyclicality estimate $\hat{\beta}_3 = -.0056$ is only slightly more procyclical and is estimated somewhat more precisely. Finally, as shown in the third column, relaxing the birth year restriction produces another big increase in sample size. The estimate $\hat{\beta}_3 = -.0042$ is less procyclical, and the estimated standard error falls to .0020. Assuming that the women's and men's estimates are approximately uncorrelated, the *t*-ratio for the difference between this women's estimate and the corresponding men's $\hat{\beta}_3 = -.0140$ is almost 4.

Our finding that women's real wages are less procyclical than men's is robust to variations in sample selection criteria and, in the largest samples, is statistically significant. Moreover, as we will see in Section IV, some similar findings are scattered around the previous literature. What could account for this gender difference in wage cyclicality?

The simplest answers come from a basic supply-and-demand analysis. Suppose that cyclical wage and employment fluctuations are driven by labor demand shifts along stable short-run labor supply

functions, which are positively sloped because of intertemporal substitution. The lesser wage procyclicality for women could be due to either or both of two factors. First, cyclical labor demand shifts might be smaller for women than for men. Second, women's short-run labor supply might be more elastic.

Both possibilities seem plausible. With regard to the first, it is well known that cyclical employment variation is concentrated in durable goods manufacturing and construction, and women are underrepresented in both those industries. With regard to the second, little is known about gender differences in intertemporal substitution, but it is widely believed that women's labor supply is generally more wage-elastic than men's.²⁰

Distinguishing between the two explanations calls for examining the gender difference in quantity as well as price variation. If women's lesser wage cyclicality were due solely to a difference in demand cyclicality, then women's hours of work also would be less procyclical. On the other hand, if a difference in labor supply elasticity were the only operative factor, women's lesser wage cyclicality would be accompanied by greater hours cyclicality.

Table VI presents some simple evidence on hours cyclicality by gender. For the period 1967-68 to 1986-87, we use OLS to estimate each gender's regression of change in the log of per capita hours of work²¹ on time and change in the overall civilian unemployment rate. Like Okun (1962) and many subsequent researchers, we find that a one-percentage-point reduction in the unemployment rate is associated with well over a one-percent increase in per capita hours of work. This occurs because of procyclicality in both labor force participation and average work hours of the employed. We estimate that men's hours increase by about 1.8 percent and women's by about 1.4 percent. The lesser quantity

²⁰See Killingsworth (1983, especially p.432), for example.

²¹Each gender's per capita work hours are calculated as the product of the gender's civilian employment/population ratio and its average work hours in nonagricultural industries.

variation for women suggests that more elastic labor supply on the part of women cannot by itself explain their lesser wage cyclicalities.

Nevertheless, although a gender difference in cyclicalities of labor demand gives a better account of why women experience less variability in hours as well as wages, a difference in labor supply elasticity may be part of the story, too. In particular, it may explain why men's cyclical hours variation is only about 25 percent greater than women's even though we have estimated their wage cyclicalities to be about three times greater. Using our previous results that β_3 is about -.014 for men and about -.005 for women, we can write the ratio of women's to men's short-run elasticity of labor supply as

$$(9) \quad \frac{d \ln H_f / d \ln W_f}{d \ln H_m / d \ln W_m} = \left(\frac{d \ln H_f / d U}{d \ln H_m / d U} \right) \left(\frac{d \ln W_m / d U}{d \ln W_f / d U} \right)$$

$$= \left(\frac{-.014}{-.018} \right) \left(\frac{-.014}{-.005} \right)$$

• 2.2 .

This back-of-the envelope calculation suggests that it takes a women's labor supply elasticity more than twice that of men's to square up the magnitudes of the gender differences in hours and wage cyclicalities.

Whatever the reason for the gender difference in cyclicalities of real wages, it poses a new difficulty for analyzing the discrepancy in real wage cyclicalities 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 cyclicalities of men and women. In the next subsection, we incorporate this aggregation problem into an analysis of the quantitative importance of composition bias.

The Importance of Composition Bias

In Section II, we found that, over the period 1967-68 to 1986-87, the regression of real growth in the BLS aggregate wage statistic on time and change in the unemployment rate U yielded an estimated $d \ln W / dU$ of $-.006$. This was far less than the $-.014$ estimate we subsequently obtained for men in the PSID, but a little more than our $-.005$ estimate for the PSID women. We now will use equation (5) from Section II to structure our analysis of how the cyclical variation of the aggregate wage statistic is influenced by both composition bias and the aggregation of men and women. Then we will illustrate the quantitative importance of composition bias by imposing on our PSID data the same sort of hours-weighting applied in the BLS wage statistics.

We begin by splitting total work hours into male hours H_m and female hours H_f . Suppressing the t subscripts, equation (5) then becomes

$$(10) \quad d \ln W / dU = (1 - S_f^*) d \ln W_m / dU + S_f^* d \ln W_f / dU + [(W_f - W_m) / W] dS_f / dU \\ = (1 - S_f^*) d \ln W_m / dU + S_f^* d \ln W_f / dU + [(W_f - W_m) / W] S_f (1 - S_f) (d \ln H_f / dU - d \ln H_m / dU)$$

where S_f is the female share of hours, S_f^* is the female share of the wage bill, and W_m and W_f are the aggregate real wage statistics for men and women.²² Equation (10) expresses the cyclical variation of the overall aggregate wage statistic as a weighted average of the cyclical variations of the male and female wage statistics plus a third term reflecting cyclical variation in the gender composition of total work hours.

Next we partition each gender's hours into those that are not cyclically marginal (H_{m1} for men, H_{f1} for women) and those that are (H_{m2} and H_{f2}). Denoting each group's real wage with the same subscripts, we assume $d \ln W_{m1} / dU = d \ln W_{m2} / dU = \beta_{3m}$, $d \ln W_{f1} / dU = d \ln W_{f2} / dU = \beta_{3f}$, and

²² W_m and W_f are not actually observed. The BLS establishment data are not broken down by gender.

$(W_f - W_m)/W_f = (W_{m2} - W_{m1})/W_m = \delta$. If, within each gender, the cyclically marginal hours are less skilled than the nonmarginal hours, $\delta < 0$. Then applying equation (6) yields

$$\begin{aligned}
 (11) \quad d \ln W_m / dU &= \beta_{3m} + \delta d(H_{m2}/H_m) / dU \\
 &= \beta_{3m} + \delta [(H_m - H_{m2})/H_m^2] dH_m / dU \\
 &= \beta_{3m} + \delta (1 - S_{m2}) d \ln H_m / dU \\
 &= \beta_{3m} + \delta d \ln H_m / dU .
 \end{aligned}$$

The second line of equation (11) comes from $dH_m/dU = dH_{m2}/dU$, which depends on H_m , not being cyclically marginal. The final approximation depends on the cyclically marginal share of total hours being small. Similarly,

$$(12) \quad d \ln W_f / dU = \beta_{3f} + \delta d \ln H_f / dU .$$

Finally, substituting (11) and (12) into (10) leads to

$$\begin{aligned}
 (13) \quad d \ln W / dU &= (1 - S_f^*) \beta_{3m} + S_f^* \beta_{3f} + [(W_f - W_m)/W] S_f (1 - S_f) (d \ln H_f / dU - d \ln H_m / dU) \\
 &\quad + \delta [(1 - S_f^*) d \ln H_m / dU + S_f^* d \ln H_f / dU] .
 \end{aligned}$$

Equation (13) shows that the cyclicity of the BLS aggregate wage statistic is approximately a weighted average of the true cyclicality of men's and women's wages plus two composition bias terms. The first of these is the already-discussed term that reflects cyclical variation in the gender composition of hours. The second reflects cyclical variation in the skill composition of each gender's hours. Despite its imposing appearance, equation (13) is practically useful for characterizing these composition biases because every element in the equation except δ either has been estimated already or is readily observable. As a result, δ itself, which is the proportional gap between the wages paid for marginal and nonmarginal hours, also can be inferred.

In Section II, we estimated the left side of equation (13) at $-.0060$. We will now decompose that cyclicity of the BLS aggregate wage statistic into its components as of 1977, the middle year of our sample period. That year's Current Population Survey statistics by gender on employment and average work hours of the employed imply a female hours share of $S_f = .36$. Given O'Neill's (1985, Table 1) 1977 female-to-male wage ratio $W_f/W_m = .648$, we also have a proportional between-gender wage gap of $(W_f - W_m)/W = -.40$ and a female wage-bill share of $S_f^* = .27$. Finally, we have estimated β_{3m} at $-.014$, β_{3f} at $-.005$, $d\ln H/dU$ at $-.014$, and $d\ln H_m/dU$ at $-.018$.

If not for composition bias, $d\ln W/dU$ would be simply a weighted average of β_{3m} and β_{3f} , the true male and female wage cyclicality. Evaluated at $S_f^* = .27$, that weighted average is $-.0116$.²³ Thus, when the different wage cyclicality of men and women are accounted for, the $-.0060$ estimate of aggregate real wage cyclicity dramatically understates the true procyclicality of the aggregate real wage. The implied countercyclical composition bias of $.0056 = -.0060 - (-.0116)$ can be further decomposed into its two components -- the term involving the cyclicity of the female share of hours and the term involving cyclical variation in the skill composition of each gender's hours. Given that women are paid less than men and have less cyclically variable hours, the gender composition term by itself imposes a procyclical bias, which is estimated at $-.0004$. Because the gender difference in hours cyclicity is not terribly large, this term turns out to be quantitatively unimportant. It is swamped by the countercyclical composition bias from the last term. Calculated as a residual, this skill composition term is estimated at $.0060$. The implied value of δ is $-.35$. This indicates a 35 percent within-gender wage gap between cyclically marginally and nonmarginal hours. This estimate is reminiscent of the many studies cited in Section II that have found that low-paid groups like young, black, and less educated workers experience greater hours cyclicity. Unlike those studies, our estimate provides an omnibus measure of the wage

²³Similarly, when we pool our balanced samples of men and women and augment equation (8) with a gender dummy and its interactions with t and X_{it} but not with ΔU_{it} , the pooled estimate of β_3 is $-.0109$.

gap between cyclically marginal and nonmarginal hours that is not restricted to particular observable characteristics of the workers.

Judging by the discrepancy between the cyclicity of the BLS and PSID wage data, cyclically marginal hours appear to be so low-paid that they induce a quantitatively important composition bias. Even though our PSID estimates imply a true aggregate wage procyclicality of $-.0116$, the BLS aggregate statistic is countercyclically biased to show a cyclicity of only $-.0060$. It is sensible to wonder, though, whether the discrepancy might be due not to composition bias, but to some other difference between the data sources. For example, the BLS wage data are restricted to production and nonsupervisory workers in private nonagricultural employment, and the PSID wage data are restricted to household heads and spouses.

We therefore perform the following mischievous exercise. We purposefully inject a BLS-like composition bias into an aggregate wage statistic based on the PSID. In particular, for each year from 1967 to 1987, we identify every PSID household head or spouse that was at least age 16 and reported positive labor income and at least 100 hours of work that year. The resulting year-by-year samples do not hold composition constant. Then we calculate an hours-weighted average real wage for each year's sample according to the same equation (3) used for the BLS aggregate wage statistic.²⁴

When we regress the change in the log of this PSID-based wage statistic on time and change in the unemployment rate, we obtain $\hat{\beta}_3 = -.0057$ (with estimated standard error $.0025$). This estimate is remarkably close to the $-.0060$ estimate for the BLS statistic, and it diverges even a little further from the more procyclical estimates we got when we used the PSID data to avoid composition bias, not inject it. This convinces us all the more that addressing the composition bias issue is crucial for the proper measurement of real wage cyclicity.

²⁴Also, to adjust for the PSID's oversampling of the low-income population, we weight observations by their inverse probabilities of selection into the sample.

In summary, our analysis of longitudinal microdata from the Panel Study of Income Dynamics has shown that real wages were considerably more procyclical during the 1967-87 period than indicated by aggregate time series data for the same period. The discrepancy is due to a composition bias in the aggregate statistics, which give more weight to low-skill workers during expansions than during recessions. Because longitudinal data have become available only since the late 1960's, it is impossible to extend our analysis to earlier periods. Nevertheless, the literature that documents the greater hours cyclical of low-skill groups of workers does extend to earlier periods, so it seems likely that the countercyclical composition bias in aggregate wage data was substantial in earlier periods, too. Consequently, although the aggregate time series evidence presented in Section II suggests that real wages were less procyclical in earlier periods than in our 1967-87 period, the importance of composition bias suggests that real wages in earlier periods were more procyclical than indicated by the available aggregate wage statistics. By the same token, if low-skill workers experience greater hours cyclical in other countries as well as in the United States, the aggregate wage statistics of other countries may be similarly misleading.

IV. Other Longitudinal Research

A rapidly growing body of empirical research has used longitudinal data from the PSID and the National Longitudinal Surveys (NLS) of labor market experience to estimate the association between real hourly wages and national business cycle indicators. These studies have avoided composition bias by

"differencing out" individual-specific wage effects, as in equation (8) in Section III.²⁵ Although most of the authors of these studies have overlooked the importance of composition bias for reconciling their results with those from aggregate wage data, by and large their results are remarkably similar to ours and consistent with our interpretation of the evidence.

Our estimates of equation (8) for men in the PSID produced a procyclical β_3 of about -.014. Deflating by the CPI instead of the GNP deflator would revise this to -.015. This is precisely the estimate obtained by Coleman (1984) using the CPI and 1968-79 data on PSID men. Using the CPI and 1967-79 PSID data on white men, Rayack (1987) estimated β_3 at -.013. Using the CPI and 1976-84 PSID data on men, Beaudry and DiNardo (1991) estimated β_3 at -.014. In our own earlier work, using the GNP deflator and 1967-84 data on a slightly different sample of PSID men, Solon and Barsky (1989) estimated β_3 at -.013. Stockman's (1983) seminal paper estimated year effects rather than a coefficient of a cycle regressor, but it is straightforward to translate his estimated year effects into an estimate of β_3 by regressing them on time and change in the unemployment rate. Doing so with his estimates based on the CPI and 1967-80 PSID data on household heads (male and female) produces a β_3 of -.014. Using the GNP deflator and 1969-81 PSID data, Blank (1989) estimated the coefficient of real GNP growth, rather than change in the unemployment rate. Her estimates of .72 for white men and .52 for black men

²⁵Some of these studies have "mean-differenced" rather than "first-differenced." A rather different approach has been pursued by Kydland and Prescott (1988), who used PSID data for 1969-82 to construct a skill-weighted index of work hours, which they then divided into the real wage bill to obtain an aggregate real wage statistic with less composition bias. Their results are not readily comparable to others in the literature, but they did find that their adjustment for skill composition led to a considerably more positive correlation between log real wages and log real GNP. Like us, they concluded that the composition bias issue is quantitatively important.

compare to our estimate of .62 in Table III.²⁶ The similarity among these estimates and our own is startling, even to us.

Turning now to estimates for the young men in the NLS, Bills (1985) used the GNP deflator and 1966-80 data to estimate β_3 at -.016 for whites and -.018 for blacks. Using the CPI and 1966-78 data, Tremblay (1990) estimated β_3 at -.015 for whites and -.016 for nonwhites. These estimates, too, are strikingly like our own. The outlier is the study by Keane, Moffitt, and Runkle (1988), which using the CPI and 1966-81 data, estimated β_3 at only -.007. This estimate was generated by an elaborate method designed to correct for sample selection bias within a two-equation model of both wages and employment status.²⁷

Estimates for women are less common and also in less uniform agreement, but they mostly do share our pattern of lesser wage cyclicality for women than for men. Stockman's estimated year effects for wives imply a $\hat{\beta}_3$ of -.011. Tremblay estimated β_3 at -.009 for white women and -.002 for nonwhite women. Blank's estimated coefficients of real GNP growth are .52 for white wives, .48 for white female household heads, .93 for black wives, and -.19 for black female household heads.

²⁶In most of the analyses in a subsequent paper, Blank (1990) did not difference out individual-specific wage effects. The exception is that, in the third row of Table 1, she reported results for the regression of average year-to-year change in real wages (not log real wages) on real GNP growth based on 1969-82 PSID data on men. Using the sample mean wages she reported in Table 2 to convert the estimated coefficients into estimates of the elasticity of real wages with respect to real GNP produces figures of .43 for whites and .26 for blacks. Blank did not discuss why these estimates are so different from those in her 1989 paper. In verbal communication with us, she has speculated that differences in the samples' age restrictions may be part of the explanation.

²⁷The identification problems associated with this sort of estimation have been discussed at length in the literature on union-nonunion wage differences (see Lewis [1986] and Freeman and Medoff [1981]). Indeed, the difficulties of identification are hinted at in the authors' own footnotes. Keane, Moffitt, and Runkle's footnote 19, for example, explained, "Our fixed effects model results in table 2 were obtained from a model with fixed effects only in the wage equation. Attempts to include fixed effects in the employment equation were unsuccessful, for they were estimated to be zero." In an admirable display of honesty in a subsequent paper, Keane and Prasad (1991b, footnote 13) acknowledged, "We discovered that FE [fixed effects] selection model estimates are very sensitive to starting values, and that the results obtained by Keane et al. were only a local maximum."

To summarize, most of the longitudinal estimates for men are about as procyclical as ours, and the estimates for women are, if anything, more procyclical than ours. Other longitudinal researchers' evidence therefore is every bit as much at odds with the aggregate time series evidence as ours is. One might have thought then that Stockman's conjecture about the importance of composition bias would be broadly accepted by now and that macroeconomists' beliefs about real wage cyclicity would have been revised accordingly.

On the contrary, as indicated by the citations in our first paragraph, macroeconomists continue to believe that real wages are only weakly cyclical. They have not had to revise their beliefs because many of the longitudinal studies since Stockman's have devoted little or no attention to composition bias in the aggregate data, and those that have discussed it have downplayed the empirical importance of Stockman's hypothesis.²⁸ Bils discussed composition bias, but concluded that "the impact is not particularly large."²⁹ Coleman (1984, p. 68) likewise declared that "the microdata do not indicate any significant aggregation bias." Reviewing the literature, Kniesner and Goldsmith (1987, p. 1257) concluded that "sample composition effects are empirically unimportant for this issue."

An examination of the numbers, however, reveals that these conclusions are unwarranted. When Bils estimated β_3 with the BLS aggregate wage statistic over the same years covered by his NLS data, he obtained $\hat{\beta}_3 = -.012$. As discussed above in footnote 4, this unusually procyclical estimate from aggregate data is idiosyncratic to the particular ten years in his sample. Nevertheless, it still is

²⁸A partial exception is Blank (1990). In the second row of her Table 1, she reported results from an hours-weighting exercise similar to the one we described in the preceding section. Like us, she found that creating a composition bias in the PSID data led to dramatically smaller estimates of real wage cyclicity. She did not stress the implications of this finding for bias in aggregate wage statistics, however, and did not mention the finding in her abstract, introduction, or conclusion. Instead, her title ("Why Are Wages Cyclical in the 1970s?"), opening paragraph, and final paragraph give the impression that the discrepancy between evidence from microdata and aggregate data is due mainly to differences in sample period.

²⁹Bils (1985, p. 668). Also see his discussion on p. 684.

considerably less procyclical than his microdata estimates of $-.016$ for whites and $-.018$ for blacks. Coleman did not present aggregate results for the 1968-79 period covered by his microdata, but his estimate of β , based on the BLS aggregate wage data for 1961-79 is $-.003$, drastically smaller than his microdata estimate of $-.015$.

So why did these authors conclude as they did? For Bils, a key factor was his estimate that the cyclically marginal young man is paid only 19 percent less than the nonmarginal young man. After emphasizing this finding in his concluding section (p. 684), he inferred, "This bias is unimportant relative to the very procyclical wage behavior found here." What Bils overlooked was that the 19-percent figure applies only within the category of young men. It therefore neglects cross-category factors such as the tendency for young workers to be paid less and have much more cyclically variable hours than more mature workers. Coleman's conclusion was based on findings that an hours-weighted real wage series is very highly correlated with a series not contaminated by composition bias and that the correlation with change in the unemployment rate is only somewhat less for the former series than for the latter. What Coleman overlooked was that these correlations in no way deny that an hours-weighted series substantially understates the amplitude of cyclical fluctuations in real wages.

Had we been the first to find that longitudinal microdata reveal substantial procyclicity of real wages, we might have been pleased by the novelty of our results, but we would have been fearful that, despite our best efforts, our results might have arisen from peculiarities of our sample or methodology. Instead, it turns out that many longitudinal researchers have come up with similar numbers, though they have not generally emphasized the importance of composition bias. Our analysis indicates that composition bias does account for the discrepancy between the aggregate time series evidence and the longitudinal evidence produced by ourselves and others.

V. Implications for Macroeconomic Theory

As the citations in this paper's first paragraph suggest, the consensus among a broad spectrum of leading macroeconomists is that real wages are only weakly cyclical. We hope by now to have convinced the reader that this consensus is built on a statistical illusion. Longitudinal evidence from both the PSID and NLS indicates that individuals' real wages are considerably more procyclical than they appear in aggregate statistics afflicted by composition bias.

As discussed in Section I, the belief that real wages are only mildly procyclical has motivated macroeconomists to devise numerous theories in which the effective labor supply curve is nearly flat. Our findings suggest that effective labor supply need not be quite so elastic after all and that theories predicting substantially procyclical real wages are not necessarily at odds with the facts. For example, we noted in Section II that, given the modest cyclicity of aggregate real wage data over our sample period, rationalizing the observed cyclical variation in hours of work as notional labor supply behavior seemingly requires a short-run labor supply elasticity between 1.2 and 2.7. Many economists view such an elasticity as implausibly large, and they accordingly have criticized intertemporal substitution theories that interpret cyclical hours variation as occurring along a short-run notional labor supply curve. The greater real wage procyclicity evident in the longitudinal data, however, implies a smaller labor supply elasticity. Estimating the regression of real wage growth as measured for our balanced sample of PSID men on time and growth in men's per capita work hours yields an estimated inverse labor supply elasticity of .67. As explained in Section II, the inverted value 1.5 is an upward-inconsistent estimate of the elasticity. Estimating the direct regression of hours growth on time and real wage growth produces

a downward-inconsistent estimate of .62.³⁰ This range of .6 to 1.5 for the short-run elasticity of labor supply strains credulity less than the 1.2-to-2.7 range implied by the aggregate wage data.³¹

Our finding that real wages and work hours are quite positively correlated in the aggregate, however, need not imply that the two variables are connected by notional labor supply behavior. For one thing, the .6-to-1.5 range for the short-run labor supply elasticity implies an intertemporal substitution elasticity in the same range only if cyclical wage innovations are highly transitory. If they are quite persistent, as suggested by Altonji and Ashenfelter (1980) among others, squaring up the observed covariation between aggregate hours growth and real wage growth requires a higher intertemporal substitution elasticity. And, if real wages were less procyclical in earlier periods, as suggested by the time series evidence in Section II, the intertemporal substitution elasticity required to fit the earlier observations is larger still. The microeconomic studies by MaCurdy (1981), Altonji (1986), Ham (1986), and Ball (1990), however, have estimated that the intertemporal substitution elasticity for men is not even as high as .6.

In any case, alternative theories of positively-sloped effective labor supply, such as the efficiency wage model of Shapiro and Stiglitz (1984), also predict positive covariation of aggregate real wages and work hours. In the Shapiro-Stiglitz model, unlike the intertemporal substitution model, workers laid off during a recession are off their notional labor supply functions, and the real wage growth associated with an expansion occurs despite excess supply of labor. Therefore, the question of why the effective labor supply curve slopes upward, like the question of why labor demand shifts cyclically along that curve, remains an open and crucially important topic for further inquiry.

³⁰This estimate is similar to estimates reported by Angrist (1991), whose aggregate wage and hours variables both were based on PSID men. Angrist acknowledged that the consistency of his estimation requires the aggregate labor supply function to be perfectly stable, i.e., to have no error term. He also noted that several previous microdata-based studies of intertemporal substitution in labor supply implicitly require the same assumption.

³¹Performing the same exercise for our balanced sample of PSID women is not very useful. It generates a range from nearly 0 to 12.

Table I
Estimates of Cyclicity of Aggregate Real Wage Variable

	1947-48 to 1986-87		1967-68 to 1986-87		
Cycle regressor					
Δ civilian unemployment rate	-.0030 (.0013)		-.0060 (.0017)		
Δ ln (real GNP)		.153 (.059)		.293 (.077)	
Δ ln (per capita hours of work)					.373 (.101)
R ²	.57	.58	.50	.54	.53
Durbin-Watson statistic	1.73	1.92	1.44	1.68	1.38

Numbers in parentheses are standard error estimates.

Table II
Yearly Averages of Real Wage Growth for Balanced Sample of PSID Men

<u>Year</u>	<u>Civilian unemployment rate</u>	<u>Δ average ln (wage/GNP deflator)</u>
1968	3.6	.032
1969	3.5	.021
1970	4.9	.038
1971	5.9	.014
1972	5.6	.052
1973	4.9	.031
1974	5.6	-.008
1975	8.5	-.032
1976	7.7	.034
1977	7.1	.027
1978	6.1	-.015
1979	5.8	.007
1980	7.1	-.002
1981	7.6	.002
1982	9.7	-.023
1983	9.6	.006
1984	7.5	.052
1985	7.2	.003
1986	7.0	-.022
1987	6.2	-.020

Table III
 Estimates of Real Wage Cyclicalities for Balanced Sample of PSID Men,
 1967-68 to 1986-87

	Dependent variable	
	Δ average ln (wage/GNP deflator)	Δ ln average (wage/GNP deflator)
Cycle regressors		
Δ unemployment rate	-.0135 (.0035)	-.0168 (.0039)
current unemployment rate	-.0114 (.0037)	
lagged unemployment rate	.0168 (.0041)	
Δ ln (real GNP)	.617 (.165)	
Δ ln (per capita hours of work)		.699 (.233)
R ²	.58	.53
Durbin-Watson statistic	1.75	1.94
		Δ average ln (wage/CPI)
		-.0147 (.0044)

Numbers in parentheses are standard error estimates.

Table IV
 Estimates of Real Wage Cyclicity for Various Samples of PSID Men

Cycle regressors	1967-68 to 1986-87		1968-69 to 1986-87	
	Estimate (Standard Error)	Estimate (Standard Error)	Estimate (Standard Error)	Estimate (Standard Error)
Δ unemployment rate	-.0135 (.0033)	-.0140 (.0015)	-.0138 (.0034)	-.0124 (.0028)
(Δ unemployment rate) x years of education			.0001 (.0005)	
Δ (unemployment rate x union dummy)				-.0007 (.0069)
Sample of workers	Balanced	Unbalanced, age \geq 16	Balanced	Balanced, but restricted to "stayer" observations
Number of workers in sample	355	7,225	355	337
Number of person-year observations in sample	7,100	64,847	6,745	4,928

Numbers in parentheses are standard error estimates.

Table V
 Estimates of Real Wage Cyclicity for PSID Women,
 1967-68 to 1986-87

	Sample of workers		
	Balanced	Unbalanced, born 1928-42	Unbalanced, age \geq 16
Cycle regressor:			
Δ unemployment rate	-.0046 (.0059)	-.0056 (.0036)	-.0042 (.0020)
Number of workers in sample	146	1,301	6,801
Number of person-year observations in sample	2,920	13,516	50,531

Numbers in parentheses are standard error estimates.

Table VI
 Cyclical Sensitivity of Men's and Women's Per Capita Hours of Work,
 1967-68 to 1986-87

	Men	Women
Estimated coefficient of ΔU_t	-.0181 (.0014)	-.0143 (.0018)
R ²	.91	.80
Durbin-Watson statistic	1.62	2.17

Numbers in parentheses are standard error estimates.

Appendix: Data Sources

Aggregate Data

The following variables are from Economic Report of the President, 1991: real GNP (Table B-2), implicit GNP deflator (Table B-3), consumer price index (Table B-58), civilian unemployment rate and employment/population ratio (Table B-32), civilian employment by gender (Table B-33), and civilian employment/population ratios by gender (Table B-36). For average hourly earnings of production or nonsupervisory workers in private nonagricultural employment, the sources are Handbook of Labor Statistics, 1975, Table 98, for 1947-63 and Employment and Earnings, April 1992, Table C-1, for 1964-87. Per capita hours of work, total and by gender, are calculated as the product of the relevant values for the civilian employment/population ratio and average work hours of the employed. The latter variable is from the January issues of Employment and Earnings.

Panel Study of Income Dynamics

Our data are drawn from the 1988 cross-year family-individual response-nonresponse file, which is documented in Survey Research Center (1991). The data were collected in annual interviews from 1968 to 1988. The responses concerning annual labor income and hours of work pertain to the preceding calendar years 1967-87, while the responses concerning union status and job tenure pertain to current employment as of the interview date. Observations with "major assignments" imputed for labor income or work hours are excluded from our sample. Except as indicated in footnote 7, we use both the Survey Research Center and Survey of Economic Opportunity components of the PSID.

Our hourly wage measure is the ratio of total annual labor income to total annual hours of work. Like the BLS average hourly earnings variable described above, this PSID measure includes work on overtime and second jobs. Thanks to assistance from the Survey Research Center, we replaced all "top-coded" values of annual labor income with their true values from the original interview forms. These

values are available on request from Gary Solon.³¹

Gender is taken from the most recent available report. Whenever possible, we measure years of education by the 1984 report of highest grade completed, with the category 17 or more assigned a value of 18. If a 1984 report is unavailable, we use the most recent available report. If that report is in "bracketed" form, we assign 3 years of education to individuals in the 0-5 category, 7 years to the 6-8 category, 10 years to the 9-11 category, 12 years to the high school graduate category, 14.5 years to the college-without-degree category, 16 years to the college degree category, and 18 years to the advanced degree category. Age is based on the most recent report of birth year if that report is in the 1983 interview or later; otherwise, it is inferred from the most recent report of age. In a number of cases, however, birth year or age obviously has been miscoded. Therefore, in every instance in which a household head or spouse with positive labor income and at least 100 work hours initially is measured as under age 16, we refer instead to the next most recent report of birth year or age. We impute years of work experience as age minus years of education minus 6. If this imputation comes out negative, we reset it to zero.

For 1968-72 and 1974-81, we classify a worker as a union member if he responded affirmatively to the question "Do you belong to a labor union?" For 1982-87, we classify him as a union member if he responded affirmatively to both "Is your current job covered by a union contract?" and "Do you belong to that labor union?" To assure that the resulting variable is sufficiently consistent over time, we have examined the years 1976-81, when both question sequences were asked, and found that the outcomes of our two classification procedures match up quite closely. In the 1973 survey, union status was not elicited. We classify a worker as a union member in 1973 if (1) he had been in his current job for at

³¹PSID users should note that some observations of the hourly wage variable on the PSID tape, even though they seem not to be top-coded, actually were computed by dividing a top-coded value of annual labor income by annual hours of work. Obtaining a correct wage value in these cases requires recovering the true value of labor income.

least a year and was a union member in 1972 or (2) he had been in his current job for less than a year, indicated in 1974 that he had been in his 1974 job for at least a year, and was a union member in 1974. Our construction of the union variable takes account of the fact that it pertains to current status as of the interview date, unlike the income and work hours question, which refer to the preceding calendar year. For example, our 1984 wage variable is based on the 1985 interview responses about labor income and work hours, while the 1984 union status variable is based on the 1984 interview.

To measure tenure with employer, we use the algorithm described in Altonji and Shakotko (1985, Appendix 2). Our use of this variable to identify "stayers" is explained in the text.

References

- Abel, Andrew B., and Bernanke, Ben S., Macroeconomics, Reading, MA: Addison-Wesley, 1992.
- Altonji, Joseph G., "Intertemporal Substitution in Labor Supply: Evidence from Micro Data," Journal of Political Economy, June 1986, 94, S176-215.
- Altonji, Joseph, and Ashenfelter, Orley, "Wage Movements and the Labour Market Equilibrium Hypothesis," Economica, August 1980, 47, 217-45.
- Altonji, Joseph, and Shakotko, Robert, "Do Wages Rise with Job Seniority?" Working Paper No. 1616, National Bureau of Economic Research, May 1985.
- Angrist, Joshua D., "Grouped-Data Estimation and Testing in Simple Labor-Supply Models," Journal of Econometrics, February-March 1991, 47, 243-66.
- Ashenfelter, Orley, "Macroeconomic Analyses and Microeconomic Analyses of Labor Supply," Carnegie-Rochester Conference Series on Public Policy, Autumn 1984, 21, 117-55.
- Ball, Laurence, "Intertemporal Substitution and Constraints on Labor Supply: Evidence from Panel Data," Economic Inquiry, October 1990, 28, 706-24.
- Barro, Robert J., and King, Robert G., "Time-Separable Preferences and Intertemporal-Substitution Models of Business Cycles," Quarterly Journal of Economics, November 1984, 99, 817-39.
- Beaudry, Paul, and DiNardo, John, "The Effect of Implicit Contracts on the Movement of Wages over the Business Cycle: Evidence from Micro Data," Journal of Political Economy, August 1991, 99, 665-88.
- Beckett, Sean, Gould, William, Lillard, Lee, and Welch, Finis, "The Panel Study of Income Dynamics after Fourteen Years: An Evaluation," Journal of Labor Economics, October 1988, 6, 472-92.
- Bils, Mark J., "Real Wages over the Business Cycle: Evidence from Panel Data," Journal of Political Economy, August 1985, 93, 666-89.
- Blanchard, Olivier Jean, and Fischer, Stanley, Lectures on Macroeconomics, Cambridge, MA: MIT Press, 1989.
- Blank, Rebecca M., "Disaggregating the Effect of the Business Cycle on the Distribution of Income," Economica, May 1989, 56, 141-63.
- _____, "Why Are Wages Cyclical in the 1970s?" Journal of Labor Economics, January 1990, 8, 16-47.
- Brown, James N., and Light, Audrey, "Interpreting Panel Data on Tenure," Journal of Labor Economics, July 1992, 10, 219-57.
- Christiano, Lawrence J., and Eichenbaum, Martin, "Current Real-Business-Cycle Theories and Aggregate

- Labor-Market Fluctuations," American Economic Review, June 1992, 82, 430-50.
- Clark, Kim B., and Summers, Lawrence H., "Demographic Differences in Cyclical Employment Variation," Journal of Human Resources, Winter 1981, 16, 61-79.
- Coleman, Thomas S., "Essays in Aggregate Labor Market Business Cycle Fluctuations," unpublished dissertation, University of Chicago, December 1984.
- _____, "A Note on the Measurement of Aggregate Effects in Panel Data," unpublished, November 1986.
- Fischer, Stanley, "Recent Developments in Macroeconomics," Economic Journal, June 1988, 98, 294-339.
- Freeman, Richard B., and Medoff, James L., "The Impact of Collective Bargaining: Illusion or Reality?" in Jack Steiber, Robert B. McKersie, and D. Quinn Mills (eds.), U.S. Industrial Relations 1950-1980: A Critical Assessment, Madison: Industrial Relations Research Association, 1981, pp. 47-97.
- Friedman, Milton, "The Role of Monetary Policy," American Economic Review, March 1968, 58, 1-17.
- Greenwald, Bruce D., and Stiglitz, Joseph E., "Examining Alternative Macroeconomic Theories," Brookings Papers on Economic Activity, 1988: 1, 207-60.
- Hall, Robert E., "Examining Alternative Macroeconomic Theories: Comment," Brookings Papers on Economic Activity, 1988: 1, 261-64.
- Hall, Robert E., and Taylor, John B., Macroeconomics: Theory, Performance, and Policy, 3rd ed., New York: W.W. Norton, 1991.
- Ham, John C., "Testing Whether Unemployment Represents Intertemporal Labour Supply Behaviour," Review of Economic Studies, August 1986, 53, 559-78.
- Keane, Michael, Moffitt, Robert, and Runkle, David, "Real Wages over the Business Cycle: Estimating the Impact of Heterogeneity with Micro Data," Journal of Political Economy, December 1988, 96, 1232-66.
- Keane, Michael J., and Prasad, Eswar S., "The Relation between Skill Levels and the Cyclical Variability of Employment, Hours, and Wages," Discussion Paper No. 41, Institute for Empirical Macroeconomics, Federal Reserve Bank of Minneapolis, May 1991a.
- _____, "The Employment and Wage Effects of Oil Price Shocks: A Sectoral Analysis," Discussion Paper No. 51, Institute for Empirical Macroeconomics, Federal Reserve Bank of Minneapolis, September 1991b.
- Keynes, John Maynard, The General Theory of Employment, Interest, and Money, London: Macmillan, 1936.

- Killingsworth, Mark R., Labor Supply, Cambridge, U.K.: Cambridge University Press, 1983.
- Kniesner, Thomas J., and Goldsmith, Arthur H., "A Survey of Alternative Models of the Aggregate U.S. Labor Market," Journal of Economic Literature, September 1987, 25, 1241-80.
- Kosters, Marvin, and Welch, Finis, "The Effects of Minimum Wages on the Distribution of Changes in Aggregate Employment," American Economic Review, June 1972, 62, 323-32.
- Kydland, Finn E., "Labor-Force Heterogeneity and the Business Cycle," Carnegie-Rochester Conference Series on Public Policy, Autumn 1984, 21, 173-208.
- Kydland, Finn E., and Prescott, Edward C., "Cyclical Movements of the Labor Input and Its Real Wage," Working Paper No. 413, Research Department, Federal Reserve Bank of Minneapolis, December 1988.
- Leamer, Edward E., "Is It a Demand Curve, or Is It a Supply Curve? Partial Identification through Inequality Constraints," Review of Economics and Statistics, August 1981, 63, 319-27.
- Lewis, H. Gregg, Union Relative Wage Effects: A Survey, Chicago: University of Chicago Press, 1986.
- Lucas, Robert E., Jr., "Understanding Business Cycles," Carnegie-Rochester Conference Series on Public Policy, 1977, 5, 7-29.
- Lucas, Robert E., Jr., and Rapping, Leonard A., "Real Wages, Employment, and Inflation," Journal of Political Economy, September-October 1969, 77, 721-54.
- MaCurdy, Thomas E., "An Empirical Model of Labor Supply in a Life-Cycle Setting," Journal of Political Economy, December 1981, 89, 1059-85.
- _____, "The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis," Journal of Econometrics, January 1982, 18, 83-114.
- Mankiw, N. Gregory, "Real Business Cycles: A New Keynesian Perspective," Journal of Economic Perspectives, Summer 1989, 3, 79-90.
- Mincer, Jacob, Schooling, Experience, and Earnings, New York: National Bureau of Economic Research, 1974.
- Mitchell, Mark L., Wallace, Myles S., and Warner, John T., "Real Wages over the Business Cycle: Some Further Evidence," Southern Economic Journal, April 1985, 51, 1162-73.
- Okun, Arthur M., "Potential GNP: Its Measurement and Significance," Proceedings of the Business and Economic Statistics Section of the American Statistical Association, 1962, 98-104.
- _____, "Upward Mobility in a High-Pressure Economy," Brookings Papers on Economic Activity, 1973: 1, 207-52.

- O'Neill, June, "The Trend in the Male-Female Wage Gap in the United States," Journal of Labor Economics, January 1985, 3, S91-116.
- Pencavel, John, and Hartsog, Catherine E., "A Reconsideration of the Effects of Unionism on Relative Wages and Employment in the United States, 1920-1980," Journal of Labor Economics, April 1984, 2, 193-232.
- Phelps, Edmund S., "Introduction: The New Microeconomics in Employment and Inflation Theory," in Edmund S. Phelps (ed.), Microeconomic Foundations of Employment and Inflation Theory, New York: W. W. Norton, 1970, pp. 1-23.
- Prescott, Edward C., "Response to a Skeptic," Federal Reserve Bank of Minneapolis Quarterly Review, Fall 1986, 10, 28-33.
- Rayack, Wendy, "Sources and Centers of Cyclical Movement in Real Wages: Evidence from Panel Data," Journal of Post Keynesian Economics, Fall 1987, 10, 3-21.
- Rotemberg, Julio J., and Woodford, Michael, "Markups and the Business Cycle," in Oliver Jean Blanchard and Stanley Fischer (eds.), NBER Macroeconomics Annual 1991, Cambridge, MA: MIT Press, 1991, pp. 63-129.
- Shapiro, Carl, and Stiglitz, Joseph E., "Equilibrium Unemployment as a Worker Discipline Device," American Economic Review, June 1984, 74, 433-44.
- Solon, Gary, and Barsky, Robert, "Real Wages over the Business Cycle," Working Paper No. 2888, National Bureau of Economic Research, March 1989.
- Stockman, Alan C., "Aggregation Bias and the Cyclical Behavior of Real Wages," unpublished, July 1983.
- Survey Research Center, A Panel Study of Income Dynamics: Procedures and Tape Codes, 1988 Interviewing Year -- Wave XXI, A Supplement, Ann Arbor: Institute for Social Research, University of Michigan, 1991.
- Topel, Robert, "Specific Capital, Mobility, and Wages: Wages Rise with Job Seniority," Working Paper No. 3294, National Bureau of Economic Research, March 1990.
- Tremblay, Carol Horton, "Wage Patterns of Women over the Business Cycle," Quarterly Review of Economics and Business, Spring 1990, 30, 90-101.
- Vroman, Wayne, "Worker Upgrading and the Business Cycle," Brookings Papers on Economic Activity, 1977: 1, 229-50.