OXFORD JOURNALS
OXFORD UNIVERSITY PRESS

Measuring the Cyclicality of Real Wages: How Important is Composition Bias
Author(s): Gary Solon, Robert Barsky and Jonathan A. Parker
Source: *The Quarterly Journal of Economics*, Vol. 109, No. 1 (Feb., 1994), pp. 1-25
Published by: Oxford University Press
Stable URL: http://www.jstor.org/stable/2118426
Accessed: 18-03-2015 11:09 UTC

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at
http://www.jstor.org/page/info/about/policies/terms.jsp

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Oxford University Press is collaborating with JSTOR to digitize, preserve and extend access to *The Quarterly Journal of Economics.*

http://www.jstor.org

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

THE
QUARTERLY JOURNAL OF ECONOMICS

Vol. CIX February 1994 Issue 1

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

GARY SOLON
ROBERT BARSKY
JONATHAN A. PARKER

In the period since the 1960s, 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 1960s. 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.

I. INTRODUCTION

Because aggregate time series on real wages display little cyclicality, macroeconomists commonly have described weak cyclicality 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

*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 advice from Joseph Altonji, the referees, and seminar participants at the University of Michigan, Harvard University, the University of Kent, the London School of Economics, McMaster University, the National Bureau of Economic Research, Princeton University, the University of Rochester, the University of Toronto, and the University of Warwick.

© 1994 by the President and Fellows of Harvard College and the Massachusetts Institute of Technology.

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

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, 454], Christiano and Eichenbaum [1992, pp. 430–31], Greenwald and Stiglitz [1988, pp. 223, 225, 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's *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

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

2. In 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.

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

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 1960s, real wages have been highly procyclical in that period even though aggregate real wage data for the same period have not been nearly 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 this 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 cyclicality 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 cyclicality 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. Section IV documents the empirical importance of composition

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

bias and demonstrates that the results in previous studies (even those that dismissed the importance of composition bias) are consistent with our interpretation of the evidence. Section V discusses implications for macroeconomic theory.

II. The Difficulty of Measuring Real Wage Cyclicality

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

$$\text{(1)} \quad \ln W_t = \gamma_1 + \gamma_2 t + \gamma_3 t^2 + \gamma_4 (U_t - \delta_1 - \delta_2 t - \delta_3 t^2) + \epsilon_t,$$

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 ϵ_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, $\gamma_4 \gtreqless 0$ as W_t is countercyclical, noncyclical, or procyclical.

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

$$\text{(2)} \quad \Delta \ln W_t = \beta_1 + \beta_2 t + \beta_3 \Delta U_t + \nu_t,$$

where

$$\nu_t = \Delta\epsilon_t, \beta_1 = \gamma_2 - \gamma_3 + \gamma_4 (\delta_3 - \delta_2),$$
$$\beta_2 = 2(\gamma_3 - \gamma_4\delta_3), \text{ and } \beta_3 = \gamma_4 \gtreqless 0$$

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.

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 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 1947–1948 to 1991–1992 with the unemployment rate as the cycle indicator. The estimated coefficient of ΔU_t, $\hat{\beta}_3 = -0.0028$, indicates mild, but

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

TABLE I
ESTIMATES OF CYCLICALITY OF BLS WAGE VARIABLE

	1947–1948 to 1991–1992		1967–1968 to 1986–1987		
Cycle regressor					
Δ unemployment rate	−0.0028 (0.0012)		−0.0060 (0.0017)		
Δ ln (real GNP)		0.146 (0.055)		0.293 (0.077)	
Δ ln (per capita hours of work)					0.373 (0.101)
R^2	0.58	0.60	0.50	0.54	0.53
Durbin-Watson statistic	1.64	1.80	1.44	1.68	1.38

Numbers in parentheses are standard error estimates.

statistically significant, procyclicality of W_t. It implies that, when the unemployment rate increases by an additional percentage point, real wage growth declines by about a quarter of a percentage point. 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 0.146 again indicates mild procyclicality.

Next, for purposes of comparability with our later analyses of longitudinal microdata, we reestimate equation (2) for only the recent subperiod 1967–1968 to 1986–1987. Doing so approximately doubles the estimated procyclicality, with the estimated coefficient of ΔU_t increasing from -0.0028 to -0.0060 and that of real GNP growth increasing from 0.146 to 0.293. For this 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 0.373 estimate of the elasticity of W_t with respect to aggregate hours of work.

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 that 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

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

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 $1/0.373 = 2.68$. 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, and if the price and quantity variables are positively correlated (as they are here), the inconsistency of our inverse estimator is in an upward direction. On the other hand, OLS estimation of the *direct* regression of hours growth on real wage growth produces a *downward*-inconsistent estimator of the supply elasticity.[3] 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 cyclicality 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, treatment of overtime pay, choice of deflator, and disaggregation by industry.[4] The point of the present paper, however, is to show that this typical finding is a spurious artifact of a composition bias in the

3. Both inconsistency results hold a fortiori when the variables are subject to classical measurement error. Given that the price and quantity variables are positively correlated, the condition of zero correlation between the supply and demand error terms is sufficient, but not necessary, for Leamer's bracketing result. The necessary condition is that the coefficient in the regression of the demand error on the supply error is less than one and greater than the ratio of the demand elasticity to the supply elasticity. If that coefficient exceeds one, both estimators are upward-inconsistent. If it is more negative than the ratio of elasticities, both estimators are downward-inconsistent.

4. See Solon and Barsky [1989] for a series of robustness checks and references to the literature. As we discuss there, disaggregating by industry makes real wages appear even less procyclical. Aggregating industries into an economywide wage measure adds some procyclicality because the employment shares of certain high-wage industries, particularly durable goods manufacturing and construction, are procyclical. This well-known effect is sometimes described as a procyclical "bias," but the greater opportunities workers have during expansions to advance into higher-paying industries are a *genuine* source of procyclicality in workers' real wages.

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

aggregate wage statistics. Although the possibility of such a bias was pointed out by Stockman [1983], the citations in our first paragraph indicate that the practical importance of this bias has not yet been taken seriously. Our empirical investigation has led us to take it very seriously, and this paper describes the evidence that has convinced us.

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 that the relevant worker population is divided into groups $j = 1, 2, \ldots, J$ with B_{jt} denoting the jth group's wage bill, H_{jt} its total work hours, $S_{jt} = H_{jt}/H_t$ its share of the population's work hours, and $W_{jt} = B_{jt}/H_{jt}$ its average hourly earnings. Then the overall wage statistic W_t can be expressed as

$$W_t = \frac{B_t}{H_t} = \sum_{j=1}^{J} \frac{B_{jt}}{H_t} = \sum_{j=1}^{J} \frac{H_{jt}W_{jt}}{H_t} = \sum_{j=1}^{J} S_{jt}W_{jt}. \tag{3}$$

As the last expression 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

$$\frac{dW_t}{dU_t} = \sum_{j=1}^{J} \frac{d(S_{jt}W_{jt})}{dU_t} = \sum_{j=1}^{J} S_{jt}\left(\frac{dW_{jt}}{dU_t}\right) + \sum_{j=1}^{J} W_{jt}\left(\frac{dS_{jt}}{dU_t}\right) \tag{4}$$

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

or, in logarithms,

$$(5)\quad \frac{d\ln W_t}{dU_t} = \left(\frac{1}{W_t}\right)\left(\frac{dW_t}{dU_t}\right) = \sum_{j=1}^{J} S_{jt}^{*}\left(\frac{d\ln W_{jt}}{dU_t}\right) + \sum_{j=1}^{J}\left(\frac{W_{jt}}{W_t}\right)\left(\frac{dS_{jt}}{dU_t}\right),$$

where $S_{jt}^{*} = S_{jt}\, W_{jt}/W_t$ 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_{jt}/W_t have procyclical hours shares, the second term contributes a countercyclical bias. For example, suppose that $J = 2$ with $W_{1t} > W_{2t}$ (i.e., group 2 is less skilled than group 1) and that both groups experience the same real wage cyclicality $\beta_3 = d\ln W_{1t}/dU_t = d\ln W_{2t}/dU_t$. Then equation (5) simplifies to

$$(6)\quad \frac{d\ln W_t}{dU_t} = \beta_3 + \left[\frac{W_2 - W_1}{W}\right]\left(\frac{dS_{2t}}{dU_t}\right).$$

Thus, if dS_{2t}/dU_t is procyclical, the measured aggregate wage cyclicality $d\ln W_t/dU_t$ is systematically less procyclical than β_3, the true wage cyclicality 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. We pursue this approach in the next section.

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 1968–1988 interviews include labor income and hours measures for 1967–1987. Our

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

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.[5] Like the BLS average hourly earnings statistic analyzed in the preceding section, this wage measure includes work on overtime and second jobs. A complete data description is given in the Appendix.

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 begin by focusing on prime-age men, the group most likely to have positive work hours in every year of the sample period. In particular, we restrict our initial 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.[6] 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 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 disadvantages are that the sample of workers is small and the population represented by this sample may experience different wage cyclicality than is experienced by other groups in the labor force. We therefore shall analyze broader samples below.

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. Now 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.

The first column of Table II 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 $\beta_3 = -0.0135$

5. In 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 Solon, Barsky, and Parker [1992].

6. Also, for comparability with subsequent analyses, we require information on years of education.

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

TABLE II
ESTIMATES OF CYCLICALITY OF WAGE VARIABLES FROM PSID, 1967–1968 TO 1986–1987

	Men's samples				Women's samples	
	Balanced			Unbalanced	Balanced	Unbalanced
Cycle regressor						
Δ unemployment rate	−0.0135			−0.0140	−0.0046	−0.0053
	(0.0035)			(0.0020)	(0.0073)	(0.0040)
Δ ln (real GNP)		0.617				
		(0.165)				
Δ ln (per capita hours of work)			0.699			
			(0.233)			
R^2	0.58	0.57	0.49	0.82	0.16	0.12
Durbin-Watson statistic	1.75	2.11	1.99	2.14	2.55	1.90

Numbers in parentheses are standard error estimates.

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

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 -0.0060, which was based on aggregate data subject to composition bias. The difference between the two estimates of β_3 is statistically significant at the 0.05 level.[7] As reported in detail in Solon, Barsky, and Parker [1992], the estimated procyclicality of the PSID wage statistic remains at least as great when we add time squared as a regressor, add leads and lags of the unemployment rate, exclude the Survey of Economic Opportunity portion of the PSID sample, measure $\ln W_t$ by the log of the sample mean wage instead of the sample mean of the log wage, or deflate by the CPI instead of the implicit GNP deflator. In addition, the second and third columns of Table II show the results from using growth in real GNP or in per capita hours as the cycle regressor. A comparison with the last two columns of Table I reveals again that the procyclicality of the PSID-based wage statistic is about double that of the BLS statistic.

These results suggest that, over the 1967–1987 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. We therefore proceed to analyze broader samples.

To do so, we 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

$$\begin{aligned}(7)\quad \ln W_{it} = \gamma_1 + \gamma_2 t + \gamma_3 t^2 + \gamma_4 (U_t - \delta_1 - \delta_2 t - \delta_3 t^2) \\ + \gamma_5' Z_i + \gamma_6 X_{it} + \gamma_7 X_{it}^2 + \gamma_8' Z_i X_{it} + \epsilon_{it},\end{aligned}$$

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. In the estimation, X_{it} is measured as age minus years of education minus 6. The second-order terms

7. As 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 0.0075 (the difference between -0.0060 and -0.0135), and its t-ratio is 2.22.

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

involving X_{it} allow year-to-year wage growth to depend on the level of X_{it} as well as on elements of Z_i. 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 $\gamma'_5 Z_i$ is impounded in the intercept, and the sample means of $\gamma_6 X_{it}$, $\gamma_7 X_{it}^2$, and $\gamma'_8 Z_i X_{it}$ 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 $\gamma'_5 Z_i + \gamma_6 X_{it} + \gamma_7 X_{it}^2 + \gamma'_8 Z_i X_{it}$ 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

$$\Delta \ln W_{it} = \beta_1 + \beta_2 t + \beta_3 \Delta U_t + \beta_4 X_{it} + \nu_{it}, \tag{8}$$

where $\nu_{it} = \gamma'_8 Z_i + \Delta\epsilon_{it}$, $\beta_2 = 2(\gamma_3 - \gamma_4 \delta_3)$ 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 $\beta_3 = \gamma_4 \gtreqless 0$ as the real wage is countercyclical, noncyclical, or procyclical. Equation (8) not only accounts for the wage growth effect of experience, but also controls implicitly for the first-order 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 analysis above, we avoided composition bias directly by restricting our sample so that it contained the exact same workers every year, but this resulted in a small and idiosyncratic sample. If we broaden the sample by relaxing the "balancedness" restriction, who is in the sample will change somewhat from year to year. The cyclical wage effects of these composition changes, however, will be accounted for by the differencing approach *provided* that the sample distribution of $\nu_{it} = \gamma'_8 Z_i + \Delta\epsilon_{it}$ is unrelated to business cycle conditions. This is precisely the way that Bils [1985] and many other researchers have addressed the composition bias issue with unbalanced longitudinal data. It is worth noting, though, that the broader samples afforded by "unbalancedness" come at a price. If the sample distribution of the first component of ν_{it} does vary cyclically, estimation of equation (8) with an unbalanced sample may be contaminated by composition bias. For example, suppose that,

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

during a recession year, the workers who exit the sample because of zero work hours are comprised disproportionately of workers with low $\gamma_8' Z_i$ (i.e., chronically slow wage growth). Then the estimation of β_3 will be subject to a countercyclical bias. Consequently, neither balanced nor unbalanced samples are clearly preferable, and it is worth viewing results from both.[8]

Equation (8) could be estimated by applying OLS to the unbalanced microdata on individuals, but as noted by Coleman [1986], the usual standard error estimators would be substantially biased. The problem is that ν_{it} is cross-sectionally correlated because different workers' error terms share common time effects. We therefore have followed a referee's advice to break the estimation into two steps. In the first step, we apply OLS to the individual-level regression of $\Delta \ln W_{it}$ on X_{it} and a vector of year dummies. The estimated coefficients of the twenty year dummies comprise a real wage time series that (subject to the caveat above) is free of composition bias. In the second step, we apply OLS to the regression of the estimated year coefficients on t and ΔU_t. The estimated coefficient of ΔU_t is an estimate of β_3 directly comparable to the estimates of β_3 from the BLS data and the balanced PSID data.

We broaden our men's sample by relaxing both the "balancedness" restriction and the age restriction. In the balanced sample, we required sample members to report positive labor income and at least 100 work hours in *every* year. Now, if an individual fails these requirements in some years, we still use his remaining observations for other years. For example, suppose that 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 eighteen observations, i.e., every year-to-year change from 1967–1968 to 1986–1987 except for 1982–1983 and 1983–1984. We also relax the 1928–1942 birth year restriction. Instead, we use all available person-year observations provided that the worker was at least age sixteen in that year. These changes dramatically increase our microdata sample size. The balanced sample contained 7100 observations of

8. If the second component of ν_{it} varies cyclically, estimates from both the balanced and unbalanced samples are subject to selection bias. As discussed in Bils [1985, pp. 676–77] and Solon and Barsky [1989, p. 20], such bias most likely would be in the direction of *under*estimating the procyclicality of real wages. If very low annual 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 procyclicality of wage opportunities.

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

year-to-year wage growth for 355 men (twenty observations per individual); the unbalanced sample contains 64,847 observations of 7225 men.

With the unbalanced sample, the second step of the estimation procedure produces the results in the fourth column of Table II. As expected, the increased sample size improves the precision in the estimation of β_3. The point estimate, however, changes only slightly from -0.0135 to -0.0140. The robustness of the result that $\hat{\beta}_3 \cong -0.014$ to the large changes in sample membership suggests that men's real wage cyclicality does not vary greatly with respect to age or labor force attachment.

As reported in detail in Solon, Barsky, and Parker [1992], we have checked directly for various types of heterogeneity in men's real wage cyclicality. In brief, like Stockman [1983], Bils [1985], and Keane and Prasad [1991], we find that estimated real wage cyclicality varies insignificantly with worker's years of education. Like Stockman, we find an insignificant contrast between union and nonunion workers. In addition, 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 who stay with the same employer. Using only observations of year-to-year wage growth for which it is clear that the worker did not change employers, we estimate β_3 at -0.0124 (with estimated standard error 0.0028). Thus, although employer-changers presumably experience even more wage procyclicality, the wage procyclicality experienced by stayers also is substantial.[9]

Our finding that β_3 is something on the order of -0.014 for men is quite robust to variations in sample selection criteria, but this estimate does not apply to women. The last two columns of Table II show, for balanced and unbalanced samples of women, results comparable to the men's results in the first and fourth columns. The 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) contains only 146 individuals. Estimation of equation (8) for this sample produces a $\hat{\beta}_3$ of -0.0046. This is noticeably less procyclical than the estimates

9. This result is somewhat different from Bils's [1985] results for young men in the National Longitudinal Surveys of labor market experience. His estimates of β_3 for stayers are -0.0064 for whites and -0.0044 for blacks. Shin's [1993] reanalysis of the NLS young men estimates the β_3 for stayers at -0.0095.

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

for men, but it is imprecisely measured with estimated standard error 0.0073. The unbalanced sample contains 50,531 observations of 6801 individuals. The larger sample reduces the estimated standard error of $\hat{\beta}_3$ to 0.0040, but changes $\hat{\beta}_3$ itself only slightly from -0.0046 to -0.0053. The *t*-ratio for the difference between this estimate and the corresponding men's estimate indicates that the discrepancy is statistically significant at the 0.05 level.[10]

What could account for the gender difference in wage cyclicality? One explanation might be that women's short-run labor supply is more elastic than men's. But then, if the cyclical shifts in labor demand are gender-neutral, women should experience greater employment variation over the cycle as well as lesser wage variation. Contrary to this prediction, results in Solon, Barsky, and Parker [1992] indicate that, over the 1967–1987 period, a one-percentage-point reduction in the overall unemployment rate is associated with a 1.8 percent increase in per capita hours of work for men and only a 1.4 percent increase for women. Explaining the gender difference in both price and quantity variation therefore requires that the cyclical shifts in labor demand are *not* gender-neutral. This seems quite plausible because cyclical employment variation is concentrated in durable goods manufacturing and construction, and women are underrepresented in both those industries. We believe further research along these lines would be worthwhile.

Whatever the reason for the gender difference in wage cyclicality, it poses a new difficulty for analyzing the discrepancy in wage cyclicality 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 section we incorporate this aggregation problem into an analysis of the quantitative importance of composition bias.

IV. The Importance of Composition Bias

In Section II we found that, over the period 1967–1968 to 1986–1987, the regression of real growth in the BLS aggregate wage statistic on time and change in the unemployment rate U

10. Similar suggestions of lesser wage cyclicality for women are scattered around the previous literature. See the literature survey in Solon, Barsky, and Parker [1992].

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

yielded an estimated $d\ln W/dU$ of -0.006. This was far less than the -0.014 estimate we subsequently obtained for men in the PSID, but a little more than our -0.005 estimate for the PSID women. We shall now use equation (5) from Section II to structure our analysis of how the cyclicality of the aggregate wage statistic is influenced by both composition bias and the aggregation of men and women. Then we shall verify the quantitative importance of composition bias by imposing on our PSID data the same sort of hours-weighting applied in the BLS wage statistics. Finally, we shall explain why the results in previous studies that dismissed the importance of composition bias are really consistent with our interpretation of the evidence.

Suppress the t subscripts in equation (5), and let the subscripts f and m, respectively, denote women and men. Thus, H_m is total work hours for men, S_f is the women's share of total work hours for the population, S_f^* is their share of the wage bill, W_m is men's average hourly earnings, and so forth. Let β_{3m} and β_{3f} denote the true wage cyclicalities for men and women. Finally, for each gender let δ denote the proportional gap between the wages paid for cyclically marginal hours of work and wages paid for nonmarginal hours. If, within each gender, the cyclically marginal hours are less skilled than the marginal hours, $\delta < 0$. Then equation (5) can be reexpressed[11] as

$$(9)\quad \frac{d\ln W}{dU} \cong (1 - S_f^*)\,\beta_{3m} + S_f^*\beta_{3f} + \left(\frac{W_f - W_m}{W}\right) S_f\,(1 - S_f) \times \left(\frac{d\ln H_f}{dU} - \frac{d\ln H_m}{dU}\right) + \delta\left[(1 - S_f^*)\,\frac{d\ln H_m}{dU} + S_f^*\frac{d\ln H_f}{dU}\right].$$

Equation (9) shows that the cyclicality of the BLS aggregate wage statistic is approximately a weighted average of the true cyclicalities of men's and women's wages plus two composition bias terms. The first of these reflects cyclical variation in the gender composition of total work hours. The second reflects cyclical variation in the skill composition of each gender's hours. Equation (9) is useful for estimating how procyclical the aggregate wage statistic would be if it were free of composition bias and for illuminating the sources of that bias.

To obtain a bias-free measure of aggregate wage cyclicality, we must estimate the weighted average of β_{3m} and β_{3f} shown as the

11. See Solon, Barsky, and Parker [1992] for the derivation.

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

first two terms of equation (9). Based on the PSID data, we have estimated β_{3m} at -0.014 and β_{3f} at -0.005. For 1977, the middle year of our sample period, Current Population Survey statistics by gender on employment and average work hours of the employed imply a female hours share of $S_f = 0.36$. Given O'Neill's [1985, Table 1] 1977 female-to-male wage ratio $W_f/W_m = 0.648$, we also have a proportional between-gender wage gap of $(W_f - W_m)/W = -0.40$ and a female wage-bill share of $S_f^* = 0.27$. The true aggregate real wage cyclicality therefore is estimated to be $0.73\ (-0.014) + 0.27\ (-0.005) = -0.0116$.

Although our PSID estimates of β_{3m} and β_{3f} imply a true $d\ln W/dU$ of -0.0116, the BLS aggregate wage statistic analyzed in Section II displayed a $d\ln W/dU$ of only -0.0060. The implied countercyclical composition bias in the BLS statistic is 0.0056. 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 who was at least age sixteen 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.[12]

When we regress the change in the log of this PSID-based wage statistic on time and change in the unemployment rate, we obtain $\hat{\beta}_3 = -0.0057$ (with estimated standard error 0.0025). This estimate is remarkably close to the -0.0060 estimate for the BLS statistic, and it diverges even a little farther from the more procyclical estimates we got when we used the PSID data to avoid composition bias, not inject it. Addressing the composition bias issue therefore appears to be very important for the proper measurement of real wage cyclicality.

Equation (9) can be used further to decompose the 0.0056 estimate of the composition bias into its two components: the term

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

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

involving the cyclicality of the female share of hours and the term involving cyclical variation in the skill composition of each gender's hours. The first of these is estimated by substituting in $(W_f - W_m)/W = -0.40$, $S_f = 0.36$, $d\ln H_f/dU = -0.0014$, and $d\ln H_m/dU = -0.0018$. Given that women are paid less than men and have less cyclically variable hours, the gender composition term by itself imposes a *pro*cyclical bias, which is estimated at -0.0004. Because the gender difference in hours cyclicality 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 0.0060. With $S_f^* = 0.27$, the implied value of δ is -0.35. 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 cyclicality. Unlike those studies, our estimate provides an omnibus measure of the wage gap between cyclically marginal and nonmarginal hours that is not restricted to particular observable characteristics of the workers. This gap is so wide that it causes a major countercyclical bias in aggregate wage statistics.

The hypothesis that aggregate wage statistics are subject to a substantial composition bias was first stated by Stockman back in 1983. Since then, numerous researchers have used longitudinal data from the PSID and the National Longitudinal Surveys of labor market experience to estimate the cyclicality of real wages. As reviewed in Section IV of Solon, Barsky, and Parker [1992], most of the estimates in the longitudinal literature are remarkably similar to ours. For example, Bils [1985] estimated β_3 at -0.016 for white men and -0.018 for black men; Coleman [1984] estimated β_3 at -0.015 for all men; and most estimates for women are at least as procyclical as 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 cyclicality 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. The reason they have not had to revise their beliefs is that many of the longitudinal studies since Stockman's have devoted little or no attention to composition bias

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

in the aggregate data, and those that have discussed it have downplayed the empirical importance of Stockman's hypothesis.[13] Bils [1985, p. 668] discussed composition bias, but concluded that "the impact is not particularly large."[14] 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 microdata, he obtained $\hat{\beta}_3 = -0.012$. This unusually procyclical estimate from aggregate data is idiosyncratic to the particular ten years in Bils's sample,[15] but it is still considerably less procyclical than his microdata estimates. Coleman did not present aggregate results for the 1968–1979 period covered by his microdata, but his estimate of β_3 based on the BLS aggregate wage data for 1961–1979 is -0.003, drastically smaller than his microdata estimate of -0.015.

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 unemploy-

13. One partial exception is Blank [1990]. In the second row of her Table 1, she reported results from an hours-weighting exercise like ours. Although she did not emphasize these results, they accord with our finding that creating a composition bias in the PSID data leads to dramatically smaller estimates of real wage cyclicality. A different approach has been pursued by Kydland and Prescott [1988], who used PSID data 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.

14. Also see Bils's [1985] discussion on p. 684.

15. See footnote 4 in Solon, Barsky, and Parker [1992].

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

ment 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.

In summary, evidence from longitudinal microdata shows that real wages have been considerably more procyclical in recent years than indicated by aggregate time series data for the same period. Our analysis reveals that 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 1960s, it is impossible to extend our analysis to earlier periods. Nevertheless, the literature that documents the greater hours cyclicality 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, the importance of composition bias implies 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 cyclicality in other countries as well as in the United States, the aggregate wage statistics of other countries may be similarly misleading.

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 Panel Study of Income Dynamics and the National Longitudinal Surveys of labor market experience indicates that 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

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

example, we noted in Section II that, given the modest cyclicality 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 procyclicality evident in the longitudinal data, however, implies a smaller labor supply elasticity. Estimating the regression of real wage growth as measured for our unbalanced sample of PSID men on time and growth in men's per capita work hours yields an estimated inverse labor supply elasticity of 0.70. As explained in Section II, the inverted value 1.4 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 0.97.

This range of about 1.0 to 1.4 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.[16] Nevertheless, our finding that real wages and work hours are quite positively correlated in the aggregate need not imply that the two variables are connected by notional labor supply behavior. For one thing, the 1.0-to-1.4 range for the labor supply elasticity implies an intertemporal substitution elasticity in the same range only if cyclical wage innovations are purely transitory. If they are quite persistent, as suggested by Altonji and Ashenfelter [1980] among others, rationalizing 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 microeconometric 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 unity.

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

16. Performing the same exercise for women turns out not to be very useful. It generates a range from 0.3 to 3.0.

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

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.

Appendix: Data Sources

Aggregate Data

Average hourly earnings of production or nonsupervisory workers in private nonagricultural employment: *Handbook of Labor Statistics,* 1975, Table 98, for 1947–1963; *Employment and Earnings,* March 1993, Table C-1, for 1964–1992.

Civilian unemployment rate, employment/population ratio, and employment by gender: *Economic Report of the President,* 1993, Tables B-30 and B-31.

Average work hours of the employed: January issues of *Employment and Earnings.*

Real GNP and implicit GNP deflator: Through 1990, real GNP (1982 dollars) is from Table B-2 of *Economic Report of the President,* 1991, and the implicit GNP deflator (1982 = 100) is from Table B-3. The 1990–1991 and 1991–1992 growth rates for each are based on the 1990, 1991, and 1992 values of real GNP (1987 dollars) and the implicit GNP deflator (1987 = 100) from Tables 1.10 and 7.3 in the Selected NIPA Tables section of *Survey of Current Business.* The 1990 figures are from the December 1992 issue, and the 1991 and 1992 figures are from the March 1993 issue.

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, and the responses concerning annual labor income and hours of work pertain to the preceding calendar years 1967–1987. Observations with "major assignments" imputed for labor income or work hours are excluded from our sample.

Our hourly wage measure is the ratio of total annual labor income to total annual hours of work. Like the BLS average hourly

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

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.[17]

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.

UNIVERSITY OF MICHIGAN
UNIVERSITY OF MICHIGAN
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

REFERENCES

Abel, Andrew B., and Ben S. Bernanke, *Macroeconomics* (Reading, MA: Addison-Wesley, 1992).
Altonji, Joseph G., "Intertemporal Substitution in Labor Supply: Evidence from Micro Data," *Journal of Political Economy,* XCIV (1986), S176–215.
Altonji, Joseph, and Orley Ashenfelter, "Wage Movements and the Labour Market Equilibrium Hypothesis," *Economica,* XLVII (1980), 217–45.
Ashenfelter, Orley, "Macroeconomic Analyses and Microeconomic Analyses of Labor Supply," *Carnegie-Rochester Conference Series on Public Policy,* XXI (1984), 117–55.

17. PSID 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.

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

Ball, Laurence, "Intertemporal Substitution and Constraints on Labor Supply: Evidence from Panel Data," *Economic Inquiry,* XXVIII (1990), 706–24.
Barro, Robert J., and Robert G. King, "Time-Separable Preferences and Intertemporal-Substitution Models of Business Cycles," *Quarterly Journal of Economics,* XCIX (1984), 817–39.
Bils, Mark J., "Real Wages over the Business Cycle: Evidence from Panel Data," *Journal of Political Economy,* XCIII (1985), 666–89.
Blanchard, Olivier Jean, and Stanley Fischer, *Lectures on Macroeconomics* (Cambridge, MA: MIT Press, 1989).
Blank, Rebecca M., "Why Are Wages Cyclical in the 1970s?" *Journal of Labor Economics,* X (1990), 219–57.
Christiano, Lawrence J., and Martin Eichenbaum, "Current Real-Business-Cycle Theories and Aggregate Labor-Market Fluctuations," *American Economic Review,* LXXXII (1992), 430–50.
Clark, Kim B., and Lawrence H. Summers, "Demographic Differences in Cyclical Employment Variation," *Journal of Human Resources,* XVI (1981), 61–79.
Coleman, Thomas S., "Essays in Aggregate Labor Market Business Cycle Fluctuations," Ph.D. thesis, University of Chicago, 1984.
——, "A Note on the Measurement of Aggregate Effects in Panel Data," unpublished, 1986.
Fischer, Stanley, "Recent Developments in Macroeconomics," *Economic Journal,* XCVIII (1988), 294–339.
Friedman, Milton, "The Role of Monetary Policy," *American Economic Review,* LVIII (1968), 1–17.
Greenwald, Bruce D., and Joseph E. Stiglitz, "Examining Alternative Macroeconomic Theories," *Brookings Papers on Economic Activity,* No. 1 (1988), 207–60.
Hall, Robert E., "Examining Alternative Macroeconomic Theories: Comment," *Brookings Papers on Economic Activity,* No. 1 (1988), 261–64.
Hall, Robert E., and John B. Taylor, *Macroeconomics: Theory, Performance, and Policy,* 3rd ed. (New York, NY: W. W. Norton, 1991).
Ham, John C., "Testing Whether Unemployment Represents Intertemporal Labour Supply Behaviour," *Review of Economic Studies,* LIII (1986), 559–78.
Keane, Michael J., and Eswar S. Prasad, "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, 1991.
Keynes, John Maynard, *The General Theory of Employment, Interest, and Money* (London: Macmillan, 1936).
Kniesner, Thomas J., and Arthur H. Goldsmith, "A Survey of Alternative Models of the Aggregate U. S. Labor Market," *Journal of Economic Literature,* XXV (1987), 1241–80.
Kosters, Marvin, and Finis Welch, "The Effects of Minimum Wages on the Distribution of Changes in Aggregate Employment," *American Economic Review,* LXII (1972), 323–32.
Kydland, Finn E., "Labor-Force Heterogeneity and the Business Cycle," *Carnegie-Rochester Conference Series on Public Policy,* XXI (1984), 173–208.
Kydland, Finn E., and Edward C. Prescott, "Cyclical Movements of the Labor Input and Its Real Wage," Working Paper No. 413, Research Department, Federal Reserve Bank of Minneapolis, 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,* LXIII (1981), 319–27.
Lucas, Robert E., Jr., "Understanding Business Cycles," *Carnegie-Rochester Conference Series on Public Policy,* V (1977), 7–29.
Lucas, Robert E., Jr., and Leonard A. Rapping, "Real Wages, Employment, and Inflation," *Journal of Political Economy,* LXXVII (1969), 721–54.
MaCurdy, Thomas E., "An Empirical Model of Labor Supply in a Life-Cycle Setting," *Journal of Political Economy,* LXXXIX (1981), 1059–85.
Mankiw, N. Gregory, "Real Business Cycles: A New Keynesian Perspective," *Journal of Economic Perspectives,* III (1989), 79–90.
Mitchell, Mark L., Myles S. Wallace, and John T. Warner, "Real Wages over the Business Cycle: Some Further Evidence," *Southern Economic Journal,* LI (1985), 1162–73.

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions

Okun, Arthur M., "Upward Mobility in a High-Pressure Economy," *Brookings Papers on Economic Activity,* No. 1 (1973), 207–52.

O'Neill, June, "The Trend in the Male-Female Wage Gap in the United States," *Journal of Labor Economics,* III (1985), S91–116.

Phelps, Edmund S., "Introduction: The New Microeconomics in Employment and Inflation Theory," *Microeconomic Foundations of Employment and Inflation Theory,* Edmund S. Phelps, ed. (New York, NY: W. W. Norton, 1970).

Prescott, Edward C., "Response to a Skeptic," *Federal Reserve Bank of Minneapolis Quarterly Review,* X (1986), 28–33.

Rotemberg, Julio J., and Michael Woodford, "Markups and the Business Cycle," *NBER Macroeconomics Annual 1991,* Olivier Jean Blanchard and Stanley Fischer, eds. (Cambridge, MA: MIT Press, 1991).

Shapiro, Carl, and Joseph E. Stiglitz, "Equilibrium Unemployment as a Worker Discipline Device," *American Economic Review,* LXXIV (1984), 433–44.

Shin, Donggyun, "Cyclicality of Real Wages among the NLS Young Men," unpublished, 1993.

Solon, Gary, and Robert Barsky, "Real Wages over the Business Cycle," NBER Working Paper No. 2888, 1989.

Solon, Gary, Robert Barsky, and Jonathan A. Parker, "Measuring the Cyclicality of Real Wages: How Important Is Composition Bias?" NBER Working Paper No. 4202, 1992.

Stockman, Alan C., "Aggregation Bias and the Cyclical Behavior of Real Wages," unpublished, 1983.

Survey Research Center, *A Panel Study of Income Dynamics: Procedures and Tape Codes, 1988 Interviewing Year—Wave XXI, A Supplement* (Ann Arbor, MI: Institute for Social Research, University of Michigan, 1991).

This content downloaded from 147.251.185.127 on Wed, 18 Mar 2015 11:09:26 UTC
All use subject to JSTOR Terms and Conditions