The Cyclical Behavior of the Price-Cost Markup

Christopher J. Nekarda
Federal Reserve Board of Governors

Valerie A. Ramey
University of California, San Diego
and NBER

First version: July 2009
This version: March 2010

Abstract

Countercyclical markups constitute the key transmission mechanism for monetary and other "demand" shocks in New Keynesian models. This paper tests the foundation of those models by studying the cyclical properties of the markup of price over marginal cost. The first part of the paper studies markups in the aggregate economy and the manufacturing sector. We use Bils's (1987) insights for converting average cost to marginal cost, but do so with richer data. We find that all measures of markups are either procyclical or acyclical. Moreover, we show that monetary shocks lead markups to fall with output. The second part of the paper merges input-output information on shipments to the government with detailed industry data to study the effect of demand changes on industry-level markups. Industry-level markups are found to be acyclical in response to demand changes.

JEL codes: E32, L16, J31
Keywords: markup, cyclicality

The views in this paper are those of the authors and do not necessarily represent the views or policies of the Board of Governors of the Federal Reserve System or its staff. We are grateful to Mark Bils, Olivier Coibion, Davide Debortoli, Graham Elliott, Neville Francis, James Hamilton, David Lebow, Robert Hall, Min Ouyang, Garey Ramey, and Gary Richardson for helpful comments. Ben Backes provided excellent research assistance. Valerie Ramey gratefully acknowledges financial support from National Science Foundation grant SES-0617219 through the National Bureau of Economic Research.

> How markups move, in response to what, and why, is however nearly terra incognita for macro. ...[W]e are a long way from having either a clear picture or convincing theories, and this is clearly an area where research is urgently needed.
>
> — Blanchard (2008), p. 18.

1 Introduction

The markup of price over marginal cost plays a key role in a number of macroeconomic models. For example, in Rotemberg and Woodford's (1992) model, an increase in government spending leads to increases in both hours and real wages because imperfect competition generates a countercyclical markup. In the standard New Keynesian model, sticky prices combined with procyclical marginal cost imply that an expansionary monetary shock or government spending shock lowers the average markup (Goodfriend and King, 1997). This result also holds in the leading New Keynesian models with both sticky prices and sticky wages, such as Erceg et al. (2000); Smets and Wouters (2003, 2007); Christiano et al. (2005). In Jaimovich and Floetotto's (2008) model, procyclical entry of firms leads to countercyclical markups, and hence to procyclical movements in measured productivity.

Estimating the cyclicality of markups, however, is one of the more challenging measurement issues in macroeconomics. Theory directs comparing price and marginal cost; however, available data typically include only average cost. Papers studying the cyclicality of marginal cost and markups either accept average cost as is (e.g., Domowitz et al., 1986) or make assumptions on how marginal cost is related to average cost (e.g., Bils, 1987; Rotemberg and Woodford, 1991, 1999; Galí et al., 2007). Using measures of price-average cost margins, Domowitz et al. (1986) find that markups are positively correlated with the growth of industry demand, suggesting that markups are procyclical. Bils (1987) estimates marginal cost in manufacturing under several assumptions about overtime and adjustment costs and concludes that markups there are countercyclical. Rotemberg and Woodford (1991, 1999) study the economy more broadly and present several mitigating reasons why seemingly procyclical patterns in measures of the average markup should be discounted.

In this paper, we present evidence that most measures of the markup are procyclical. Moreover, they increase in response to positive monetary and government spend-

ing shocks. The first part of the paper presents the evidence using aggregate data. Markups based on average wages are procyclical, hitting troughs during recessions and reaching peaks in the middle of expansions. Because of concerns that average wages do not adequately capture marginal costs, we use insights from Bils (1987) to make adjustments to convert average wages into marginal costs. In contrast to Bils, we find that all measures of the markup remain procyclical or acyclical even after adjustment. We trace the source of the difference in part to Bils's study of the manufacturing sector and in part to his use of annual data. We also consider generalizations of the standard Cobb-Douglas production function and find little effect. We then consider the response of our various markup measures to monetary policy shocks. We find that a contractionary monetary policy shock leads to a fall in both output and the markup. This result raises questions about the basic propagation mechanism of the New Keynesian model: If the markup does not move countercyclically, how can money have short-run real effects?

In the second part of the paper we analyze the markup in a panel of industries. Building on the ideas of Shea (1993) and Perotti (2008), we match detailed input-output (IO) data on government demand and its downstream linkages with data on employment, hours, and output. We argue that the government demand variable we construct is an excellent instrument for determining the effects of shifts in demand on markups. We find that an increase in output associated with higher government spending has essentially no effect on markups at an annual frequency.

2 Theoretical Framework

In this section we review the theory that guides our empirical investigation. We first derive the marginal cost of increasing output by raising average hours per worker. We then derive an expression for converting data on average wages to marginal wages. Finally, we consider a more general production function.

The theoretical markup, M, is defined as

$$M = \frac{P}{MC}, \tag{1}$$

where P is the price of output and MC is the nominal marginal cost of increasing output. The inverse of the right hand side of equation 1, MC/P, is also known as the

real marginal cost.

A cost-minimizing firm should equalize the marginal cost of increasing output across all possible margins for varying production. Thus, it is valid to consider the marginal cost of varying output by changing a particular input. As in Bils (1987) and Rotemberg and Woodford (1999), we focus on the labor input margin, and in particular on hours per worker. We assume that there are potential costs of adjusting the number of employees and the capital stock, but no costs of adjusting hours per worker.[1]

Focusing on the static aspect of this cost-minimization problem, consider the problem of a firm that chooses hours per worker, h, to minimize

$$\text{Cost} = W_{\text{A}} hN + \text{other terms not involving } h, \tag{2}$$

subject to $\bar{Y} = F(AhN, \ldots)$. W_{A} is the average hourly wage, N is the number of workers, Y is output, and A is the level of labor-augmenting technology. Bils (1987) argues that the average wage paid by a firm may be increasing in the average hours per worker because of the additional cost of overtime hours. We capture this assumption by specifying the average wage as:

$$W_{\text{A}} = W_{\text{S}} \left(1 + \rho\theta \frac{v}{h}\right). \tag{3}$$

where W_{S} is the straight-time wage, ρ is the premium for overtime hours, θ is the fraction of overtime hours that command a premium, and v/h is the ratio of average overtime hours to total hours. The term $\rho\theta v/h$ captures the idea that firms may have to pay a premium for hours worked beyond the standard workweek.[2] Bils did not include the θ term in his specification because he used manufacturing data from the Current Employment Statistics (CES), in which overtime hours are *defined* as those hours commanding a premium, where $\theta = 1$. For the entire economy, we define overtime hours as those hours in excess of 40 hours per week. Because overtime premium regulations do not apply to all workers, we must allow for the possibility that θ is less than unity and my vary over time.

1. Hamermesh and Pfann's (1996) summary of the literature suggests that adjustment costs on the number of employees are relatively small and that adjustment costs on hours per worker are essentially zero.

2. It would also be possible to distinguish wages paid for part-time work versus full-time work. However, Hirsch (2005) finds that nearly all of the difference in hourly wages between part-time and full-time workers can be attributed to worker heterogeneity rather than to a premium for full-time work.

We assume that the firm takes the straight-time wage, the overtime premium, and the fraction of workers receiving premium pay as given, but recognizes the potential effect of raising h on overtime hours v. Letting λ be the Lagrange multiplier on the constraint, we obtain the first-order condition for h as:

$$W_{\mathrm{S}}\left[1+\rho\theta\left(\frac{\mathrm{d}v}{\mathrm{d}h}\right)\right]=\lambda F_1(AhN,\ldots)A, \tag{4}$$

where $\mathrm{d}v/\mathrm{d}h$ is the amount by which average overtime hours rise for a given increase in average total hours and F_1 is the derivative of the production function with respect to effective labor, AhN. The multiplier λ is equal to marginal cost, so the marginal cost of increasing output by raising hours per worker is given by:

$$MC=\lambda=\frac{W_{\mathrm{S}}\left[1+\rho\theta\left(\frac{\mathrm{d}v}{\mathrm{d}h}\right)\right]}{AF_1(AhN,\ldots)}. \tag{5}$$

The denominator of equation 5 is the marginal product of increasing hours per worker; the numerator is the marginal cost of increasing average hours per worker. As discussed above, this marginal cost should also be equal to the marginal cost of raising output by increasing employment or the capital stock. If there are adjustment costs involved in changing those factors, the marginal cost would include an adjustment cost component. Focusing on the hours margin obviates the need to estimate adjustment costs.

Equation 5 makes it clear that the marginal cost of increasing hours per worker is not equal to the average wage, as is commonly assumed. Following Bils (1987), we call the term in the numerator the "marginal wage" and denote it by W_{M}:

$$W_{\mathrm{M}}=W_{\mathrm{S}}\left[1+\rho\theta\left(\frac{\mathrm{d}v}{\mathrm{d}h}\right)\right]. \tag{6}$$

To the extent that the marginal wage has different cyclical properties from the average wage, markup measures that use the average wage may embed cyclical biases. Bils (1987) used approximation to the marginal wage itself to substitute for marginal cost in his markup measure. We instead use an adjustment that does not require approximation. In particular, we combine the expressions for the average wage and the marginal

wage to obtain their ratio:

$$\frac{W_{\mathrm{M}}}{W_{\mathrm{A}}} = \frac{1+\rho\theta\left(\frac{\mathrm{d}v}{\mathrm{d}h}\right)}{1+\rho\theta\left(\frac{v}{h}\right)}. \tag{7}$$

This ratio can be used to convert the observed average wage to the theoretically-correct marginal wage required to estimate the markup. We show in section 3 that the ratio of overtime hours to average hours, v/h, is procyclical, and that θ is roughly constant. Thus, the denominator in equation 7 is procyclical. How $W_{\mathrm{M}}/W_{\mathrm{A}}$ evolves over the business cycle depends on the relative cyclicality of $\mathrm{d}v/\mathrm{d}h$. The fact that v/h increases with h does *not* imply that $\mathrm{d}v/\mathrm{d}h$ increases with h. It can be shown that for a constant θ, $\mathrm{d}^2v/\mathrm{d}h^2 > 0$ is a necessary, but not sufficient, condition for the wage ratio to be increasing in h. Thus, it is possible for v/h to be procyclical, but $W_{\mathrm{M}}/W_{\mathrm{A}}$ to be countercyclical.

A second complication with estimating the markup is estimating the marginal product of labor. If the production function is Cobb-Douglas, then the marginal product of labor is proportional to the average product. Consider a more general case in which the production function has constant elasticity of substitution (CES):

$$Y = \left[\alpha(AhN)^{\frac{\sigma-1}{\sigma}} + (1-\alpha)K^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}, \tag{8}$$

where σ is the elasticity of substitution between capital and labor. The derivative with respect to effective labor (the F_1 needed for equation 5) is

$$\frac{\partial Y}{\partial(AhN)} = \alpha\left(\frac{Y}{AhN}\right)^{\frac{1}{\sigma}}. \tag{9}$$

The exponent in equation 9 is the reciprocal of the elasticity of substitution. If the elasticity of substitution is unity, this specializes to the Cobb-Douglas case. On the other hand, if the elasticity of substitution is less than unity, then the exponent will be greater than unity.

In the simple case where the marginal wage is equal to the average wage and the marginal product of labor is proportional to the average product (as in the Cobb-

Douglas case), then the markup is given by

$$M_{\mathrm{A}}^{\mathrm{CD}} = \frac{P}{W_{\mathrm{A}}/\left[\alpha\left(Y/hN\right)\right]} = \frac{\alpha}{s}, \tag{10}$$

where α is the exponent on labor input in the production function and s is the labor share. In the case where the wage is increasing in average hours, the markup can be written as

$$M_{\mathrm{M}}^{\mathrm{CD}} = \frac{P}{W_{\mathrm{M}}/\left[\alpha\left(Y/hN\right)\right]} = \frac{\alpha}{s\left(W_{\mathrm{M}}/W_{\mathrm{A}}\right)}, \tag{11}$$

where we use equation 7 to convert average wages to marginal wages. Finally, allowing for a CES production function, we obtain the markup

$$\begin{aligned} M_{\mathrm{M}}^{\mathrm{CES}} &= \frac{P}{W_{\mathrm{M}}/\left[A\alpha\left(Y/AhN\right)^{\frac{1}{\sigma}}\right]} \\ &= \frac{\alpha}{s\left(W_{\mathrm{M}}/W_{\mathrm{A}}\right)}\left(\frac{Y}{AhN}\right)^{\frac{1}{\sigma}-1}. \end{aligned} \tag{12}$$

One important issue is that Y should be *gross* output. As Basu and Fernald (1997) argue, "value added is not a natural measure of output and can in general be interpreted as such only with perfect competition."[3] Because any nonconstant, nonunity markup requires imperfect competition, markup measures should be based on gross output, not value added. Unfortunately, no measures of gross output are available for the broad aggregates studied in section 4 so we must use value added. We use gross output when studying the markup at the industry level in section 7.

3 Estimating the Marginal-Average Wage Adjustment

This section describes the estimation of the four components of the marginal-average wage adjustment factor. For expositional purposes, we repeat equation 7:

$$\frac{W_{\mathrm{M}}}{W_{\mathrm{A}}} = \frac{1+\rho\theta\left(\frac{\mathrm{d}\nu}{\mathrm{d}h}\right)}{1+\rho\theta\left(\frac{\nu}{h}\right)}. \tag{7}$$

3. Basu and Fernald (1997), p. 251.

To construct the ratio of marginal to average wages, we require (1) estimates of the ratio of overtime hours to average hours, v/h; (2) estimates of the marginal change in overtime hours with respect to a change in average total hours, $\mathrm{d}v/\mathrm{d}h$; (3) the fraction of overtime hours that command a premium, θ; and (4) the premium for overtime hours, ρ.

3.1 Measuring v/h

Most researchers analyzing the cyclical behavior of overtime hours have focused on production workers in manufacturing, since this group is the only one for whom data on overtime hours are readily available. The manufacturing sector is not representative of the entire U.S. economy, however. Even at its post-World War II peak, manufacturing accounted for only 26 percent of GDP and one-third of nonfarm employment; it now accounts for only 11 percent of output and 9 percent of employment.

In order to obtain results that are more representative, we exploit an overlooked data source that we use to construct a time series on the distribution of hours worked in the entire civilian economy. The Bureau of Labor Statistics (BLS)'s *Employment and Earnings* publication provides information on persons at work by hours of work, total persons at work, and average hours worked by persons at work, derived from the Current Population Survey (CPS). These data are available monthly beginning in May 1960. The number of persons at work are available only within particular ranges of hours worked, such as 35–39 hours per week, 40 hours per week, 41–48 hours per week, etc. To approximate the distribution of hours worked, we use data from the CPS to calculate a time-varying average of actual hours worked for each published range. We seasonally adjust the monthly data and then take the quarterly average to smooth the considerable month-to-month volatility. Appendix A contains a complete list of data sources and additional details of our methodology.

We also study the manufacturing sector using data from the CES database. The BLS publishes monthly data on average weekly hours and average weekly overtime hours of production and nonsupervisory workers in manufacturing. There are two key differences between the CES data and those for all workers that we construct from the CPS. First, the CES data cover only production and nonsupervisory workers, whereas the CPS data cover all civilian workers. Second, the CES defines overtime hours as any hours that are paid a premium, whereas we consider any hours worked above 40 hours

per week to be overtime hours. As with the CPS data on all workers, we use seasonally-adjusted series and take quarterly averages of monthly data. Using quarterly averages is particularly important in the CES data, where the series are only published to one decimal place. At this precision, a substantial amount of the change from one month to the next may reflect measurement error; the change in the quarterly average is less subject to this bias.

Figure 1 shows three hours series for all workers in the aggregate economy. The top panel plots average weekly hours per worker, the middle panel plots average weekly hours in excess of 40 ("overtime hours"), and the bottom panel plots the fraction of total hours worked that are overtime hours. All three series show pronounced low-frequency movements, decreasing throughout the 1960s and 1970s and rising over the 1980s and 1990s. They also appear to be weakly procyclical, although the low-frequency movements dominate. Figure 2 shows the same hours series for production and nonsupervisory workers in manufacturing. All three series are noticeably procyclical. In particular, a procyclical fraction of overtime hours implies that average overtime hours varies more than average hours.

The series shown in the bottom panels of figures 1 and 2 are the first element needed to estimate the marginal-average wage adjustment factor. How overtime hours vary with changes in average hours is an important determinant of the cyclicality of marginal cost.

3.2 Estimating $\mathrm{d}v/\mathrm{d}h$

To estimate the second element, the marginal change in overtime hours, we consider a difference approximation such as

$$\Delta v_t = \alpha + \eta_t \Delta h_t + \xi_t.$$

Such a specification relates the marginal change in overtime hours resulting from an increase in average hours, whatever the source of the increase in average hours. The error term, ξ_t, captures other sources of variations in overtime hours.

Bils (1987) specifies η_t as a polynomial function of average hours and time trends. We explore similar parametric specifications and find that a combination of a parametric and nonparametric, time-varying coefficient captures both the time variation from average hours and the time trend.

We estimate η_t as a combination of an unobserved, time-varying component, β_t, and a parametric component, $\gamma(h_{t-1}-40)$. We estimate β_t using the state-space model:

$$\Delta v_t = \alpha + \beta_t \Delta h_t + \gamma \Delta h_t (h_{t-1} - 40) + \xi_t \tag{13}$$

$$\beta_t = \beta_{t-1} + \psi_t, \tag{14}$$

where equation 14 is the transition equation for the unobserved state β_t and disturbances ξ_t and ψ_t are mutually-uncorrelated normally-distributed errors with mean zero and variances σ^2_ξ and σ^2_ψ. We estimate the model by maximum likelihood and use the Kalman smoother to obtain our final estimates of β_t.

The solid line in figure 3 plots the estimate of η_t when β_t is allowed to vary over time and the dashed line shows the estimate of η_t where β_t is constant. The coefficient for all workers (top panel) exhibits little variation over time. The variance of ψ that maximizes the likelihood function, expressed as a ratio of the variance of ξ, is 0.036.[4] We use the time-varying estimate, β_t, although our results are similar if we hold β_t constant, as indicated by the similar paths of the two estimates of η_t.

The estimate of η for manufacturing production workers (lower panel) varies much more than for all workers. It displays both low frequency movements and significant procyclicality. The relative variance of ψ in manufacturing is 0.167, considerably greater than in the aggregate data. In manufacturing, allowing β to vary over time captures important changes in the relationship between dv and dh, conditional on lagged average hours, as evidenced by the difference between the solid and dashed line.

3.3 Estimating θ

We now turn to the estimation of θ, which is the fraction of hours worked over 40 hours per week that command a premium. For the CES data on production workers in manufacturing, $\theta = 1$ by definition, since overtime hours are defined as those hours that are paid an premium.[5]

Estimating θ is more difficult for the aggregate data. The only direct information is from the May supplements to the CPS in 1969–81, which asked workers whether they

4. Stock and Watson's (1998) median unbiased estimator also yields a relative variance of essentially 0.

5. See, however, Deleire et al. (2002) for evidence that the establishment survey overcounts overtime hours.

received higher pay for hours over 40 hours per week.[6] Using this question, along with information on hours worked, we construct a measure of the percent of hours worked over 40 that are paid an overtime premium.[7]

Unfortunately, the key question on premium pay was dropped from the May supplement after 1985. A potential alternative source of information is the BLS's Employee Costs for Employee Compensation (ECEC) survey which provides information on total compensation, straight time wages and salaries, and various benefits, such as overtime pay, annually from 1991 to 2001 and quarterly from 2002 to the present. If one assumes a particular statutory overtime premium, then one can construct an estimate of θ from these data. We assume that the statutory premium is 50 percent and construct a θ accordingly.

Figure 4 shows annual estimates of θ based on these two sources. From 1969 to 1981, θ averages 0.33, meaning that only one-third of hours over 40 command a premium. From 1991 to 2009, θ averages 0.27. Although it appears that the estimate of θ from the CPS falls during recessions, regressing θ on average hours does not yield a significant relationship.[8] On the other hand, the fraction of hours paid a premium is slightly countercyclical in the ECEC data.[9] It is difficult to tell whether the structure of the economy actually changed or whether the two surveys are simply not comparable. Because there is little cyclical variation in θ in either survey, we assume that θ is a constant equal to the average across the two surveys of 0.3.[10]

3.4 Value of ρ

The last element needed to construct the marginal-average wage conversion factor is the premium paid for overtime hours. The Fair Labor Standards Act requires that employers pay a 50 percent premium for hours in excess of 40 per week for covered employees. Evidence from Carr (1986) indicates that in 1985, 92 percent of those who earned premium pay received a 50 percent premium.[11]

6. There were some anomalies with this question in the 1985 survey, so we did not use it.

7. See appendix A.

8. The coefficient from this regression is 0.02 and has a t statistic of 1.40.

9. Regressing θ estimated from the ECEC on CPS average hours yields a coefficient of -0.03 with a t statistic of -3.2.

10. If we instead assume that θ is procyclical with the coefficient of 0.024 on average hours, our estimates of the marginal-average wage factor are still either countercyclical or acyclical.

11. See Wetzel (1966) and Taylor and Sekscenski (1982) for other estimates.

Trejo (1991) has questioned whether the true cost of an extra overtime hour for those covered is actually 50 percent. He shows that the implicit cost of overtime hours is lower than 50 percent because straight-time wages are lower in industries that offer more overtime. Hamermesh (2006) updates his analysis and finds supporting results: the implicit overtime premium is 25 percent, not 50 percent. The following theory is consistent with these results. Workers and firms bargain over a job package that involves hours and total compensation. If workers' marginal disutility of overtime hours is less than the statutory premium for overtime hours, then workers and firms in industries with higher average overtime hours will adjust to overtime regulations by agreeing to a lower straight-time wage. This implicit contract means that when the firm pays the worker a 50 percent premium for an overtime hour, part of that premium compensates the worker for the true marginal disutility of overtime, but part is simply a payment on a debt incurred to the worker because the contracted straight-time wage is lower.

In light of these results, we construct our markups under two alternative assumptions. First, we assume that the effective marginal overtime premium is equal to the statutory premium, so that the ρ in both the numerator and denominator of equation 7 are equal to 0.50. Second, we assume that the effective marginal overtime premium is equal to Hamermesh's (2006) estimate of 0.25. This means that the ρ in the numerator of equation 7 is equal to 0.25; however, since the average wage data includes the 50 percent statutory premium, the ρ in the denominator remains at 0.50.

3.5 Marginal-Average Wage Adjustment Factor

Figure 5 plots our wage adjustment factors, $W_{\mathrm{M}}/W_{\mathrm{A}}$ for both the aggregate economy and manufacturing with a 25-percent and a 50-percent effective marginal overtime premium. In the aggregate economy (upper panel), the adjustment factor is above unity for both values of the overtime premium, indicating that the marginal wage is greater than the average wage, and increases over most of the sample. However, the marginal-average wage adjustment factor is countercyclical. The correlation of the cyclical components of the log of real GDP and the log of $W_{\mathrm{M}}/W_{\mathrm{A}}$ is -0.7 for both values of the overtime premium.[12] The adjustment factor is countercyclical because the numerator varies little over the business cycle whereas the denominator is procyclical. In manufacturing (lower panel) the adjustment factor is procyclical. The correlation of

12. Throughout the paper, we extract cyclical components using a Hodrick-Prescott (HP) filter with a standard smoothing parameter of $1,600$.

the cyclical component of the adjustment factor in manufacturing with real GDP is 0.5 for the 25 percent premium and 0.6 for the 50 percent premium.

4 The Cyclical Behavior of Markups

As discussed in section 2, the average markup is proportional to the inverse of the labor share. We study four measures of the labor share covering several broad aggregates. Our primary measure is the BLS's index of labor share for the private business sector. This is the broadest aggregate measure. We also consider three labor share measures constructed from the U.S. national income and product accounts (NIPA). Labor shares for private business and manufacturing are constructed by dividing wage and salary disbursements by total income. We also consider a measure for the nonfinancial corporate business sector that divides labor compensation by gross value added less indirect taxes.[13] Appendix A provides additional details.

4.1 The Markup Measured with Average Wages

Figure 6 displays measures of the markup based on average wages, as defined in equation 10 (i.e., assuming Cobb-Douglas production). The top series is the markup in nonfinancial corporate business, the measure favored by Rotemberg and Woodford (1999). The middle two series plot the markup in private business using NIPA and BLS data. The bottom series shows the markup in manufacturing. The most salient characteristic of all four measures is the propensity to trough during a recession and to peak in the middle of an expansion. That is, they all appear to be procyclical.

We assess the cyclicality of the markups more systematically in three ways. First, we test whether the markups are indeed lower during recessions by regressing the log markup on a quadratic time trend and a dummy variable for recessions:

$$\ln M_t \equiv \mu_t = \beta_0 + \beta_1 t + \beta_2 t^2 + \beta_3 \mathrm{I}(\text{recession}_t) + \omega_t, \tag{15}$$

where ω_t is the error term. We consider a recession to start in the quarter following the business cycle peak, as dated by the National Bureau of Economic Research (NBER),

13. The first three measures do not subtract indirect taxes because of data availability. However, in nonfinancial corporate business we find that the tax-adjusted markup is *more* procyclical than the unadjusted markup. Thus, the measures that do not adjust for indirect taxes have a countercyclical bias.

and end in the quarter of the trough. As a second measure of cyclicality, we calculate the contemporaneous correlation between the cyclical components of the log of real GDP and the log of the markup, where the cyclical components are extracted using an HP filter. Finally, we study the contemporaneous correlation using a first-difference filter.

The first four rows in table 1 report our three measures of cyclicality for the markup using average wages. The first column shows that in every case the markup is estimated to be significantly lower relative to trend during a recession, verifying the visual impression from the graphs. The second and third columns show that the cyclical component of the markup is positively correlated with the cyclical component of GDP, whether we use an HP filter or take first differences. The correlations range from 0.20 to 0.50. Thus, the average markup in all four broad sectoral aggregates is procyclical.

4.2 The Markup Measured with Marginal Wages

We next assess the cyclicality of markups measured using marginal wages. Our estimate of the markup is the product of the average markup and the marginal-average wage adjustment factor. We focus our analysis on the markup based on the BLS labor share for private business. We prefer this measure for two reasons. First, the BLS and NIPA data have different sampling frames and the resulting aggregates sometimes differ. Since our hours data are from the BLS, we prefer the BLS measure of labor share. Second, in one of our specifications we require additional data that is available for the BLS private business sector but not for the other aggregates.

Figure 7 shows the average markup in private business along with two measures of the marginal markup that differ only in the effective overtime premium. The top panel plots the markup series and the bottom panel highlights only the cyclical component of the series extracted using the HP filter. The bottom panel shows that there is essentially no cyclical difference between the average and marginal measures of the markup. Rows 5 and 6 of table 1 show that the marginal markup for the private business sector is significantly lower during recessions and is positively correlated with GDP. Hence, both measures of the marginal markup for the broadest aggregate are significantly procyclical.

Figure 8 plots the cross-correlations of the cyclical components of the markup measures with GDP. The correlations for private business, shown in the upper panel, are

positive for all leads and current values, indicating that an increase in the markup signals a forthcoming increase in GDP. The peak correlation occurs at a lead of two quarters. The correlations only become negative for lagged values of two quarters or greater, meaning that a current decrease in GDP signals an upcoming increase in markups.

Figure 9 shows the average and marginal markups for manufacturing. The most noticeable feature is the extraordinary run-up between 2001 and 2006 in all three measures. The bottom panel removes these low frequency movements in order to focus on behavior over the business cycle. Adjusting for the marginal-average wage factor makes the markup in manufacturing less procyclical. Rows 7 and 8 of table 1 show that the correlation with the cyclical component of GDP falls from 0.4 to nearly zero when we assume that the effective overtime premium is 50 percent. In no case, though, is there evidence of countercyclical markups.

The bottom panel of figure 8 plots the cross-correlations of the cyclical components of the manufacturing markup measures with GDP. The average markup in manufacturing is positively correlated with real GDP for all leads and current values and has a slightly higher contemporaneous correlation than for all workers. Adjusting for the marginal-average wage factor reduces the correlation at leads and lags of one year or less and draws out the cyclical response relative to the average markup.

4.3 CES Production Function

We also consider a measure of the aggregate markup under the assumption that the production function has a lower elasticity of substitution between capital and labor than the Cobb-Douglas production function. This markup is based on the expression in equation 12. The extra term in this markup measure consists of output divided by the product of total hours and the level of technology (Y/AhN) raised to the inverse of the elasticity of substitution. Thus, this measure requires an estimate of the level of technology (A), as well as the elasticity of substitution.

We consider two methods of estimating A. The first estimates technology as the HP trend in labor productivity. This method assumes that all business cycle variation in labor productivity is due to factors other than technology. The second uses Galí's (1999) structural vector autoregression (VAR) method for identifying a technology shock as the only shock that has a permanent effect on labor productivity. To do this, we es-

timate a bivariate VAR with productivity growth and the growth of hours per capita using quarterly data for 1948–2008. We define technology as the part of labor productivity explained by the technology shock. The estimates from this method imply that much of the cyclical variation in labor productivity is due to technology.

Chirinko (2008) surveys the literature estimating the elasticity of substitution between capital and labor and concludes that the elasticity is in the range of 0.4 to 0.6. Thus, we use an elasticity of substitution between capital and labor of 0.5, along with the marginal wage factor with an overtime premium of 25 percent. Rows 9 and 10 of table 1 report the cyclicality of this markup with the two different estimates of technology. The first column shows that both measures are significantly lower during a recession. The third column shows that both measures are even more procyclical than the one that assumes Cobb-Douglas; the correlation with the cyclical component of GDP is 0.4.[14]

To summarize our results, we find that markups measured using average wages are procyclical in the aggregate economy and in manufacturing. Adjusting the markup for the difference between marginal and average wage has little effect in the aggregate data, but makes the markup less procyclical in the manufacturing data.

4.4 Comparison to Bils's (1987) Estimates

Despite building on his insights, we reach the opposite conclusion from Bils concerning the cyclicality of markups. Thus, in this section we investigated the potential sources of the differences.

There are numerous differences in the way the theory is implemented. We estimate a mixed parametric and nonparametric specification using aggregate manufacturing data for 1956–2009 and investigate the cyclical correlation between the estimated markup and real GDP; Bils estimates a polynomial parametric specification using annual 2-digit Standard Industrial Classification (SIC) manufacturing data for 1956–83, and investigates the cyclical correlation between the estimated markup and labor input. Thus, candidates for the source of the different conclusions are (i) a nonparametric specification versus a polynomial specification; (ii) industry data versus aggregate manufacturing data; (iii) the time period; (iv) the frequency of the data; and (v) whether output or labor input is used as the cyclical indicator.

14. We do not explore non-Cobb-Douglas markups for manufacturing because there are no quarterly data on labor productivity before 1987.

To determine the source of the differing conclusions, we explore the effects of these implementation details. To bring our analysis closer to Bils, we replicate Bils's parametric approach by using his polynomial specification for the estimation of $\mathrm{d}v/\mathrm{d}h$. In particular, we estimate:

$$
\begin{aligned}
(16) \qquad \Delta v_{it} = & \Big\{ b_{i0} + b_{i1}t + b_2 t^2 + b_3 t^3 + c_1 \left[h_{i(t-1)} - 40 \right] + c_2 \left[h_{i(t-1)} - 40 \right]^2 \\
& + c_3 \left[h_{i(t-1)} - 40 \right]^3 \Big\} \Delta h_{it} + a_{i0} + a_1 t + a_2 t^2 + a_3 t^3 \\
& + d_{i1} \ln \left[N_{it}/N_{i(t-1)} \right] + d_{i2} \Delta \ln \left[N_{it}/N_{i(t-1)} \right] + e_{it}.
\end{aligned}
$$

In this equation, all parameters listed as a function of i indicate that the parameters are allowed to differ across industries. The interaction term with Δh includes an industry-specific mean, an industry-specific linear time trend, a common quadratic and cubic function of time, as well as a cubic function of the deviation of the starting level of average hours from 40. The terms outside the interaction with Δh allow for further industry effects and time trends. We also follow Bils in including the growth and change in the growth rate of employment.[15]

We use monthly CES data for 2-digit SIC manufacturing industries. All hours and employment data are for production and nonsupervisory workers. We seasonally adjust the monthly data for each industry and remove outlier observations from holidays, strikes, and bad weather.[16] The annual series we use is the annual average of not-seasonally-adjusted data.

When we estimate this equation on monthly or quarterly data, we use average hours in the previous month or quarter for h_{t-1}. When we estimate this equation on annual data, we follow Bils and use the average of average hours in the previous and current year for h_{t-1}. When we aggregate the 2-digit data, we take a weighted average of h, v, and $\mathrm{d}v/\mathrm{d}h$, using the industry's share of total hours as the weight. For employment, we simply sum across industries.

Table 2 shows the effects of data frequency and industry aggregation on the estimates of $\mathrm{d}v/\mathrm{d}h$. All estimates are for Bils's sample of 1956–83. Consider first using 2-digit industry data and varying the frequency of the data (upper panel). The table shows that monthly and quarterly data give very similar estimates of the slope of $\mathrm{d}v/\mathrm{d}h$ relative to average hours. In contrast, the annual data imply a steeper slope. Thus,

15. See Bils (1987), p. 844, for his motivation for including these terms.
16. See appendix A for details.

time aggregation appears to bias the slope estimate upward. The lower panel shows the same experiment using data aggregated to the level of overall manufacturing. Two results are noteworthy here. First, we continue to see that the annual estimates of the slope are much higher. Second, for the monthly and quarterly data, the industry results are similar to those obtained using the aggregated data. From this we conclude that our use of aggregate industry data does not impart a bias. However, Bils's use of time-aggregated annual data appears make $\mathrm{d}v/\mathrm{d}h$ more procyclical.

Table 3 shows the effect of changing frequencies and using different cyclical indicators on the inferences about the cyclicality of markups. We focus only on the marginal markup measure that assumes a 50 percent premium to compare it to Bils. Because the markup data are not available on a monthly basis, we consider only quarterly and annual data. Row 1 shows the results of using quarterly data to estimate $\mathrm{d}v/\mathrm{d}h$ and applying it to quarterly markups for Bils's sample from 1956–83. The correlation with HP filtered GDP is 0.3. The next column shows the correlation with the cyclical component of output in manufacturing, measured using the index of industrial production (IP) in manufacturing. The correlation is half that for GDP, but is still positive. The last column shows the effect of using total hours in manufacturing as the cyclical measure, which is closer to what Bils did. In this case, the correlation is near zero. Thus, it appears that in Bils's sample, the procyclicality of the markup is attenuated both by using industry-specific output measures and by using industry-specific labor input measures. When cyclicality is measured relative to industry output, markups are still mildly procyclical for quarterly data. However, when cyclicality is measured using total manufacturing hours, markups become acyclical.

The second row of table 3 shows the results when we continue to use quarterly data to estimate $\mathrm{d}v/\mathrm{d}h$ but then time aggregate it and apply it to annual data. In each case, the correlations drop. The correlation is 0.2 when real GDP is used, but essentially zero when either manufacturing output or hours is used. The third row shows the results when we use annual data to estimate $\mathrm{d}v/\mathrm{d}h$ and calculate the markup, which is what Bils's did. In this case, the markup is acyclical or countercyclical for all three indicators of the business cycle. The markup is most countercyclical (a correlation of -0.3) when using total hours—the indicator Bils used—as a cyclical indicator.

To gain insight into the effect of the sample, the fourth row replicates the procedure used in previous row, but extends the sample through 2002.[17] The fourth row

17. The sample does not run through 2009, as our aggregate analysis does, because the SIC-based

shows that the cyclicality falls to near zero when the sample is extended for an additional nineteen years. As in the shorter sample, the markup is more countercyclical when measuring the cycle with manufacturing IP, and still more so using total hours in manufacturing, than when using real GDP.

Finally, because of our finding that different cyclical indicators can lead to different implied cyclical properties of the markup, we explore the effects on the results that use our mixed parametric and non-parametric procedures on quarterly data for 1956–2009. The fifth row shows these results. The first column reproduces our main result from table 1: with a 50 percent effect overtime premium, the cyclical correlation of the markup with real GDP is 0.05. The second column shows that the correlation with the cyclical component of manufacturing IP is 0.01 and with that of total hours is −0.05. Thus, unlike when we use Bils's sample (row 1), we find little difference in the cyclicality of the markup across cyclical indicators when we use the full sample to 2009.

In sum, Bils's use of time-aggregated annual data and his choice of cyclical indicator for his sample period were all necessary conditions for finding a countercyclical markup. In contrast, we find that markups in manufacturing are procyclical to acyclical in quarterly data, whether we use our mixed, parametric and time-varying parameter model or Bils's parametric model and even with a 50 percent overtime premium.

5 Effect of a Monetary Policy Shock on Markups

In many New Keynesian models, such as those by Goodfriend and King (1997) and Smets and Wouters (2003, 2007), money is nonneutral because all prices do not adjust immediately. A contractionary monetary policy shock raises the markup because marginal cost falls more than price. Thus, the markup should move countercyclically if the business cycle is driven by monetary policy. However, because these models also imply that the markup increases in response to a technology shock, a procyclical markup does not, by itself, necessarily invalidate the models.

To test the mechanism of these models more directly, we investigate the response of our markup measures to a monetary policy shock. To do this, we add the markup to a standard monetary VAR. The VAR consists of the log of real GDP, the log of commodity prices, the log of the GDP deflator, the log markup measure, and the federal

industry data are not available after 2002.

funds rate.[18] We include four lags of each variable and a linear time trend. Following standard practice, we identify the monetary policy shock as the shock to the federal funds rate when it is ordered last. We estimate the VAR using quarterly data over 1960:3–2009:4.

We consider the average markup in private business and several measures of the marginal markup in the private business and manufacturing sectors. Figure 10 shows the impulse response of these markups to a positive shock to the federal funds rate—a contractionary monetary shock. The impulse response functions for the other variables (not shown) are similar to those in the literature; in particular, output falls and stays below trend for about four years.

In every case the markup *falls* in response to a contractionary monetary shock. Furthermore, the responses are below zero at conventional significance levels. Interestingly, the markup in manufacturing falls more than the markup for the overall economy. Thus, the behavior of markups is contrary to the mechanism of the New Keynesian model. However, one should not confuse statistical significance with economic significance. The results imply that a monetary shock that leads GDP to fall by one percent leads the markup to fall by just under one percent. Thus, if the markup starts out at 1.20, then it falls to just 1.19.

One possible omission of our markup measure is its failure to capture the cost channel. Barth and Ramey (2002), Christiano et al. (2005), and Ravenna and Walsh (2006) have argued that a contractionary monetary policy shock might raise firms' costs by raising interest rates. If firms must finance working capital, then an increase in interest rates raises their marginal cost. To include this effect, we multiply the wage measure of marginal cost by the gross nominal interest rate. For this we use a quarterly average of the prime interest rate. The middle right panel of figure 10 shows the aggregate marginal markup with an effective overtime premium of 25 percent, including the gross nominal interest rate as a part of marginal cost. Allowing for a cost channel does little to change the procyclicality of the markup.

To summarize, the New Keynesian model requires markups to rise in response to a contractionary monetary shock in order to generate monetary nonneutrality. In the data, none of the markup measures rises in response to a monetary shock. In fact, all estimates of markups fall. Thus, our finding of procyclical markups in the previous section cannot be explained away with technology shocks.

18. Details of the data sources are in appendix A.

6 Discussion of the Aggregate Results

In one sense, our results should not be surprising since procyclical markups are a direct consequence of countercyclical labor shares. However, the New Keynesian literature seems to have overlooked the potential contradiction between the transmission mechanism required by the theory and the variables that perform best empirically.[19] For example, Galí and Gertler (1999) and Sbordone (2002) improved the performance of the New Keynesian Phillips curve by substituting a measure of real marginal cost for the output gap. Previously, the output gap had been used to predict inflation, but the estimated coefficient was negative, in contradiction to the theory. Galí and Gertler (1999) and Sbordone (2002) argued that "real marginal cost" or "unit labor cost" was the more theoretically-correct variable. Their measure of real marginal cost was, in fact, the labor share. This variable predicted inflation very well, and entered with the correct positive sign. The output gap and the labor share perform so differently because they are negatively correlated with each other, i.e., the output gap is procyclical and the labor share is countercyclical. This negative relationship is in direct contradiction to the theory.[20] Unfortunately for the model, the countercyclical labor shares that work so well in the New Keynesian Phillips curve imply procyclical markups (measured using average wages).

Our results are also not surprising in the sense that procyclical markups imply a procyclical capital share. As Hall (2004) argues, cyclical patterns in firm rents are linked to adjustment costs on capital. Although he estimates very low adjustment costs, most estimates in the literature suggest higher adjustment costs that are completely consistent with procyclical capital share.[21] Moreover, as Christiano et al. (2005) note, profits increase significantly in response to an expansionary monetary shock in the data. Because Christiano et al. (2005) allow for a cost channel of monetary policy in their theoretical model, an expansionary monetary shock can lead profits to increase slightly. However, as their figure 1 shows, the greatest gap between the data and model is in the behavior of profits; profits rise much more in the data than in the model. The evidence on markups that we present provides further cause for concern about the extent to which the current New Keynesian models capture the transmission of demand shocks.

Why then do Rotemberg and Woodford (1999) argue that markups are counter-

19. We are indebted to Olivier Coibion for pointing this out to us.
20. See, for example, Woodford's (2003) equation 2.7 on page 180.
21. See Cooper and Haltiwanger (2006) and Gourio and Kashyap (2007).

cyclical? Figure 2 of their chapter shows that labor share rises significantly during every recession. Their table 1 also shows that labor share is countercyclical, with most correlations between labor share and cyclical indicators being negative.[22] Since labor share is the inverse of the average markup, their figure and table imply a very procyclical average markup. Their ultimate conclusion that the markup is countercyclical is thus somewhat puzzling. After showing a procyclical markup using labor share data, Rotemberg and Woodford (1999) consider more general specifications, such as non-Cobb-Douglas production functions and overhead labor, and find some evidence that markups move countercyclically. However, many of their calculations were based on educated guesses about parameters that were not well-measured at the time, such as the elasticity of substitution between capital and labor and the fraction of labor that is overhead labor. We use richer data, and more recent estimates of key parameters, to analyze the key generalizations and show that the markup continues to be procyclical.

Our results do support one of Rotemberg and Woodford's (1999) conclusions. In the first part of the introduction they argue that "there exists of great deal of evidence in support of the view that marginal cost rises more than prices in economic expansions, especially late in expansions."[23] Although our results contradict the first part of that statement, we find that markups begin to fall during the last part of the economic expansion. However, it is a mistake to infer that markups are countercyclical from this one feature of the data. The rise in markups in the first half of an expansion might be linked to adjustment costs on capital. During the second half of the expansion, capital has adjusted so that rents are dissipated.

One key assumption made in our work, as well as in virtually all of the New Keynesian models, is that wages are allocative and that firms are on their labor demand curves. If wages include insurance aspects, as suggested by Baily (1974) and Hall (1980), then our measures of marginal costs based on wages may not indicate the true marginal cost of increasing output. Also, while our method allows for adjustment costs on the number of workers, if firms engage in labor hoarding and are prevented from lowering hours per worker below some threshold, then the true marginal cost of an extra hour of labor may fall much more in a recession than suggested by our measure. In

22. The main exceptions are when they detrend hours using a linear trend. However, Francis and Ramey (2009) show that hours worked per capita exhibit a U-shape in the post–World War II period because of the effects of the baby-boom and sectoral shifts. Thus a linear time trend does not adequately capture the low frequency movements.

23. Rotemberg and Woodford (1999), p. 1053.

any case, these same considerations apply to any analysis, including the New Keynesian models, that use wage data to estimate parameters of the model.

Unfortunately, other methods for inferring the cyclicality of marginal cost or markups are fraught with their own weaknesses. For example, Bils and Kahn (2000) argued that one could infer the cyclicality of markups from the cyclicality of the inventory-sales ratio. They presented a stock-out avoidance model in which sales depended positively on the level of inventories. Based on that model, they showed that the countercyclicality of the inventory-sales ratio implied countercyclical markups.

However, it turns out that this result is model specific. For example, in a standard linear-quadratic production-smoothing model with a target inventory-sales ratio and a monopolist facing linear demand, a positive demand shock can easily lead to an increase in the markup and a decrease in the inventory-sales ratio. In Khan and Thomas's (2007) general-equilibrium S,s model, the inventory-sales ratio is strongly countercyclical, even though the markup is constant in their model. It is difficult to choose one model over another because most of them are rejected when subjected to formal econometric tests. Moreover, Bils (2005) tested the hypothesis of Bils and Kahn (2000) using microdata from the consumer price index (CPI). According to the Bils-Kahn model, stock-outs should be procyclical. For a sample of durable goods, Bils (2005) found that stock-outs are completely acyclical. In the context of the Bils-Kahn model, this implies that markups are acyclical. Thus, the behavior of inventories is not particularly informative about the markup.

In sum, we have found that measures of the markup based on the inverse of the labor share are procyclical in both the aggregate economy and the manufacturing sector. This measure is identical to the inverse marginal cost measure used in New Keynesian Phillips curve models. Adjustments that convert the average wage to the marginal wage do not mitigate the procyclicality at the aggregate level. At the manufacturing level, the adjustments create a markup that is less procyclical—but not countercyclical.

7 Industry Analysis

We now turn to an analysis of a panel of 4-digit SIC manufacturing industries. Using detailed industry-level data has several advantages. First, since sectoral shifts might drive the aggregate results, it is useful to examine the cyclicality of the markup at the industry level. Second, the industry data allow us to use gross output rather than

value-added output. As discussed above, Basu and Fernald (1997) argue that standard value-added measures are only valid under perfect competition. Thus, it is inconsistent to use value-added measures to explore the cyclicality of markups. Third, the industry data allows us to distinguish between production workers and other workers. Some have argued that overhead labor is an important factor in estimating markups; we test whether our results are sensitive to including nonproduction workers, who are more likely to be overhead labor. Fourth, the industry data allow us to create much richer variables for testing New Keynesian explanations of the effects of aggregate demand shocks. In particular, we are able to use detailed industry-specific changes in government spending as instruments for studying the behavior of markups. The New Keynesian model predicts that the markup falls in response to an increase in government spending. Thus, it is particularly interesting to study this potential mechanism in detail.

There are two disadvantages of this data source, however. One is that the data are only available at annual frequency. As discussed in the previous section, it appears that time-averaging the data tends to induce more countercyclicality. Thus, using annual data biases the results toward finding countercyclical markups. The second disadvantage is the lack of information to estimate our marginal-average wage adjustment factors. Thus, we must use markups based on average wages, which we show in section 4 are more procyclical than those that are adjusted for marginal wages. Thus, this limitation of the data imparts a bias in the opposite direction. To the extent that these two biases had similar magnitudes in the aggregate manufacturing data, it is likely that our analysis using annual 4-digit industry data will not be strongly biased in either direction.

7.1 Data and Variable Construction

Building on the ideas of Shea (1993) and Perotti (2008), we match data at the 4-digit SIC level on government spending and its downstream linkages calculated from benchmark input-output (IO) accounts to the NBER–Center for Economic Studies (CEcS) Manufacturing Industry Database (MID).[24] The available years for the benchmark IO tables are 1963, 1967, 1972, 1977, 1982, 1987, and 1992.[25] Merging manufacturing

24. Technical documentation for the NBER–CEcS MID is available in Bartelsman and Gray (1996).

25. More recent years of the IO benchmark tables are available. However recent tables use the North American industrical classification system (NAICS) and thus merging industries becomes considerably

SIC industry codes and IO industry codes yields 272 industries. Appendix B details how we merged the two data sets.

Because it is difficult to distinguish nondefense from defense spending when calculating indirect effects, we use total federal government spending. Using total federal spending should not be problematic because most of the level and variation in federal government purchases is defense spending and related spending, such as for the National Aeronautics and Space Administration. Also, we use a measure that includes both direct and indirect government spending. We include both direct and indirect government spending since both are potentially important demand shifters. For example, if the military purchases more airplanes from aircraft manufacturers, which in turn buy more aircraft parts, then indirect government purchases are an important component of demand shifts for aircraft parts.

Perotti (2008) defined his government demand variable as the change in an industry's shipments to the government between two benchmark years divided by total initial shipments of the industry, i.e., $\frac{G_{it}-G_{i(t-5)}}{S_{i(t-5)}}$. As Nekarda and Ramey (2010) argue, there is reason to believe that Perotti's measure is correlated with industry-specific technological change, and hence is invalid as a demand instrument. We use the Nekarda-Ramey measure since it should not be correlated with industry-specific technology.

In particular, we define the government demand growth for industry i, Δg_{it}, as

$$\Delta g_{it} = \phi_i \Delta g_t, \tag{17}$$

where ϕ_i is the average share of an industry's total nominal shipments that go to the government and g_t is log of aggregate real federal spending (based on NIPA data). We calculate the share of industry i's total nominal shipments that go to the government in year t as

$$\phi_{it} = \frac{G_{it}}{\sum_{j=1}^{J} S_{ijt}}, \tag{18}$$

where S_{ijt} is the value of inputs produced by industry i shipped to industry j in year t and G_{it} is total government demand for industry i in year t.

We then average over the time period the input-output data are available (1963–92) to calculate the industry's average dependence on the government, ϕ_i. Thus, this

more difficult and fraught with potential error.

measure converts the aggregate government demand variable into an industry specific variable using the industry's long-term dependence on the government as a weight. Nekarda and Ramey (2010) discuss the properties of this variable in more detail.

The remaining variables are constructed using data from the MID. This database provides information on total employment, as well as employment in the subcategories of production and nonproduction workers. Unfortunately, the MID provides information on annual hours only for production workers. To create an hours series for all workers, we constructed two measures of total hours. We consider two extreme assumptions: (a) nonproduction workers always work 1,960 hours per year and (b) nonproduction workers always work the same number of annual hours as production workers. The results are similar under both assumptions; we report the results using the conservative assumption that nonproduction workers' hours are not cyclical.

We next create series on average weekly hours for production workers and for all workers in the industry data by dividing total annual hours by 49 times production worker employment. We use 49 weeks rather than 52 weeks because the MID does not include vacation and sick leave in its accounting of hours. Our assumption yields a series on average hours for production workers in the industry database with a mean of 40.7, very close to the mean of 40.8 in the CES manufacturing data over 1958–96.

The average wage for production workers is the production worker wage bill divided by production worker hours. The average wage for all workers is the total wage bill divided by total hours.

7.2 Empirical Specification and Results

Our goal is to estimate how the markup responds to a change in output induced by shifts in demand. To construct our markup measure, we add industry (i) and time (t) subscripts to equation 10 and take annual log differences. The log change in the markup over average cost can be written as

$$\Delta\mu_{it} = -\Delta \ln\left(s_{it}\right), \tag{19}$$

which is the negative log change in the labor share, defined as the wage bill divided by the value of shipments.

Our estimation involves regressing the change in the markup, $\Delta\mu$, on the change

in the natural logarithm of real shipments, $\Delta \ln Y$. In particular, we estimate:

$$\Delta\mu_{it} = \gamma_{0it} + \gamma_1 \Delta \ln Y_{it} + \varepsilon_{it}, \tag{20}$$

where $\Delta\mu_{it}$ is now given by equation 19 and ε_{it} is the error term. The coefficient γ_0 depends on both i and t because we include industry and year fixed effects in some specifications. The coefficient γ_1 describes how the markup responds to a change in output. The regressions are run using a panel with 10,135 observations.

Table 4 reports estimates of $\hat{\gamma}_1$ under several specifications. The first row uses markups calculated for production workers only and the second row uses markups calculated including both production and nonproduction workers.

The first column reports coefficients from OLS regressions without industry and year fixed effects. The coefficient on log shipments is negative, but not statistically significant in either specification. The small negative unconditional correlation is consistent with the effects of time aggregation we discussed in previous sections. Controlling for industry and year fixed effects makes the coefficient slightly positive for production workers and more positive for all workers.

These OLS regressions summarize how the markup responds to a change in shipments on average. However, shipments can change both due to demand shifts and technological change. Thus, to isolate changes in markups due to demand shifts, we use our government spending variable as an instrument for total shipments. To be a valid demand instrument, it must be both uncorrelated with technology and relevant. On the first feature, Ramey (2009) presents evidence that most fluctuations in aggregate federal spending are driven by military buildups, which in turn are driven by political, rather than economic, events. A key additional question, however, is whether the allocation of government spending by industry might be affected by technological change. As discussed in Nekarda and Ramey (2010), there is evidence that Perotti's (2008) definition of industry-specific government demand is correlated with industry-specific technological change. In contrast, our instrument is constructed so that it cannot be correlated with the time variation in industry-specific technological change. To control for long-run differences in technological change across industries, we also include industry fixed effects. In addition, we include year fixed effects to control for any year-specific aggregate influences. The second feature required is that the industry be relevant. Even after controlling for industry and year fixed effects, the first-stage re-

gression of the industry-specific annual growth of shipments on the government instrument produces an F statistic is 198. Thus, the instrument clearly passes the relevance test.

The third column of table 4 reports the instrumental variables (IV) regressions. This model instruments $\Delta \ln Y_{it}$ with Δg_{it} in the specification with industry and year fixed effects. For production workers, the estimate is positive, but not different from zero, either statistically or economically. For all workers, the estimate is positive and is marginally statistically different from zero. It is economically small, though.

Rotemberg and Woodford (1991) argue that the standard markup measure might be biased toward being procyclical if overhead labor is important. As Ramey (1991) argues, however, production worker hours are much less likely than total worker hours to include overhead labor. Thus, a comparison of our results for production workers to those for all workers might shed light on the potential bias. For most specifications shown in table 4, the markup measured using production workers is less procyclical than the markup measured using all workers. This result supports the notion that including overhead workers may make the markup more procyclical. Nevertheless, the markup measured using production workers is still acyclical.

To summarize, we find that markups are essentially acyclical in response to demand shocks in disaggregated industry-level manufacturing data at the annual frequency. These results are not inconsistent with our findings of procyclical markups at the aggregate level, since our comparison with Bils revealed that annual correlations tend to bias the results toward countercyclicality as compared to quarterly or monthly correlations.

8 Conclusion

This paper has presented evidence that markups are largely procyclical. Whether we look at broad aggregates or detailed manufacturing industries, average wages or marginal wages, we find that most measures of the markup are significantly procyclical and no evidence of countercyclical markups. These results hold even when we confine our analysis to changes in output driven by monetary policy or government spending. In both cases, an increase in demand leads to higher markups, although the size of the increase is not economically large.

Our results call into question the basic mechanism of the leading New Keynesian

models. These models assume that monetary policy and government spending affect the economy through their impact on markups. If prices are sticky, an increase in demand should raise prices less than marginal cost, resulting in a fall in markups. Even with sticky wages, most New Keynesian models still predict a fall in markups. Our empirical evidence suggests that the opposite is true.

Appendix A
Aggregate Analysis

Data Sources

Item		*Frequency and source*
Average weekly hours	M	Aggregate economy: BLS *Employment and Earnings*; see "Marginal Wage" below
	M	Manufacturing: BLS series `CEU3000000007`
Average weekly overtime hours	M	Aggregate economy: BLS *Employment and Earnings*; see "Marginal Wage" below
	M	Manufacturing: BLS series `CEU3000000009`
Labor share, nonfinancial corporate business	Q	`Compensation of employees` divided by `Gross value added of nonfinancial corporate business` less `Taxes on production and imports less subsidies`; all series from NIPA table `1.14`
Labor share, private business (NIPA) & manufacturing	Q	`Wage and salary disbursements` divided by `National income without capital consumption adjustment`; wages from NIPA table `2.2A/B`, national income from NIPA table `6.1B/C/D`
Labor share, private business (BLS)	Q	BLS series `PRS84006173`
Labor productivity	Q	BLS series `PRS84006093`
Real GDP	Q	NIPA table `1.1.6`
Implicit GDP price deflator	Q	NIPA table `1.1.9`
Commodity price index	M	BLS series `WPU00000000`
Federal funds rate	M	Board of Governors of the Federal Reserve System `H.15 release`
Prime loan rate	M	Board of Governors of the Federal Reserve System `H.15 release`

Seasonal Adjustment

Both the Current Population Survey (CPS) and Current Employment Statistics (CES) surveys ask respondents to report actual hours worked during the week of the month containing the twelfth. Two holidays, Easter and Labor day, periodically fall during the reference week. When one of these holidays occurs during the reference week, actual hours worked falls substantially. However, because there this pattern is not regular, the seasonally-adjusted series published by the Bureau of Labor Statistics (BLS) do not account for this reference-week holiday effect.

We seasonally adjust the CPS data we construct on employment and average hours for all workers in the aggregate economy using use the Census Bureau's X-12-ARIMA program. We allow the program to remove all large outliers from the final seasonally-adjusted series. Because the seasonally-adjusted series on production workers in manufacturing are seasonally adjusted at a highly-disaggregated level, we do not perform our own seasonal adjustment. Instead, we use X-12-ARIMA to identify outliers from the seasonally-adjusted series and remove the outlier component from the seasonally-adjusted series.

Average Markup

We construct measures of the price-average cost markup as the inverse of the labor share. Details of the data are provided in the table above.

For the private business sector using BLS data, the markup is 100 divided by the index of labor share. For the private business sector using national income and product accounts (NIPA) data, the markup is wage and salary disbursements divided by national income without capital consumption adjustment. The markup in manufacturing is constructed analogously to NIPA private business. Finally, the markup for nonfinancial corporate business is constructed from the NIPA by dividing compensation of employees by gross value less taxes on production and imports less subsidies.

Marginal Wage

Data for estimating the marginal wage come from the BLS's monthly *Employment and Earnings* publication. We use three categories of data: persons at work by hours worked, total persons at work, and average hours worked by persons at work.

Data on total persons at work and average hours worked by persons at work are from Cociuba et al. (2009). Data on persons at work by hours worked are from two sources. Data are available on the BLS web site for June 1976 to present, but earlier data are taken from *Employment and Earnings*. Presently we are missing data for May 1963. This missing value is interpolated using Gómez and Maravall's (1996) Time Series Regression with ARIMA Noise, Missing Observations, and Outliers (TRAMO).

The table below lists the BLS series for persons at work for June 1976 to December 2009.

Category	*Series*
1–4 hours	LNU02010044
5–14 hours	LNU02010385
15–29 hours	LNU02010726
30–34 hours	LNU02011067
35–39 hours	LNU02011501
40 hours	LNU02011797
41–48 hours	LNU02012093
49–59 hours	LNU02012389
60+ hours	LNU02012685
Total at work	LNU02005053
Average hours, persons at work	LNU02005054

The hour bins into which the BLS classifies workers change in 1967. The table below shows the concordance of which bins are combined to create categories that are consistent for the entire sample.

1960–66	*1967–2009*	*Universal*
1–14	1–4 5–14	1–14
15–21 22–29	15–29	15–29
30–34	30–34	30–34
35–39	35–39	35–39
40	40	40
41–47 48	41–48	41–48

1960–66	*1967–2009*	*Universal*
49–54 55–59	49–59	49–59
60–69 70+	60+	60+

Average Hours Worked per Bin

We use data from the CPS's Annual Social and Economic Study ("March") supplement to calculate the average actual number of hours worked for the bins above. This allows us to calculate the total number of hours worked and total overtime hours worked. We use data from the March supplement rather than monthly data because the annual supplement data are available back to 1962.

Using the March supplement instead of a monthly series is relatively innocuous. Average hours worked per week in March is generally representative of average hours worked over the entire calendar year. Also, hours calculated from the monthly Basic CPS data are very similar to actual hours reported in the March supplement.

Figure A1 plots the annual average of hours worked per bin calculated from the monthly Basic CPS together with average hours calculated from the March CPS. Because the annual average and the March level are not always the same, we scale the March CPS data to the annual average. To account for the difference following the 1994 redesign, we calculate the adjustment factors over two subsamples: 1976–93 and 1994–2009. We apply the pre-1994 factors to all data before the redesign. The adjustment factor is simply the difference between the period average of the annual average and of the March CPS. The table below reports the mean adjustment.

Bin	*Pre-1994*	*Post-1994*
1–14	−0.063	−0.028
15–29	0.148	0.080
30–34	0.048	0.034
35–39	−0.030	−0.012
40	0.000	0.000
41–48	0.009	0.001
49–59	−0.024	−0.004
60+	0.389	0.219

Because the March CPS data begin in 1962, we must estimate the average hours per bin for 1960 and 1961. We estimate a linear trend over the 1962–67 period and project that trend backward to obtain estimates for 1960 and 1961. The remaining years use the actual value from the March CPS for average hours per bin. The table below reports summary statistics of the adjusted average hours per bin in the March CPS over 1960–2009.

Bin	*Mean*	*Standard deviation*	*Minimum*	*Maximum*
1–14	7.9	0.2	7.6	8.4
15–29	21.1	0.1	20.9	21.3
30–34	31.3	0.1	31.1	31.5
35–39	36.3	0.1	36.1	36.4
40	40.0	0.0	40.0	40.0
41–48	45.3	0.4	44.8	46.0
49–59	51.8	0.3	51.5	52.7
60+	68.2	1.2	66.8	70.4

We calculate total hours worked in bin b, H_t^b, using the number of persons at work by bin, N_t^b, and the number of hours worked per bin, h_t^b. We seasonally adjust N_t^b and apply the annual average of hours (h_t^b) to create the seasonally adjusted series of total hours per bin as $H_t^b = h_t^b N_t^b$.

We calculate total hours as the sum of all bins:

$$H_t = H_t^{1-14} + H_t^{15-29} + H_t^{30-34} + H_t^{35-39} + H_t^{40} + H_t^{41-48} + H_t^{49-59} + H_t^{60+} \tag{A.1}$$

and overtime hours as

$$H_t^{OT} = \left(h_t^{41-48} - 40\right) N_t^{41-48} + \left(h_t^{49-59} - 40\right) N_t^{49-59} + \left(h_t^{60+} - 40\right) N_t^{60+}. \tag{A.2}$$

Share of Overtime Hours That Are Paid a Premium

We calculate share of overtime hours that are paid a premium using data from CPS May extracts provided by the National Bureau of Economic Research (NBER) (http://www.nber.org/data/cps_may.html). The overtime variable (x174) is a dummy for whether an individual receives higher pay for work exceeding 40 hours in a week.

(Note that the value 0 indicates that a worker received premium pay.)

We drop all individuals that do not report total hours (variable x28). We calculate overtime hours as hours worked at primary job (variable x182) less 40 when this is reported; otherwise, overtime hours is calculated as total hours worked less 40. An individual's paid overtime hours is the product of overtime hours and the indicator for whether overtime hours are paid a premium. We aggregate overtime hours, paid overtime hours, and total hours by year using the individual sampling weights (variable x80). For a given year, the share of overtime that is paid a premium is the ratio of paid overtime hours to total overtime hours.

Appendix B
Industry Analysis

This section describes the data sources for the industry analysis, in particular how we merge the input-output (IO) data with the NBER–Center for Economic Studies (CEcS) Manufacturing Industry Database (MID) data.

Input-Output Data

We use data from the Bureau of Economic Analysis (BEA)'s benchmark IO tables to construct measure of industry-level government spending. The benchmark revisions the IO tables are conducted every five years. Benchmark tables are available for 1963, 1967, 1972, 1977, 1982, 1987, 1992. Starting in 1997, the IO tables moved from a Standard Industrial Classification (SIC)-based classification to one based on the North American industrical classification system (NAICS); we do not use the NAICS–based IO tables.

The data are available on the U.S. Census Bureau's web site (http://www.bea.gov/industry/io_benchmark.htm). The table below lists the file names for the transactions and total requirements benchmark IO tables.

Year	*Source file*
Transactions	
1963	`1963 Transactions 367-level Data.txt`
1967	`1967 Transactions 367-level Data.txt`
1972	`1972 Transactions 367-level Data.txt`
1977	`1977 Transactions 366-detail Data.txt`
1982	`82-6DT.DAT`
1987	`TBL2-87.DAT`
1992	`SICUSE.TXT`
Total requirements	
1963	`1963 Transactions 367-level Data.txt`
1967	`1967 Transactions 367-level Data.txt`
1972	`1972 Total Requirements 365-level Data.txt`
1977	`1977 Total Req Coeff 366-level Data.txt`
1982	`82CCTR.TXT`
1987	`TBL4-87.DAT`
1992	`CXCTR.TXT`

Except for 1963, the IO data are available at a 6-digit level (537 industries); in 1963, the data are available at 4-digit level (367 industries). All calculations are performed at the most disaggregated level available.

Let S_{ijt} be the value of inputs produced by industry i shipped to industry j in year t, measured in producers' prices. Direct government demand for industry i is the value of inputs from industry i used by the federal government ($j = g$):

$$G^d_{it} = S_{igt}. \tag{B.1}$$

Although the IO tables distinguish between defense and nondefense federal purchases, we take the sum of both categories. The table below lists the IO codes for the government for each benchmark IO table.

	Industry code	
Year	*Defense*	*Nondefense*
1963	9710	9720
1967	971000	972000
1972	960000	970000
1977	960000	970000
1982	960000	970000

	Industry code	
Year	*Defense*	*Nondefense*
1987	960000	970000
1992	9600I0, 9600C0	9700I0, 9700C0

Indirect government demand is calculated using commodity-by-commodity unit input requirement coefficients. Let r_{ijt} be the commodity i output required per dollar of each commodity j delivered to final demand in year t. The indirect government demand for industry i's output is the direct government purchases from industry j times the unit input requirement of industry i for industry j's output:

$$G_{it}^{n} = \sum_{j=1}^{J_t} G_{jt}^{d} \times r_{ijt}. \tag{B.2}$$

Total government demand for industry i in year t is the sum of direct and indirect demand:

$$G_{it} = G_{it}^{d} + G_{it}^{n}. \tag{B.3}$$

After calculating direct and indirect government shipments, the IO data are aggregated (see below) to merge with the MID data.

NBER–CEcS Manufacturing Industry Database

Data on manufacturing industries comes from the NBER–CEcS MID database.[26] The MID database contains annual data on 459 manufacturing industries from 1958–96. The data are compiled from the Annual Survey of Manufacturers and the Census of Manufactures and adjust for changes in industry definitions over time.

We use industry-level information on employment, hours, and the wage bill for production and nonproduction workers; gross shipments; and price deflators.

26. Bartelsman and Gray (1996).

IO-SIC Correspondence

We create a cross-walk between the IO data at various aggregations and years and the 1987 4-digit SIC codes used in the MID data. The ultimate correspondence is between 4-digit IO and 4-digit SIC, but aggregation was required in both data sets to achieve a one-to-one correspondence. A correspondence was developed between the 6-digit IO code–based IO data and the 4-digit SIC code–based MID data. The merged database contains 272 industries at the 4-digit the SIC level.

A supplemental appendix is posted on the authors' web site. It contains the lengthy crosswalk between combined IO and combined SIC codes. It also includes tables detailing how 6-digit IO codes were assigned to a 4-digit IO industry before aggregating to the 4-digit IO level. This assignment changes for each benchmark IO table.

Some SIC codes also required combination. This aggregation occurs mostly at the 3-digit SIC level, which roughly corresponds to the 4-digit IO level, but it was ultimately tailored to preserve the best correspondence with the IO data. A complete list is available in the supplemental appendix.

References

Baily, Martin Neil, 1974. "Wages and Employment under Uncertain Demand." *Review of Economic Studies* 41(1): 37–50.

Bartelsman, Eric J. and Wayne Gray, 1996. "The NBER Manufacturing Productivity Database." Technical Working Paper 205, National Bureal of Economic Research, Cambridge, MA.

Barth, Marvin J. and Valerie A. Ramey, 2002. "The Cost Channel of Monetary Transmission." In *NBER Macroeconomics Annual 2001*, edited by Ben S. Bernanke and Julio Rotemberg, pp. 199–256. University of Chicago Press.

Basu, Susanto and John G. Fernald, 1997. "Returns to Scale in US Production: Estimates and Implications." *Journal of Political Economy* 105(2): 249–83.

Bils, Mark, 1987. "The Cyclical Behavior of Marginal Cost and Price." *American Economic Review* 77(5): 838–55.

Bils, Mark, 2005. "Deducing Markup Cyclicality from Stockout Behavior." Unpublished paper, University of Rochester.

Bils, Mark and James A. Kahn, 2000. "What Inventory Behavior Tells Us About Business Cycles." *American Economic Review* 90(3): 458–81.

Blanchard, Olivier J., 2008. "The State of Macro." Working Paper 14259, National Bureau of Economic Research, Cambridge, MA.

Carr, Darrell E., 1986. "Overtime Work: An Expanded View." *Monthly Labor Review* 109(11): 36–9.

Chirinko, Robert S., 2008. "σ: The Long and Short of It." *Journal of Macroeconomics* 30(2): 671–86.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, 2005. "Nominal Rigidities and the Dynamic Effects of A Shock to Monetary Policy." *Journal of Political Economy* 113(1): 1–45.

Cociuba, Simona, Edward C. Prescott, and Alexander Ueberfeldt, 2009. "U.S. Hours and Productivity Behavior Using CPS Hours Worked Data: 1947-III to 2009-II." Working paper, Federal Reserve Bank of Dallas.

Cooper, Russell W. and John C. Haltiwanger, 2006. "On the Nature of Capital Adjustment Costs." *Review of Economic Studies* 73(3): 611–33.

Deleire, Thomas, Jayanta Bhattacharya, and Thomas MaCurdy, 2002. "Comparing Measures of Overtime Across BLS Surveys." *Industrial Relations* 21(2): 362–69.

Domowitz, Ian R., Glenn Hubbard, and Bruce C. Petersen, 1986. "Business Cycles and the Relationship between Concentration and Price-Cost Margins." *The RAND Journal of Economics* 17(1): 1–17.

Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin, 2000. "Optimal Monetary Policy with Staggered Wage and Price Contracts." *Journal of Monetary Economics* 46(2): 281–313.

Francis, Neville and Valerie A. Ramey, 2009. "Measures of per Capita Hours and Their Implications for the Technology-Hours Debate." *Journal of Money, Credit and Banking* 41(6): 1071–1097.

Galí, Jordi, 1999. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review* 89(1): 249–71.

Galí, Jordi and Mark Gertler, 1999. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics* 44(2): 195–222.

Galí, Jordi, Mark Gertler, and J. David López-Salido, 2007. "Markups, Gaps, and the Welfare Costs of Business Fluctuations." *Review of Economics and Statistics* 89(1): 44–59.

Gómez, Victor and Agustín Maravall, 1996. "Programs TRAMO and SEATS: Instructions for the User." Working paper 9628, Banco de España.

Goodfriend, Marvin and Robert G. King, 1997. "The New Neoclassical Synthesis and the Role of Monetary Policy." In *NBER Macroeconomics Annual 1997*, edited by Ben S. Bernanke and Julio Rotemberg, pp. 231–83. MIT Press.

Gourio, François and Anil K. Kashyap, 2007. "Investment Spikes: New Facts and a General Equilibrium Exploration." *Journal of Monetary Economics* 54(Supplement 1): 1–22.

Hall, Robert E., 1980. "Employment Fluctuations and Wage Rigidity." *Brookings Papers on Economic Activity* 1: 91–123.

Hall, Robert E., 2004. "Measuring Factor Adjustment Cost." *Quarterly Journal of Economics* 119(3): 899–927.

Hamermesh, Daniel S., 2006. "Overtime Laws and the Margins of Work Timing." Unpublished paper, University of Texas, Austin.

Hamermesh, Daniel S. and Gerard A. Pfann, 1996. "Adjustment Costs in Factor Demand." *Journal of Economic Literature* 34(3): 1264–92.

Hirsch, Barry T., 2005. "Why Do Part-Time Workers Earn Less? The Role of Worker and Job Skills." *Industrial and Labor Relations Review* 58(4): 525–51.

Jaimovich, Nir and Max Floetotto, 2008. "Firm Dynamics, Markup Variations, and the Business Cycle." *Journal of Monetary Economics* 55(7): 1238–52.

Khan, Aubhik and Julia K. Thomas, 2007. "Inventories and the Business Cycle: An Equilibrium Analysis of (S,s) Policies." *American Economic Review* 97(4): 1165–88.

Nekarda, Christopher J. and Valerie A. Ramey, 2010. "Industry Evidence on the Effects of Government Spending." Working paper 15754, National Bureal of Economic Research.

Perotti, Roberto, 2008. "In Search of the Transmission Mechanism of Fiscal Policy." In *NBER Macroeconomics Annual 2007*, edited by Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, pp. 169–226. University of Chicago Press.

Ramey, Valerie A., 1991. "Discussion of 'Markups and the Business Cycle'." In *NBER Macroeconomics Annual 1991*, edited by Olivier Jean Blanchard and Stanley Fischer, pp. 134–40. MIT Press.

Ramey, Valerie A., 2009. "Identifying Government Spending Shocks: It's All in the Timing." Unpublished paper, University of California, San Diego.

Ravenna, Federico and Carl E. Walsh, 2006. "Optimal Monetary Policy with the Cost Channel." *Journal of Monetary Economics* 53(2): 199–216.

Rotemberg, Julio J. and Michael Woodford, 1991. "Markups and the Business Cycle." In *NBER Macroeconomics Annual 1991*, edited by Olivier Jean Blanchard and Stanley Fischer, pp. 63–129. MIT Press.

Rotemberg, Julio J. and Michael Woodford, 1992. "Oligopolistic Pricing and the Effects of Aggregate Demand on Economic Activity." *Journal of Political Economy* 100(6): 1153–207.

Rotemberg, Julio J. and Michael Woodford, 1999. "The Cyclical Behavior of Prices and Costs." In *Handbook of Macroeconomics*, edited by John B. Taylor and Michael Woodford, pp. 1051–135. Elsevier.

Sbordone, Argia, 2002. "Price and Unit Labor Costs: A New Test of Price Stickiness." *Journal of Monetary Economics* 49(2): 265–92.

Shea, John, 1993. "Do Supply Curves Slope up?" *Quarterly Journal of Economics* 108(1): 1–32.

Smets, Frank and Rafael Wouters, 2003. "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area." *Journal of the European Economic Association* 1(5): 1123–75.

Smets, Frank and Rafael Wouters, 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97(3): 586–606.

Stock, James H. and Mark W. Watson, 1998. "Median Unbiased Estimation of Coefficient Variance in a Time-Varying Parameter Model." *Journal of the American Statistical Association* 93(441): 349–58.

Taylor, Daniel E. and Edward S. Sekscenski, 1982. "Workers on Long Schedules, Single and Multiple Jobholders." *Monthly Labor Review* 105(5): 47–53.

Trejo, Stephen J., 1991. "The Effects of Overtime Pay Regulation on Worker Compensation." *American Economic Review* 81(4): 719–40.

Wetzel, James R., 1966. "Overtime Hours and Premium Pay." *Monthly Labor Review* 89(9): 973–7.

Woodford, Michael, 2003. *Interest and Prices: Foundations of A Theory of Monetary Policy*. Princeton, N.J.: Princeton University Press.

Table 1. Cyclicality of Price-Cost Markup

		Correlation with real GDP[b]	
Markup Measure	*Recession indicator*[a]	*HP filter*	*First difference*
Average markup			
1. Nonfinancial corporate business	-0.020^{***}	0.279^{***}	0.325^{***}
2. Private business (NIPA)	-0.022^{***}	0.434^{***}	0.352^{***}
3. Private business (BLS)	-0.020^{***}	0.346^{***}	0.496^{***}
4. Manufacturing	-0.037^{***}	0.396^{***}	0.405^{***}
Marginal Markup[c]			
5. Private business, $\rho = 0.25$	-0.022^{***}	0.326^{***}	0.545^{***}
6. Private business, $\rho = 0.50$	-0.024^{***}	0.351^{***}	0.550^{***}
7. Manufacturing, $\rho = 0.25$	-0.039^{***}	0.238^{***}	0.308^{***}
8. Manufacturing, $\rho = 0.50$	-0.034^{***}	0.051	0.210^{***}
Marginal Markup, CES production function[d]			
9. HP-filtered technology[e]	-0.039^{***}	0.390^{***}	0.633^{***}
10. SVAR technology[f]	-0.034^{***}	0.410^{***}	0.644^{***}

Source: Authors' calculations using quarterly data from the BLS and NIPA. Average markup data are for 1947:1–2009:4; marginal markup data are for 1960:2–2009:4 in aggregate economy and for 1956:1–2009:4 in manufacturing.

Notes: a. Reports $\widehat{\beta}_3$ from regression $\mu_t = \beta_0 + \beta_1 t + \beta_2 t^2 + \beta_3 \mathrm{I}(\mathrm{rec}_t) + \omega_t$ (equation 15), where μ_t is the logarithm of the markup, $\mathrm{I}(\mathrm{rec}_t)$ is an indicator for a recession in quarter t, and ω_t is the error term. A recession begins in the quarter following the business cycle peak, as determined by the NBER, and ends in the quarter of the trough. Standard errors (not reported) incorporate a Newey-West adjustment, allowing for 2 lags in the autocorrelation structure.

b. Contemporaneous correlation of cyclical components of log real GDP and log markup, $\mathrm{corr}(y_t, \mu_t)$. For HP filter, detrended using $\lambda = 1,600$. Significance calculated using the Fisher transformation.

c. Markup measure is BLS private business. ρ is effective overtime premium.

d. Elasticity of substitution between capital and labor $\sigma = 0.5$; see equation 12. Markup measure is BLS private business with ρ=0.25.

e. Technology is the HP trend in labor productivity.

f. Technology is identified using a bivariate VAR with labor productivity and hours per capita in first-differences.

Table 2. Effect of Time and Industry Aggregation on Slope of $\mathrm{d}v/\mathrm{d}h$, 1956–83

Coefficient $\hat{\eta}$

	Change in average hours	
Frequency	*39 to 41*	*36 to 43*
Disaggregated		
1. Monthly	0.12	0.26
2. Quarterly	0.11	0.33
3. Annual	0.24	0.81
Aggregated		
4. Monthly	0.15	
5. Quarterly	0.18	
6. Annual	1.53	

Source: Authors' regressions using 2-digit CES manufacturing data.

Notes: Reports coefficient on Δh from regression $\Delta v_{it} = \left\{ b_{i0} + b_{i1}t + b_2 t^2 + b_3 t^3 + c_1 \left[h_{i(t-1)} - 40 \right] + c_2 \left[h_{i(t-1)} - 40 \right]^2 + c_3 \left[h_{i(t-1)} - 40 \right]^3 \right\} \Delta h_{it} + a_{i0} + a_1 t + a_2 t^2 + a_3 t^3 + d_{i1} \ln \left[N_{it}/N_{i(t-1)} \right] + d_{i2} \Delta \ln \left[N_{it}/N_{i(t-1)} \right] + e_{it}$ (equation 16). Data are annual and cover 1956–83.

Table 3. Effect of Time Aggregation on Markup, 50 Percent Overtime Premium

		Correlation of markup with		
Frequency of $\mathrm{d}v/\mathrm{d}h$	*Frequency of markup*	*Real GDP*	*Industrial production*	*Total hours*
2-digit industry data, 1956–83				
1. Quarterly	Quarterly	0.307	0.140	0.069
2. Quarterly	Annual	0.200	0.010	−0.047
3. Annual	Annual	−0.004	−0.205	−0.245
2-digit industry data, 1956–2002				
4. Annual	Annual	0.011	−0.049	−0.084
Aggregate manufacturing data, 1956–2009				
5. Quarterly	Quarterly	0.052	0.008	−0.051

Source: Authors' calculations using 2-digit CES manufacturing data.

Notes: Contemporaneous correlation of cyclical components of log markup and cyclical indicator, where cyclical component is extracted using HP filter. Industrial production and total hours are for manufacturing.

Table 4. Regression of Markup on Gross Shipments

Coefficient $\widehat{\gamma}_1$

Specification	*OLS* *(1)*	*OLS-FE* *(2)*	*IV-FE* *(3)*
Production workers	−0.133	0.052	0.021
	(0.101)	(0.012)	(0.048)
All workers	−0.134	0.175	0.071*
	(0.095)	(0.111)	(0.043)

Source: Author's regressions using data from MID and BEA benchmark IO account.

Notes: Regression of $\Delta\mu_{it} = \gamma_{0it} + \gamma_1 \Delta \ln Y_{it} + \varepsilon_{it}$ (equation 20) for industry i in year t. Average markup is given by equation 19. Sample contains 272 4-digit SIC industries from 1960–96 for a total of 10,135 observations. FE indicates inclusion of industry and year fixed effects. IV regressions instrument $\Delta \ln Y_{it}$ by Δg_{it}; see section 7.1. *** indicates significance at 1-percent, ** at 5-percent, and * at 10-percent level.

Figure 1. Average Weekly Hours per Worker, All Workers

Source: Authors' calculations using data from *Employment and Earnings* and the CPS covering 1960–2009.

Notes: Quarterly average of seasonally-adjusted monthly data. Shaded areas indicate recessions as defined by the NBER.

Figure 2. Average Weekly Hours per Worker, Manufacturing Production Workers

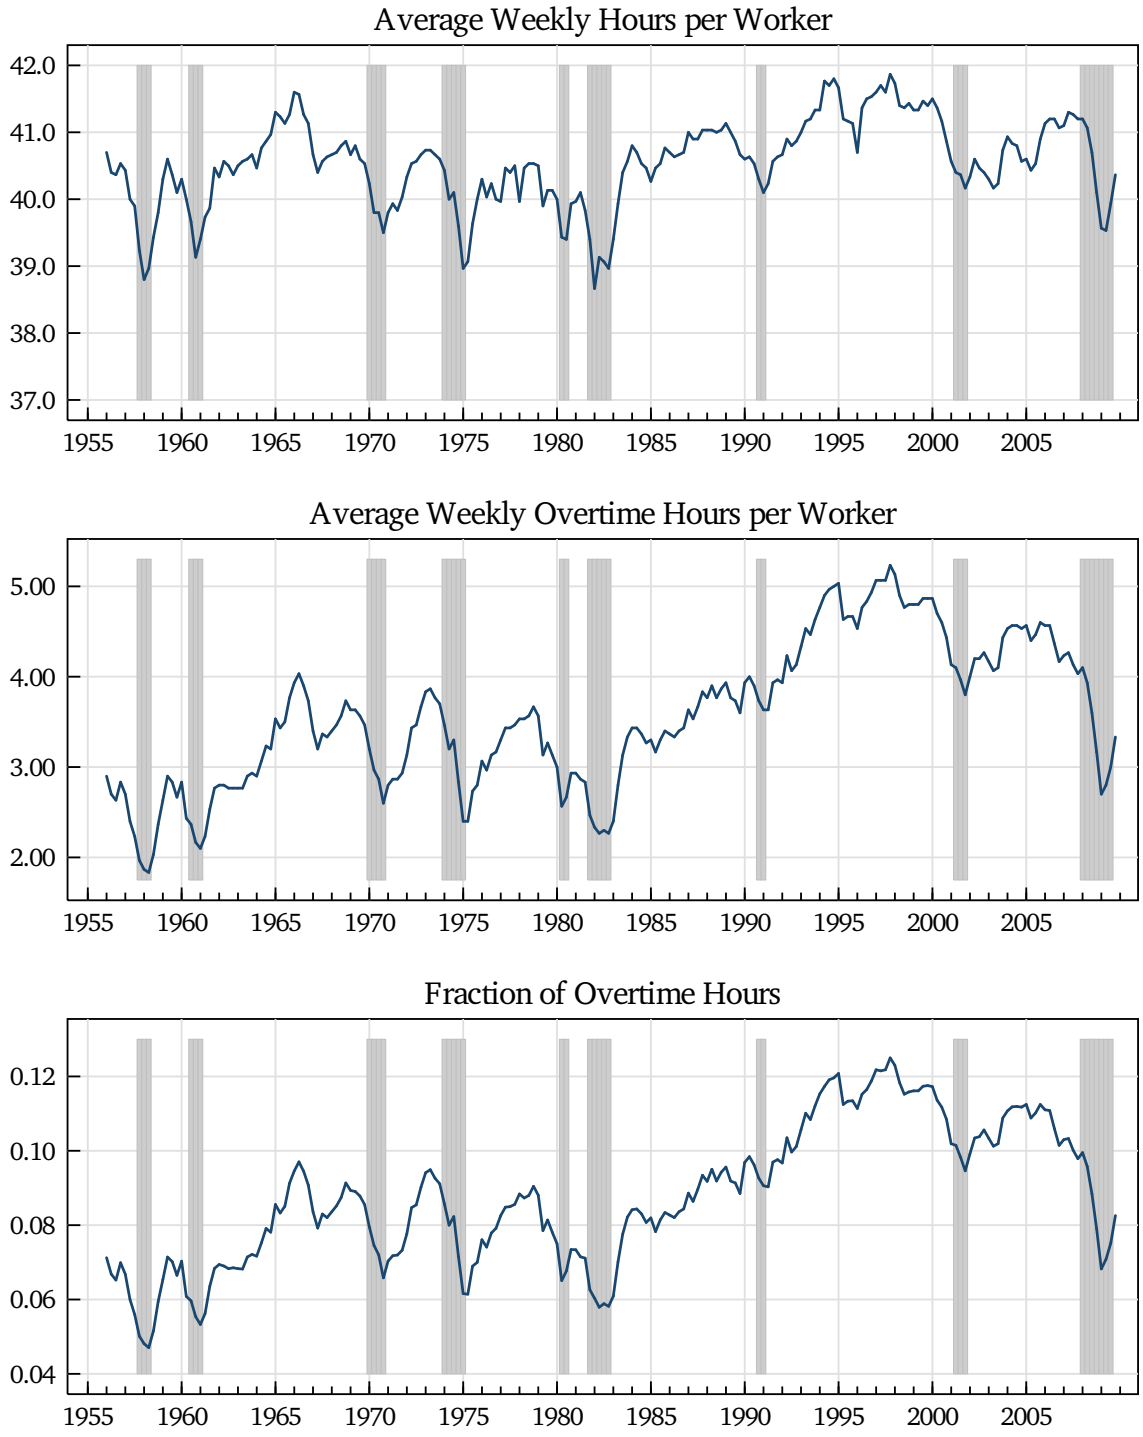

Source: Authors' calculations using CES data covering 1956–2009.

Notes: Quarterly average of seasonally-adjusted monthly data. Shaded areas indicate recessions as defined by the NBER.

Figure 3. Estimated Relationship between Change in Overtime Hours and Change in Average Hours

Coefficient $\widehat{\eta}_t$

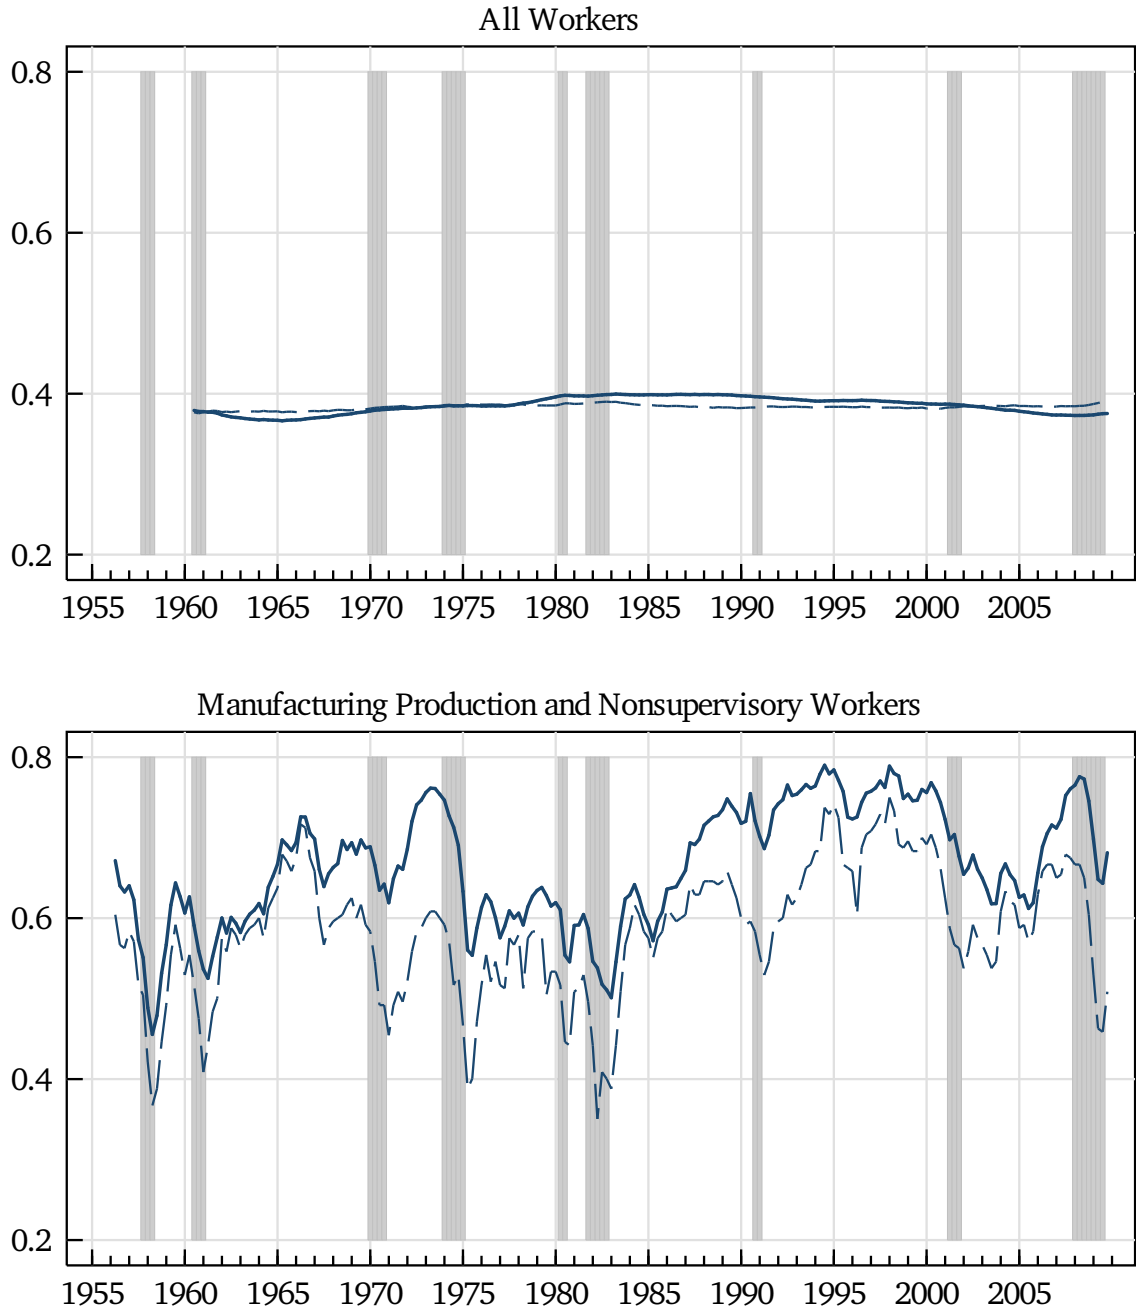

Source: Authors' regressions. Data for all workers are from *Employment and Earnings* and cover 1960:2–2009:4; data for manufacturing production workers are from the CES and cover 1956:1–2009:4.

Notes: $\eta_t = \beta_t + \gamma(h_{t-1} - 40)$ is estimated from the regression $\Delta v_t = \alpha + \Delta h_t[\beta_t + \gamma(h_{t-1} - 40)] + \epsilon_t$. Shaded areas indicate recessions as defined by the NBER.

Figure 4. Fraction of Overtime Hours Worked Paid a Premium

Source: Authors' calculations using data from May CPS extracts (NBER) and Employer Costs for Employee Compensation (BLS).

Notes: The implied θ for the early sample is based on individual worker reports on hours and whether they are paid a premium from the May CPS extract. The implied θ for the later sample is based on aggregated data on wages and salaries and overtime compensation from the Employer Cost survey, coupled with our constructed measure of v/h.

Figure 5. Marginal-Average Wage Adjustment Factor

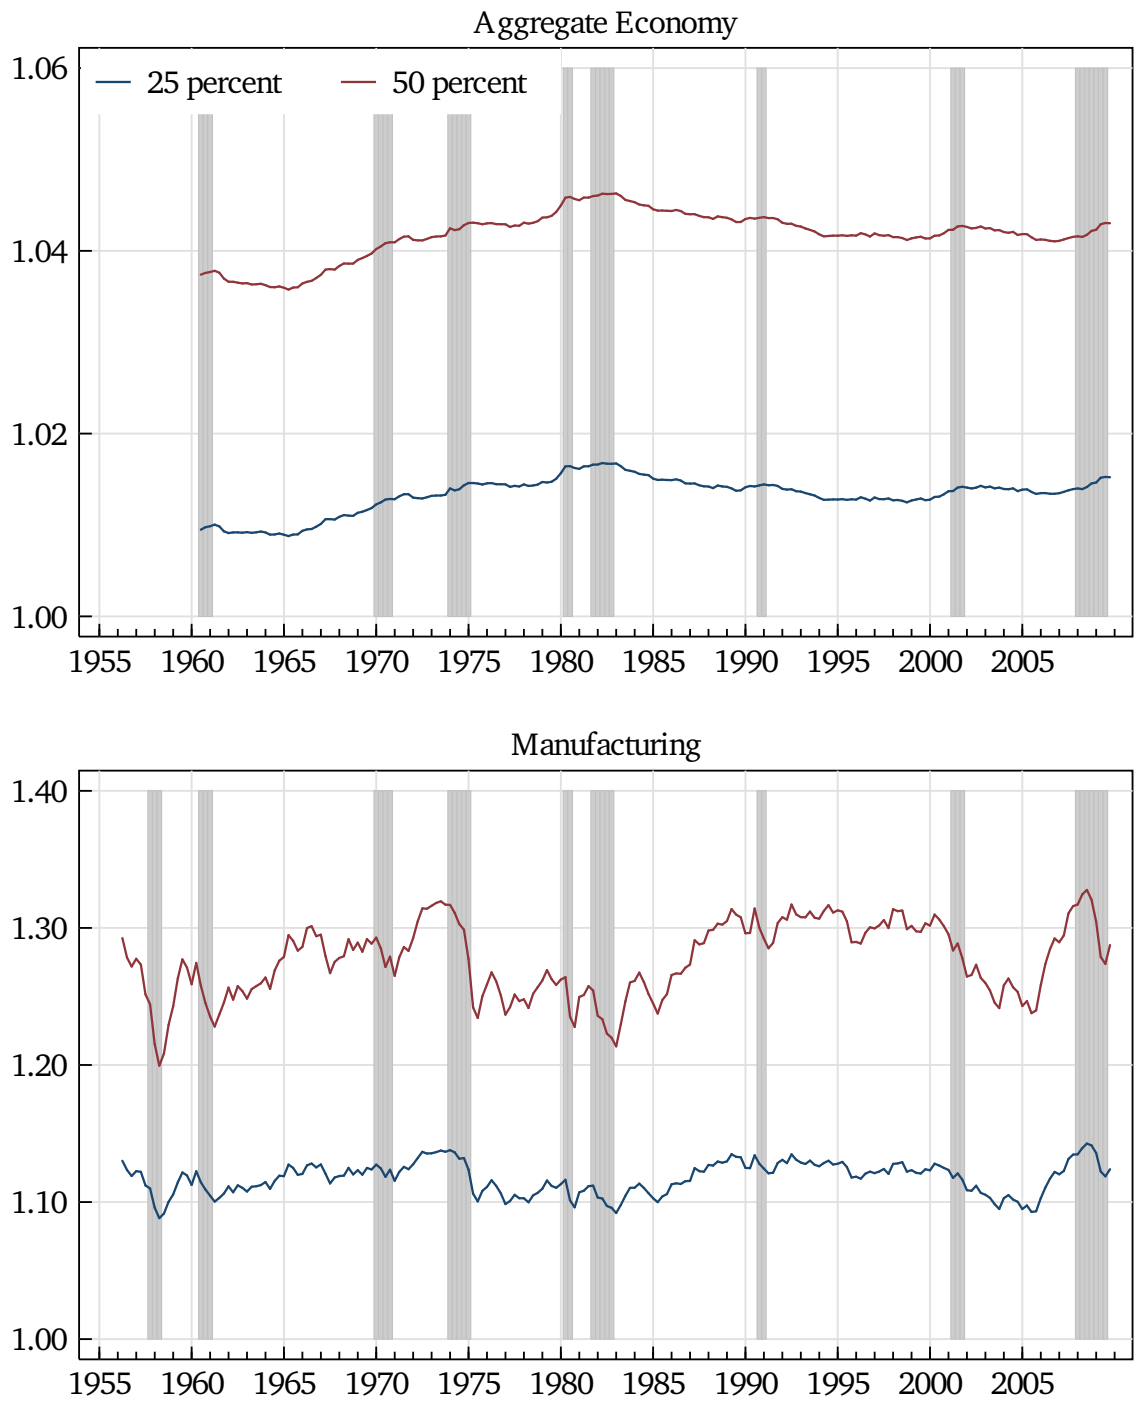

Source: Authors' calculations using quarterly data from the BLS. Data for aggregate economy are from *Employment and Earnings* and cover 1960:3–2009:4; data for manufacturing are from the CES and cover 1956:1–2009:4.

Notes: Adjustment factor is $\frac{W_M}{W_A}{}_t = \frac{1+\rho(dv/dh)_t}{1+\rho(v/h)_t}$; see equation 7. Shaded areas indicate recessions as defined by the NBER.

Figure 6. Aggregate Price-Cost Markup

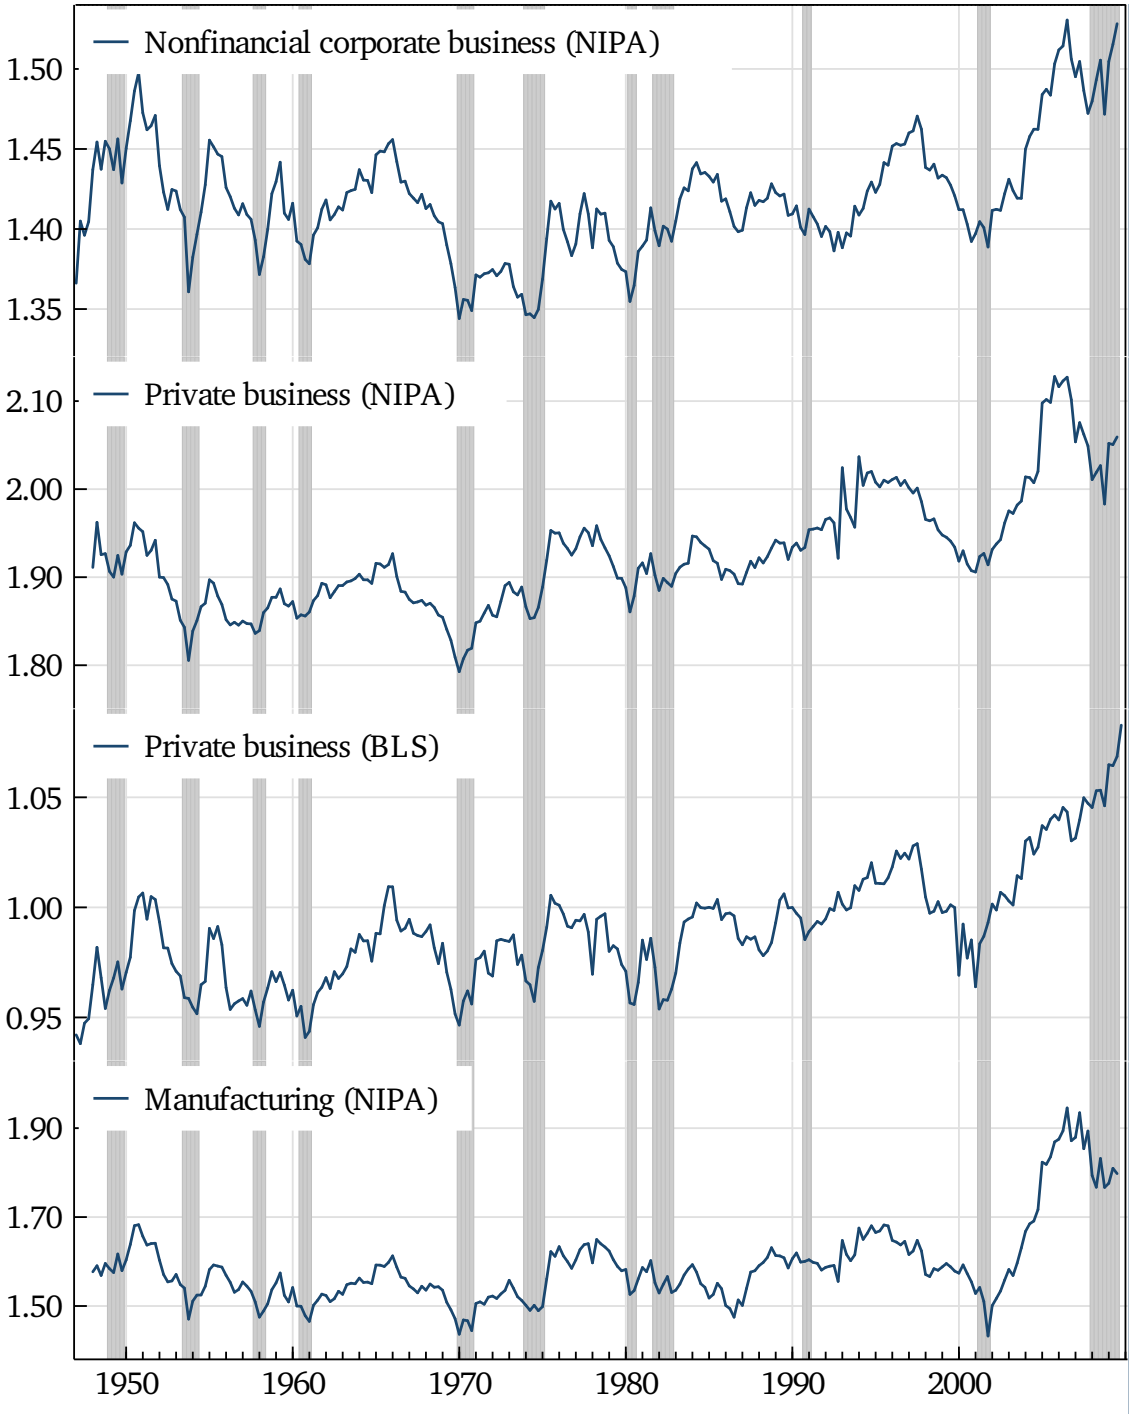

Source: Authors' calculations using quarterly data from the NIPA and the BLS.

Notes: Markup in nonfinancial corporate business is compensation divided by gross value added less taxes on production; other NIPA markups are wage and salary disbursements divided by income without capital consumption adjustment. BLS markup is 100/index of labor share. Shaded areas indicate recessions as defined by the NBER.

Figure 7. Marginal Price-Cost Markup, Aggregate Economy

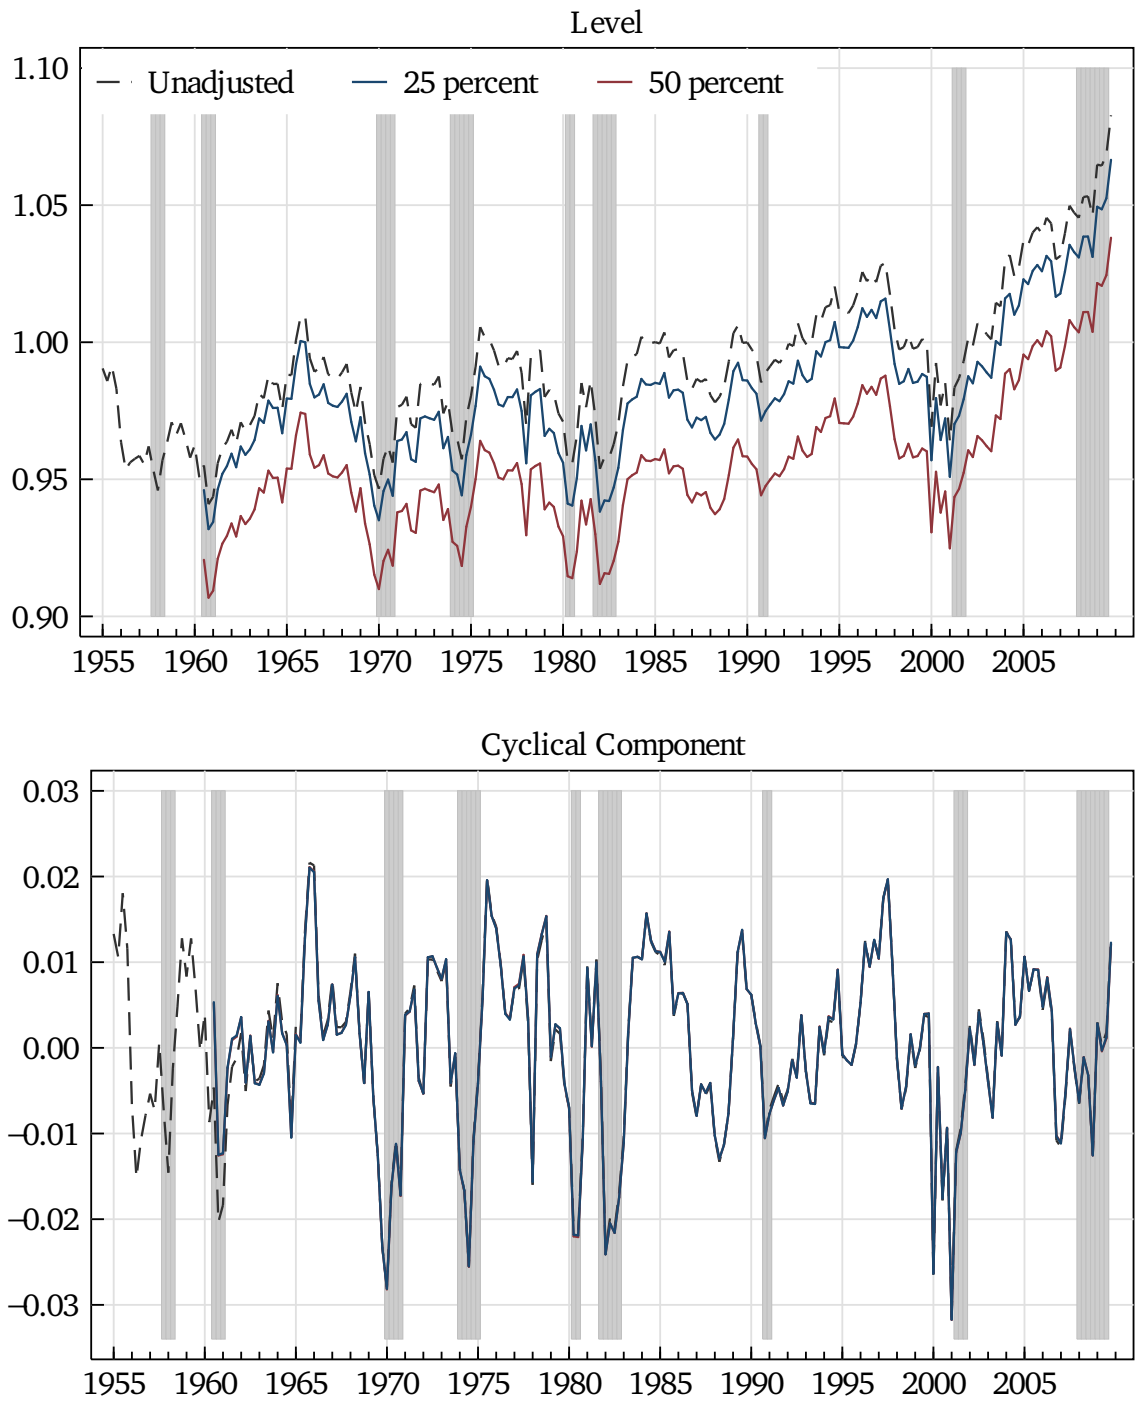

Source: Authors' calculations using quarterly BLS data on labor share for private business.

Notes: Unadjusted plots average markup; remaining series are marginal markup with indicated overtime premium. Cyclical component extracted using HP filter ($\lambda = 1,600$). Shaded areas indicate recessions as defined by the NBER.

Figure 8. Cross-Correlations of Markups with Real GDP

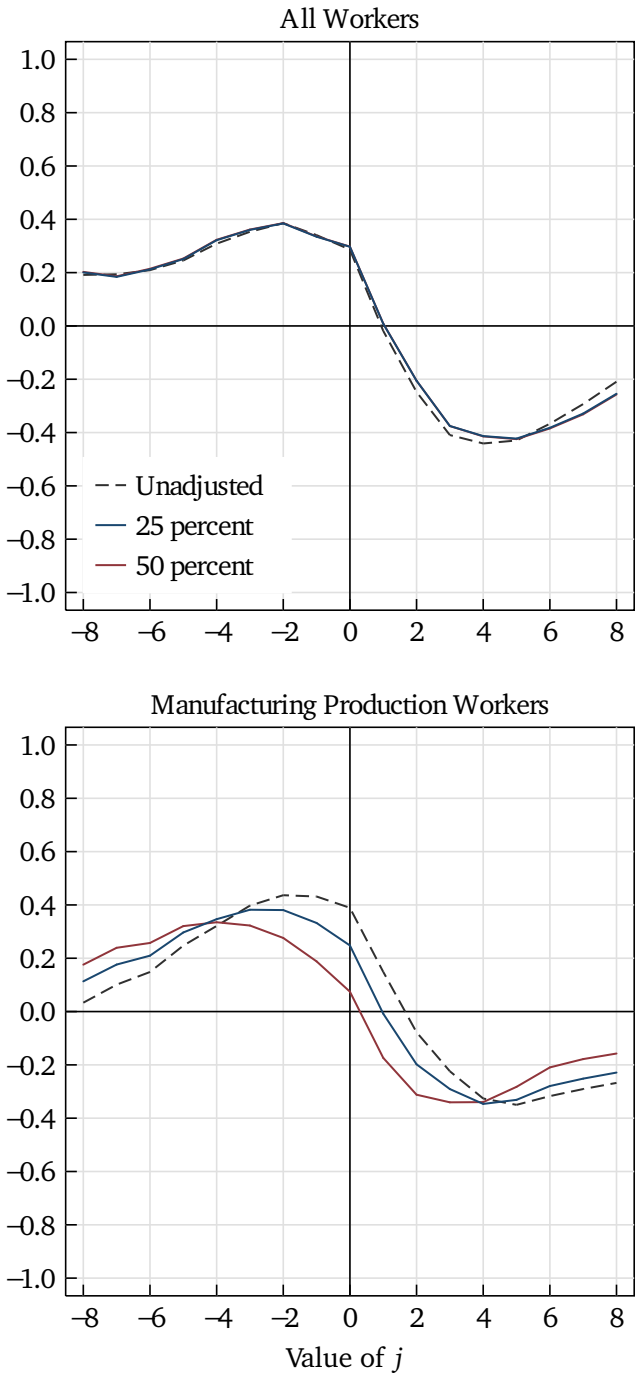

Source: Authors' calculations using quarterly data. Markup for aggregate economy is inverse labor share for private business over 1960:3–2009:4 from the BLS; markup for manufacturing is inverse labor share over 1956:1–2009:4 from the NIPA; real GDP is from the NIPA.

Notes: Correlation of cyclical components of y_t and μ_{t+j}. Detrended using HP filter ($\lambda = 1,600$). Unadjusted is average markup; remaining series are marginal markup with indicated overtime premium.

Figure 9. Marginal Price-Cost Markup, Manufacturing

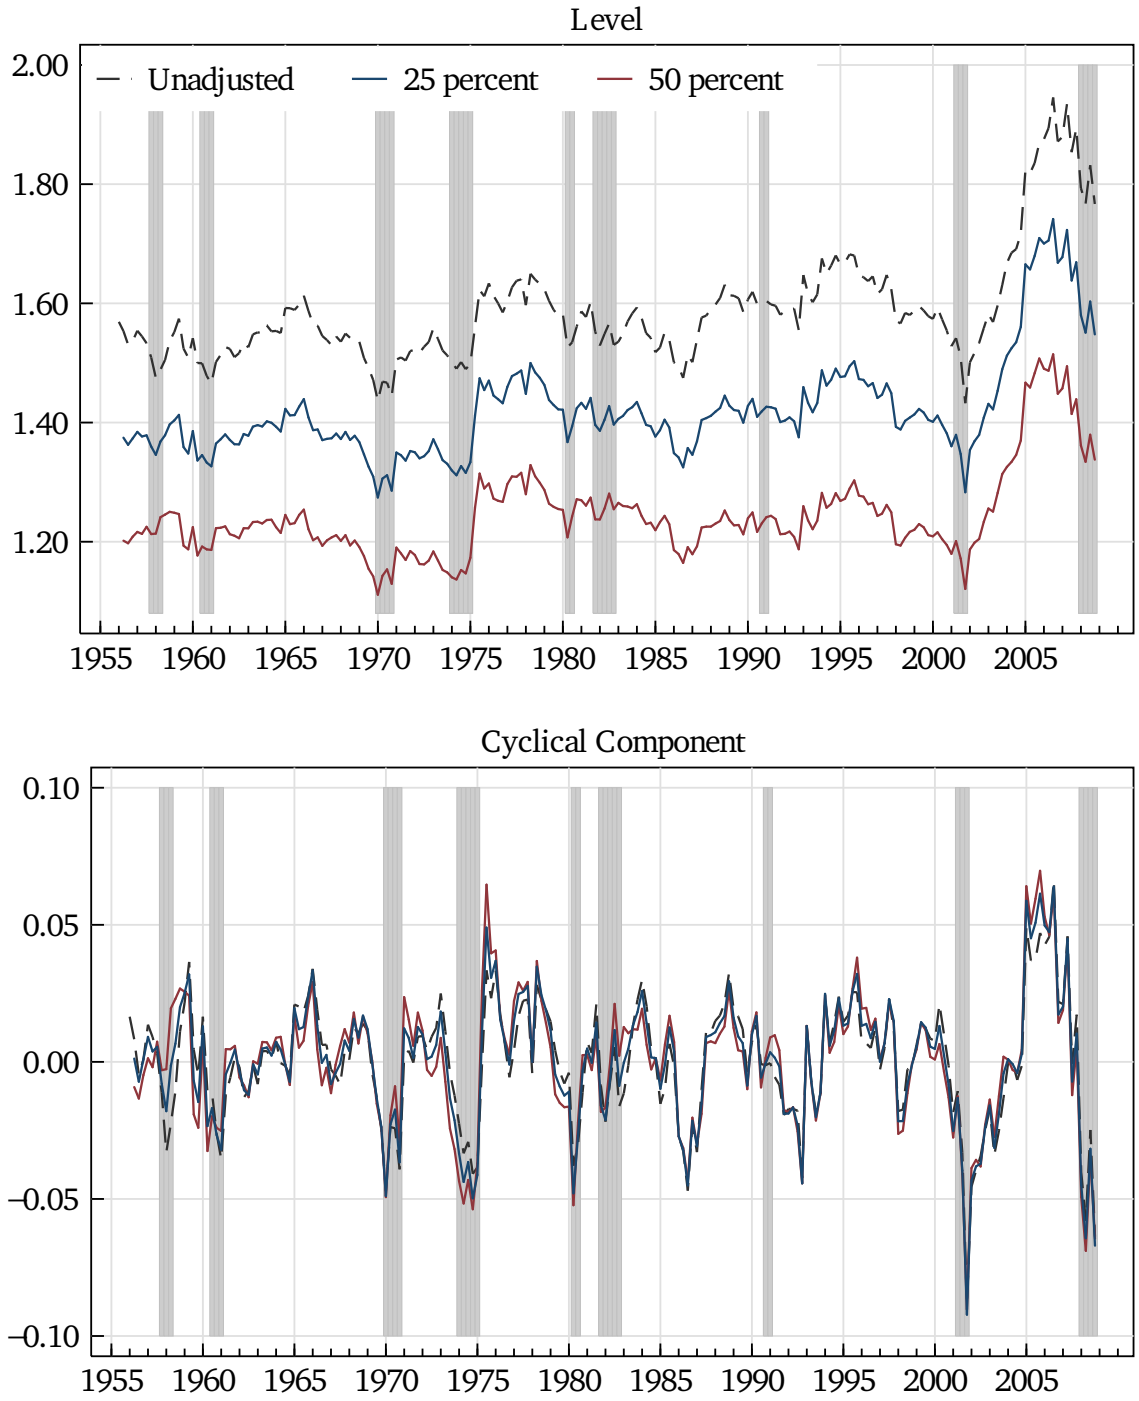

Source: Authors' calculations using quarterly NIPA data on labor share for manufacturing.

Notes: Unadjusted plots average markup; remaining series are marginal markup with indicated overtime premium. Cyclical component extracted using HP filter ($\lambda = 1,600$). Shaded areas indicate recessions as defined by the NBER.

Figure 10. Response of Markup to a Contractionary Monetary Policy Shock

100 × coefficient

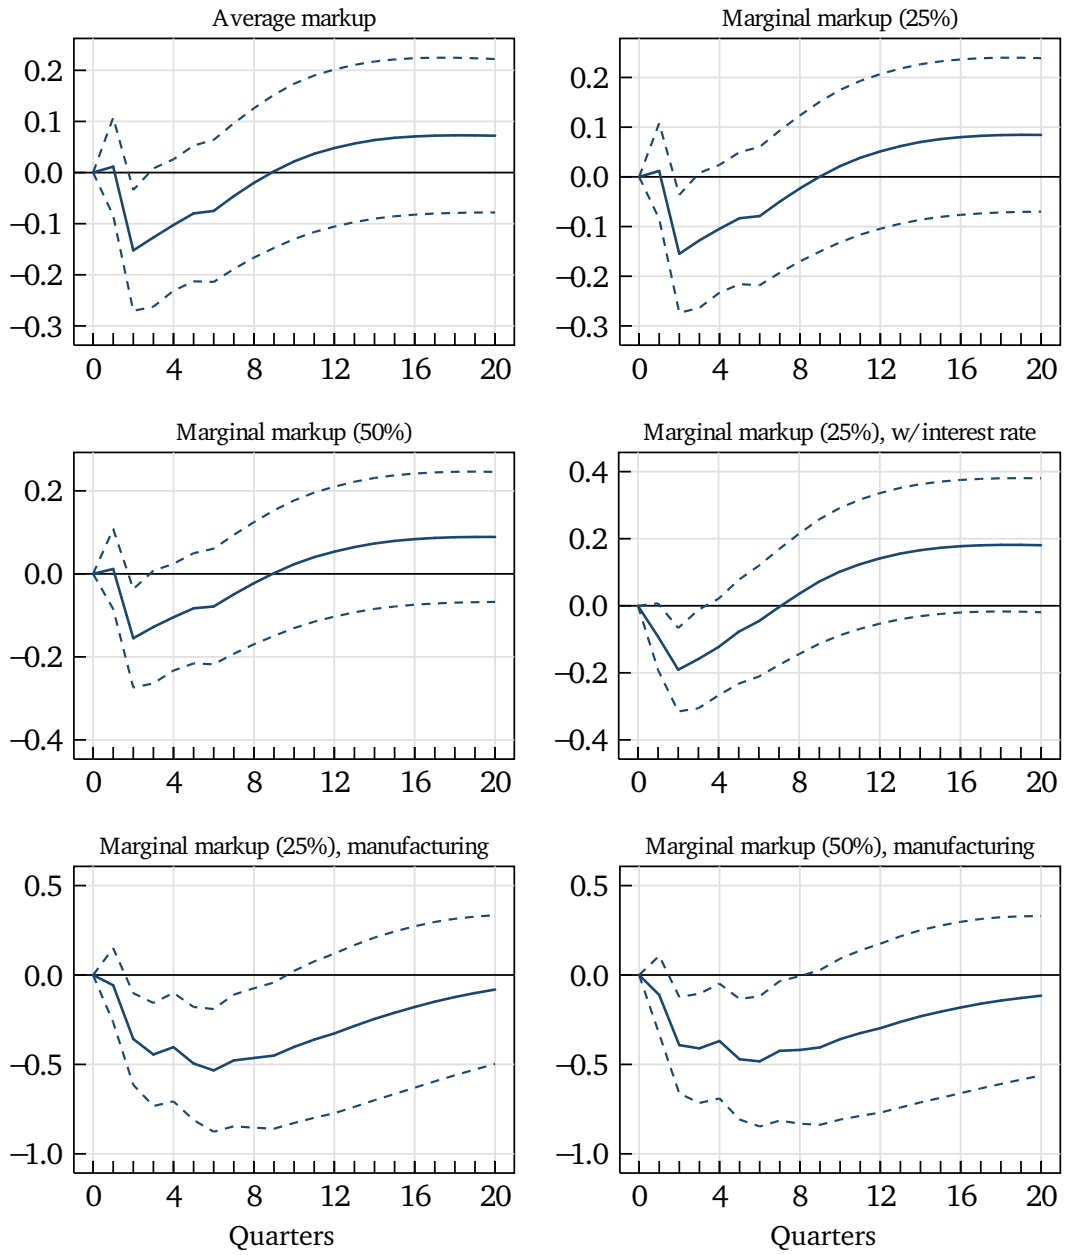

Source: Authors' calculations using quarterly data for 1960:3–2009:4. Markup is inverse of labor share for private business from the BLS; real GDP and GDP deflator are from the NIPA; commodity prices are from the BLS; federal funds rate and prime rate are from the Board of Governors of the Federal Reserve System.

Notes: Impulse responses estimated from VAR(4) with log real GDP, log commodity prices, log GDP deflator, markup measure, and federal funds rate; also includes a linear time trend. Monetary policy shock identified as shock to federal funds rate when ordered last. Specification with interest rate includes the prime rate in marginal cost. Dashed lines indicate 95-percent confidence interval.

Figure A1. March and Annual Average Hours per Bin

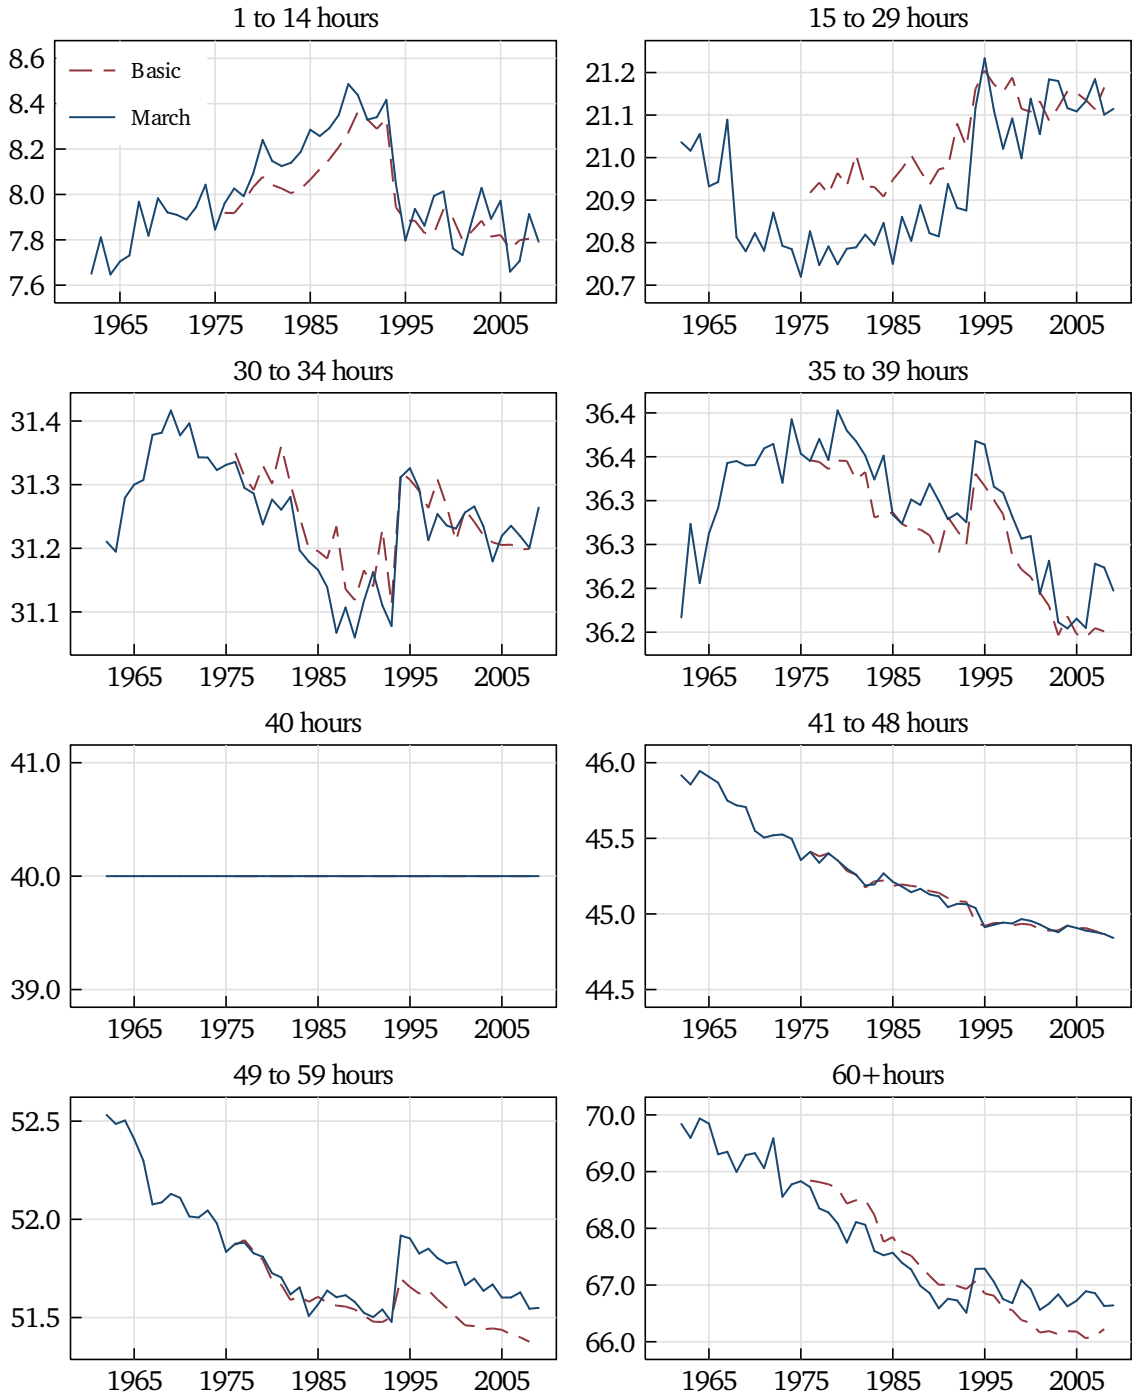

Source: Authors' calculations using monthly Basic CPS data for 1976–2009 and March CPS data for 1962–2009.

Notes: Annual average of monthly CPS data.