NBER WORKING PAPER SERIES

THE CYCLICAL BEHAVIOR OF THE PRICE-COST MARKUP

Christopher J. Nekarda
Valerie A. Ramey

Working Paper 19099
http://www.nber.org/papers/w19099

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2013

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 Susanto Basu, Mark Bils, Larry Christiano, Olivier Coibion, Steve Davis, Davide Debortoli, Martin Eichenbaum, Robert Hall, Garey Ramey, Sergio Rebelo, Harald Uhlig, Raf Wouters, and participants at numerous seminars for helpful comments. Ben Backes and Myungkyu Shim provided excellent research assistance. Valerie Ramey gratefully acknowledges financial support from National Science Foundation grant SES-0617219 through the National Bureau of Economic Research. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

The Cyclical Behavior of the Price-Cost Markup
Christopher J. Nekarda and Valerie A. Ramey
NBER Working Paper No. 19099
June 2013
JEL No. E32,L16

ABSTRACT

A countercyclical markup of price over marginal cost is the key transmission mechanism for demand shocks in textbook New Keynesian (NK) models. This paper re-examines the foundation of those models. We study the cyclicality of markups in the private economy as well as in detailed manufacturing industries. First, we show that frameworks for measuring markups that have produced the strongest evidence for countercyclicality produce the opposite result when we substitute new methods and data. Second, because the NK model's predictions differ by the nature of the shock, we present evidence on the cyclicality of the markup conditional on various types of shocks. Consistent with the NK model, we find that markups are procyclical conditional on a technology shock. However, we find that they are either procyclical or acyclical conditional on demand shocks. Thus, the textbook NK explanation for the effects of government spending or monetary policy is not supported by the behavior of the markup.

Christopher J. Nekarda
Federal Reserve Board of Governors
Research and Statistics Division
20th St. & C St. NW
Washington, DC 20551-0001
Christopher.J.Nekarda@frb.gov

Valerie A. Ramey
Department of Economics, 0508
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093-0508
and NBER
VRAMEY@UCSD.EDU

> 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 textbook 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; Woodford, 2003). This result also holds in the leading New Keynesian models with both sticky prices and sticky wages, such as Erceg, Henderson and Levin (2000); Smets and Wouters (2003, 2007); Christiano, Eichenbaum and Evans (2005). In Jaimovich and Floetotto's (2008) model, procyclical entry of firms leads to countercyclical markups, and hence to procyclical movements in measured productivity.

The dependence of Keynesian models on countercyclical markups is a feature only of the models formulated since the early 1980s. From the 1930s through the 1970s, the Keynesian model was founded on the assumption of sticky wages (such as Keynes, 1936; Phelps, 1968; Taylor, 1980). Some researchers believed that the implications of this model were at odds with the cyclical properties of real wages, leading to a debate known as the "Dunlop-Tarshis" controversy.[1] In response to the perceived disparity between the data and predictions of the traditional Keynesian model, the literature shifted in the early 1980s to the assumption of sticky prices rather than sticky wages (such as Gordon, 1981; Rotemberg, 1982). This type of model emerged as the leading textbook New Keynesian model. Virtually all current New Keynesian models incorporate the notion that markups fall in response to positive demand shifts.

Estimating the cyclicality of markups is one of the more challenging measurement is-

1. In fact, Dunlop (1938) and Tarshis (1939) were repeatedly misquoted by the literature as showing that real wages were procyclical. Neither of them showed this. Both authors showed that money wages and real wages were positively correlated, and Tarshis went on to show that real wages were in fact negatively correlated with aggregate employment. Dunlop (1998) discusses the debate in his retrospective article.

sues in macroeconomics. Theory prescribes a comparison of price and marginal cost; however, available data typically include only average cost. As we will discuss, researchers have used a variety of techniques to measure markups directly to assess their cyclicality, or have inferred the cyclicality of markups using indirect evidence. Despite the very mixed results in this literature, most current research has proceeded as if there were firm evidence of a countercyclical markup.

In this paper, we present evidence that most measures of the markup are procyclical or acyclical. After developing the theoretical framework for measuring markups, we first analyze the cyclicality of markups for the broad economy. We find that markups are generally procyclical or acyclical, hitting troughs during recessions and reaching peaks in the middle of expansions. These results hold both for the baseline measure in which markups are inversely proportional to labor share, as well as the generalizations advocated by Bils (1987) and Rotemberg and Woodford (1999). Moreover, both monetary shocks and government spending shocks cause markups and output to comove positively. We then turn to the analysis of markups in detailed manufacturing industry data. We show that markups are procyclical unconditionally, as well as conditional on technology shocks. Conditional on demand shocks, markups are slightly procyclical or acyclical. Our results differ from some of those in the previous literature because we are able to use new data sources rather than relying on steady-state calibrations or proxy equations. Our new results raise questions about the basic propagation mechanism of the current versions of the New Keynesian model.

2 Relationship to the Literature

Industrial organization economists have a long history of studying the cyclicality of price-cost margins. Macroeconomists only began studying this issue in the mid-1980s when macro models started to emphasize price setting behavior of firms. Four principal methods have been used to measure markups directly and two additional methods have been used to assess the cyclicality of markups indirectly.

The first of the direct methods uses the standard industrial organization concept of a price-cost margin constructed from revenues and average variable costs. Domowitz, Hubbard and Petersen (1986) use this method in a panel of four-digit Standard Industrial Classification (SIC) manufacturing industries and find that margins are significantly procyclical.

The second method builds on Hall's (1986) generalization of the Solow residual to estimate the cyclicality of markups. For example, Haskel, Martin and Small (1995) extend Hall's framework to allow for time-varying markups and apply it to a panel of two-digit U.K. manufacturing industries. They find that markups are markedly procyclical. Marchetti (2002) applies a similar framework to two-digit manufacturing industries in Italy. He finds no clear pattern of cyclicality of markups; in only 2 of 13 industries does he find consistent evidence across specifications of countercyclical markups.

The third method uses generalized production functions with quasi-fixed factors to estimate markups relative to marginal cost estimated from stochastic Euler equations. Using this type of approach, Morrison (1994) finds weakly procyclical markups in Canadian manufacturing, and Chirinko and Fazzari (1994) find acyclical or procyclical markups in firm-level data. Galeotti and Schianterelli (1998) test the Rotemberg and Saloner (1986) game-theoretic hypothesis and find that, consistent with this hypothesis, markups depend negatively on the current level of output but positively on the growth of output.

The fourth method is the only one that routinely finds evidence for countercyclical markups. This method uses the labor input margin to estimate marginal cost. Under standard assumptions, such as Cobb-Douglas production functions and no overhead labor, this method implies that the markup is inversely proportional to the labor share. Since the labor share is countercyclical, this measure of markups implies that markups are procyclical. Most of the papers using this method thus apply adjustments to the standard model to account for reasons why marginal labor costs might be more procyclical than average labor costs. For example, Bils (1987) argues that the marginal hourly wage paid to workers should be more procyclical than the average wage. He constructs a measure of marginal cost based on estimates of the marginal wage and finds that his markup series has a negative correlation with industry employment in a panel of two-digit industries, suggesting countercyclicality. Rotemberg and Woodford (1991), Rotemberg and Woodford (1999), Oliveira Martins and Scarpetta (2002) and Galí, Gertler and López-Salido (2007) apply additional adjustments to the standard model, such as substituting a constant elasticity of substitution (CES) production function for Cobb-Douglas and allowing for overhead labor.[2] Their applications of these adjustments typically convert procyclical markups (based on standard assumptions) into countercyclical markups.

The two indirect methods for assessing the cyclicality of the markup use entirely different frameworks. Bils and Kahn (2000) present a model of inventories and stockouts in

2. The appendix of Galí, Gertler and López-Salido (2007) gives a particularly clear and concise summary of the adjustments.

which the joint cyclicality of the ratio of sales to inventories and the discounted growth rate of output prices reveals the cyclicality of markups. They use this framework to conclude that markups are countercyclical in several two-digit U.S. manufacturing industries. Hall (2012) exploits standard advertising theory to show that countercyclical markups imply that advertising should also be highly countercyclical. He shows, in fact, that advertising is somewhat procyclical.

Overall, this literature has used a host of innovative and clever ways to measure markups and analyze their cyclicality. Most of the papers have tended to find procyclical or acyclical markups. The exceptions are those papers using the methods and adjustments advocated by Bils (1987) and Rotemberg and Woodford (1999) or the inventory research of Bils and Kahn (2000). Perhaps for this reason, many researchers have concluded that markups are countercyclical.

In this paper, we will revisit the method that has produced the strongest evidence of countercyclical markups and has been cited the most by the New Keynesian literature: the hours margin method employed by Bils (1987), Rotemberg and Woodford (1999), and Galí, Gertler and López-Salido (2007), among others. We improve on their implementation in several ways. First, because modern theories predict that markups should move in opposite directions in response to supply and demand shocks, we will analyze conditional cyclicality as well as unconditional cyclicality. Most of the early papers simply used output or shipments as their demand indicator. Second, because the importance of instrument relevance was only beginning to become apparent in the late 1980s and 1990s, many of the methods relied on estimates based on instruments with very low first-stage F statistics. We now know that even in large samples, instruments with low relevance can lead to very misleading results [3] In our industry analysis, we will use instruments with much better properties. Third, because of progress in data availability and computational capability, we will be able to use richer data, rather than relying on steady-state calibrations and proxy equations. Finally, as a robustness check, we will also conduct the Bils and Kahn (2000) inventory test on our industry data.

Thus, the contribution of this paper is essentially to revisit the empirical work behind the stylized facts that are at the foundation of modern New Keynesian models, but to do so with updated empirical methods and data.

3. See Bound, Jaeger and Baker (1995).

3 Theoretical Framework

We begin by presenting the theoretical framework that forms the basis of our main estimates of markups. The theoretical markup, $\mathcal{M}$, is defined as

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

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.[4]

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

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

subject to $\bar{Y} = F(ZhN, \ldots)$. W_{A} is the average hourly wage (which is potentially a function of average hours), N is the number of workers, Y is output, and Z is the level of labor-augmenting technology. Letting λ be the Lagrange multiplier on the constraint, we obtain the first-order condition for h as:

$$W'_{\mathrm{A}}(h)h + W_{\mathrm{A}}(h) = \lambda F_1(ZhN, \ldots)Z, \tag{3}$$

where W'_{A} is the derivative of the average wage with respect to h and F_1 is the derivative of the production function with respect to effective labor, ZhN. The multiplier λ is equal to marginal cost, so the marginal cost of increasing output by raising hours per worker is

4. 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.

given by:

$$MC = \lambda = \frac{W'_{\mathrm{A}} h + W_{\mathrm{A}}}{ZF_1(ZhN, \ldots)}. \tag{4}$$

The denominator of equation 4 is the marginal product of increasing hours per worker; the numerator is the marginal increase in the wage bill (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.

3.1 Production Function Assumptions

The formula for the markup above requires an estimate of the marginal product of labor, necessitating assumptions about the production function. Under the standard assumptions that the production function takes the form that output is Cobb-Douglas (denoted by a superscript "CD") in total hours and that the marginal wage is equal to the average wage (denoted by a subscript "A"), the markup is given by

$$\mathcal{M}_{\mathrm{A}}^{\mathrm{CD}} = \frac{P}{W_{\mathrm{A}} / \left[\alpha \left(Y/hN\right)\right]} = \frac{\alpha}{s}, \tag{5}$$

where α is the exponent on labor input in the production function and s is the labor share.

Rotemberg and Woodford (1999) note several reasons why the standard assumption of a production function that is Cobb-Douglas in total hours may lead to estimates of the markup that are biased toward being procyclical. We now consider the most plausible generalizations.

The first reason is overhead labor. In this generalization, the labor term in the production function is instead $\left(ZhN - Z\bar{h}\bar{N}\right)^{\alpha}$ where $\bar{h}\bar{N}$ represents overhead labor hours. In the Cobb-Douglas case, the markup is given by:

$$\mathcal{M}_{\mathrm{A}}^{\mathrm{CD,\,OH}} = \frac{P}{W_{\mathrm{A}} / \left[\alpha \left(Y/hN\right)\right]} = \frac{\alpha}{s'}, \tag{6}$$

where

$$s' = \frac{W_{\mathrm{A}} \left(hN - \bar{h}\bar{N}\right)}{PY}, \tag{7}$$

is the labor share of non-overhead labor.

A second generalization allows the elasticity of substitution between inputs to deviate from unity. For example, consider the following CES production function:

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

and σ is the elasticity of substitution between labor and capital. The derivative with respect to effective labor (the F_1 needed for equation 4) is

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

The exponent in equation 9 is the reciprocal of the elasticity of substitution between capital and labor. 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 this case, we obtain a markup:

$$\mathcal{M}_{\mathrm{A}}^{\mathrm{CES}} = \frac{P}{W_{\mathrm{A}}/\left[Z\alpha\,(Y/ZhN)^{\frac{1}{\sigma}}\right]} = \frac{\alpha}{s}\left(\frac{Y}{ZhN}\right)^{\frac{1}{\sigma}-1}. \tag{10}$$

This equation is derived assuming that output is measured as value added. When output is measured as gross output, we obtain the same expression for the markup as long as the production function is either (1) a generalized CES in which the elasticities of substitutions are equal across all inputs or (2) a nested CES in which σ is the elasticity of substitution between the labor input and a composite of the other inputs.

Rotemberg and Woodford (1999) and Galí, Gertler and López-Salido (2007) implement these two generalizations using log-linear approximations around a steady-state and then calibrating parameters based on zero profit conditions and assumptions on steady-state markups. As our alternative derivation shows, it is unnecessary to use approximations around a steady state. Instead, our equations require additional data and estimates of elasticities of substitution. We discuss the sources for these shortly.

3.2 Marginal vs. Average Labor Cost

The standard New Keynesian Phillips curve literature assumes that the average wage is the appropriate measure of the marginal increase in hours. Bils (1987) argues to the contrary 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. Following Bils, we capture this possibility by specifying the average wage as:

$$W_{\mathrm{A}}(h) = W_{\mathrm{S}}\left(1 + \rho\theta\frac{v(h)}{h}\right). \tag{11}$$

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.[5] Bils did not include the θ term in his specification because he used data for manufacturing from the BLS's establishment survey, in which overtime hours are *defined* as those hours commanding a premium (that is, $\theta = 1$). In our data, 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.

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. With this functional form, the marginal cost of increasing output by raising hours per worker is given by:

$$MC = \lambda = \frac{W_{\mathrm{S}}\left[1 + \rho\theta\left(\frac{\mathrm{d}v}{\mathrm{d}h}\right)\right]}{ZF_1(ZhN, \ldots)}. \tag{12}$$

Equation 12 makes it clear that the marginal cost of increasing hours per worker is not necessarily 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}:

$$W_{\mathrm{M}} = W_{\mathrm{S}}\left[1 + \rho\theta\left(\frac{\mathrm{d}v}{\mathrm{d}h}\right)\right]. \tag{13}$$

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 approximations to the marginal wage itself to substitute for marginal cost in his markup

5. 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.

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:

$$\frac{W_{\mathrm{M}}}{W_{\mathrm{A}}} = \frac{1 + \rho\theta\left(\frac{\mathrm{d}v}{\mathrm{d}h}\right)}{1 + \rho\theta\left(\frac{v}{h}\right)}. \tag{14}$$

This ratio can be used to convert the observed average wage to the theoretically-correct marginal wage required to estimate the markup. We show below that the ratio of overtime hours to average hours, v/h, is procyclical. Thus, the denominator in equation 14 is procyclical. How $W_{\mathrm{M}}/W_{\mathrm{A}}$ evolves over the business cycle depends on the relative cyclicality of $\mathrm{d}v/\mathrm{d}h$.

In the case where the wage is increasing in average hours, the markup in any of the previous formulations can be adjusted by multiplying W_{A} by $W_{\mathrm{M}}/W_{\mathrm{A}}$. For example in the Cobb-Douglas case, the markup is given by:

$$\mathcal{M}_{\mathrm{M}}^{\mathrm{CD}} = \frac{P}{W_{\mathrm{M}}/\left[\alpha\left(Y/hN\right)\right]} = \frac{\alpha}{s\left(W_{\mathrm{M}}/W_{\mathrm{A}}\right)}, \tag{15}$$

where we use equation 14 to convert average wages to marginal wages.

4 Overview of the Empirical Analysis

The rest of the paper uses the theory from the last section to derive new measures of the markup in order to assess cyclicality. We perform the empirical analysis on two data sets. One of the data sets covers the entire private economy and the other consists of a panel of four-digit SIC manufacturing industries. Each data set has advantages and disadvantages.

The aggregate data has the advantage of covering a much broader segment of the U.S. economy and having a higher frequency (quarterly). Moreover, we are able to use an auxiliary data set to calculate the factor needed to construct the theoretically correct marginal wage. In addition, the data are ideal for exploring the effects of aggregate monetary shocks and government spending shocks on markups. On the other hand, the aggregate data have several disadvantages: we only have measures of value added, not gross output; and the standard macroeconomic shocks we study have the usual low relevance as instruments.

Thus, we also explore the cyclicality of markups in annual four-digit manufacturing data from 1958 to 2009. Analysis of this data set has several advantages. First, because sectoral shifts might drive aggregate results, it is useful to examine the cyclicality of the markup at the disaggregated industry level. Second, the industry data allow us to use gross output rather than value-added output. As Waldmann (1991), Norrbin (1993) and Basu and Fernald (1997) argue, using value-added data can introduce errors in the measurements of markups. Third, we are able to construct highly relevant industry-specific instruments to identify demand and supply shocks.

The detailed industry data set has some disadvantages, though. First, the data are only available at the annual frequency, which masks some business cycle effects. Second, as we will discuss below, the data are not as well suited for estimating the marginal-average wage factor. Third, the manufacturing sector is not representative of the entire U.S. economy: Even at its post-World War II peak, manufacturing accounted for only 25 percent of employment; it now accounts for only 9 percent of employment.

5 Aggregate Analysis

5.1 Baseline Cobb-Douglas

As discussed in section 3, the markup is proportional to the inverse of the labor share if the production function is Cobb-Douglas. We first explore the markup based on three measures of the labor share covering several broad aggregates. Our first measure is the labor share in the private business sector published by the Bureau of Labor Statistics (BLS). This is the broadest aggregate measure and it covers all compensation in this sector; Galí, Gertler and López-Salido (2007) use the nonfarm business version of this measure. We also consider two measures of labor share for nonfinancial corporate business, constructed from the U.S. national income and product accounts (NIPA). The first includes total compensation of labor and the second includes only wages and salaries, excluding fringe benefits. Since parts of compensation might be considered more a fixed cost per worker than a payment per hour, we want to see whether the cyclicality changes noticeably when we switch to just wages and salaries. The nonfinancial corporate business measure using total compensation is the measure favored by Rotemberg and Woodford (1999), from the BLS. The appendix provides additional details.

Figure 1 displays measures of the baseline markup, as defined earlier in equation 5. The data are quarterly and extend from 1947 through 2012. The most salient characteristic of all three measures is the propensity to trough during a recession and to peak in the middle

Figure 1. Aggregate Price-Cost Markup

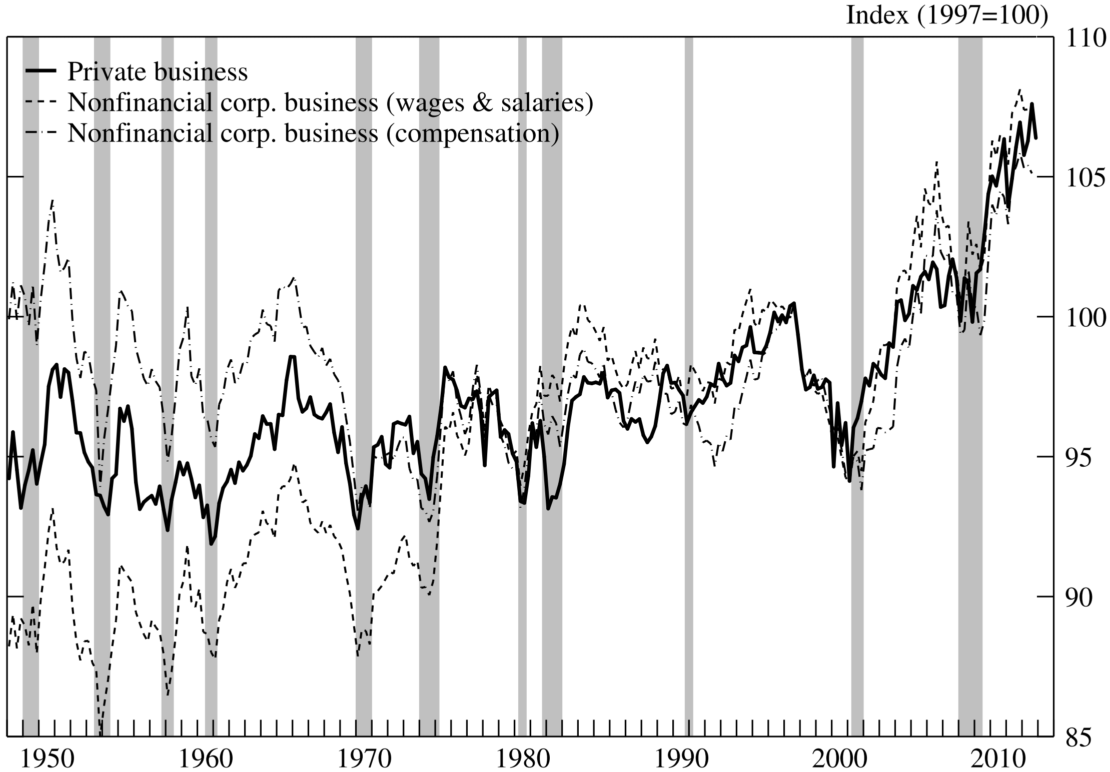

Source: Authors' calculations using quarterly data from the BLS and BEA.

Notes: The BLS markup is the inverse of labor share in private business. The markups for nonfinancial corporate business are constructed by dividing NIPA data on either total compensation or wage and salary disbursements by income without capital consumption adjustment less indirect business taxes. Shaded areas represent periods of business recession as determined by the National Bureau of Economic Research.

of an expansion. That is, they all appear to be procyclical, though the peaks appear to lead the business cycle. Also evident from figure 1 is a pronounced upward trend. Although prior to the mid-2000s the markup was considered stationary, its rise over the past decade looks to be an acceleration of an upward trend that began in the 1980s.

We assess the cyclicality statistically by computing the correlation of markups with GDP.[6] To extract the cyclical components from each series, we consider three different filters: Hodrick-Prescott (HP), Baxter-King, and first-differences of the logarithms of the

6. Hall (2012) assesses cyclicality with respect to labor market variables rather than GDP. Because the cyclical behavior of productivity changed dramatically in the mid-1980s and because some shocks, such as technology shocks, are often found to drive output and labor in opposite directions, we chose GDP as the best measure of cyclicality.

variables.[7] The first three rows in table 1 report our three measures of cyclicality for the baseline markup over 1947 to 2012. In every case, the correlation is positive, with correlations ranging from 0.23 to 0.5, depending on the markup measure and the filtering method used. These results should not be a surprise to anyone who has studied the cyclicality of labor share. In fact, Table 1 of Galí, Gertler and López-Salido (2007) report a correlation of the price-cost markup with GDP of 0.28 for their sample and data.

As figure 1 showed, and as Rotemberg and Woodford (1999) also noted, there are interesting dynamics involved as well—in particular, the markup peaks well before the peak of the business cycle. The left panel of figure 2 plots the cross-correlations of the cyclical components of real GDP and the markup, with the cyclical components derived using the HP filter. The correlations are positive for all leads and current values, indicating that an increase in the markup signals a current and forthcoming increase in GDP. The peak correlation occurs at a lead of three quarters. The correlations become negative for lagged values, though, meaning that a current decrease in GDP signals an upcoming increase in markups. The right panel plots the dynamic correlations of the markup with the unemployment rate. Note that because the unemployment rate moves opposite of GDP, the pattern is inverted. Additionally, because the unemployment rate tends to lag GDP, it is not surprising to see that the dynamic correlations are shifted so that the contemporaneous correlation with unemployment is roughly zero.[8]

5.2 Production Function Generalizations

The markup based on a Cobb-Douglas production function with no overhead labor was procyclical. This section relaxes some of those assumptions. We focus on generalizations of the markup for private business, since the additional series required match that sector best.

Unfortunately, there is no data series on "overhead labor." Ramey (1991) has argued

7. Because our sample ends so close to a long and deep recession, there are end-point problems with the statistical filters. Thus, we also assessed the cyclical correlations with GDP detrended using the Congressional Budget Office's potential output series. The markup was a bit less procyclical using this cyclical indicator, but never countercyclical.

8. These dynamic correlations provide a clue as to why Hall (2012) finds procyclical labor share. He uses a nonstandard filter that is based on regressing labor share on both the current value of the unemployment rate as well as the difference between the current value of the unemployment rate and the average of four lags of the unemployment rate; see equation 23 of his paper. The impulse response functions we present later, which show the full dynamics, reveal the standard stylized fact of countercyclical labor share (and hence procyclical markups).

Table 1. Cyclicality of the Price-Cost Markup

Correlation with real GDP

	Filter		
Measure	*Hodrick Prescott*	*Baxter King*	*First difference*
Cobb-Douglas (CD), 1947:Q1–2012:Q4			
1. Private business (BLS)	0.331	0.337	0.498
2. Nonfin. corp. business - compensation	0.285	0.311	0.427
3. Nonfin. corp. business - wages	0.233	0.255	0.405
CD, Overhead labor, 1964:Q1–2012:Q4			
4. No overhead labor	0.267	0.212	0.497
5. Production worker labor share	−0.035	−0.125	0.416
CES, 1949:Q2–2012:Q4			
6. Fernald unadjusted TFP	−0.464	−0.546	0.012
7. Fernald utilization-adjusted TFP	0.488	0.515	0.444
8. Technology estimated from SVAR	0.456	0.407	0.642
CES, Overhead labor, 1964:Q1–2012:Q4			
9. Fernald utilization-adjusted TFP	0.225	0.239	0.392
10. Technology estimated from SVAR	0.146	0.074	0.568
CD, 1976:Q1–2012:Q4			
11. Average wage	0.266	0.304	0.503
12. Marginal wage	0.179	0.179	0.448
CES, Overhead labor, 1976:Q1–2012:Q4			
13. Marginal wage, SVAR technology	0.192	0.198	0.585

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

Notes: Contemporaneous correlation of cyclical components of log real GDP and log markup. Rows 4–14 use data for private business (BLS). "CD" stands for Cobb-Douglas. For CES production function, elasticity of substitution between capital and labor $\sigma = 0.5$; see equation 10. See the text for discussion of the three versions of the CES. Baxter-King correlations exclude three years at start and end of sample. Lines 2 and 3 use data through 2012:Q3. Markup based on marginal wage is given by equation 15.

that the number of nonproduction or supervisory workers is probably an upper bound on the number of overhead workers, though even this category of workers shows significant

Figure 2. Cross-Correlations of Markup with Real GDP and Unemployment Rate

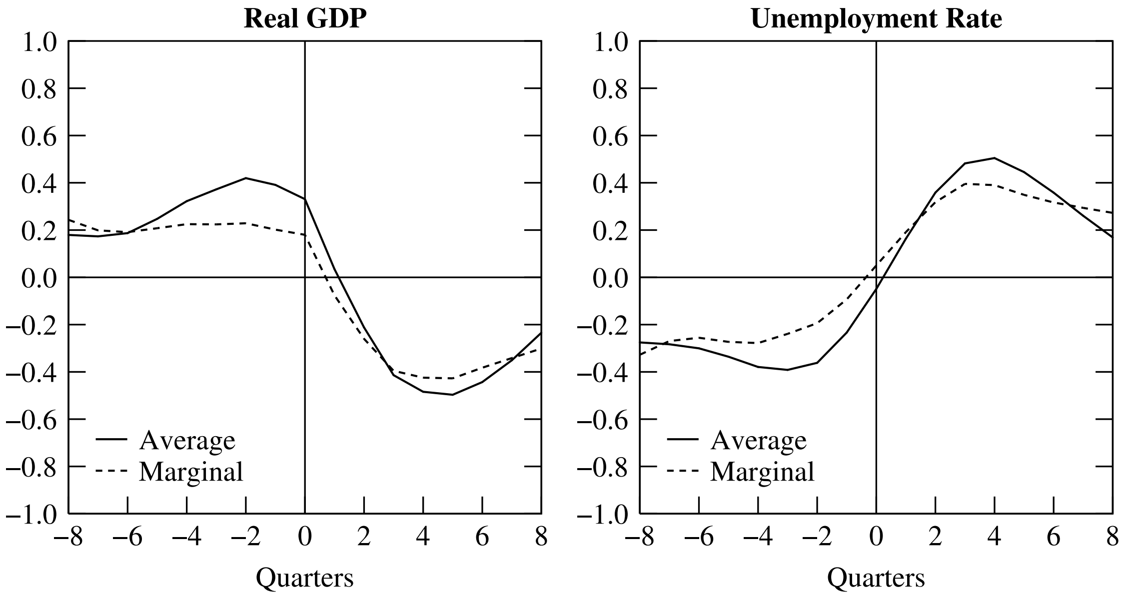

Source: Authors' calculations using quarterly BEA and BLS data.

Notes: Markup is inverse labor share for private business (BLS) over 1948:1–2012:4. Correlation of cyclical components of μ_{t+j} with y_t and u_t; detrended using HP filter ($\lambda = 1,600$).

cyclicality of employment.[9] To establish a *lower bound* on the cyclicality of markups due to the correction for overhead labor, we assume that all nonproduction workers constitute overhead labor. Using equation 6, we define the markup as:

$$\mu_{\mathrm{A}t}^{\mathrm{CDOH}} = -\ln s_t', \tag{16}$$

where s' is the labor share of production workers. This labor share is constructed by multiplying BLS data on employment, average weekly hours and average hourly wages of production workers in the private sector and then dividing by current dollar output in private business. The production worker series is only available starting in 1964.

The fourth row of table 1 shows the results for the baseline markup for the shorter sample from 1964 and row 5 shows the markup that assumes that all nonproduction and supervisory workers are overhead labor. Using the HP filter, the correlation falls from 0.27 to just below zero. Using the Baxter-King filter, the correlation becomes slightly more

9. Using HP filtered data, we find that the elasticity of the log of employment of nonproduction workers to GDP is positive and statistically significant and is about half of the elasticity of production workers with respect to GDP.

negative, -0.13. In contrast, the first-difference filter still shows very positive correlation of the markup with GDP. Thus, for the aggregate private sector, even the lower bound estimates do not support the idea of significant countercylicality. These results are different from those of Rotemberg and Woodford (1999) and Galí, Gertler and López-Salido (2007), who use steady-state approximations and calibrations that result in adjusting the standard markup by subtracting 0.4 times the cyclical variation in labor from the measure. Since labor input is very procyclical, it is easy to see how their adjustment would make markups look much more countercyclical.

We next consider the generalization that allows for a CES production function and a lower elasticity of substitution between capital and labor. The logarithm of the CES measure of the markup is

$$\ln\left(\mathcal{M}_{\mathrm{A}}^{\mathrm{CES}}\right) = -\ln s + \left(\frac{1}{\sigma} - 1\right)\left[\ln Y - \ln(ZhN)\right]. \tag{17}$$

To construct this measure of the markup, we require a value of the elasticity of substitution (σ) and a measure of the level of technology (Z). Chirinko (2008) surveys the literature estimating the elasticity of substitution between capital and labor and concludes that it is in the range of 0.4 to 0.6. Thus, we use an elasticity of substitution of 0.5. We consider three alternative methods for creating a series for the level of technology. The first two use Fernald's (2012) new quarterly series on TFP growth to back out Z. The first uses his unadjusted series, which is just a standard Solow residual. The second uses his utilization-adjusted TFP growth series. This series applies the Basu and Kimball (1997) framework to correct for unobserved variations in utilization. The third uses Galí's (1999) structural vector autoregression (SVAR) method to estimate technology shocks that can be used to create a technology level series. This SVAR identifies technology shocks as those shocks that have permanent effects on labor productivity in the long-run; thus any movements in labor productivity due to cyclical variations in utilization are excluded from this series. We use a simple bivariate SVAR in productivity growth and per capita hours growth, allowing for four lags.[10]

The third panel of table 1 shows the results. Line 6 shows that when we use the standard Solow residual, we find that the measured markup is markedly countercyclical. As seen in equation 17, anything that induces procyclicality of Z will make the markup more coun-

10. Of course, there is a long-standing debate on whether hours should appear in levels, first-differences, or detrended. Since we do not study the response of hours to technology shocks, the particular choice we make is not as important.

tercyclical. In stark contrast, Line 7 shows that when we use Fernald's utilization-adjusted technology series, markups are procyclical once again. The reason is that more of the procyclicality of labor productivity is attributed to variable factor utilization than to true technology changes. Finally, the SVAR approach, which uses completely different methodology and data than the Fernald approach, gives results very similar to his utilization-adjusted series.[11]

The big differences between the results using the unadjusted Solow residual and the other two methods shed light on why we continue to find procyclical markups whereas Rotemberg and Woodford (1999) and Galí, Gertler and López-Salido (2007) found countercyclical markups when they used a CES generalization. As Appendix B of Galí, Gertler and López-Salido (2007) outlines, they operationalize the CES production function assumption differently. To understand how they reached their approximation, first note that the expression for the markup in our earlier equation 10 can be written an alternative way:

$$\mathcal{M}_{\mathrm{A}}^{\mathrm{CES}} = \frac{\alpha}{s}\left(\frac{Y}{ZhN}\right)^{\frac{1}{\sigma}-1} \equiv \frac{1}{s}\left[1-(1-\alpha)\left(\frac{Y}{K}\right)^{\frac{1}{\sigma}-1}\right]. \tag{18}$$

where K is capital services and again s is the standard labor share. Thus, this alternative way adjusts the standard markup with a function of the output-capital ratio rather than with the ratio of output to the technology-adjusted labor input. However, in order to write the expression in logs of the key variables, Galí, Gertler and López-Salido (2007) must take approximations around the steady state. Thus, they specify the log markup as

$$\mu \approx -\ln s + \theta(\ln Y - \ln K),$$

Rotemberg and Woodford (1999) and Galí, Gertler and López-Salido (2007) set the value of θ based on a nonlinear combination of several steady-state elasticities. In addition to an elasticity of substitution between capital and labor of 0.5, their value of θ also depends on their calibrations of steady-state labor share of 0.7 and a steady-state average gross markup near unity. These values imply a value of θ equal to -0.4. They also assume that capital services are proportional to the stock of capital, so there can be no variation in capital utilization. Since output tends to rise relative to the capital stock during booms and their

11. In an analysis of the determinants of inflation, McAdam and Willman (2012) also find procyclical markups when they generalize from Cobb-Douglas to CES, but find more evidence of countercyclicality when they adjust for "labor utilization." However, their measure of labor utilization is simply hours per worker, which is the measure we use in our baseline. Thus, it is not clear to us why their results differ.

value of θ is negative, the second term in the equation above induces significantly more countercyclicality of the markup.

Thus, it is the differing assumptions about variable capital utilization that explain the gap between our results and theirs. The two methods we use that produce a procyclical markup allow for variable utilization of capital. Rotemberg and Woodford's (1999) and Galí, Gertler and López-Salido's (2007) findings are due to their implicit assumption that there is no variable utilization of capital. This assumption is at odds with the leading New Keynesian models, which rely on variable utilization of capital to fit the data. For example, Christiano, Eichenbaum and Evans (2005) find that allowing for significant variation in capital utilization is crucial for matching their data. In response to a monetary shock that results in a peak response of output just above 0.5, capital utilization rises to a peak of about 0.4.[12]

Finally, rows 9-10 of table 1 show the results when we combine the two generalizations, allowing for both overhead labor and CES production functions. As long as we use estimates of Z that allow for variable capital utilization, we still do not find countercyclical markups. The correlations suggest slight procyclicality or acyclicality.

5.3 The Marginal Wage Distinction

We next consider the cyclicality of the markup when we allow for the marginal-average wage distinction emphasized by Bils (1987). In this case, the measured markup (in natural logarithms) is given by

$$\mu_{\mathrm{M}}^{\mathrm{CD}} = -\ln s - \ln\left(\frac{W_{\mathrm{M}}}{W_{\mathrm{A}}}\right), \tag{19}$$

where $\mu \equiv \ln \mathcal{M}$. The last term is the log of the wage factor used in the average-marginal wage adjustment factor (equation 14).

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

The series for $\mathrm{d}v/\mathrm{d}h$ is the most challenging to measure. Bils (1987) speculated that

12. In contrast, Smets and Wouters (2007) find that capital utilization is not crucial to their results. The key reason is that they are able to match the cyclicality of labor productivity by estimating that fixed costs of production are 60 percent of GDP.

dv/dh was procyclical because a given increase in average hours would be more likely to come from an increase in overtime hours if the starting level of average hours was higher. He implemented this idea by regressing the change in average overtime hours, Δv, on the change in average total hours, Δh, in annual two-digit SIC manufacturing data, and allowing the coefficient in the regression to be a polynomial of average hours.

Average hours based on industry or aggregate data are not ideal for measuring this component for several reasons. As Bils pointed out, higher moments of the average hours distribution could matter because all workers do not work the same average hours. For example, it matters for the marginal wage whether average hours are increasing because more workers are moving from 38 to 39 hours per week or more workers are moving from 40 to 41 hours per week. Ideally, we want to compute the ratio of the change in overtime hours to the change in average hours at the level of the individual worker and then average over all workers at each point in time. That is, we want to construct the "average marginal" change in overtime hours with respect to a change in average hours. The ideal way to do this is to use panel data on individual workers.[13] To construct this series, we use Nekarda's (2013) Longitudinal Population Database, a monthly panel data set constructed from the Current Population Survey (CPS) microdata that matches individuals across all months in the survey. The data are available starting in 1976. We measure overtime hours as any hours worked above 40 per week. For each matched individual j who was employed two consecutive months, we compute the change in average hours Δh_{jt} and the change in overtime hours Δv_{jt}. By studying only those employed two consecutive months, we isolate the intensive margin, consistent with the theory. We construct the ratio dv_{jt}/dh_{jt} for each individual and compute the average of this ratio for all individuals each month. We then take the quarterly average of the monthly series to match our other aggregate data. Additional details are provided in the appendix.

The top panel of figure 3 shows the value of dv/dh for from 1976 through 2012. The series shows obvious procyclicality: it tends to rise during expansions and fall during recessions. It also exhibits some low frequency movements, rising from the mid-1970s to late 1990s and then trending lower thereafter.[14] Because dv/dh appears in the numerator of the wage factor, its procyclicality makes the wage factor more procyclical. But because the wage factor appears in the denominator of the markup, procyclicality of dv/dh has a countercyclical influence on the markup.

13. We are indebted to Steve Davis for suggesting this method for calculating dv/dh.

14. The two downward spikes during the mid-1990s are due to unusually disruptive winter storms.

Figure 3. Ratio of Marginal to Average Wages and Selected Components

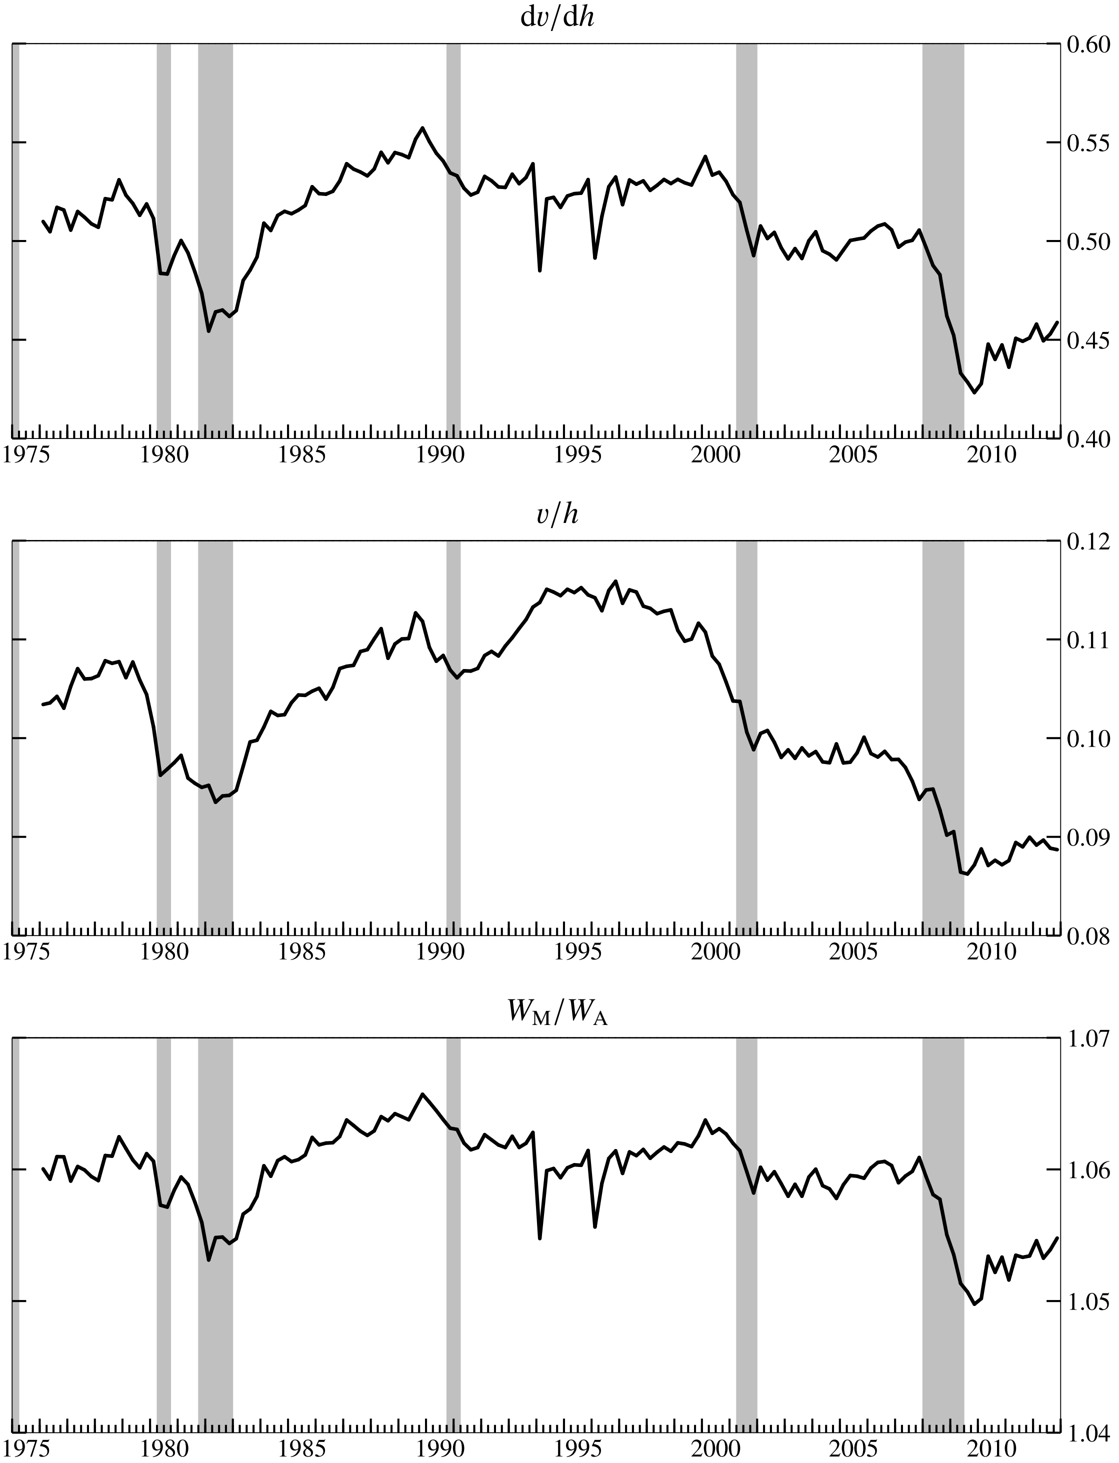

Source: Authors' calculations from Nekarda (2013).

Notes: Shaded areas represent periods of business recession as determined by the National Bureau of Economic Research.

The second series required for the wage factor is the average series v/h that appears in the denominator of equation 14. To be consistent in our data sources, we also use the LPD to measure this series. In particular, we calculate time series for v and h based on all individuals, and then compute their ratio.

The middle panel of figure 3 shows the fraction v/h. It is procyclical as well, though it tends to peak a bit before the peak of the business cycle. Like $\mathrm{d}v/\mathrm{d}h$, it also displays low frequency movements, although the decline since the late-1990s is more pronounced. Thus, the wage factor in equation 14 contains a procyclical series in both the numerator and denominator. Hence, the cyclicality of the factor depends in large part on the relative cyclicality of $\mathrm{d}v/\mathrm{d}h$ versus v/h.

Two more parameters are also required to construct the marginal-average wage factor. One is the fraction of overtime hours that command a premium, θ. We define as overtime hours, any hours worked greater than 40 hours per week. As some of those hours may come from salaried workers or persons with second jobs, not all hours over 40 are paid a premium. The only direct information is from the May supplements to the CPS in 1969–81, which asked workers whether they received higher pay for hours over 40 hours per week. We also use the BLS's Employee Costs for Employee Compensation 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. Based on the information from these two data sources, we use a value of $\theta = 0.3$ for the private economy. Additional details are provided in the appendix.

The final input required for the wage factor is the premium paid for overtime hours, ρ. 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.[15] Although there is considerable evidence that the implicit premium could be closer to 0.25, we use a ρ of 0.50 to reflect the statutory premium.[16] Using the higher overtime premium will bias the analysis toward finding countercyclical markups.

The bottom panel of figure 3 shows the marginal-average wage factor. Although the

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

16. 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 results using a 25 percent premium lie between those using the average wage and those using the 50 percent premium.

movements in the wage factor look procyclical, the variation is so small that it is unlikely to change substantially the cyclicality of the markup. This intuition is verified in the second from the bottom panel of table 1, shown previously. Because our estimate of the adjustment factor begins in 1976, row 11 reports the markup based on average wage over this shorter sample. Row 12 shows the markup over the marginal wage. Taking into account the procyclicality of marginal wages relative to average wages reduces the correlations by between 10 and 40 percent, depending on the filtering method. In no case, however, does the correlation become negative. The dashed lines in figure 2 show the effect of adjusting for the marginal wage on the dynamic correlation. As expected, the markup using marginal wages is less procyclical.

These results stand in contrast to those of Bils (1987), who found countercyclical markups in two-digit annual manufacturing data from 1956 to 1983. As the appendix shows, Bils' results are due to the combination of details in the implementation of his method for estimating $\mathrm{d}v/\mathrm{d}h$. We show that even within his framework, small adjustments in the method eliminate the finding of countercyclicality.

Finally, Row 13 shows the results when we combine the adjustments for overhead labor, CES production function (using the SVAR to estimate TFP), and marginal wages. The estimates continue to point to procyclicality.

To summarize, this section explored the unconditional cyclicality of the markup. Using three different measures of the baseline markup and three different methods for isolating the cyclical component, we found that the correlation with GDP was always positive. Some of the generalizations reduced the cyclicality. For example, assuming that all nonproduction workers are overhead labor, we were able to produce a correlation that was slightly negative. In contrast, the CES production function generalization did not reduce the procyclicality, as long as the estimate of technology was purged of variable factor utilization. Adjusting for marginal wages decreased the correlations somewhat, but they remained positive. However, even with the baseline markup, we found interesting dynamic correlations, which while positive for contemporaneous values and leads, were negative for lags. These results are linked to the propensity for markups to peak well in advance of the peak of the business cycle.

5.4 The Conditional Cyclicality of Markups

We now consider the cyclicality of the markup conditional on three types of shocks: technology shocks, government spending shocks, and monetary shocks. To capture the full dynamics, we investigate the comovement of GDP and markups using impulse responses estimated from VARs.

Because of sample constraints on key variables, as well as a desire not to include too many variables in one nested VAR, we use three separate off-the-shelf VARs to analyze the effects of our three shocks.[17] We estimate the effects of a technology shock using the same SVAR we used to estimate the technology level in section 5.2, but with the first-difference of the log markup included as a third variable. This system is estimated from 1948:Q1 through 2012:Q4. To investigate the effects of a government spending shock, we use the Ramey (2011) military news variable, which is measured as the present value of changes in expectations about future military spending, divided by lagged nominal GDP. The VAR also includes log real government spending, log real GDP, the three-month treasury bill rate, the Barro and Redlick (2011) average marginal tax rate, and the log of the markup.[18] We estimate the aggregate monetary shock using a standard VAR where shocks to the federal funds rate are the monetary policy shocks (such as in Christiano, Eichenbaum and Evans, 1999). We include quarterly log real GDP, the log of the GDP deflator, the log of the price index for commodities, the log of the markup, and the federal funds rate.[19]

The top three panels of Figure 4 show the impulse responses for log real GDP and the log of the baseline markup in response to the shocks.[20] For ease of comparison, we consider expansionary shocks in all three cases, and scaled such that the peak effect on GDP is unity. According to the estimates, a positive technology shock raises output. The markup rises on impact and remains positive, but response is small and not statistically different from zero. The remaining graphs show that both an expansionary government spending shock and an expansionary monetary shock raise both real GDP and the markup temporarily. Moreover, the positive responses are statistically different from zero for at least several

17. As argued by Ramey (2011) there is not enough variation in aggregate government purchases to identify shocks in the post-Korean War sample. Standard measures of monetary shocks include shocks to the federal funds rate, which is only available starting after the end of the Korean War.

18. The military news variable is ordered first. The sample extends from 1948:Q1 through 2008:Q4 and the specification includes four lags and deterministic trends.

19. The federal funds rate is ordered last. Four lags are included and a quadratic time trend are included and the VAR is estimated from 1954:Q3 through 2012:Q4.

20. The baseline markup is the inverse of labor share in private business, and thus assumes Cobb-Douglas and no overhead labor

periods. Thus, even conditional on classic demand shocks such as monetary policy or government spending shocks, markups appear to be procyclical. We find no evidence of countercyclical markups in response to demand shocks for our baseline measure. The bottom panel shows the responses in a monetary VAR where we use the markup based on the overhead generalization, which produced the lowest unconditional contemporaneous correlations. Even in this case, the response is generally positive, though not different from zero in the short-run.

In their analysis of inefficiency gaps and markups, Galí, Gertler and López-Salido (2007) estimate a monetary VAR and find that output and price-cost markups move in opposite directions in response to a federal funds rate shock (see figure 5 of their paper). They use the inverse of the labor share in nonfarm business as their measure of the price markup and they estimate their model from 1960:Q1 through 2004:Q4. To see why our results were different, we obtained their RATS programs and data. We were able to exactly replicate their results when we used their data and sample period. However, when we used the revised version of the data on their sample period, we found that markups moved in the same direction as GDP, and significant so. Thus, it appears that data revisions since 2004 change the sign of markup response.[21]

In summary, the aggregate data suggests that markups are either procyclical or acyclical. The only case in which we could produce a significant negative correlation between GDP and the markup was when we used a CES production function and did not allow for variable utilization of capital. Once we used Fernald's utilization-adjusted series, the correlation became very positive.

6 Industry Analysis

We now turn to the analysis of disaggregated manufacturing industry data. As discussed earlier, this data set allows us to study the cyclicality of markups within narrow industries, as well as to construct markups based on gross output and to use better instruments.

21. Monacelli and Perotti (2008) also explore the effects of government spending shocks on the markup. When they use the standard Blanchard and Perotti (2002) SVAR, they find that markups fall in response to an increase in government spending. When they use the Ramey and Shapiro (1998) war dates, they find that markups initially rise in response to a rise in government spending. None of their responses are statistically different from zero, though.

Figure 4. Conditional Cyclicality of Markups

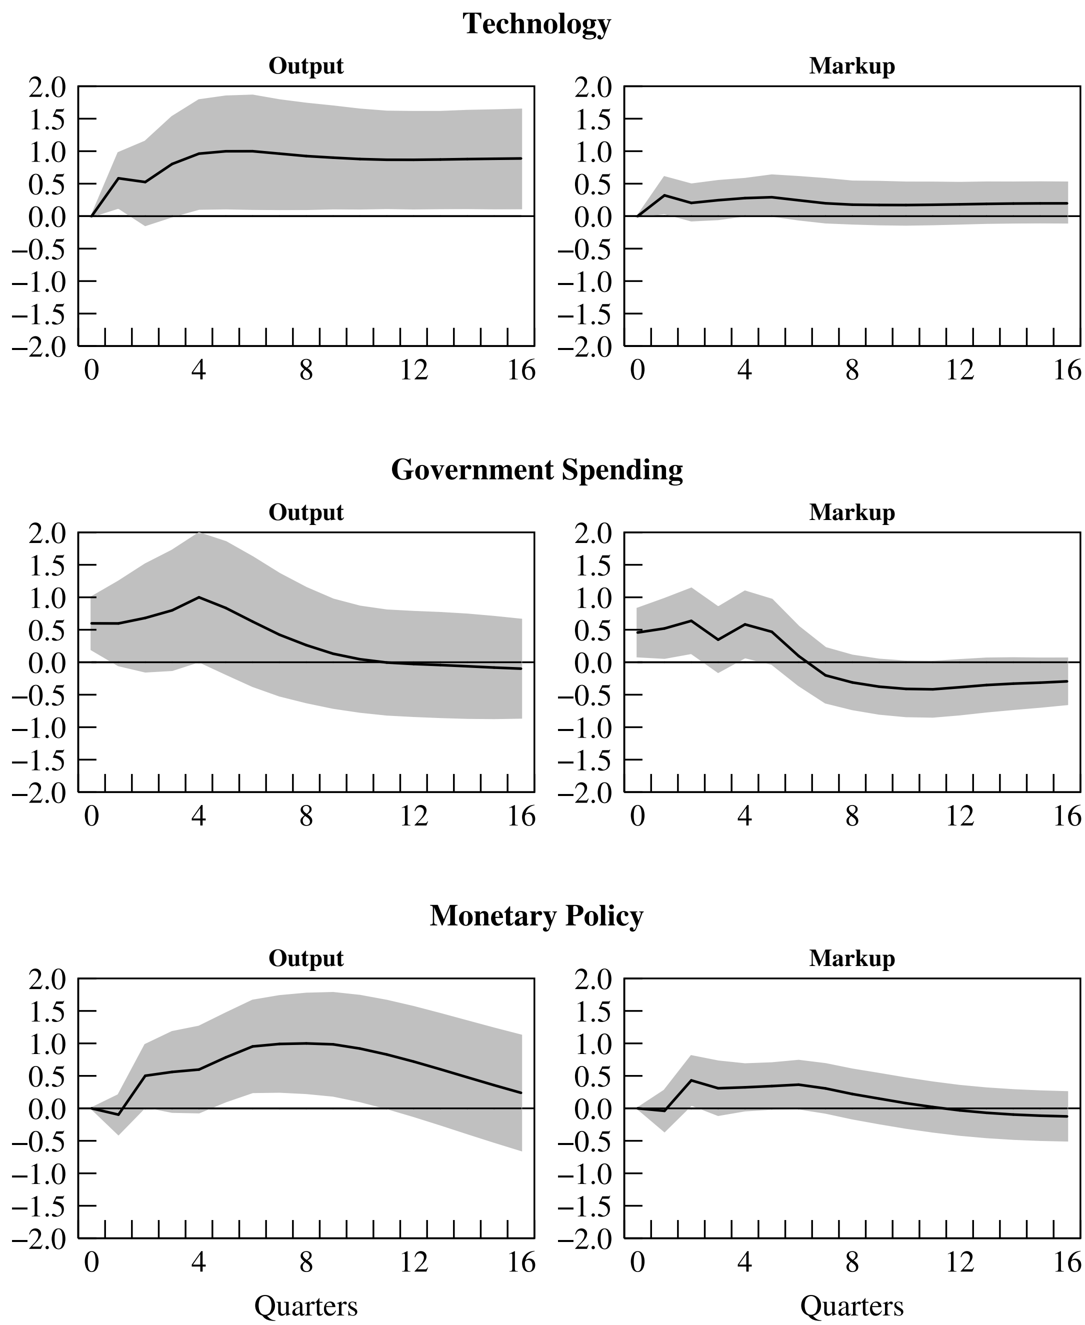

Source: Authors' calculations using quarterly BEA and BLS data.

Notes: Impulse response of real GDP and markup to a shock to variable indicated in heading. Shaded area indicates 95 percent confidence intervals. For technology shock, VAR includes first differences of logs of productivity, hours, and markup, with four lags. For government spending shock, VAR estimated as in Ramey (2011), but adding the log of the markup. For monetary policy shock, VAR includes log real GDP, the log of the GDP deflator, the log of the price index for commodities, the log of the markup, and the federal funds rate, with four lags.

6.1 Data and Econometric Specification

The main part of the data we use is an updated version of the data set constructed by Nekarda and Ramey (2011). This data set matches four-digit SIC level on government spending and its downstream linkages calculated from the Bureau of Economic Analysis's benchmark input-output accounts to the Manufacturing Industry Database (MID) published by the National Bureau of Economic Research and the Census Bureau's Center for Economic Studies. The new data extend from 1958 to 2009. Merging manufacturing SIC industry codes and input-output industry codes yields 274 industries. The web appendix of Nekarda and Ramey (2011) gives full details.

Most of the variables are constructed using data from the MID. This database provides information on variables such as shipments, inventories, price deflators, employment, hours and payroll. The appendix gives full details concerning the construction of the variables.

Our goal is to estimate how the markup responds to changes in real shipments generally, and to changes induced by either shifts in demand or technology. To construct our markup measure, we add industry (i) and time (t) subscripts and take annual log differences. Our estimation involves regressing the change in the logarithm of the markup, $\Delta\mu$, on the change in the natural logarithm of real shipments, $\Delta \ln Y$. In particular, we estimate:

$$\Delta\mu_{it} = \alpha_i + \alpha_t + \beta\Delta \ln Y_{it} + \varepsilon_{it}. \tag{20}$$

where α_i is industry fixed effects, α_t is year fixed effects, and ε is the error term. The estimate of β in the OLS regression indicates the unconditional cyclicality of the markup. To assess the cyclicality conditional on demand or technology shocks, we instrument for shipments with the appropriate instrument.

We construct three instruments in order to assess the conditional cyclicality of markups. The first instrument is estimated using long-run restrictions to identify technology shocks. Following Kiley (1998) and Chang and Hong (2006), we apply Galí's technique to the industry data. Because the MID's calculations for total factor productivity (TFP) assume a constant markup, we use labor productivity to identify shocks, as in Kiley (1998). (Chang and Hong used TFP instead of labor productivity.) We estimate SVARs with long-run restrictions on each industry separately to generate series of technology shocks for each industry. Following Chang and Hong (2006), we estimate for each industry a bivariate VAR in productivity growth and hours growth, and allow for one lag. We also use these estimates to construct the index of technology required for the CES production function generalization.

We also consider two demand instruments. The first is based on Nekarda and Ramey (2011). This instrument uses government demand for an industry's output as an instrument and is constructed by linking the MID to the input-output tables. The government demand instrument is defined as:

$$\Delta g_{it} = \bar{\theta}_i \cdot \Delta \ln G_t, \tag{21}$$

where $\bar{\theta}_i$ is the time average of the share of an industry's shipments that are sent to the federal government and G_t is aggregate real federal purchases from the NIPA. Thus, this 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. As discussed in Nekarda and Ramey (2011), this measure purges the demand instrument of possible correlation between industry-specific technological change and the distribution of government spending across industries. Since all regressions will include industry and year fixed effects, this instrument should be uncorrelated with industry-specific changes in technology or aggregate changes in technology.

The second demand instrument is a monetary shock instrument. We hypothesize that industries that produce more durable goods should find that demand for their products is more sensitive to monetary policy shocks. We combine data gathered by Bils and Klenow (1998) with information from the Los Angeles HOA Management "Estimating Useful Life for Capital Assets" to assign a service life to the product of each four-digit industry. We estimate the aggregate monetary shock using a standard VAR with quarterly log real GDP, GDP deflator, the price index for commodities, and the federal funds rate.[22] To create the industry specific shock, we multiply the aggregate monetary shock by the industry-specific service life:

$$m_{it} = \kappa_i \cdot \xi_t^M, \tag{22}$$

where κ_i is the service life for industry i and ξ_t^M is the aggregate monetary shock identified from the VAR.[23]

22. The system is the same one estimated for the aggregate analysis, except that it excludes the markup. We converted the quarterly shocks to annual data using the value of the shock in the first quarter since it had a higher first-stage F statistic than the average of the shocks over the year.

23. Since we wrote the earlier version of our paper, Bils, Klenow and Malin (2012) have used industry durability measures in another way to argue that markups are countercyclical. As Reis (2012) points out in

Table 2. First-Stage F Statistics in Manufacturing Industry Analysis

Period	*Technology*	*Government spending*	*Monetary policy*
1961–2009	843.3	62.3	8.6
1976–2009	562.6	23.8	17.1

Source: Authors' estimates using four-digit MID data.

Notes: Reported F statistics are from the regression of the change in log real shipments on the instrument, with year and industry fixed effects included. F statistics are based on Newey-West standard errors with two lags.

Table 2 reports the first-stage F statistics from the regression of $\Delta \ln Y_{it}$ on each of the instruments, with fixed effects included. Because there is evidence of serial correlation, the F statistics are based on Newey and West (1987) standard errors allowing for two lags. In the full sample, two of the three instruments have first-stage F statistics well above the Staiger and Stock (1997) threshold of 10, suggesting that these instruments are highly relevant. The first-stage F statistic for the monetary policy shock is 8.5 for the full sample, just under the threshold. However, in the restricted sample starting in 1976, all instruments have F statistics above the threshold.

6.2 Baseline Industry Results using the Standard Markup Measure

Our baseline measure of the markup is similar to the one used in the aggregate analysis. This measure assumes that (1) the average wage is equal to the marginal wage; (2) the production function is Cobb-Douglas; and (3) there is no overhead labor. This measure is given by

$$\mu_{\mathrm{A}it}^{\mathrm{CD}} = -\ln s_{it}, \tag{23}$$

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

We estimate equation 20 using the panel of four-digit manufacturing industries. We lose some observations at the beginning when creating the technology shock instruments, so our sample extends from 1961 through 2009 for those industries with data from 1958 to 2009. Because a few industries do not exist at the beginning or end of the sample, the panel

his discussion, though, their method only indicates relative cyclicality and not absolute cyclicality.

is not balanced. All told, the baseline regressions include a total of 13,307 industry-year observations. To account for serial correlation within industries, we report Newey-West standard errors using two lags.

The first row of table 3 shows results from the baseline specification. The first column shows the simple regression of markup growth on real shipment growth, controlling for year and industry fixed effects. This regression reveals the sign of the unconditional cyclicality of the standard markup measure. The estimate of β is 0.27 and is very precisely estimated. Thus, these results indicate that the standard measure of the markup is significantly procyclical in detailed manufacturing industry data. As with the aggregate data, these results on the unconditional cyclicality should come as no surprise to anyone familiar with the time series properties of the labor share. The labor share is known to be countercyclical, and since the standard measure of the markup is proportional to the inverse of the labor share, this measure of the markup is naturally procyclical.[24]

As discussed earlier, the New Keynesian model predicts that the cyclicality of the markup should differ based on the type of shock. To assess cyclicality conditional on shocks, we estimate equation 20 using instrumental variables. As shown in the second column, the markup is strongly procyclical conditional on a technology shock; the coefficient is 0.76 and is precisely estimated. This result is consistent with New Keynesian model, which predicts that a positive technology shock will raise the markup because prices are slow to adjust. Contrary to the predictions of the New Keynesian model, though, Column 3 shows that the markup does not respond to an industry-specific shock to government demand; the coefficient estimate is small, 0.06, and is not statistically different from zero. Finally, the fourth column reports results from using the industry-specific monetary policy shock as an instrument. The coefficient is 0.07 and is also not statistically significant from zero. Thus, the markup appears to be slightly procyclical or acyclical with respect to both the government spending and monetary shock demand instruments.

These baseline results are generally consistent with most of the literature's results that use dynamic factor demand methods or that generalize Hall's method for measuring markups. We now explore how the cyclicality changes when we apply Bils's (1987) and Rotemberg and Woodford's (1999) adjustments using our techniques and data.

24. We also investigated the robustness of the results to two variations in the specification. First, we estimated the model industry-by-industry and found that the markups were procyclical in all but a handful of industries. Second, to determine the sensitivity of the results to our use of first-differences, we instead HP filtered the data. These results also indicated procyclical markups. These results are available upon request.

Table 3. Industry Results

		Instrument for shipments		
Specification	*OLS*	*Technology*	*Government spending*	*Monetary policy*
1961–2009 (13,307 observations)				
Baseline	0.267**	0.757**	0.057	0.070
	(0.012)	(0.036	(0.063)	(0.245)
Overhead labor	0.244**	0.733**	0.030	0.021
	(0.015)	(0.035	(0.065)	(0.269)
CES production, SVAR	0.316**	0.740**	−0.008	0.170
	(0.012)	(0.033	(0.074)	(0.261)
1976–2009 (8,927 observations)				
Baseline	0.306**	0.806**	−0.117	0.091
	(0.015)	(0.046)	(0.152)	(0.160)
Marginal wage	0.303**	0.806**	−0.144	0.107
	(0.015)	(0.047)	(0.157)	(0.158)

Source: Author's regressions using data from MID, BEA benchmark IO accounts, and the BLS.

Notes: Regression of $\Delta\mu_{it} = \alpha_i + \alpha_t + \beta\Delta \ln Y_{it} + \varepsilon_{it}$ (equation 20) for industry i in year t. Newey-West standard errors (2 lags) are reported in parentheses; *** indicates significance at 1-percent, ** at 5-percent, and * at 10-percent level.

6.3 Generalizing the Production Function

We now redefine the markup to allow for the generalizations of the production function discussed in section 3 above. Following our aggregate analysis, we bound the effects of overhead labor by assuming all nonproduction workers are overhead labor. Thus, we define the markup as:

$$\mu^{\mathrm{CDOH}}_{\mathrm{A}it} = -\ln s'_{it}, \tag{24}$$

where s' is the labor share of production workers.

The second generalization allows for a CES production function. We do not have an industry-specific measure of utilization-adjusted TFP, so we just use the SVAR method to construct technology for each industry.

The second row of table 3 shows the results for the markup that allows for overhead labor. All of the results are very close to those in the baseline case. Both the OLS results and the IV results conditional on the technology instrument show substantial procycliality, whereas the IV results using the government spending instrument and the monetary shock instrument indicate slight procyclicality or acyclicality. Thus, excluding non-production workers has little effect in the industry data in contrast to the more noticeable effect in the aggregate data. One possible reason is that the industry data count only workers at the establishment level, so they already exclude nonproduction workers at headquarters of firms.

The third row shows the results when we allow for a CES production function with elasticity of substitution equal to 0.5. The results are similar to the Cobb-Douglas case. When the cyclicality is measured conditional on government spending, the parameter estimate becomes slightly negative (−0.01) but is not different from zero.

In short, none of the production function generalizations imply countercyclical markups, even conditional on demand shocks.

6.4 The Marginal Wage Distinction

We now consider the cyclicality of the markup when we allow for the marginal-average wage distinction. In this case, the measured markup is given by:

$$\mu^{\mathrm{CD}}_{\mathrm{M}it} = -\ln s_{it} - \ln\left(\frac{W_{\mathrm{M}}}{W_{\mathrm{A}}}\right)_{it}. \tag{25}$$

The last term is the log of the wage factor used in the average-marginal wage adjustment factor.

Our technique for constructing the marginal-average wage is similar to the one used in the aggregate data, but has two main differences. First, to ensure sufficient cell sizes, we constructed $\mathrm{d}v/\mathrm{d}h$ and v/h at the two-digit industry level rather than the four-digit industry level. We then assigned the two-digit value to each four-digit industry.[25]

The second difference is the value of θ, which gives the fraction of hours above 40 that are paid a premium. Because manufacturing data are richer, we were able to calibrate this parameter by comparing our estimates of overtime hours in the two-digit LPD to overtime

25. A second reason we did not calculate $\mathrm{d}v/\mathrm{d}h$ and v/h at the four-digit level is the difficulty of compiling a crosswalk of detailed industries across time. Separately, we explored using an alternative procedure for imputing the values to the four-digit data. We regressed the two-digit estimates of $\mathrm{d}v_{jt}/\mathrm{d}h_{jt}$ and v/h on h in the two-digit data and then used the estimated coefficients along with h from the four-digit MID data to create the two series for the four-digit data. The results we report below are little changed by this alternative procedure. Moreover, Nekarda and Ramey (2011) show that applying a Bils' type cubic polynomial adjustment, but estimated on quarterly data rather than annual data, also gives similar results to those reported below.

hours from the two-digit manufacturing data from the BLS's establishment survey, since this latter data set defines overtime hours as those hours that are paid a premium. We found that overtime hours in the establishment survey were, on average, only slightly lower than our constructed overtime hours series, so we set $\theta = 1$ for the manufacturing sample.

Because the LPD starts in 1976, we restrict our analysis to start in that year. The bottom panel of table 3 shows the results when we adjust the markup measure. We continue to make the baseline assumption of Cobb-Douglas in total hours. For reference, the first row of the lower panel reports the results using the average wage for the shorter sample. As in the longer sample, markups are procyclical in the baseline specification both unconditionally and conditional on technology shocks. The markup is negatively related to shipments conditional on the government spending shock and positively related to shipments conditional on the monetary shock. However, the coefficients are near zero in magnitude and statistical significance. The second row of the lower panel shows the markup constructed with the marginal wage assuming an overtime premium of 50 percent. In every case, the results differ little from those of the baseline case. The coefficient in the case of the government spending shock falls to -0.14 but it is not statistically different from zero. We also explored the effects of combining the production function generalizations with the marginal wage adjustment and found little effect.

Thus, the industry results give the same message as the aggregate results. Using richer data rather than calibrations based on steady-state approximations and parameters indicates that adjusting the markup for production function generalizations or marginal wage considerations has a minor effect on the estimated cyclicality of markups. Moreover, there is no evidence of significant countercyclicality of markups even conditional on demand shocks.

6.5 Robustness Check: The Behavior of Inventories

We have shown that neither the baseline measure of the markup nor standard generalizations exhibit significant countercyclicality. These measures depend, however, on various assumptions. 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.

Thus it is useful to provide a robustness check based on a framework that uses completely different assumptions. To this end, we check our results by subjecting the data to the Bils and Kahn's (2000) inventory test. They introduce a model in which the joint behavior of the inventory-sales ratio and output price growth has implications for the cyclicality of markups. In their model, if the ratio of sales to inventory stock is procyclical, then either the markup must be countercyclical or the discounted growth of marginal costs must be countercyclical. The intuition is straightforward. In their model, firms hold inventories in order to raise demand. When markups are constant and there are no short-run intertemporal considerations at play, firms find it optimal to hold a constant sales-inventory ratio. In the short-run, firms will allow the sales-inventory ratio to rise above its long-run average only if the benefit of holding inventories (i.e., the markup) falls or if current marginal costs of production are high relative to expected discounted marginal costs next period.

Estimating marginal costs sufficiently precisely to uncover subtle variations across periods is difficult. One can use this theory, however, to assess qualitatively whether markups are countercyclical without having to estimate marginal costs.[26] In particular, the model implies that if both the sales-inventory ratio and the discounted growth of prices are procyclical, then markups must be countercyclical.[27]

To assess the cyclicality of these two series in our data, we estimate the following equations:

$$\Delta \ln\left(\frac{Y_{it}}{stock_{it}}\right) = \alpha_i + \alpha_t + \gamma\Delta \ln Y_{it} + \varepsilon_{it}. \tag{26}$$

$$\Delta \ln\left(\frac{p_{it}}{p_{it-1}}\right) = \alpha_i + \alpha_t + \phi\Delta \ln Y_{it} + \varepsilon_{it}. \tag{27}$$

where *stock* is the stock available for sale (equal to beginning of period inventories plus current production), and p is the output price. The appendix gives details on how we construct the various series.

According to the Bils and Kahn (2000) model, if both γ and ϕ are positive, then markups

26. Bils and Kahn (2000) implement both the qualitative version and the structural version that requires the estimation of marginal costs. Their test statistics indicate that the structural model is rejected by the data for most industries.

27. Several assumptions are required to derive this implication; see Bils and Kahn (2000) for more details.

Table 4. Inventory Results

		Instrument for shipments		
Dependent variable	*OLS*	*Technology*	*Government spending*	*Monetary policy*
Sales-inventory ratio growth	0.111**	0.103**	0.013	0.032
	(0.005)	(0.013)	(0.029)	(0.070)
Price change growth	−0.144**	−0.299**	0.051	0.177
	(0.021)	(0.037)	(0.027)	(0.342)

Source: Author's regressions using data from MID, BEA benchmark IO accounts, and the BLS.

Notes: Regressions of $\Delta \ln X = \alpha_i + \alpha_t + \gamma \Delta \ln Y_{it} + \varepsilon_{it}$, where $X = Y_{it}/stock_{it}$ and p_{it}/p_{it-1}, respectively. Newey-West standard errors (2 lags) are reported in parentheses; *** indicates significance at 1-percent, ** at 5-percent, and * at 10-percent level.

must be countercyclical. Table 4 shows the coefficient estimates. The first column shows OLS regressions, which indicate the unconditional correlations. The estimates imply that the sales-inventory ratio depends positively on shipments, while the growth rate of prices depends negatively on shipments. Thus, the two key variables have opposite signs on their correlations, so the results do not imply countercyclical markups. The remaining columns show the results when the three instruments are used for shipments growth in order to judge the conditional correlations. The technology shock instrument produces correlations that are opposite in sign, and hence does not point to countercyclical markups. For the government spending and monetary shock instruments, both the sales-inventory ratio and the growth rate of prices are acyclical. All coefficients are positive, but they are not statistically different from zero and they are small.

Thus, in the two cases where the sales-inventory ratio is procyclical with respect to shipments, the growth of prices is countercyclical. In the other two cases, both variables are acyclical. Thus, none of these results implies countercyclical markups. The results are therefore consistent with those we found using our earlier method.

7 Conclusion

This paper has presented evidence that markups are largely procyclical or acyclical. Whether we look at broad aggregates or detailed manufacturing industries, average wages or marginal wages, or generalize the production function for lower elasticities of substitution or overhead labor, we find that all measures of the markup are procyclical or acyclical. We find no evidence of significantly countercyclical markups. These results hold even when we confine our analysis to changes in output driven by monetary policy or government spending.

Our results call into question the basic mechanism of the leading textbook 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.

Recently, some researchers have begun to focus more on the wage markup (such as Galí, Gertler and López-Salido, 2007) and to study sticky wages in more detail (such as Baratierri, Basu and Gottschalk, 2010). It is possible that a return to this traditional focus of Keynesian models on sticky wages might render these models more consistent with the microeconomic evidence.

References

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

Baratierri, Alessandro, Susanto Basu, and Peter Gottschalk. 2010. "Some Evidence on the Importance of Sticky Wages." National Bureau of Economic Research Working Paper 16130.

Barro, Robert J., and Charles J. Redlick. 2011. "Macroeconomic Effects from Government Purchases and Taxes." *Quarterly Journal of Economics*, 126(1): 51–102.

Bartelsman, Eric J., Randy A. Becker, and Wayne Gray. 2000. "The NBER-CES Manufacturing Industry Database." National Bureal of Economic Research Working Paper, Cambridge, MA.

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

Basu, Susanto, and Miles S. Kimball. 1997. "Cyclical Productivity with Unobserved Input Variation." National Bureau for Economic Research Technical Working Paper 5915, Cambridge, MA.

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

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

Bils, Mark, and Peter J. Klenow. 1998. "Using Consumer Theory to Test Competing Business Cycle Models." *Journal of Political Economy*, 106(2): 233–61.

Bils, Mark, Peter J. Klenow, and Benjamin A. Malin. 2012. "Testing for Keynesian Labor Demand."

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

Blanchard, Olivier J., and Roberto Perotti. 2002. "An Empirical Characterization of the Dynamic Effects of Changes in Government Spending and Taxes on Output." *the Quarterly Journal of economics*, 117(4): 1329–68.

Bound, John, David A. Jaeger, and Regina Baker. 1995. "Problems with Instrumental Variables Estimation When the Correlation Between the Instruments and the Endogenous Explanatory Variables is Weak." *Journal of the American Statistical Association*, 90(430): 443–50.

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

Chang, Yongsung, and Jay H. Hong. 2006. "Do Technological Improvements in the Manufacturing Sector Raise or Lower Employment?" *American Economic Review*, 96(1): 352–68.

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

Chirinko, Robert S., and Steven M. Fazzari. 1994. "Economic Fluctuations, Market Power, and Returns to Scale: Evidence from Firm-Level Data." *Journal of Applied Econometrics*, 9(1): 47–69.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 1999. "Monetary Policy Shocks: What Have We Learned and to What End?" In *Handbook of Macroeconomics*. , ed. John B. Taylor and Michael Woodford, 65–148. Elsevier.

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.

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.

Dunlop, John T. 1938. "The Movement of Real and Money Wage Rates." *The Economic Journal*, 48(191): 413–34.

Dunlop, John T. 1998. "Retrospectives: Real and Money Wage Rates." *Journal of Economic Perspectives*, 12(2): 223–34.

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.

Fernald, John. 2012. "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity." Federal Reserve Bank of San Francisco Working Paper 2012-19, San Francisco, CA.

Galeotti, Marzio, and Fabio Schianterelli. 1998. "The Cyclicality of Markups in a Model with Adjustment Costs: Econometric Evidence for US Industry." *Oxford Bulletin of Economics and Statistics*, 60(2): 121–142.

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

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.

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

Gordon, Robert J. 1981. "Output Fluctuations and Gradual Price Adjustment." *Journal of Economic Literature*, 19(2): 493–530.

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

Hall, Robert E. 1986. "Market Structure and Macroeconomic Fluctuations." *Brookings Papers on Economic Activity*, 2: 285–322.

Hall, Robert E. 2012. "The Cyclical Response of Advertising Refutes Counter-Cyclical Profit Margins in Favor of Product Market Frictions." National Bureau of Economic Research Working Paper 18370, Cambridge, MA.

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

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

Haskel, Jonathan, Christopher Martin, and Ian Small. 1995. "Price, Marginal Cost and the Business Cycle." *Oxford Bulletin of Economics and Statistics*, 57(1): 25–41.

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.

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

Kiley, Michael. 1998. "Labor Productivity in U.S. Manufacturing: Does Sectoral Comovement Reflect Technology Shocks?" Federal Reserve Board of Governors Unpublished paper.

Marchetti, Domenico J. 2002. "Markups and the Business Cycle: Evidence from Italian Manufacturing Branches." *Industrial and Labor Relations Review*, 13: 87–103.

McAdam, Peter, and Alpo Willman. 2012. "Technology, Utilization, and Inflation: What Drives the New Keynesian Phillips Curve?" European Central Bank working paper.

Monacelli, Tommaso, and Roberto Perotti. 2008. "Fiscal Policy, Wealth Effects, and Markups." National Bureau of Economic Research Working Paper 14584.

Morrison, Catherine J. 1994. "The Cyclical Nature of Markups in Canadian Manufacturing: A Production Theory Approach." *Journal of Applied Econometrics*, 9(3): 269–282.

Nekarda, Christopher J. 2013. "The Longitudinal Population Database." Federal Reserve Board of Governors Unpublished paper.

Nekarda, Christopher J., and Valerie A. Ramey. 2011. "Industry Evidence on the Effects of Government Spending." *American Economic Journal: Macroeconomics*, 3(1): 36–59.

Newey, Whitney K., and Kenneth D. West. 1987. "A Simple, Positive, Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55(3): 703–08.

Norrbin, Stefan C. 1993. "The Relationship between Price and Marginal Cost in U.S. Industry: A Contradiction." *Journal of Political Economy*, 101(6): 1149–64.

Oliveira Martins, Joaquim, and Stefano Scarpetta. 2002. “Estimation of the Cyclical Behavior of Markups: A Technical Note.” *OECD Economic Studies*, 34(1): 173–88.

Phelps, Edmund S. 1968. “Money-Wage Dynamics and Labor Market Equilibrium.” *Journal of Political Economy*, 76(4): 678–711.

Ramey, Valerie A. 1991. “Discussion of ‘Markups and the Business Cycle’.” In *NBER Macroeconomics Annual 1991.* , ed. Olivier Jean Blanchard and Stanley Fischer, 134–40. MIT Press.

Ramey, Valerie A. 2011. “Identifying Government Spending Shocks: It’s All in the Timing.” *Quarterly Journal of Economics*, 126(1): 1–50.

Ramey, Valerie A., and Matthew D. Shapiro. 1998. “Costly Capital Reallocation and the Effects of Government Spending.” *Carnegie-Rochester Conference Series on Public Policy*, 48: 145–94.

Reis, Ricardo. 2012. “Discussion of "Testing for Keynesian Labor Demand".”

Rotemberg, Julio J. 1982. “Monopolistic Price Adjustment and Aggregate Output.” *Review of Economic Studies*, 49(4): 517–31.

Rotemberg, Julio J., and Garth Saloner. 1986. “A Super-Game Theoretic Model of Price Wars During Booms.” *American Economic Review*, 76(2): 659–75.

Rotemberg, Julio J., and Michael Woodford. 1991. “Markups and the Business Cycle.” In *NBER Macroeconomics Annual 1991.* , ed. Olivier Jean Blanchard and Stanley Fischer, 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.* , ed. John B. Taylor and Michael Woodford, 1051–135. Elsevier.

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.

Staiger, Douglas, and James H. Stock. 1997. “Instrumental Variables Regression with Weak Instruments.” *Econometrica*, 65(3): 557–86.

Tarshis, Lorie. 1939. "Changes in Real and Money Wages." *The Economic Journal*, 49(193): 150–4.

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

Taylor, John B. 1980. "Aggregate Dynamics and Staggered Contracts." *Journal of Political Economy*, 88(1): 1–23.

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

Waldmann, Robert J. 1991. "Implausible Results or Implausible Data? Anomalies in the Construction of Value-Added Data and Implications for Estimates of Price-Cost Markups." *Journal of Political Economy*, 99(6): 1315–28.

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.

Appendix

Aggregate Data

Item		*Frequency and source*
Labor share, private business (BLS)	Q	BLS series `PRS84006173`
Gross value added, nonfinancial corporate business (NIPA)	Q	NIPA table `1.14`, line 17
Taxes on production and imports less subsidies, nonfinancial corporate business (NIPA)	Q	NIPA table `1.14`, line 23
Compensation of employees, nonfinancial corporate business (NIPA)	Q	NIPA table `1.14`, line 20
Wage and salary accruals, nonfinancial corporate business (NIPA)	Q	NIPA table `1.14`, line 21
Real GDP	Q	NIPA table `1.1.6`
Civilian unemployment rate	M	BLS series `LNS14000000`
Production and nonsupervisory employees, private business	M	BLS series `CES0500000006`
Average weekly hours of production and nonsupervisory employees, private business	M	BLS series `CES0500000007`
Average hourly earnings of production and nonsupervisory employees, private business	M	BLS series `CES0500000008`
Current dollar output, private business	Q	BLS series `PRS84006053`
TFP growth - unadjusted, private business	Q	Fernald (2012) series `dtfp`
TFP growth - adjusted for utilization, private business	Q	Fernald (2012) series `dtfp_util`
Labor productivity, private business	Q	BLS series `PRS84006093`
Hours, private business	Q	BLS series `PRS84006033`
Civilian noninstitutional population ages 16+	M	BLS series `LNS10000000`, adjusted to smooth revisions to population controls
Military news	Q	Ramey (2011)
Nominal GDP	Q	NIPA table `1.1.5`

Item		*Frequency and source*
Nominal government purchases	Q	NIPA table `1.1.5`
Implicit GDP price deflator	Q	NIPA table `1.1.9`
Three month treasury bill, secondary market rate	M	Board of Governors of the Federal Reserve System `H.15 release`
Average marginal tax rate	A	Barro and Redlick (2011)
Commodity price index	M	Commodity Research Bureau Spot Commodity Price Index `PZALL`
Federal funds rate	M	Board of Governors of the Federal Reserve System `H.15 release`
Nominal federal government purchases	Q	NIPA table `1.1.5`
Implicit price deflator for federal purchases	Q	NIPA table `1.1.9`

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 nonfinancial corporate business, the markup is Gross valued added less taxes on production and imports less subsidies divided by either compensation of employees or wage and salary accruals.

For the markup using private sector production workers, we divide the index of current dollar output in private business by the product of employment, average hours, and average hourly earnings of production and nonsupervisory workers in the private sector.

Longitudinal Population Database

We used individual-level data from Nekarda's (2013) Longitudinal Population Database, a monthly panel data set constructed from CPS microdata that matches individuals across all months, available for 1976 to 2012. In order to match the BLS private business data, we limit the sample to private-sector workers. We calculate v/h as follows. For all employed workers in each month we sum average weekly overtime hours (defined as those hours in excess of 40 per week) and average weekly hours. We seasonally adjust these two series separately (as discussed below) and then form our series as $\sum v / \sum h$.

To calculate $\mathrm{d}v/\mathrm{d}h$, for each matched individual i who is employed in two consecutive months we calculate

$$\left(\frac{\Delta v}{\Delta h}\right)_{it} = \frac{v_{it} - v_{i(t-1)}}{h_{it} - h_{i(t-1)}}.$$

Then for each month t we take the average over all individuals:

$$\left(\frac{\Delta v}{\Delta h}\right)_{t} = \frac{1}{P_t} \sum_{i=1}^{P_t} \left(\frac{\Delta v}{\Delta h}\right)_{it}.$$

Ideally, we would limit the matches to individuals employed in the same job over the two consecutive months, but the same-job measure does not exist prior to 1994. However, we found that the matched same-job measure was nearly identical to the matched employment measure after 1994, so we used the matched measure for individuals employed in consecutive months for the entire sample.

The raw data have significant seasonal variation. The CPS asks 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.

We seasonally adjust the series we calculate from the LPD (h, v, and dv/dh) using the Census Bureau's X-12-ARIMA program. We treat Easter and Labor day holidays that fall during the reference period as additive outliers and let the program impute the values in those months.

Share of Overtime Hours That Are Paid a Premium

We calculate the share of overtime hours that are paid a premium using data from CPS May extracts provided by the NBER.[28] 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.

Unfortunately, the key question on premium pay was dropped from the May supplement

28. http://www.nber.org/data/cps_may.html

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 A1 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.[29] On the other hand, the fraction of hours paid a premium is slightly countercyclical in the ECEC data.[30] 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.[31]

Industry Data

Hours, Shipments, and Markups

The main data on the four-digit manufacturing industries come from the NBER-CES Manufacturing Industries Database (MID).[32] The MID contains annual data on 459 manufacturing industries from 1958 to 2009. 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 the version based on the 1987 SIC codes.

In order to create industry-specific government spending instruments, we merge these data with input-output tables, as described in the web appendix to Nekarda and Ramey (2011). Merging these two data sources in a way that creates consistently defined industries results in 274 total industries.

We use MID measures of gross shipments, employment, annual hours worked, the wage bill for production and nonproduction workers, inventories, and the price deflator for shipments. We construct real shipments by dividing nominal shipments by the shipments price deflator.

The database provides information on annual hours only for production workers. We created two measures of total hours using two extreme assumptions: nonproduction workers always work 1,960 hours per year and nonproduction workers always work as much as production workers. The constant-hours value is slightly less than the usual 2000 hours

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

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

31. 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 change little.

32. Bartelsman, Becker and Gray (2000)

Figure A1. Fraction of Overtime Hours Worked Paid a Premium

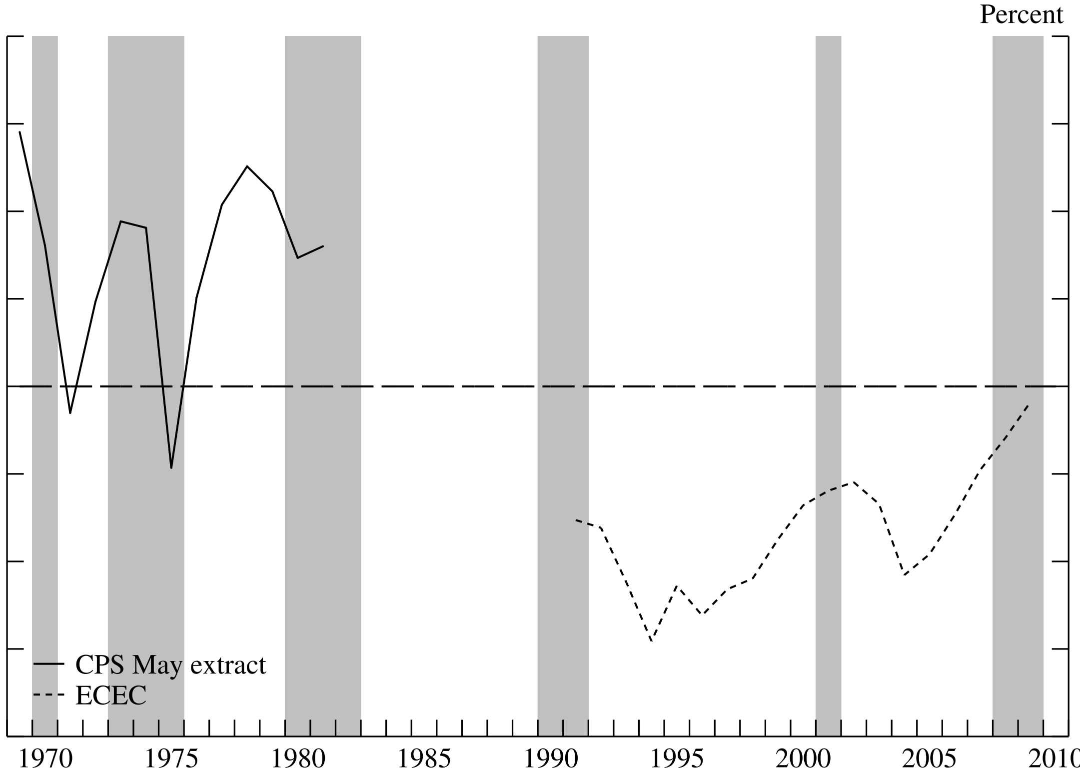

Source: Authors' calculations using data from May CPS extracts (NBER) and Employer Costs for Employee Compensation (BLS).

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.

per year because it allows for vacations and holidays, which are not included in production worker hours measures. The results were very similar using both measures, so we only report the results using the conservative assumption that nonproduction workers' hours are constant.

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 40.6, equal to the mean in the CES manufacturing data over 1958–2009.

dv/dh

We estimated v, h, and $\mathrm{d}v/\mathrm{d}h$ using individual-level data from the Longitudinal Population Database, as described in the earlier section. For application to the manufacturing industries, we estimated the series separately for individuals in each 2-digit manufacturing industry. To compare with the MID, we computed the annual average based on not-seasonally adjusted observations from March, May, August, and November.

Instruments

The technology and nontechnology shocks are estimated using a Galí (1999) type SVAR. We estimate a bivariate model separately for each industry. The variables are the log change in labor productivity, defined as real shipments divided by total hours, and the log change in total hours, allowing for one annual lag. We also use the estimated technology shocks to create a series for the level of technology used for the production function generalization.

The government spending variables are estimated as in Nekarda and Ramey (2011). The aggregate monetary variable is estimated from a VAR that includes log real GDP, GDP deflator, the price index for commodities, and the federal funds rate. The federal funds rate is ordered last. Four lags and a quadratic time trend are included and the VAR is estimated from 1954:Q3 through 2012:Q4. We then convert to annual data using the value of the shock in the first quarter since it had a higher first-stage F statistic than the average of the shocks over the year. To create the industry specific shock, we multiply the aggregate monetary shock by the industry-specific service life. We use data gathered by Bils and Klenow (1998) as well as the Los Angeles HOA Management "Estimating Useful Life for Capital Assets" to assign a service life to each industry.[33]

Comparison to Bils (1987) Estimates

Despite building on his insights about marginal wages, we reach the opposite conclusion from Bils concerning the cyclicality of markups. In this section, we show that the differences are due to implementation details in the estimation of $\mathrm{d}v/\mathrm{d}h$.

In order to replicate Bils's results, we construct a data set of the same industries, sample, and data sources used by Bils. In particular, we use monthly CES data for two-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. The annual series we use is the annual average of not-seasonally-adjusted data.

Recall that we had access to data that allowed us to calculate $\mathrm{d}v/\mathrm{d}h$ using individual data. Bils did not have these data, so he was forced to rely on an alternative method to generate a time-varying series on $\mathrm{d}v/\mathrm{d}h$. We replicate Bils's approach by using his polynomial specification for the estimation of $\mathrm{d}v/\mathrm{d}h$; his paper gives the reasoning behind this

33. http://www.homeownersassociationmanagementla.com/Estimating-Useful-Lives-of-Building-Components.html.

Table A1. Effect of Time Aggregation on the Slope of $\mathrm{d}v/\mathrm{d}h$, 1956–83

	Change in average hours	
Frequency	*39 to 41*	*36 to 43*
1. Monthly	0.12	0.26
2. Quarterly	0.11	0.33
3. Annual	0.24	0.81

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

Notes: Reports coefficient on Δh from regression $\Delta v_{it} = \{b_{i0} + b_{i1}t + b_2t^2 + b_3t^3 + c_1[h_{i(t-1)} - 40] + c_2[h_{i(t-1)} - 40]^2 + c_3[h_{i(t-1)} - 40]^3\}\Delta h_{it} + a_{i0} + a_1t + a_2t^2 + a_3t^3 + d_{i1}\ln[N_{it}/N_{i(t-1)}] + d_{i2}\Delta\ln[N_{it}/N_{i(t-1)}] + e_{it}$ (equation A.1). Data are annual and cover 1956–83.

specification. In particular, we estimate:

$$\text{(A.1)}\qquad \begin{aligned}\Delta v_{it} = \Big\{ & b_{i0} + b_{i1}t + b_2t^2 + b_3t^3 + c_1[h_{i(t-1)} - 40] + c_2[h_{i(t-1)} - 40]^2 \\ & + c_3[h_{i(t-1)} - 40]^3\Big\}\Delta h_{it} + a_{i0} + a_1t + a_2t^2 + a_3t^3 \\ & + d_{i1}\ln[N_{it}/N_{i(t-1)}] + d_{i2}\Delta\ln[N_{it}/N_{i(t-1)}] + e_{it}.\end{aligned}$$

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.[34]

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 two-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 A1 shows the effects of data frequency on the estimates of $\mathrm{d}v/\mathrm{d}h$ on the two-digit data. All estimates are for Bils's sample of 1956–83. 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, time aggregation appears to bias the slope estimate upward. From this we conclude that Bils's use of time-aggregated annual data appears to make $\mathrm{d}v/\mathrm{d}h$ more procyclical.

Table A2 shows the effect of changing frequencies and using different cyclical indica-

34. See Bils (1987), p. 844, for his motivation for including these terms.

Table A2. Effect of Time Aggregation on the Cyclicality of the Markup

		Correlation of markup with		
Frequency of dv/dh	*Frequency of markup*	*Real GDP*	*Industrial production*	*Total hours*
Two-digit industry data, 1956–83, 50 Percent Premium				
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

Source: Authors' calculations using two-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.

tors on the inferences about the cyclicality of markups. 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 dv/dh and applying it to quarterly markups. 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 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 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 A2 shows the results when we continue to use quarterly data to estimate dv/dh 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 dv/dh and to calculate the markup, which is a close replication of Bils's procedure. 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 −0.25) when using total hours—the indicator Bils used—as a cyclical indicator.

In sum, Bils's use of time-aggregated annual data to estimate dv/dh and his choice of cyclical indicator for his sample period were all necessary conditions for finding a countercyclical markup.