

Resurrecting the Role of the Product Market Wedge in Recessions

Mark Bilis

University of Rochester and NBER

Peter J. Klenow

Stanford University and NBER

Benjamin A. Malin*

Federal Reserve Bank of Minneapolis

September 24, 2014

Abstract

Employment and hours appear far more cyclical than dictated by the behavior of productivity and consumption. This puzzle has been labeled “the labor wedge” — a cyclical wedge between the marginal product of labor and the marginal rate of substitution of consumption for leisure. The wedge can be broken into a product market wedge (price markup) and a labor market wedge (wage markup). Based on the wages of employees, the literature has attributed the wedge almost entirely to labor market distortions. Because employee wages may be smoothed versions of the true cyclical price of labor, we instead examine the self-employed, intermediate inputs, and work-in-process inventories. Looking at the past quarter century in the U.S., we find that price markup movements are at least as important as wage markup movements — including in the Great Recession and its aftermath. Thus sticky prices and other forms of countercyclical markups deserve a central place in business cycle research, alongside sticky wages and matching frictions.

*The views expressed here are those of the authors and do not necessarily reflect the views of the Federal Reserve System. We are grateful to Corina Boar, Cian Ruane and Zoe Xie for excellent research assistance.

1. Introduction

Employment and hours are more cyclical than can be explained by real labor productivity under conventional preferences for consumption and leisure – see Hall (1997), Mulligan (2002), and Chari, Kehoe and McGrattan (2007), among others. More recently, there is a growing consensus that this “labor wedge” reflects labor market frictions. Examples include Gali, Gertler, and Lopez-Salido (2007), Shimer (2009), Hall (2009), and Karabarbounis (2014). This has helped drive an explosion of work on search, matching, and wage setting in the labor market.

The consensus that the labor wedge reflects labor market frictions is based on measuring the price of labor using average hourly earnings. The gap between average hourly earnings and labor productivity is acyclical, suggesting price markup movements are not cyclical. But it is not clear whether the marginal cost of labor to firms is well-measured by average hourly earnings. Employee wages may not reflect the true marginal cost of labor to the firm. Wages may be smoothed versions of the shadow cost due to implicit contracting (e.g., for salaried workers). One obtains a very different picture of the cyclical price of labor from using the wages of new hires, as measured by Pissarides (2009) and Haefke, Sonntag, and van Rens (2013), or from the user cost of labor, as measured by Kudlyak (2013).

In this paper, we seek evidence on cyclical distortions in the product market that do not rely on measures of the shadow price of labor. We attack this problem from three directions. First, we estimate the labor wedge for the self-employed. If we observe significant cyclicity in the labor wedge for the self-employed, it cannot be ascribed to wage rigidities or other labor market frictions. Second, we estimate the product market wedge from intermediate inputs (energy, materials, and

services). Intermediate prices should provide a truer measure of that input's cyclical price than do average hourly earnings for labor. Third, we estimate cyclical product market distortions by comparing movements in the marginal utility of consumption to those in the marginal product of work-in-process inventories. Intuitively, because the marginal utility of consumption is high in recessions, firms should sacrifice work-in-process inventories to put more finished products on the market in the absence of product market distortions.

Our evidence is for the U.S. from 1987 onward. For the self-employed, we look at the Current Population Survey and the Consumer Expenditure Survey, both conducted by the Bureau of Labor Statistics (BLS). For intermediates we use the BLS Multifactor Productivity Database covering 60 industries. For work-in-process inventories we draw on data from the Bureau of Economic Analysis. Our consistent finding is that, contrary to the emerging consensus, product market distortions are at least as important as labor market distortions in recent recessions.

Our findings speak directly to the puzzle of unemployment's high cyclical relative to that in labor productivity — the Shimer (2005) puzzle. A highly countercyclical product-market wedge translates into strongly procyclical labor demand, beyond what might be attributed to labor productivity. It provides a rationale for firms to create less employment in recessions without a decline in productivity, and even absent important wage stickiness.

The cyclical wedge we see for all inputs is compatible with firms whose sales are constrained in recessions by a (too high) sticky price. Given the wedges' strong persistence, it is also consistent with firms purposefully choosing a higher markup over marginal cost in recessions. As a recent example, Gilchrist, Schoenle, Sim, and Zakrajsek (2014) find that financially-constrained firms chose higher markups rather than

investing in market share during the Great Recession. Any model where expanding production has a component of investment (e.g., learning-by-doing) should have similar implications. Additionally, the product market wedge could reflect greater uncertainty, or aversion to uncertainty, in recessions, e.g., as in Arellano, Bai, and Kehoe (2012).

Our study is related to several earlier efforts. Bils (1987), Rotemberg and Woodford (1999), Bils, Klenow and Malin (2013), and Nekarda and Ramey (2013) rely on average hourly earnings to estimate the marginal cost of labor. Basu (1995) and Nakamura and Steinsson (2010) look at intermediate inputs, in particular how sticky intermediate prices affect business cycles. Galeotti and Schiantarelli (1998) infer the cyclicalities of price markups from the cyclicalities of capital investment, Bils and Kahn (2000) and Kryvtsov and Midrigan (2012) based on the cyclicalities of final-goods inventories, and Hall (2014) from the cyclicalities of advertising.

The paper proceeds as follows. Section 2 revisits the standard labor wedge calculations as a point of comparison. Section 3 looks at the self-employed. Section 4 investigates intermediate input use. Section 5 examines work-in-process inventories. Section 6 concludes.

2. The Aggregate Labor Wedge

We begin by constructing the standard *representative-agent* labor wedge (RAW), defined as the (log) ratio of the marginal product of labor (mpn) to the tax-adjusted marginal rate of substitution of consumption for leisure (mrs). Shimer (2009) provides a thorough derivation of the RAW, starting from the maximization problems of a representative household and firm. Constructing the wedge requires assumptions on preferences and technology; our baseline case follows Hall (1997) and Gali, Gertler, and Lopez-Salido (2007). Production features a constant elasticity with

respect to hours.¹ Preferences are separable in consumption and hours, and over time.² They feature a constant intertemporal elasticity for consumption, and a constant Frisch elasticity of labor supply. These assumptions yield a log-linear labor wedge:

$$\begin{aligned} RAW_t &\equiv \ln(mpn_t) + \ln(1 - \tau_t) - \ln(mrs_t) \\ &= \ln(y_t) + \ln(1 - \tau_t) - \left[\frac{1}{\sigma} \ln(c_t) + \frac{1}{\eta} \ln(n_t) \right], \end{aligned} \quad (1)$$

where y_t is output per hour, c_t is nondurables and services consumption per adult equivalent, n_t is hours per capita, and $\tau_t \equiv \frac{\tau_t^c + \tau_t^n}{1 + \tau_t^c}$ is a combination of average marginal tax rates on consumption and labor.³

For our baseline case, we use an intertemporal elasticity of substitution (IES) of $\sigma = 0.5$ following Hall (2009), and a Frisch elasticity of labor supply of $\eta = 1.0$, based on Chang, Kim, Kwon, and Rogerson (2014). The latter argue, based on a heterogeneous-agent model with both intensive and extensive labor margins, that a *representative-agent* Frisch elasticity of 1 (or slightly higher) is reasonable.

To gauge the cyclicity of the RAW, we project it on real GDP and hours worked. (All variables in logs and HP-filtered.) We use quarterly data from 1987 through 2012. Table 1 reports the cyclical elasticity of the wedge and its components: labor productivity, hours worked, consumption, and taxes. The wedge is strongly countercyclical (elasticity with respect to GDP: -2.69), reflecting mildly *countercyclical* productivity (-0.10), procyclical consumption (0.61) and very procyclical hours (1.40). In recessions, the RAW increases as the *mrs* plummets but the *mpn*

¹We later entertain CES production in capital, labor, and intermediate inputs.

²Nonseparable utility in consumption and leisure would not alter our results significantly. Shimer (2009) and Karabarbounis (2014) found this as well. We find this if we calibrate the nonseparability to how consumption responds to retirement (Aguiar and Hurst, 2013) or unemployment (Saporta-Eksten, 2014).

³See the Appendix for a precise description of all variables used.

changes little. Using the results in Table 1, it is straightforward to recalculate the wedge's cyclicality for alternative calibrations of σ and η .

Table 1: Representative Agent Wedge

	Elasticity wrt	
	GDP	Total Hours
Representative Agent Wedge	-2.69 (0.20)	-2.00 (0.06)
Labor Productivity	-0.10 (0.08)	-0.28 (0.04)
Hours per capita	1.40 (0.07)	0.99 (0.01)
Consumption per capita	0.61 (0.03)	0.36 (0.02)
Tax Rates	0.02 (0.07)	-0.01 (0.04)

Note: Each entry is from a separate regression. Sample covers 1987Q1-2012Q4. All variables in logs and HP filtered. Wedge calculation uses $\sigma = 0.5$ and $\eta = 1.0$.

As shown, the contribution of marginal tax rates to the cyclicality of the RAW is small. Because our tax measures have little impact on our results, we drop them in the remainder of the paper.⁴

In these baseline calculations we ignore cyclical fluctuations in the quality of the workforce and a role for overhead labor. We calculate that the declines in average workforce quality in expansions imply that we understate the cyclical elasticity of labor's marginal product by perhaps 0.1 to 0.2 (causing cyclicalities in the wedge to be overstated). Ignoring overhead labor, conversely, causes one to overstate the procyclicality of labor's marginal product (Rotemberg and Woodford, 1999). For an overhead labor component of the magnitude suggested by Nekarda and Ramey (2013), 10 to 20 percent, the impacts of composition and

⁴Mulligan (2012) contends that changes in effective marginal tax rates influenced labor market behavior in the Great Recession. His focus is on how lower income workers have been affected by the expansion of means-tested assistance programs.

overhead labor on cyclicity of labor's marginal versus average product (and on the estimated labor wedge) should approximately offset.

2.1. Extensive- and Intensive-Margin Wedges

We next construct separate wedges on the extensive margin (EMW) and the intensive margin (IMW). These distinguish between the two components of hours worked, employment and hours per worker. We make this distinction for four reasons. First, we can calibrate the Frisch elasticity of labor supply to micro estimates at the hours margin. Second, we can compare the intensive margin here to the intensive margin for the self-employed (in Section 3). Third, product market distortions should impact the wedge on both margins. If the labor wedge is only important at one margin, it would suggest the product market wedge has little cyclical importance. Finally, although the EMW appears in many theoretical models, to our knowledge it has not been constructed empirically.

In order to analyze the extensive margin, we make some additional assumptions. We consider a representative household that consists of many members. Consumption is perfectly shared across members, and labor supply decisions are made on both the extensive and intensive margins. Preferences are given by

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \frac{c_t^{1-1/\sigma}}{1-1/\sigma} - \nu \left(\frac{h_t^{1+1/\eta}}{1+1/\eta} + \psi \right) e_t \right\},$$

where c_t denotes per-capita consumption, e_t employment, and h_t hours worked per employee. ψ is a fixed cost of employment, which guarantees an interior solution for the choice of hours versus employment. The marginal disutility of employment is $\nu \left(\frac{h_t^{1+1/\eta}}{1+1/\eta} + \psi \right) \equiv \nu \Omega_t h_t$, while the marginal disutility of an extra hour per worker is $\nu h_t^{1/\eta} e_t$.

For firms, we assume (i) a constant output elasticity with respect to labor, and (ii) employment and hours per worker are perfect substitutes (i.e., production depends on total hours, $n_t = e_t h_t$). The mpn with respect to an additional hour of work is thus proportional to output per hour y_t , while the marginal product of employment is therefore $mpn_t^e = y_t h_t$ and the marginal product of hours per worker is $mpn_t^h = y_t e_t$.

There are frictions in finding employment. Firms post vacancies at the beginning of the period, and matches form and produce during the period.⁵ The matching technology is $m_t = v_t^\phi f(u_t)$, where m_t are matches, v_t vacancies, and u_t unemployment. κ denotes the opportunity cost of creating a vacancy, expressed in labor input as the fraction of the steady-state workweek h . δ is the exogenous per-period separation rate; and γ is the fraction of the initial period of employment devoted to training.

In this environment, the intensive margin wedge is given by

$$IMW_t \equiv \ln(mp_n^h) - \ln(mrs_t^h) = \ln(y_t) - \left[\frac{1}{\sigma} \ln(c_t) + \frac{1}{\eta} \ln(h_t) \right]. \quad (2)$$

The IMW, equation (2), differs from the standard RAW, equation (1), in two ways: hours per worker h_t replaces hours per capita n_t , and we calibrate $\eta = 0.5$. A lower η is appropriate given it now reflects the Frisch elasticity only at the intensive (hours) margin (e.g., Chetty et al., 2013).⁶

On the extensive margin, consider creating one more vacancy in period t and reducing vacancies in $t + 1$ just enough to keep employment unaffected in $t + 1$ forward. Spending $\kappa y_t h$ to create an additional vacancy generates $\phi m_t / v_t$ additional matches, of which $(1 - \delta)$ survive to $t + 1$. The perturbation thus requires lower spending on vacancies at

⁵Blanchard and Gali (2010) and Gali (2011) use this timing, although it is more conventional for matches to start producing in the following period. The former timing gives cleaner results, but it could be altered without changing our analysis substantially.

⁶Pescatori and Tasci (2012) point out that the labor wedge is less variable when calculated using hours per worker rather than hours per capita. They, however, hold the Frisch elasticity fixed across the workweek and representative agent calculations.

$t + 1$ by $(1 - \delta)\kappa y_{t+1}h \frac{m_t/v_t}{m_{t+1}/v_{t+1}}$. A social planner would set:

$$\begin{aligned} & \frac{\phi m_t}{v_t} \left[u'(c_t) \left(1 - \gamma \frac{h}{h_t} \right) y_t h_t - \Omega_t h_t \right] - u'(c_t) \kappa y_t h \\ & + \beta(1 - \delta) \mathbb{E}_t \left\{ u'(c_{t+1}) \left(\kappa y_{t+1} h \frac{m_t/v_t}{m_{t+1}/v_{t+1}} + \frac{\phi m_t}{v_t} \gamma y_{t+1} h \right) \right\} = 0. \end{aligned} \quad (3)$$

I.e., the marginal benefit of an extra vacancy (utility from consuming increased output today) equals its marginal cost (the disutility of employment plus adjusting consumption for the resource cost of creating an added vacancy today, while creating fewer vacancies in the future). To derive the EMW, we rearrange equation (3) to get a (log) ratio of the marginal benefit to the marginal cost of an additional unit of labor on the extensive margin. We get

$$\begin{aligned} EMW_t &= \ln(y_t) - \left[\frac{1}{\sigma} \ln(c_t) + \ln(\Omega_t) \right] - S_t, \\ S_t &\approx \frac{\frac{h}{h_t} \left(\frac{\kappa v}{\phi m} \frac{v_t}{m_t} \left[1 - (1 - \delta) \mathbb{E}_t \left\{ \frac{1}{1+r_{t+1}} \frac{y_{t+1}}{y_t} \frac{v_{t+1}}{m_{t+1}} \right\} \right] + \gamma \left[1 - (1 - \delta) \mathbb{E}_t \left\{ \frac{1}{1+r_{t+1}} \frac{y_{t+1}}{y_t} \right\} \right] \right)}{1 - \frac{r+\delta}{1+r} \left[\frac{\kappa v}{\phi m} + \gamma \right]}, \end{aligned} \quad (4)$$

where Ω_t is the marginal disutility of employment (per hour worked).⁷

The EMW, like the IMW, reflects movements in labor productivity, $\ln(y_t)$, and the marginal utility of consumption, $\frac{1}{\sigma} \ln(c_t)$. But there are differences from the IMW. Whereas the IMW reflects the marginal disutility of an extra hour, which is highly procyclical for reasonable Frisch elasticities, the extensive margin reflects the average disutility of adding a worker. We find this average disutility to be nearly acyclical. The term S_t reflects the efficacy of spending on vacancies, and it is specific to the EMW. In recessions S_t declines as vacancies are more

⁷Our derivation uses $\frac{1}{1+r_{t+1}} \equiv \frac{\beta u'(c_{t+1})}{u'(c_t)}$ and $\ln(1+x_t) \approx \text{constant} + \frac{x_t}{1+x}$. The Appendix provides more details on the EMW (and IMW) construction.

likely to yield a match. This lends a countercyclical component to the EMW. The cyclical nature of the EMW vis-a-vis the IMW essentially reduces to whether cyclical variation in the hiring term S_t exceeds that in the marginal disutility of working a longer workweek.

Constructing the EMW requires data on vacancies (v_t), matches (m_t), real interest rates (r_t), and additional parameters. A quarterly separation rate of $\delta = 0.105$ matches the average rate of quits, layoffs, and discharges in JOLTS. $r = 0.004$ yields an annual real interest rate of 1.6%, the average of the 3-month T-bill rate less core PCE inflation. Hiring costs per match, $\frac{\kappa v}{m}$, are set to 0.4 quarters of output, and training costs to $\gamma = 0.16$, consistent with estimates by Barron et al. (1999). Finally, the elasticity of matches to vacancies is $\phi = 0.5$. These parameters imply a steady-state ratio of mrs to mpn on the extensive margin of about 0.90.

Table 2 shows the elasticities of the EMW and IMW with respect to real GDP and hours worked. The EMW and IMW elasticities are similar, and smaller than for the RAW. An aggregate Frisch elasticity of 2.3 would make the RAW behave similarly to the IMW and EMW.

Table 2: Extensive & Intensive Margin Wedges

	Elasticity wrt	
	GDP	Total Hours
Extensive Margin Wedge	-1.89 (0.28)	-1.54 (0.15)
Intensive Margin Wedge	-1.91 (0.13)	-1.38 (0.05)

Note: Each entry is from a separate regression. Sample covers 1987Q1-2012Q4. All variables in logs and HP-filtered. Wedge calculations use $\sigma = 0.5$ and $\eta = 0.5$, and EMW expectation terms constructed by a VAR.

2.2. Decomposing the Wedge

We now empirically decompose the labor wedge into product-market (i.e., price markup) and labor-market (i.e., wage markup) components. This requires a measure of the marginal cost of labor to firms. As stressed by Gali, Gertler, and Lopez-Salido (2007) and others, the assumption that any particular wage measure reflects labor's true marginal cost is controversial. We will show that alternative wage measures lead to vastly different conclusions about the relative importance of the product- and labor-market wedges. This motivates our subsequent analysis, which decomposes the labor wedge without using wage data.

The IMW decomposition is standard and given by

$$\begin{aligned} IMW_t &= \left[\ln(y_t) - \ln\left(\frac{w_t}{p_t}\right) \right] + \left[\ln\left(\frac{w_t}{p_t}\right) - \frac{1}{\sigma}\ln(c_t) - \frac{1}{\eta}\ln(h_t) \right] \\ &= pmw_t^h + lmw_t^h, \end{aligned} \quad (5)$$

where $\frac{w_t}{p_t}$ is the (real) marginal cost of labor to firms. The intensive product market wedge (pmw^h) is the gap between the firm's marginal product and marginal cost of labor. The intensive labor market wedge (lmw^h) is the gap between the firm's marginal cost and the household's cost of providing an additional hour. The EMW decomposition is

$$\begin{aligned} EMW_t &= \left[\ln(y_t) - \tilde{S}_t - \ln\left(\frac{w_t}{p_t}\right) \right] + \left[\ln\left(\frac{w_t}{p_t}\right) + \tilde{S}_t - S_t - \frac{1}{\sigma}\ln(c_t) - \ln(\Omega_t) \right] \\ &= pmw_t^e + lmw_t^e, \end{aligned} \quad (6)$$

where \tilde{S}_t takes the same form as S_t (see equation (4)) but with $\phi = 1$. For intuition, temporarily let $\tilde{S}_t = S_t$ in equation (6). Doing so, it's apparent that the extensive labor-market wedge (lmw^e) mirrors the intensive (lmw^h), but with the household's *mrs* measured along the employment margin. For the extensive product-market wedge (pmw^e), the benefit to

the firm of an additional employee is $\ln(y_t) - S_t$; that is, our decomposition treats firms as paying the vacancy costs (S_t) and adjusts pmw^e accordingly. Finally, using \tilde{S}_t rather than S_t reflects the fact that firms do not internalize the congestion effects of their decision to post another vacancy; each firm views the probability of filling a vacancy as $\frac{m}{v}$, whereas the social planner knows one more vacancy generates $\frac{\phi m}{v}$ additional matches.

Table 3 decomposes the EMW and IMW into product- and labor-market wedges using average hourly earnings (AHE) as the (nominal) measure of the firm's marginal cost of labor (w_t). This wage measure is conventional and would reflect the true marginal cost if all workers were employed in spot markets. In this case, $\ln(y_t) - \ln(w_t/p_t)$ is the (log) inverse labor share. The product-market wedge accounts for between 2 and 6% of the wedge cyclical on the intensive margin and between 17 and 23% on the extensive margin. Thus, the results are in line with Karabarbounis's (2014) conclusion that the product-market wedge is relatively unimportant.

Alternative frameworks for understanding the labor market, however, emphasize the durable nature of the firm-worker relationship and imply the contemporaneous wage plays no allocative role. For example, in matching models with search frictions, what matters is the (expected) economic surplus generated over the life of the match and not the wage payment at any one time.⁸ In light of this, Table 4 shows how alternative measures of firms' marginal cost of labor affect the EMW decomposition. We consider the wages of new hires (NH) and Kudlyak's (2013) user cost of labor (UC), which have been argued to be more relevant for job formation in search frameworks.⁹ We find that,

⁸Implicit contracting models also imply the spot wage is not allocative.

⁹See Pissarides (2009), Haefke, Sonntag, and van Rens (2013), and Kudlyak (2013). Kudlyak (2013) finds that the new hire wage falls 1.2% relative to all workers' wages for

Table 3: **Wedge Decomposition: Average Wage**

	Elasticity wrt	
	GDP	Total Hours
Extensive Margin Wedge	-1.89 (0.28)	-1.54 (0.15)
Product Market (AHE)	-0.32 (0.13)	-0.35 (0.09)
Intensive Margin Wedge	-1.91 (0.13)	-1.38 (0.05)
Product Market (AHE)	-0.04 (0.08)	-0.08 (0.05)

Note: Each entry is from a separate regression. Sample is from 1987Q1 through 2012Q4. All variables in logs and HP filtered. Expectation terms in EMW constructed using a VAR. The extensive PMW follows equation (6), and the intensive PMW follows equation (5).

depending on the wage measure used, the pmw can account for almost none or essentially all of the EMW (i.e., from 15% to 115%). We presume the pmw to be similar at both the intensive and extensive margins — so this implies the pmw could be anywhere from nearly zero to nearly all of the intensive margin as well.

3. The Self-Employed Wedge

We now consider the cyclicity of hours specifically for self-employed workers, including measures of their labor wedge. If we observe significant cyclicity in the wedge for these workers, presumably it cannot be attributed to wage rigidities or other labor market distortions.

As a starting point, we note that self-employment has been as cyclical as total employment, at least for recent recessions. Hipple (2010) reports

each percentage point rise in the unemployment rate, while the relative fall in the user cost of labor is 3.4%. The results in Table 4 reflect adjusting the time series of average hourly earnings by these cyclical factors.

Table 4: **Wedge Decomposition: Alternative Wage Measures**

	Elasticity wrt	
	GDP	Total Hours
Extensive Margin Wedge	-1.89 (0.28)	-1.54 (0.15)
Product Market (AHE)	-0.32 (0.13)	-0.35 (0.09)
Product Market (NH)	-0.98 (0.16)	-0.81 (0.09)
Product Market (UC)	-2.17 (0.21)	-1.65 (0.09)

Note: Each entry is from a separate regression. Sample is from 1987Q1 through 2012Q4. All variables in logs and HP filtered. The product market wedge follows equation (6).

annual rates of self-employment in the U.S. for 1994 to 2009 based on the monthly CPS, and these series are extended through 2012 (based on Hipple input) by Heim (2014). The Hipple series reflect both incorporated and unincorporated self-employed. (Incorporated self-employed have constituted about one third of total self-employed.) The *share* of self-employed in nonagricultural industries declined slightly during each of the past two NBER-defined recessions: from 10.1 to 10.0 percent during 2001 and from 10.5 to 10.4 percent from 2007 to 2009. So, for 1994 to 2012, self-employment appears as cyclical as overall employment. The self-employment share exhibits lower-frequency fluctuations, but, if we HP-filter, the resulting cyclical series is completely acyclical with respect to GDP or aggregate hours.¹⁰

¹⁰These numbers are for nonagriculture, which represents about 94 percent of all self-employment. For agriculture, self-employment (again from Hipple) is acyclical—so the share is clearly countercyclical. We focus on nonagriculture workers for several reasons. For one, top and bottom coding of earnings in the CPS is extreme for farmers. Secondly, for farmers it is implausible to treat realized income as known at the time labor input is chosen, an assumption implicit in calculations of the labor wedge. Finally, we presume farmers face a competitive market.

Becoming self-employed requires starting a business; so fluctuations in self-employment could be affected by financing costs and constraints. In particular, the decline in self-employment during the Great Recession may partly reflect financing constraints. Going forward, we thus focus on the intensive (hours) margin for examining the self-employed.

We base our analysis on the Annual Social and Economic supplements to the CPS, typically referred to as the March CPS. In the March supplement household members report their hours and earnings for the previous calendar year. They also report the earnings and class of worker at their primary job – the job held longest during the prior year. The class-of-worker variables allow us to distinguish the self-employed separately for agriculture and non-agriculture. We begin our sample in 1987, the first year that data on primary-job earnings are available. Advantages of the March supplement are: (i) it is large; (ii) its top-coding of earnings is less extreme than in the monthly surveys; and (iii) some households can be matched across two consecutive March surveys, allowing us to examine year-over-year changes for a given set of workers.

Our unmatched sample contains 197,723 self-employed workers combining calendar years 1987 to 2012 (1,901,936 wage earners). Some earnings and hours responses are top coded. For each survey we trim the top and bottom 9.6% of workers by earnings on primary job. 9.6% is just large enough to remove top-coding of earnings for self-employed in all 26 years of surveys. We trim the bottom for symmetry; this also serves to remove all negative earnings. Usual hours are top coded at 99 per week. We trim the top 1.2 percent of workers by weekly hours. This is the minimal trimming that removes top-coded hours for all years. In lieu of trimming at low reported hours, we require that respondents worked at least 10 hours per week and at least 10 weeks during the year. We additionally require that workers be between ages 20 and 70.

Figure 1: Hours for Self-employed vs. Wage-earners

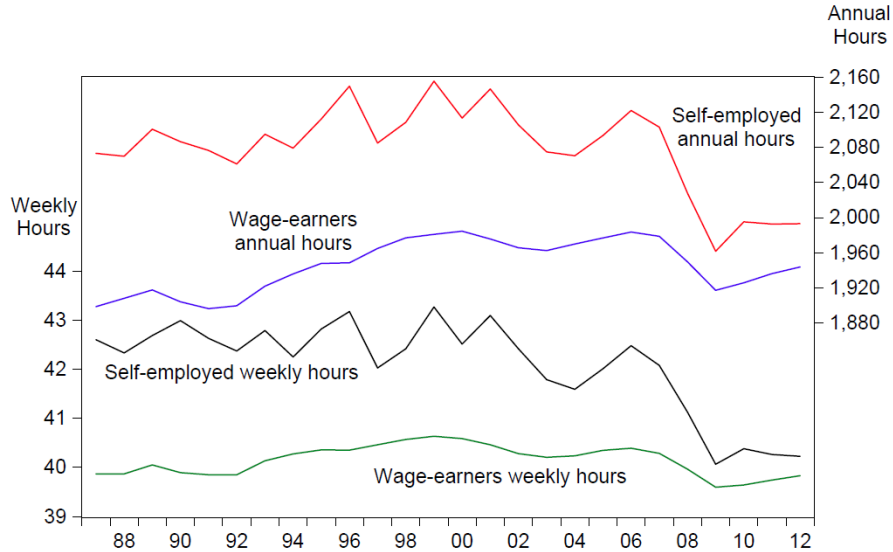


Figure 1 reports usual weekly hours and total annual hours worked separately for self-employed (nonagricultural) and those earning wages and salaries for 1987 to 2012. The intensive margin is clearly more cyclical for the self-employed. During the Great Recession (2007-2009) the workweek for self-employed declined by 4.9 percent (2 full hours) compared to only 1.7 percent for wage earners. Similarly, annual hours declined by 6.9 percent for self-employed compared to 3.2 percent for wage earners. More generally, if we regress hours per week on real GDP (both series in logs and hp-filtered), the impact for self-employed at 0.37 (standard error 0.14) is nearly twice that of 0.20 (s.e. 0.02) for wage earners. For annual hours, the impact for the self-employed is 0.57 (s.e. 0.18) compared to 0.39 (s.e. 0.04) for wage earners.

Figure 1 might be influenced by composition bias. For example, if workers becoming self-employed in expansions work more hours than the typical self-employed worker, then hours from Figure 1 will be biased

in a procyclical direction. For this reason, we match self-employed workers across consecutive March supplements, constructing growth rates for their hours and earnings.¹¹ Using these growth rates, we obtain hours and earnings relative to 1987. We are not able to match workers across the 1994 and 1995 calendar years due to the CPS sample redesign to reflect the 1990 U.S. Census. For 1994-1995 we impute to a series its mean growth rate. We create the corresponding level series indexed to 1987, then HP-filter that series. In all subsequent statistics, we exclude years 1994 and 1995. Our matched sample includes 39,306 self-employed workers, prior to trimming to deal with top coding.¹²

We construct indices for hours per week for both the self-employed and wage-earners. As with the unmatched data, the workweek is more cyclical for the self-employed. The elasticity of the workweek with respect to real GDP (both variables HP-filtered) is 0.28 (s.e. 0.07) for self-employed versus 0.17 (s.e. 0.03) for wage-earners. We did the same for annual hours and found them to be slightly more cyclical for wage-earners, with an elasticity with respect to real GDP of 0.57 (s.e. 0.07), compared to 0.54 (s.e. 0.13) for self-employed workers. (Similar remarks apply if we measure the cycle by HP-filtered aggregate hours.)

In Table 5 we report the cyclicity of the intensive-margin labor wedge. The first column is estimated for all workers, not just the self-employed. It repeats the analysis from Section 2, but uses workweek fluctuations constructed from the matched-CPS surveys. It is also annual, rather than quarterly, and excludes years 1994 and 1995 because

¹¹We follow standard matching procedures for the March CPS. Respondents are matched across years based on household and person identifiers and conformity of respondent's sex, race, and age.

¹²The March CPS responses for weeks worked and usual hours per week are for all jobs during the prior year, whereas class of worker and earnings refer to the primary (longest-held) job. To achieve a more earnings-compatible measure of hours growth, we restrict our self-employed sample to those who received 95 percent of their earnings from the primary self-employed job. (The average of earnings shares across the two years must be at least 0.95.)

we are unable to match those years in the CPS. It dispenses with the tax wedge, which we found to have little impact. As in Section 2, we find a strongly countercyclical wedge. The elasticities of the wedge with respect to real GDP and aggregate hours, -1.87 and -1.20, are modestly smaller than in Section 2 (estimates there being -1.91 and -1.38), with the difference largely reflecting a slightly less procyclical workweek.

Table 5: Cyclicalities of Labor Wedge, All Workers vs. Self-Employed

Elasticity wrt	(1)	(2)	(3)	(4)
Real GDP	-1.87 (0.10)	-2.06 (0.17)	-1.97 (0.25)	-3.23 (1.00)
Total Hours	-1.20 (0.05)	-1.41 (0.10)	-1.29 (0.16)	-1.93 (0.61)

Hours	All workers	SE	SE	SE
MPN	Agg. y/n	Agg. y/n	SE earn/hr	SE earn/hr
Consumption	NIPA PCE	NIPA PCE	NIPA PCE	+ CE Adj.

Notes: Sample is based on matched-March CPS self-employed outside government and agriculture. CPS observations are weighted. Each cell represents a separate regression. Regressions have 24 annual observations, 1987-1993 and 1996-2012. Newey-West standard errors are in parentheses. Hours are weekly. NIPA PCE refers to aggregate real expenditures on nondurables and services. CE adjustment incorporates consumption for the self-employed vs. all persons from the Consumer Expenditure Surveys.

Columns 2-4 of Table 5 show how the wedge's cyclicalities change as we sequentially replace the estimates of cyclicalities in hours, productivity,

and consumption for all workers with estimates for the self-employed.¹³ Column 2 constructs the wedge using fluctuations in the workweek just for self-employed workers, maintaining the same aggregate series for productivity and consumption. Not surprising, given the greater cyclicity of the workweek for self-employed described previously, this results in a slightly more cyclical labor wedge. For instance, the elasticity with respect to real GDP goes from -1.87 (s.e. 0.10) to -2.06 (s.e. 0.17).¹⁴

We next replace aggregate labor productivity with earnings per hour of the self-employed as a measure of the marginal product.¹⁵ (We deflate earnings by the nondurables and services PCE deflator.) This assumes that (i) a constant elasticity of output with respect to the self-employed worker's labor (as in Section 2); and (ii) earnings of the self-employed worker are proportional to output. Self-employed earnings per hour could overstate the procyclicality of labor's marginal product by ignoring the overhead component of self-employed labor. That could be especially important for self-employed production, given its small scale of operations. Another concern is that reported earnings could misstate actual earnings. The self-employed tend to understate earnings. Hurst, Li, and Pugsley (2014) show that the ratio of consumption to income is higher in survey data for the self-employed, consistent with the self-employed understating income. The concern for us would be if these workers underreport at a lower rate in recessions.

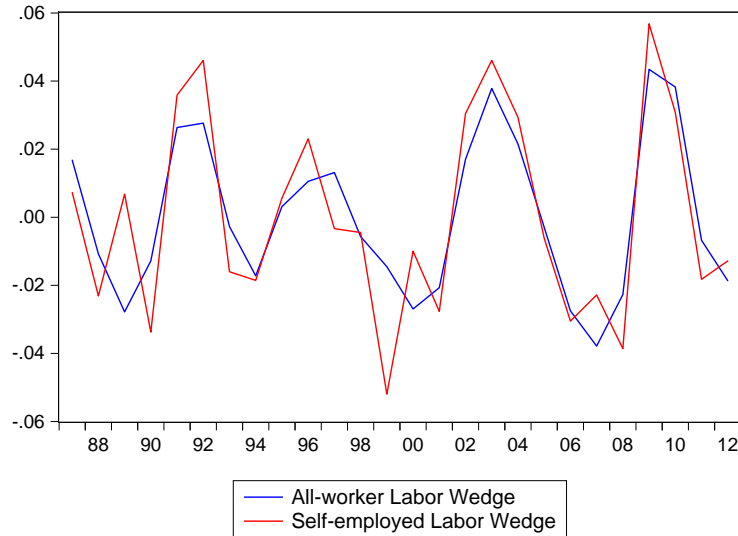
Going from Column 2 to 3 of Table 5, we replace aggregate labor productivity with self-employed earnings per hour. Aggregate labor productivity has been modestly countercyclical since 1987, with an

¹³Figures 8 - 11 in the Appendix plot the series that underlie these estimates.

¹⁴If we used fluctuations in annual hours, rather than weekly hours, then the wedge in column 1 with respect to real GDP, maintaining a Frisch elasticity of 0.5, would become -2.66 (s.e. 0.15), while that in column 2 would become -2.60 (s.e. 0.24).

¹⁵We use the midpoint formula to calculate the percentage change in a worker's earnings across matched years. This avoids extreme values for workers with very low earnings in one of the matched years.

Figure 2: Labor Wedge for Self-Employed vs All Workers



elasticity with respect to real GDP of -0.21 (s.e. 0.07). Self-employed earnings per hour have been less cyclical (elasticity -0.13 , s.e. 0.19). Thus, the estimated labor wedge becomes slightly more procyclical, with an elasticity of -1.97 (s.e. 0.25) with respect to real GDP. In summary, the wedge calculated based on earnings *and* hours for the self-employed is just as cyclical as that for all workers. Figure 2 plots the time series of these two wedges.

We have assumed the cyclicality of consumption for self-employed workers is the same as for consumption per capita. For robustness, we estimate self-employed consumption relative to aggregate consumption based on quarterly growth rates in household spending on nondurables and services in the Consumer Expenditure Surveys (CE). We add these estimates of relative consumption to aggregate consumption to obtain an estimate of consumption for the self-employed.

The elasticity of aggregate consumption with respect to real GDP is 0.64 (s.e. 0.04). Self-employed consumption is even more procyclical,

with an elasticity of 1.27. But the standard error is too large, at 0.56, to reliably infer that the self-employed have more procyclical consumption. The big standard error reflects the small number of self-employed observations in the CE. If we do use this measure of consumption in constructing the labor wedge, however, we get an even more cyclical wedge for the self-employed. This is illustrated in the last column of Table 5. The self-employed labor wedge now exhibits an elasticity of -3.23 (s.e. 1.00) with respect to real GDP. In this section's subsequent exercises, we revert to measuring self-employed consumption by aggregate consumption, rather than adopting such a noisy measure.

Table 6 presents two robustness exercises. First, self-employed who are incorporated might take income in the form of corporate profits rather than business earnings. It is not obvious how incorporated self-employed treat these profits in answering the CPS earnings question. For this reason, as an alternative measure of labor productivity, we consider earnings per hour excluding the incorporated self-employed. This series is more procyclical than earnings per hour for all self-employed. Its elasticity with respect to real GDP is 0.28 (s.e. 0.26), whereas the measure for all self-employed is slightly countercyclical. The first two columns of Table 6 repeat Columns 1 and 3 from Table 5, while Table 6, Column 3 employs earnings per hour for those not incorporated. The wedge becomes less countercyclical, with an elasticity with respect to real GDP of -1.57 (s.e. 0.24). Nevertheless, the self-employed labor wedge remains extremely cyclical, and still nearly as cyclical as the wedge estimated for all workers (Column 1).¹⁶

A second robustness exercise considers that self-employed workers are distributed differently across industries than are wage earners. For

¹⁶Because the workweek is slightly less cyclical for unincorporated self-employed, a wedge constructed solely for those not incorporated is slightly less cyclical than in Column 3. (Its elasticity with respect to real GDP is -1.39, standard error 0.38.)

Table 6: **Cyclicalilty of the Wedge, All vs. Self-Employed, Alternatives**

Elasticity wrt	(1)	(2)	(3)	(4)
Real GDP	-1.87 (0.10)	-1.97 (0.25)	-1.57 (0.24)	-1.64 (0.32)
Total Hours	-1.20 (0.05)	-1.29 (0.16)	-1.03 (0.20)	-1.03 (0.19)
Hours	All workers	SE	SE	SE
MPN	Agg. y/n	Agg. y/n	SE earn/hr	SE earn/hr
CPS weighting	CPS weights	CPS weights	CPS weights	+ Ind. shares

Notes: Sample is based on matched-March CPS self-employed outside government and agriculture. CPS observations are weighted. Each cell represents a separate regression. Regressions have 24 annual observations, 1987-1993 and 1996-2012. Newey-West standard errors are in parentheses. Hours are weekly. NIPA PCE consumption.

instance: Self-employment is about twice as frequent among workers in construction, a highly cyclical industry, than overall; it is considerably less common in durable manufacturing, which is also highly cyclical. We constructed a self-employed labor wedge reweighting observations by industry so that the weighted shares of self-employed workers by industry mimics that for all workers. (We do this for a breakdown of 12 major industries.) For example, if self-employment is twice as frequent in construction, then self-employed workers in construction are down-weighted by a factor of one-half. The results are given in Table 6, Column 4. The cyclicalilty of the self-employed wedge is modestly

reduced. The elasticity is now -1.64 (s.e. 0.32) with respect to real GDP. Again, however, it remains extremely cyclical, nearly as cyclical as that calculated for all workers.

Thus we conclude that self-employed workers exhibit a highly countercyclical labor wedge. Depending on specification choices, it is either as cyclical as the wedge calculated for all workers or nearly as cyclical. Because this wedge is presumably not driven by wage or other labor market distortions, it is evident of a highly countercyclical product market wedge. By extension, we find it suggestive of a countercyclical product market wedge for the overall economy.

4. Intermediate Inputs

The conventional way to estimate the Product Market Wedge (PMW) is based on the cyclicality of labor's share of income, e.g., Karabarbounis (2014). But, in principle, any input with a well measured marginal product and marginal price can be used to infer marginal cost and therefore price markups. Here we investigate the cyclicality of spending on intermediate inputs – materials, energy, and services – relative to gross output.

Intermediate inputs are promising for several reasons. First, intermediates are used by all industries, so evidence on them goes beyond just manufacturing or the self-employed. Second, adjustment costs for intermediates are believed to be low relative to adjustment costs for capital or even labor. See, for example, Basu (1995) or Levinsohn and Petrin (2003). Third, the wedge calculations based on cyclicality in a factor's share of output assumes that factor has no overhead component. This assumption seems much more defensible for intermediates than for labor.

A key question is whether intermediate prices reflect the marginal cost

of intermediate inputs. Long term relationships between firms and their suppliers could raise the same implicit contracting issues that arise with labor. Still, intermediates offer an independent piece of evidence *vis a vis* labor. And, as with labor, one would expect price smoothing relative to true input costs to impart a procyclical bias to the estimated PMW.

4.1. Technology for Gross Output

We assume a C.E.S. production function for gross output in an industry:

$$y = \left[\theta m^{1-\frac{1}{\varepsilon}} + (1-\theta) \left[z_v \left[(1-\alpha)k^{1-\frac{1}{\omega}} + \alpha(z_n n)^{1-\frac{1}{\omega}} \right]^{\frac{\omega}{\omega-1}} \right]^{1-\frac{1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}}$$

where y denotes gross output, m intermediate inputs, k capital, and n labor. Technology shocks can be specific to value added (z_v) or labor (z_n). The elasticity of substitution between intermediates and value added is ε and between capital and labor (within value added) is ω .

This technology implies the marginal product of output with respect to intermediate inputs is

$$\frac{\partial y}{\partial m} = \theta \left(\frac{y}{m} \right)^{\frac{1}{\varepsilon}}.$$

Based on this marginal product, we can estimate the PMW as

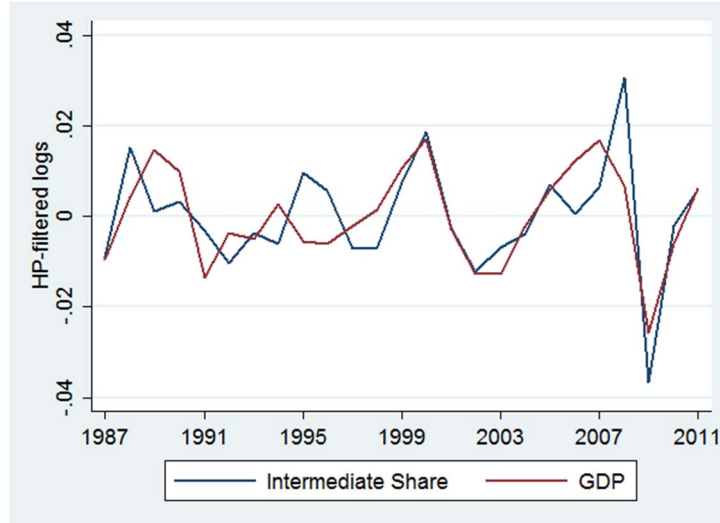
$$\mu^p = \frac{p}{p_m / \frac{\partial y}{\partial m}} = \frac{p \theta \left(\frac{y}{m} \right)^{\frac{1}{\varepsilon}}}{p_m}.$$

In the special case of Cobb-Douglas aggregation of intermediates and value added ($\varepsilon = 1$), the PMW is the inverse of intermediates' share:

$$\mu^p = \frac{\theta p y}{p_m m}$$

A higher price-cost markup boosts gross output relative to spending on

Figure 3: Cyclicity of Intermediate Share



intermediates in an industry. This is analogous to using inverse labor's share to measure price markup movements. A countercyclical markup would show up as a procyclical intermediate inputs share.

4.2. Evidence on the Cyclicity of Intermediate Inputs

We use the Multifactor Productivity Database from the U.S. Bureau of Labor Statistics on industry gross output and KLEMS inputs (capital, labor, energy, materials and services). It contains annual data from 1987–2011 and covers 60 industries (18 in manufacturing).¹⁷

Figure 3 plots the weighted-average industry intermediate share against GDP, where both variables are in logs and HP-filtered. As shown, spending on intermediates relative to gross output is highly procyclical. This is also true if we define the cycle in terms of hours worked.

We next run regressions of the inverse intermediate share on the cycle.

¹⁷The Appendix provides more details. KLEMS intermediate inputs come from BEA annual input-output accounts. These reflect purchases during the year minus inventory accumulation. (For details, see www.bea.gov/papers/pdf/IOmanual.092906.pdf.)

The specification is

$$\log \left(\frac{p_{it} y_{it}}{p_{mit} m_{it}} \right) = \alpha_i + \beta^p \log(cyc_t) + \epsilon_{it}$$

where cyc_t is either real GDP or hours worked, and all variables are HP filtered. The industry fixed effects (α_i) should take out changes in the aggregate share due to shifting industry composition over the cycle. We weight industries by the average share of their value added in all industry value added from 1987–2011. Standard errors are clustered by year.

Table 7 presents the results. The inverse intermediate share, a proxy for the price-cost markup, is systematically countercyclical. This is true for both measures of the cycle, for all industries together, and separately for manufacturing and non-manufacturing industries. It is also true if we weight industry-years by Tornqvist value added shares rather than industry shares over the entire sample (not reported in the table).¹⁸

What share of the total wedge might be accounted for by the PMW? To answer this, we construct an industry-specific total wedge that is consistent with the gross-output production function we consider. We replace aggregate labor productivity ($\frac{v}{n}$) with nominal gross output per hour in each BLS industry (relative to the consumption deflator). We also consider preferences that allow for an industry-specific marginal rate of substitution.¹⁹ The industry- i (intensive-margin) total wedge is thus

$$\ln(\mu_i^h) = \ln \left(\frac{p_i m p n_i^h}{p m r s_i^h} \right) = \ln \left(\frac{p_i v_i}{p \frac{v}{n}} \right) + \ln \left(\frac{y_i}{v_i} \right) - \frac{1}{\eta} \ln \left(\frac{h_i}{h} \right) + \ln \left(\frac{m p n^h}{m r s^h} \right),$$

¹⁸For manufacturing, we can break intermediate inputs into materials, energy, and services. The inverse shares for materials and energy are both countercyclical, and significantly so. The inverse share of spending on services, in contrast, is acyclical. Perhaps services are contracted less in spot markets than materials, or especially, energy.

¹⁹Preferences are $\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \frac{c_t^{1-1/\sigma}}{1-1/\sigma} - \sum_i \left[\nu \left(\frac{h_{it}^{1+1/\eta}}{1+1/\eta} + \psi \right) e_{it} \right] \right\}$. The Appendix shows $\ln \left(\frac{m r s_{it}^h}{m r s_t^h} \right) = \frac{1}{\eta} \ln \left(\frac{h_{it}}{h_t} \right) + \ln \left(\frac{e_{it}}{e_t} \right)$ for these preferences and considers alternatives. It also constructs the industry-specific total wedge on the extensive margin and for $\varepsilon \neq 1$.

Table 7: **Cyclicality of Inverse Intermediate Share**

	Elasticity wrt	
	GDP	Total Hours
All Industries	-0.86 (0.23)	-0.56 (0.14)
Manufacturing	-0.80 (0.30)	-0.55 (0.19)
Non-Manufacturing	-0.88 (0.22)	-0.56 (0.14)

Note: Each entry is from a separate regression. Annual data from 1987-2011 for 60 industries (1500 industry-year observations), 18 manufacturing and 42 non-manufacturing. All variables in logs and HP filtered. Regressions include industry fixed effects and use industry average value added shares as weights. Standard errors are clustered by year.

which differs from the aggregate labor wedge in three possible ways. Value-added per hour could be more or less cyclical in the BLS industries than for the average industry (the first term on the right-side). The cyclicality of gross output may differ from value added (the second term). Finally, hours worked per worker, and thus the *mrs*, could be more (or less) cyclical for the BLS industries (the third term).

Table 8 presents the cyclical elasticities. The all-industry wedge has a smaller elasticity (-1.11 wrt GDP) than the aggregate IMW (-1.91) from Section 2. Why? Gross output is more procyclical than value-added (0.33) and nominal value-added labor productivity is more procyclical in our BLS industries than in the aggregate (0.36). Industry-specific workweeks also affect the all-industry, manufacturing and nonmanufacturing wedges separately.²⁰ Comparing Tables 7 and 8, the

²⁰Because some of our industries only have workweek data starting in 1990, we use the aggregate average workweek from 1987 through 2011, which had an elasticity with respect to GDP of 0.32 (s.e. 0.03), adjusted by the *relative* elasticity of industry-specific workweeks to the aggregate from 1990 through 2011. These elasticities are 0.34 (s.e. 0.03) for the aggregate, 0.29 (s.e. 0.03) for all KLEMS industries, 0.45 (0.04) for manufacturing, and 0.24 (0.03) for non-manufacturing. Appendix Table 10 reports the cyclicality of the

intermediates-based PMW accounts for most of the cyclical labor wedge. Figure 4 provides visual corroboration by plotting the weighted-average industry (inverse) intermediate share against the total wedge.²¹

Table 8: Cyclicity of Intensive-Margin Total Wedge

	Elasticity wrt	
	GDP	Total Hours
All Industries	-1.11 (0.24)	-0.72 (0.12)
Manufacturing	-0.73 (0.39)	-0.39 (0.19)
Non-Manufacturing	-1.20 (0.22)	-0.80 (0.11)

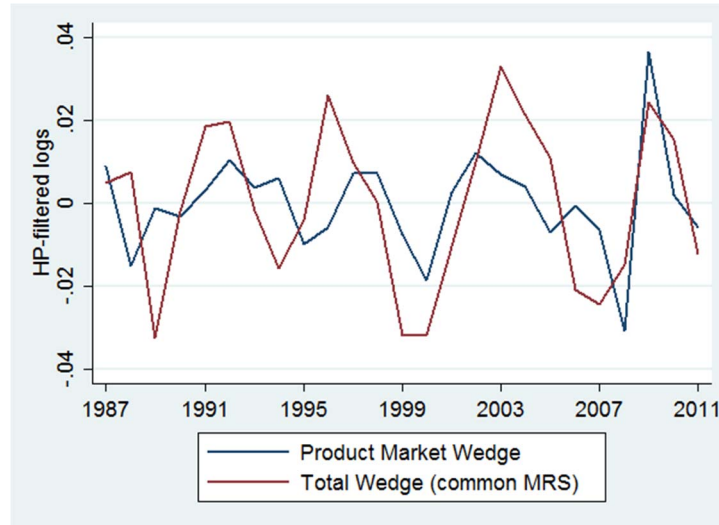
Note: Each entry is from a separate regression. Annual data from 1987-2011 for 60 industries (1500 industry-year observations), 18 manufacturing and 42 non-manufacturing. All variables in logs and HP filtered. Regressions include industry fixed effects and use industry average value added shares as weights. Standard errors are clustered by year.

It is often argued that is tough to substitute between intermediates and value added. Bruno (1984) and Rotemberg and Woodford (1996) estimate elasticities of 0.45 and 0.69, respectively, for U.S. manufacturing. Oberfield and Raval (2014) obtain estimates ranging from 0.63 to 0.90 by looking across regions in U.S. manufacturing. Atalay (2014) estimates even smaller elasticities (below 0.1). A smaller elasticity makes the PMW based on intermediates *more* countercyclical. Because firms are shifting toward intermediates in booms, the marginal product of intermediates will fall faster if substitutability is more limited, making marginal cost more procyclical. Thus the price-cost markups implied by

total wedge using a common workweek for all industries (i.e., we omit the industry-specific adjustments).

²¹The total wedge in Figure 4 (also Figure 6) is constructed using the aggregate average workweek. Using industry-specific workweeks produces similar plots, just with three fewer years.

Figure 4: Intermediates-based PMW vs Total Wedge



intermediate inputs becomes more countercyclical. In fact, for $\varepsilon = 0.78$, the PMW accounts for 100% of the cyclical variation in the total wedge.²²

5. Work-in-Process Inventories

Our third approach to decomposing the labor wedge into product-market and labor-market components exploits data on work-in-process (WIP) inventories. We posit a role for inventories in production. With inventories, firms' production and sales decisions can be separated intertemporally. We then derive a relationship linking a firm's real marginal revenue to the marginal utility of consumption and the expected path of the firm's inventory-to-output ratios. Finally, we infer the PMW from the markup of price over marginal revenue.

Following Christiano (1988), we assume a production function that

²²Note that a smaller ε also makes the total wedge *less* countercyclical because firms shift away from value-added (labor) in booms.

uses WIP inventories as one of its inputs. For a firm in industry i ,

$$y_{it} = g(z_{it}, n_{it}, k_{it})q_{it}^{\varphi_{it}},$$

where y_{it} denotes output, q_{it} is beginning-of-period inventories, and z_{it} , n_{it} , and k_{it} are TFP, hours worked and capital, respectively. The elasticity of output with respect to inventories, φ_{it} , is allowed to vary both across industry and time. The law of motion for inventories is assumed to be

$$q_{i,t+1} = (1 - \delta_q)q_{it} + y_{it} - s_{it},$$

where δ_q is the depreciation rate of inventories, s_{it} is the quantity of finished goods sold, and $q_{i,t+1} \geq 0$. As shown, output can be sold or can augment next period's WIP inventories.²³

A perturbation of the firm's profit-maximizing strategy would be to shift a unit of output out of WIP inventory investment and into current sales. The benefit is extra marginal revenue today, evaluated at the marginal utility from consuming that extra revenue. The cost is curtailing the next period's sales and consumption in order to restore the WIP path. The resulting first-order condition is

$$\frac{mr_{it}}{p_t} = \mathbb{E}_t \left[M_{t,t+1} \frac{mr_{i,t+1}}{p_{t+1}} \left(1 - \delta_q + \varphi_{i,t+1} \frac{y_{i,t+1}}{q_{i,t+1}} \right) \right], \quad (7)$$

where $M_{t,t+1}$ is the firm's (real) discount factor, mr_{it} its marginal revenue, and p_t the price of consumption. The term $\left(1 - \delta_q + \varphi_{i,t+1} \frac{y_{i,t+1}}{q_{i,t+1}} \right)$ equals the reduction in sales next period for each unit pushed to market today. It reflects not only putting back the unit pushed to sales (minus depreciation), but also making up for any reduced production caused by

²³Our discussion omits finished-goods inventories for expositional purposes only. More exactly, the perturbation we consider has a firm shift output toward finishing and selling one more unit, holding finished-goods inventories fixed.

economizing on WIP.

Recall the definition of the (industry) product-market wedge: $\mu_{it}^p \equiv \frac{p_{it}}{mc_{it}} = \frac{p_{it}}{mr_{it}}$, where p_{it} is the price of output and mc_{it} is marginal cost. Using this definition in (7) we have

$$\frac{p_{it}/p_t}{\mu_{it}^p} = \mathbb{E}_t \left[M_{t,t+1} \frac{p_{i,t+1}/p_{t+1}}{\mu_{i,t+1}^p} \left(1 - \delta_q + \varphi_{i,t+1} \frac{y_{i,t+1}}{q_{i,t+1}} \right) \right]. \quad (8)$$

Note that marginal revenue may differ from the price of consumption ($mr_i/p \neq 1$) because of price markups (μ_i^p) or differences between the firm's output price and the price of consumption (p_i/p). Under certain assumptions on the variables in the expectation term of equation (8), we can take logs and iterate forward to obtain

$$\ln(\mu_{it}^p) = -\frac{1}{\sigma} \ln(c_t) + \ln \left(\frac{p_{it}}{p_t} \right) - \mathbb{E}_t \sum_{s=1}^{\infty} \frac{\varphi_{i,t+s} y_{i,t+s}}{1 - \delta_q q_{i,t+s}} + \text{constant terms}, \quad (9)$$

where $M_{t,t+1} = \beta \frac{u'(c_{t+1})}{u'(c_t)}$ and $u'(c_t) = c_t^{-1/\sigma}$.²⁴

The intuition for equation (9) is as follows. Suppose the economy is in a recession in period t , so the log marginal utility of consumption, $-\ln(c_t)/\sigma$, is high. If the firm's price markup and relative price are not cyclical, then (9) says the path of future output-to-inventory ratios must be high. That is, the firm should be depleting future WIP inventories in order to push output out the door today and boost consumption.

Alternatively, if the expected path of output-to-inventory ratios is not cyclical, then for equation (9) to hold, the firm's real marginal revenue (mr_{it}/p_t) must be low in recessions. In turn, either the product-market

²⁴The Appendix contains the derivation of equation (9). To do so, we assume the joint conditional distribution of $\frac{1}{\mu_{i,t+1}^p}$, $\frac{p_{i,t+1}}{p_{t+1}}$, $M_{t,t+1}$, and $1 - \delta_q + \varphi_{i,t+1} \frac{y_{i,t+1}}{q_{i,t+1}}$ is log-normal and homoskedastic. As explained by Campbell (2003), log-normality implies the log of an expectation can be expressed as an expectation of the log plus a variance term. The conditional homoskedasticity means the variance term is not time-varying. The Appendix also explains the empirical construction of each term in equation (9). The expectational term is based on estimated VARs.

wedge (μ_{it}^p) is high or the firm's relative price (p_{it}/p_t) is low in recessions. That is, if firms do not deplete inventory investment in recessions, one explanation is that product market distortions keep the firm's price high relative to its marginal revenue and marginal cost.²⁵

To measure the PMW according to equation (9), we turn to NIPA, which provides quarterly and monthly measures of inventories, sales, and sales price deflators by industry. We define industry output as sales plus the change in (total) inventories, and we use quarterly data from 1987-2011 for comparison to previous sections. WIP inventories are available for 22 (roughly 2-digit) manufacturing industries, but the industry classification changed from the SIC to NAICS system in 1997. To bridge that year and create consistent industry definitions, we aggregate some industries, leaving 14 sectors.

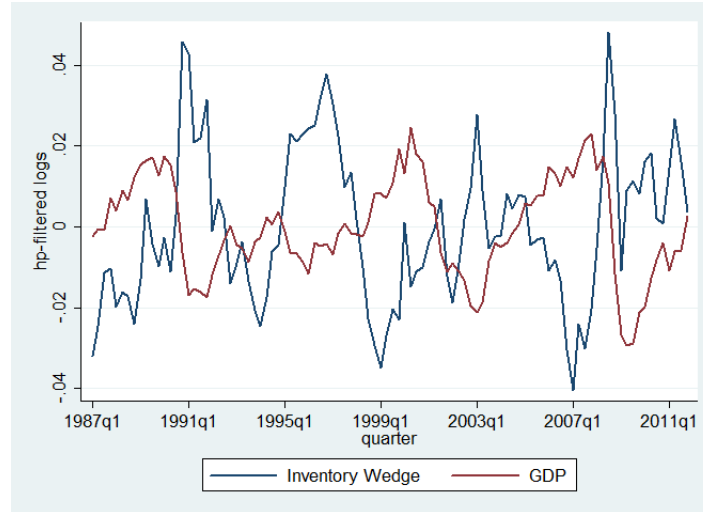
To calibrate the parameters in equation (9), we first note that inventory-to-output ratios exhibited significant low-frequency movement over our sample period. We thus let φ_{it} vary over time and set $\varphi_{it} = \left[\frac{1}{\beta} - (1 - \delta_q) \right] \frac{\bar{q}_{it}}{\bar{y}_{it}}$, where $\frac{\bar{q}_{it}}{\bar{y}_{it}}$ is a quadratic trend fitted to the inventory-output ratio.²⁶ Our quarterly calibration sets $\beta = 0.996$ and $\delta_q = 0.01$. As a result φ_{it} , which measures the share of output attributable to inventories, is quite low, about 0.2%, on average.

Figure 5 plots the weighted-average industry PMW against GDP. As shown, the wedge is quite countercyclical. This is also true if we define the cycle in terms of hours worked. Figure 6 plots the PMW again, but now aggregated to an annual frequency and plotted against the weighted-average manufacturing-industry total wedge constructed in

²⁵In the Appendix, we consider a different perturbation: producing one more unit of output today, increasing WIP investment, and reducing output tomorrow to restore the WIP path, all while leaving the path of finished goods (and sales) unchanged. This replaces equation (7) with one involving marginal cost rather than marginal revenue, but leaves equations (8) and (9) unchanged. It also shows that our PMW derivation does not require assuming $mr = mc$.

²⁶This specification for φ_{it} assures equation (8) holds in (detrended) steady-state.

Figure 5: Cyclicality of Inventory-based PMW



Section 4 rather than against GDP. The PMW accounts for most of the cyclical variation in the total wedge.

We next run regressions of the industry-level wedge on the cycle

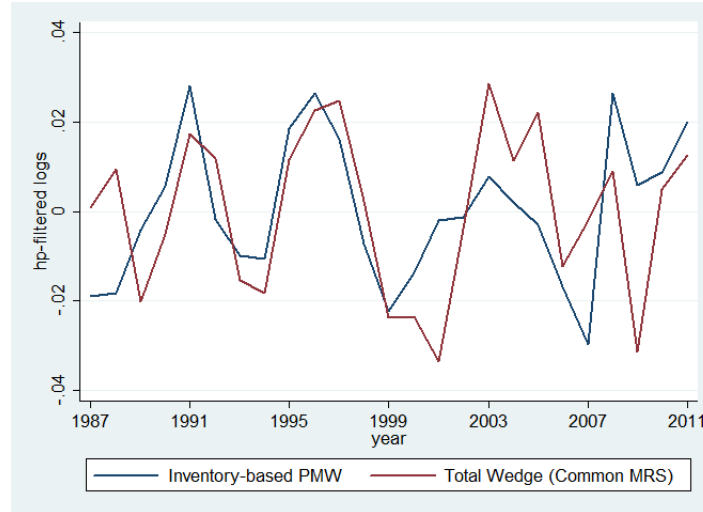
$$\log(\mu_{it}^p) = \alpha_i + \beta^p \log(cyc_t) + \varepsilon_{it},$$

where the weights are the industry's average share of output and standard errors are clustered by period. Table 9 displays the results at an *annual* frequency for comparison to the total wedge (Table 8).²⁷ The strongly countercyclical PMW (-0.70 elasticity with respect to GDP) accounts for nearly all of the cyclical variation in the total wedge (-0.73).

What about the wedge components? The marginal utility of consumption is countercyclical (-1.33 wrt GDP), while relative price movements are procyclical (0.83); in a recession, manufacturers' prices decrease relative to the price of final consumption. Expected future

²⁷The quarterly elasticities are more precisely estimated: -0.80 (s.e. 0.12) with respect to GDP and -0.33 (0.08) with respect to Hours.

Figure 6: Inventory-based PMW vs Total Wedge



output-to-inventory ratios are mildly procyclical: in recessions, manufacturers shift more output into inventories. Rather than depleting their inventory stock when the marginal utility of consumption is high, manufacturers tend to build up inventories. For this to be consistent with optimization, it must be that the PMW is countercyclical.

Finally, we have used WIP inventories for our calculations because these align most closely with the theory, which posits a role for inventories in *production*. Christiano (1988) argues for total inventories (i.e., including materials, WIP, and final goods inventories), noting that labor inputs can be conserved by transporting materials in bulk and holding finished inventories. For robustness, we redo our calculations using total inventories instead of WIP inventories. The results are fairly similar to those reported in Table 9: the cyclical elasticity of the product-market wedge is -0.56 wrt GDP and -0.22 wrt Hours.²⁸

²⁸Using total inventories enables one to consider industries outside of manufacturing (e.g., wholesale and retail trade). In related work, Kryvtsov and Midrigan (2012) focus on finished good inventories and find an important role for countercyclical price markups.

Table 9: **Cyclicality of Inventory-based PMW**

	Elasticity wrt	
	GDP	Total Hours
Product Market Wedge	-0.70 (0.22)	-0.26 (0.13)
Marginal Utility of Consumption	-1.33 (0.06)	-0.76 (0.08)
Relative Price	0.83 (0.20)	0.60 (0.11)
Expected Output/Inventory Path	0.21 (0.05)	0.10 (0.04)

Note: Each entry is from a separate regression. Annual data from 1987-2011 for 14 manufacturing industries (350 industry-years). Variables in logs and HP filtered. Regressions include industry fixed effects and use industry average value added shares as weights. Standard errors clustered by year. See equation (9) for the wedge components.

6. Conclusion

Hours worked fall more in recessions than can be explained by optimal changes in labor supply in response to real labor productivity. This “labor wedge” could reflect frictions in the labor market (e.g. sticky wages and matching problems), frictions in the product market (e.g. sticky prices), or some combination.

Research has increasingly focused on problems in labor markets, in particular for firms hiring workers. Using average hourly earnings, the labor wedge seems to arise between the cost of labor to firms and the value of jobs to workers. But this inference could be mistaken if the true cost of labor to firms is more cyclical than average hourly earnings. If labor’s price is measured by the wages of new hires or a user cost of labor, instead of by average hourly earnings, the labor wedge arises as much between the cost of labor and real labor productivity.

To bring new evidence to bear on this debate, we estimate the

product market component of the labor wedge in three ways that do not rely on workers' wages. First, we look only at the self-employed. A cyclical labor wedge appears for the self-employed as much as for workers, even though sticky wages and matching frictions should not be barriers to the self-employed working more hours. The hours of the self-employed appear to fall in recessions in no small part because of difficulty, or reluctance, in selling their output (for example due to sticky prices). The second piece of evidence we present is for intermediate inputs. We find that output prices rise relative to the marginal cost of producing by buying more intermediates in recessions. Again, this suggests that firms face difficulty converting production into revenue in recessions. Finally, we present evidence on the cyclical nature of work-in-process (WIP) inventories. In recessions, firms build up WIP inventories relative to output, which implies that marginal revenue is low relative to price in recessions.

We estimate that at least three-quarters of the labor wedge's cyclical variation reflects product market distortions as opposed to labor market distortions. While labor market distortions matter, they are less important than has been inferred using average hourly earnings. Our evidence is consistent with a price of labor that is at least as cyclical as the new hire wage.

Our evidence does not determine the exact nature of these product market distortions, which is critical for informing stabilization policy. Our results do suggest, however, that in recessions firms have trouble selling their output — as if their shadow value of output is low relative to its market price. Why would firms' shadow value of output be pushed down in recessions? One explanation is price stickiness that constrains production from translating into added sales. Output's shadow value also falls relative to price in models of countercyclical desired markups.

(See Rotemberg and Woodford, 1999, for a review.) And in any setting where producing at the margin has an investment component (e.g., the customer base model in Gilchrist et al, 2014). Our evidence is also consistent with models where expanding production puts firms in a riskier position, and risk (or risk avoidance) heightens during recessions (e.g., Arellano et al, 2012).

7. Appendix

This appendix provides a detailed description of our data, a more thorough explanation of some calculations, and some robustness results.

7.1. Representative-Agent Wedge

Variables include:

- y_t : (Real) Output per hour; BLS, Labor Productivity & Costs, Business Sector.
- c_t : (Real) Nondurables and services consumption per adult equivalent; NIPA consumption data, adjusted for indirect taxes following Prescott (2004, Mpls Fed QR) and McDaniel (2007). Adult-equivalent population = Population ≥ 16 + 0.5 · (Population ≤ 15).
- n_t : Hours worked per capita; Hours worked from BLS (LPC, Business Sector), and population is Civilian Pop 16+.
- $\tau_t = ((\tau_t^c + \tau_t^n)/(1 + \tau_t^c))$, where τ_t^c is the average tax rate on consumption, following McDaniel (2007), and τ_t^n is the average marginal labor tax rate, using NBER TaxSim to extend Barro-Redlick (2011) through 2012.

7.2. Extensive- and Intensive-Margin Wedges

To construct the IMW and EMW, some variables (e.g., y_t and c_t) are the same as used for the RAW. Additional variables include:

- h_t : Average weekly hours worked (per worker); BLS, LPC, Business Sector.
- r_{t+1} : (Ex-post) Real Interest Rate; 3-mo Tbill at t less (realized) Core PCE inflation at $t + 1$.
- v_t : Vacancies (per capita); Pre-1995 is help-wanted index, and post-1995 is Barnichon's (2010) spliced series of help-wanted and JOLTS. Population is 16+.
- m_t : Matches (per capita); Post-1994 from Fallick-Fleischman (2004), and pre-1994 is backcast using data on unemployment and vacancies, following Blanchard and Diamond (1989).

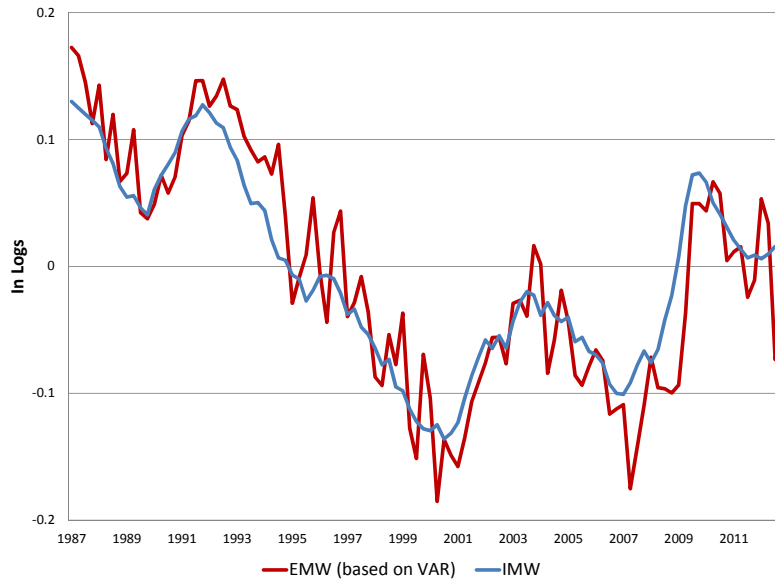
Note that all variables, except the 3-mo Tbill, are seasonally adjusted.

The calibration is described in the text with the exception of ψ , the fixed (utility) cost of employment. One can derive an expression for ψ by combining the steady-state optimality conditions for the extensive and intensive margins and assuming the EMW and IMW are the same in steady state. The result is that $\psi \equiv \frac{h^{1+1/\eta}}{\eta+1} \left[1 - (\eta + 1) \frac{r+\delta}{1+r} \left[\frac{\kappa v}{\phi m} + \gamma \right] \right]$.

Finally, the EMW includes expectational terms in S_t , e.g., $\mathbb{E}_t \left\{ \frac{1}{1+r_{t+1}} \frac{y_{t+1}}{y_t} \right\}$. We construct these using 3-variable, 4-lag VARs consisting of real GDP growth, aggregate (log) hours worked, and the respective expectational term. We estimate the VAR using data over the entire sample period, and then use the estimated coefficients to construct time series of the expectational terms.

Figure 7 displays the unfiltered extensive and intensive margin wedges from 1987–2012.

Figure 7: EMW vs IMW (Unfiltered)



7.3. Aggregate Wedge Decomposition

The decomposition requires wage measures. For our baseline (labeled AHE), we assume $\frac{w_t}{p_t y_t}$ is the labor share of income as measured in the BLS's LPC Business Sector. Because we also have a series for labor productivity y_t , we can back out the average real wage in the economy.

Kudlyak (2013) estimated the semi-elasticities of average hourly earnings, new hire wages, and the user cost of labor, respectively, to the unemployment rate. We use these estimated elasticities, along with the time series of unemployment and our (baseline) average wage measure, to construct time series for new hire wages and the user cost of labor.

7.4. Self-Employed

Figures 8–11 display time series of the variables that underlie our estimates of the cyclicity of the all-worker and self-employed wedges, respectively. Figure 8 displays HP-filtered (log) indices for hours per

Figure 8: Weekly Hours: Self-employed vs. Wage-earners



week for both the self-employed and wage-earners, while Figure 9 presents the same comparison for annual hours. Figure 10 presents HP-filtered aggregate labor productivity, self-employed earnings per hour, and earnings per hour for the unincorporated self-employed. We construct the time series in these three figures using data from the March CPS, as described in Section 3 of the paper.

Figure 11 presents our consumption series for the self-employed together with aggregate consumption, both HP-filtered. To construct consumption for the self-employed, we use the Consumer Expenditure Surveys (CE) from 1987 through 2012 to get a quarterly series for the growth rate of consumption of self-employed workers relative to that for a representative sample of households in the CE. The relative growth rate, in turn, is integrated to obtain a series for relative consumption of the self-employed, indexed to the beginning of 1987. We add this relative estimate to NIPA aggregate consumption to arrive at an estimate of the

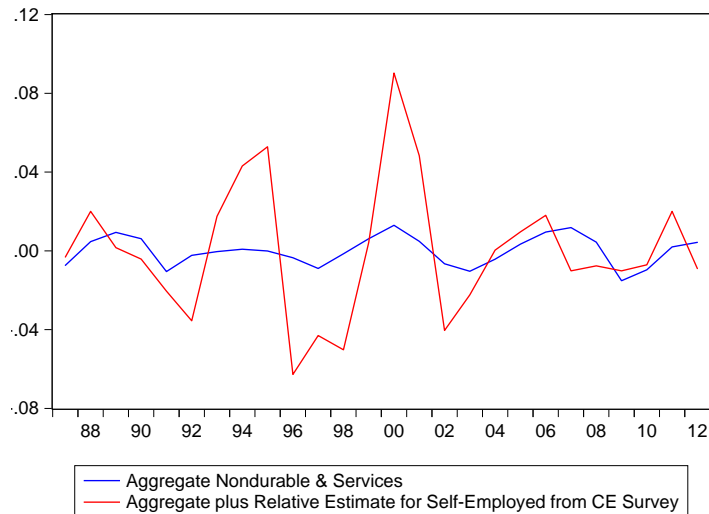
Figure 9: Annual Hours: Self-employed vs. Wage-earners



Figure 10: Alternative Productivity Measures



Figure 11: Alternative Consumption Measures



cyclicality of consumption for the self-employed. The following paragraphs describe the construction of the quarterly growth rates of consumption for self-employed workers.

The CE has been an ongoing quarterly survey since 1980, with about 5000 households interviewed each quarter. Households are surveyed about their detailed expenditures for the previous three months. Each household is surveyed for up to four consecutive quarters, allowing for construction of up to three observations on quarterly growth. We focus on expenditures on non-durables and services, which we construct by aggregating those individual categories that are clearly not durables by NIPA standards. We include expenditures on housing: for renters this is captured by household rent; for home-owners it reflects the owner's estimate of its rental value (rental equivalence). The categories we can classify as nondurables and services constitute about two thirds of household expenditures. We deflate these expenditures by the GDP deflator for nondurables and services. Individual growth rates across any

two quarters are calculated by the midpoint formula to reduce the impact of extreme values.

During each of the first and fourth quarterly interviews on expenditures, households are surveyed about the work experience of its household members during the past 12 months. We focus on the work history in the latter survey, as the work history over the prior 12 months conforms to the time-frame for reported expenditures. (For a small number of households, we fill in for missing employment information from responses collected in earlier quarters.) We create a sample of workers from the CE households, including all members that meet our sample requirements. These requirements are chosen to mimic our treatment of the CPS data: (i) Individuals must be between ages 20-70; (ii) They must report working at least 10 weeks during the year, at a workweek of 10 hours or more when working; (iii) We exclude workers in the top or bottom 9.6 percent of the earnings distribution and the top 1.2 percent of hours per week. These last exclusions are chosen to match those we made on the CPS data, dictated by its top coding of earnings and hours. We make two other sample restrictions in order to measure quarterly growth rates of household consumption. We exclude households in the top and bottom one percent of expenditures in any quarter in order to eliminate top-coded expenditures and outliers. We exclude households that exhibited a change in household size across the quarters that are the basis for the growth rate. In all calculations we employ the CE sampling weight that is designed to make the sample representative of the U.S. civilian non-institutionalized population.

We classify workers as self-employed, as opposed to wage-earners, if they report that the job for which they received most earnings was self-employment and, in fact, at least 95 percent of their reported earnings over the past 12 months is from (non-farming)

self-employment. This conforms well to our definition in Section 3 based on CPS data. We do not observe consumption at the individual level (e.g., for a self-employed member versus a wage-earning member). Thus we have to make the simplifying assumption that households equate consumption across members. For example, if a household has one self-employed worker and one wage earner, then that household contributes two members to our overall sample, and one member to our self-employed sample. But the growth rate in consumption in any quarter will be the same for both members of that household. We have 11,849 quarterly observations on consumption growth that apply for self-employed workers, which equals 115 per quarter on average.

7.5. Intermediates

We first derive an industry-level, intensive-margin total wedge using more general technology and preferences than we used for our baseline results. We then derive the industry-level, extensive-margin total wedge.

The *gross-output* production function implies a marginal product of labor on the intensive margin of

$$mpn_{it}^h = \alpha(1 - \theta) \left(\frac{y_{it}}{v_{it}} \right)^{\frac{1}{\varepsilon}} \left(\frac{v_{it}}{n_{it}} \right)^{\frac{1}{\omega}} (z_{v,it} z_{n,it})^{\frac{\omega-1}{\omega}} e_{it}.$$

For our baseline, $\varepsilon = \omega = 1$, this simplifies to $mpn_{it}^h = \alpha(1 - \theta) \frac{y_{it}}{n_{it}} e_{it}$.

Our baseline used the following preferences:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \frac{c_t^{1-1/\sigma}}{1-1/\sigma} - \nu \sum_i \left[\left(\frac{h_{it}^{1+1/\eta}}{1+1/\eta} + \psi \right) e_{it} \right] \right\},$$

so the marginal rate of substitution of consumption for an extra hour per worker in industry i is $mr s_{it}^h = \nu h_{it}^{1/\eta} c_t^{1/\sigma} e_{it}$.

Thus, our baseline industry- i (intensive-margin) labor wedge is (up to

an additive constant)

$$\begin{aligned} \ln(\mu_i^h) &= \ln\left(\frac{p_i m p n_i^h}{p m r s_i^h}\right) = \ln\left(\frac{p_i y_i}{p n_i}\right) - \left[\frac{1}{\sigma} \ln(c) + \frac{1}{\eta} \ln(h_i)\right] \\ &= \ln\left(\frac{p_i \frac{v_i}{n_i}}{p \frac{v}{n}}\right) + \ln\left(\frac{y_i}{v_i}\right) - \frac{1}{\eta} \ln\left(\frac{h_i}{h}\right) + \ln\left(\frac{m p n^h}{m r s^h}\right), \end{aligned} \quad (10)$$

where $m p n_t^h \equiv \alpha(1 - \theta) \frac{v_t}{n_t} e_t$ and $m r s_t^h \equiv \nu h_t^{1/\eta} c_t^{1/\sigma} e_t$ are based on aggregate data. For $\varepsilon, \omega \neq 1$, it's straightforward to see how the total wedge would be altered. Specifically, for $\varepsilon < 1$, the total wedge becomes less countercyclical if gross output is more procyclical than value-added.

Note that our preferences assume separability across labor supply in different industries. This seems reasonable because workweeks are person-specific. But, we could consider alternative preferences, say

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \frac{c_t^{1-1/\sigma}}{1-1/\sigma} - \nu \left(\frac{h_t^{1+1/\eta}}{1+1/\eta} + \psi \right) e_t \right\},$$

where $h_t \equiv \frac{\sum_i h_{it} e_{it}}{\sum_i e_{it}}$ and $e_t \equiv \sum_i e_{it}$. In this case, $m r s_{it}^h = \nu h_t^{1/\eta} c_t^{1/\sigma} e_{it}$. The industry- i labor wedge is thus (for baseline technology, $\varepsilon = \omega = 1$)

$$\ln(\mu_i^h) = \ln\left(\frac{p_i \frac{v_i}{n_i}}{p \frac{v}{n}}\right) + \ln\left(\frac{y_i}{v_i}\right) + \ln\left(\frac{m p n^h}{m r s^h}\right).$$

Under these preferences, labor supply is perfectly substitutable across industries and only aggregate labor supply, e and h , matters. The labor wedge no longer needs an adjustment for industry-specific workweeks. This matters little for our “all-industry” results, but because manufacturing industries exhibit more procyclical workweeks than other industries, the labor wedge is noticeably less countercyclical for manufacturing, as shown in Table 10 (compare to Table 8).

On the extensive margin, $m p n_{it}^e = m p n_{it}^h \frac{h_{it}}{e_{it}}$ and $m r s_{it}^e = m r s_{it}^h \frac{\Omega_{it}}{h_{it}^{1/\eta}} \frac{h_{it}}{e_{it}}$.

Table 10: **Cyclicality of (Common-MRS) Intensive-Margin Total Wedge**

	Elasticity wrt	
	GDP	Total Hours
All Industries	-1.21 (0.25)	-0.79 (0.13)
Manufacturing	-0.51 (0.42)	-0.32 (0.22)
Non-Manufacturing	-1.39 (0.22)	-0.90 (0.11)

Note: Each entry is from a separate regression. Annual data from 1987 - 2011 for 60 industries (1500 industry-year observations), 18 manufacturing and 42 non-manufacturing. All variables in logs and HP filtered. Regressions include industry fixed effects and use industry average value added shares as weights. Standard errors are clustered by year.

We can thus write the industry- i extensive-margin labor wedge as

$$\ln(\mu_i^e) = \ln\left(\frac{p_i m p n_i^e}{p m r s_i^e}\right) - S_i = \ln(\mu_i^h) - \ln\left(\frac{\Omega_i}{h_i^{1/\eta}}\right) - S_i,$$

or, for our baseline case, as

$$\ln(\mu_i^e) = \ln\left(\frac{p_i \frac{v_i}{n_i}}{p \frac{v}{n}}\right) + \ln\left(\frac{y_i}{v_i}\right) - \ln\left(\frac{\Omega_i}{\Omega}\right) + \ln(\mu^e) - (S_i - S). \quad (11)$$

Due to data limitations (i.e., vacancies and matches are not available at the industry-level), S_{it} differs across industry only because of industry-specific workweek movements. That is, $S_{it} = \left[\frac{h_i/h_{it}}{h/h_t}\right] S_t$. Table 11 displays the cyclicality of the extensive-margin total wedge.

To construct the industry-level total wedge and the intermediates-based product market wedge, some variables (e.g., c_t) are the same as used earlier in the paper. Additional variables include:

- p_t : Price deflator for nondurables and services consumption;
- Tornqvist index of NIPA implicit price deflators for nondurables

Table 11: **Cyclicality of (Common-MRS) Extensive-Margin Total Wedge**

	Elasticity wrt	
	GDP	Total Hours
All Industries	-1.18 (0.51)	-0.89 (0.27)
Manufacturing	-0.48 (0.69)	-0.43 (0.38)
Non-Manufacturing	-1.35 (0.47)	-1.01 (0.25)

Note: Each entry is from a separate regression. Annual data from 1987-2011 for 60 industries (1500 industry-year observations), 18 manufacturing and 42 non-manufacturing. All variables in logs and HP filtered. Regressions include industry fixed effects and use industry average value added shares as weights. Standard errors are clustered by year.

consumption and services consumption.

- $p_{it}, y_{it}, n_{it}, p_{mit}m_{it}$: Respectively, the gross output deflator, real gross output, hours worked, and expenditures on intermediates (tornqvist index of materials, energy and services) by industry from BLS KLEMS.
- h_{it} : Average weekly hours worked (per worker); Ratio of hours worked (from BLS KLEMS) to industry-specific employment (calculated with data underlying BLS LPC dataset).

7.6. Inventories

We first present an alternative perturbation exercise as a basis for our WIP-inventory-based PMW. It clarifies that we do not need to assume that firms can increase sales at the margin or that marginal revenue equals marginal cost. These assumptions were made in the main text for expositional purposes only. We then provide a step-by-step derivation of the PMW, filling in details left out of the main text.

We begin with a more general expression for the law of motion for WIP inventories

$$q_{i,t+1} = (1 - \delta_q)q_{it} + y_{it} - y_{it}^f,$$

where $y_{it}^f \geq 0$ is output of finished goods. That is, total output y_{it} is the sum of gross investment in WIP inventories and finished-good output, which includes both final sales and (gross) investment in finished-goods inventories. For our analysis, it will not be necessary to separate these latter two.

An optimizing firm minimizes the expected present discounted cost of producing a given path of finished goods. One perturbation on its cost-minimizing strategy would be to produce an additional unit of output in the form of WIP inventories at time t and then reduce production just enough at $t + 1$ — i.e., by $\left(1 - \delta_q + \varphi_{i,t+1} \frac{y_{i,t+1}}{q_{i,t+1}}\right)$ — to keep inventories unaffected at $t + 2$ forward. At an optimum,

$$\frac{mc_{it}}{p_t} = \mathbb{E}_t \left[M_{t,t+1} \frac{mc_{i,t+1}}{p_{t+1}} \left(1 - \delta_q + \varphi_{i,t+1} \frac{y_{i,t+1}}{q_{i,t+1}} \right) \right], \quad (12)$$

where $\frac{mc_{it}}{p_t}$ is the firm's (real) marginal cost of production and $M_{t,t+1}$ its real discount factor. In words, the cost-minimizing firm equates the marginal cost of output to the marginal benefit, which is reduced future production costs.²⁹ Because the industry product-market wedge is defined as $\mu_{it}^p \equiv \frac{p_{it}}{mc_{it}}$, equation (12) provides the same expression for the inventory-based

²⁹Note that an optimizing firm will *always* produce to the point that the marginal value of an extra unit of output equals its marginal cost. This principle is more general than the assumption used in the main text; namely, that marginal revenue equals marginal cost. If the firm can adjust sales at the margin (as considered in the main text), then the marginal value of output is simply the marginal revenue generated by an additional sale. But, if the firm cannot sell an additional unit, the additional unit of output is held as an inventory and is valued accordingly. The marginal value of a finished-good inventory is the expected discounted revenue it will generate when it is sold in the future. Finally, if the output takes the form of WIP, then it cannot be sold in the current period but it still has value — the firm enters the next period with a larger stock of WIP inventories.

PMW that was derived in the main text.

A step-by-step derivation of the inventory-based PMW is as follows. We start from

$$\frac{1}{\mu_{it}^p} \frac{p_{it}}{p_t} = \mathbb{E}_t \left[\beta \frac{u'(c_{t+1})}{u'(c_t)} \frac{1}{\mu_{i,t+1}^p} \frac{p_{i,t+1}}{p_{t+1}} \left(1 - \delta_q + \varphi_{i,t+1} \frac{y_{i,t+1}}{q_{i,t+1}} \right) \right].$$

Assuming that the joint conditional distribution of $\frac{u'(c_{t+1})}{u'(c_t)}$, $\frac{1}{\mu_{i,t+1}^p}$, $\frac{p_{i,t+1}}{p_{t+1}}$, and $1 - \delta_q + \varphi_{i,t+1} \frac{y_{i,t+1}}{q_{i,t+1}}$ is log-normal and homoskedastic, we can take logs and get (up to a constant)

$$\ln(\mu_{it}^p) \approx \ln\left(\frac{p_{it}}{p_t}\right) + \mathbb{E}_t \left\{ -\ln\left(\frac{u'(c_{t+1})}{u'(c_t)}\right) + \ln(\mu_{i,t+1}^p) - \ln\left(\frac{p_{i,t+1}}{p_{t+1}}\right) - \frac{\varphi_{i,t+1}}{1 - \delta_q} \frac{y_{i,t+1}}{q_{i,t+1}} \right\},$$

where $\frac{\varphi_{i,t+1}}{1 - \delta_q} \frac{y_{i,t+1}}{q_{i,t+1}} \approx \ln\left(1 + \frac{\varphi_{i,t+1}}{1 - \delta_q} \frac{y_{i,t+1}}{q_{i,t+1}}\right)$. Iterating forward for $\ln(\mu_{i,t+s}^p)$ yields the inventory-based PMW in the paper:

$$\ln(\mu_{it}^p) \approx -\frac{1}{\sigma} \ln(c_t) + \ln\left(\frac{p_{it}}{p_t}\right) - \mathbb{E}_t \sum_{s=1}^{\infty} \frac{\varphi_{i,t+s}}{1 - \delta_q} \frac{y_{i,t+s}}{q_{i,t+s}} + \text{constant terms.} \quad (13)$$

How do we construct $\ln(\mu_{it}^p)$? Some variables (e.g., c_t and p_t) are the same as earlier in the paper. Additional variables include:

- $\frac{y_{it}}{q_{it}}$: (Real) Output-to-WIP-Inventory Ratio; NIPA Underlying Detail Tables, Real Inventories and Sales. Output = Sales + change in (total) inventories.
- p_{it} : Price deflator for (industry) sales; NIPA Underlying Detail Tables, Real Inventories and Sales.

Work-in-process inventories are only available for manufacturing industries, 11 durable-producing and 10 nondurable-producing. NIPA's industry classification switched from SIC to NAICS in 1997. The inventory data is reported for both classifications in 1997, but there is no

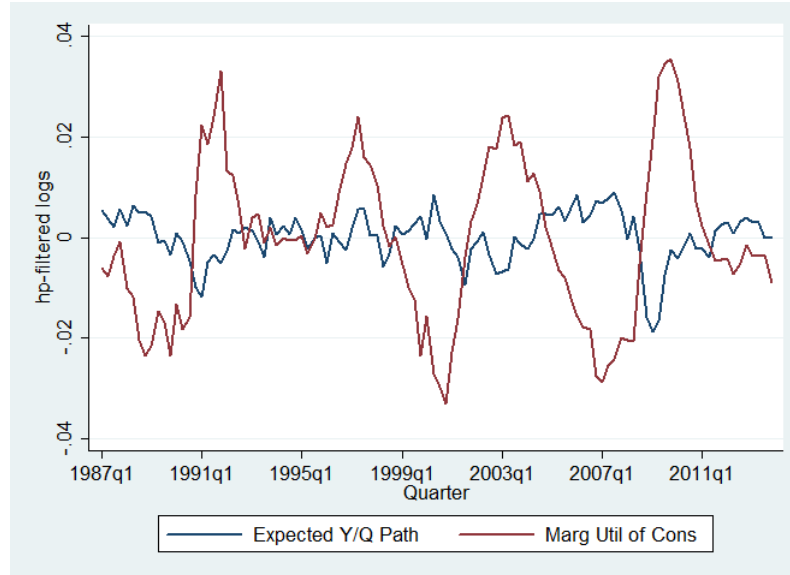
overlap for sales. To have consistently defined categories across time, we combine some industries, resulting in 14 sectors. We use a Tornqvist index to construct chain-weighted growth rates of real sales, real inventories, and price deflators for the combined industries.

For bridging across the 1996-97 break, we made two assumptions. For inventories, we assume the industry shares of nominal inventories don't change between Dec 96 and Jan 97. For sales, we assume the growth rate in the nominal inventory-to-shipments ratio is the same as that of the real inventory-to-shipments ratio (in Jan 97). The former is constructed using data from the Census M3 survey, which has a consistent NAICS industry classification across 96-97.

Constructing the inventory-based wedge requires computing, at each point in time, the sum of expected future output-to-inventory ratios. We estimate industry-specific, 3-variable, 12-(monthly)-lag VARs consisting of real GDP growth, aggregate (log) hours worked, and the industry-specific output-to-inventory ratio. The latter two variables are quadratically detrended. We estimate the VARs using data over the entire sample period, and then use the estimated coefficients to produce a time series for the expected sum of future output-to-inventory ratios.

Figure 12 plots the hp-filtered time series of the two components we use to infer real marginal cost: namely, the (log) marginal utility of consumption, $-\frac{1}{\sigma} \ln(c_t)$, and (log) expected path of (gross) returns to WIP-inventory investment, $\mathbb{E}_t \sum_{s=1}^{\infty} \frac{\varphi_{i,t+s}}{1-\delta_q} \frac{y_{i,t+s}}{q_{i,t+s}}$. See equation (13). We find the cyclicity of real marginal cost is largely determined by that of the marginal utility of consumption, which is strongly countercyclical. Future returns to WIP-inventory investment are actually slightly *procyclical* – i.e., although inventories decline in recessions, output declines more so that the marginal product of inventories declines – thus reinforcing the countercyclicity of real marginal cost. The calculation

Figure 12: Components of Real Marginal Cost



of the price-cost markup, described by equation (13), also requires accounting for relative price movements which are procyclical (as described in the paper).

Finally, we also considered a second approach to calculating the expected sum of future output-to-input ratios, which involved truncating the sum at either 4 or 8 quarters and calculating the (ex-post) realized sum. Because we project the constructed wedge on the time- t business cycle, using the (ex-post) realized values is valid for our purposes. It does require using a 1-sided hp-filter for the business cycle, so the difference between expected and realized values of the output-to-inventory ratios is orthogonal to the time- t cycle. This second approach produced results for the wedge that were very similar to the VAR approach.

References

- Aguiar, Mark and Erik Hurst, “Deconstructing Lifecycle Expenditure,” *Journal of Political Economy*, 2013, 121 (3), 437 – 492.
- Arellano, Cristina, Yan Bai, and Patrick Kehoe, “Financial Markets and Fluctuations in Volatility,” July 2012.
- Atalay, Engin, “How Important are Sectoral Shocks?,” June 2014.
- Barnichon, Regis, “Building a composite Help-Wanted Index,” *Economic Letters*, 2010, 109, 175 – 178.
- Barro, Robert J. and Charles J. Redlick, “Macroeconomic Effects from Government Purchases and Taxes,” *Quarterly Journal of Economics*, February 2011.
- Barron, John M, Mark C Berger, and Dan A Black, “Do Workers Pay for On-the-Job Training?,” *Journal of Human Resources*, 1999, 34 (2), 235–252.
- Basu, Susanto, “Intermediate Goods and Business Cycles: Implications for Productivity and Welfare,” *American Economic Review*, 1995, 85 (3), 512–531.
- Bils, Mark, “The Cyclical Behavior of Marginal Cost and Price,” *American Economic Review*, 1987, 77 (5), 838 – 855.
- and James A. Kahn, “What Inventory Behavior Tells Us About Business Cycles,” *American Economic Review*, 2000, 90 (3), 458–481.
- , Peter J. Klenow, and Benjamin A. Malin, “Testing for Keynesian Labor Demand,” in D. Acemoglu, M. Woodford, and J. Parker, eds., *NBER Macroeconomics Annual 2012*, pp. 515–563.
- Blanchard, Olivier and Jordi Gali, “Labor Markets and Monetary Policy: A New Keynesian Model with Unemployment,” *American Economic Journal: Macroeconomics*, 2010, 2 (2), 1–30.

- and Peter Diamond, “The Beveridge Curve,” *Brookings Papers on Economic Activity*, 1989, 1, 1–76.
- Bruno, Michael, “Raw Materials, Profits and the Productivity Slowdown,” *Quarterly Journal of Economics*, 1984, 99 (1), 1–30.
- Campbell, John Y., “Consumption-Based Asset Pricing,” in G.M. Constantinides, M. Harris, and R. Stulz, eds., *Handbook of the Economics of Finance*, Vol. 1B, 2003, chapter 13, pp. 803–887.
- Chang, Yongsung, Sun-Bin Kim, Kyooho Kwon, and Richard Rogerson, “Individual and Aggregate Labor Supply in a Heterogeneous Agent Economy with Intensive and Extensive Margins,” April 2014.
- Chari, V.V., Patrick Kehoe, and Ellen McGrattan, “Business Cycle Accounting,” *Econometrica*, 2007, 75 (3), 781 – 836.
- Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber, “Does Indivisible Labor Explain the Difference between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities,” in D. Acemoglu, M. Woodford, and J. Parker, eds., *NBER Macroeconomics Annual 2012*, pp. 1–56.
- Christiano, Lawrence J., “Why Does Inventory Investment Fluctuate so Much?,” *Journal of Monetary Economics*, 1988, 21, 247 – 280.
- Fallick, Bruce and Charles A. Fleischman, “Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows,” May 2004. FEDS 2004-34.
- Galeotti, Marzio and Fabio Schiantarelli, “The Cyclicalities of Markups in a Model with Adjustment Costs: Econometric Evidence for US Industry,” *Oxford Bulletin of Economics and Statistics*, 1998, 60 (2), 121 – 142.
- Galí, Jordi, “Monetary Policy and Unemployment,” in Benjamin M. Friedman and Michael Woodford, eds., *Handbook of Monetary Economics*, Vol. 3A, 2011, chapter 10, pp. 487 – 546.

- , Mark Gertler, and J. David López-Salido, “Markups, Gaps, and the Welfare Costs of Business Fluctuations,” *Review of Economics and Statistics*, 2007, 89 (1), 44–59.
- Gilchrist, Simon, Raphael Schoenle, Jae Sim, and Egon Zakrajsek, “Inflation Dynamics During the Financial Crisis,” January 2014.
- Haefke, Christian, Marcus Sonntag, and Thijs van Rens, “Wage Rigidity and Job Creation,” *Journal of Monetary Economics*, 2013, 60 (8), 887–899.
- Hall, Robert E., “Macroeconomic Fluctuations and the Allocation of Time,” *Journal of Labor Economics*, 1997, 15 (1), S223–250.
- , “Reconciling Cyclical Movements in the Marginal Value of Time and the Marginal Product of Labor,” *Journal of Political Economy*, 2009, 117 (2), 281–323.
- , “What the Cyclical Response of Advertising Reveals about Markups and other Macroeconomic Wedges,” April 2014.
- Heim, Bradley T., “Understanding Self-Employment Dynamics Among Individuals Near Retirement,” April 2014.
- Hipple, Steven F., “Self-employment in the United States,” *BLS Monthly Labor Review*, 2010, 133 (9), 17 – 32.
- Hurst, Erik, Geng Li, and Ben Pugsley, “Are Household Surveys Like Tax Forms? Evidence from Income Underreporting of the Self Employed,” *Review of Economics and Statistics*, 2014, 96 (1), 19 – 33.
- Karabarbounis, Loukas, “The Labor Wedge: MRS vs. MPN,” *Review of Economic Dynamics*, 2014, 17 (2), 206 – 223.
- Kryvtsov, Oleksiy and Virgiliu Midrigan, “Inventories, Markups, and Real Rigidities in Menu Cost Models,” *Review of Economic Studies*, 2012, 80 (1), 249–276.

- Kudlyak, Marianna, “The Cyclical of the User Cost of Labor,” 2013. Richmond Fed Working Paper 09-12.
- Levinsohn, James and Amil Petrin, “Estimating Production Functions Using Inputs to Control for Unobservables,” *Review of Economic Studies*, 2003, 10 (2), 317–342.
- McDaniel, Cara, “Average Tax Rates on Consumption and Investment Expenditures and Labor and Capital Income in the OECD 1950-2003,” 2007.
- Mulligan, Casey B., “A Century of Labor-Leisure Distortions,” 2002. NBER WP 8774.
- , *The Redistribution Recession: How Labor Market Distortions Contracted the Economy*, Oxford University Press, 2012.
- Nakamura, Emi and Jón Steinsson, “Monetary Non-Neutrality in a Multi-Sector Menu Cost Model,” *Quarterly Journal of Economics*, 2010, 125 (3), 961–1013.
- Nekarda, Christopher J. and Valerie A. Ramey, “The Cyclical Behavior of the Price-Cost Markup,” 2013. NBER WP 19099.
- Oberfield, Ezra and Devesh Raval, “Micro Data and Macro Technology,” May 2014.
- Pescatori, Andrea and Murat Tasci, “Search Frictions and the Labor Wedge,” August 2012. Federal Reserve Bank of Cleveland Working Paper 11-11R.
- Pissarides, Christopher A., “The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?,” *Econometrica*, 2009, 77 (5), 1339–1369.
- Prescott, Edward C., “Why Do Americans Work So Much More Than Europeans,” *Federal Reserve Bank of Minneapolis Quarterly Review*, 2004, 28 (1), 2 – 13.
- Rotemberg, Julio J. and Michael Woodford, “Imperfect Competition and the Effects of Energy Price Increases on Economic Activity,” *Journal of Money, Credit and Banking*, 1996, 28 (4), 550–77.

— and —, “The Cyclical Behavior of Prices and Costs,” in J.B. Taylor and M. Woodford, eds., *Handbook of Macroeconomics*, Vol. 1B, 1999, pp. 1051–1135.

Saporta-Eksten, Itay, “Job Loss, Consumption and Unemployment Insurance,” January 2014.

Shimer, Robert, “The Cyclical Behavior of Equilibrium Unemployment and Vacancies,” *American Economic Review*, 2005, 95 (1), 25–49.

—, “Convergence in Macroeconomics: The Labor Wedge,” *American Economic Journal: Macroeconomics*, 2009, 1 (1), 280–297.