

Is the labor wedge due to rigid wages? Evidence from the self-employed *

Trevor Gallen[†]

May 2017

Abstract

A central goal of labor economics is explaining cyclical variation in hours worked. Procyclical hours can always be explained in a market clearing model with a residual tax wedge, the “labor wedge.” Convincing progress has been made in reducing the cyclical volatility in the labor wedge and therefore explaining movements in hours worked by amplifying technology shocks with endogenously rigid wages (Hall, 2005; Shimer, 2010). This paper establishes a problem with this method: the self-employed, who are not vulnerable to the same frictions as wage and salary workers, and who are not subject to the same degree of wage stickiness, have labor hours which are of comparable volatility as those of wage-workers, even conditional on wages, consumption, and occupation-industry worker composition.

JEL Classification: E24, E32, J20, L26

Keywords: Self-employment, labor wedge, business cycles, labor markets

*I'd like to thank Andrew Compton for excellent research assistance, and seminar participants at Purdue for helpful suggestions and comments. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

[†]Correspondence: Purdue University Krannert School of Business, West Lafayette, IN 47906. Tel.: (765) 496-2458. Email: tgallen@purdue.edu. Web: web.ics.purdue.edu/~tgallen

1 Introduction

By examining the labor market behavior of the self-employed, this paper challenges recent research arguing rigid wages and labor market search matching frictions are the leading contender to explain cyclical behavior in labor hours. Some, Hall (2005) and Shimer (2010) prominently among them, make the case that endogeneously rigid wages along with labor market matching frictions explain a large part of procyclical movements in labor hours. This paper uses the fact that many self-employed are residual claimants on firm profits and therefore are not as subject to wage rigidity to test this mechanism. I find the self-employed have per-capita labor hour fluctuations of comparable magnitude and cyclicity as wage workers. I find the strong procyclicality in self-employed labor hours is robust across datasets, present not only in the Current Population Survey (CPS), but also during the Great Recession in the American Community Survey (ACS). The labor hours of the self-employed are as cyclical as those of wage workers even conditional on industry-occupation composition and are not explained by differences in shocks to permanent income as measured by consumption and income behavior.

An important question in most business-cycle labor market models is to what degree the marginal rate of substitution between consumption and leisure (MRS) is less than or equal to the marginal product of labor (MPL). The difference between log MRS and log MPL defines the labor wedge—the implied tax wedge necessary for labor market equilibrium, in the same way that the better known Solow residual reconciles capital, labor, and production using the production function equation. Measured as a residual, the procyclical movements in the labor wedge unexplained by actual tax rates is a measure of failure in models of the labor market at business cycle frequencies.

Because the labor wedge is a residual measure of unexplained aspects of the labor market, a reduction in the labor wedge translates to a better explanation of labor market fluctuations. After documenting the strong procyclicality of the labor wedge, Shimer (2010) offers a partial solution in the form of endogenously rigid wages, (as in Hall (2005)), combined with search-and-matching models in the style of Mortensen and Pissarides (1994) that reduces the cyclical volatility of the labor wedge to a quarter its initially, naively-calibrated size. By deviating from the common Nash equilibrium wage-setting process to one with a rigid wage, the response of labor hours to productivity fluctuations increases by an order of magnitude (Shimer, 2004; Hall, 2005). These more cyclical hours better match the data. Thus Hall’s endogenously rigid

rigid wage bargaining model substantially reduces the market hours fluctuations which are left unexplained by technology shocks. More generally, there is a large literature of other papers using Calvo-style rigid wages to explain cyclical patterns in labor hours (for instance, Erceg et al. (2000)) or other backward looking wages (Christiano et al., 2016).

In principle, the self-employed are not subject to such frictions. It would be surprising to find the same cyclical in the labor wedge among the self-employed as is found among wage workers if wage rigidity was the driving force behind the cyclical in the labor wedge. This paper documents exactly this phenomenon.

Unfortunately, while many self-employed are residual claimants, there is clearly an ambiguity in self-employment: some self-employed workers may be self-employed in name only, and are largely or even completely subordinate to a client firm (Parker, 2004). Fortunately Lowe and Schellenberg (2001) attempt to quantify this with survey data, finding that only 15% of single-worker self-employed firms and 8% of self-employed persons with employees in Canada have potentially high overlap with paid employment. The self-employed also display markedly lower wage rigidity, suggesting that these “dependent” self-employed are not significant enough to threaten inference. There is ample evidence that monetary wages within a job are rigid (Barattieri et al., 2014; Bihan et al., 2012). The central point of this paper is that the self-employed are less subject to such rigidities. Since the self-employed are typically the residual claimant on firm profits, their full wages are more flexible than the fixed contracts of wage-workers. While the self-employed are still subject to goods-market firm matching frictions (such as the price frictions of Calvo (1983)), they are not subject, or less subject, to employer-employee matching frictions.

There has been a large literature extending representative agent models to capture secular aspects of the labor wedge (Cociuba and Ueberfeldt, 2015). Similarly, the countercyclical labor wedge has been subject to various challenges. For instance, Gourio and Rudanko (2014) challenges calculations of the labor wedge, while others write down different flavors of representative agents to reduce the wedge (Karabarbounis, 2010; Gourio and Rudanko, 2014; Atesagaoglu and Elgin, 2015). Supporting this paper’s exercise, Karabarbounis (2014) notes that explanations for the labor wedge are unlikely to be generated by the firm side of the labor market, while Cheremukhin and Restrepo-Echavarria (2014) point to changes matching efficiency. This paper’s central contribution is to highlight that the behavior of a specific group (the self-employed)

gives important information about the degree to which labor market frictions play a role in generating the labor wedge.

This basic idea of comparing self-employed and wage and salary workers is also present in Carrington et al. (1996), who document the similar cyclical variation of both wages and hours worked per worker. I differ in four important ways. First, I examine total hours worked, which are more likely to differ if wage and salary workers act on the extensive margin while self-employed act on the intensive (which they do). Second, I control for the effect of wages on work hours, an idea at the heart of modern macroeconomic studies of labor supply, and supplement this analysis with consumption data. Finally, I document that the self-employed do in fact have less sticky wages.

Another contemporaneous paper uses the labor market behavior of the self-employed to think about cyclical variation in hours. Bils et al. (2015) examines intermediate inputs and the intensive margin of self-employed labor market behavior to reach the conclusion that cyclical variation in price markups are as important as wage markups. While it shares the basic idea that the self-employed labor wedge is just as cyclical, this paper contributes in four important ways. First, I incorporate total hours, rather than only the intensive margin for self-employed. If the self-employed are better able to use the extensive margin to adjust labor hours than wage/salary workers (for instance, due to a lack of search frictions for self-employed) this will give a better measure of their labor wedge. Second, I use published aggregate statistics on self-employed and wage/salary hours at a quarterly frequency rather than CPS microdata at a yearly frequency: this higher frequency potentially allows me to more accurately examine movements at business-cycle frequencies. Third, I use self-employed wages in the calculation of the labor wedge: because these wages mitigate the self-employed labor wedge, my findings are strengthened by the assumption.¹ Finally, and perhaps most importantly, this paper examines the degree to which the significantly different industrial-occupational differences of jobs held by worker class plays a role in generating the main result, finding that they do not substantially contribute to the cyclical variation in hours.

If the self-employed are subject to lower wage or search frictions but display the same measured wedge between MRS and MPL, it is unlikely that wage and search frictions are the proper

¹An identifying assumption in Bils et al. (2015) is that there are no wage rigidities or distortions causing there to be a wedge between MRS and MPL for the self-employed, so that the MRS can be used to measure the shadow wage. Because I use self-employed income, this is not an identifying assumption in my model, but a motivating fact.

explanation for the labor wedge, in spite of its relative success for wage and salary workers. While I confirm the intuition that the self-employed have less sticky wages using the National Longitudinal Survey of Youth (1979), I also discuss the existing literature on the income and hours elasticity of the self-employed which corroborates this common intuition. I use BLS published series on nonagricultural employment of self-employed and wage/salary workers, and average nonagricultural working hours data and hours data of wage and salary workers to construct two series of total hours, one for the self-employed and one for wage and salary workers. I combine these series with measurements from the Consumer Expenditure Survey (CEX) to construct the labor wedges by worker class, and analyze cyclical aspects of the two labor wedges, I find that while both have strong cyclical correlations with GDP (-0.85 for the wage/salary labor wedge, -0.49 for the self-employed labor wedge), the standard deviation of the self-employed labor wedge is 45% higher than that of the wage/salary labor wedge (3.5% vs. 2.4%).

There are many other potential sources for the labor wedge that this method does not exclude. Constructively, any friction that hits the self-employed in the same way as wage and salary workers is a more likely candidate, given this papers results. This elevates models that study common shocks, such as implicit marginal taxes (Mulligan, 2012b), final good price rigidities (Christiano et al., 2016; Bils et al., 2015), or household production (Atesagaoglu and Elgin, 2015), while casting doubt on matching frictions or labor contracting rigidities.

The rest of this paper is organized as follows: section 2 describes the model I use to measure the labor wedges, section 3 discusses wage stickiness of self-employed and wage/salary workers, describes the parameters and main data sources I use to construct the labor wedge, depicts these series and corroborates my main findings using secondary data sources. Section 4 contains my main results, while section 5 concludes.

2 Model

The labor wedge is the difference between the log marginal product of labor and the log marginal rate of substitution between consumption and leisure: a residual that explains, given consumption and income, what taxes must have been in order for labor markets to be in equilibrium. This type of “wedge accounting” is not new: while Parkin (1988) or Hall (1997) treat the parameter governing the steady-state labor-leisure tradeoff as stochastic, Mulligan (2002), Cole and Ohanian (2002) and Chari et al. (2007) treat it as a time-varying tax.

In most specifications of the labor wedge, time series on consumption, income, and labor are necessary. Consumption informs us about the marginal utility of consumption, (a necessary component of the marginal rate of substitution between consumption and leisure), and income given labor informs us about wages (a function of the marginal product of labor). Labor hours worked, along with given information on wages, gives the marginal utility of leisure. Consumption given income contains information about what a household believes about future wages: when the ratio of consumption and income is high, households believe wages are currently low compared to the future. Consequently, when that ratio is high, it gives us a non-tax reason to justify low labor hours. Because empirically the national consumption-income ratio is countercyclical, it will play a role in diminishing the labor wedge.

The intuition behind the labor wedge is a simple accounting for intertemporal substitution of labor. Empirically, we observe procyclical labor hours, and look to explain these with procyclical wages. Because permanent changes in wages have offsetting income effects on labor supply while temporary changes do not, when explaining labor hours with wages, we need to make a distinction between temporary and permanent changes. Because households intertemporally smooth consumption, the consumption-income ratio gives information on the permanence of wage shocks. When wages rise permanently, so will consumption, and consumption-income will not change. For instance, assuming a Frisch elasticity of one, if wages fall by 1% and consumption falls by 1%, it suggests a permanent change in wages, and any change in labor hours are evidence of change in taxes, given offsetting income and substitution effects. On the other hand, if wages fall by 1% and consumption does not fall, it suggests a temporary change in wages, if labor hours then drop by 1% with a Frisch elasticity of one, the implicit tax must not have changed. If they drop by less, then in spite of labor hours falling, we expect the tax wedge has fallen, while if they drop by more, then the tax wedge has risen. A countercyclical labor wedge is an admission that this intertemporal substitution channel is not enough to explain changes in labor hours over the business cycle.

I start with the simple observation that the setup most likely to minimize the self-employment labor wedge allows households no substitutability between wage/salary work and self-employment. The idea behind this is simple: if self-employment is less subject to the frictions that create the labor wedge, then households in a model that allows substitutability will move into self-employment. Insofar as they don't (labor typically falls more for the self-employed), this is

evidence of a higher countercyclical self-employment labor wedge. In other words, in order to explain why people don't move into self-employment, the model would have to have a higher self-employment labor wedge. By removing this requirement, (completely separating the labor choices of both groups and allowing no labor market risk-sharing between employment groups) I minimize the countercyclicality of the self-employment labor wedge, but still find it to be comparable in size to the wage and salary labor wedge. In addition, if resources were pooled (in a representative-family/complete markets framework) then consumption would be less correlated with class-specific wages, because the procyclical savings rate mitigates the labor wedge. Because I will be arguing the magnitude of the self-employment labor wedge calls into question the ability of sticky wages to explain labor market movements, I consider complete separation to be a conservative model. To re-iterate, the assumption that both self-employment and wage and salary work is not possible for the same household works against countercyclicality of the self-employed labor wedge. Removing this assumption only strengthens my main conclusions.

The formulas for the two labor wedges therefore follow Shimer (2009) and Shimer (2010). With worker of class $i \in \{W, SE\}$ (where W denotes wage/salary workers and SE denotes the self-employed) having period utility over consumption $c_{i,t}$ and labor $L_{i,t}$, disutility of labor ψ_i and constant Frisch elasticity of labor supply ϵ_i :

$$U_i(c_{i,t}, L_{i,t}) = \log(c_{i,t}) - \psi_i \frac{\epsilon_i}{1 + \epsilon_i} L_{i,t}^{\frac{1 + \epsilon_i}{\epsilon_i}}$$

And period budget constraint with investment $i_{i,t}$, depreciation rate of capital δ , interest rate r_t , wage rate $w_{i,t}$, tax rate τ and productive capital determined the previous period $k_{i,t-1}$:

$$c_{i,t} + k_{i,t} = (1 - \delta + r_t)k_{i,t-1} + (1 - \tau_t)w_{i,t}L_{i,t}$$

Taking the ratio of the first order conditions with respect to consumption and labor gives that the marginal rate of substitution between labor and leisure is equal to the after-tax wage. After some rearranging, I write implied labor tax rate (the labor wedge) $\tau_{i,t}$ in equation (1).

$$\tau_{i,t} = 1 - \psi_i \frac{c_{i,t}}{w_{i,t}} L_{i,t}^{\frac{1}{\epsilon_i}} \quad (1)$$

With constant returns to scale Cobb-Douglas production technology and competitive firms,

profit maximization gives the wage as a constant fraction of output $Y_{i,t}$ equal to the output elasticity $1 - \alpha_i$:

$$w_{i,t} = (1 - \alpha_i) \frac{Y_{i,t}}{L_{i,t}}$$

Plugging this into equation (1), we get the simple expression for the market-clearing labor wedge:

$$\tau_{i,t} = 1 - \frac{\psi_i}{1 - \alpha_i} \frac{c_{i,t}}{Y_{i,t}} L_{i,t}^{\frac{1+\epsilon_i}{\epsilon_i}} \quad (2)$$

We therefore need data on $c_{i,t}$, $Y_{i,t}$, $L_{i,t}$, and need to calibrate α_i , ψ_i , and ϵ_i . In a simple representative family, complete markets framework that allows income pooling and has separable preferences over consumption, the equation would simply substitute $c_{i,t}$ in each equation for c_i , the joint household consumption, as in equation (3), an alternative model of the labor wedge, in which $c_{i,t}$ of equation 2 is replaced by a common c_t :

$$\tau_{i,t} = 1 - \frac{\psi_i}{1 - \alpha_i} \frac{c_t}{Y_{i,t}} L_{i,t}^{\frac{1+\epsilon_i}{\epsilon_i}} \quad (3)$$

This reflects the fact that now the marginal utility of consumption controlled by $c_t = c_{SE,t} = c_{W,t}$ is the same across the family, even as wages, controlled by $Y_{i,t}$, differ. I report results for both of the extreme cases: my main calibration in which I allow no substitution (equation 2) and an alternative calibration that allows family labor income pooling (equation 3) in the Appendix, with the same qualitative result of the self-employed labor wedge's cyclicalty. By using income data for the self-employed in equations 2 and 3, I break with Bils et al. (2015), which assumes no wedge between MRS and MPL for the self-employed.

3 Data

To construct the labor wedge, I need data on six data series: $c_{i,t}$, $Y_{i,t}$, $L_{i,t}$ for both the self-employed and wage/salary workers, and need to calibrate labor's shares α_i , the disutilities of labor ψ_i , and the Frisch elasticities of labor ϵ_i . Before I calibrate my labor wedge and examine the component data series, I first address the identifying assumption in my model: do the self-employed experience lower frictions than do wage and salary workers? After establishing that the self-employed display less rigid wages and discussing additional evidence, I turn to my parameter choices for α_i , ψ_i , and ϵ_i , then to the data.

3.1 Wage Stickiness by Worker Class

In this section I document wage stickiness by worker class. Most studies that use publicly-available data on rigid wages use the Current Population Survey. Because of its short panel nature, with four consecutive months of interviews (with the last interview asking detailed questions about wages), and a follow up a year later, the CPS allows researchers insight into both labor market flows and wage stickiness (see, for instance, Daly and Hobijn (2016) and Daly and Hobijn (2014)). Unfortunately, the outgoing rotation group (ORG) with data on wages excludes the self-employed from wage and hours questions.

With detailed information about the employment nature of each job held and the hourly wage earned at each job, the NLSY79 provides a wealth of information about rigid wages. The NLSY79 has repeatedly surveyed approximately thirteen thousand individuals who were born between 1957 and 1965, annually from 1979 to 1994, and biennially thereafter. Pooling changes by year, the data provide 78,000 usable observations, of which nearly 6,000 are self-employed and 73,000 are wage or salary workers. An alternative dataset sometimes used for sticky wages is the Survey of Income and Program Participation (see for instance Barattieri et al. (2014)): unfortunately, due to seam bias and income imputation, the 2004 panel yields less than 170 usable self-employed observations, and I consequently focus on the NLSY.

There are two potential problems of measuring the self-employed as compared to wage and salary workers. First, if wage workers report their actual wage each year, while the self-employed worker wages are calculated with error using their income and hours, our calculations will be predisposed to find larger fluctuations in wages for the self-employed due to misreporting. Second, some self-employed pay themselves a salary and also receive income from their business. If they pay themselves a fixed salary and allow all residual income to be reported as firm profits, this will incorrectly increase the self-employed wage rigidity. This is a particular problem if households use firm capital stock to smooth temporary fluctuations in labor.

To deal with these problems and generate data on the degree of wage stickiness, I examine the absolute change in wages as measured by reported labor income divided by reported hours in a given year. This avoids the problem of more accurate wage reporting by wage workers (using the reported wage for wage workers does indeed increase wage stickiness by a factor of four), and would require differentially bad recollection about income per hour worked to generate the patterns I observe. Second, and as one might expect, when I include net profits in the calculation

of the self-employed’s wages, then wages get even less sticky. I exclude net profits as part of wage to be conservative in how much less sticky self-employed wages are.

Figure 1 shows the histogram of wage changes for self-employed and wage/salary workers in the pooled 1980-1994 NLSY (the years for which annual observations are available). The same figure for the available biennial observations (1995-2012) is provided as Appendix Figure 1 and shows the same pattern. Within-job wage for wage and salary workers is noticeably more sticky than for self-employed workers. Summary statistics for the distribution of wage changes between two thresholds are available in Figure 1. The difference in wage changes is not simply a function of small measurement errors for the self-employed: the pattern is present throughout the distributions.

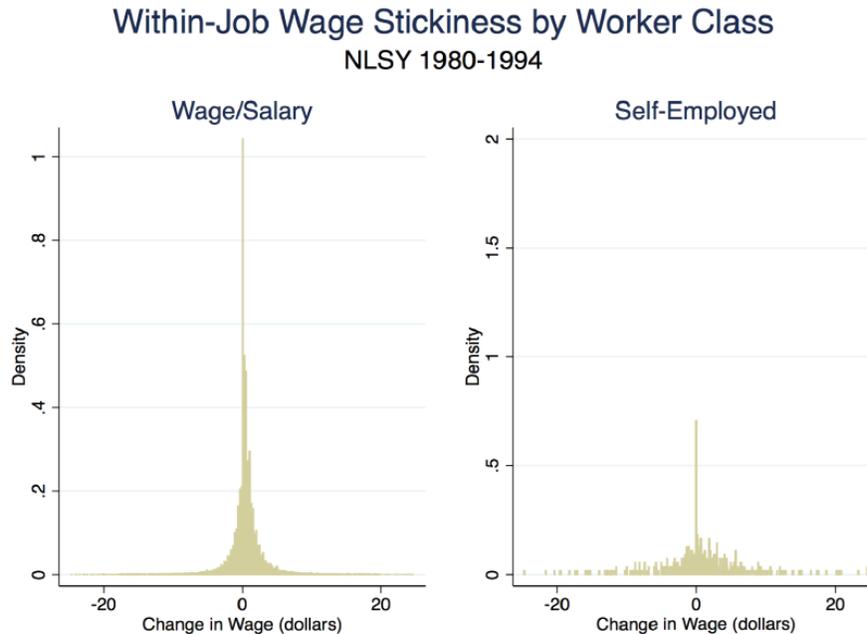


Figure 1: This figure displays, using the NLSY from 1980-1994, the distribution of wage changes (in dollars) for wage and salary workers and the self-employed. Comparable figures can be produced using NLSY’s biennial observations from 1996-2012. Wages are calculated as total labor income divided by hours worked.

I conclude this section by noting that if we take the idea that the self-employed are truly not subject to any contracting frictions, then post-tax wage should be equal to marginal rate of substitution, and given a measure of the marginal rate of substitution no wage measurements are necessary, as in [Bils et al. \(2015\)](#). Given the weak evidence in Figure 1 that the self-employed

Table 1: Distribution of Wage Changes

Statistic	Wage and Salary	Self-Employed
Mean	0.51	.91
Standard deviation	3.14	10.12
Percent at zero	9.6%	6.6%
Percent between -\$0.05 and \$0.05 change	11.1%	7.1%
Percent between -\$0.10 and \$0.10 change	13.5%	7.9%

Table 1: This table displays the differences in annual wage changes from 1980-1993 in the NLSY79 for wage/salary workers and the self-employed. All statistics are calculated for wage changes below \$50.

may be subject to some sticky wages, and that the wage behavior of the self-employed appears to weakly mitigate the labor wedge (see below), measuring wages directly may be conservative, reducing the self-employed labor wedge.

3.2 The Frisch elasticity, disutility of labor, and labor income share

While a large literature has been devoted to estimating the Frisch elasticity of labor supply, it generally commingles the deep preference parameter ϵ_i with optimization frictions (Chetty, 2012). Put another way, the frequently-estimated parameter ϵ_i is a reduced form parameter describing how people actually behaved in response to some wage change subject to frictions, but does not necessarily give the optimal frictionless response, and likely underestimates ϵ_i . While less concerning with an aggregated agent, it is possible estimated Frisch elasticities may represent a deep preference parameter for the self-employed but a mongrel coefficient for wage and salary workers.

Complicating the problem of drawing taxable income elasticities from the literature for each group, increased responsiveness between groups may come from a different Frisch elasticity, a differential ability to misreport income due to the lack of third-party reporting for the self-employed, or to lower frictions. For instance, Saez (2010) finds strong bunching at the points where the Earned Income Tax Credit changes its implicit tax, but only for the self-employed. While consistent with his model of income misreporting (corroborated by the random audit study of Kleven et al. (2010), which found misreporting rates of 37% for the self-employed and 0.3% of returns subject to third-party reporting) it is also consistent with lower frictions. Similarly, le Maire and Schjerning (2013) examine the same Danish data and find that self-employed bunch under tax kinks, though they are unable to distinguish between effort and reporting

issues. Fortunately, some studies examine hours reported on surveys which are not subject to the misreporting or income-shifting motivations that administrative tax data is. Showalter and Thurston (1997) find that self-employed physician hours are much more responsive to marginal tax rates than those employed by hospitals, finding an intensive labor elasticity with respect to marginal tax rate of -0.30 for all self-employed, -0.57 for solo/sole proprietor physicians and an insignificant point estimate of -0.09 for wage/salary physicians. It is also worth noting that wage and salary workers may have a much lower intensive elasticity but a correspondingly larger extensive elasticity of labor supply compared to the self-employed if the primary way wage and salary workers control their hours is by choosing when to start and stop jobs while the self-employed have finer (intensive) control. Because there are no estimates of the Frisch elasticity of total labor hours for the self-employed, especially in light of Chetty (2012), I assume that the two are equal, setting the aggregate Frisch elasticity of both groups to 0.75, a value for representative agent models consistent with micro evidence in light of optimization frictions (Chetty et al., 2011).

As in Shimer (2010), I choose the value for $\frac{\psi_i}{1-\alpha_i}$ so the average measured tax wedge for both groups equal to 0.4, as in Prescott (2004)². This allows me to sidestep the potentially difficult question of how to measure α_i for the self-employed, as normally the share of labor is calibrated by assuming unobserved income shares to labor and capital for sole proprietors is the same as shares for wage and salary workers (Cooley and Prescott, 1993).

3.3 Consumption, income, and labor hours

Since 1980, the CEX has continuously surveyed focusing on household expenditures. Households are interviewed quarterly for five quarters. Questions about income and employment, including employment class, are available in the second and fifth quarter interviews, while expenditure information has two components. In the first quarterly interview component, households are asked about major durable goods purchases. In the second daily diary survey (lasting two consecutive weeks) households are asked about their daily purchases.

Not all consumption is the same: major durable goods may be classified as investment, rather than consumption (Kuznets, 1942). In my primary specifications, I therefore follow Dynan

²Prescott (2004)'s estimates of 0.4 are specifically for tax rates measured from 1970-1974 and 1993-1996, his periods of interest. Nevertheless, changes in the level of τ_i would only matter for my core results insofar as $\frac{\psi_i}{1-\alpha_i}$ have changed differentially by income class over time.

et al. (2004) in constructing consumption from the CEX by taking total household expenditures excluding home capital improvements (including major appliances), life insurance payments, and spending on new and used vehicles. For those who own their own home, I add the difference between their imputed rent and their mortgage payments³. It's theoretically unclear whether or not educational spending should be included as investment, in the vein of Becker (1962), or consumption, a point central to Schultz (1961). Because its inclusion does not matter for my results, I included it as consumption.

To construct my two time series, I use the quarterly CEX public use microdata. While this quarterly data series is noisy, annual tables are produced by the BLS summarizing purchases by employment class. I temper the levels of my quarterly data series so they match the annual savings rates displayed in the CEX data. These annual series are depicted in Figure 2. They show a different secular pattern to NIPA's C/Y data, potentially because the CEX's coverage is different. The different coverage of the two series is not a concern if the CEX is simply more or less countercyclical than the data, because it will increase the cyclicity of both labor wedges, but will not affect their difference (or lack thereof), which is the focus of this analysis. A greater concern is if misstatement of cyclicity of the CEX differs by income class. The stronger countercyclical nature of the CEX's series on consumption as a fraction of income for the self-employed compared to wage and salary workers is evident from Figure 2: this will act to reduce the volatility of the self-employed labor wedge. Because my results find labor wedge cyclicalities of equal magnitude, I consider my results to be conservative, in light of potential CEX mismeasurement.

Cyclical fluctuations in labor hours are at the heart of the labor wedge. My primary source of data on labor hours comes from published quarterly series on employment and hours generated from the BLS using the CPS. The BLS publishes series on nonagricultural employment by worker class: for wage and salary workers and for unincorporated self-employed workers⁴. Unfortunately, the BLS does not publish a separate quarterly series for incorporated self-employed workers, who make up about one third of all self-employed workers. The BLS also publishes series on average hours worked by persons at work for nonagricultural industries as a whole (including both self-employed and wage and salary workers), as well as a separate series on

³Mortgage payments are broadly defined, and include fees and taxes.

⁴BLS series LNS12032187Q and LNS12032192Q

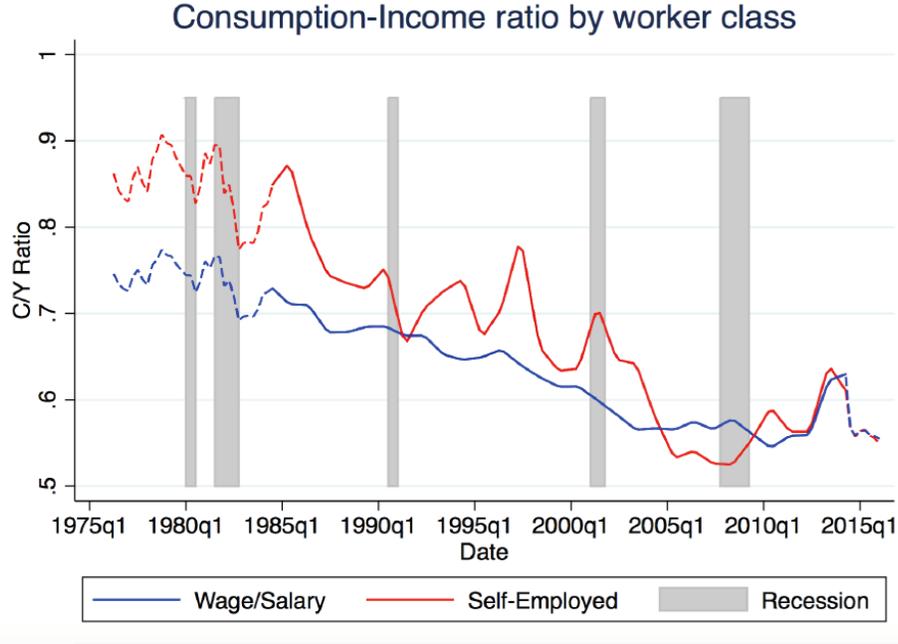


Figure 2: This figure depicts the ratio of consumption and income by worker class reported in the CEX annual tables from 1984-2014. Dashed values are where the CEX is unavailable: fitted values from 1976-1983 and for 2015 (displayed with a dash) are generated by regressing the CEX's class-specific consumption-income ratio on NIPA's consumption-income ratio and plotting the fit for unobserved years.

(nonagricultural) hours worked by workers for only wage and salary workers⁵. I generate a measure of hours worked by self-employed workers by noting that the total of all hours at work in nonagricultural industries (the product of hours worked and employment) must be equal to the sum of the product of employment and hours worked by wage and salary workers and the product of employment and hours worked by the self-employed. Because I have each data series but the last, letting E_i , $i \in \{All, W, SE\}$ denote total employment, wage employment, and (unincorporated) self-employment respectively, and H_i be hours conditional on work for the same series, I can solve for the last using equation 4.

$$H_{SE} = \frac{E_{All} \cdot H_{All} - E_W \cdot H_W}{E_{SE}} \quad (4)$$

The benefits of using these published BLS data series along with equation 4 are that they allow for both better measurement and for more frequent observation than other public use data series

⁵BLS series LNS12033120 and LNS12033251

do. However, there are two potential issues with using these BLS data series. First, the data is only on nonagricultural workers. Because agricultural workers are only 1.2% of all employment, it is unlikely they are driving the cyclical movements in the labor wedge. Indeed, using the CPS to look at all industries at an annual frequency, including the agricultural workers decreases the cyclical volatility of log labor hours for the self-employed from 1.98% to 1.81%, while it decreases the cyclical volatility of wage and salary workers from 1.23% to 1.25%, suggesting that their exclusion is unlikely to be important.

A second issue is that the data suffers from rounding error. Because hours of those employed are only published with precision of tenths of an hour, a 0.026% change in hours worked, from 38.34 to 38.35 hours per worker, may be calculated as a 0.26% change in hours worked, from 38.3 to 38.4. Indeed, the effects of these changes are evident from indexed total hours per working-age person series calculated using equation 4 and depicted in Figure 4. There is high-frequency noise created by the BLS rounding error: fortunately, the annual CPS hours series, which does not suffer this problem, shows similar (but more muted) cyclical behavior. To deal with the primarily one-quarter high-frequency noise, I use the band-pass Baxter-King (BK) filter (Baxter and King, 1999), which filters out both high and low frequencies, rather than the more common high-pass Hodrick-Prescott (HP) filter (Hodrick and Prescott, 1997), which focuses on only low frequencies. Using the HP filter leaves my basic cyclical findings unchanged, but includes more high-frequency quarterly noise.

I depict the (Baxter-King) cyclical component of log total hours per working-age person by worker class in Figure 5. makes clear the core hours data that produces the procyclical labor wedges: the percentage change in labor hours during times of low output are as significant (or more) for the self-employed as they are for wage and salary workers. While at first glance it's surprising that labor hours fluctuate so much during non-recession years, the same potentially troubling patterns (drops during 1986, 1995) are comparable to those used in Shimer (2010), replicated in Figure A.1 in Appendix A.⁶

⁶It may be noted that my use of total hours opens up the possibility that some other friction that acts only on the extensive margin and impacts the self-employed more than wage/salary workers may coincidentally cause their labor wedge to look like the wedge of wage/salary workers. The intensive margin findings of Carrington et al. (1996) and Bils et al. (2015) suggest this is unlikely.

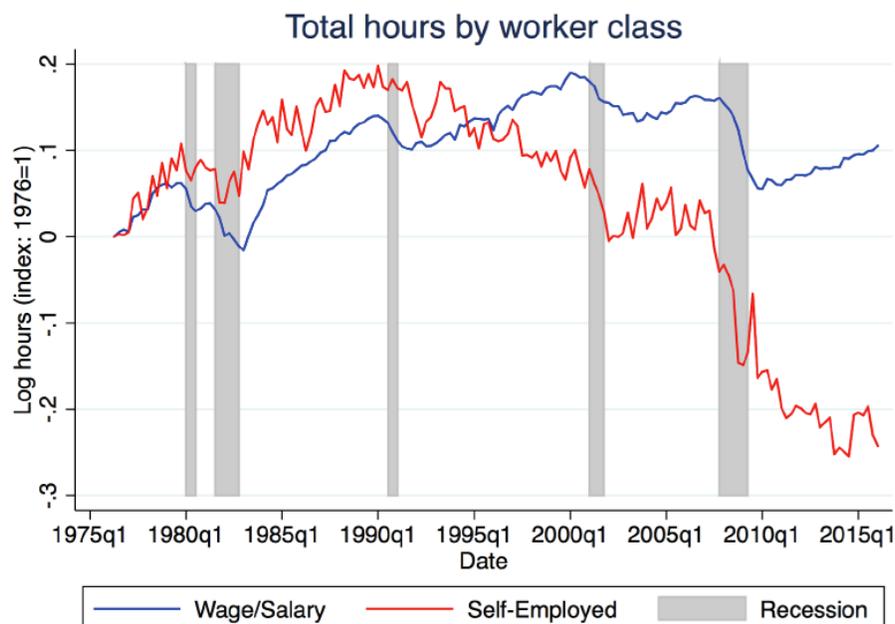


Figure 3: This figure depicts indexed log hours per working age person by worker class, with self-employed hours per worker generated by equation 4 and combined with employment to calculate total hours. As discussed in the text, the noise in self-employed hours is partly generated by rounding error and motivates my use of the band-pass BK filter rather than the high-pass HP filter.

4 Results

Using the data and parameters discussed in Section 3, I calculate the labor wedges with equation 2. The level of my labor wedge is depicted in Figure A.2, and is comparable to Shimer (2010), replicated in Figure A.3 in Appendix A. While the two series have different time coverages reflecting different availability of sources for total hours, there are clear (and different) secular trends in both the wage and salary labor wedge and the self-employed labor wedge. Reflecting the increase in worked hours per capita in the U.S. over the 1980's and 1990's, the labor wedge fell (except during the recession in the early 1990's). It rose and never fully recovered from the recession in the early 2000's, with a similar pattern after the recession of 2007-2009. Because it is noisier, some subtle patterns are harder to pick out of the self-employed labor wedge, except for its dramatic rise from the early 1990's through 2012. This clear secular trend in the tax wedge of the self-employed reflects the decades-long decline in self-employment in the United States (Hipple and Hammond, 2016). Secular changes in self-employment of this nature are likely to

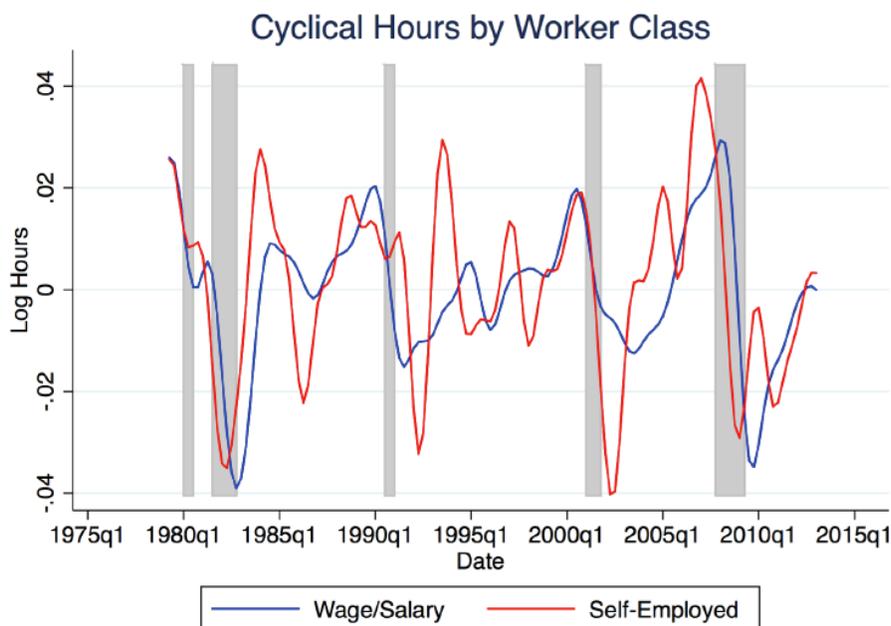


Figure 4: This figure depicts hours per worker by worker class, with self-employed hours per worker generated by equation 4. As discussed in the text, the noise in self-employed hours is partly generated by rounding error and motivates my use of the band-pass BK filter rather than the high-pass HP filter.

follow secular changes in management and communication technology, industrial composition, and tax rates, and are outside the scope of this paper.⁷

While the secular changes are interesting, they are not the focus of this paper. Similarly, Figure 2 displays a great deal of high-frequency noise. To exclude both from my comparisons, I BK-filter each tax wedge and display them in Figure 6. The wedges display similar increases during business cycles, giving the core result of this paper. From it, we learn that the two cyclical components are largely the same. Of the five post-1980 recessions observed in my data, the self-employed labor wedge climbs more than the wage and salary labor wedge. The exception is the 2007-2009 recession, which saw the self-employed labor wedge jump earlier, and fall as the wage at the tail end of the wage/salary labor wedge's rise. While this is puzzling, it reflects the hours data of Figure 5. More specifically, self-employed labor hours saw a sharp spike in the first two quarters 2010 before a sharp fall back to levels comparable to those of wage/salary workers.

⁷For an older analysis of the drivers of secular changes in self-employment, see Blau (1987).

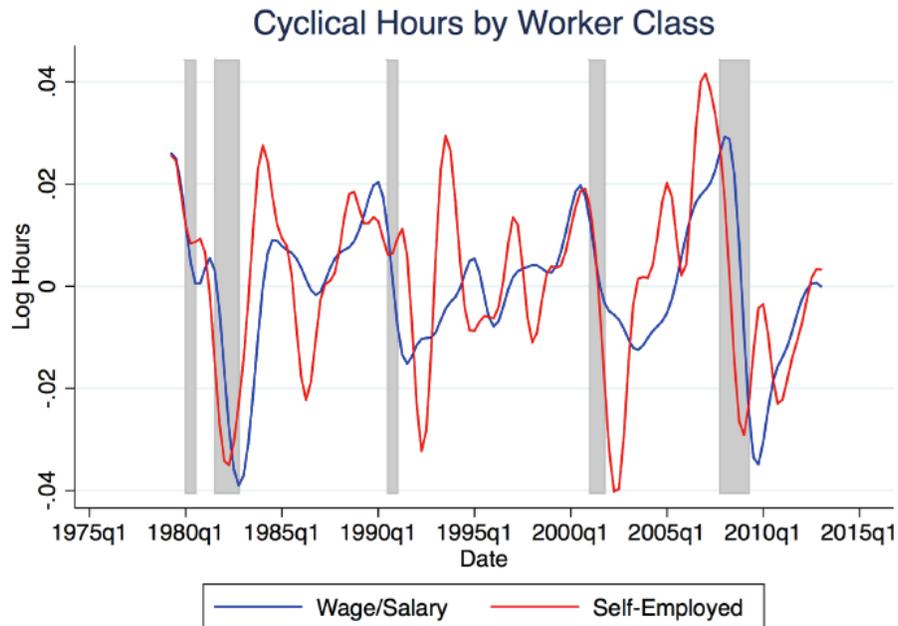


Figure 5: This figure depicts the core data that generates the labor wedge: the cyclical hours of wage/salary workers and the self employed. Importantly, the magnitude of fluctuations are highly comparable even though the self-employed are not, or less, subject to wage rigidity.

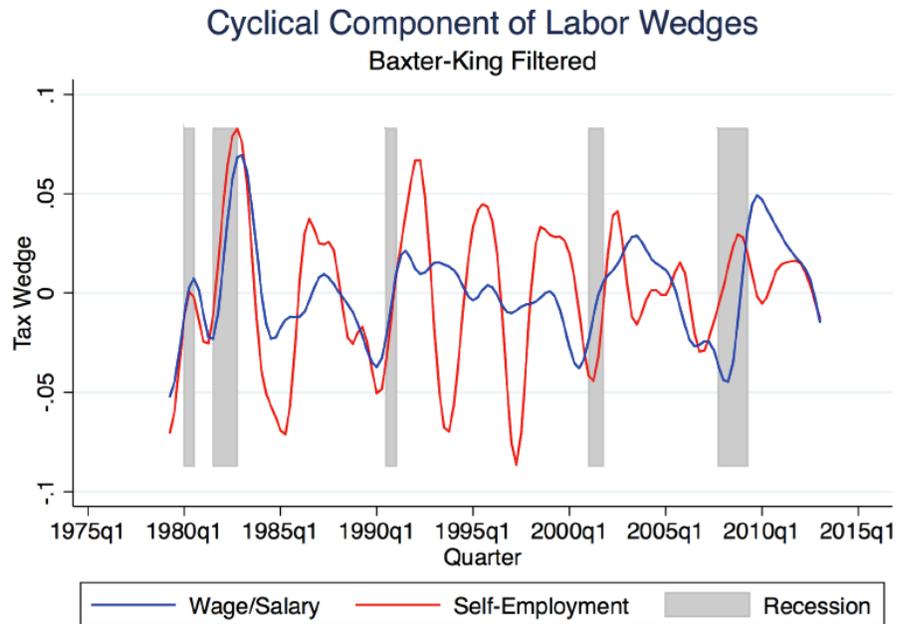


Figure 6: This figure depicts the BK-filtered version of the self-employed and wage and salary labor wedges whose levels are shown in Appendix Figure A.2.

There are a few periods when the two do not broadly mimic one another. In the early and mid-1990's after the recession, the self employed labor wedge falls and then rises while the wage and salary labor wedge does not. Similarly, there is a prolonged rise in the wage and salary labor wedge in the post-recession early 2000's where the self-employed labor wedge saw a sharp fall. To better understand the time series patterns of the two labor wedges, including their relationship to GDP, I display the standard deviations of the BK-filtered labor wedges along with BK-filtered log GDP in Table 2. There is a strong (-0.85) correlation between the wage labor wedge and GDP, and a strong but weaker relation with the self-employed labor wedge (-0.49). The two labor wedges themselves are strongly correlated (0.51). The self-employed labor wedge has a dramatically higher BK-filtered standard deviation. While some of this doubtlessly reflects noise in measurement, it presents a challenge for any wage rigidity explanation.

Table 2: Summary Statistics

Variable	Standard Deviation	
	Wage & Salary	Self-Employed
GDP	1.34%	
Consumption-Income Ratio	0.66%	2.5%
Tax Wedge	2.4%	3.5%
Pair	Correlation	
$Corr(GDP_t, \tau_{W,t})$	-0.85	
$Corr(GDP_t, \tau_{SE,t})$	-0.49	
$Corr(\tau_{W,t}, \tau_{SE,t})$	0.51	

Table 2: This table displays summary statistics on BK-filtered log GDP and the tax wedges for wage/salary and self-employed workers.

As is evident from Figure 2, CEX data are very noisy, and are generate some of the mis-measurement in the labor wedge. The representative-household model of equation 3 makes it possible to use National Income and Product Accounts (NIPA) data for measurement of the savings rate. In equation 3, we no longer need consumption by class, though we need incomes by class. Consequently, I use NIPA's consumption and income tables. While the NIPA does not break up income by worker class, it separates both proprietor's income and compensation of employees: if these are proportionally related to self-employed and wage/salary employee income, then they allow for less noisy measurement of the consumption-income ratio. The result of using this different consumption and income measure are displayed in Appendix Figure A.4. Using NIPA data reduces the self-employed labor wedge's fluctuations in the recessions in the early 1980's and early 1990's, while amplifying them in the most recent two recessions in the early

and late 2000's. I prefer the use of the CEX, as the assumptions underlying the representative household are clearly rejected in Figure 2.

4.1 Occupation-Industry Composition

A frequent concern when thinking about the self-employed is that the composition of their labor differs from the composition of wage and salary workers. This is true by gender, race, occupation, and industry (Georgellis and Wall, 2005). Because the root of the similarities in the labor wedge are driven by similarities in the cyclicality and magnitude of hours, the question is whether or not differential occupational and industrial composition lends itself to different exposures to cyclical volatility. For instance, if the self-employed were primarily in industries that were highly cyclical, such as the capital goods or durable goods manufacturing, while wage/salary workers were employed primarily in the acyclical industries, such as services, suggest the self-employed would have had lower hours volatility if they had had similar occupational-industrial composition as wage/salary workers.

To examine this hypothesis, I use the American Community Survey from 2000-2014⁸ Fortunately, for the ACS, unemployed workers are classified according to their last occupation or industry, so hours measures do not capture only intensive changes in hours worked, but total hours. On self-employed workers, I run a regression of hours worked on industry-occupation-year fixed effects, extracting the occupation-industry fixed effect for each year, as in equation 5, where $X_{i,t}^{SE}$ denotes a matrix of occupation-industry dummies by person for the self-employed population. I then apply those occupation-industry fixed effects $\hat{\beta}_t^{SE}$ to the wage and salary population⁹. In other words, I take the occupation-industry trends for the self-employed workers and apply them to wage and salary workers. Figure 7 compares the predicted hours change with those of the self-employed. The fact that the two are very similar suggests that the actual pattern of hours displayed by the self-employed closely follows the pattern of hours predicted if they had been wage and salary workers. Insofar as the prediction line is flatter than the actual data, it suggests that industrial-occupation effects amplified hours differences, while higher volatility would suggest composition effects muted differences. The BK-filtered version is depicted in Figure 8 and shows that industrial-occupational differences play a minimal role in cyclical hours

⁸Including the 2000 Census.

⁹Some occupations have no self-employed workers: for these, I use results from the same regression on wage workers. Excluding these workers completely does not impact my results.

differences.

$$\text{Hours}_{i,t} = \alpha + \sum_{t=2000}^{2014} \beta_t^{SE} X_{i,t}^{SE} \quad (5)$$

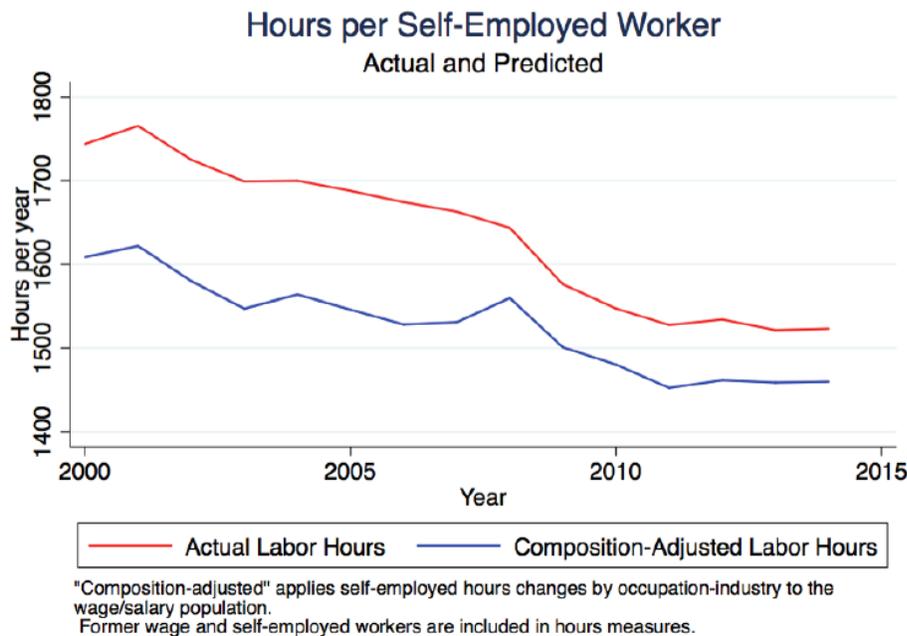


Figure 7: This figure depicts actual hours per self-employed worker and hours predicted if self-employed workers had the same industrial composition as wage and salary workers. Insofar as the two are similar we may conclude that occupation-industrial composition does not play a large role in explaining labor hours. In the language of equation 5, it graphs both $\hat{\beta}_t^{SE} X_{i,t}^{SE}$ as well as $\hat{\beta}_t^{SE} X_{i,t}^W$, where $X_{i,t}^W$ is a matrix with the occupational composition of wage and salary workers.

5 Conclusion

This paper argues that the movements of self-employed worker's hours inform us about the degree to which frictions (and what frictions) drive cyclical movements in labor markets. These results cast doubt on the suggestion in Hall (2005) and Shimer (2010) that rigid wages can explain labor market fluctuations at business cycle frequencies. While data on the self-employed is relatively sparse, and efforts must be made to cobble together a picture from several sources, the core point of this paper remains clear: the self-employed are subject to lower wage rigidity but their hours are highly variable and their labor-wedge is as cyclical as that of wage workers.

Some papers have tried to account for labor market behavior during recessions without rigid

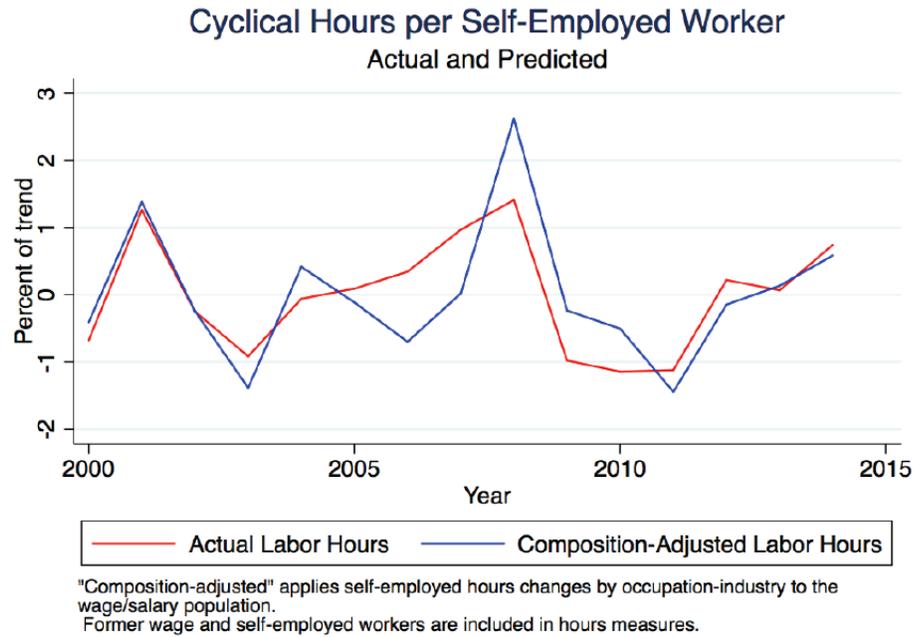


Figure 8: This figure displays the BK-filtered hours of self-employed workers and BK-filtered hours of self-employed workers if they displayed the same occupation-industry pattern as wage and salary workers. The figure shows that in spite of the secular decline in η , the cyclical drop in hours per worker for the self-employed is very similar to that of wage and salary workers.

wages. For instance, Christiano et al. (2015) use non-wage price rigidities to explain (among other things) labor market behavior during the Great Recession. While others have used large changes in implicit marginal tax rates by group to explain labor market behavior (Mulligan, 2012a). One point of this paper is to suggest that these may be more fruitful avenues in explaining labor market behavior.

References

- Orhan Erem Atesagaoglu and Ceyhun Elgin. Cyclicalilty of labor wedge and informal sector. *Economics Letters*, 136(C):141–146, 2015.
- Alessandro Barattieri, Susanto Basu, and Peter Gottschalk. Some Evidence on the Importance of Sticky Wages. *American Economic Journal: Macroeconomics*, 6(1):70–101, January 2014.
- Marianne Baxter and Robert G. King. Measuring Business Cycles: Approximate Band-Pass Filters For Economic Time Series. *The Review of Economics and Statistics*, 81(4):575–593, November 1999.
- Gary S. Becker. Investment in Human Capital: A Theoretical Analysis. *Journal of Political Economy*, 70:9, 1962.
- Herv Le Bihan, Jirmi Montorns, and Thomas Heckel. Sticky Wages: Evidence from Quarterly Microeconomic Data. *American Economic Journal: Macroeconomics*, 4(3):1–32, July 2012.
- Mark Bils, Peter J. Klenow, and Benjamin A. Malin. Resurrecting the Role of the Product Market Wedge in Recessions. Staff Report 516, Federal Reserve Bank of Minneapolis, October 2015.
- David M Blau. A Time-Series Analysis of Self-employment in the United State. *Journal of Political Economy*, 95(3):445–67, June 1987.
- Guillermo A. Calvo. Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, 12(3):383–398, September 1983.
- William J Carrington, Kristin McCue, and Brooks Pierce. The Role of Employer-Employee Interactions in Labor Market Cycles: Evidence from the Self-Employed. *Journal of Labor Economics*, 14(4):571–602, October 1996.
- V. V. Chari, Patrick J. Kehoe, and Ellen R. McGrattan. Business Cycle Accounting. *Econometrica*, 75(3):781–836, 05 2007.
- Anton A. Cheremukhin and Paulina Restrepo-Echavarria. The labor wedge as a matching friction. *European Economic Review*, 68(C):71–92, 2014.

- Raj Chetty. Bounds on Elasticities With Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply. *Econometrica*, 80(3):969–1018, 05 2012.
- Raj Chetty, Adam Guren, Day Manoli, and Andrea Weber. Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins. *American Economic Review*, 101(3):471–75, May 2011.
- Lawrence J. Christiano, Martin S. Eichenbaum, and Mathias Trabandt. Understanding the Great Recession. *American Economic Journal: Macroeconomics*, 7(1):110–67, January 2015.
- Lawrence J. Christiano, Martin S. Eichenbaum, and Mathias Trabandt. Unemployment and Business Cycles. *Econometrica*, 84:1523–1569, 07 2016.
- Simona Cociuba and Alexander Ueberfeldt. Heterogeneity and long-run changes in aggregate hours and the labor wedge. *Journal of Economic Dynamics and Control*, 52(C):75–95, 2015.
- Harold L. Cole and Lee E. Ohanian. The Great U.K. Depression: A Puzzle and Possible Resolution. *Review of Economic Dynamics*, 5(1):19–44, January 2002. URL <https://ideas.repec.org/a/red/issued/v5y2002i1p19-44.html>.
- Thomas F. Cooley and Edward C. Prescott. Economic growth and business cycles. In Thomas F. Cooley, editor, *Frontiers of Business Cycle Research*, chapter 1, pages 1–38. Princeton University Press, Princeton, New Jersey, 1993.
- Mary C. Daly and Bart Hobijn. Downward Nominal Wage Rigidities Bend the Phillips Curve. *Journal of Money, Credit and Banking*, 46(S2):51–93, October 2014.
- Mary C. Daly and Bart Hobijn. The intensive and extensive margins of real wage adjustment. Working Paper Series 2016-4, Federal Reserve Bank of San Francisco, March 2016.
- Karen E. Dynan, Jonathan Skinner, and Stephen P. Zeldes. Do the Rich Save More? *Journal of Political Economy*, 112(2):397–444, April 2004.
- Christopher J. Erceg, Dale W. Henderson, and Andrew T. Levin. Optimal monetary policy with staggered wage and price contracts. *Journal of Monetary Economics*, 46(2):281–313, October 2000.

- Yannis Georgellis and Howard Wall. Gender differences in self-employment. *International Review of Applied Economics*, 19(3):321–342, 2005.
- Francois Gourio and Leena Rudanko. Can Intangible Capital Explain Cyclical Movements in the Labor Wedge? *American Economic Review*, 104(5):183–88, May 2014.
- Robert E Hall. Macroeconomic Fluctuations and the Allocation of Time. *Journal of Labor Economics*, 15(1):S223–50, January 1997.
- Robert E. Hall. Employment Fluctuations with Equilibrium Wage Stickiness. *American Economic Review*, 95(1):50–65, March 2005.
- Steven F. Hipple and Laurel A. Hammond. Spotlight on statistics: Self-employment in the united states. Technical Report 2, U.S. Bureau of Labor Statistics, The address of the publisher, 7 2016. An optional note.
- Robert J Hodrick and Edward C Prescott. Postwar U.S. Business Cycles: An Empirical Investigation. *Journal of Money, Credit and Banking*, 29(1):1–16, February 1997.
- Loukas Karabarbounis. Labor wedges and open economy puzzles. MPRA Paper 31370, University Library of Munich, Germany, October 2010.
- Loukas Karabarbounis. The Labor Wedge: MRS vs. MPN. *Review of Economic Dynamics*, 17 (2):206–223, April 2014.
- Henrik J. Kleven, Martin B. Knudsen, Claus T. Kreiner, Sren Pedersen, and Emmanuel Saez. Unwilling or Unable to Cheat? Evidence from a Randomized Tax Audit Experiment in Denmark. NBER Working Papers 15769, National Bureau of Economic Research, Inc, February 2010.
- Simon Kuznets. *Uses of National Income in Peace and War*. Number kuzn42-1 in NBER Books. National Bureau of Economic Research, Inc, 07 1942.
- Daniel le Maire and Bertel Schjerning. Tax bunching, income shifting and self-employment. *Journal of Public Economics*, 107(C):1–18, 2013.
- Graham S. Lowe and Grant Schellenberg. *What's a Good Job? The Importance of Employment Relationships*. Renouf Publishing Co. Ltd., 2001.

- Dale T. Mortensen and Christopher A. Pissarides. Job Creation and Job Destruction in the Theory of Unemployment. *Review of Economic Studies*, 61(3):397–415, 1994.
- Casey B. Mulligan. A Century of Labor-Leisure Distortions. NBER Working Papers 8774, National Bureau of Economic Research, Inc, February 2002.
- Casey B. Mulligan. Recent marginal labor income tax rate changes by skill and marital status. Working Paper 18426, National Bureau of Economic Research, September 2012a.
- Casey B. Mulligan. Recent Marginal Labor Income Tax Rate Changes by Skill and Marital Status. NBER Working Papers 18426, National Bureau of Economic Research, Inc, September 2012b.
- Simon Parker. *The Economics of Self-Employment and Entrepreneurship*. Cambridge University Press, 2004.
- Michael Parkin. A method for determining whether parameters in aggregative models are structural. *Carnegie-Rochester Conference Series on Public Policy*, 29(1):215–252, January 1988.
- Edward C. Prescott. Why do Americans work so much more than Europeans? *Quarterly Review*, (Jul):2–13, 2004.
- Emmanuel Saez. Do Taxpayers Bunch at Kink Points? *American Economic Journal: Economic Policy*, 2(3):180–212, August 2010.
- Theodore W. Schultz. Investment in human capital. *American Economic Review*, 51(1):1–17, Mar. 1961.
- Robert Shimer. The Consequences of Rigid Wages in Search Models. *Journal of the European Economic Association*, 2(2-3):469–479, 04/05 2004.
- Robert Shimer. Convergence in Macroeconomics: The Labor Wedge. *American Economic Journal: Macroeconomics*, 1(1):280–97, January 2009.
- Robert Shimer. *Labor Markets and Business Cycles*. Number 9217 in Economics Books. Princeton University Press, 2010.
- Mark H. Showalter and Norman K. Thurston. Taxes and labor supply of high-income physicians. *Journal of Public Economics*, 66(1):73–97, October 1997.

A Appendix Figures

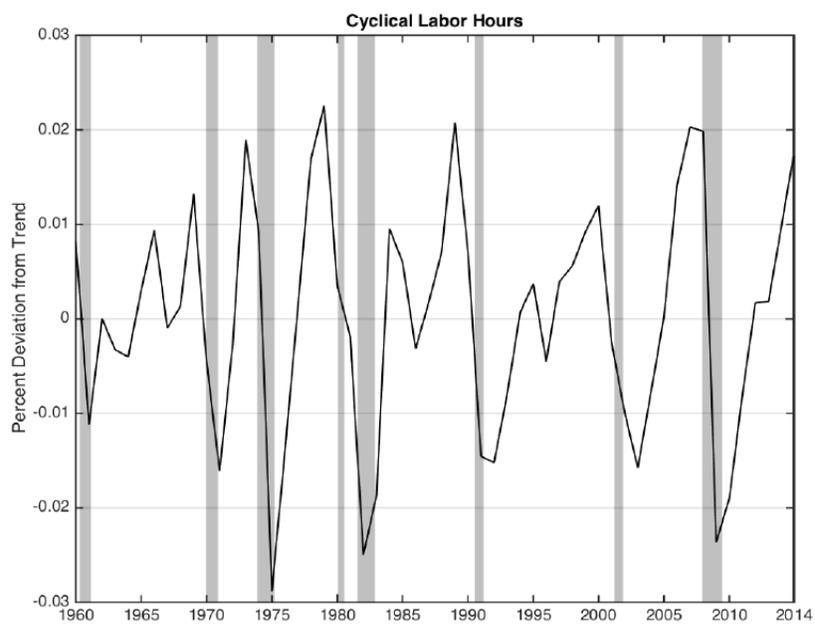


Figure A.1: This figure depicts the baxter-king filtered labor hours over time, generated as in Shimer (2010) from OECD data on hours worked and employment, extended to 2014.



Figure A.2: This figure depicts the labor wedges over time, as calculated from equation 2 and data in Figures 2 and 4.

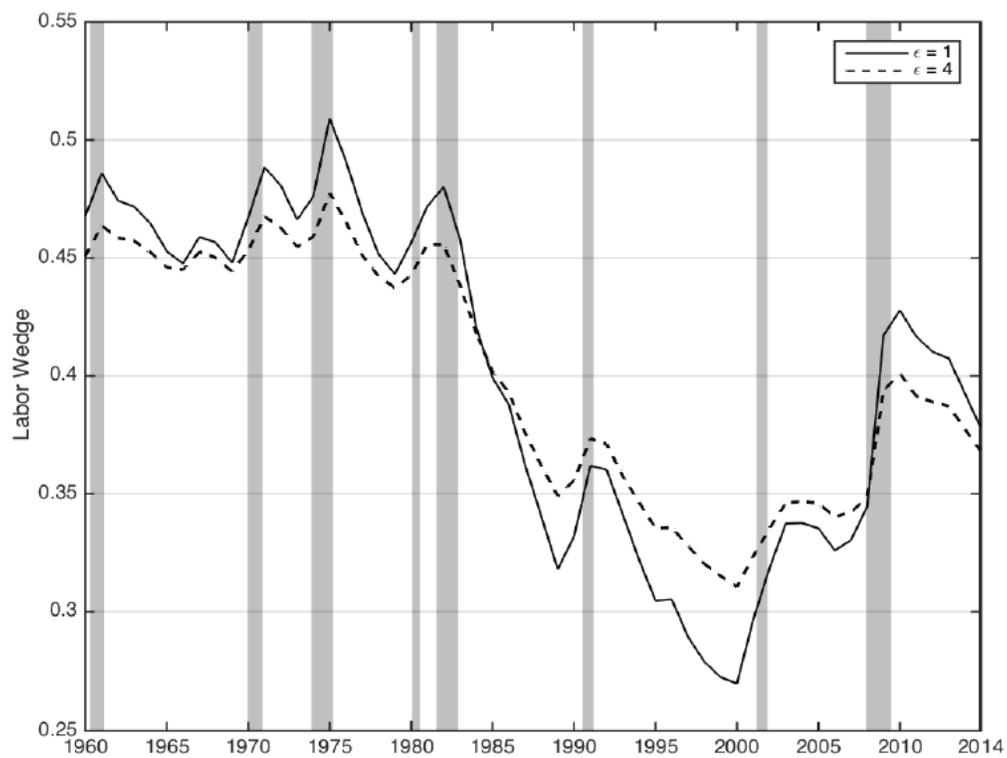


Figure A.3: This figure depicts the labor wedge over time, as measured (and extended from) Shimer (2010).

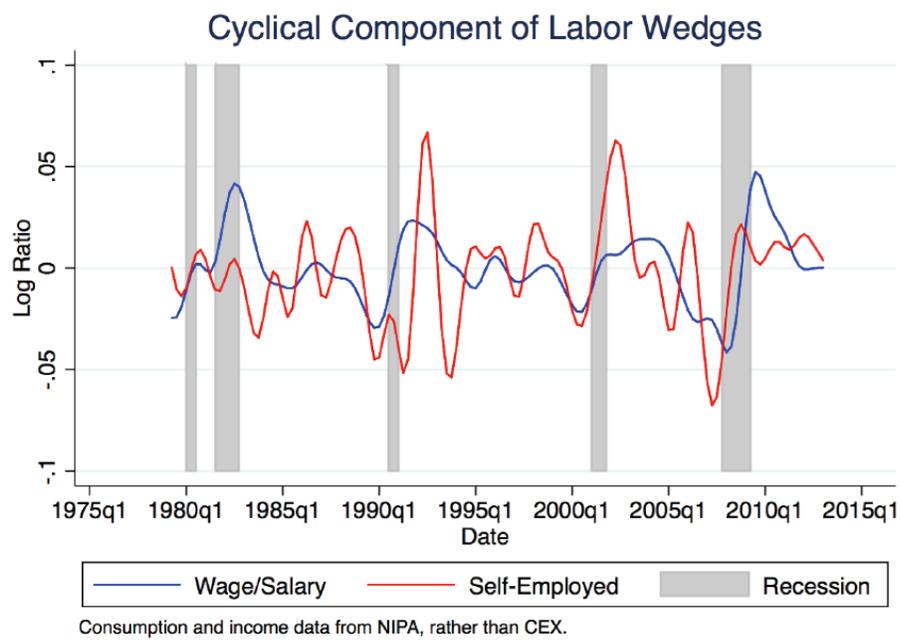


Figure A.4: This figure depicts the cyclical deviations of the tax wedge, using equation 3 and NIPA consumption data, rather than CEX consumption data, and so corresponding to a single household pooling self-employed and wage and salary income. It is comparable to Figure 6 in the main text.