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

Trevor Gallen†

September 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), and rigid wages in general. This paper demonstrates that the cyclical component of labor hours for the self-employed, who are not vulnerable to such frictions, is of comparable cyclicality and volatility as the cyclical component of labor hours for wage and salary workers, even conditioning on wages, consumption, and occupation-industry composition. This finding suggests that explaining the labor hours of the self-employed presents a new test for amplification mechanisms.

JEL Classification: E24, E32, J20, L26

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

*I'd like to thank Yana Gallen and Cathy Zhang for useful comments and suggestions. Andrew Compton for excellent research assistance, and seminar participants at Purdue University and the National Tax Association meetings 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 endogneously 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 employ themselves and therefore are not subject to wage rigidity or employer-employee bargaining frictions. To test this mechanism, I find the self-employed have per-capita labor hour fluctuations of comparable magnitude and cyclicality 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 correlated with those of wage workers, even conditional on industry-occupation composition and are not entirely 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 MPL and log MRS 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, countercyclical 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's cyclical fluctuations unexplained by actual tax-rate measurement translates to a better explanation of labor market fluctuations. After documenting the strong procylicality 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. By amplifying the impact of technology shocks on labor hours, Hall's endogenously 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. If the results of employer-employee bargaining outcomes caused the large swings in the labor hours of wage and salary employees over the business cycle, it would be surprising to find comparable cyclicality in the labor wedge among the self-employed. This paper documents exactly this phenomenon. Given these results, there are a few possibilities. Speculatively, a natural conclusion is that both labor wedges have a significant common amplifying component, such as sticky prices. A second alternative is that sticky wages are indeed the correct amplification mechanism for wage and salary workers, but a separate, parallel mechanism which wage and salary workers are not subject to, but the self-employed are, explains the self-employed labor wedge.[1] A final possibility is that our measurement of the labor wedge is incorrect: for instance, if the self-employed had a dramatically higher labor supply elasticity with respect to temporary wage shocks, or different changes in MRS (perhaps due to different household production technologies). Whatever the explanation, a test of any model of the labor wedge is its ability to convincingly explain both the self-employed and wage/salary labor wedges.

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. There is ample evidence that monetary wages within wage and salary jobs are rigid (Barattieri et al., 2014; Bihan et al., 2012). The central assumption of this paper is that the self-employed are less subject to the sorts of bargaining, contracting, and principle-agent problems that wage and salary workers (and their firms) are subject to. While the self-employed are still subject to goods-market firm matching frictions (such as the price frictions of Calvo (1983)), they are not

[1] For instance, if self-employed hours were more affected by financial frictions that increased during recessions.

subject, or are 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.[2] 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). 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 cyclicality of both wages and hours worked per worker. I differ in three important ways. First, I examine total hours worked.[3] 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 examine hours at a quarterly, rather than yearly, frequency, allowing for a better cyclical measure.

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 of comparable cyclicality, this paper contributes in four important ways. First, I incorporate total hours, rather than only the intensive margin for self-employed. This is important because the primary driver of labor hours at quarterly frequency is the extensive employment margin.[4] 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:

[2]Some heterogeneous agent models have also had success in explaining the representative-agent's labor wedge through the combination of incomplete capital markets and indivisible labor shocks (Chang and Kim, 2007).

[3]The lion's share of total hours fluctuations at the quarterly frequency comes from fluctuations in employment, rather than hours per worker (for instance, Cho and Cooley (1994) finds 75%). Moreover, some contend that much of the bargaining frictions between workers and employers affect the extensive margin (Bewley, 2007). Using total hours avoids these potential measurement issues by using the full variation in hours worked.

[4]Section 2.1 of Bils et al. (2015) examines the intensive and extensive labor wedges and find them of comparable volatility.

this higher frequency potentially allows me to more accurately examine movements at business-cycle frequencies. Third, I extend the self-employed labor wedge back to 1980, rather than 1987 onwards, including two more recessions. Finally, and perhaps most importantly, while Bils et al. do adjust for the worker class composition of twelve industrial categories, this paper uses the vastly larger American Community Survey to examine in detail the degree to which the industrial-occupational differences of jobs held by worker class plays a role in generating similar labor hours movements between the two classes, finding that they do not substantially contribute to the cyclical variation in hours.[5]

If the self-employed are subject to lower wage or search frictions but display a comparable wedge between MRS and MPL, it is unlikely that wage and search frictions are the sole explanation for the labor wedge, in spite of its relative success for wage and salary workers. 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 and nonagricultural average hours worked for both self-employed and wage/salary workers to construct a measure of total hours by worker class. 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 the cyclical component of the self-employed labor wedge is strongly related to both the cyclical components of both the wage/salary labor wedge and GDP.

There are many other potential sources for the labor wedge that my findings do not exclude. Constructively, any friction that hits the self-employed in the same way as wage and salary workers is a more likely candidate to explain their comovements, 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 as the sole mechanism.

The rest of this paper is organized as follows: section 2 describes the model I use to measure the labor wedges. Section 3, 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

[5]Because consumption, and even income data is relatively more sparse and noisy for the self-employed, if my extremely-detailed controls for occupation-industry proxies well for consumption and income changes, cyclical variation in hours alone, adjusting for occupation-industry, informs us about the similarity of the labor wedges.

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 measurements of the labor wedge, time series on consumption, income, and labor are used. Consumption informs us about the marginal utility of consumption, (a necessary component of the marginal rate of substitution between consumption and leisure), and earnings given labor informs us about the marginal product of labor. Labor hours worked gives the marginal utility of leisure. Consumption given income contains information about what a household believes about future MPL: when the ratio of consumption and income is high, households believe MPL is 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 the 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 change little, if at all. 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 for two reasons. First, with substitutability between labor types, a decline in wage/salary labor hours ceteris paribus reduces marginal disutility of labor supply. If the self-employed labor wedge hasn't risen, self-employment should also rise. Second, when the wage/salary labor wedge increases and income declines, the marginal utility of income rises and self-employment should increase. With both effects tending to increase self-employment, the only way for our model to reconcile a decrease in wage/salary employment with no change in self-employment would be for the self-employment labor wedge to also have increased.

By 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. 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. Removing this assumption only strengthens my main conclusions. Appendix B formalizes this intuition mathematically.

The formulas for the two labor wedges therefore are similar to Shimer (2009) and Shimer (2010). With worker of class $i \in \{W, SE\}$ (where W denotes wage/salary workers and SE denotes the self-empoyed) 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_t 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).

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

Multiplying top and bottom by $L_{i,t}$, so that labor earnings ($E_{i,t} = w_{i,t}L_{i,t}$) are in the denominator, we get the simple expression for the market-clearing labor wedge in terms of two parameters and three data series:

$$\tau_{i,t} = 1 - \psi_i \frac{c_{i,t}}{E_{i,t}} L_{i,t}^{\frac{1+\epsilon_i}{\epsilon_i}} \tag{2}$$

I may therefore use data on $c_{i,t}$, $E_{i,t}$, $L_{i,t}$, need to calibrate ψ_i and ϵ_i to calculate $\tau_{i,t}$.[6] 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_t, the joint household consumption. I do so in equation (3), which provides an alternative "representative family" model of the labor wedge:

$$\tau_{i,t} = 1 - \psi_i \frac{c_t}{E_{i,t}} L_{i,t}^{\frac{1+\epsilon_i}{\epsilon_i}} \tag{3}$$

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 $w_{i,t}$, 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), with the same qualitative result of the self-employed labor wedge's cyclicality. I also calculate the same wedge using E_t (rather than $E_{i,t}$) as well as c_t, which would be the case if total earnings better measured the MRS of each type.

3 Data

To construct the labor wedge, I use data on six data series: $c_{i,t}$, $E_{i,t}$, and $L_{i,t}$ for both the self-employed and wage/salary workers, and need to calibrate ψ_i and the Frisch elasticities of

[6] An important assumption embedded in this formulation is that wages reflect the marginal product of each type of worker.

labor ϵ_i. I turn to my parameter choices for ψ_i and ϵ_i, then to the data.

3.1 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. 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 or 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 compared to 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 ψ_i so the average measured tax wedge for both groups equal to 0.4, as in Prescott (2004)[7].

3.2 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 construct consumption from the CEX by taking total household expenditures excluding spending on education, health, insurance and pensions, and new vehicle purchases. 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 (educational spending makes up about 2% of spending for both types of households I analyze), I include it as investment.

To construct my two time series, I use the quarterly CEX public use microdata. I depict the hp-filtered quarterly series of C/E in Appendix Figure A.1. The two series show a different secular pattern to the National Income and Product Account's C/E 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 cyclicality 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 cyclicality of the CEX differs by income class. The stronger countercyclical nature of the CEX's series on consumption as a fraction

[7]Prescott (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 ψ_i has changed differentially by income class over time.

of labor earnings for the self-employed compared to wage and salary workers is evident from Figure A.1: this will act to reduce the volatility of the self-employed labor wedge.[8] Because my results find strong commonalities between labor wedge cyclicalities, 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 both wage and salary workers, and for unincorporated self-employed workers.[9] Similarly, it publishes series on average hours at work by worker class.[10] I calculate total hours per capita for each group by taking the product of employed workers and average hours at work per worker and dividing it by the working-age population. 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 increases 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 depicted in Figure 1. 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. While my primary results use the more common high-pass Hodrick-Prescott (HP) filter (Hodrick and Prescott, 1997), and as a consequence retain this high-frequency noise, I also display the results using the band-pass Baxter-King (BK) filter (Baxter and King, 1999), which filters out both high and low frequencies, in the Appendix.

[8]The cyclical elasticities of hp-filtered wage and salary consumption/earnings is 0.74, while the cyclical elasticity of self-employed consumption rate is -0.41, with large standard errors.

[9]BLS series LNS12032187 and LNS12032192.

[10]BLS series LNS12033251 and LNU02033377.

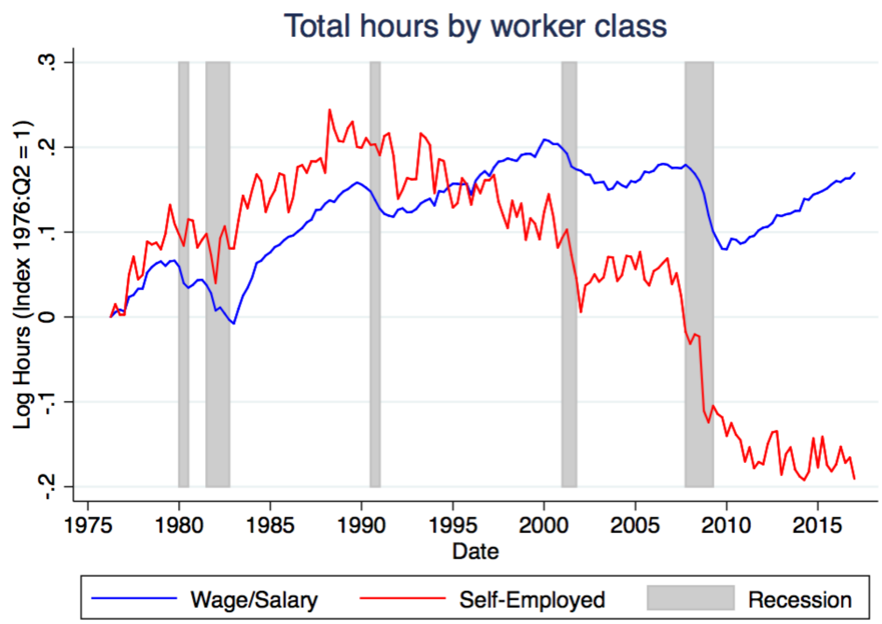

Figure 1: This figure depicts indexed log hours per working age person by worker class, combining average hours per worker with employment to calculate total hours for each worker class.

I depict the (HP-filtered) cyclical component of log total hours per working-age person by worker class in Figure 2. Figure 2 makes clear the core hours data that produces the countercylical 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).[11]

4 Results

Using the data and parameters discussed in Section 3, I calculate the labor wedges with equation 2, HP-filter each tax wedge and display them in Figure 3. The wedges have a clear common component over the business cycle, the core result of this paper. The self-employed labor wedge peaked higher in the 1981-1982 recession and the 2001 recession. The wage/salary labor wedge

[11]It 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.

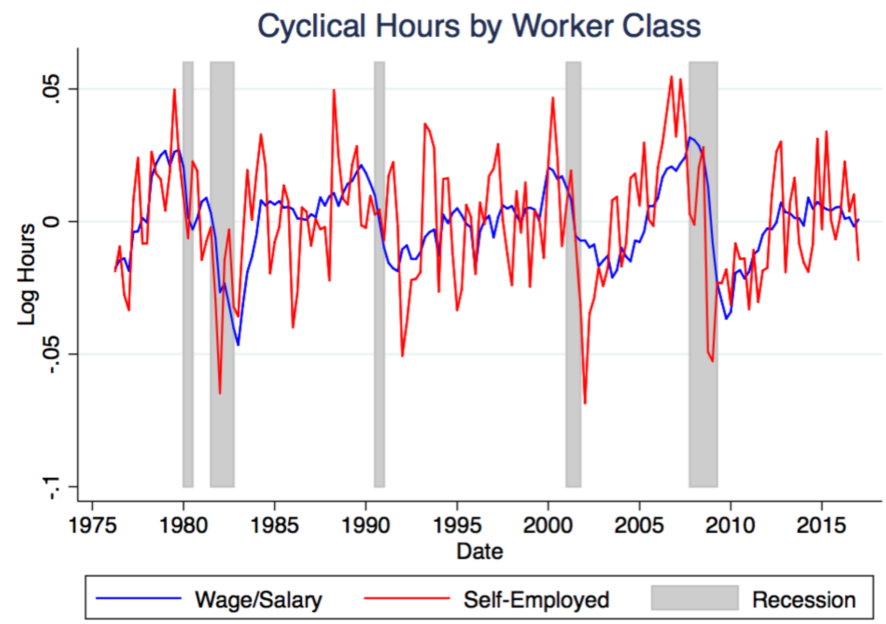

Figure 2: This figure depicts the core data that generates the labor wedge: the cyclical (HP-filtered) hours of wage/salary workers and the self employed. Importantly, the magnitude of fluctuations are highly comparable even though the self-employed are not, or less, subject to wage rigidity.

peaked more during the 1980 recession, (during which the self-employed labor wedge fell), and while both wedges rose from -0.065 in 2006:Q4 to 0.019 for wage and salary workers and 0.038 for self-employed workers in 2008:Q4, the wage and salary labor wedge continued to rise thereafter, while the self-employed wedge fell.

There are a few periods when the two do not broadly mimic one another. The self-employed labor wedge clearly is far more volatile at high frequencies, with a standard deviation of 0.066, compared to 0.037 for wage and salary workers. Data limitations doubtlessly play a role: beyond the significant rounding limitations of the self-employed hours data I use, the Consumer Expenditure Survey typically surveys 21,000 households, less than 10% of which are self-employed, resulting in noisy measures of both consumption and income. Indeed, because income and employment questions are only asked of the 40% of the sample in their second or fifth months, misstatement of a small set of families can result in mismeasurement at frequencies longer than a quarter.

To help understand what underlies the quantitative similarities in Figure 3, I regress a transformed, HP-filtered labor wedges and their components on HP-filtered log GDP. I transform

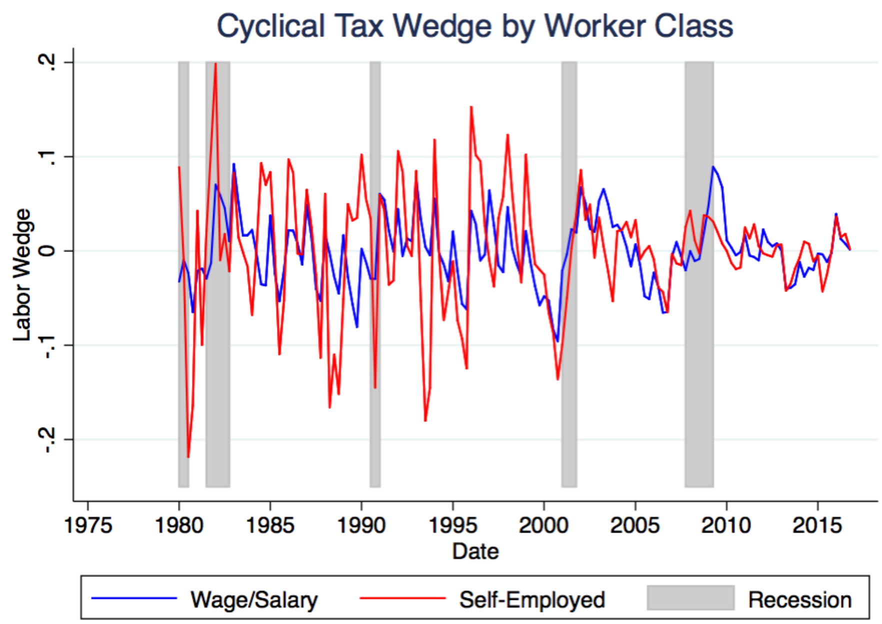

Figure 3: This figure depicts HP-filtered self-employed and wage and salary labor wedges. Comparable band-pass filtered set of wedges is shown in Figure A.2 in the Appendix.

the labor wedge and components by subtracting both sides of equation 2 from one, logging, and multiplying by negative one, which decomposes the labor wedge into four linear contributing parts: the constant, log consumption, log income, and log total hours, weighted by the Frisch Elasticity:

$$-\log(1-\tau_{i,t}) = -\log(\psi_i) - \log(c_{i,t}) + \log(E_{i,t}) - \frac{1+\epsilon}{\epsilon}\log(L_{i,t}) \tag{4}$$

Moreover, I offer several different alternative specifications of the labor wedges in the same table. The first column in the first and second panels of Table 1 reports the cyclical elasticities of each transformed labor wedge (or its components, hp-filtered $-\log(c_{i,t})$, $\log(E_{i,t})$, and $-\frac{1+\epsilon}{\epsilon}\log(L_{i,t})$). The third panel reports statistics of interest: the coefficient of a regression of the transformed self-employed tax wedge on the transformed wage and salary tax wedge, and their correlation. The third panel is particularly informative if the two wedges comove together in a way not completely captured by GDP, as appears to be the case. Because the HP-filter induces autocorrelation, Newey-West standard errors with a six-year window are used in all regressions (Newey and West, 1987).

The second column of Table 1 uses only intensive hours for the same regression: because intensive hours elasticities are lower than total hours elasticities, I use an intensive elasticity of 0.54, consistent with Chetty et al. (2011). This regression would be more informative than the total hours regression if, for instance, financial frictions affected the extensive margin of labor choice for the self-employed more than they do for wage and salary workers. The third column of Table 1 uses jointly-defined consumption, consistent with the alternative specification in equation 3. Because visually the two wedges in Figure 3 appear to be more correlated post-2000 (0.61 rather than 0.53), the final column displays the main regressions restricted to the first quarter of 2000 and after.[12]

The third panel of Table 1 quantitatively summarizes the optical results of Figure 3, suggesting that the two wedges comove together: a 1% increase in the wage and salary labor wedge typically is coincident with a 0.82% increase in the self-employed labor wedge. The significant correlation between the two cyclical wedges (0.53) is comparable in magnitude to the correlation between cyclical GDP and the cyclical wage and salary labor wedge (-0.64). Moreover, the correlation of their quarterly difference is even higher (0.57), even as the correlation with GDP's quarterly difference and the wage/salary labor wedge falls markedly (-0.26). However, the first panel also makes clear that the covariance of the cyclical wage and salary labor wedge with cyclical GDP is twice the covariance of the cyclical self-employed labor wedge with cyclical GDP. Together, the two suggest that the two wedges are moving together in a way not completely captured by GDP. The component regressions and alternative specifications shed light on why the two wedges diverge in terms of their cyclicality with GDP: while self-employed labor is more cyclical than wage and salary labor, and self-employed consumption is more cyclical, both of which tend to make self-employed labor wedge more cyclical, self-employed income in the CEX is strongly countercyclical (albiet with very large standard errors), unlike wage and salary income.

Telling a similar story, while the use of intensive hours in column two reduces the cyclical responsiveness of both tax wedges, causing the self-employed labor wedge to display a cyclical elasticity around 40% of the wage and salary labor wedge, their correlation and cyclical relationship increases markedly. Defining consumption jointly as in the alternative specification of equation 3 changes the cyclical elasticities little, and decreases their correlation and the elasticity

[12]Figures analogous to Figure 3 but for intensive hours and joint consumption can be found in Appendix Figures A.3 and A.4

Table 1: Cyclical relationships of the labor wedges and their Components

$\log(\xi_{i,t}^{HP}) = \beta \log(Y_t^{HP}) + \epsilon_{i,t}$

	Main Specification	Intensive Hours	Joint Consumption	Post-2000 (Main)
WS Wedge	-3.04	-1.59	-3.09	-3.33
	(0.34)	(0.33)	(0.33)	(0.49)
WS Consumption	-0.50	-0.50	-0.59	-0.77
	(2.83)	(2.83)	(2.83)	(1.45)
WS Income	-0.40	-0.40	-0.40	-0.26
	(5.21)	(5.21)	(5.21)	(2.39)
WS Labor	-2.73	-0.72	-2.73	-3.01
	(1.34)	(0.30)	(1.34)	(1.01)
SE Wedge	-1.45	-0.68	-1.47	-2.53
	(0.62)	(0.50)	(0.49)	(0.87)
SE Consumption	-0.83	-0.83	-0.59	-1.64
	(3.05)	(3.05)	(2.83)	(1.28)
SE Income	1.75	1.75	1.75	2.66
	(7.55)	(7.55)	(7.55)	(3.85)
SE Labor	-2.99	-1.40	-2.99	-4.25
	(2.43)	(1.32)	2.43	(3.31)
Corr(τ_{SE},τ_{WS})	0.53	0.59	0.46	0.61
β SE on WS	0.82	0.98	0.75	0.86

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 1: This first two panels of this table depict the results of regressions of the cyclical labor wedges and their components on corresponding cyclical controls. All standard errors are calculated using Newey-West standard errors with a six-year window (three year lag). The first and fifth rows are coefficients of regressing the hp-filtered log $1 - \tau_{i,t}$ on hp-filtered log GDP. The successive rows are the results of replacing the independent variable in the regression with $-\log(c_{i,t})$, $\log((E_{i,t}))$, and $-\frac{1+\epsilon}{\epsilon}\log(L_{i,t})$, where each series is HP-filtered after logging. The third panel of this row displays the correlation between the two HP-filtered wedges, and the cyclical elasticity of the self-employed labor wedge with respect to the wage and salary labor wedge. The first column includes the main specification and data of equation 2. The second column replaces the labor hours measure with the intensive hours of each worker class, and reduces the elasticity used in calculations to 0.54. The third column examines the "representative household" model of equation 3. The final column replicates the first, but only includes dates after 1999:Q4. All statistics in this table use transformed tax wedges.

of self-employment wedge with the wage and salary wedge. The last column, which replicates the main regression but restricts it to the first quarter of 2000 and beyond, suggests that the difference in cyclicality of the two wedges decreased markedly after 2000. In general, the cyclicality of the self-employed wedge has increased over time, much more than the wage and salary labor wedge.

In conjunction with Table 1, equation 4 makes clear that the contribution of labor hours to the labor wedge, which is substantial, is strongly contingent on the elasticity of labor supply. If, for instance, the labor supply elasticity of the self employed was actually 1.5, rather than 0.75, the cyclical coefficient on the self-employed wedge would fall significantly (to -0.9 from -1.54), reflecting its reduced cyclical elasticity. Intuitively, a higher elasticity means a smaller cyclical wedge is needed to explain labor market movements, which are now more responsive to temporary changes in wage.

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 then 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 2000-2015[13] Because unemployed workers are classified according to their last occupation or industry, 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 per worker 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

[13] Including the 2000 Census.

population. I then apply those occupation-industry fixed effects $\hat{\beta}_t^{SE}$ to the wage and salary population[14]. In other words, I take the occupation-industry trends for the self-employed workers and apply them to wage and salary workers, "adjusting" the self-employed occupation-industry composition. Figure 4 compares the predicted hours change with those of the self-employed.

It would be troubling if composition-adjusted self-employed hours in Figure 4 were much less volatile, as it would suggest that self-employed hours volatility was a product of occupation-industrial composition, rather than reflecting a true similarity with wage and salary workers. 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 HP-filtered version is depicted in Figure 5 and shows that industrial-occupational differences play a minimal role in cyclical hours differences.

$$\text{Hours}_{i,t} = \alpha + \sum_{t=2000}^{2015} \beta_t^{SE} X_{i,t}^{SE} \tag{5}$$

When we acknowledge the limitations of the Consumer Expenditure Survey's consumption and income data for the self-employed, Figure 5 also highlights another possibility. The tax wedge is the residual "error" of hours after controlling for income and consumption. If savings within occupation-industry groups is very similar across worker class, then it may be that controlling for occupation-industry controls for consumption and income. There is economic reason to think that they should be the same: if specific skill sets, industries, or tasks are affected with permanent and transitory shocks, we might expect savings within occupation-industry categories to be very similar across worker class.

Unfortunately the CEX does not collect detailed data on worker industry, and only has 18 occupational codes (compared to 57 industries and 480 occupations in the 2015 ACS). However, I offer suggestive evidence for this hypothesis by examining how much of the between-group variation in savings occupation explains. Using the 1998-2000 CEX, I find that introducing 18 occupation dummies to year and worker class controls explains 37% of the between-group variance. It is reasonable to expect more detailed industry-occupational controls may explain

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

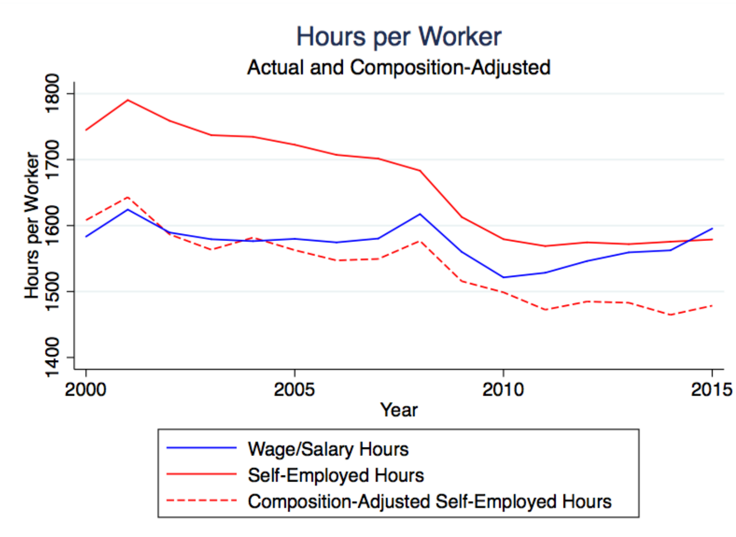

Figure 4: Figure 4 depicts actual hours per self-employed worker and wage/salary worker, along with hours predicted if self-employed workers had the same industrial composition as wage and salary workers. Insofar as the two are similar we may conclude that occupation-industrial composition does not play 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.

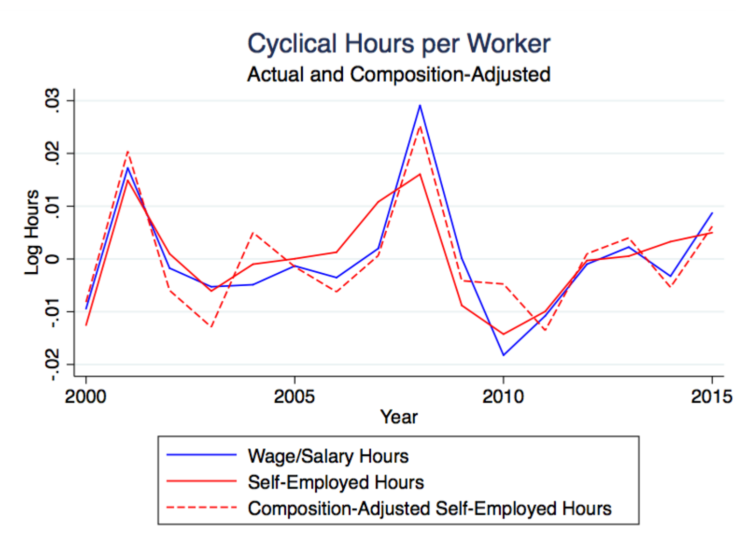

Figure 5: This figure displays the HP-filtered hours of self-employed and wage/salary workers, along with HP-filtered hours of self-employed workers if they displayed the same occupation-industry pattern as wage and salary workers. The figure shows that in spite of the secular decline in 4, the cyclical drop in hours per worker for the self-employed is very similar to that of wage and salary workers.

a significant portion of savings behavior. Because labor hours are better and more frequently measured than consumption and income, establishing the degree to which occupational variables can control for income shocks across groups is an important step for future research, and may help in the measurement of the labor wedge for specific subgroups which is currently difficult due to data availability in the CEX.

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, whose labor hours are not subject to employer-employee bargaining frictions, display a similar time series pattern as wage and salary hours, and display a labor wedge of comparable cyclicality as that of wage and salary workers. More constructively, periods of time in which the self-employed labor wedge does not increase, while the wage and salary labor wedge does, such as after the first quarter of 2009 onwards, may help us diagnose when employer-employee frictions are likely to be the cause.

If the the business cycle fluctuations in hours of the self-employed look similar to the fluctuations in hours of wage and salary workers, what does this mean for the labor wedge? First, it may mean nothing: if the self-employed labor wedge is mismeasured (perhaps due to a higher Frisch elasticity) or if the self-employed savings rate falls in a way that explains labor movements, then the similarity between labor hours may be a red herring. Second, an alternative amplification mechanism, common between types of workers, may explain both labor wedges simultaneously. Two particularly fruitful examples used to explain the Great Recession might be non-wage price rigidities (Christiano et al., 2015), or changes in implicit marginal tax rates by group (Mulligan, 2012a). A third possibility is that wage rigidities are indeed a good explanation for wage/salary workers, and that the self-employed have another explanation that wage/salary workers are not subject to, but which gives rise to the comparable fluctuations in labor hours.

At the most basic level, this paper's contribution suggests that examining the parallel labor behavior of wage and salary workers with the self-employed is a way to fruitfully select among competing explanations of the labor wedge.

References

Orhan Erem Atesagaoglu and Ceyhun Elgin. Cyclicality 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.

T.F. Bewley. *Why Wages Don't Fall During a Recession*. Harvard University Press, 2007.

Herv Le Bihan, Jrmi 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.

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.

Yongsung Chang and Sun-Bin Kim. Heterogeneity and aggregation: Implications for labor-market fluctuations. *American Economic Review*, 97(5):1939–1956, December 2007. doi: 10.1257/aer.97.5.1939.

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.

Jang-Ok Cho and Thomas Cooley. Employment and hours over the business cycle. *Journal of Economic Dynamics and Control*, 18(2):411–432, 1994.

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.

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.

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.

Whitney K. Newey and Kenneth D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, 1987. ISSN 00129682, 14680262.

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

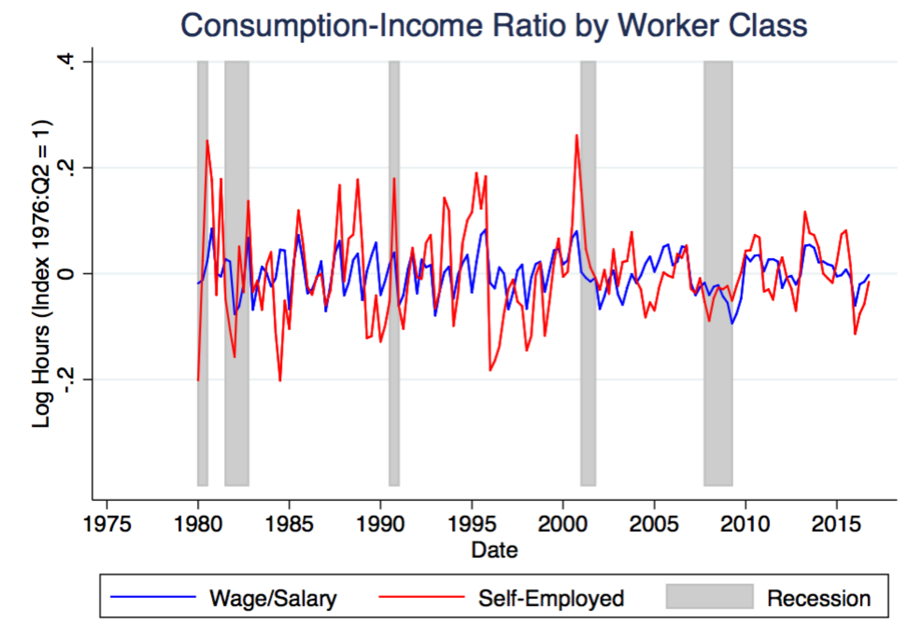

Figure A.1: This figure depicts the ratio of consumption and income by worker class calculated from CEX public use microdata from 1980-2016.

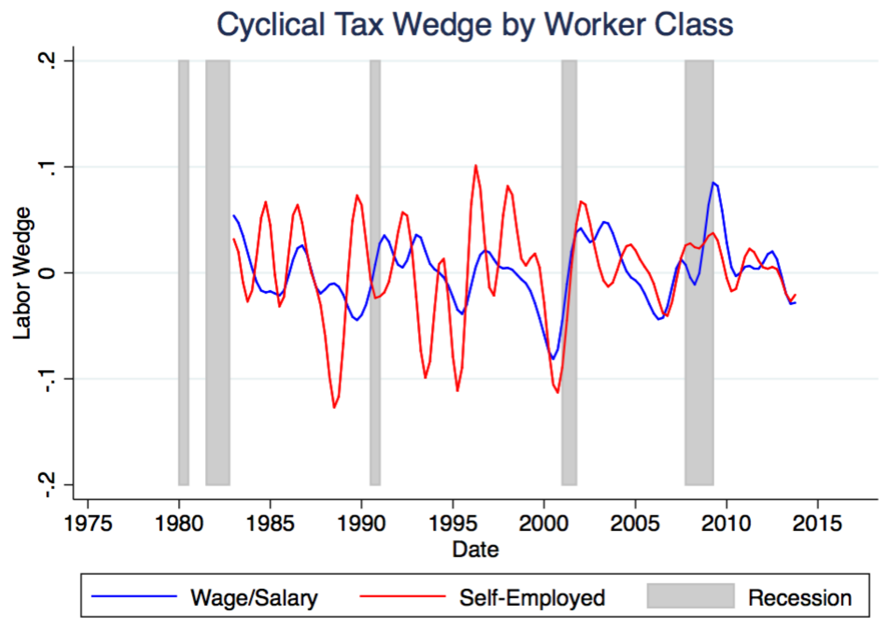

Figure A.2: This figure depicts the same labor wedge as in Figure 3, but using the Baxter-King band-pass filter rather than the Hodrick-Prescott high-pass filter. This has the advantage of reducing some, but not all of the noise generated by the Consumer Expenditure Survey's low number of observations for the self-employed.

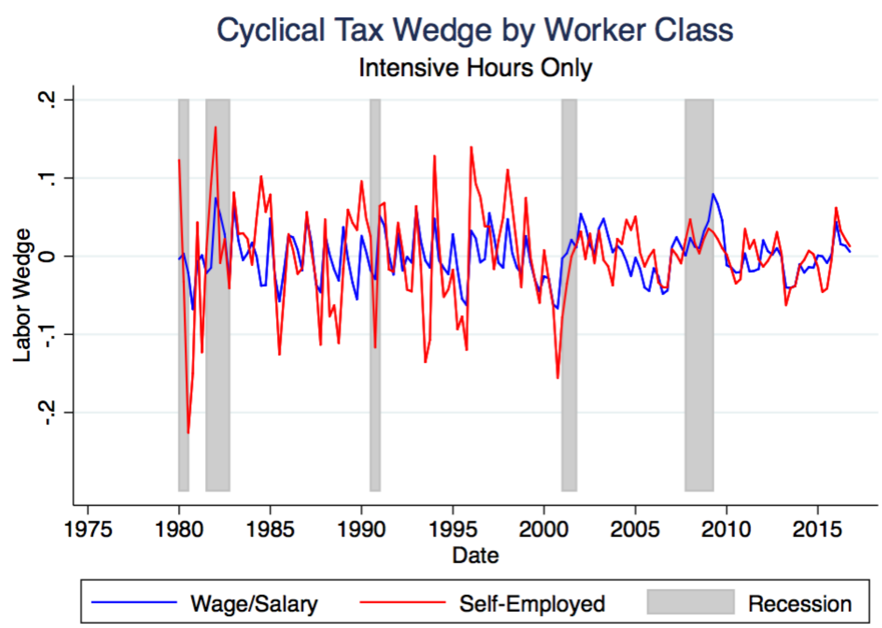

Figure A.3: This figure depicts the same labor wedge as in Figure 3, but uses only the intensive margin of hours and intensive Frisch elasticity. The two figures are very similar, with labor wedge correlations of 0.94 and 0.91 between the two wage/salary and self-employed respectively.

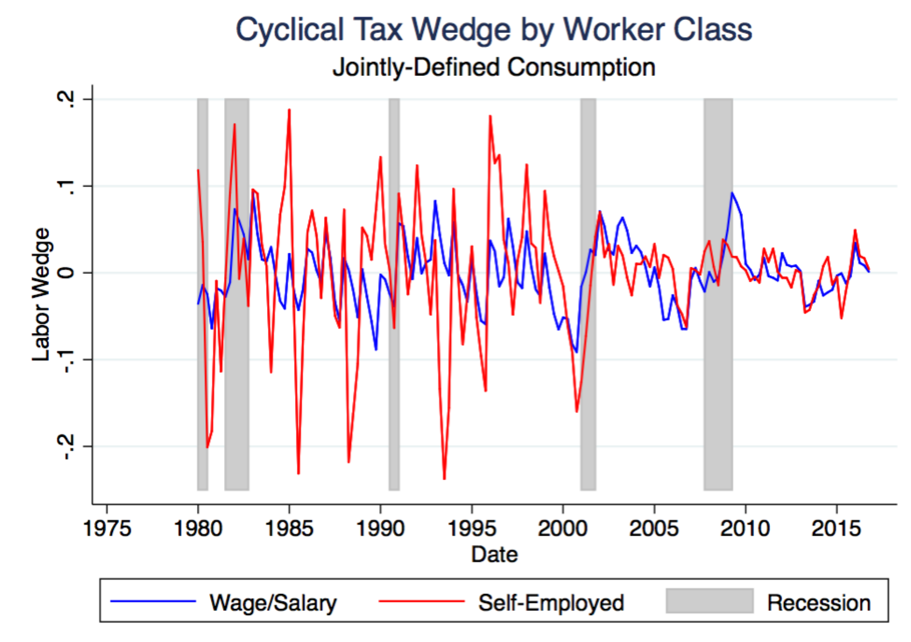

Figure A.4: This figure depicts the cyclical deviations of the tax wedge, using equation 3 and using average consumption of both wage and salary workers, which corresponds to a single household pooling self-employed and wage and salary income. It is comparable to Figure 3 in the main text.

B Mathematical Appendix

This section establishes that if we allow for substitutability between labor types, a decrease in wage/salary labor hours will, ceteris paribus, cause the calculated self-employed labor wedge to increase.

I allot the representative household two workers: a self-employed worker and a wage/salary worker. Total utility from leisure is generated by a combination of the two worker's leisure hours, ℓ_W and ℓ_{SE}, where the two types of leisure are substitutes. For simplicity, I assume each type of worker has total amount of time equal to one, and that wages are unity, so that utility is given by:

$$u(c) + v(\ell_W, \ell_{SE})$$

Subject to the budget constraint:

$$c = (1 - \tau_W)(1 - \ell_W) + (1 - \tau_{SE})(1 - \ell_{SE})$$

Plugging in for c, first-order conditions are:

$$v_{\ell_W} = (1 - \tau_W)u'(c)$$

$$v_{\ell_{SE}} = (1 - \tau_{SE})u'(c)$$

We're interested in the business-cycle behavior of the labor wedge, in temporary shifts in τ_W and τ_{SE}. Such temporary changes largely shut off the income effect on labor supply, as the effect on lifetime income of temporary changes in wages is small. With this in mind, I set the marginal utility of consumption equal to a constant λ, as would be true, for instance, in a continuous time model if τ_W changed for an infintesimally small amount of time. So, the first order conditions can be written:

$$v_{\ell_W} = (1 - \tau_W)\lambda$$

$$\frac{v_{\ell_{SE}}}{v_{\ell_W}} = \frac{1 - \tau_{SE}}{1 - \tau_W}$$

Totally differentiating both equations (and noting $d\lambda = 0$ by assumption):

$$\lambda d\tau_W + v_{\ell_W,\ell_{SE}} d\ell_{SE} + v_{\ell_W,\ell_W} d\ell_W = 0 \tag{6}$$

$$\frac{1}{1-\tau_W} d\tau_{SE} - \frac{(1-\tau_{SE})}{(1-\tau_W)^2} d\tau_W + \left(\frac{v_{\ell_W,\ell_{SE}}}{v_{\ell_W}} - \frac{v_{\ell_{SE}} v_{\ell_W,\ell_W}}{v_{\ell_W}^2} \right) d\ell_W + \left(\frac{v_{\ell_{SE},\ell_{SE}}}{v_{\ell_W}} - \frac{v_{\ell_{SE}} v_{\ell_W,\ell_{SE}}}{v_{\ell_W}^2} \right) d\ell_{SE} = 0 \tag{7}$$

Theorem B.1. *If self-employed leisure and wage/salary leisure are substitutes in utility, then increases in the wage/salary labor wedge without an increase in self-employment (a decrease in self-employed leisure) implies that the self-employed labor wedge must have increased. (i.e. when* $d\tau_W > 0$ *and* $d\ell_W > 0$ *commensurate with equation 6, then if* $d\ell_{SE} = 0$*,* $d\tau_{SE} > 0$*.*

Proof. Taking equation 6 and solving for $d\tau_W$, plugging in $d\ell_{SE} = 0$

$$d\tau_W = -\frac{v_{\ell_W,\ell_W}}{\lambda} d\ell_W$$

Plugging this into equation 7, and plugging in for $v_{\ell_W} = (1-\tau_W)\lambda$ and $v_{\ell_{SE}} = (1-\tau_{SE})\lambda$

$$\frac{1}{1-\tau_W} d\tau_{SE} + \frac{(1-\tau_{SE})}{(1-\tau_W)^2} \left(\frac{v_{\ell_W,\ell_W}}{\lambda} \right) d\ell_W + \left(\frac{v_{\ell_W,\ell_{SE}}}{(1-\tau_W)\lambda} - \frac{(1-\tau_{SE})\lambda v_{\ell_W,\ell_W}}{(1-\tau_W)^2\lambda^2} \right) d\ell_W = 0$$

Becomes:

$$d\tau_{SE} = -\frac{v_{\ell_W,\ell_{SE}}}{\lambda} d\ell_W$$

Given $d\ell_{SE} = 0$, $d\ell_W > 0$ and $\lambda > 0$, then $d\tau_{SE} > 0 \iff v_{\ell_W,\ell_{SE}} < 0$, i.e. self-employed leisure and wage/salary leisure are substitutes in utility. ■

As $v_{\ell_W,\ell_{SE}} \to 0$ (separability in utility, as in the main model of the paper), then $\frac{d\tau_{SE}}{d\ell_W} = 0$: when wage/salary leisure increases, but self-employed leisure doesn't increase, there is no evidence that the self-employed labor wedge increases. But as soon as two types become more substitutable, so $v_{\ell_W,\ell_{SE}} \to v_{\ell_W,\ell_W} = v_{\ell_{SE},\ell_{SE}} < 0$, then $\frac{d\tau_{SE}}{d\ell_W} > 0$.

C Data Appendix

This appendix discusses the data series I use and their construction.

C.1 Consumer Expenditure Survey

To obtain average consumption and income within a worker class, I use the standard CEX "weighted mean" methodology outlined in each codebook. However, there are several assumptions and choices that go into defining a household and comparing surveys (income in particular) across years, which I delineate in detail here.

C.1.1 Defining a household

I define three household types: "self-employed," "wage and salary workers" and "other," only using the first two categories. In defining each type, I use the income variable "INCOMEY1," which is the "employer from which reference person recieved the most earnings in past 12 months." In most years defines six types of businesses a worker: 1) private company (business or individual), 2) Federal government 3) State government 4) Local government 5) Self-employed in own business, professional practice or farm 6) Family business or farm, working without pay. I define the first four categories as "wage/salary" and the last two as "self-employed."[15]

If the reference person has no response, I examine their spouse, and conditional on a working spouse with a valid "INCOMEY2" variable, allocate the household using their employment. This increases the average self-employed sample by approximately 5% of the sample (for instance, in 1996, it adds 105 self-employed and 1,032 wage/salary observations onto 1,847 self-employed and 16,310 wage/salary families).

C.1.2 Defining consumption

Fortunately, while CEX Universal Classification Codes (UCCs) change over time and are retired or created, each year the CEX defines sets of broad categories each UCC code goes into (such as "Alcoholic Beverages," "Health," "Insurance and Pensions, " "Apparel and Services" and so on. From 1980-2016, I use the following broad categories:

- Average annual expenditures

[15] There are typically only 5 or 6 of the last category, approximately 7% of the total "self-employed sample).

- Health spending
- Education spending
- Vehicle Purchases (net outlay)
- Contributions to Insurance and Pensions

Where I define "consumption" as the difference between average annual expenditures (which includes all of the subsequent categories) and the four categories below. This is in keeping with the general idea that health and education spending are investments in human capital, rather than consumption, that vehicle purchases are investments in a durable good, and that contributions to insurance and pensions are savings, not consumption.

C.1.3 Defining income

In order to use (after-tax) earnings $E_{i,t}$, I exclude retirement and survivors disability income (including social security, private, and other government retirement sources), interest and dividends, room rental income, royalty, estate, and trust income, public assistance, supplemental security income and other government transfer programs, unemployment and workers compensation, and "other income" from my income definition. I only use wage and salary income and self-employment income, and subtract total personal taxes, which includes federal, state, local, and other miscellaneous tax expenditures (such as vehicle taxes).

C.1.4 Changes in the CEX

In 2003, there was a major change in the methodology used to calculate income. Specifically, the CEX started imputing values for nonresponses to the income questions. As a consequence, income series before and after 2003 are not directly comparable: indeed, income jumps by approximately 23% for wage/salary workers and by 18% for the self-employed. To fix the issue, I anneal the two data series by proportionately decreasing post-2003 income series by a constant ω (a different ω for each income type).[16] I estimate ω by estimating the linear trend for each of the four dates in 2003 and choosing ω such that the first quarter of 2004 fits this time trend. The unadjusted and adjusted income series can be seen in Figure C.5.

[16]Because the level of C/E is unimportant due to my normalization of each average labor wedge to 0.4, proportionately increasing pre-2003 income would yield the same results.

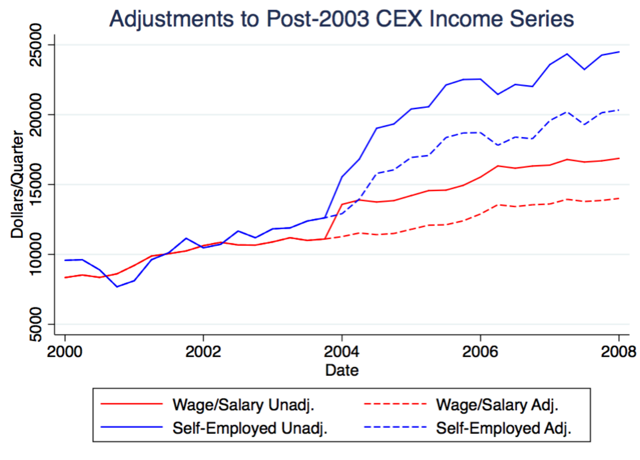

Figure C.5: This figure displays the adjusted- and unadjusted income series after the CEX's 2003/2004 change in income imputation procedure. Adjustment is taken by taking the time trend implied by the four quarters in 2003 and predicting the first quarter of 2004 with them, and proportionately multiplying all quarters after Q1 2004 by the ratio of Q1 2004's fitted value with its actual value.

C.2 Bureau of Labor Statistics

I use several Bureau of labor statistics series to construct my data. While data creation is discussed in the paper, I delineate the series here, and what they are:

- LNS12032187: Employment level: nonagricultural industries: wage/salary workers
- LNS12032192: Employment level: nonagricultural industries: unincorporated self-employed
- LNS12033251: Average hours at work in nonagricultural industries: wage and salary workers
- LNU02033377: Average hours at work in nonagricultural industries: self-employed workers
- LNU00000000: Civilian non-institutional population: 16 years and over
- LNU00000097: Civilian non-institutional population: 65 years and over