NBER WORKING PAPER SERIES

THE LABOR WEDGE: MRS VS. MPN

Loukas Karabarbounis

Working Paper 19015
http://www.nber.org/papers/w19015

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

The Labor Wedge: MRS vs. MPN
Loukas Karabarbounis
NBER Working Paper No. 19015
May 2013
JEL No. E2,E20,E24,E32

ABSTRACT

Do fluctuations of the labor wedge, defined as the gap between the firm's marginal product of labor (MPN) and the household's marginal rate of substitution (MRS), reflect fluctuations of the gap between the MPN and the real wage or fluctuations of the gap between the real wage and the MRS? For many countries, fluctuations of the labor wedge reflect predominantly fluctuations of the gap between the real wage and the MRS. As a result, business cycle theories of the labor wedge should primarily focus on improving the household side of the labor market. Explanations of the labor wedge based on departures of the representative firm's MPN from the real wage are rejected by the data because the labor share of income is not strongly procyclical.

Loukas Karabarbounis
University of Chicago
Booth School of Business
5807 S. Woodlawn Avenue
Chicago, IL 60637
and NBER
loukas.karabarbounis@chicagobooth.edu

1 Introduction

The neoclassical growth model predicts that, along any optimal path, the tax-adjusted marginal rate of substitution between leisure and consumption (MRS) equals the marginal product of labor (MPN). The labor wedge is defined as the gap between these two objects (Hall, 1997; Chari, Kehoe, and McGrattan, 2007; Shimer, 2009). Contrary to the prediction of the neoclassical growth model, the labor wedge, when measured under standard aggregate production function and utility function of a representative household, varies significantly over the business cycle and in a countercyclical way to output. Most of the times, variations of the labor wedge are interpreted as evidence that allocative inefficiencies increase during recessions. Less often, variations of the labor wedge are interpreted as specification errors in the neoclassical growth model. Whatever the interpretation, the intriguing business cycle variation of the labor wedge has recently spurred a significant amount of research.

There are two classes of models trying to explain the large business cycle variation of the labor wedge. Some papers achieve this by departing from the efficiency condition that the measured MPN of the representative firm equals the real wage. Examples include models with price markups, models with labor adjustment costs, models with financial frictions, models with firm heterogeneity that do not aggregate to a representative firm, and models with aggregate production functions different from Cobb-Douglas.[1] Other papers explain the large business cycle variation of the labor wedge by departing from the efficiency condition that the real wage equals the measured MRS of the representative household. Examples include models with wage markups, models with household heterogeneity that do not aggregate to a representative household, and models that allow for substitution possibilities toward home production.[2]

These two classes of explanations differ fundamentally. The first class requires a modification of the labor demand side of the neoclassical growth model (firm's problem). The second class requires a modification of the labor supply side of the neoclassical growth model

[1] See Rotemberg and Woodford (1999) for a discussion of price markups and alternative production functions. See Jermann and Quadrini (2012) and Arellano, Bai, and Kehoe (2012) for examples of recent research that generates labor wedges from firm-level financial frictions.

[2] See Cole and Ohanian (2004) for shocks to labor bargaining power that increase the gap between the real wage and the MRS. See Chang and Kim (2007) for a model with heterogeneity that generates variations of the labor wedge through the household side. See Karabarbounis (2012) for an explanation of the labor wedge based on home production.

(household's problem). Understanding whether fluctuations of the labor wedge manifest failures of the labor demand side of the model or failures of the labor supply side of the model is important for building successful models of the business cycle.

This paper makes progress to understanding this fundamental difference by decomposing the labor wedge into a gap between the measured MPN and the real wage and a gap between the real wage and the measured MRS. Remarkably, in the United States the gap between the measured MPN and the real wage explains by its own at most 3 percent of the movements of the labor wedge over the business cycle. By contrast, the gap between the real wage and the measured MRS explains by its own more than 70 percent of the movements of the labor wedge over the business cycle. Although this difference is more profound in the United States than in other countries, rarely does the gap between the measured MPN and the real wage account for more than 10 percent of the movements of the labor wedge. In addition, for most countries in the sample, the gap between the measured MPN and the real wage actually becomes smaller and not larger when output falls below its trend.

While these findings do not necessarily advance any particular explanation within the second class of explanations, they do narrow down the potential explanations of the labor wedge to those that operate through the household's MRS. Future research trying to generate movements of the labor wedge must do so by explaining why the measured MRS deviates so much from the real wage. Generating volatile and countercyclical labor wedges by increasing the discrepancy of the measured MPN from the real wage in recessions is grossly at odds with the data.

There is a simple logic behind these results. The labor wedge is volatile over the business cycle and countercyclical. Under Cobb-Douglas production function, the gap between the measured MPN and the real wage is a decreasing function of the labor share of income (real wages divided by the average product of labor). On average, real wages do not rise more than the average product of labor in booms and real wages do not fall more than the average product of labor in recessions. Equivalently, the labor share of income does not fluctuate in a procyclical way to output.[3] As a result, the firm's first-order condition that the measured MPN equals the real wage needs to be augmented by a relatively smooth and procyclical wedge in order to make this condition hold exactly in the data. If the firm's wedge is

[3] This is consistent with results reported elsewhere, see for example Rotemberg and Woodford (1999).

important over the business cycle, then we would observe a relatively smooth and procyclical labor wedge. Alternatively, theoretical models modifying the firm's side of the neoclassical growth model in order to generate volatile and countercyclical labor wedges, necessarily also generate the counterfactual prediction of a strongly procyclical labor share of income.

On the other hand, the measured MRS is much more volatile than real wages and is strongly procyclical to output. Recessions appear to be times when a representative worker's perceived value of time falls dramatically. As a result, the household's condition that the measured MRS equals the real wage needs to be augmented by a volatile and countercyclical wedge in order to make this condition hold exactly in the data. If the household's wedge is important over the business cycle, then we would observe a volatile and countercyclical labor wedge. Alternatively, theoretical models modifying the household's side of the neoclassical growth model in order to generate volatile and countercyclical labor wedges, necessarily generate in recessions a significant decline in the marginal value of time as perceived by a representative worker. This prediction can not be rejected by the data.

A strong implication of these findings is that researchers working on the labor wedge should develop models that better explain the household side of the labor market, while keeping the firm side close to the neoclassical growth model. This is because departures from the firm's side of the neoclassical growth model either generate a strongly procyclical labor share of income (when the departure is made consistent with fluctuations of the labor wedge), or miss the business cycle properties of the labor wedge (when the departure is made consistent with fluctuations of the labor share).

Importantly, the result that the gap between the MPN and the real wage is insignificant in explaining business cycle movements of the labor wedge is meant to guide positive aspects of business cycle modeling. The result does not necessarily have implications about the welfare consequences of economic fluctuations. For example, the insignificant fluctuation of the gap between the measured MPN and the real wage could simply mask offsetting forces that collectively have important welfare implications.

Two earlier papers have come close to reaching these conclusions. Cole and Ohanian (2002) measure deviations from the first-order conditions of the neoclassical growth model during and after the Great Depressions in the United States and the United Kingdom. They show that the great majority of the slow recovery from the recessions is accounted for by

the deviation in the household's first-order condition and not by the deviation in the firm's first-order condition. Consistently with these results, they propose explanations based on increased bargaining power of workers in the United States and a combination of workweek length declines and increases of unemployment benefits in the United Kingdom.

Galí, Gertler, and López-Salido (2007) also decompose the labor wedge into a firm-based gap and a household-based gap. They primarily interpret the gap between the measured MPN and the real wage as a price markup and the gap between the real wage and the measured MRS as a wage markup. The main difference relative to this paper is that they do not use the decomposition to advocate household-based explanations of the labor wedge. This is because, under the interpretation that the gaps reflect markups, the overall welfare consequences of business cycle fluctuations depend on the sum of the gaps (the labor wedge) and not on the individual components. Consistently with my conclusion that the distinction between the two gaps is informative for positive analyses and not necessarily for welfare, I adopt a more general interpretation of these gaps as reflecting either failures of the firm side or failures of the household side of the neoclassical growth model.

The result that we should be focusing on the household's MRS in explaining movements of the labor wedge is informative about the model environment and its deeper parameters and not necessarily about the sources of business cycle fluctuations. A view among business cycle theorists is that mostly labor demand shocks and less labor supply shocks drive the business cycle. This view is consistent with the results here as long as the model is augmented with a wedge that alters significantly the key condition that the measured MRS equals the real wage. As an example, a model driven by productivity shocks in which consumption and leisure are highly substitutable but preferences are stable over time is consistent with the findings here. Whereas, a model driven by productivity shocks in which input financing frictions drive a wedge between the measured MPN and the real wage is not consistent with the findings here.

2 Labor Wedge Decomposition: MRS vs. MPN

In the stochastic neoclassical growth model there is a representative household with preferences:

$$E_0 \sum_{t=0}^{\infty} \beta^t U(c_t, l_t), \tag{1}$$

where U is a utility function, c_t denotes consumption, l_t denotes leisure, and β is a discount factor. In every date the household has one unit of time to allocate between leisure and work, $l_t + n_t = 1$. The household supplies labor to the firm for a real wage w_t per unit of time. The household owns the capital stock and rents it to the firm at a real rental rate R_t. The household takes the path of wages and rental rates as given.

Denoting by q_t^c consumption taxes and by q_t^n labor income taxes, the household maximizes its utility subject to a sequence of budget constraints:

$$(1 + q_t^c)c_t + x_t = (1 - q_t^n)w_t n_t + R_t k_t + \Pi_t + T_t, \tag{2}$$

where Π_t denotes firm's profits and T_t denotes lump sum transfers from the government. The capital stock k_t accumulates according to the rule $k_{t+1} = (1-\delta)k_t + x_t$, where x_t is investment spending and δ is the depreciation rate of capital goods.

A perfectly competitive firm hires labor and capital to produce output y_t according to an aggregate production function $y_t = F(k_t, n_t) = c_t + x_t + g_t$. In the aggregate resource constraint, g_t denotes the amount of final goods purchased by the government. An exogenous state vector s_t drives the economy's fluctuations. For instance, the state s_t could include government spending shocks, tax shocks, and shocks to the production function.

Standard arguments imply that, along any equilibrium path of this economy, the firm chooses sequences of labor demand such that in every date the marginal product of labor (MPN) equals the real wage:

$$\text{MPN}_t := F_n(k_t, n_t) = w_t. \tag{3}$$

Similarly, the household chooses sequences of consumption and leisure such that in every date the tax-adjusted marginal rate of substitution (MRS) equals the real wage:

$$\text{MRS}_t := \left(\frac{1 + q_t^c}{1 - q_t^n}\right)\left(\frac{U_l(c_t, l_t)}{U_c(c_t, l_t)}\right) = w_t. \tag{4}$$

Therefore, along any equilibrium path, labor market clearing implies:

$$\text{MPN}_t = \text{MRS}_t. \tag{5}$$

It is important to note that the equalization of the tax-adjusted marginal rate of substitution to the marginal product of labor does not rest on many of the assumptions made here just for simplicity. For example, this result does not rest on whether capital markets are efficient or inefficient or on the specific shocks driving the economy's fluctuations.

Macroeconomists who have measured the two components of equation (5) usually find a significant and volatile wedge between the measured MPN and the measured MRS. Since condition (5) does not hold in data, the typical way to proceed is to augment one of the two efficiency conditions, (3) or (4), with a residual and then define the labor wedge as the resulting gap between the measured MPN and the measured MRS.

For the purposes of this paper, it is important to differentiate between a labor wedge due to failures of the firm's first-order condition (3) and a labor wedge due to failures of the household's first-order condition (4). Therefore, assume that:

$$\exp(-\tau_t^f)\text{MPN}_t = w_t, \tag{6}$$

$$\exp(\tau_t^h)\text{MRS}_t = w_t, \tag{7}$$

where τ_t^f denotes the component of the labor wedge due to the fact that the measured MPN may be higher than the real wage and τ_t^h denotes the component of the labor wedge due to the fact that the measured MRS may be lower than the real wage.

The labor wedge, defined as the gap between the measured MPN and the measured MRS, simply equals the sum of these two components:

$$\tau_t := \log(\text{MPN}_t) - \log(\text{MRS}_t) = \tau_t^f + \tau_t^h. \tag{8}$$

That is, the MPN deviates from the MRS (τ_t) either because the MPN deviates from the real wage (τ_t^f) or because the real wage deviates from the MRS (τ_t^h) or both.

As Chari, Kehoe, and McGrattan (2007) have shown, there are many underlying models that equivalently yield a gap τ_t between the MPN and the MRS. The idea here is to explore which of the two components, τ_t^f or τ_t^h, is mostly responsible for business cycle movements of the labor wedge τ_t. While there are various models that equivalently yield a condition similar to (6) for the firm, these models are fundamentally different than models that equivalently yield a condition similar to (7) for the household.

The difference is that the former class of models requires modifications of the firm side of the model whereas the latter class of models requires modifications of the household side of the model. As a result, if the firm component τ_t^f turns out to be mostly responsible for the cyclical movements of the labor wedge, then successful models of the business cycle should primarily focus on understanding why the measured MPN deviates so much from the real

wage. If the household component τ_t^h turns out to be mostly responsible for the cyclical movements of the labor wedge, then successful models of the business cycle should primarily focus on understanding why the measured MRS deviates so much from the real wage.

To measure the labor wedge, one needs to make specific assumptions about the production function and the utility function. I begin by following Chari, Kehoe, and McGrattan (2007) and assume that the production function is given by:

$$F(k_t, n_t) = A_t k_t^{\alpha_m} n_t^{1-\alpha_m}, \tag{9}$$

and the utility function is given by:

$$U(c_t, l_t) = \left(\frac{1}{1-\gamma}\right) \left[c_t^{1-\alpha_l}(1-n_t)^{\alpha_l}\right]^{1-\gamma}. \tag{10}$$

Note that certain forms of mis-specification of the production function or of the utility function will show up in τ_t^f and τ_t^h respectively. Based on these functional forms, we take:

$$\tau_t^f = \log(1-\alpha_m) - \log(s_t^n), \tag{11}$$

$$\tau_t^h = \log(1/\alpha_l - 1) + \log(s_t^n) + \log(1/n_t - 1) + \log(y_t/c_t) + \log((1-q_t^n)/(1+q_t^c)), \tag{12}$$

where $s_t^n = w_t n_t / y_t$ denotes the labor share of income.

The intuition developed in this paper can be easily summarized by inspecting equation (11). This equation shows an inverse relationship between the firm's component of the labor wedge τ_t^f and the labor share of income s_t^n (or the ratio of the average product of labor to real wages). If fluctuations of the labor wedge τ_t primarily reflect fluctuations of the firm's component τ_t^f, then generating a strongly countercyclical and volatile labor wedge necessarily implies strongly procyclical and volatile movements of the labor share of income. As it turns out, the labor share of income is not procyclical. Since the firm and the household component add up to the labor wedge, fluctuations of the labor wedge must primarily reflect fluctuations of the household's component τ_t^h.

3 Results

I collect quarterly and annual data for various OECD countries between 1970 and 2010. For some countries the analysis is restricted to a subset of these years due to missing observations.

The source of data is the OECD Quarterly and Annual National Accounts. Output is defined as constant-price GDP and consumption is defined as constant-price final consumption expenditure of households and non-profits.

Measurement of the overall labor wedge, τ_t, requires data on the labor input, consumption, and output, but not on the labor share of income (or wages). Here, data on the labor share of income are required in order to be able to differentiate between the two components of the labor wedge, τ_t^f and τ_t^h. The labor share of income is defined as total labor costs divided by GDP. The target variable for total labor costs in the OECD is compensation of employees compiled according to the SNA 93. For some countries this variable is not available. In this case the OECD uses, in order of preference, gross wages and salaries, labor cost indices, or average earnings to construct total labor costs. Below, I discuss robustness with respect to the measurement of the labor share of income.

I use data on population, employment, and hours of market work from Ohanian and Raffo (2012). The labor input is defined as the product of the number of employed and hours worked per employed person. Tax data are obtained from McDaniel (2007) who provides effective labor (social security and labor income) and consumption taxes. Given that the data come at annual frequency, I begin by assuming that effective tax rates are equal across all quarters of a given year. Below, I report results at annual frequency and results when omitting the tax terms from the measurement of the MRS.

The quarterly data are HP-filtered with smoothing parameter of 1600 and the annual data are HP-filtered with a smoothing parameter of 6.25. Denote by asterisks trend values of variables. Define the cyclical component of the labor wedge as $\widehat{\tau}_t = \tau_t - (\tau_t)^*$, the cyclical component of the firm component of the labor wedge as $\widehat{\tau}_t{}^f = \tau_t^f - \left(\tau_t^f\right)^*$, and the cyclical component of the household component of the labor wedge as $\widehat{\tau}_t{}^h = \tau_t^h - \left(\tau_t^h\right)^*$. By definition we have $\widehat{\tau}_t = \widehat{\tau}_t{}^f + \widehat{\tau}_t{}^h$. For any other variable X_t, define the cyclical component of the variable as $\widehat{X_t} = \log(X_t) - \log(X_t^*)$.

Figure 1 plots the cyclical component of the labor wedge $\widehat{\tau}_t$ together with the cyclical component of the firm's wedge $\widehat{\tau}_t{}^f$ and the cyclical component of the household's wedge $\widehat{\tau}_t{}^h$ in the United States. There is a tight association between the household component of the labor wedge and the overall labor wedge. During recessions, the labor wedge increases because the gap between the real wage and the MRS increases.

Figure 1: Labor Wedge and Components in the United States

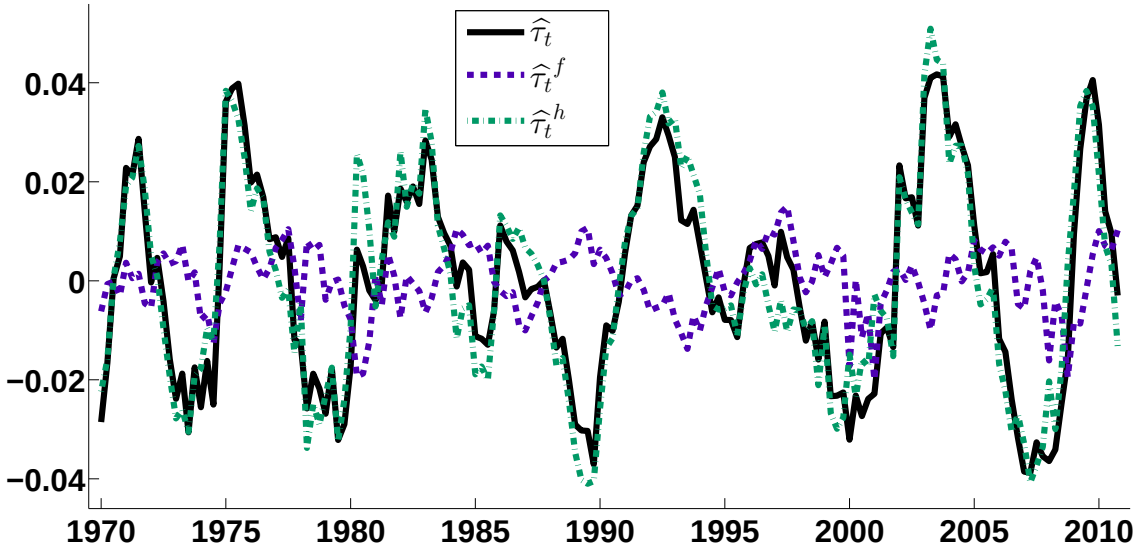

On the the other hand, the gap between the MPN and the real wage is disconnected from the labor wedge. This result has strong implications for theoretical models departing from the neoclassical model by modifying the firm's side of the model. These models can generate labor wedges as volatile and as countercyclical as in the data only by introducing a strongly countercyclical ratio of measured labor productivity to the real wage or, equivalently, a strongly procyclical labor share of income. However, in the data measured labor productivity relative to the real wage is not countercyclical and the labor share is not procyclical. In Figure 1, this shows up as a weak relationship between the firm's component of the labor wedge and the overall labor wedge.

Figures 2 to 5 show the two components of the labor wedge for the United Kingdom, France, Germany, and Canada respectively. Consistently with the results in the United States, in most cases the gap in the household's first-order condition drives the great majority of the business cycle movements of the labor wedge. Germany is a counterexample as in some instances the firm's component appears to contribute in a non-trivial way to the business cycle variation of the labor wedge.

To assess more systematically the contribution of each component to business cycle movements of the labor wedge, I calculate R-squared coefficients from the regressions:

$$\widehat{\tau}_t = \beta_0^f + \beta_1^f \widehat{\tau}_t^f + u_t^f, \tag{13}$$

Figure 2: Labor Wedge and Components in the United Kingdom

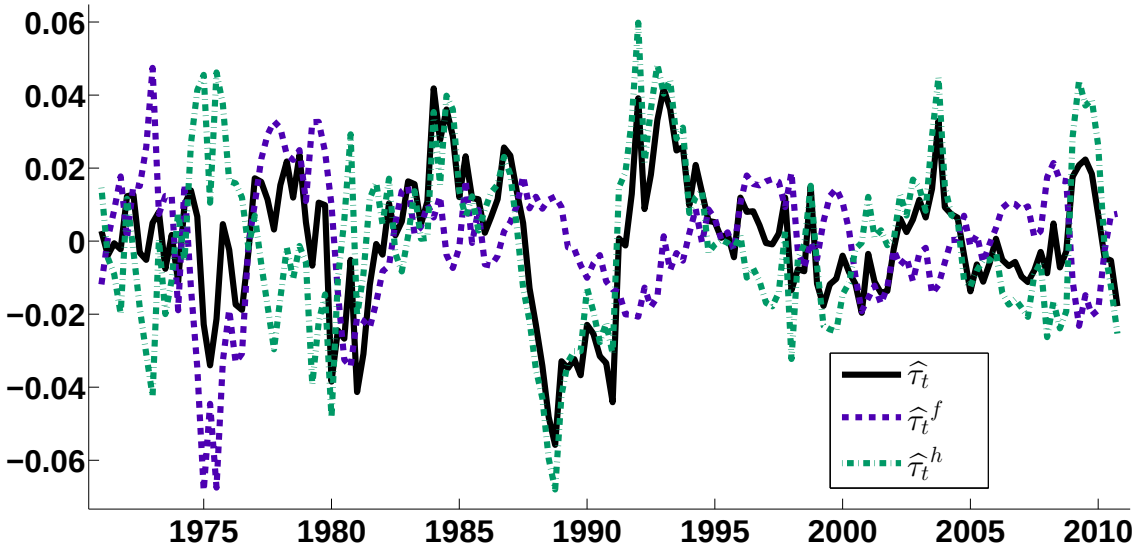

Figure 3: Labor Wedge and Components in France

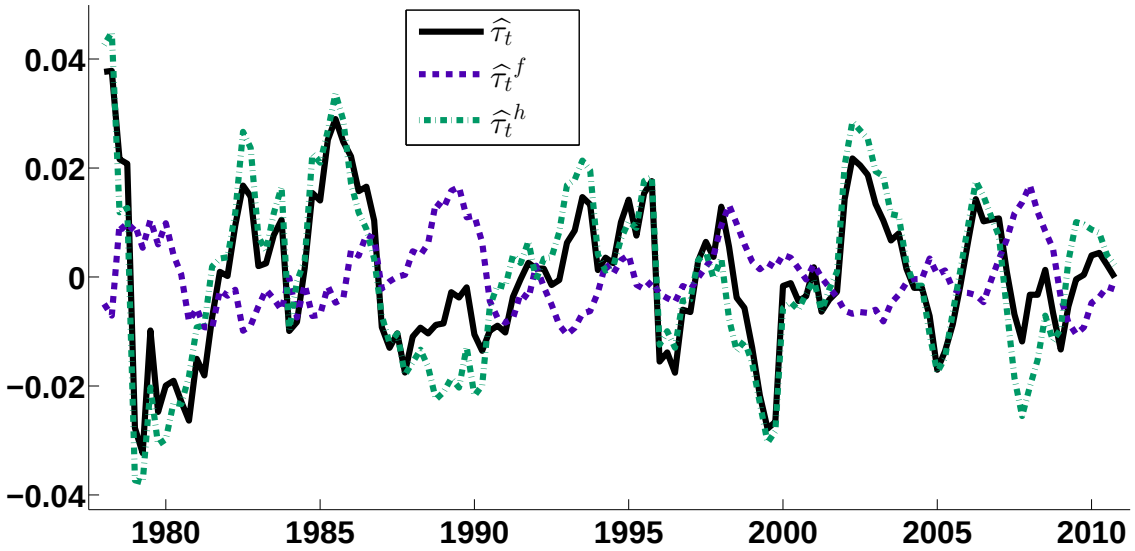

Figure 4: Labor Wedge and Components in Germany

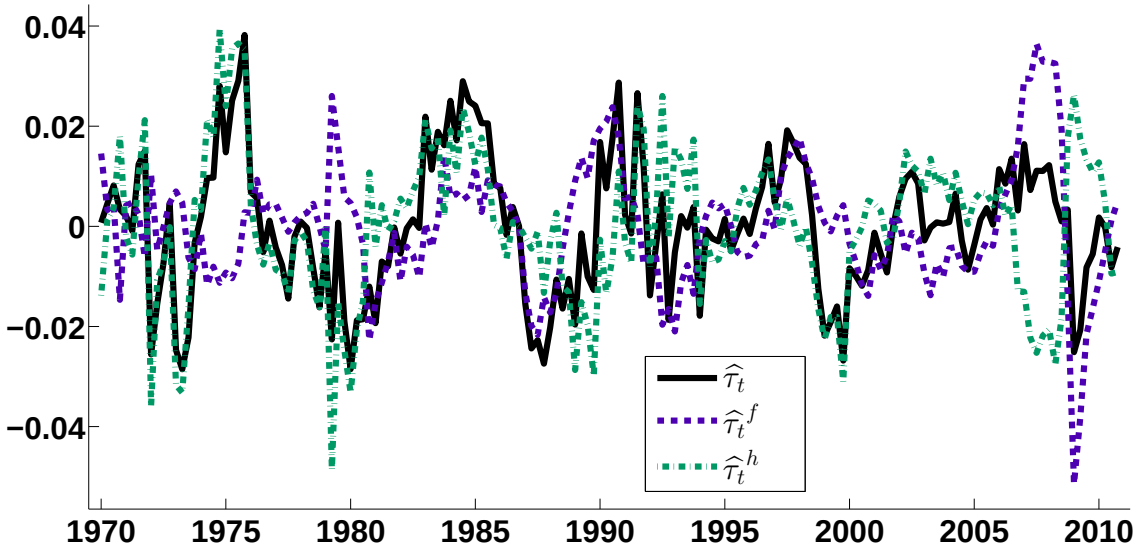

Figure 5: Labor Wedge and Components in Canada

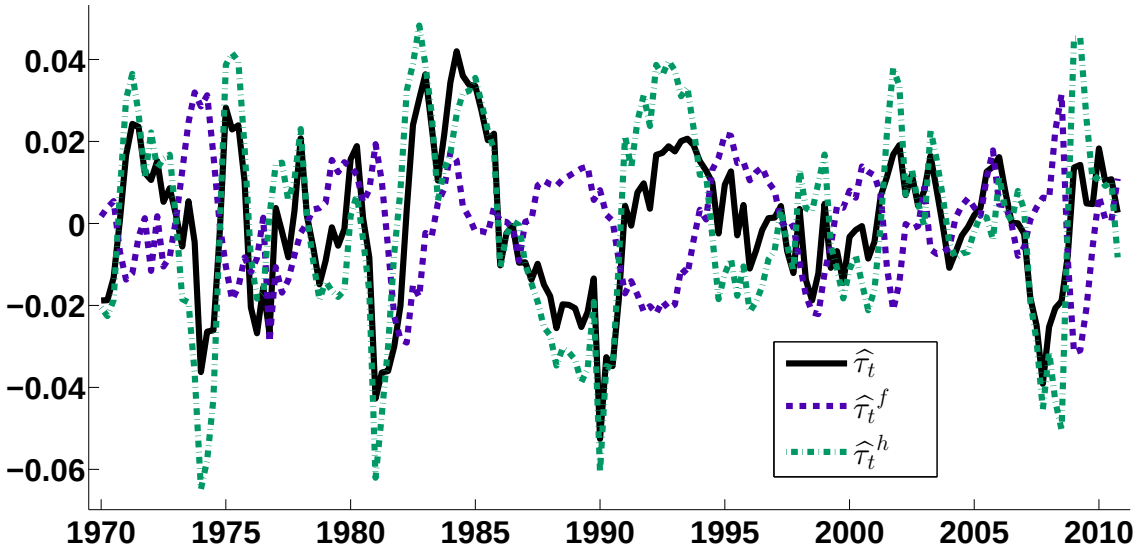

$$\widehat{\tau}_t = \beta_0^h + \beta_1^h \widehat{\tau}_t^h + u_t^h. \tag{14}$$

The sum of the R-squared coefficients need not equal 100 percent because the two components are not orthogonal. Still, the R-squared coefficient is useful because it captures the thought experiment of attributing fractions of the business cycle variation of the labor wedge to one component, assuming that the other component does not vary systematically.

The first column of Table 1 shows percent R-squared coefficients from these regressions for each country in the sample. Consistently with Figure 1, the cyclical component of the gap between the real wage and the MRS explains roughly 90 percent of the movements of the cyclical component of the labor wedge in the United States. The cyclical component of the gap between the MPN and the real wage explains less than 1 percent of the variance of the cyclical component of the labor wedge. This striking pattern holds in the majority of other countries in the sample. A notable exception is Germany in which the two components are strongly negatively correlated and, therefore, neither of them by their own can capture the majority of the business cycle variation of the labor wedge.

The other columns of Table 1 show some robustness checks. Column 2 repeats these regressions at the annual frequency (the same frequency at which we observe the effective tax rates). Column 3 repeats these regressions at quarterly frequency by omitting taxes. That is, the last term in equation (12) is always set equal to zero. With few exceptions only, the results are similar to the baseline results.

In columns 4 and 5, I consider an alternative utility function:

$$U(c_t, n_t) = \log c_t - \left(\frac{\chi}{1 + 1/\epsilon}\right) n_t^{1+1/\epsilon}, \tag{15}$$

where ϵ denotes the Frisch elasticity of labor supply. The measured MRS will now be a function of the elasticity ϵ, with higher values of the elasticity causing the measured MRS to be less volatile. In column 4, I set $\epsilon = 1$. The results remain similar to the baseline results. In column 5, I set $\epsilon = 5$, a value higher than the upper bound of the estimates that macroeconomists find reasonable for the Frisch elasticity of labor supply. Again, the results do not change significantly relative to the baseline results. This finding implies that, in the separable utility function case, mis-specifications of the utility function do not manifest as a volatile MRS due to a low assumed value of the Frisch elasticity of labor supply.

Finally, in column 6 of Table 1, I add government spending to consumption when measuring the labor wedge. That is household consumption is now defined as $c_t + g_t$, where g_t

Table 1: Fraction of Cyclical Variation of Labor Wedge Explained By Individual Components

		Baseline	**Annual**	**No Taxes**	$\epsilon = 1$	$\epsilon = 5$	**Government**
		(1)	(2)	(3)	(4)	(5)	(6)
Australia	$\widehat{\tau}_t^f$	1.10	0.02	0.33	0.91	0.65	2.23
	$\widehat{\tau}_t^h$	50.68	68.65	49.04	69.53	49.23	45.38
Canada	$\widehat{\tau}_t^f$	5.42	4.40	12.00	11.07	3.67	1.29
	$\widehat{\tau}_t^h$	72.80	73.75	74.88	84.95	66.62	62.73
Finland	$\widehat{\tau}_t^f$	0.84	0.63	3.66	0.28	1.41	4.30
	$\widehat{\tau}_t^h$	65.35	67.26	55.37	79.54	61.04	61.19
France	$\widehat{\tau}_t^f$	5.53	5.48	1.92	8.81	5.49	2.03
	$\widehat{\tau}_t^h$	84.72	86.34	80.62	91.70	82.62	81.71
Germany	$\widehat{\tau}_t^f$	11.38	16.71	6.99	3.34	16.70	21.23
	$\widehat{\tau}_t^h$	39.22	37.08	39.73	58.28	29.51	30.77
Italy	$\widehat{\tau}_t^f$	7.02	10.85	6.97	6.54	8.17	12.00
	$\widehat{\tau}_t^h$	73.44	78.27	61.25	81.94	70.50	71.70
Japan	$\widehat{\tau}_t^f$	1.09	5.58	1.57	0.06	1.54	4.09
	$\widehat{\tau}_t^h$	58.28	59.77	52.81	74.00	53.41	52.86
Norway	$\widehat{\tau}_t^f$	1.80	12.36	0.65	1.05	2.00	3.53
	$\widehat{\tau}_t^h$	52.93	22.53	48.05	70.06	50.59	47.52
Spain	$\widehat{\tau}_t^f$	18.48	22.45	8.31	9.74	16.00	15.36
	$\widehat{\tau}_t^h$	95.60	96.62	89.78	96.96	95.22	94.46
Sweden	$\widehat{\tau}_t^f$	1.61	3.78	7.73	0.10	1.79	2.80
	$\widehat{\tau}_t^h$	48.24	56.25	19.09	62.30	45.71	46.12
U.K.	$\widehat{\tau}_t^f$	3.07	6.17	0.00	0.11	3.77	11.83
	$\widehat{\tau}_t^h$	44.69	40.03	56.13	66.94	40.77	30.38
U.S.	$\widehat{\tau}_t^f$	0.18	0.11	0.00	0.27	0.91	0.82
	$\widehat{\tau}_t^h$	89.60	91.53	85.99	95.03	87.53	86.51

Notes: The table shows percent R-squared coefficients from regressions (13) and (14).

denotes constant-price government final expenditure. As the Table shows, the results remain roughly unchanged.

To summarize, for most countries in the sample and most forcefully for the United States, business cycle movements of the labor wedge, $\widehat{\tau}_t$, reflect primarily movements of the gap between the real wage and the MRS, $\widehat{\tau}_t{}^h$. The fraction of the variation of the labor wedge explained by only the gap between the MPN and the real wage, $\widehat{\tau}_t{}^f$, is always trivially small. In a minority of cases the covariation of the two components in shaping labor wedge movements dominates any individual component.

A justified concern with these results might be that the analysis rests heavily on measuring the labor share of income correctly. The primary source of data for the labor share is the OECD in order to increase cross-country coverage at quarterly frequency. However, the OECD data on the overall labor share may be less than ideal. First, the overall labor share of income includes income generated in the government sector. The government's optimization problem may be quite different from that of the representative firm in the neoclassical growth model as some goods and services produced in the government sector are not sold in the market. Second, the overall labor share includes income generated by sole proprietors. As a result, the measurement of the labor share is sensitive to how one allocates sole proprietor's income between labor income and capital income.

An alternative that is not subject to these measurement issues is to focus on the labor share within the corporate sector. In Figure 6, the gap between the measured MPN and the real wage is constructed using BEA data on the corporate labor income share. The corporate labor income share is defined as corporate compensations to employees (line 4 of NIPA Table 1.14) divided by the gross value added of the corporate sector (line 1 of NIPA Table 1.14). The correlation between the OECD-based measure of the gap $\widehat{\tau}_t{}^f$ using the overall labor share and the BEA-based measure of the gap $\widehat{\tau}_t{}^f$ using the corporate labor share is 0.75.

These differences do not affect the main conclusions. The firm's component remains disconnected from the labor wedge whether measured with the OECD data or with the BEA data. The R-squared of a regression of $\widehat{\tau}_t$ on $\widehat{\tau}_t{}^f$ was found to be less than 1 percent in the United States in Table 1. Using the corporate labor share from the BEA to calculate the gap between the MPN and the real wage increases the R-squared to roughly 2.9 percent. The R-squared of a regression of the cyclical component of $\widehat{\tau}_t$ on $\widehat{\tau}_t{}^h$ was found to be roughly 89

Figure 6: Gap Between MPN and Real Wage in United States and Labor Share Measurement

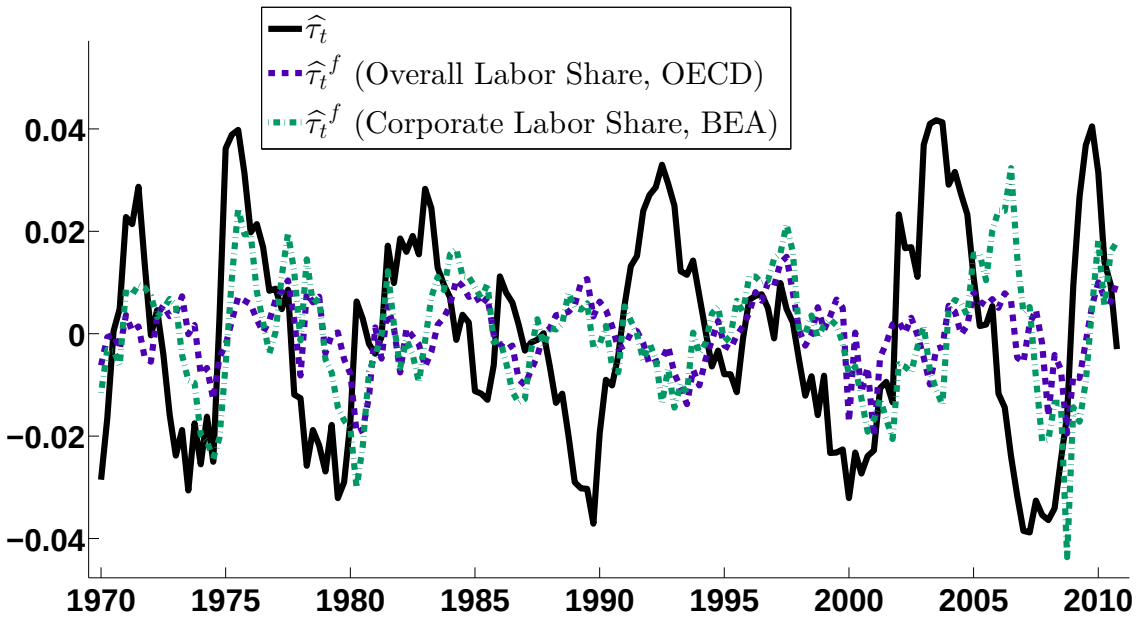

percent in the United States in Table 1. Using the corporate labor share from the BEA to calculate the gap between the real wage and the MRS decreases the R-squared to roughly 71 percent.

4 Discussion

There are three possible reactions to the finding that the gap between the MPN and the real wage explains a tiny fraction of the business cycle variation of the labor wedge. The first is to exclusively focus on models that explain the labor wedge through a gap between the real wage and the MRS. In this case, the firm's side of the model would be kept close to the neoclassical growth model.

A possible objection to this conclusion could be that, even if the gap between the MPN and the real wage does not explain any of the movements of the labor wedge over the business cycle, variations of the gap still contain useful information about overall macroeconomic fluctuations. That is, one could argue that the facts above present only a disconnect of the labor wedge from the the gap between the MPN and the real wage, but nothing in the facts excludes the possibility that the gap between the MPN and the real wage comoves in an interesting way with other macroeconomic aggregates.

Table 2 shows statistics for the cyclical components of the gap between the MPN and the real wage and the gap between the real wage and the MRS. For almost all countries, the gap between the MPN and the real wage does not move countercyclically to output, consumption, and labor. For instance, in the United States the correlation of the cyclical component of the gap $\widehat{\tau}_t^{\,f}$ with the cyclical component of output $\widehat{y}_t$ is roughly 0.20 in the OECD data.[4] As a result, in recessions the marginal value of labor as perceived by the representative firm does not increase relative to the real wage. Models predicting an increasing discrepancy between the MPN and the real wage in recessions are clearly at odds with the data.

On the other hand, the gap between the real wage and the MRS moves countercyclically to output, consumption, and labor. That is, in recessions the marginal value of time as perceived by the representative worker falls even more relative to the real wage. Models predicting an increasing discrepancy between the real wage and the measured value of time in recessions are consistent with this feature of the data.

These summary statistics make it difficult to imagine how variations in the gap between the MPN and the real wage can be an important element of the aggregate business cycle. Admittedly, there could be other macroeconomic variables that interact in an interesting way with the gap between the MPN and the real wage. An example of such a variable could be income inequality, to the extent that labor share variations over the business cycle are associated with significant changes in the distribution of income across households.

Another way to interpret the insignificance of the gap between the MPN and the real wage, is that the disconnect of the gap with the business cycle could be masking two or more interesting forces that happen to offset each other. To illustrate this possibility with an example, suppose that firms are imperfectly competitive and set prices as a markup μ_t over marginal costs. In addition, suppose that firms operate a CES production function with an elasticity of substitution between capital and labor equal to σ:

$$y_t = A_t \left(\alpha_m k_t^{\frac{\sigma-1}{\sigma}} + (1-\alpha_m) n_t^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}. \tag{16}$$

In this case a firm's first-order condition becomes:

$$(1-\alpha_m) A_t^{\frac{\sigma-1}{\sigma}} \left(\frac{y_t}{n_t} \right)^{\frac{1}{\sigma}} = \mu_t w_t. \tag{17}$$

[4] The correlation between the cyclical component of output and the cyclical component of the gap between the MPN and the real wage, constructed with the corporate labor share from the BEA, drops to 0.03. The correlation between $\widehat{\tau}_t^{\,h}$ and $\widehat{y}_t$ is -0.76 using the OECD-based measure of the overall labor share to measure $\widehat{\tau}_t^{\,h}$ and -0.70 using the BEA-based measure of the corporate labor share to measure $\widehat{\tau}_t^{\,h}$.

Table 2: Business Cycle Statistics of Individual Components

		$\mathbf{sd}(\widehat{\tau}_t)/\mathbf{sd}(\widehat{y}_t)$	$\mathbf{corr}(\widehat{\tau}_t, \widehat{y}_t)$	$\mathbf{corr}(\widehat{\tau}_t, \widehat{c}_t)$	$\mathbf{corr}(\widehat{\tau}_t, \widehat{n}_t)$
		(1)	(2)	(3)	(4)
Australia	$\widehat{\tau}_t^f$	1.26	-0.03	-0.01	-0.08
	$\widehat{\tau}_t^h$	1.78	-0.17	-0.41	-0.52
Canada	$\widehat{\tau}_t^f$	0.89	0.63	0.44	0.51
	$\widehat{\tau}_t^h$	1.66	-0.73	-0.68	-0.85
Finland	$\widehat{\tau}_t^f$	0.90	0.53	0.25	0.06
	$\widehat{\tau}_t^h$	1.53	-0.59	-0.52	-0.68
France	$\widehat{\tau}_t^f$	0.66	0.50	0.29	0.34
	$\widehat{\tau}_t^h$	1.65	-0.34	-0.31	-0.75
Germany	$\widehat{\tau}_t^f$	0.81	0.48	0.04	0.26
	$\widehat{\tau}_t^h$	0.98	-0.41	-0.48	-0.52
Italy	$\widehat{\tau}_t^f$	0.93	0.41	0.16	-0.12
	$\widehat{\tau}_t^h$	1.75	-0.30	-0.58	-0.71
Japan	$\widehat{\tau}_t^f$	0.68	0.63	0.51	0.27
	$\widehat{\tau}_t^h$	1.18	-0.42	-0.63	-0.61
Norway	$\widehat{\tau}_t^f$	2.40	0.32	0.14	-0.01
	$\widehat{\tau}_t^h$	3.48	-0.33	-0.58	-0.69
Spain	$\widehat{\tau}_t^f$	0.59	0.04	0.22	0.06
	$\widehat{\tau}_t^h$	2.56	-0.56	-0.66	-0.67
Sweden	$\widehat{\tau}_t^f$	0.98	0.55	0.30	0.23
	$\widehat{\tau}_t^h$	1.35	-0.38	-0.40	-0.37
U.K.	$\widehat{\tau}_t^f$	1.09	0.52	0.26	0.23
	$\widehat{\tau}_t^h$	1.43	-0.61	-0.61	-0.68
U.S.	$\widehat{\tau}_t^f$	0.41	0.19	0.13	0.02
	$\widehat{\tau}_t^h$	1.29	-0.75	-0.72	-0.88

Notes: The table shows business cycle statistics of the cyclical components of τ_t^f and τ_t^h.

Combining equation (17) with equation (6) which was used to define the firm's component of the labor wedge, we take:

$$\tau_t^f = \log(1-\alpha_m) - \log(s_t^n) = \log(\mu_t) + \left(\frac{1-\sigma}{\sigma}\right)\log(A_t y_t/n_t)\,. \tag{18}$$

This stylized example shows, for instance, that the relative smoothness of s_t^n and τ_t^f could reflect fluctuations in which negative productivity shocks are accompanied by simultaneous increases in price markups in an environment with capital-labor complementarity ($\sigma < 1$).

The general lesson is that the underlying forces of the gap between the MPN and the real wage sum up to explaining a tiny only fraction of the variability of the labor wedge over the business cycle. This finding is useful in guiding positive analyses of economic fluctuations. However, as the stylized example above shows, the gap between the measured MPN and the real wage could reflect opposing forces that happen to cancel out. If this is the case, the welfare consequences can differ dramatically depending on what these forces are exactly. As a result, the finding that the firm's wedge is unable to explain business cycle variations of the labor wedge is useful for positive analyses of economic fluctuations but it does not make strong predictions about the welfare consequences of these fluctuations.

5 Conclusions

Recent research has focused on fluctuations of the gap between the representative firm's measured marginal product of labor and the representative household's measured marginal rate of substitution as a useful moment against which to test theoretical models. Fluctuations of the labor wedge do not reflect fluctuations of the gap between the representative firm's measured marginal product and the real wage. As a result, models that generate volatile and countercyclical labor wedges by modifying the firm side of the neoclassical growth model are rejected by the data. The most promising explanations of the labor wedge should be able to generate large deviations between the real wage and the household's measured marginal rate of substitution.

References

Arellano, Cristina, Yan Bai, and Patrick Kehoe. 2012. "Financial Frictions and Fluctuations in Volatility." Federal Reserve Bank of Minneapolis Staff Report 466.

Chang, Yongsung, and Sun-Bin Kim. 2007. "Heterogeneity and Aggregation: Implications for Labor Market Fluctuations." *American Economic Review,* 97(5): 1939–1956.

Chari, Varadarajan, Patrick Kehoe, and Ellen McGrattan. 2007. "Business Cycle Accounting." *Econometrica,* 75(3): 781–836.

Cole, Harold, and Lee Ohanian. 2002. "The U.S. and U.K. Great Depressions through the Lens of Neoclassical Growth Theory." *American Economics Review,* 92(2): 28–32.

Cole, Harold, and Lee Ohanian. 2004. "New Deal Policies and the Persistence of the Great Depression: A General Equilibrium Analysis." *Journal of Political Economy,* 112(4): 779–816.

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

Hall, Robert. 1997. "Macroeconomic Fluctuations and the Allocation of Time." *Journal of Labor Economics,* 15(1): 223–250.

Jermann, Urban, and Vincenzo Quadrini. 2012. "Macroeconomic Effects of Financial Shocks." *American Economic Review,* 102(1): 238–271.

Karabarbounis, Loukas. 2012. "Home Production, Labor Wedges, and International Real Business Cycles." NBER Working Paper No. 18366.

McDaniel, Cara. 2007. "Average tax rates on consumption, investment, labor and capital in the OECD 1950-2003." Working Paper, Arizona State University.

Ohanian, Lee, and Andrea Raffo. 2012. "Aggregate Hours Worked in OECD Countries: New Measurement and Implications for Business Cycles." *Journal of Monetary Economics,* 59(1): 40–56.

Rotemberg, Julio, and Michael Woodford. 1999. "The Cyclical Behavior of Prices and Costs." in John B. Taylor, and Michael Woodford (ed.), *Handbook of Macroeconomics.*

Shimer, Robert. 2009. "Convergence in Macroeconomics: The Labor Wedge." *American Economic Journal: Macroeconomics,* 1(1): 267–297.