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

I thank Bob Hall, Casey Mulligan, and especially Gianluca Violante and an anonymous referee for very useful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research. All data and codes used in the analysis are available in my webpage at http://faculty.chicagobooth.edu/loukas.karabarbounis/research/index.html.

The Labor Wedge: MRS vs. MPN
Loukas Karabarbounis
NBER Working Paper No. 19015
May 2013, Revised July 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 and most forcefully for the United States, fluctuations of the labor wedge predominantly reflect 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; Mulligan, 2002; 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.

There are two classes of models trying to explain the intriguing 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 than Cobb-Douglas.[1] Other papers explain the 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, models with policies that distort household's incentive to work, 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 (household's problem). Understanding whether fluctuations of the labor wedge manifest failures of

[1] Rotemberg and Woodford (1999) discuss price markups and alternative production functions. Jermann and Quadrini (2012) and Arellano, Bai, and Kehoe (2012) are two examples of recent research that generates labor wedges from firm-level financial frictions. In Bigio and La'O (2013) the labor wedge is a weighted average of sectoral labor wedges that arise because of pledgeability constraints.

[2] Cole and Ohanian (2004) show how shocks to labor's bargaining power increase the gap between the real wage and the MRS. Chang and Kim (2007) present a model with heterogeneity that generates variations of the labor wedge through the household side. Mulligan (2012) demonstrates a relationship between the expansion of social safety nets and the increase of the household gap during the Great Recession. Hall (2009) accounts for the cyclical behavior of the labor wedge in a model that features complementarity between consumption and labor and Karabarbounis (2012) offers an explanation of the labor wedge based on home production.

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.

Building on influential work by Galí, Gertler, and López-Salido (2007), 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 (the firm component of the labor wedge) and a gap between the real wage and the measured MRS (the household component). Remarkably, in the United States, the gap between the measured MPN and the real wage explains on its own at most 2 percent of the variation of the labor wedge over the business cycle. By contrast, the gap between the real wage and the measured MRS explains on its own roughly 80 percent of the cyclical variation of the labor wedge.

The disconnect of the firm component from fluctuations of the labor wedge is a stylized fact characterizing 14 other OECD economies. Although the difference between the household component and the firm component is more profound in the United States than in other countries, the gap between the measured MPN and the real wage is quantitatively unimportant in explaining cyclical movements of the labor wedge in all countries in the sample. For almost all countries, the gap between the measured MPN and the real wage actually becomes smaller, and not larger, when output falls below its trend.

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 component of the labor wedge is important over the business cycle, then we would observe a relatively smooth and procyclical

[3] The lack of procyclicality in the labor share is consistent with the previous findings of Gomme and Greenwood (1995) and Rotemberg and Woodford (1999) in earlier samples. It is also found in more recent samples such as in Rios-Rull and Santaeulalia-Llopis (2010) and Nekarda and Ramey (2013).

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 that the labor share of income is strongly procyclical.[4]

On the other hand, the measured MRS is strongly positively correlated with 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 component of the labor 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 measured marginal value of time as perceived by a representative worker. This prediction can not be rejected by the data.

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. As a result, business cycle theories of the labor wedge must focus on improving the household side of the neoclassical growth model. The insignificance of the firm component in explaining business cycle movements of the labor wedge can be interpreted in two ways. The first is that successful models of the labor wedge should keep the firm side of their model close to the neoclassical growth model. Alternatively, the results are also consistent with the possibility that two or more deviations from the firm side of the model happen to offset each other. However, in this case deviations from the firm side of the neoclassical growth model should be introduced jointly in a way that still produces a relatively constant gap between the measured MPN and the real wage.

The dominance of the household relative to the firm component in explaining movements of the labor wedge does not have strong implications about the welfare effects of economic

[4] There is a large literature linking the labor share of income to price markups under Cobb-Douglas production function and no overhead labor. However, my results do not require such interpretation of the labor share movements. Bils (1987) and Rotemberg and Woodford (1999) discuss generalizations of the production function that decouple the price markup from the measured labor share in the data. Hall (2013) shows that cyclical movements in profit margins can be disconnected from cyclical movements in the labor share in the presence of other product-market wedges.

fluctuations. For instance, the welfare implications of the insignificance of the firm component depend on whether one adopts the view that there are no deviations from the firm side of the neoclassical growth model or the view that two or more deviations happen to offset each other. In the latter interpretation, the offsetting forces could collectively have important welfare implications even though they cannot explain business cycle movements of the labor wedge.

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) use the labor wedge to measure the welfare costs of fluctuations in the United States. They also decompose the labor wedge into a firm-based gap and a household-based gap and primarily interpret the former as a price markup and the latter as a wage markup. They do not use the decomposition to advocate household-based explanations of the labor wedge because, under the interpretation that the gaps reflect markups, the welfare effects of business cycles 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 analysis in this paper expands and improves upon the decomposition presented first in Galí, Gertler, and López-Salido (2007) along several dimensions. First, I carefully document the facts for the United States by using both establishment-level and household-level survey data for hours worked and by adjusting the labor wedge for time-varying tax wedges. Second, given difficulties in measuring the share of income accruing to labor, I demonstrate the robustness of my conclusions to four alternative measures of the labor share. Third, I consider both Cobb-

Douglas and separable preferences between consumption and labor as baseline preferences for measuring the labor wedge and show the robustness of my findings to the assumed preference specification. Fourth, I demonstrate that the insignificance of the firm component of the labor wedge does not reflect compositional changes over the business cycle, where economic activity happens to reallocate toward industries with high labor share levels during recessions. Finally, I document the pervasiveness of the dominance of the household relative to the firm component for fluctuations of the overall labor wedge by performing decompositions in a cross section of 15 OECD countries.

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 allocates time between leisure and work, $l_t + n_t = T$. 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 $y_t = F_t(k_t, n_t) = c_t + x_t + g_t$. In the aggregate resource constraint, g_t denotes the amount of final goods purchased by the government. An exogenous state vector z_t drives the economy's fluctuations. For instance, the state z_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:

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

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:

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

and the utility function is given by:

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

Note that certain forms of misspecification 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:

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

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

where $s_t = w_t n_t / y_t$ denotes the labor share of income, defined as the ratio of real wages to the average product of labor.

The intuition developed in this paper can be easily summarized by inspecting equation (11). This equation shows an inverse relationship between the firm component of the labor wedge τ_t^f and the labor share of income s_t. If fluctuations of the labor wedge τ_t primarily reflect fluctuations of the firm 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 strongly 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 component τ_t^h.

3 Measurement

In this section, I discuss the sources of data in the United States, the filtering of variables, and various measurement issues. Measurement of the overall labor wedge, τ_t, requires data on aggregate hours worked n_t and leisure l_t, consumption c_t, output y_t, and tax rates q_t^c and q_t^n. The decomposition of the labor wedge into its two components, τ_t^f and τ_t^h, additionally requires data on the labor share of income, s_t.[5]

I begin by introducing some notation to separate cyclical components from trends. Let asterisks denote 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 $\widehat{X_t} = \log(X_t) - \log(X_t^*)$. To extract cyclical components and trends, I use either the Hodrick-Prescott (HP) filter or the Baxter-King (BK) filter.[6]

[5] All terms in equations (11) and (12) are ratios of variables. Therefore, it is not necessary to distinguish between aggregate and per capita variables in order to measure the two components of the labor wedge.

[6] For the HP filter, I use a smoothing parameter of 1600 for quarterly data and a smoothing parameter of 6.25 for annual data. For the BK filter, I isolate fluctuations with a period of length ranging between 6 to 32 periods for quarterly data and between 2 and 8 periods for annual data.

Consumption c_t is defined as the sum of real expenditures on non-durables and real expenditures on services. Output y_t is defined as real GDP. Data for these variables in current prices are obtained from the Bureau of Economic Analysis (BEA) in NIPA Table 1.1.5. All current-price variables are deflated by their respective price deflator, available in NIPA Table 1.1.9, to obtain real variables.

Recent research shows that measured labor productivity's behavior during the Great Recession is sensitive to the choice of hours data (see, for instance, Hagedorn and Manovskii, 2011; Ramey, 2012). To examine the sensitivity of my results, I use four sources of data to construct aggregate hours worked n_t.[7] The first source of data is Cociuba, Prescott, and Ueberfeldt ("CPU", 2012). The authors use Employment and Earnings publications from the Current Population Survey (CPS) to estimate total hours for the United States and hours worked per person aged 16 to 64. The population adds military personnel worldwide to civilian non-institutional population. Accordingly, the authors estimate military hours and add them to the CPS hours to construct their total hours variable.

The second source of data for n_t is an index of aggregate weekly hours of production and non-supervisory employees provided by the Bureau of Labor Statistics (BLS; Series CES0500000034). The third source is an index of total hours worked in the business sector provided by the BLS, Major Sector Productivity and Costs Program (Series PRS84006033). I call these series "CES/PRO" and "CES/CPS" respectively. The index CES/CPS is based on the Current Establishment Survey (CES), which collects paid hours of non-supervisory workers in the business sector. These estimates are adjusted with the ratio of hours worked to hours paid. Finally, the index uses information from the CPS to include supervisory and non-production workers and to estimate hours for sectors other than the business sector.

For the analysis of countries other than the United States, I use hours data from Ohanian and Raffo ("OR", 2012). For the United States, their hours per employee data comes from the

[7] I define leisure as discretionary time less than aggregate hours worked, $l_t = T - n_t$. Using data from the American Time Use Survey between 2003 and 2010, Aguiar, Hurst, and Karabarbounis (2013) show that sleeping, eating, and personal care are the three most acyclical major time uses. These "non-discretionary" time uses account for 53.5 percent of time not spent working for the average person. In Ohanian and Raffo (2012) and Cociuba, Prescott, and Ueberfeldt (2012), work per person averages roughly 26 hours per week between 1964(1) and 2011(4). As a result, I set discretionary time available for work and leisure equal to 92 hours per week per person.

same sources used in the CES/CPS index. Ohanian and Raffo (2012) complement hours per employee data with data on total employment and civilian non-institutional population aged 16 to 64 from the BLS. Based on their data, I construct aggregate hours worked for the United States and hours worked per person. I benchmark the CES/PRO and CES/CPS indices to the Ohanian and Raffo (2012) series of hours per worker and total aggregate hours, since the two series rely mostly on CES data. It is convenient to normalize the two series in 1974(1), since this is the date when the Ohanian and Raffo (2012) series on hours per person displays the minimum deviation from the Cociuba, Prescott, and Ueberfeldt (2012) series on hours per person.

Figure 1 shows hours per person across the four data sources between 1964(1) and 2011(4). The figure shows that there are significant differences across the four datasets for the trends in hours worked per person, with the Cociuba, Prescott, and Ueberfeldt (2012) series displaying a higher trend increase between 1975 and 2000. Differences between the various series are expected since they measure hours for non-overlapping class of workers. Frazis and Stewart (2010) show how differences in coverage and concepts can explain level differences between CPS-based and CES-based measures of hours per employee. However, the authors conclude that it is difficult to explain the increasing divergence between CPS-based and CES-based estimates of hours per employee over time.

As demonstrated below, these differences in levels and trends of hours across datasets do not affect significantly my conclusions. Conclusions regarding the role of the two components of the labor wedge rely on the cyclicality of the labor input and the cyclical components of the various series are almost perfectly correlated. The upper panel of Table 1 shows pairwise correlations between the cyclical components of the four measures of hours per person at quarterly frequency using the HP filter. The lower panel displays these correlations using the BK filter. All correlations exceed 0.90.

Next, I turn to the measurement of the labor share of income s_t. Krueger (1999) discusses conceptual difficulties in measuring the labor share related to the treatment of proprietor's income. For this reason, I use four alternative measures of the labor share. The first measure simply takes compensation of employees (wages, salary accruals, and supplements such as employer's contributions) and divides it by national income (lines 1 and 2 of NIPA Table 1.12). I call this series "BEA Unadjusted" because it assumes that only income unambiguously accruing to labor is part of labor's compensation.

The second series for the labor share, "BEA Adjusted," treats compensation of labor as unambiguous labor income, proprietor's income and net taxes on production and imports as ambiguous income, and all other categories such as rental income, corporate profits, and business transfers as unambiguous capital income. "BEA Adjusted" differs from "BEA Unadjusted" in that it allocates a fraction of the ambiguous income to labor. Specifically, I allocate to labor income a fraction of the ambiguous income equal to the ratio of unambiguous labor to unambiguous labor and capital income. All data necessary to perform these calculations are from NIPA Table 1.12.

An alternative that does not require imputations of the labor earnings of sole proprietors is to focus on the labor share within the corporate sector. The third measure of the labor share, "BEA Corporate," 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). An additional benefit of using the corporate labor share is that the overall labor share of income includes income generated in the government sector. The government's optimization problem, however, 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. Finally, the fourth labor share variable, "BLS Corporate," comes from the BLS (Series PRS88003173) and is defined as the labor share in the non-financial corporate sector.

The upper panel of Table 2 shows pairwise correlations between the cyclical components of the four measures of the labor share at quarterly frequency using the HP filter. The lower panel displays these correlations using the BK filter. All correlations exceed 0.85.

Following Mendoza, Razin, and Tesar (1994), I use national income accounts to measure effective average tax rates on consumption q_t^c and labor q_t^n.[8] To measure consumption taxes, I

[8] Effective tax rates do not capture in a satisfactory way all policy-induced distortions in the labor market. For instance, Mulligan (2012, Figure 7) shows that between 2007(4) and 2009(2) there is a tight association of the increasing gap between the real wage and the MRS with the declining rate of self-reliance due to expansions in safety net benefits. If the expansion of benefits is deficit-financed or financed with corporate taxes, then there need not be any movement in the tax wedge that I consider. But even if labor, social security,

use data on net taxes on production and imports (lines 19 and 20 of NIPA Table 1.12). These taxes include items such as federal excise taxes, state sale taxes, and property taxes, and therefore affect both consumption spending and investment spending. I calculate consumption taxes as a fraction of net taxes on production and imports. The fraction equals the ratio of personal consumption expenditures to the sum of personal consumption expenditures and gross private domestic investment from NIPA Table 1.1.5. Then, q_t^c is defined as the ratio of consumption taxes divided by personal consumption expenditures less than consumption taxes.

I define the tax rate on labor income q_t^n as the sum of the tax rate on personal income and the tax rate on social insurance. The tax rate on personal income is defined as the ratio of personal current taxes from NIPA Table 3.1 (including items such as federal income taxes, state income taxes, state and local property taxes) divided by GDP less than taxes on production and imports. The assumption implicit in this calculation is that labor and capital income are taxed at the same rate. The tax rate on social insurance is defined as the ratio of contributions for government social insurance from NIPA Table 3.1 (which predominantly reflects contributions to OASDI) divided by the product of the labor share with GDP less than net taxes on production and imports.[9]

Figure 2 plots the resulting tax wedge, $(1 + q_t^c)/(1 - q_t^n)$. As the figure shows, there has been an increase in the tax wedge from the mid 1960s to 2000. This finding implies that taxes cannot explain the decline of the tax-unadjusted labor wedge between 1980 and 2000 found in the United States data. Because income tax rates have increased only modestly over time and consumption tax rates have actually decreased, the increase of the tax wedge mostly reflects the increase of the tax rate on social insurance from roughly 5 percent to more than 10 percent.

In the same figure, I compare my tax wedge variable to the variable provided by McDaniel (2011) at annual frequency. McDaniel (2011) uses OECD data to update the tax rate series

and consumption taxes are used to finance increased benefits, these benefits may distort labor supply choices beyond what it may be captured by effective tax rates. As a result, changes in benefits should be thought as a component of the tax-adjusted gap between the real wage and the MRS.

[9] Because I use four different sources to measure the labor share, this results in four different series for social insurance tax rates. Figures 2, 3, and 4 below plot the tax wedge, the labor wedge, and the cyclical component of the labor wedge under the “BEA Corporate” labor share of income. The results are similar when using the other labor share measures.

of Mendoza, Razin, and Tesar (1994). Both my series using BEA data and McDaniel's series for the United States using OECD data exhibit a roughly parallel increase until 2000 followed by a noticeable decline in the early 2000s. The level difference between the two series is partly accounted for by the level of social security tax rates, since the labor share assumed in McDaniel (2011) differs from the level of the labor share that I calculate. Importantly for my analysis of fluctuations below, the two tax wedge series comove strongly at higher frequencies. Later, I use McDaniel's series on the tax wedge to decompose the labor wedge into its two components for countries other than the United States.

4 The Labor Wedge in the United States

Figure 3 depicts the labor wedge τ_t in the United States for each of the four different hours series between 1964(1) and 2011(4). The labor wedge is normalized to 0.30 in 2000(1). The trend decline in the labor wedge between 1980 and 2000 is mostly attributed to the increase in hours worked since the consumption to output ratio did not display any significant trend during this period.[10] Figure 4 shows the cyclical component of the labor wedge, $\widehat{\tau}_t$, for the different hours series and the cyclical component of real GDP per person. From this figure it is apparent that the labor wedge is volatile and strongly countercyclical to output.

4.1 Labor Wedge Decomposition

Figure 5 presents the cyclical components of the labor wedge, $\widehat{\tau}_t$, of the firm component of the labor wedge, $\widehat{\tau}_t^f$, and of household component of the labor wedge, $\widehat{\tau}_t^h$, between 1964(1) and 2011(4). This figure uses the "CES/PRO" series for hours and the "BLS Corporate" labor share series to construct the two components. Under this particular combination of hours and labor share series, the contributions of each component to the cyclical variation of the labor wedge are representative of the average contributions documented below across all different combinations of hours and labor share series.

[10] The figure depicts the tax-adjusted labor wedge. The decline in the tax-adjusted labor wedge between 1980 and 2000 is larger than the decline in the tax-unadjusted labor wedge because the tax wedge term increased over this period. Most of the variation of the labor wedge, however, at this low frequency is accounted for by changes in hours per person.

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. On the other hand, fluctuations of the gap between the MPN and the real wage are disconnected from fluctuations of 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 strongly countercyclical and the labor share is not strongly procyclical. In Figure 5, this shows up as a weak relationship between the firm component of the labor wedge and the overall labor wedge at business cycle frequencies.

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:

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

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

The regressions are performed at quarterly data covering the period of 1964(1) to 2011(4). 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.

Table 3 shows percent R-squared coefficients from these regressions when I use the HP filter to detrend variables. The upper panel of the table shows percent R-squared coefficients from the regression of the labor wedge on the firm component and the lower panel shows percent R-squared coefficients from the regression of the labor wedge on the household component. Each row corresponds to a different hours series and each column corresponds to a different labor share series.

Consistently with the message conveyed in Figure 5, the cyclical component of the gap between the real wage and the MRS explains on average roughly 80 percent of the variation of the cyclical component of the labor wedge. Averaged across different hours and labor share measures, the cyclical component of the gap between the MPN and the real wage explains no more than 2 percent of the variance of the cyclical component of the labor wedge. This striking pattern also holds when using the BK filter to detrend variables. Table 4 shows these results.

4.2 Labor Wedge Components: Comovement and Volatility

A possible objection to these conclusions 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 firm component 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 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 5 shows statistics for the cyclical component of the gap between the MPN and the real wage between 1964(1) and 2011(4). Across various definitions of the labor share and detrending methods, the gap between the MPN and the real wage moves either procyclically or acyclically relative to real output per person, real consumption on non-durables and services per person, and work hours per person. This reflects the fact that the labor share of income is in general mildly countercyclical.[11] Put it differently, in recessions the measured 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 measured MPN and the real wage in recessions are clearly at odds with the data.

Table 6 presents statistics for the gap between the real wage and the MRS between 1964(1) and 2011(4). The household component of the labor wedge moves countercyclically to output, consumption, and labor. That is, in recessions the measured marginal value of time as perceived by the representative worker falls even more relative to the real wage. In addition, the

[11] This result is consistent with the results reported in Rios-Rull and Santaeulalia-Llopis (2010). The authors additionally stress the overshooting property of the labor share, in which the labor share initially falls below trend and then rises above trend in response to a positive innovation in output and TFP. The R-squared coefficients presented in Tables 3 and 4 imply that the overall dynamics of the labor share (including the overshooting property) are disconnected from the cyclical variation of the labor wedge.

household component of the labor wedge is more than twice as volatile as the firm component. Models predicting an increasing discrepancy between the real wage and the measured marginal value of time in recessions are consistent with this feature of the data.

These summary statistics make clear that understanding the business cycle properties of the labor wedge requires modifications of the household side of the neoclassical growth model. At the same time, it is difficult to imagine how variations of the firm component can be an important element of the 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. At minimum, however, the facts documented in this paper present a challenge for models that rely on deviations from the firm's first-order condition (as opposed to the household's) to generate links between the labor wedge and basic macroeconomic aggregates such as output, consumption, and labor.

4.3 Alternative Preferences

This section considers the sensitivity of my conclusions to the utility function used to generate the labor wedge. To calculate the labor wedge, some authors in the literature (e.g. Shimer, 2009) consider separable preferences of the form:

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

where ϵ denotes the Frisch elasticity of labor supply. Under these preferences the household component of the labor wedge becomes:

$$\tau_t^h = -\log(\chi) + \log(s_t) - \left(1 + \frac{1}{\epsilon}\right) \log(n_t) + \log(y_t/c_t) - \log\left((1+q_t^c)/(1-q_t^n)\right), \tag{16}$$

where n_t denotes hours worked per person. The household component of the labor wedge in equation (12) under Cobb-Douglas preferences differs from the household component of the labor wedge in equation (16) under separable preferences only up to terms that involve hours n_t. Since the firm component of the labor wedge is not affected by preferences, the overall labor wedge similarly differs between Cobb-Douglas preferences and separable preferences only up to terms that involve hours n_t.

To understand the difference between Cobb-Douglas and separable preferences for the measurement of the labor wedge, I replace terms involving hours in equations (12) and (16) with their first-order approximations around the mean level of hours $\bar{n}$. In each of the expressions below, I omit a constant that can be normalized to equalize the labor wedges in some arbitrary period. For the Cobb-Douglas utility function in equation (10) we take:

$$\tau_t^h \approx \log\left(s_t\right) + \log(y_t/c_t) - \log\left((1+q_t^c)/(1-q_t^n)\right) - \left(\frac{1}{1-\bar{n}/T}\right)\left(\frac{n_t-\bar{n}}{\bar{n}}\right), \tag{17}$$

while for the separable utility function in equation (15) we take:

$$\tau_t^h \approx \log\left(s_t\right) + \log(y_t/c_t) - \log\left((1+q_t^c)/(1-q_t^n)\right) - \left(1+\frac{1}{\epsilon}\right)\left(\frac{n_t-\bar{n}}{\bar{n}}\right). \tag{18}$$

Equations (17) and (18) imply that the household and the overall labor wedge are approximately equal under the two preference specifications for all periods t when:

$$\frac{1}{1-\bar{n}/T} = 1+\frac{1}{\epsilon}. \tag{19}$$

Solving this equation for the Frisch elasticity of labor supply ϵ, one can calculate the value of the elasticity that makes the two utility functions approximately equivalent for the measurement of the labor wedge. For example, in the baseline results presented above, labor over discretionary time $\bar{n}/T$ equals 0.28, so the two specifications are equivalent for the measurement of the labor wedge when the Frisch elasticity of labor supply is approximately equal to 2.5.

The higher is the Frisch elasticity of labor supply ϵ in the separable utility function or the lower is the average labor to discretionary time ratio $\bar{n}/T$ in the Cobb-Douglas utility function, the lower is the responsiveness of the measured labor wedge to a given change in hours. When the responsiveness of the measured labor wedge to changes in hours is low, both the measured labor wedge and the household component of the labor wedge become less volatile over the business cycle.

To examine the sensitivity of my results to the assumed preferences, Figure 6 shows percent R-squared coefficients from regressions (13) and (14) under various values of the Frisch elasticity ϵ. This figure uses the "CES/PRO" series for hours and the "BLS Corporate" labor share series because this particular combination of hours and labor share series produces representative contributions of the two components to the cyclical variation of the labor wedge. As

Figure 6 shows, increasing the Frisch elasticity implies a smaller contribution of the household component to the business cycle variation of the labor wedge. However, even for $\epsilon = 5$, a value higher than the upper bound of the estimates that macroeconomists find reasonable for the Frisch elasticity of labor supply, the household component continues to explain more than 75 percent of the cyclical variation of the labor wedge. Similarly, even for extreme values of ϵ, the firm component continues to explain at most 2 percent of variation of the labor wedge at business cycle frequencies.[12]

4.4 Labor Shares Across Industries

The insignificance of the firm component for fluctuations of the overall labor wedge reflects the lack of strong procyclicality in the labor share of income. However, the lack of strong procyclicality in the movements of the aggregate labor share could simply reflect compositional shifts over the business cycle. Imagine, for instance, that the labor share is indeed procyclical in most industries but economic activity happens to reallocate towards industries with higher labor shares during recessions. In that case, the aggregate labor share could appear to be acyclical or even countercyclical despite a significant procyclicality in industry-level labor shares.

To address this issue, I use industry-level data on labor shares from two sources. First, I use BEA's GDP-by-industry accounts to construct labor shares at the industry level, called "BEA Unadjusted" and "BEA Adjusted." The unadjusted labor share is defined as compensation of employees divided by industry value added. The adjusted labor share is defined as compensation of employees plus a fraction of net taxes on production and imports divided by industry value added. I define 16 non-overlapping industries, with the value added of all industries adding up to aggregate value added. The data come at annual frequency between 1987 and 2011.

The second source of data on industry-level labor shares is the EU KLEMS dataset, for years between 1977 and 2007. Here I define 10 non-overlapping industries whose value added sum up to aggregate value added. The labor share "KLEMS Unadjusted" is defined as com-

[12] When the Frisch elasticity of labor supply becomes infinite, the household component explains 68.4 percent and the firm component explains 2.3 percent of the cyclical variation of the labor wedge.

pensation of employees divided by industry value added. The labor share "KLEMS Adjusted" is defined as labor compensation divided by industry value added. Labor compensation equals compensation of employees plus a fraction of other taxes on production and an imputation for the income of self-employed using the category operating surplus / mixed income.

The aggregate labor share is a weighted average of industry-level labor shares:

$$s_t = \sum_{j=1}^{J} \omega_{j,t} s_{j,t}, \tag{20}$$

where $s_{j,t}$ denotes the labor share of industry j and $\omega_{j,t} = VA_{j,t}/VA_t$ denotes the nominal value added share of industry j in total value added in period t. The change of the labor share between period $t-1$ and period t can be decomposed as:

$$\Delta s_t = \underbrace{\sum_{j=1}^{J} \left(\frac{\omega_{j,t}}{2} + \frac{\omega_{j,t-1}}{2}\right)(s_{j,t} - s_{j,t-1})}_{Within-Industry} + \underbrace{\sum_{j=1}^{J} \left(\frac{s_{j,t}}{2} + \frac{s_{j,t-1}}{2}\right)(\omega_{j,t} - \omega_{j,t-1})}_{Between-Industry}. \tag{21}$$

Changes in the within component of the labor share reflect changes in the aggregate labor share due to changes in industry-level labor shares, holding constant the industrial composition of economic activity. Changes in the between component of the labor share reflect changes in the aggregate labor share due to changes in economic activity across industries, holding constant the labor share of each industry.

Denote by $t = 0$ the first observation in the sample. Based on the decomposition of the change in the labor share into the two components, for any $t \geq 1$, I define the "within" and the "between" components of the labor share as:

$$s_t^w = \frac{s_0}{2} + \sum_{t=1}^{T} \left(\sum_{j=1}^{J} \left(\frac{\omega_{j,t}}{2} + \frac{\omega_{j,t-1}}{2}\right)(s_{j,t} - s_{j,t-1}) \right), \tag{22}$$

$$s_t^b = \frac{s_0}{2} + \sum_{t=1}^{T} \left(\sum_{j=1}^{J} \left(\frac{s_{j,t}}{2} + \frac{s_{j,t-1}}{2}\right)(\omega_{j,t} - \omega_{j,t-1}) \right). \tag{23}$$

By construction we have $s_t = s_t^w + s_t^b$.

Figures 7 and 8 show percentage point deviations of the labor share from its trend and decompositions of these deviations into a within and a between component in the BEA and the KLEMS data respectively. The figures show that virtually all of the cyclical fluctuation of the aggregate labor share is accounted for by fluctuations of industry-level labor shares. Cyclical fluctuations of the aggregate labor share cannot be explained by changes in industrial composition over the business cycle. In these figure the definition of the labor share is the unadjusted one and the trend is computed with the HP filter, but similar results are obtained using the adjusted labor share series and when detrending the data with the BK filter.

Table 7 shows statistics of the cyclical component of the labor share at the industry level. The majority of industries have countercyclical labor shares, with finance and business services being leading examples of sectors with procyclical labor shares. In addition, most industry-level labor shares comove strongly with the aggregate labor share. These results hold broadly across all definitions of the labor share, detrending methods, and sources of data.[13]

To summarize, fluctuations of the aggregate labor share reflect fluctuations of industry-level labor shares rather than a reallocation of economic activity towards industries with different labor share levels. Further, the cyclical patterns documented for the aggregate labor share hold broadly at the industry level. Therefore, the insignificance of the firm component for fluctuations of the labor wedge does not reflect compositional changes across industries over the business cycle.

5 Labor Wedges Across Countries

This section demonstrates that the dominance of the household component and the insignificance of the firm component for fluctuations of the labor wedge is a pervasive stylized fact across OECD countries. The data necessary to measure the two components of the labor wedge in a broader sample of countries are mostly available at annual frequency. To maximize country coverage I conduct the analysis at annual frequency. For most countries the results reported below cover the period of 1970 to 2010.

Data on population, employment, and hours of work per employee are taken from Ohanian and Raffo (2012). The authors have compiled the most comprehensive dataset currently available for total hours of work at the international level. The dataset covers 15 OECD countries

[13] All these results continue to hold when using an alternative disaggregation available in the KLEMS data with 23 industries. Results at this more disaggregated level can be obtained using the codes and data made available online.

mostly between 1970 and 2010. To construct the dataset, Ohanian and Raffo (2012) draw on a variety of sources and adjust national estimates to take into account cross-country differences in sick leave policies and holidays. As discussed before, effective labor (social security and labor income) and consumption taxes for various countries are obtained from McDaniel (2011). For two countries (Ireland and Korea) with insufficient tax data to perform the calculations, I assume that the tax wedge does not change over time.

Using data from OECD Annual National Accounts, output y_t is defined as constant-price GDP and consumption c_t is defined as constant-price consumption expenditure of households and non-profits on non-durables and services. Since the OECD does not provide data on durable expenditures for all countries, I impute spending on durables for countries with missing data using the cyclical properties of durables in countries with the required data. Then, I subtract spending on durables from total final consumption expenditure to obtain consumption of non-durables and services.[14]

I use four series to measure the labor share of income for the various countries. "OECD Unadjusted" is defined as the ratio of current-price total labor costs divided by current-price GDP (from "OECD Benchmarked Unit Costs: Total Costs and Real Output"). 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. The labor share "OECD Adjusted" adjusts for the self-employed and is available at the "OECD Unit Labor Costs - Annual Indicators" database. Finally, from the EU KLEMS dataset, I use the two labor share measures defined above, "KLEMS Unadjusted" and "KLEMS Adjusted," for the total economy.

[14] Let d_t denote spending on durables, c_t denote spending on non-durables and services, and $c_t + d_t$ denote total final consumption spending of households and non-profits. For each of the 8 countries with at least 25 years of data, I regress the change in the ratio of spending on durables to total consumption expenditure, $\Delta(d_t/(d_t + c_t))$, on a constant and the growth of real GDP per person g_t. I impute the ratio of spending on durables to total consumption spending in the remaining 7 countries using the formula $d_t/(d_t + c_t) = 0.084 + 0.110(g_t - 0.023)$, where 0.084 denotes the mean level of durables to total consumption spending in the 8 countries with the required data, 0.110 denotes the average slope coefficient of the regression in the 8 countries, and 0.023 is the average growth rate of GDP per capita in these 8 economies. Using instead the corresponding numbers in the United States (0.067, 0.095, and 0.015) does not change significantly the results. With few exceptions, the results do not change significantly when I use total consumption spending instead of consumption on non-durables and services to proxy for c_t.

Table 8 presents percent R-squared coefficients from regressions (13) and (14) when detrending is performed with the HP filter. Table 9 shows percent R-squared coefficients when detrending in performed with the BK filter. In each table, countries are presented in rows and labor share series are presented in columns. The qualitative patterns documented before for the United States continue to hold in the great majority of countries, with the household component almost always explaining a significantly larger fraction of the cyclical variance of the labor wedge than the firm component.

In 11 out of the 15 countries, including 6 out of the 8 largest OECD economies, the household component explains at least two-thirds of the cyclical variation of the labor wedge. Across the two tables, the median contribution of the household component to cyclical variations of the labor wedge is roughly 70 percent. This number is smaller than the contribution of the household component documented for the United States before (80 percent), but the household component is still the dominant component for the median country. The firm component is never quantitatively important on its own for understanding business cycle movements of the labor wedge. In the remaining 4 countries, including Germany and the United Kingdom, the covariation of the household component with the firm component becomes important in understanding cyclical movements of the labor wedge.

The insignificance of the firm component for fluctuations of the labor wedge can be also understood by examining its contemporaneous correlation with real GDP per capita. Table 10 shows this correlation for the 15 countries (rows) across different labor share series and detrending methods (columns). As the table shows, almost all firm components are procyclical or, equivalently, most labor shares are countercyclical. Consistently with my conclusion for the United States, this finding makes it difficult to imagine how labor wedges induced by departures from the firm's side of the neoclassical model can be an important element of the business cycle.

6 A Caveat

The main finding of this paper is that the gap between the real wage and the MRS explains the majority of the cyclical variation of the labor wedge. By contrast, the gap between the

MPN and the real wage explains a minor fraction of the cyclical variation of the labor wedge. A possible reaction to this finding is to exclusively focus on models that modify the household side of the neoclassical growth model in order to generate fluctuations of the labor wedge. In this case, the firm side of the model would be kept close to the neoclassical growth model.

An alternative way to interpret the insignificance of the gap between the MPN and the real wage, is that the disconnect of the gap from 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 production function with a constant elasticity of substitution between capital and labor equal to σ:

$$y_t = A_t \left(\alpha_m k_t^{\frac{\sigma-1}{\sigma}} + (1-\alpha_m) n_t^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}. \tag{24}$$

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

$$(1-\alpha_m) A_t^{\frac{\sigma-1}{\sigma}} \left(\frac{y_t}{n_t} \right)^{\frac{1}{\sigma}} = \mu_t w_t. \tag{25}$$

Combining equation (25) with equation (6), which was used to define the firm component of the labor wedge, we take:

$$\tau_t^f = \log(1-\alpha_m) - \log(s_t) = \log(\mu_t) + \left(\frac{1-\sigma}{\sigma} \right) \log(A_t y_t / n_t). \tag{26}$$

This stylized example shows, for instance, that the relative smoothness of s_t 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 message of my analysis is that the underlying forces of the gap between the MPN and the real wage sum up to explaining an insignificant fraction of the variation of the labor wedge over the business cycle. 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. This finding is useful because it implies that models with mechanisms explaining the cyclical behavior of the labor wedge by modifying the firm's problem also leave out of the analysis some other force that necessarily offsets the proposed mechanism. However, 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 component 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.

7 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

Aguiar, Mark, Erik Hurst, and Loukas Karabarbounis. 2013. "Time Use During the Great Recession." *American Economic Review,* 103(5).

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

Bigio, Saki, and Jennifer La'O. 2013. "Financial Frictions in Production Networks." University of Chicago Booth Working Paper.

Bils, Mark. 1987. "The Cyclical Behavior of Marginal Cost and Price." *American Economic Review,* 77(5): 838–855.

Chang, Yongsung, and Sun-Bin Kim. 2007. "Heterogeneity and Aggregation: Implications for Labor Market Fluctuations." *American Economic Review,* 97(5): 1939–56.

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

Cociuba, Simona, Edward Prescott, and Alexander Ueberfeldt. 2012. "U.S. Hours and Productivity Behavior Using CPS Hours Worked Data: 1947III-2011IV." Arizona State University Working Paper.

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.

Frazis, Harley, and Jay Stewart. 2010. "Why Do BLS Hours Series Tell Different Stories About Trends in Hours Worked?" BLS Working Paper No. 433.

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.

Gomme, Paul, and Jeremy Greenwood. 1995. "On the Cyclical Allocation of Risk." *Journal of Economic Dynamics and Control,* 19: 91–124.

Hagedorn, Marcus, and Iourii Manovskii. 2011. "Productivity and the Labor Market: Comovement over the Business Cycle." *International Economic Review,* 52(3): 603–19.

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

Hall, Robert. 2009. "Reconciling Cyclical Movements in the Marginal Value of Time and the Marginal Product of Labor." *Journal of Political Economy,* 117(2): 281–323.

Hall, Robert. 2013. "What the Cyclical Response of Advertising Reveals about Markups and Other Macroeconomic Wedges." NBER Working Paper No. 18370.

Jermann, Urban, and Vincenzo Quadrini. 2012. "Macroeconomic Effects of Financial Shocks." *American Economic Review,* 102(1): 238–71.

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

Krueger, Alan. 1999. "Measuring Labor's Share." *American Economic Review,* 89(2): 45-51.

Mendoza, Enrique, Assaf Razin, and Linda Tesar. 1994. "Effective Tax Rates in Macroeconomics: Cross-Country Estimates of Tax Rates on Factor Incomes and Consumption." *Journal of Monetary Economics,* 34: 297–323.

McDaniel, Cara. 2011. "Forces Shaping Hours Worked in the OECD, 1960-2004." *American Economic Journal: Macroeconomics,* 3(4): 27–52.

Mulligan, Casey. 2002. "A Century of Labor-Leisure Distortions." NBER Working Paper No. 8774.

Mulligan, Casey. 2012. "Do Welfare Policies Matter For Labor Market Aggregates? Quantifying Safety Net Work Incentives Since 2007." NBER Working Paper No. 18088.

Nekarda, Christopher and Valerie A. Ramey. 2013. "The Cyclical Behavior of the Price-Cost Markup." NBER Working Paper No. 19099.

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.

Ramey, Valerie. 2012. "The Impact of Hours Measures on the Trend and Cycle Behavior of U.S. Labor Productivity." University of California at San Diego Working Paper.

Rios-Rull, Jose-Victor, and Raul Santaeulalia-Llopis. 2010. "Redistributive Shocks and Productivity Shocks." *Journal of Monetary Economics,* 57: 931–48.

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): 267–97.

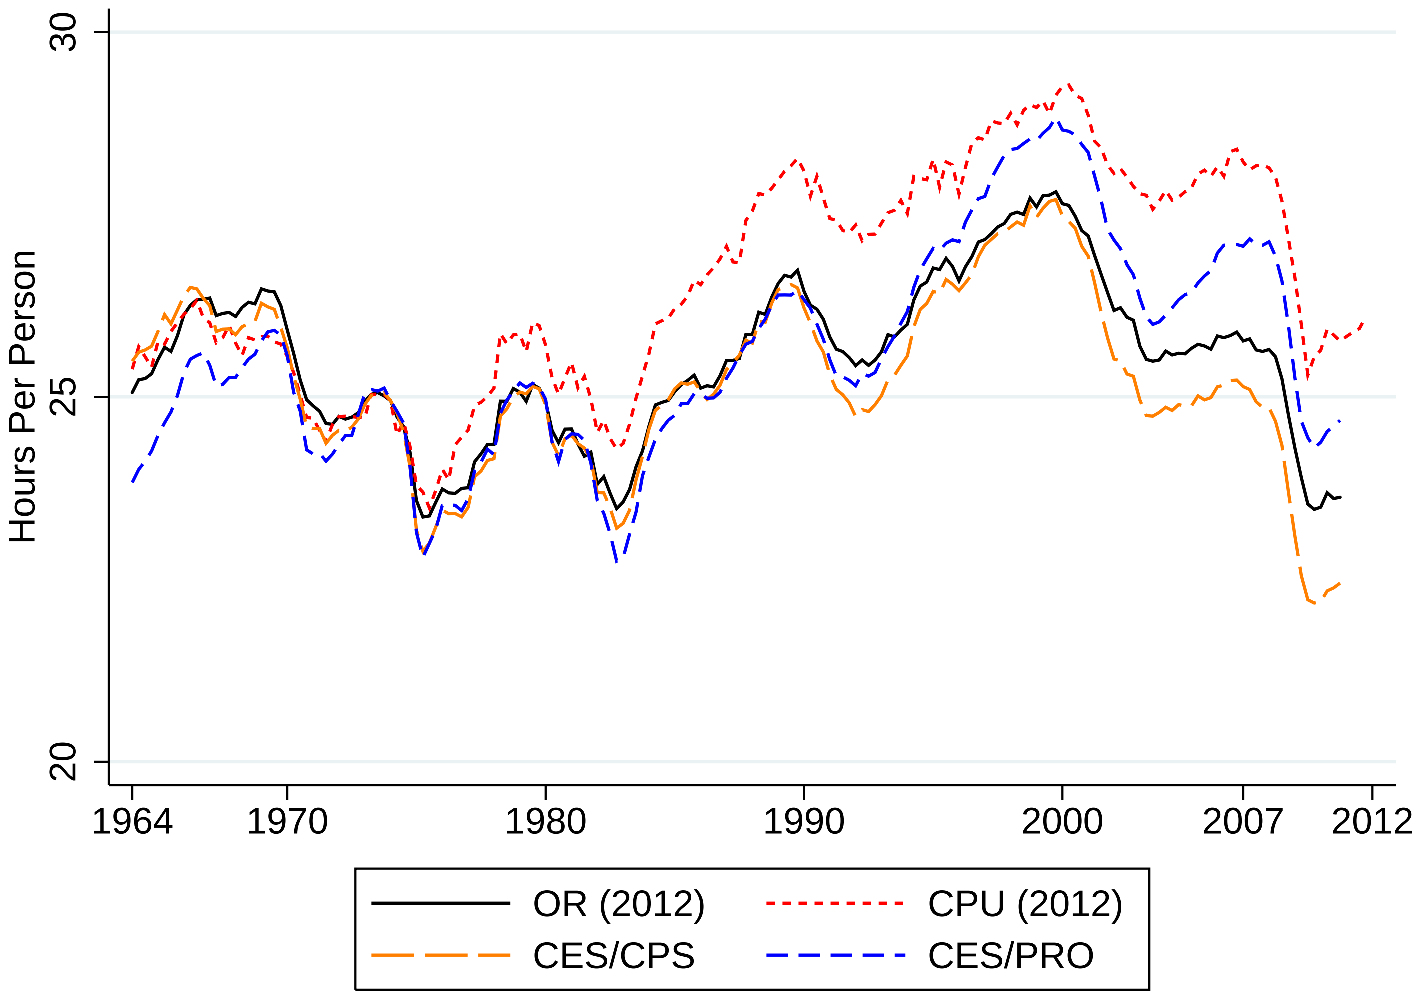

Figure 1: Hours Per Person in the United States

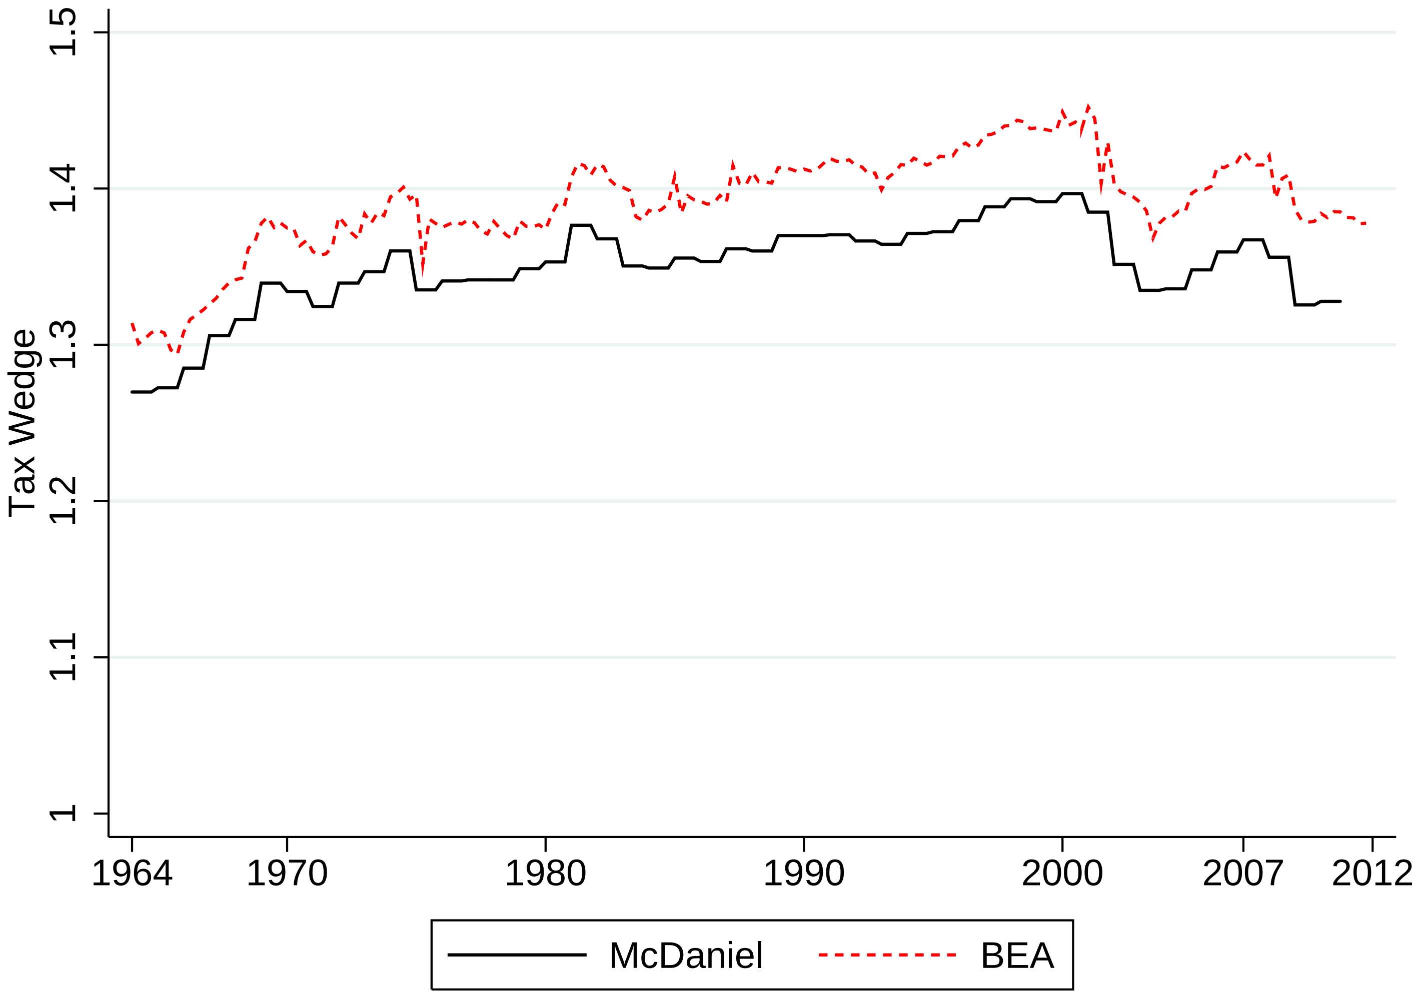

Figure 2: Tax Wedge in the United States

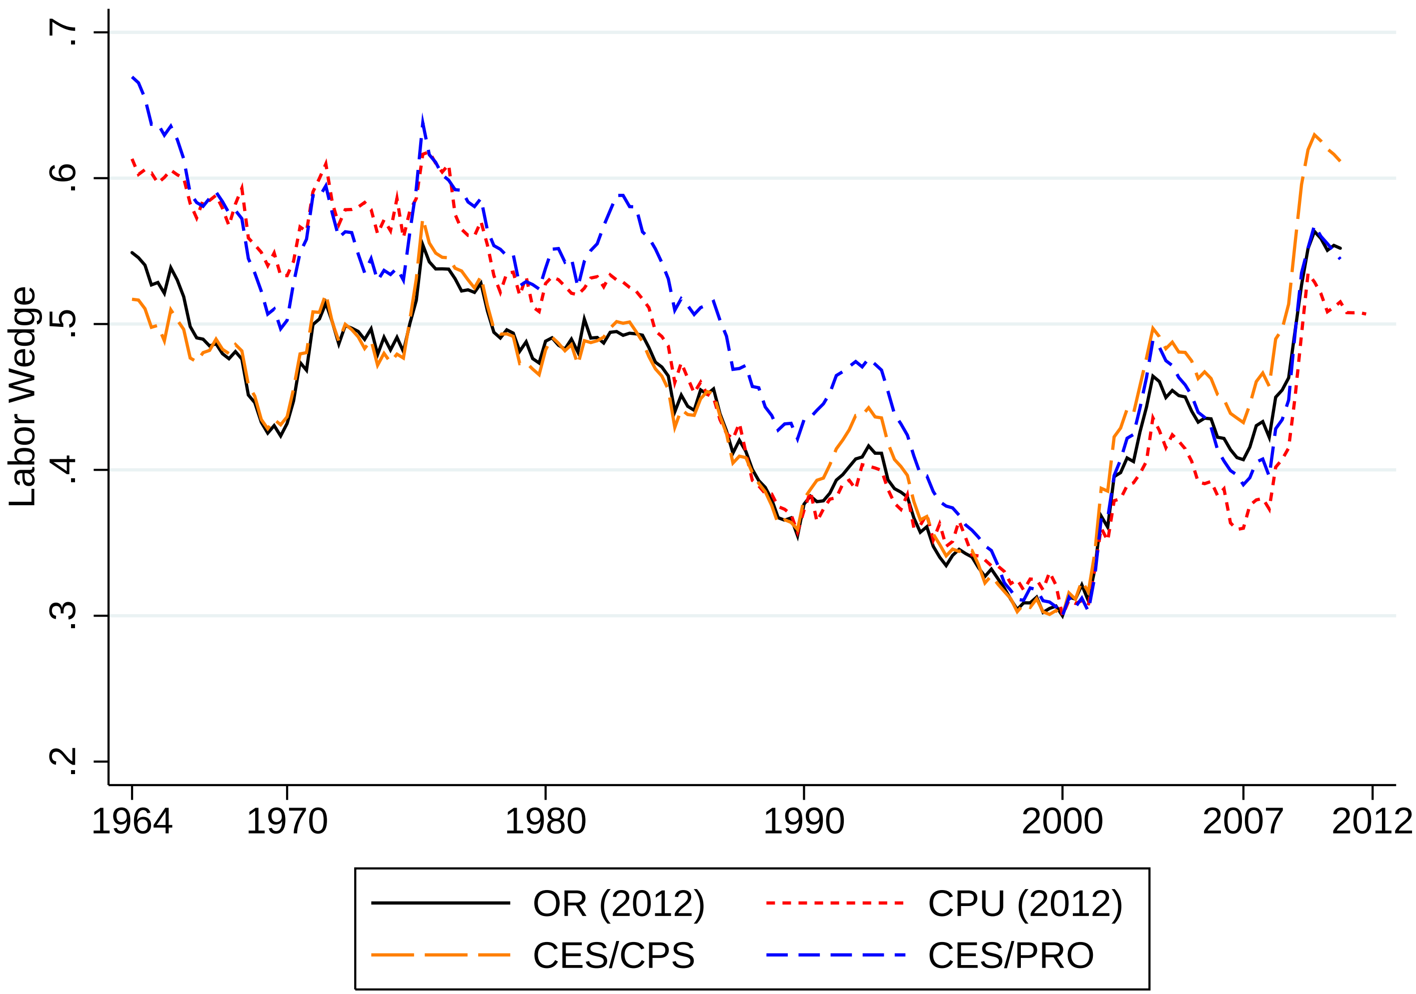

Figure 3: Labor Wedge in the United States

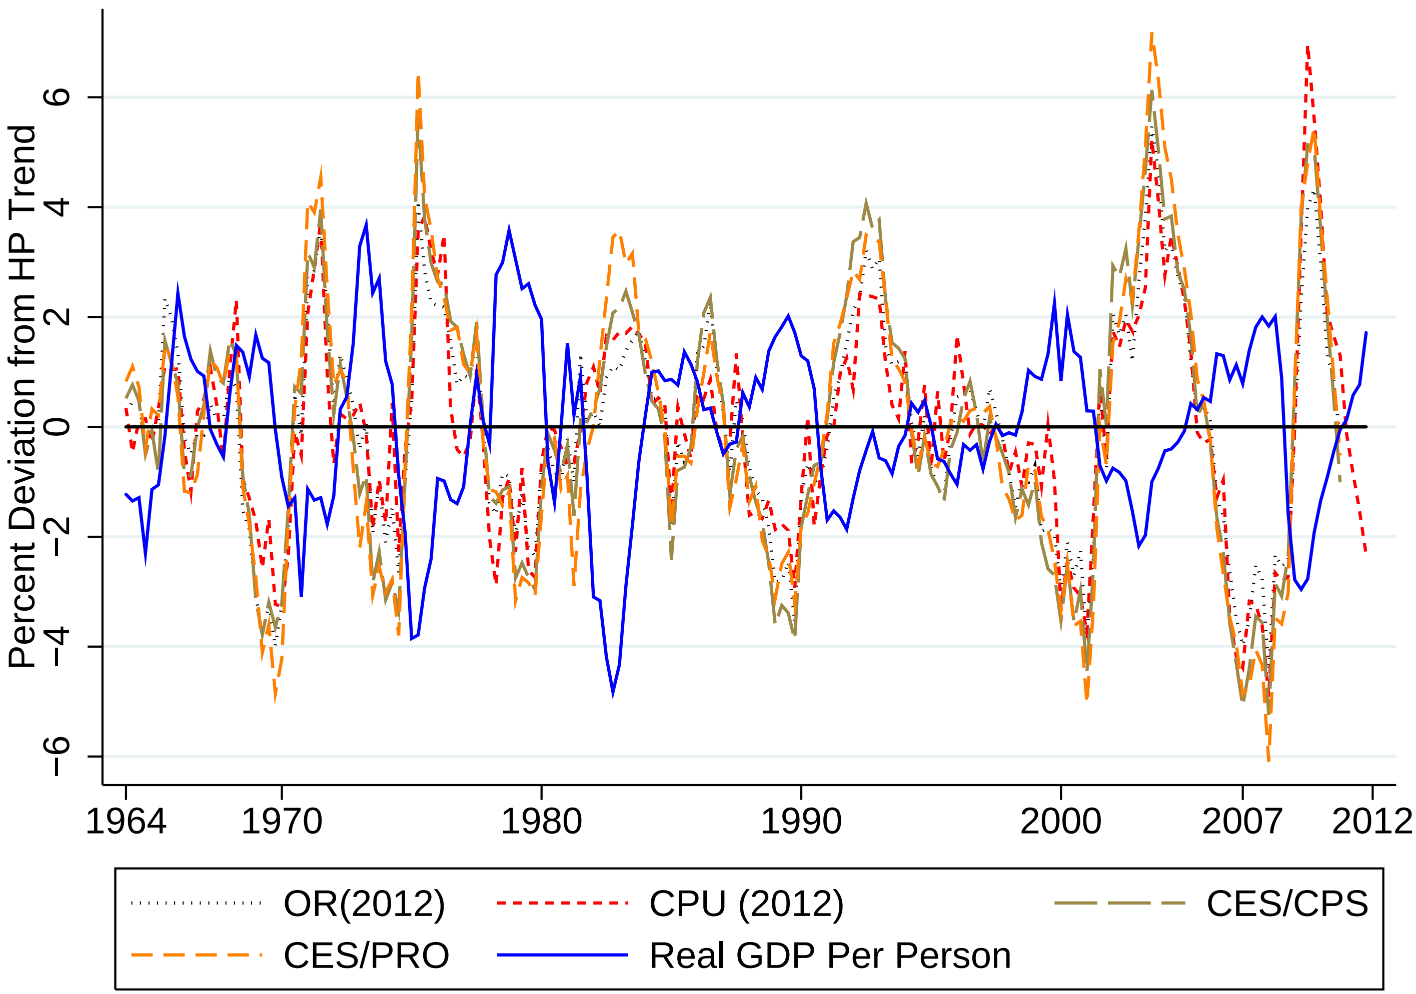

Figure 4: Cyclical Components of Labor Wedge and Output in the United States

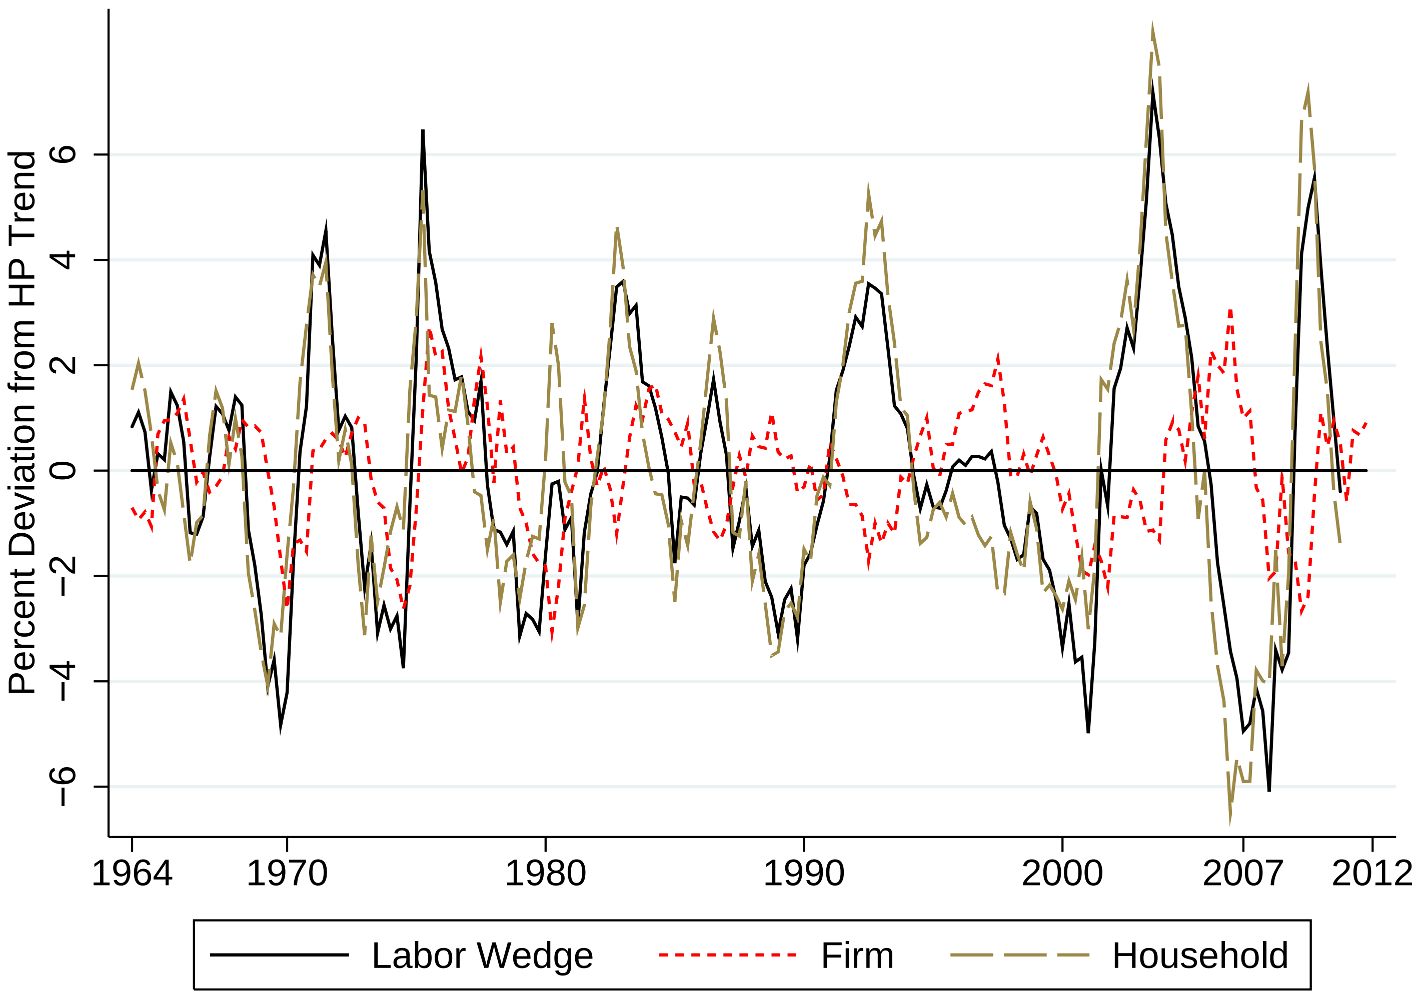

Figure 5: Decomposition of Cyclical Component of Labor Wedge in the United States

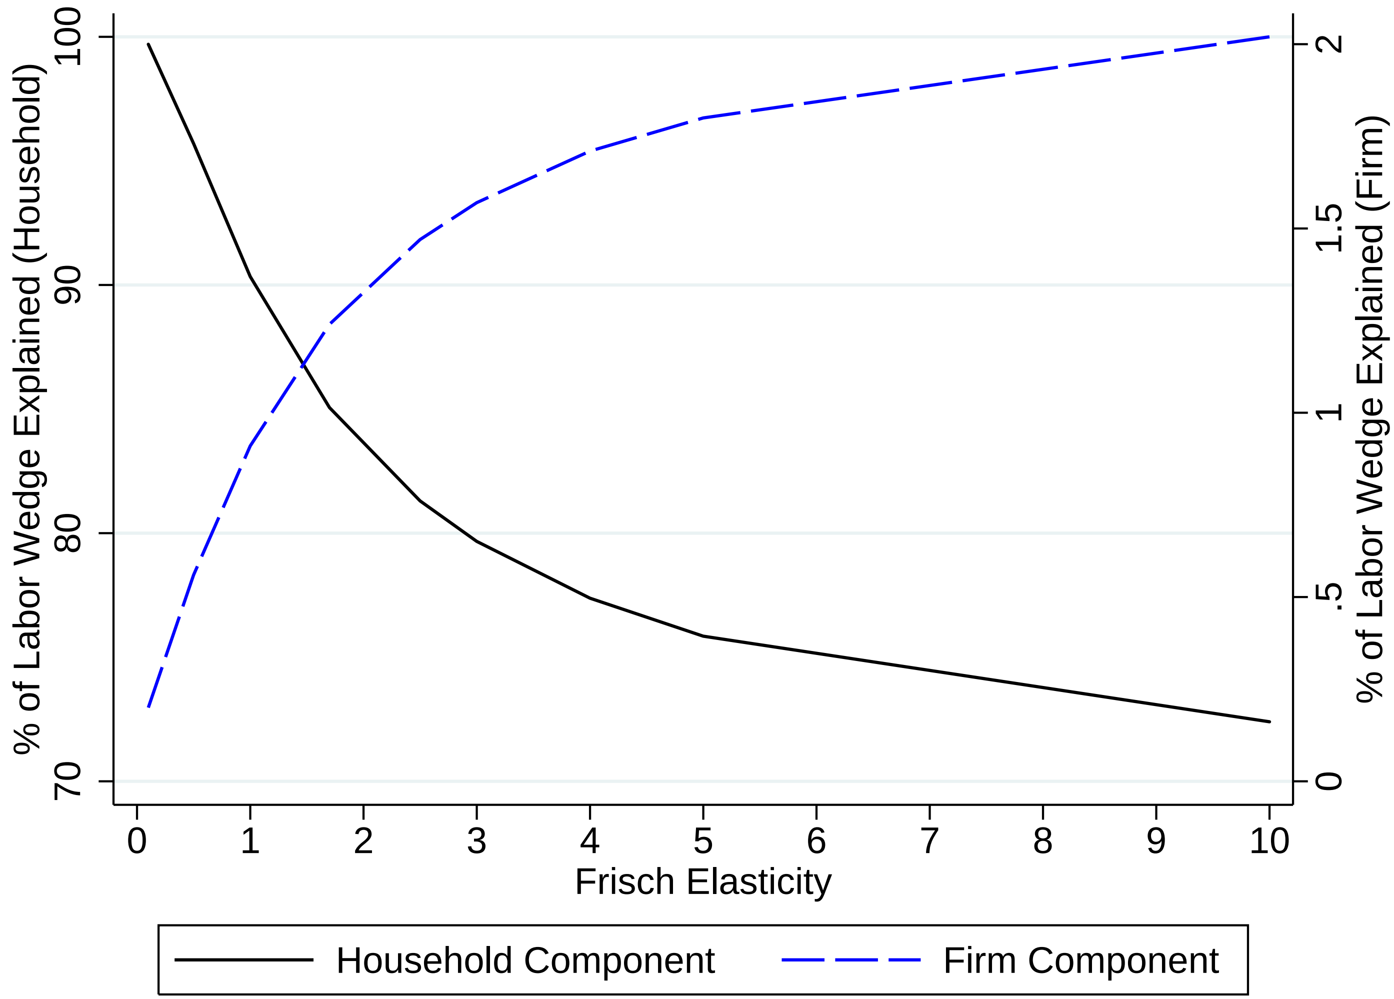

Figure 6: Decomposition of Labor Wedge Under Separable Preferences

Figure 7: Decomposition of Cyclical Fluctuations of Labor Share: BEA

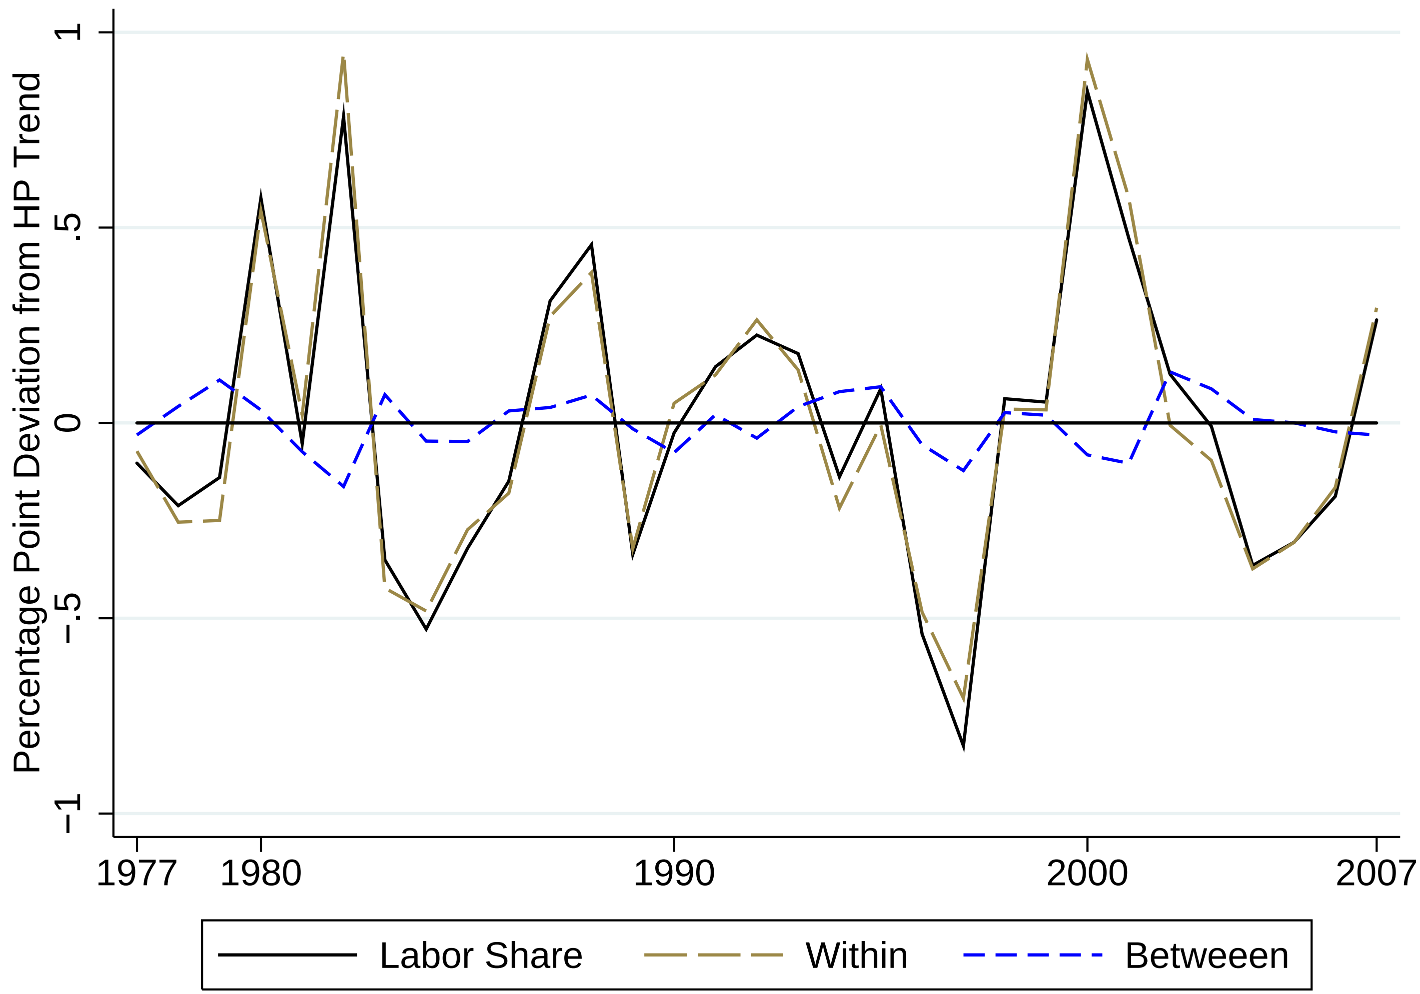

Figure 8: Decomposition of Cyclical Fluctuations of Labor Share: KLEMS

A. HP Filter:	OR (2012)	CPU (2012)	CES/CPS	CES/PRO
OR (2012)	1.00			
CPU (2012)	0.94	1.00		
CES/CPS	0.98	0.93	1.00	
CES/PRO	0.97	0.93	0.98	1.00

B. BK Filter:	OR (2012)	CPU (2012)	CES/CPS	CES/PRO
OR (2012)	1.00			
CPU (2012)	0.96	1.00		
CES/CPS	0.99	0.95	1.00	
CES/PRO	0.98	0.96	0.98	1.00

Table 1: Pairwise Correlations of Cyclical Components of Hours Per Person: 1964(1)-2011(4)

A. HP Filter:	BEA Unadjusted	BEA Adjusted	BLS Corporate	BEA Corporate
BEA Unadjusted	1.00			
BEA Adjusted	0.96	1.00		
BLS Corporate	0.85	0.86	1.00	
BEA Corporate	0.86	0.87	0.95	1.00
B. BK Filter:	BEA Unadjusted	BEA Adjusted	BLS Corporate	BEA Corporate
BEA Unadjusted	1.00			
BEA Adjusted	0.97	1.00		
BLS Corporate	0.85	0.86	1.00	
BEA Corporate	0.85	0.86	0.97	1.00

Table 2: Pairwise Correlations of Cyclical Components of Labor Share: 1964(1)-2011(4)

Labor Share	BEA Unadjusted	BEA Adjusted	BLS Corporate	BEA Corporate	
A. Firm Component:					**Average**
OR (2012) Hours	0.02	0.11	1.73	4.35	1.55
CPU (2012) Hours	0.09	0.03	1.10	4.67	1.47
CES/CPS Hours	1.50	1.83	0.39	2.05	1.44
CES/PRO Hours	0.59	0.81	1.41	4.04	1.72
Average	0.55	0.70	1.16	3.78	**1.55**
B. Household Component:					**Average**
OR (2012) Hours	80.89	83.32	70.01	68.37	75.65
CPU (2012) Hours	81.82	83.83	71.55	69.17	76.59
CES/CPS Hours	87.37	89.00	79.19	78.07	83.41
CES/PRO Hours	88.44	89.96	80.92	80.07	84.85
Average	84.63	86.53	75.42	73.92	**80.13**

Table 3: Percent Cyclical Variation of Labor Wedge Explained By Individual Components: HP Filter

Labor Share	BEA Unadjusted	BEA Adjusted	BLS Corporate	BEA Corporate	
A. Firm Component:					**Average**
OR (2012) Hours	0.40	1.13	2.05	2.54	1.53
CPU (2012) Hours	0.00	0.02	4.10	6.48	2.67
CES/CPS Hours	2.41	3.57	0.82	1.29	2.02
CES/PRO Hours	1.37	1.94	2.06	2.80	2.04
Average	1.05	1.66	2.26	3.28	**2.06**
B. Household Component:					**Average**
OR (2012) Hours	83.36	85.96	70.19	71.70	77.80
CPU (2012) Hours	80.58	83.19	66.24	65.22	73.81
CES/CPS Hours	89.22	90.91	79.67	80.72	85.13
CES/PRO Hours	90.23	91.69	81.57	82.62	86.53
Average	85.85	87.94	74.42	75.07	**80.82**

Table 4: Percent Cyclical Variation of Labor Wedge Explained By Individual Components: BK Filter

Labor Share	BEA Unadjusted	BEA Adjusted	BLS Corporate	BEA Corporate
A. HP Filter:				
corr $\left(\widehat{\tau}^f, \widehat{y}\right)$	0.46	0.44	0.14	0.10
corr $\left(\widehat{\tau}^f, \widehat{c}\right)$	0.32	0.28	0.07	0.02
corr $\left(\widehat{\tau}^f, \widehat{n}\right)$ (OR)	0.27	0.28	-0.02	-0.07
corr $\left(\widehat{\tau}^f, \widehat{n}\right)$ (CPU)	0.27	0.26	0.00	-0.08
sd $\left(\widehat{\tau}^f\right)$ /sd $(\widehat{y})$	0.60	0.55	0.74	0.73
B. BK Filter:				
corr $\left(\widehat{\tau}^f, \widehat{y}\right)$	0.44	0.39	0.04	0.01
corr $\left(\widehat{\tau}^f, \widehat{c}\right)$	0.32	0.27	0.02	-0.01
corr $\left(\widehat{\tau}^f, \widehat{n}\right)$ (OR)	0.29	0.30	-0.07	-0.08
corr $\left(\widehat{\tau}^f, \widehat{n}\right)$ (CPU)	0.24	0.22	-0.12	-0.16
sd $\left(\widehat{\tau}^f\right)$ /sd $(\widehat{y})$	0.57	0.51	0.72	0.71

Table 5: Cyclical Properties of Firm Component of Labor Wedge

Hours Per Person	OR (2012)	CPU (2012)	CES/PRO	CES/CPS
A. HP Filter:				
corr $\left(\widehat{\tau}^h, \widehat{y}\right)$	-0.54	-0.56	-0.62	-0.66
corr $\left(\widehat{\tau}^h, \widehat{c}\right)$	-0.56	-0.61	-0.60	-0.64
corr $\left(\widehat{\tau}^h, \widehat{n}\right)$ (OR)	-0.74	-0.69	-0.78	-0.81
corr $\left(\widehat{\tau}^h, \widehat{n}\right)$ (CPU)	-0.67	-0.75	-0.72	-0.76
sd $\left(\widehat{\tau}^h\right)$ /sd $\left(\widehat{y}\right)$	1.35	1.39	1.63	1.70
B. BK Filter:				
corr $\left(\widehat{\tau}^h, \widehat{y}\right)$	-0.57	-0.59	-0.64	-0.68
corr $\left(\widehat{\tau}^h, \widehat{c}\right)$	-0.56	-0.61	-0.61	-0.63
corr $\left(\widehat{\tau}^h, \widehat{n}\right)$ (OR)	-0.74	-0.69	-0.79	-0.81
corr $\left(\widehat{\tau}^h, \widehat{n}\right)$ (CPU)	-0.66	-0.71	-0.73	-0.76
sd $\left(\widehat{\tau}^h\right)$ /sd $\left(\widehat{y}\right)$	1.30	1.21	1.58	1.65

Table 6: Cyclical Properties of Household Component of Labor Wedge

A. HP Filter:	BEA Unadjusted	BEA Adjusted	KLEMS Unadjusted	KLEMS Adjusted
Fraction of corr $(\widehat{s}_{j,t}, \widehat{y}_t) > 0$	0.25	0.31	0.20	0.20
Median of corr $(\widehat{s}_{j,t}, \widehat{y}_t)$	-0.14	-0.18	-0.26	-0.31
Fraction of corr $(\widehat{s}_{j,t}, \widehat{s}_t) > 0$	0.88	0.94	1.00	1.00
Median of corr $(\widehat{s}_{j,t}, \widehat{s}_t)$	0.56	0.51	0.49	0.46
B. BK Filter:	BEA Unadjusted	BEA Adjusted	KLEMS Unadjusted	KLEMS Adjusted
Fraction of corr $(\widehat{s}_{j,t}, \widehat{y}_t) > 0$	0.31	0.37	0.10	0.10
Median of corr $(\widehat{s}_{j,t}, \widehat{y}_t)$	-0.05	-0.12	-0.21	-0.25
Fraction of corr $(\widehat{s}_{j,t}, \widehat{s}_t) > 0$	0.81	0.94	1.00	1.00
Median of corr $(\widehat{s}_{j,t}, \widehat{s}_t)$	0.54	0.49	0.50	0.52

Table 7: Cyclical Properties of Industry-Level Labor Shares

Labor Share	OECD Unadjusted		OECD Adjusted		KLEMS Unadjusted		KLEMS Adjusted		**Average**	
	Firm	Household	Firm	Household	Firm	Household	Firm	Household	Firm	Household
Australia	1.30	57.21	8.82	75.75	10.09	72.37	0.62	79.16	5.21	71.12
Austria	27.83	42.30	23.97	11.45	32.31	61.10	35.88	64.19	29.99	44.76
Canada	4.05	72.75	0.07	63.34	—	—	—	—	2.06	68.05
Finland	2.17	62.61	3.50	58.90	4.29	70.61	5.66	68.10	3.91	65.06
France	4.14	83.57	0.18	70.69	6.74	85.98	1.12	82.26	3.05	80.63
Germany	24.01	40.33	27.90	36.64	24.29	58.10	19.59	61.14	23.95	49.05
Ireland	1.78	78.36	0.32	59.07	0.98	73.21	0.08	78.45	0.79	72.27
Italy	9.25	75.49	22.08	75.43	21.91	76.10	18.30	74.49	17.89	75.38
Japan	6.66	70.98	1.41	58.53	6.73	66.90	4.20	67.49	4.75	65.98
Korea	6.51	90.07	19.86	86.06	30.21	86.42	13.96	84.67	17.64	86.81
Norway	12.64	17.38	17.78	8.66	—	—	—	—	15.21	13.02
Spain	15.81	95.52	1.01	98.24	0.20	58.44	0.10	59.67	4.28	77.97
Sweden	0.49	30.28	13.62	78.57	14.49	79.12	16.88	79.94	11.37	66.98
United Kingdom	7.70	45.33	9.99	47.38	13.37	39.64	17.08	41.02	12.04	43.34
United States	3.00	88.04	0.09	90.69	2.00	79.30	0.55	80.54	1.41	84.64

Table 8: Percent Cyclical Variation of Labor Wedge Explained By Individual Components: HP Filter

Labor Share	OECD Unadjusted		OECD Adjusted		KLEMS Unadjusted		KLEMS Adjusted		**Average**	
	Firm	Household	Firm	Household	Firm	Household	Firm	Household	Firm	Household
Australia	3.20	58.11	8.00	80.36	16.20	78.94	0.00	84.80	6.85	75.55
Austria	41.78	59.63	22.19	15.39	22.60	67.76	26.91	70.83	28.37	53.40
Canada	1.13	73.41	0.88	63.66	—	—	—	—	1.01	68.54
Finland	0.08	60.50	1.41	57.16	0.91	67.04	1.09	65.26	0.87	62.49
France	0.67	75.96	0.18	71.26	1.58	87.16	0.05	84.60	0.62	79.75
Germany	22.53	40.87	26.70	41.19	24.09	60.10	19.76	63.34	23.27	51.38
Ireland	0.02	87.78	0.41	63.45	2.81	75.74	0.06	83.38	0.83	77.59
Italy	11.63	75.72	22.94	71.75	23.15	73.45	18.60	72.66	19.08	73.40
Japan	0.46	68.09	1.29	55.71	0.37	55.01	2.39	53.10	1.13	57.98
Korea	7.35	91.42	22.49	88.11	46.96	89.36	19.45	87.91	24.06	89.20
Norway	18.30	12.18	24.81	4.52	—	—	—	—	21.55	8.35
Spain	15.87	96.07	0.65	97.19	2.79	97.02	5.74	97.73	6.26	97.00
Sweden	3.52	68.03	11.11	79.20	13.25	80.85	15.20	81.92	10.77	77.50
United Kingdom	14.64	43.51	16.84	45.73	14.68	39.36	19.43	40.17	16.40	42.19
United States	0.87	87.74	0.78	90.59	1.29	84.34	0.23	85.29	0.79	86.99

Table 9: Percent Cyclical Variation of Labor Wedge Explained By Individual Components: BK Filter

Labor Share	OECD Unadjusted		OECD Adjusted		KLEMS Unadjusted		KLEMS Adjusted	
	HP Filter	BK Filter	HP Filter	BK Filter	HP Filter	BK Filter	HP Filter	BK Filter
Australia	0.26	0.26	0.33	0.38	0.40	0.45	0.41	0.47
Austria	0.60	0.56	0.56	0.51	0.50	0.61	0.41	0.56
Canada	0.68	0.63	0.55	0.57	—	—	—	—
Finland	0.48	0.41	0.51	0.43	0.24	0.25	0.33	0.33
France	0.56	0.47	0.65	0.61	0.35	0.24	0.24	0.27
Germany	0.45	0.26	0.44	0.24	-0.04	-0.03	0.07	0.04
Ireland	0.38	0.39	0.40	0.49	0.22	0.28	0.15	0.25
Italy	0.38	0.33	0.53	0.49	0.46	0.43	0.57	0.57
Japan	0.77	0.75	0.02	0.00	0.63	0.44	0.62	0.44
Korea	0.01	-0.01	0.00	-0.04	-0.41	-0.37	-0.06	-0.03
Norway	0.21	0.18	0.08	0.11	—	—	—	—
Spain	-0.01	0.06	0.02	0.07	-0.21	-0.10	-0.09	-0.01
Sweden	0.59	0.39	0.39	0.31	0.19	0.18	0.15	0.15
United Kingdom	0.55	0.50	0.54	0.52	0.48	0.51	0.46	0.50
United States	0.23	0.25	0.30	0.36	0.19	0.22	0.28	0.31

Table 10: Correlation of Firm Component with Real Output Per Person (Cyclical Components)