

Long-Term Changes in Labor Supply and Taxes: Evidence from OECD Countries, 1956-2004*

Lee Ohanian

University of California, Los Angeles

Andrea Raffo

Federal Reserve Bank of Kansas City

Richard Rogerson

Arizona State University

December 2006

Abstract

We document large differences in trend changes in hours worked across OECD countries over the period 1956-2004. We then assess the extent to which these changes are consistent with the intratemporal first order condition from the neo-classical growth model. We find large and trending deviations from this condition, and that the model can account for virtually none of the changes in hours worked. We then extend the model to incorporate observed changes in taxes. Our findings suggest that taxes can account for much of the variation in hours worked both over time and across countries.

*We would like to thank Andy Atkeson, Ariel Burstein, Hal Cole, Chuck Carlstrom, Mark Doms, Christian Haefke, Gary Hansen, Ed Prescott and participants at the Research Department at the Federal Reserve Bank of Atlanta and Kansas City, 2006 Midwest Macroeconomic Meetings, 2006 NBER Summer Institute, 2006 ECB-IMOP conference, Society for Economic Dynamics 2004 Annual Meeting, and UC Davis for comments. We thank Cara McDaniel for supplying us with her tax rate series. Sonal Hate provided excellent research assistance. Rogerson thanks the NSF for financial support. The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Kansas City or the Federal Reserve System.

1 Introduction

Macroeconomists have long been interested in understanding the time series behavior of aggregate hours of work. Much effort has been devoted to one particular context – changes in hours worked at business cycle frequencies in the US economy – and assessing the role of various model features in accounting for these changes. The starting point for this paper is the observation that *trend* changes in hours of work across OECD countries over the last 50 years exceed *business cycle* changes in hours of work by roughly an order of magnitude. We believe that these large trend changes warrant a systematic analysis similar to that which has occurred for business cycle movements. This paper takes a first step in this direction. In particular, we assess the ability of several versions of the neoclassical growth model to account for trend changes in hours of work for a sample of 21 OECD countries over the period 1956 - 2004. One goal is to isolate those episodes, i.e., time periods in specific countries, which seem most puzzling from the perspective of this theoretical framework. To our knowledge, this is the most comprehensive analysis of long-run labor supply using the growth model.

Our main findings follow. First, the standard version of the growth model can account for only a small fraction of the large, trend decreases in hours of work observed over this period. This conclusion is robust to the value of the labor supply elasticity assumed, the inclusion of a subsistence consumption term in preferences, and to assumptions regarding how individuals value government provision of goods and services. Second, when the analysis is extended to incorporate taxes on consumption and labor income, we find that the model can account for almost all of the average decrease in hours of work over the sample period. Additionally, the model with consumption and labor taxes represents an improvement in accounting for changes at the individual country level. Third, although the model with distorting taxes can largely account for the average change in hours of work over time, there are several episodes that the model does not account for. Some of

these episodes entail periods in which hours do not fall enough in specific countries, while others entail periods in which hours fall too much in specific countries. Importantly, those episodes which are singled out as puzzling relative to this extension of the standard model are significantly different from those episodes which stand out as puzzling relative to the standard model without distorting taxes. Fourth, we present statistical evidence to show that taxes are important even after controlling for a variety of other features that have been suggested as important determinants of cross-country differences in labor supply.

Given that the one-sector growth model has become the standard framework for interpreting aggregate economic data, it is the natural starting point for trying to understand trend changes in hours across countries. To assess the ability of this framework to account for the large changes in hours worked, we employ a methodology that has been found to be very useful in the business cycle literature. Specifically, we focus on the static first order condition implied by equilibrium and assess the extent to which this condition holds at each point in time in the data. One feature of this exercise is that it does not require us to assume that preferences or technology are the same across countries. We implement this exercise both for a benchmark model and for several extensions of the benchmark model, and assess how various extensions influence the extent to which the condition holds. Isolating those episodes where the condition is farthest from being satisfied serve to identify those episodes on which we need to focus and on the direction in which the model fails.

This same methodology has been profitably employed in the business cycle literature by Parkin (1988), Bencivenga (1992), Ingram, Kocherlakota, Savin (1994), Hall (1997), Gali, Gertler and Lopez-Salido (2002), and Chari, Kehoe and McGrattan (2002). Mulligan (2002) uses this method to analyze changes in hours of work in the US over the 20th century, while Cole and Ohanian (2004) use it to shed light on changes in hours worked during the U.S. Great Depression.

In studying the effects of taxes on labor supply, the study most similar to ours is Prescott (2004). He assesses the extent to which the standard growth model with taxes can account

for the changes in hours of work for a small set of countries between two particular points in time. Relative to Prescott, our contribution is fourfold. First, we extend the analysis to a much larger set of countries. Second, we consider a much longer time period. Third, we assess the equilibrium condition for all years in our data set and not just at two distinct points in time. Fourth, we assess how several other modifications of the theory affect the findings.

The paper is organized as follows. In the next section we document the large and persistent changes in total hours worked across OECD countries over the last fifty years. In particular, we document the large reduction over time in the cross-sectional mean of hours worked and the large dispersion in the extent of the decrease across countries. In Section 3 we describe the benchmark model and the methodology that we employ to construct labor wedges from the static optimality condition. Section 4 presents our findings for the benchmark model, and Section 5 considers several extensions. Section 6 extends the analysis to consider distorting taxes. Section 7 carries out a statistical analysis of the wedges from our benchmark analysis and several factors thought to be important in influencing labor market outcomes. Section 8 concludes.

2 Hours Worked in 21 OECD Countries: 1956-2004

In this section we describe some key features of the distribution of hours worked across 21 OECD countries for the period 1956-2004. Our measure of aggregate hours worked is the product of total civilian employment and annual hours worked per person in employment, divided by the size of the population aged 15-64.¹ Specifically, we document the following three features:

(1) On average, labor supply is falling: mean hours of work in these countries have decreased substantially between 1956-2004.

(2) The magnitude of the decline in hours worked varies significantly across countries

¹Information about data sources is contained in the Appendix.

between 1956 and 2004.

(3) Differences in trend changes in hours worked account for much of the variance in hours worked across countries and over time.

We begin with the behavior of average hours worked over time. Figure 1 plots the cross-country average of hours worked for each year from 1956 to 2004.

[Insert here Figure 1. Mean Hours Worked]

As is easily seen, this time period has witnessed a dramatic decrease in mean hours worked, with the 2004 value almost 20% less than the 1956 value. This decline occurs at a relatively steady pace over the period 1956-1985, at which point mean hours worked becomes relatively flat.

The second property that we establish is that changes in hours of work between 1956 and 2004 have been far from uniform across countries. Table 1 displays the value of hours worked in 2004 relative to 1956 for all 21 countries in our sample.

[Insert here Table 1. Hours Worked in 2004 Relative to 1956]

While the average change is a 20% decrease, the changes range from an increase of 8% for Canada to a decrease of 40% for Germany, with other countries distributed throughout.

It is also of interest to examine the time series changes for each of the countries. We classify countries into 4 groups based on the trend behavior of hours worked between 1956 and 2004. Group 1 includes Austria, Belgium, Denmark, Finland, France, Germany, and Italy. These countries experienced the steepest decline in hours worked in the postwar period, and exhibit a monotonic decline over time until hours worked level off around 1990. Group 2 includes Japan, Norway, Portugal, Sweden, and the UK. Similar to Group 1, these countries experienced a monotonic decrease in hours prior to leveling off at the

end of the period, but the overall decrease is of a smaller magnitude than for Group 1. Group 3 includes those countries for which there is no major trend in hours worked, namely Australia, Canada, New Zealand, and the US. The fourth and final group consists of Greece, Ireland, Netherlands, Spain and Switzerland. A distinctive feature of this fourth group is that the trend behavior in hours worked is not monotone. For simplicity, Figure 2 depicts the behavior of mean hours for Groups 1, 2 and 3².

[Insert here Figure 2. Mean Hours Worked by Group]

Consistent with the definition of the three groups, Groups 1 and 2 both display substantial decreases over time, with Group 1 showing the larger decline, while Group 3 displays relatively little change over time.

The final property that we document is the importance of the differences in trend changes in hours worked in accounting for the dispersion in hours worked in the panel of countries. We first note that the dispersion of hours worked across countries is large at all points in time during the period 1956-2004. In particular, the cross-sectional standard deviation of log hours averages roughly 0.12, and never falls below 0.09. In contrast, the standard deviation of the cyclical component of log hours worked in the US time series is only 0.02. To indicate the importance of differences in trends in accounting for this dispersion, we run a panel regression of log hours on a common constant term and a country specific linear time trend:

$$\log h_{it} = a + b_it + \varepsilon_{it}$$

where h_{it} is hours worked for country i in period t , and t runs from 1 to 49. The R-squared from this regression is 0.76.³ In contrast, if we run the same regression imposing that the trend coefficient is the same across all countries then the R-squared drops to .26.

²We do not include average hours worked for Group 4 because of the high heterogeneity within the group. The figure is available upon request.

³If one adds a country specific intercept the R-squared for this regression increases to 0.89.

To summarize, the most striking features of this international comparison of hours worked is that hours have fallen substantially in most countries. In addition, this process of declining hours is not uniform, but differs considerably across countries. Finally, these differences in trend rates of decline are the dominant source of dispersion of hours worked in the panel of countries over the period 1956-2004. We next develop a benchmark growth model to evaluate these long-run patterns.

3 Wedges in a Benchmark Model

Given that the one-sector growth model has become the standard framework for organizing and interpreting aggregate data, it is also the natural starting point for our analysis of changes in aggregate hours worked across countries. In this section we describe our benchmark model and define the concept of a wedge in the static first order condition governing the choice of the household's allocation of time that serve as the focus of our analysis.

3.1 Model

While the details of this model are standard, we describe them briefly in order to introduce the notation that is needed subsequently. The economy consists of a single household with utility defined over streams of (private) consumption (C_t) and leisure time ($\bar{H} - H_t$) given by:

$$\sum_{t=0}^{\infty} \beta^t U(C_t, \bar{H} - H_t) \tag{1}$$

where $0 < \beta < 1$ is the discount factor. The household is endowed with \bar{H} units of time each period, and H_t represents time devoted to market work. As is standard, we restrict the form of the utility function U to be consistent with balanced growth.⁴ We further

⁴One may interpret the trend decline in hours worked as evidence against balanced growth preferences. See, for example, Blanchard (2005). Our exercise should be interpreted as assessing the ability of a model with balanced growth preferences to account for the long-term changes in hours worked.

restrict the utility function to be of the form:

$$U(C_t, \bar{H} - H_t) = \alpha \log C_t + (1 - \alpha) \frac{(\bar{H} - H_t)^{1-\gamma} - 1}{1 - \gamma} \quad (2)$$

where $\gamma \geq 0$, and $0 \leq \alpha \leq 1$. This form of the utility function is of particular interest because varying the parameter γ changes the elasticity of substitution between leisure and consumption, which is known to be an important parameter in analyzing how various factors influence hours of work in equilibrium.⁵

Technology is specified by a Cobb-Douglas production function:

$$Y_t = A_t K_t^\theta H_t^{1-\theta} \quad (3)$$

where A_t is a measure of technology at time t , and K_t and H_t represent inputs of capital and labor services at time t . Output at time t must be divided between consumption and investment (X_t) and the capital stock evolves according to the standard law of motion:

$$K_{t+1} = (1 - \delta)K_t + X_t \quad (4)$$

where $0 < \delta < 1$ is the depreciation rate.

3.2 Wedges

We focus on the competitive equilibrium allocation for this economy. Combining the first order conditions from the consumer and firm problems, one derives the so-called “static” first order condition as one of the conditions characterizing the competitive equilibrium allocation:

$$\frac{U_2(C_t, \bar{H} - H_t)}{U_1(C_t, \bar{H} - H_t)} = F_2(K_t, H_t, A_t). \quad (5)$$

This condition states that in equilibrium, the marginal rate of substitution between consumption and leisure must equal the marginal product of labor at each point in time. Given

⁵Conditional on restricting attention to balanced growth preferences, the assumption of separability plays no substantive role in the analysis.

our functional form assumptions it reduces to:

$$\frac{(1 - \alpha)}{\alpha} \frac{C_t}{(\bar{H} - H_t)^\gamma} = (1 - \theta) \frac{Y_t}{H_t} \quad (6)$$

While this is only one of the conditions imposed by equilibrium, it has proven to be a useful and widely-used diagnostic both to assess the model's ability to account for observed changes in hours, and to provide information as to what types of additional factors would allow the model to better account for the data. As noted in the introduction, this analysis has typically examined the changes in hours worked within a given country at business cycle frequencies. Our goal is to use this condition to help shed light on the trend changes documented in the previous section.

Because our interest is in assessing the extent to which this first order condition does not hold in the data, it is convenient to introduce some notation to capture this. For notational convenience we abstract from the country index in what follows. Let H_t , Y_t , and C_t denote actual time series data for a given country. We define a time series for Δ_t for this same country by requiring that the following equation hold at each point in time:

$$\frac{(1 - \alpha)}{\alpha} \frac{H_t}{(\bar{H} - H_t)^\gamma} = (1 - \Delta_t)(1 - \theta) \frac{Y_t}{C_t} \quad (7)$$

We will refer to Δ_t as the wedge in period t . Note that given our definition, a value of $\Delta_t = 0$ implies that the equilibrium first order condition holds exactly in the data. Written somewhat more compactly, Δ_t can be expressed as:

$$\Delta_t = 1 - \frac{MRS_t}{MPL_t} \quad (8)$$

where MRS_t is the marginal rate of substitution between leisure and consumption as implied by the data in period t , and MPL_t is the marginal product of labor as implied by the data in period t . Thus, the wedge measures the percentage deviation between the marginal rate of substitution and the marginal product of labor. If the benchmark model is successful, then the wedge factors will be small over time and across countries. For future

reference, we note that if the wedge is positive, this implies that hours in the model are “too low” relative to hours in the data.

4 Results for the Benchmark Model

In this section we compute time series for the wedge for each of the 21 countries in our sample, and analyze their properties. We report results for two different calculations. In the first calculation we focus on the time series for each country considered individually. The advantage of this approach is that it does not require us to assume that preference and technology parameters are the same across countries. In the second calculation we assume that technology and preference parameters are the same across countries and compute the wedges required for the model to account for the cross-country dispersion in hours worked.

4.1 Country-Level Analysis

Letting i denote country and t denote year, define B_{it} by:

$$B_{it} = \frac{H_{it}}{(\bar{H} - H_{it})^\gamma} \frac{C_{it}}{Y_{it}}. \quad (9)$$

It follows directly from equation (7) that:

$$\frac{1 - \Delta_{it+s}}{1 - \Delta_{it}} = \frac{B_{it+s}}{B_{it}} \quad (10)$$

It is apparent that one can compute the time series values for the wedge relative to some benchmark year without any information about the values of α and θ .

Given that one of our interests is in understanding what factors are able to account for the large changes in hours worked across time, we focus on the change in the wedge over time as well, without trying to assess the absolute level of the wedge. In view of this we normalize the wedge factor to equal 0 in what we call a normalization year for all countries and then use equation (10) to compute the time series for the wedge relative to the normalization year in all the other years. In what follows we choose 1980 as the

normalization year for all countries. Although it is not necessary to supply values for α and θ to carry out this exercise, it is necessary to assume values for γ and \bar{H} in order to compute a series for Δ_{it} . The results presented here are for the case in which preferences are log in consumption and leisure, i.e., the limiting case as γ tends to one. We discuss sensitivity of the results to the choice of γ later on. The value of \bar{H} is set to $14 \times 365 = 5110$ for these calculations.

Given values for γ and \bar{H} , the time series for Δ can be computed given data for consumption, output and hours worked.⁶ Our measure of output is GDP. For our benchmark results in this section our measure of consumption is private consumption expenditures on nondurables, services, and durables⁷.

The first property we report is the time series for the cross-sectional average of wedge factors across countries. This is shown in Figure 3, along with a line showing the change in mean log hours relative to 1980.

[Insert Figure 3. Mean Wedge and Mean Hours]

The average wedge increases at a fairly steady rate from 1956 though to the mid 1980s, at which point it levels off. Overall, the increase is about 0.4. Given our normalization of the wedge to be zero in all countries in 1980, it is only the change in the wedge that has any significance. The negative comovement between the hours and wedge series is striking. In fact, the correlation between the two series is nearly -1.00.

Before examining the series at the country level, we contrast the behavior of the average wedge across the three groups defined earlier in the paper. This is done in Figure 4.

[Insert Figure 4. Mean Wedge by Group]

⁶Details of the data used are provided in the Appendix.

⁷While it would be of interest to use a measure of consumption which incorporates the flow of services from durables rather than expenditures on durables, we do not pursue it due to data limitations for a number of countries.

Recalling the figure that showed the behavior of mean hours for these three groups, we note that there is again a striking pattern between the behavior of mean hours and mean wedges even at the group level. Group 1 exhibits a larger change (in absolute value) for both hours and the wedge as compared to Group 2. Hours and wedges are relatively flat after 1985 for both groups as well. And Group 3 displays little trend for either hours or the wedge.

We next examine the correlation between hours and wedges over time at the country level. Because our main goal is to understand the changes in the trend component of hours worked rather than the business cycle component, we isolate the trend component of each series using the Hodrick-Prescott (HP) filter with a smoothing parameter of 100. Table 2 presents the correlations between the trend components of the two series.

[Insert Table 2. Correlation Hours and Wedges: Trend Components]

As the table indicates, the two series exhibit a very strong negative correlation across virtually all countries. The majority of the values exceed 0.97 in absolute value, and only three values are less than 0.80. Though we will not be interested in cyclical movements, for completeness we report the correlations between the cyclical components in Table 3.

[Insert Table 3. Correlation Hours and Wedges: Cyclical Components]

Consistent with earlier studies, we find a very large and negative correlation for the US. The table indicates that a negative correlation is found for all the other countries, though the correlation is not as strong as for the US. The mean correlation across countries is equal to -0.79 .

Table 1 documented that trend changes in hours worked differed significantly across countries. In light of this property of the data, we examine how trend changes in wedges differ across countries, and how they are correlated with trend changes in hours worked.

Figure 5 plots the country-level pairs for the change in log hours worked from 1956 to 2004 and the change in wedges from 1956 to 2004.

[Insert Figure 5. Change in Wedges and Hours]

As the figure indicates, changes in hours and changes in wedges are strongly negatively correlated. In fact, the correlation between the two is -0.83 .

The above statistics provide a fairly strong characterization of the evolution of the wedges over time, together with their relation with the evolution of hours worked. For completeness, given that we did not display hours worked in Group 4 (countries with non-monotonic hours), we examine the wedges for these countries. In the interest of space we show the results for Netherlands and Spain in Figures 6 and 7. Most noteworthy here is the fact that in both cases, the turning points in trend hours are associated with turning points in the wedge.

[Insert Figure 6. Wedges and Hours: Netherlands]

[Insert Figure 7. Wedges and Hours: Spain]

4.2 Cross-Country Comparison

To this point we have only compared the changes in wedges over time for different countries. However, given the large dispersion in hours of work across countries at any point in time, it is also of interest to assess the wedges associated with the cross-section of countries. In order to make such a calculation one must be explicit about the relative values of α and θ across countries. In this section we assume that these values are the same for all countries. In this sense our calculations should be understood as measuring the wedges that are implied by requiring that all countries have the same values for γ , α , and θ .⁸

⁸Note that no assumption is required regarding the relative value of the TFP parameter A , either across time or across countries. Similarly, the analysis does not require any assumption regarding the values of the discount factor β , or the depreciation rate δ , since they do not enter the static first order condition.

With all parameters constant across countries, one can treat a cross-section of observations across countries analogously to how we previously treated observations across time for a given country. In particular, we do this exercise for the 2004 cross-section and normalize the wedge factor for the US to be equal to zero. Figure 8 shows the pairs of wedges and log deviation of hours from the US value for the 21 countries in our sample.

[Insert Figure 8. Hours and Wedges Relative to US (2004)]

Once again, there is a striking negative correlation between the values. The correlation is equal to -0.82 . This calculation can be repeated for each year of the sample. The results are summarized in Figure 9, which simply reports the time series of the correlation between hours worked and wedges for each cross-section.

[Insert Figure 9. Correlation of Hours and Wedges]

As Figure 9 shows, there is nothing special about the 2004 cross-section. At all points in time, there is a strong negative correlation between hours of work and the wedge ranging between -0.77 and -0.89 .

4.3 The Role of Consumption/Output

The fact that our preceding analysis finds large wedges implies that there are large differences in hours that the model does not account for, both in the time series and the cross-section. While it is revealing that there is much that the model does not account for, it is also of interest to learn how much of the variation in hours worked *is* accounted for by the model. It is useful to first note that our finding of a strong negative correlation between hours and wedges does not necessarily imply that the model accounts for little of the variation in hours. To see why, consider an example in which the model accounts for half of all changes in hours worked. In this case one would still find large wedges and they would still be strongly negatively correlated with hours, even though the model accounts

for half of the variation in hours. In this section we show that reality is quite far from this hypothetical example. Specifically, we show that not only is the model unable to account for the variation in hours worked, but that in many cases it actually predicts differences in hours of the opposite sign from that found in the data.

To begin our analysis, note that if $\gamma = 1$, the expression for the static first-order condition can be rearranged to yield:

$$H_t = \frac{\frac{(1-\theta)\alpha}{1-\alpha}}{\frac{(1-\theta)\alpha}{1-\alpha} + \frac{C_t}{Y_t}} \bar{H} \quad (11)$$

This condition makes it clear that the model predicts changes in H_t only to the extent that there are changes in C_t/Y_t , with the implied changes being negatively correlated. Assessing the extent to which differences in hours of work (either over time for a given country or across countries at a point in time) are correlated with differences in C/Y is a natural first step to address the extent to which the model can account for the changes in hours worked.

From the perspective of accounting for the change in hours worked in a given country between 1956 and 2004, the relevant calculation is the change in C/Y for that country between those two years. Table 4 shows the change in C/Y for each country in our sample over this time period.

[Insert Table 4. Changes in C/Y : 1956–2004]

Two points emerge. First, many of the changes are negative, while in order to account for a decrease in hours one would require an increase in C/Y . Second, the correlation of these changes with changes in hours worked is 0.12. The low correlation is evidenced by Figure 10, which is a scatter plot for the changes in C/Y and the changes in (log) hours worked between 1956 and 2004.

[Insert Figure 10. Changes in C/Y and Log Hours]

Similar findings emerge when we examine cross-sectional differences in the value of C/Y at a given point in time. For example, in 2004, the correlation between differences in log hours and C/Y relative to the US is 0.23, and the scatter plot of values is shown in Figure 11.

[Insert Figure 11. C/Y and Hours in 2004]

Note that many counties have a C/Y ratio that differs substantially from that in the US, but since the difference is negative the benchmark model predicts that this should be associated with larger values for hours worked rather than smaller, so that differences in C/Y actually contribute negatively to accounting for differences in hours of work. Incidentally, in the 2004 cross-section the US has the highest value of C/Y .

The preceding analysis has examined the relationship between H and C/Y at a qualitative level. The main conclusion that emerges is that from the perspective of the benchmark model, the differences in C/Y either over time for a given country or across countries at a given point in time are typically of the wrong sign in order to account for the difference in hours of work.

Nonetheless, it is also of interest to ask what the model implies for the quantitative relation between C/Y and hours of work. To assess the quantitative implications of equation (11) we need to specify values for α and θ . We assume $\alpha = 0.37$, $\theta = 1/3$ and an initial value of C/Y equal to 0.8. These values would be consistent with an investment to output ratio of 0.2 and time devoted to market work equal to $1/3$ of the time endowment in the steady state of the benchmark model. Changes in C/Y to 0.7 or 0.9 yield changes in H of 9.1% and -7.7% respectively. These results are affected very little by reasonable changes in α , θ , or the initial value of C/Y .

It follows that in order to explain a drop of hours worked of 30 percent as was observed in several countries, one would need to observe an increase in C/Y of nearly 0.4. Moreover,

in order to explain differences in hours worked in the 2004 cross-section one would need the countries with low hours of work to have C/Y ratios that are as much as 0.5 higher than in the countries with high hours of work.

To summarize, the main finding of this subsection is that the magnitude of differences in C/Y found either in the time series or the cross-section are not large enough to be a quantitatively significant factor in accounting for the variation in hours found in the data. In addition, qualitatively the differences in C/Y move in the direction opposite to that required to account for the differences in hours.

5 Sensitivity

In this section we examine the sensitivity of the previous findings along three dimensions. First, we ask to what extent the results are affected by the value of the labor supply elasticity parameter γ . Second, we ask to what extent the results are affected if we allow for the possibility that government consumption is a substitute for private consumption. Lastly, we examine how the results are affected by considering Stone-Geary preferences that include a subsistence consumption term.

5.1 Sensitivity to Labor Supply Elasticity

Our previous calculations were for the case in which we assigned the labor supply elasticity parameter, γ , equals 1 (log utility over leisure). In this subsection we examine how our previous results are affected by the value of this parameter. The key observation is that subject to normalizing the wedge to equal zero in a particular year, different choices of γ serve simply to either increase or decrease all wedges in absolute value. While the change is not exactly proportionate, it is sufficiently close to proportionate that none of our earlier conclusions is significantly affected. Specifically, our findings regarding relative magnitudes of changes and correlations all continue to hold.

To see the effect of changes in γ recall that Δ is defined by:

$$1 - \Delta_t = \frac{(1 - \alpha)}{\alpha(1 - \theta)} \frac{H_t}{(\bar{H} - H_t)^\gamma} \frac{C_t}{Y_t} \quad (12)$$

For simplicity, assume that there is no variation in C/Y . Letting $t = b$ denote the base year, and following our earlier procedure that normalizes the wedge to be zero in period b , the wedge in period t is given by:

$$1 - \Delta_t = \frac{H_t}{H_b} \left(\frac{\bar{H} - H_b}{\bar{H} - H_t} \right)^\gamma \quad (13)$$

It follows that changing the value of γ changes the value of the wedge solely through its effect on the second term on the right-hand side. This term, however, has the property that it is strictly increasing in H_t and is equal to zero when $H_t = H_b$, independently of γ . Hence, the parameter γ affects only the slope of the relation between H_t and Δ_t .

To illustrate the effect of changes in γ on the implied series for Δ , Figure 12 shows the implied time series for the wedge for France, assuming values of γ equal to 0, 1, and 3.

[Insert Figure 12. Wedge and Labor Supply Elasticity]

As the figure indicates, in all three cases the series exhibits the same upward sloping behavior. The only effect is to alter the effect of the slope. As γ increases the increase in the wedge between 1956 and 2004 increases, being equal to 0.47, 0.60 and 0.80 for the cases of $\gamma = 0, 1$ and 3 respectively.

5.2 Government Consumption

Our benchmark calculations assumed that only private consumption entered into the household's utility function. All of the economies that we are studying, however, have significant government expenditure on consumption, and for most countries this has increased over time. Our analyses thus far are consistent with two different assumptions about government spending on goods and services. One is that households do not attach any value to the goods and services provided by the government. The second is that households do

value these goods and services but that they enter the utility function in a manner that is separable with respect to private consumption and leisure. In contrast, if government consumption is at least partially a substitute for private consumption, our findings could be affected. Given that governments provide many different types of goods and services, some of which are at least partial substitutes for private consumption, and that the particular activities in which they are involved vary considerably both across time and across countries, it is likely to be quite challenging to account for government consumption in the most appropriate way.

However, we can gauge the extent to which incorporating government expenditures into the analysis will affect our results by examining the polar extreme case to that previously considered, i.e., we now consider the case in which households consider all government expenditures on goods and services to be a perfect substitute for private consumption. To do this, we adjust our previous calculations by replacing C/Y with $(C + G)/Y$, where G is government expenditures on goods and services. This latter specification is clearly an overestimate of how much households value government spending, as certain categories of government spending (e.g., military expenditures) provide little substitution for private consumption.

A very simple finding emerges when we do this exercise. The effect of this change is to reduce the time series change in the average wedge by a small amount, and to leave all of the previous properties unchanged. This is well-illustrated by looking at the time series for the mean value of the wedge. Figure 13 shows the two mean wedge series based on whether government expenditures on goods and services are included.

[Insert Figure 13. Wedge and Government Consumption]

This figure shows that the two series move together very closely, with the only difference being that the original series starts at a slightly lower value.

This change also does not affect our results concerning the extent to which differences in C/Y can account for differences in hours of work, either across time within a country, or across countries at a point in time. The main effect is that differences in $(C + G)/Y$ are much less than differences in C/Y both in the cross-section and the time series, but the correlations of differences in $(C + G)/Y$ and H remain weakly positive in both cases. In fact, the reason that the wedges decrease when one uses $(C + G)/Y$ is that by decreasing the differences in this ratio we are decreasing the impact of a factor that for the most part goes in the wrong direction.

One may conjecture that very different results might be obtained if the nature of government consumption changes over time and differs across countries (in other words, the extent to which it substitutes for private consumption varies). While such a possibility certainly increases the potential for variation in $(C + G)/Y$ to play a greater role, we report two simple calculations that suggest that it is somewhat unlikely for this role to be quantitatively important.

The first calculation is concerned with the time series changes within countries. Consider a country such as France, which exhibits one of the larger decreases in hours of work between 1956 and 2004, equal to 35%. In order to be consistent with this decrease, the model would require that C/Y increases during this period. An extreme lower bound for C/Y in 1956 for France would be the ratio of private consumption to GDP. And an extreme upper bound for C/Y in 2004 would be the ratio of private plus government consumption to GDP. But even with this extreme case the increase in C/Y is only from 0.58 to 0.76. With the values of α and θ used earlier, the predicted decrease in H is only 13%.

The second calculation carries out a similar exercise in the context of the 2004 cross-section. Again, an extreme lower bound on C/Y for the US is the ratio of private consumption to GDP, while an extreme upper bound for other countries is the ratio of private plus government consumption to GDP. This lower bound for the US is equal to 0.65, while the extreme upper bound is less than 0.8 except for three countries: Greece, the UK, and

Portugal. Again, even extreme assumptions do not suggest that variation in C/Y plays a large role. These calculations show that government spending does not eliminate the large wedges in the first-order condition governing labor supply.

5.3 Subsistence Consumption

One of the features in the data is that almost all countries experience a decrease in hours between 1956 and 1985. There is substantial evidence to indicate that the process of development is initially associated with a decrease in the workweek, after which it appears to level off.⁹ Stone-Geary preferences of the form:

$$U(C_t, \bar{H} - H_t) = \alpha \log(C_t - \bar{C}) + (1 - \alpha) \frac{(\bar{H} - H_t)^{1-\gamma} - 1}{1 - \gamma} \quad (14)$$

can account for this phenomenon. With preferences of this form, a country that starts from a low initial level of productivity and experiences ongoing technological change that is labor augmenting will asymptotically converge to a balanced growth path in which hours of work are constant, but during the early part of the transition to this asymptotic balanced growth path the economy will experience a decrease in hours of work. It is only in the limit as C_t becomes large compared to \bar{C} that the balanced growth path is achieved. In this regard it is noteworthy that the four countries that do not experience any significant drop in hours worked in the early part of the time period being studied, Australia, Canada, New Zealand and the US, are all among the most productive countries in the sample as of the beginning of the period.¹⁰ In view of this we explore the possibility that the dynamics associated with subsistence consumption might be relevant in accounting for some part of the decrease in hours of work during our sample period.

⁹See for example, Maddison (1995) for estimates of changes in the annual hours of work per person in employment.

¹⁰Rogerson (2006a) used this observation to argue that differences in productivity in the early part of the time period studied are a potentially important factor in accounting for differences in hours worked across countries.

Repeating the previous analysis, we obtain the following expression for the wedge:

$$1 - \Delta_t = \frac{(1 - \alpha)}{\alpha(1 - \theta)} \frac{H_t}{(\bar{H} - H_t)^\gamma} \frac{C_t - \bar{C}}{Y_t}$$

In order to implement this extension one needs to assume a value of \bar{C} . We report results here that correspond to a value of \bar{C} equal to 10% of total US consumption in 1956. For these results we assume that the measure of consumption is the sum of private consumption plus government expenditure on goods and services. The results were quite similar when \bar{C} was set to 5% of total US consumption in 1956. We begin by showing the effect that this has on the mean wedge series, as compared to our original calculations. Figure 14 shows the results for the trend components of the two series for average wedges.

[Insert Figure 14. Effect of Subsistence on Δ]

Three features are worth noting. First, the two curves have the same qualitative features. They both exhibit relatively steady increases up until the mid 1980s, and are relatively flat thereafter. Second, the magnitude of the change is significantly larger for the case without subsistence consumption. Third, in the post 1980 period, the two curves are virtually identical. This last finding is perhaps not too surprising—since the effect of the subsistence term diminishes as consumption increases, we expect that the effect of subsistence will decrease over the course of the time period being studied.

Once again, although this extension to the benchmark model does reduce the magnitude of the wedge, all of the properties that we previously established continue to hold: differences in hours worked are associated with large differences in wedges, both in the time series and in the cross-section. As one additional illustration of this, in Figure 15 we plot the pairs of changes in log hours and changes in wedges between 1956 and 2004 for the 21 countries in our sample.

[Insert Figure 15. Change in Hours and Wedges with Subsistence]

As before, this plot shows a strong negative correlation between the two variables.

5.4 Summary

Having explored sensitivity along three dimensions, and presented a substantial amount of additional information, we summarize the main message from these exercises. Although there are variations/extensions that serve to reduce the magnitude of the wedges that we found earlier, we find that the standard growth model is unable to account for a large part of the changes in hours worked during the period 1956-2004.

6 The Role of Taxes

A simple conclusion emerges from the preceding analysis. Any extension to the standard growth model that is going to successfully account for the large changes in hours of work over time and across countries must be capable of producing a “static” first order condition that in equilibrium contains a term like the wedge that we introduced earlier. It is relatively well-known that various factors can give rise to such wedges, including distorting taxation, product market regulation, non-competitive wage setting and labor market regulation. In this section we extend the original model to include a government that taxes labor income and uses the proceeds to both purchase goods and services and fund a transfer program. Our motivation for focusing on taxes is two-fold. First, prior work by Prescott (2004) and Rogerson (2005) has argued that taxation may be a quantitatively important factor in accounting for changes in hours worked over time for a small set of countries, and Davis and Henrekson (2004) argue that taxes help to account for differences in hours worked in a cross-section of countries. Second, although there are issues associated with measuring tax rates, from the list of factors just offered, taxes may be the factor that is easiest to measure, making them a natural first choice.

In this section we therefore repeat the earlier analysis to see how incorporating taxes will affect the findings. We first modify the theoretical framework to include taxes into the wedge calculation. We then graphically present the wedges from the model with taxes, and

compare them to the wedges without taxes. We also conduct a complementary analysis that graphically compares hours as predicted by the model to actual hours. We finally examine particular episodes/countries that are puzzling from the perspective of the theory and show how our views on these puzzles are shaped by incorporating taxes into the analysis.

6.1 Model

The only change that we make to the earlier model is to add a government sector that places a proportional tax of τ_{ht} on labor income and τ_{ct} on consumption expenditures. Government spending on goods and services in period t is given by G_t , and the remaining tax revenues are used to finance a lump sum transfer T_t so as to balance the budget each period.

We modify preferences to explicitly allow for the possibility that consumers value the goods and services purchased by the government. In particular, we assume a one period utility function of the form:

$$U(C_t, \bar{H} - H_t, G_t) = \alpha \log(C_t + \lambda G_t - \bar{C}) + (1 - \alpha) \frac{(\bar{H} - H_t)^{1-\gamma} - 1}{1 - \gamma} \quad (15)$$

While the assumption that private and public consumption are additive is somewhat restrictive, as previously noted, consideration of the two extreme cases of $\lambda = 0$ and $\lambda = 1$ will be sufficient for us to bound the importance of G for the analysis. While we incorporate the subsistence term \bar{C} in the preferences, we will present results for the case in which this value is set to zero as well as to a positive value.

Solving for a competitive equilibrium in this case, and following the same steps as before, one can derive the following modified version of the “static” first order condition:

$$\frac{U_2(C_t, \bar{H} - H_t, G_t)}{U_1(C_t, \bar{H} - H_t, G_t)} = \frac{(1 - \tau_{ht})}{(1 + \tau_{ct})} F_2(K_t, H_t, A_t) \quad (16)$$

To simplify notation it is convenient to define a single tax term τ_t by combining labor and consumption taxes as follows:

$$(1 - \tau_t) \equiv \frac{(1 - \tau_{ht})}{(1 + \tau_{ct})}. \quad (17)$$

Given our functional forms, this equation now reduces to:

$$\frac{(1 - \alpha)}{\alpha} \frac{H_t}{(\bar{H} - H_t)^\gamma} = (1 - \tau_t)(1 - \theta) \frac{Y_t}{(C_t + \lambda G_t - \bar{C})} \quad (18)$$

and if we introduce the wedge as before we now have:

$$\frac{(1 - \alpha)}{\alpha} \frac{H_t}{(\bar{H} - H_t)^\gamma} = (1 - \tilde{\Delta}_t)(1 - \tau_t)(1 - \theta) \frac{Y_t}{(C_t + \lambda G_t - \bar{C})} \quad (19)$$

where $\tilde{\Delta}_t$ denotes the after-tax wedge.

6.2 Results with Taxes

We now quantitatively evaluate how distorting taxation influences the size of the wedge over time and across the countries. Again we consider the $\gamma = 1$ case as our benchmark. In order to compute the wedge we need series for both labor and consumption taxes, plus a value for λ . We use series for average taxes from McDaniel (2006), who extends the procedure of Prescott (2004) to a larger set of countries and a longer time period.¹¹ We note that unlike Prescott, we do not make any adjustment to turn average tax rates into marginal tax rates. Due to data limitations we do not have tax series for all countries, so we will only present results for 15 countries in what follows.¹² Consistent with our earlier results, it turns out that the value of λ matters very little for most of the countries in our sample. For the results presented below we assume that $\lambda = 1$.

We begin by showing how the inclusion of taxes affects the mean series for wedges, first for the case where $\bar{C} = 0$, and next for the case in which \bar{C} is set to 10 percent of US

¹¹Following Mendoza et al (1994), Carey and Rabesona (2005) have produced series for average labor tax rates for a large set of countries from 1970 on. The series from McDaniel produce very similar values for differences in tax rates across time and countries for the overlapping time periods and countries. We use the McDaniel series because there are large changes in hours in the early part of the period that is not covered by the other available series. Our findings are virtually identical when we use the other tax series for the larger sample of countries for the smaller time period.

¹²The countries not included are Denmark, Greece, Ireland, New Zealand, Norway, and Portugal. We include Australia, although the tax series for this country begins in 1960 rather than 1956.

consumption in 1956. Figure 16a shows the average wedge series for the first case, and for comparison includes the comparable series derived assuming no taxes.¹³

[Insert Figure 16a. Mean Wedge and Taxes]

The contribution of taxes is striking. Whereas the original wedge calculated without taxes exhibited a large positive trend, the new series displays virtually no trend. It follows that this model is able to account for the fact that average hours of work have decreased about 20%.

Figure 16b reports the mean after-tax wedge calculated under Stone-Geary preferences, i.e., the case in which \bar{C} is positive. The new series not only eliminate the large positive trend from mean wedges, but actually induces a downward trend. We will return to this point later when we analyze group patterns.

[Insert Figure 16b. Mean Wedge and Taxes (Stone-Geary)]

Next we examine the patterns at a more disaggregated level. We begin by displaying the mean wedge series for Groups 1-3 as before, suitably modified due to the sample restriction. Again, we present figures both for the case of $\bar{C} = 0$ and \bar{C} equal to ten percent of US consumption in 1956. These are contained in Figure 17a and 17b.

[Insert Figure 17. Mean Wedges, Groups 1-3]

Several interesting patterns emerge. We begin by discussing Figure 17a, which does not allow for subsistence consumption. The average wedge for the Group 1 countries still displays a period of steady increase, but this period of an increasing wedge ends in the early 1970s, as opposed to the mid 1980s without taxes. Moreover, while in our initial analysis the mean wedge increased by about 0.5 for this group, in this figure the increase is only

¹³The no-tax series in this figure is also computed for the 15 country sample, and so is not identical to that depicted earlier.

about 0.15. For Group 2, the profile has switched from having a substantial upward trend without taxes to being relatively flat with taxes. The implication is that once the model is augmented to allow for taxes, it is now able to explain the bulk of the decrease in hours in these countries. Finally, the series for Group 3 shows a substantial downward trend, and one that accelerates around 1980. Overall the decrease is about 0.15. In contrast, the wedge series for this group from the original analysis was relatively flat. For this group, the issue becomes how to explain why hours worked have not fallen given the changes in taxes.

When subsistence consumption is incorporated, the results are further modified. Perhaps not surprisingly given our earlier findings on the effects of adding subsistence consumption, this modification induces a negative trend in the wedge series. The overall result is that the wedge series for Group 1 is now relatively flat while the trend series for Groups 2 and 3 display downward trends.

[Insert Figure 17b. Mean Wedges, Groups 1-3 (Stone-Geary)]

Based on this analysis, we conclude that once one allows for subsistence consumption, taxes are able to account for the bulk of the change in hours worked in the countries belonging to Group 1. For Groups 2 and 3 the puzzle is to account for the fact that hours of work did not decrease by more than they did in the data.

6.3 Predicted and Actual Hours

Thus far we have presented our results by reporting properties of the wedge series. Another useful calculation is to solve for the series of hours of work for each country that is consistent with no wedge given observed values for C , G , Y and taxes. Comparing this series to the actual series for hours of work provides some perspective on the extent to which the model is able to account for the data. The advantage of this approach is that it is perhaps easier to interpret a deviation between actual hours and model hours than it is to interpret the significance of a wedge of a given magnitude.

Because our wedges are all normalized to be zero in a particular year, when we do the calculation for hours we also need to normalize model hours in a particular year. In what follows we have normalized model hours so that hours worked in 1956 are the same as in the data. For simplicity, we present a few representative countries to illustrate our results, namely the Netherlands, Germany, Sweden and the US. Figures 18-21 report the actual and predicted series for these countries, for the specification which includes subsistence consumption.

[Insert Figure 18. Actual and Predicted: Netherlands]

[Insert Figure 19. Actual and Predicted: Germany]

[Insert Figure 20. Actual and Predicted: Sweden]

[Insert Figure 21. Actual and Predicted: US]

As noted, we chose these four countries because they illustrate the variety of outcomes. The Netherlands is a case in which actual and predicted hours are very similar. We obtain this result for Finland as well. Note, in particular, that in the case of the Netherlands the model is able to capture the non-monotonic behavior of hours. Germany is a case in which the model is able to account for a large share of the decline in hours worked, but not all of it. Other countries fitting this pattern are Austria, Belgium, France, Italy, Japan and UK. Switzerland and Spain fall also in this category, but while for the former the model replicates qualitatively the non-monotonic behavior displayed by hours worked, for the latter this is not so. We find Sweden to be a particularly interesting case in that the model predicts a larger drop than what we observe in the data. Finally, the US is a country which shows little trend in hours of work, but for which the model predicts that there should be a small decrease. Other countries fitting this pattern are Australia and Canada.

6.4 Puzzling Episodes

One of the useful products of our analysis is a comprehensive analysis of which time periods in which countries are “puzzling” from the perspective of the growth model that incorporates various features. Simply put, if the change in the wedge for a particular country in a given time period is “large”, then this represents an episode which is puzzling in the context of the specification in which the wedge is derived. Since the growth model is the standard benchmark model in macroeconomics, this information should be very useful in helping to focus future research efforts in the area.

In this spirit, we contrast how incorporating taxes alters one’s views about which episodes are puzzling, and what the nature of the puzzle is. In particular, we compute changes in wedges over five time periods for our sample of 15 countries, comparing pre-tax changes in the wedge to after-tax changes (both with and without a subsistence term). The five time periods are 1956-1964, 1965-1974, 1975-1984, 1985-1994 and 1995-2003. In order to summarize this information, Table 5 shows the frequency count for the distribution of changes in wedges across the time periods for the three different cases.

[Insert Table 5. Distribution of Changes in Wedges by Decade]

As a rough guide to interpreting the sizes of these changes in wedges, we note that with $\gamma = 1$ a 10 percent increase in the wedge corresponds to roughly a gap between actual and predicted hours between 6 and 7 percent. Given that business cycle fluctuations may involve swings in hours on the order of 6 – 8 percent from peak to trough, we consider a criterion in which we isolate those episodes experiencing changes in the wedge exceeding 0.125 in absolute value as reasonable. When taxes are not included in the analysis, the number of such episodes is 25, and all but two of them are cases in which the wedge increases. Not surprisingly, most episodes are concentrated in countries belonging to Group 1, with France and Belgium experiencing large changes in three out of five decades. The two negative changes in the wedge, on the other side, refer to Spain and Netherlands during

the last decade. Again, this finding does not come as a surprise given that hours in these two countries display remarkable non-monotonicities.

After including taxes in the analysis, we obtain a different picture. When we do not consider subsistence consumption, the number of episodes with large changes is only 14, with 9 of these displaying an increase in the wedge. More generally, Table 5 indicates that the distribution of changes in wedges is more concentrated around 0, and is much less skewed toward positive values. When we consider after-tax wedge changes with Stone-Geary preferences, instead, we obtain 13 episodes with changes larger than 0.125 in absolute value. However, in 9 out of 13 cases we measure a large *decrease* in the wedge: Spain alone accounts for 4 of these episodes (and 1 with positive changes) while Italy accounts for 2 of them.

7 Statistical Analysis of Other Factors

There are many institutional, policy, and regulatory factors other than taxes that are typically thought to influence the determination of hours of work. Incorporating many of them into the analysis requires substantial extensions of the model, and do not produce simple analytical expressions. We therefore leave this task for future work.

However, we felt that a simple statistical analysis might be valuable in giving a crude sense of the possible importance of these other factors. The idea of this analysis is to check which factors might be highly correlated with the wedges that are produced from our benchmark analysis. Loosely speaking, any factor that is not highly correlated with the wedges is probably not likely to be of great importance in accounting for changes in hours worked over time. Our strategy is to perform panel regressions to investigate the importance of taxes and indices of labor market institutions in explaining the evolution in wedges. We stress that our exercise is not intended to provide estimates of deep parameters, but simply some *prima facie* evidence regarding the potential importance of various factors.

With respect to the institutional variables, we refer the reader to the Labour Market

Institutions Database constructed by Nickell and Nunziata (2001) and collect the following variables:

1. EP = index for Employment Protection
2. UDNET = Net Union Density
3. CO and COW = measures of bargaining coordination
4. BRR = Benefit Replacement Ratio
5. BD = Benefit Duration

These variables have been used extensively in the literature to capture institutional differences that are of potential importance in accounting for labor market outcomes. Our specification is

$$\log(1 - \Delta_{it}) = a_i + b \log(1 - \tau_{it}) + \gamma' X_{it} + \varepsilon_{it}$$

where Δ_{it} is the time series of wedges, a_i is a country fixed effect, τ_{it} is the tax rate used in our analysis, and X_{it} includes the institutional regressors¹⁴. Data on the institutional variables restricts the sample period to the years 1975-1995 and our sample includes all 15 countries considered in our previous analysis.¹⁵ The wedge series that we use in these regressions corresponds to the values calculated when we include government consumption in our measure of household consumption, but assumes that there is no subsistence consumption.¹⁶ Table 6 reports our regression results.

¹⁴Including a country fixed effect a_i is consistent with the exercise in which we did not impose that the preference and technology parameters α and θ were the same across countries. We also ran the regressions with the variables in levels instead of logarithm. Results are not affected.

¹⁵Although the average hours series is relatively flat after the early 1980s, there is still a great deal of variation across countries during the period 1975-1995.

¹⁶We note that subsistence consumption is mostly relevant in the period 1956-1970, when consumption levels are lower.

[Insert Figure Table 6. Regression Results. Dependent Variable: Wedge]

There are three main findings. First, taxes are able to account for a significant fraction of the variation in wedges. Second, this finding is unaffected by the inclusion of any of the other factors, either individually or collectively. Third, although many of the other factors are statistically significant, they add relatively little explanatory power either individually or collectively.

While these regressions can only be interpreted as identifying correlations and partial correlations, we find it interesting that of the factors considered, taxes appear to be the most highly correlated.

8 Summary and Conclusion

We have used the neoclassical growth model to shed light on the large reduction in hours worked observed in 21 OECD countries over the last 50 years, and the large variation in the magnitude of this reduction across countries. The standard growth model without distortions is able to account for virtually none of the observed reduction in hours worked. In contrast, a model that allows for tax distortions as observed in the data is able to account for virtually all of the average reduction across countries. Looking at individual countries, we find that in some cases the model implies too large of a reduction in hours, while in others it implies too small of a reduction. Future work should focus on understanding what additional factors might account for these discrepancies between theory and data.

There are several avenues which we think are likely to be important. First, our analysis has abstracted from potentially important features of the tax and transfer schemes. For example, Ragan (2005) and Rogerson (2006b) argue that the effect of tax distortions in Scandinavian countries is partly undone by the fact that tax revenues are used to subsidize market activities such as child and elderly care. Second, our analysis has abstracted from home production. Recent work by Aguiar and Hurst (2005) and Francis and Ramey (2005)

shows that the increase in market work in the US over the last 40 years reflects a decrease in time spent in home production and not a reduction in leisure. Additionally, Freeman and Schettkat (2006) show that time devoted to home production in many European countries is higher than in the US. The lack of time series data on home production precludes one carrying out our exercise in a model that explicitly allows for home production, but assessing the impact of home production on our findings remains a relevant issue.¹⁷ Third, it is important to assess the role of other factors in accounting for the wedges that we have measured.

¹⁷Ingram et al (1994) show how to infer a time series for time spent in home production for a given country assuming that the wedge series is identically zero. This calculation is of interest, but cannot be used to determine the effect of adding home production on the size of the wedge.

9 Appendix : Data

Data on consumption, government consumption and GDP are taken from PWT.

Data on employment and population are from OECD's *Economic Outlook* and *Main Economic Indicators*. Series for hours worked are from Groningen Growth and Development Centre (GGDC) and The Conference Board. For population we consider working age population (between 15 and 64 years old), while hours is the total number of hours worked over the year divided by the average numbers of people in employment.

In our regressions, we use as measures of institutional variables the indices available from the Nickell-Nunziata Labor Market Institutional database (2001). We invite the reader to consult the original source for the exact details behind the construction of these variables.

References

- [1] Aguiar, M. and AE. Hurst, (2006), "Measuring Trends in Leisure: The Allocation of Time over Five Decades", *NBER Working Paper* 12082.
- [2] Bencivenga, Valerie R., (1992), "An Econometric Study of Hours and Output Variation with Preference Shocks", *International Economic Review*, 33(2), pp. 449-71.
- [3] Blanchard, O. (2005), "The Economic Future of Europe," *Journal of Economic Perspectives*.
- [4] Chari, V.V., Kehoe Patrick J. and McGrattan Ellen R., (2004), "Business Cycle Accounting", *NBER Working Paper* 10351.
- [5] Cole, Harold and Ohanian, Lee E. (2002), "The Great U.S. and U.K. Depressions through the Lens of Neoclassical Theory, *American Economic Review*.
- [6] Davis, S., and Henrekson, M. (2004), "Tax Effects on Work Activity, Industry Mix and Shadow Economy Size: Evidence from Rich Country Comparisons", *NBER Working Paper* #10509.
- [7] Francis. N. and V. Ramey (2005) "A Century of Work and Leisure," mimeo, UCSD, 2005.
- [8] Freeman, R., and Schettkat R., "Marketization of Household Production and the EU-US Gap in Work", *Economic Policy* 20 (2005), 6-50.
- [9] Gali', Jordi, Gertler, Mark and Lopez-Salido, J. David, (2002), "Mark-ups, Gaps, and the Welfare Cost of Business Fluctuations", *NBER Working Paper* 8850.
- [10] Hall, Robert E, (1997), "Macroeconomic Fluctuations and the Allocation of Time", *Journal of Labor Economics*, January, vol. 15, no. 1, pp. 223-250.

- [11] Ingram, Beth, Kocherlakota, Narayana and Savin, N.E., (1994), "Explaining Business Cycle: A Multiple Shock Approach", *Journal of Monetary Economics*, 34 (3), 415-428.
- [12] Maddison, A., *Monitoring the World Economy 1820-1992*, OECD, Paris, 1995.
- [13] McDaniel, Cara, (2006), "Effective Tax Rates for 15 OECD Countries: 1950-2003", mimeo, Arizona State University.
- [14] Mulligan, Casey B., (2002), "A Century of Labor-Leisure Distortions", *NBER Working Paper 8774*.
- [15] Nickell, S. and L. Nunziata (2001), "Labor Market Institutions Database", Center for Economic Performance.
- [16] Parkin, Michael, (1988), "A Method for Determining Whether Parameters in Aggregative Models are Structural", *Carnegie-Rochester Conference Series on Public Policy*, 29, pp.215-52.
- [17] Prescott, Edward C., (2004), "Why Do Americans Work So Much More than Europeans?", *Federal Reserve Bank of Minneapolis Quarterly Review*, Vol. 28, No. 1, July, pp. 2-13.
- [18] Ragan, K., "Taxes, Transfers and Time Use: Fiscal Policy in a Household Production Model," mimeo, University of Chicago, 2005.
- [19] Rogerson, Richard, (2005), "Structural Transformation and the Deterioration of European Labor Market Outcomes," mimeo, Arizona State University.
- [20] Rogerson, Richard (2006a), "Understanding Differences in Hours Worked," *Review of Economic Dynamics* 9, 365-409.
- [21] Rogerson, Richard, (2006b), "Taxation and Market Work: Is Scandinavia an Outlier?", forthcoming, *Economic Theory*.

Table 1. Hours Worked in 2004 Relative to 1956

Australia	0.97	Germany	0.60	Norway	0.82
Austria	0.74	Greece	0.99	Portugal	0.84
Belgium	0.70	Ireland	0.70	Spain	0.86
Canada	1.08	Italy	0.74	Sweden	0.80
Denmark	0.71	Japan	0.85	Switzerland	0.83
Finland	0.75	Netherlands	0.81	United Kingdom	0.79
France	0.67	New Zealand	1.05	United States	1.01

Table 2. Correlation Hours and Wedges: Trend Components

Australia	-0.88	Germany	-1.00	Norway	-0.99
Austria	-1.00	Greece	0.07	Portugal	-0.82
Belgium	-0.98	Ireland	-0.97	Spain	-0.99
Canada	-0.57	Italy	-0.98	Sweden	-0.99
Denmark	-0.99	Japan	-0.96	Switzerland	-0.98
Finland	-0.98	Netherlands	-0.98	UK	-0.95
France	-1.00	New Zealand	-0.79	US	-0.96

Table 3. Correlation Hours and Wedges: Cyclical Components

Australia	-0.80	Germany	-0.83	Norway	-0.75
Austria	-0.78	Greece	-0.39	Portugal	-0.74
Belgium	-0.76	Ireland	-0.72	Spain	-0.93
Canada	-0.88	Italy	-0.82	Sweden	-0.88
Denmark	-0.70	Japan	-0.65	Switzerland	-0.76
Finland	-0.81	Netherlands	-0.83	UK	-0.90
France	-0.81	New Zealand	-0.92	US	-0.94

Table 4. Changes in C/Y : 1956–2004

Australia	-0.04	Germany	0.09	Norway	-0.19
Austria	-0.03	Greece	0.01	Portugal	-0.20
Belgium	-0.11	Ireland	-0.35	Spain	-0.09
Canada	-0.08	Italy	0.15	Sweden	-0.21
Denmark	-0.24	Japan	-0.11	Switzerland	-0.03
Finland	-0.10	Netherlands	0.00	UK	0.12
France	-0.04	New Zealand	-0.02	US	0.10

Table 5. Changes in Wedges by Decade

Change in Wedges	<i>Number of Observations</i>		
	Pre-Tax Wedges	After-Tax Wedges	
		Standard Pref.	Stone-Geary Pref.
$\leq -.125$	2	5	9
$(-.125, -.075]$	4	7	14
$(-.075, -.025)$	7	11	16
$[-.025, +.025]$	13	18	13
$(+.025, +.075)$	17	15	11
$[+.075, +.125)$	9	9	7
$\geq +.125$	23	9	4

Table 6. Regression Results. Dependent Variable:Wedge

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Tax	1.051 (0.053)	1.010 (0.054)	1.101 (0.054)	1.073 (0.053)	1.035 (0.053)	1.076 (0.061)	0.925 (0.057)	1.000 (0.064)
EP		-0.045 (0.016)						-0.043 (0.017)
UD			0.239 (0.065)					0.202 (0.066)
CO				0.075 (0.032)				0.070 (0.035)
COW					0.028 (0.013)			0.006 (0.014)
BRR						0.029 (0.036)		0.053 (0.037)
BD							-0.155 (0.029)	-0.139 (0.029)
R²	0.470	0.479	0.485	0.476	0.476	0.471	0.501	0.528

Figure 1. Mean Hours Worked

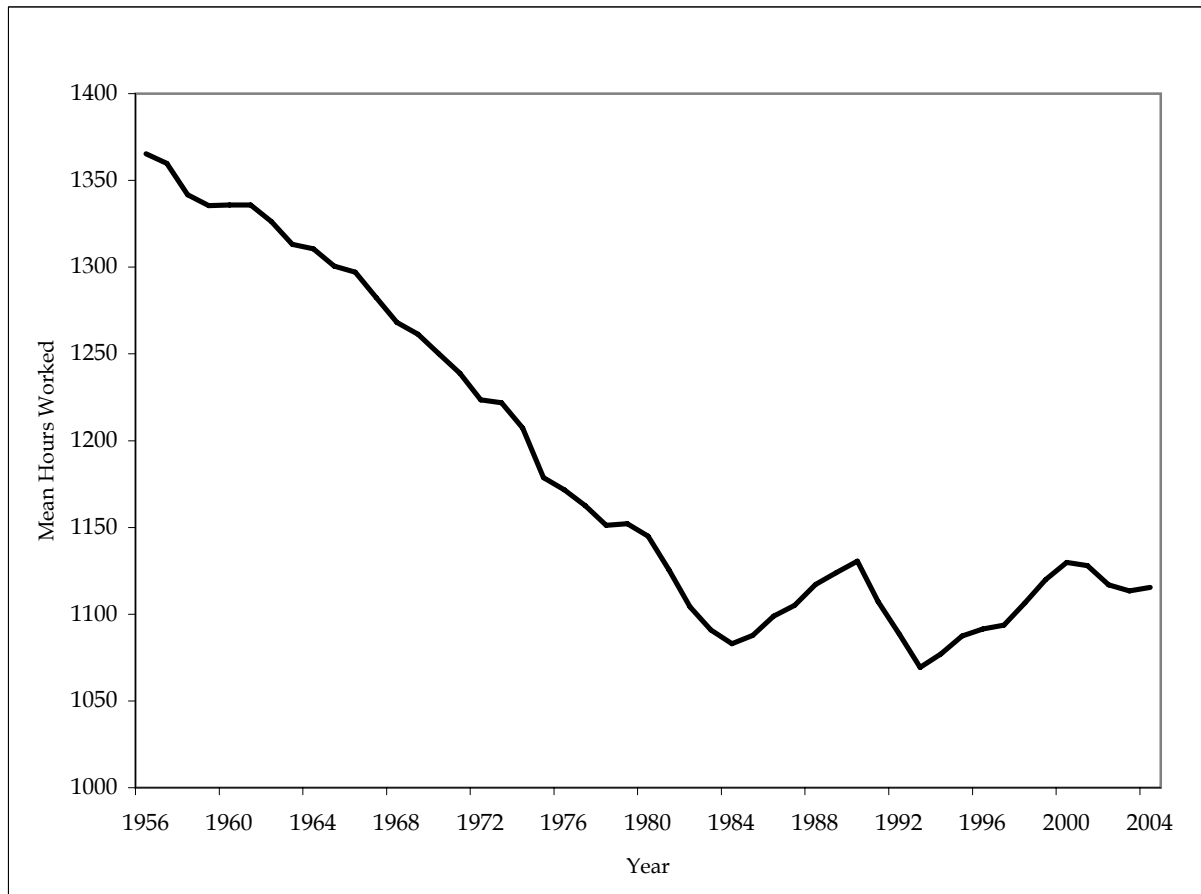


Figure 2. Mean Hours Worked by Group

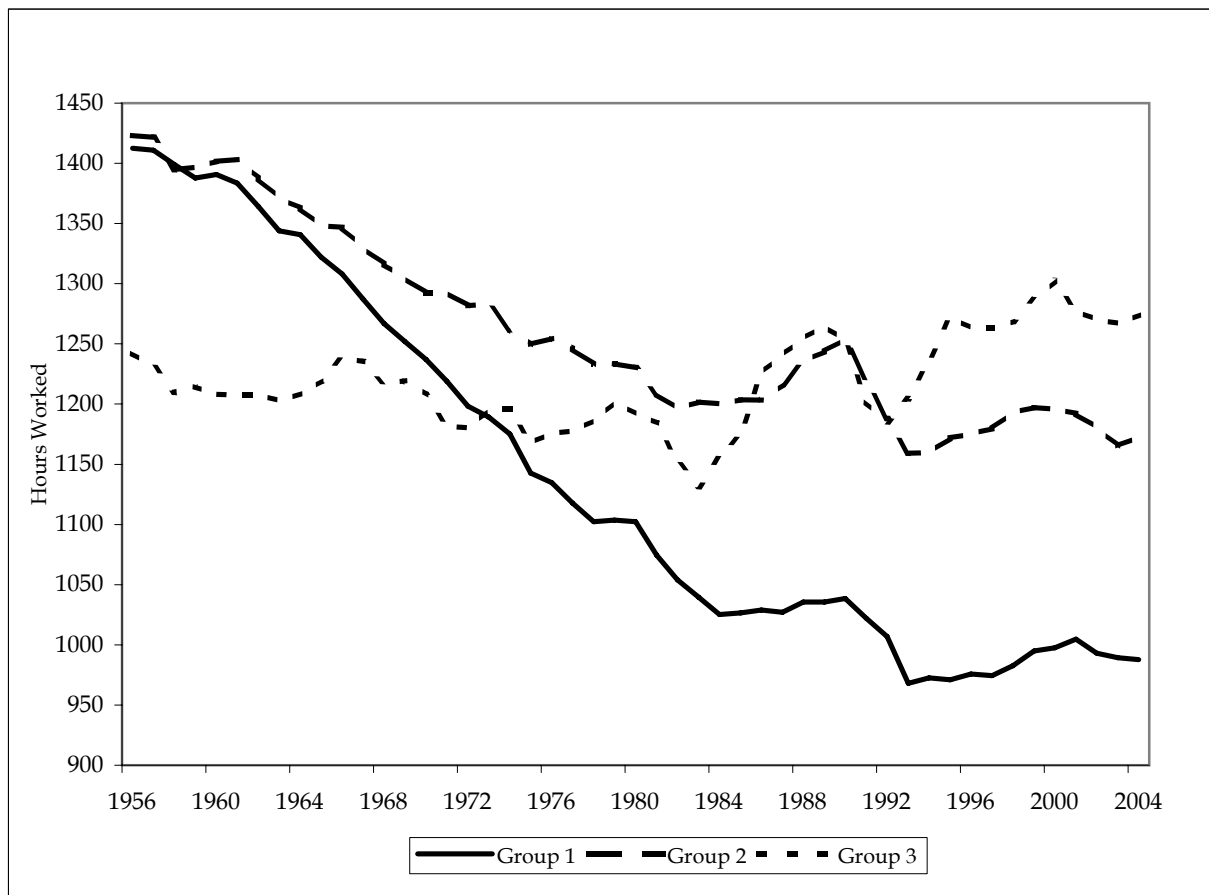


Figure 3. Mean Wedge and Mean Hours

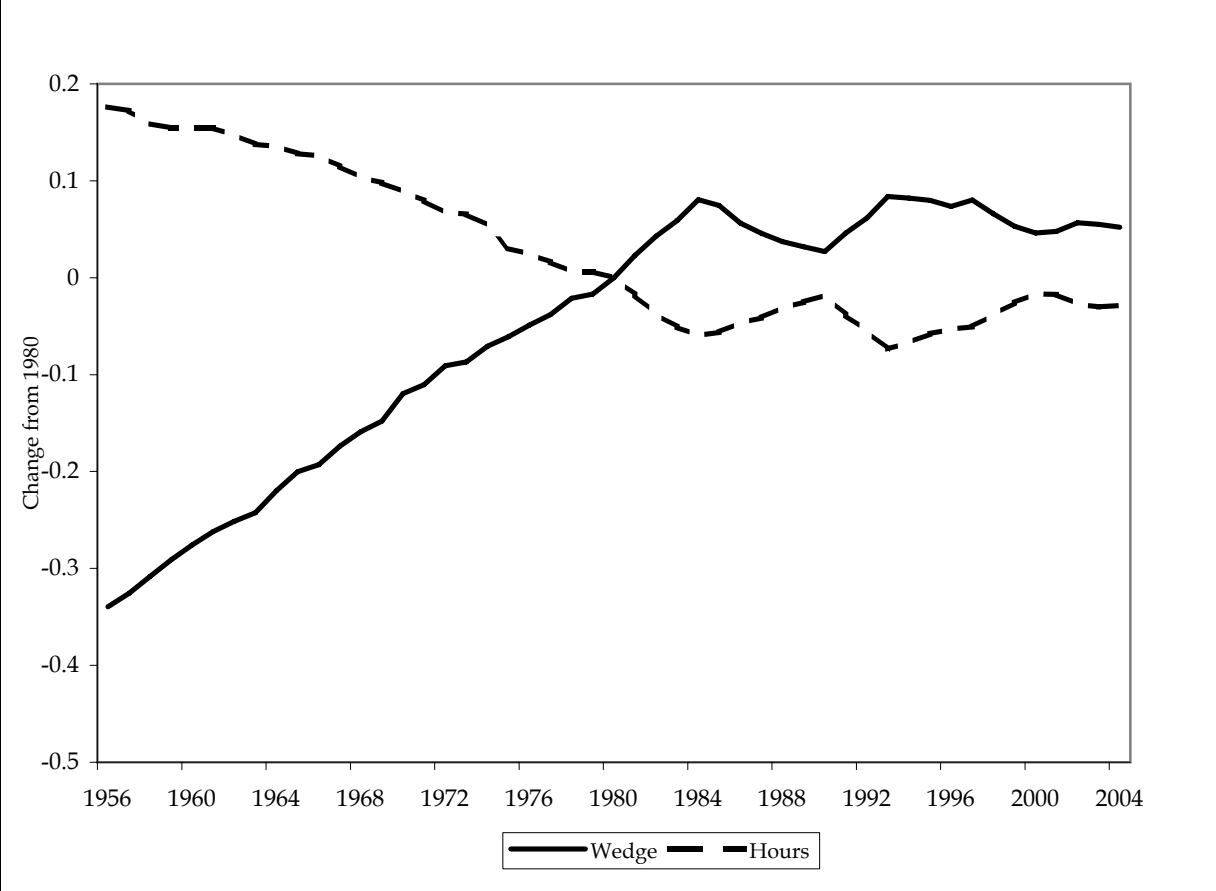


Figure 4. Mean Wedge by Group

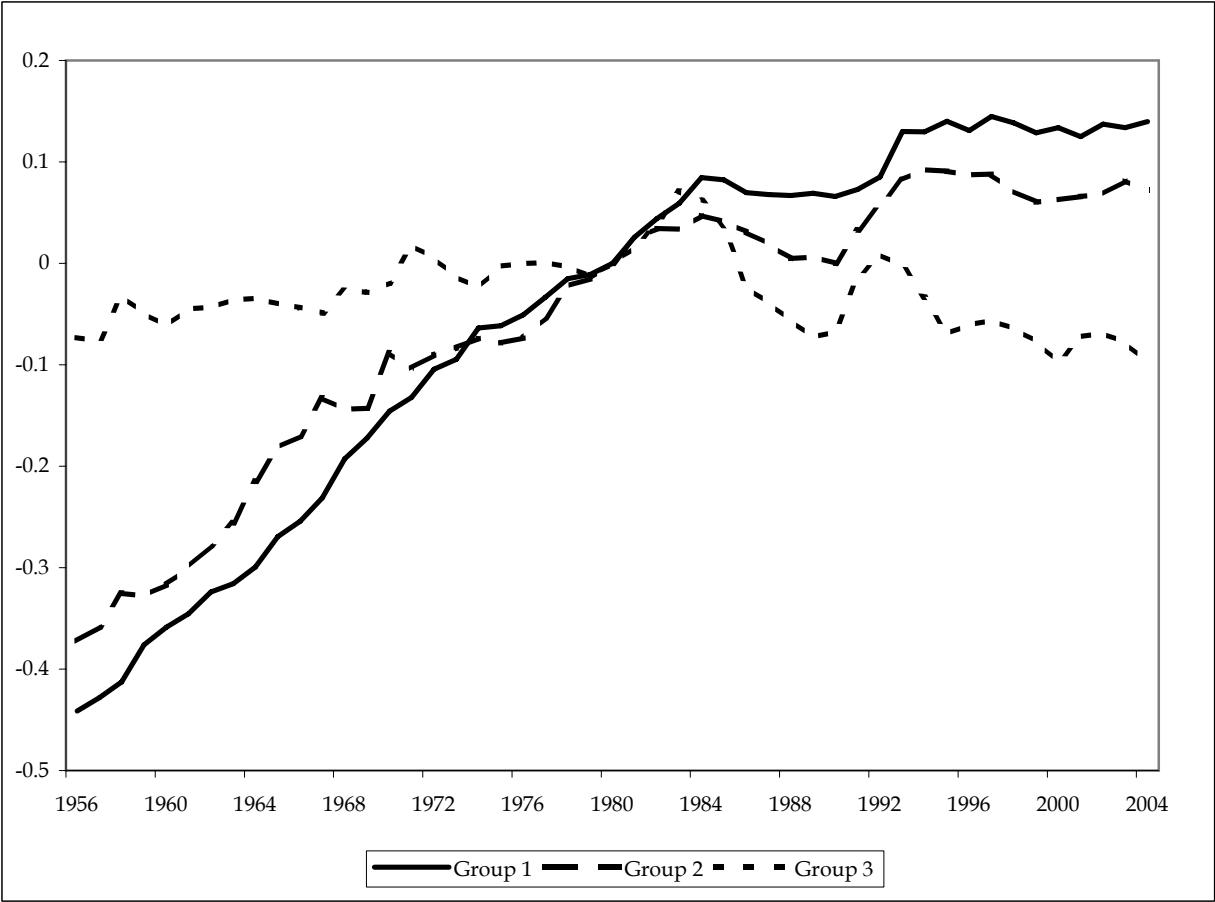


Figure 5. Change in Wedges and Hours

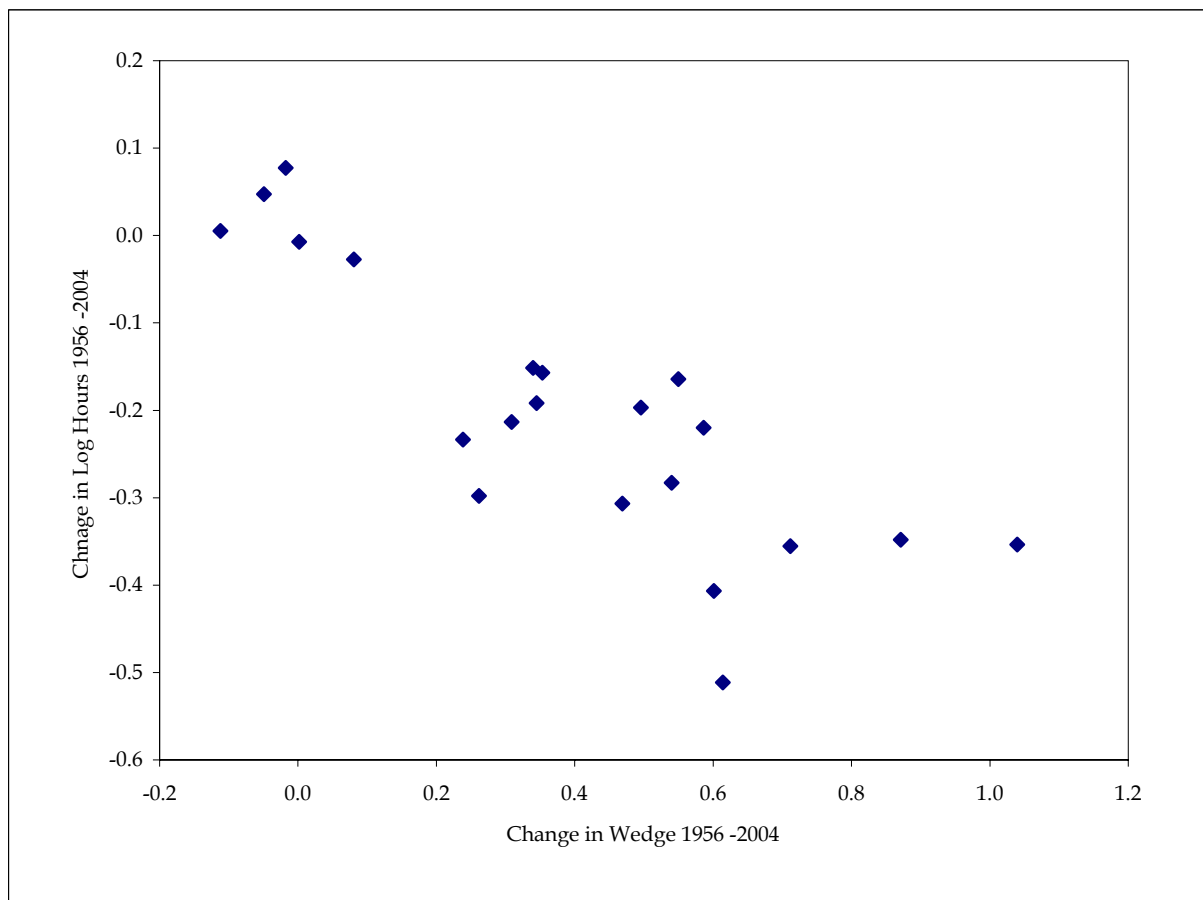


Figure 6. Wedge and Hours: Netherlands

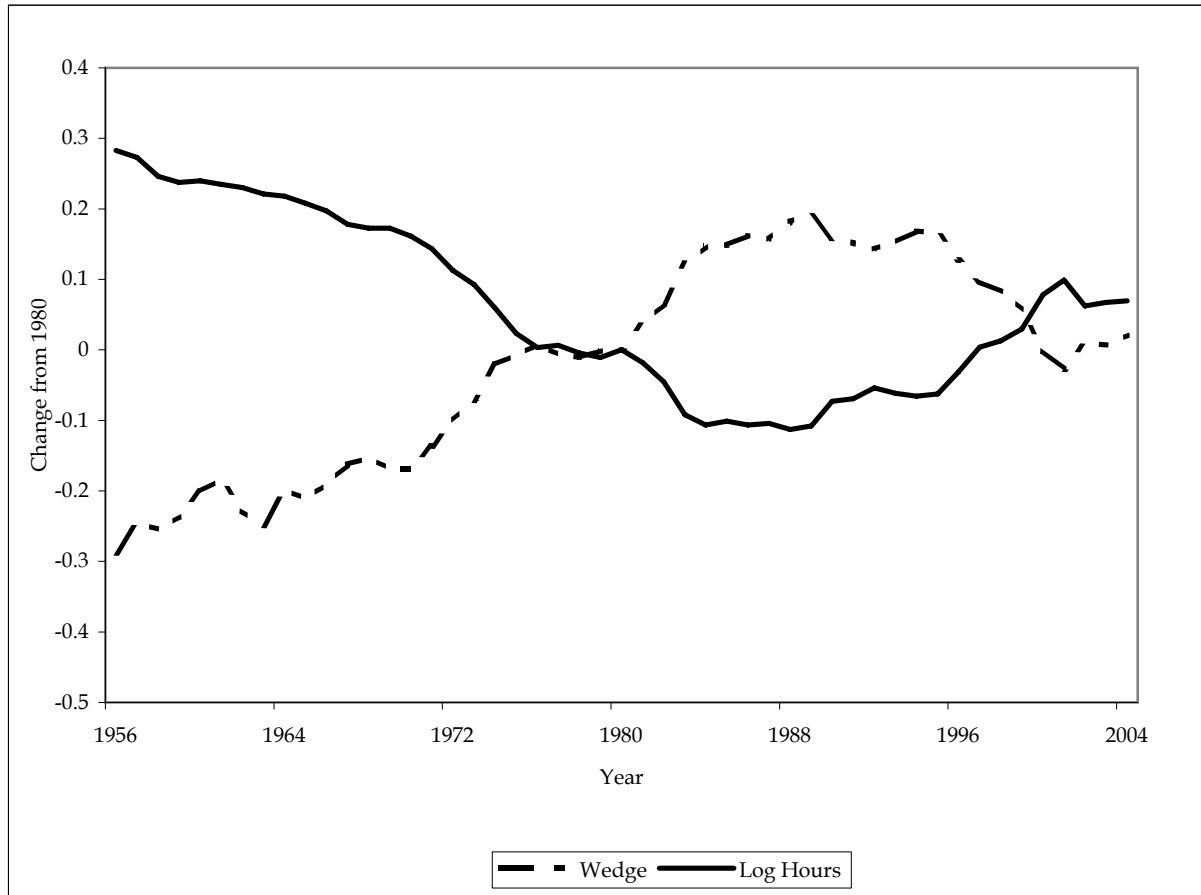


Figure 7. Wedge and Hours: Spain

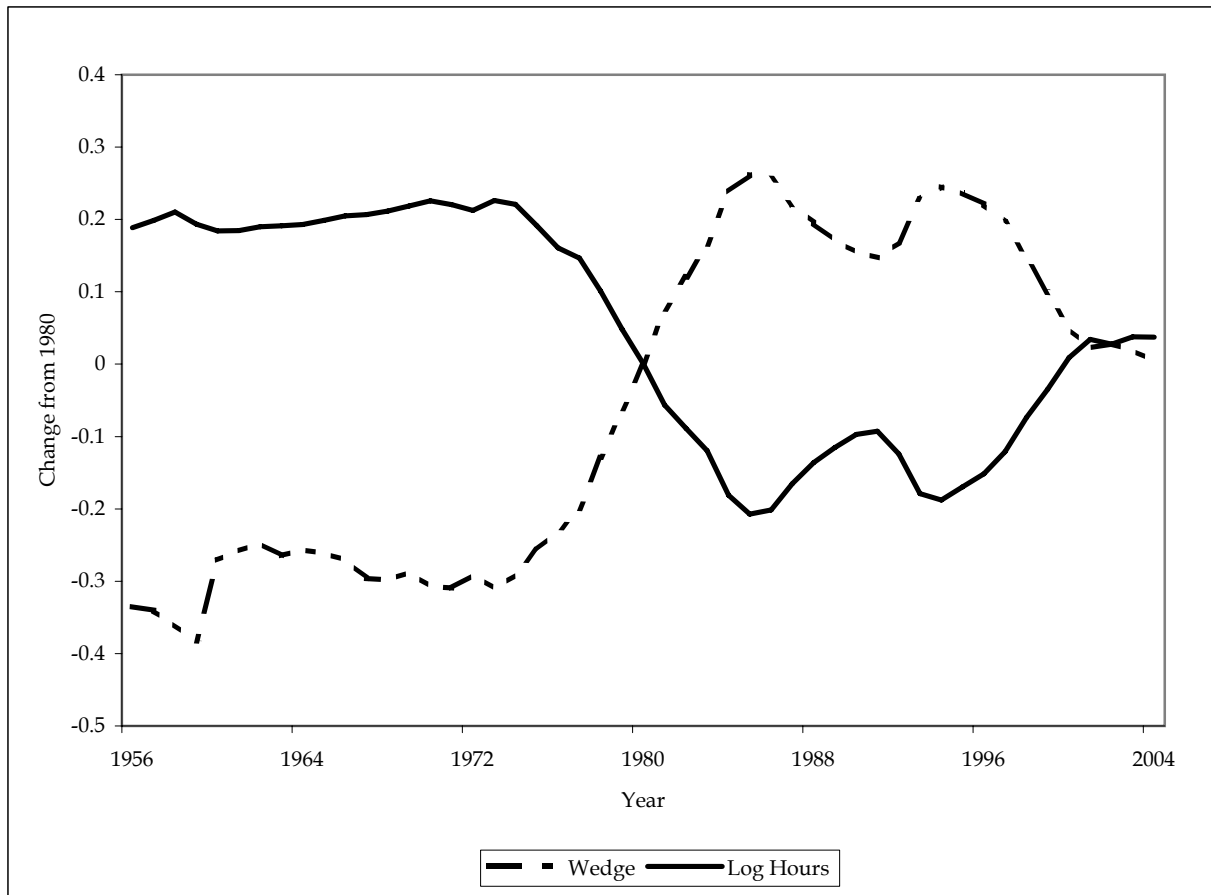


Figure 8. Hours and Wedge Relative to US (2004)

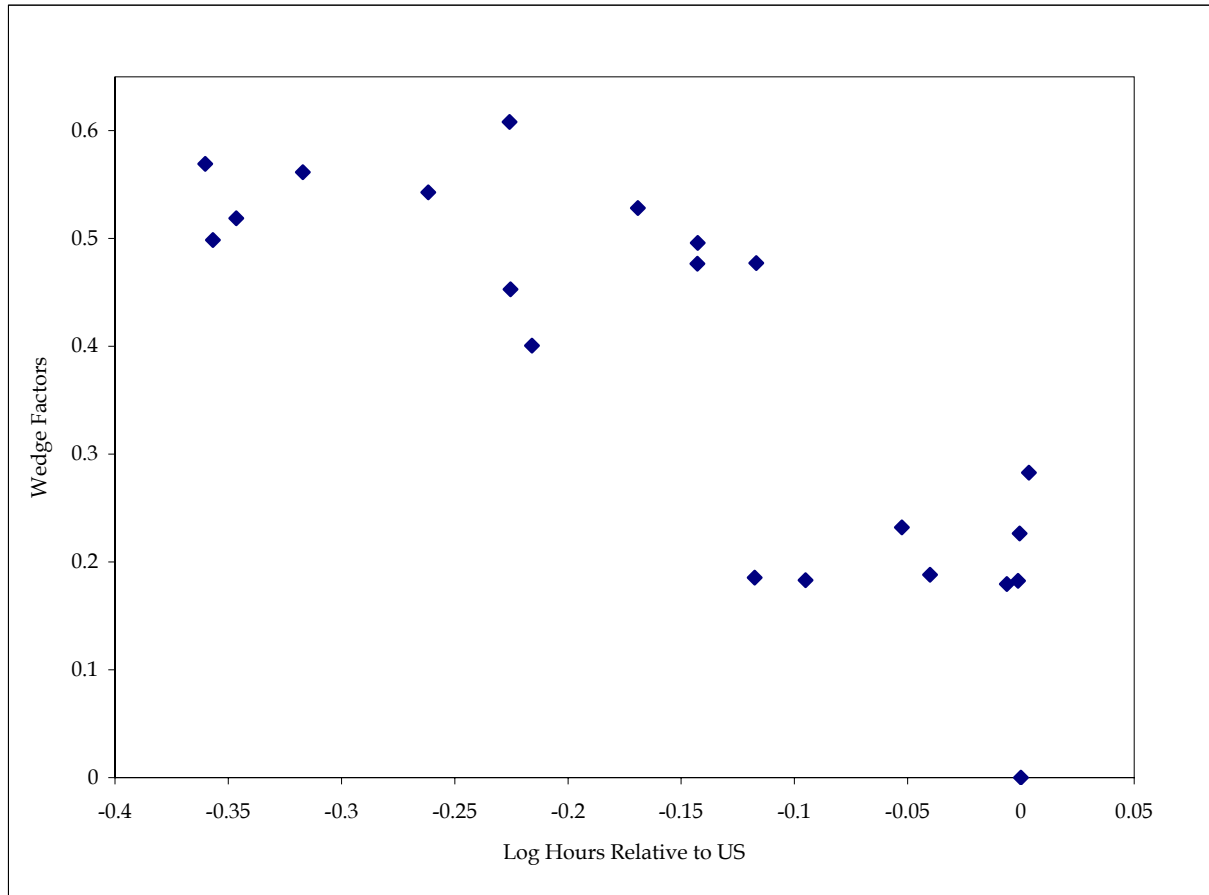


Figure 9. Correlation of Hours and Wedges

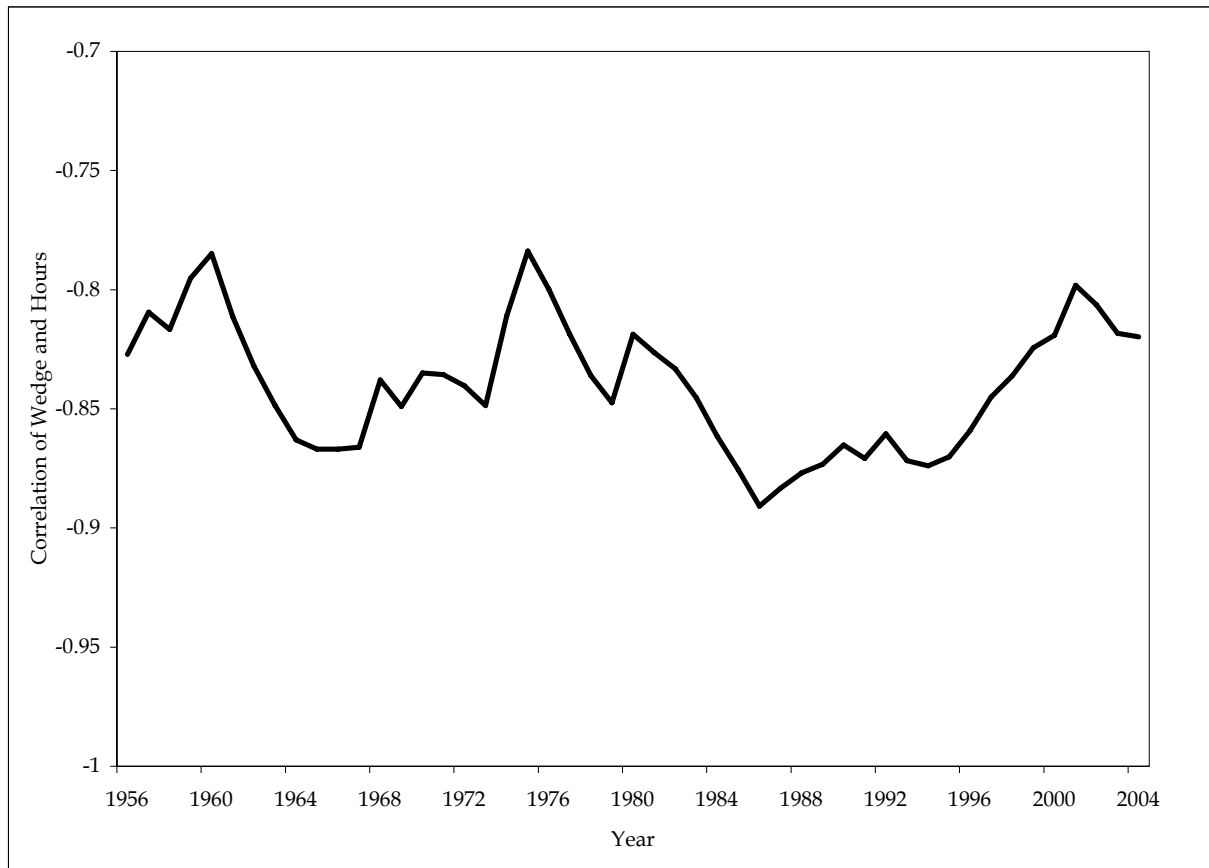


Figure 10. Changes in C/Y and Hours

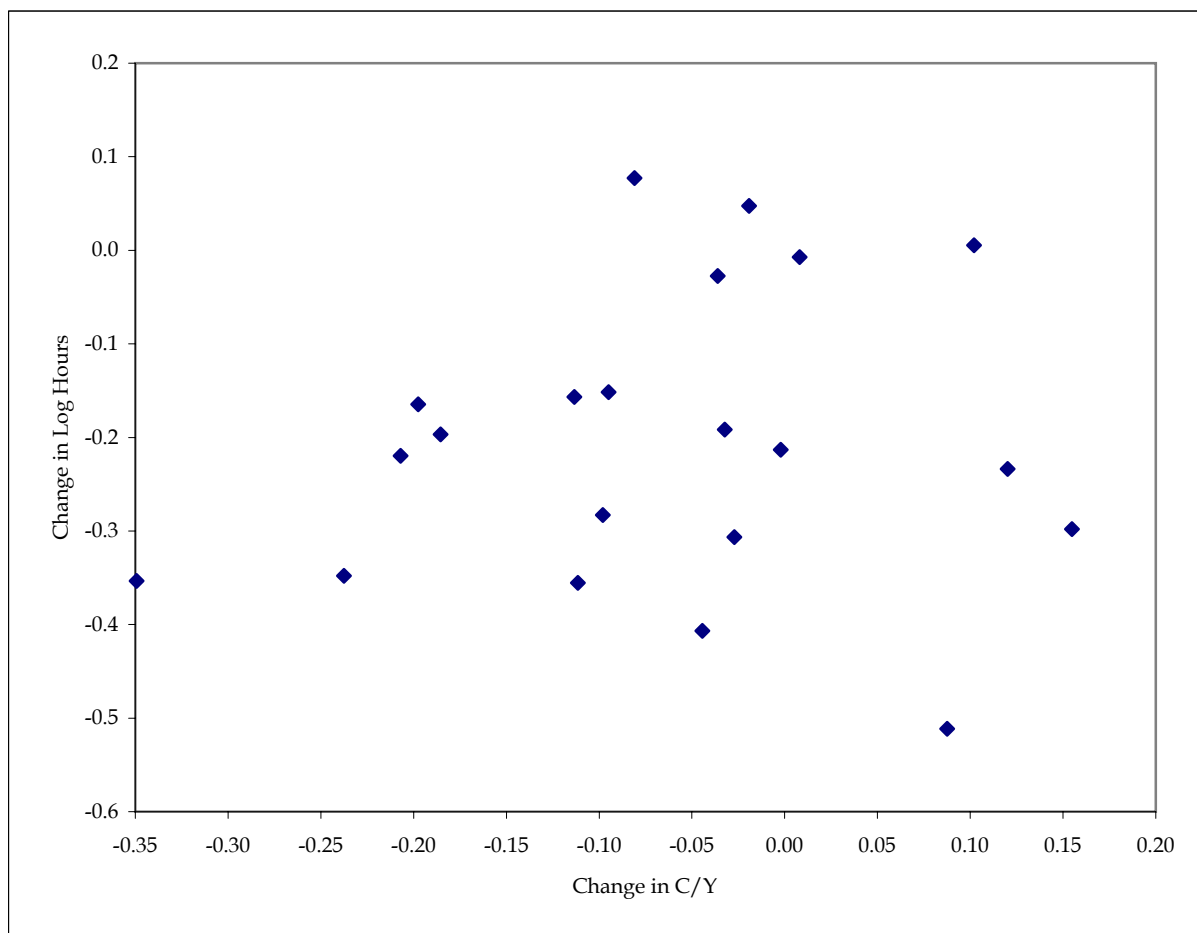


Figure 11. Changes in C/Y and Hours in 2004

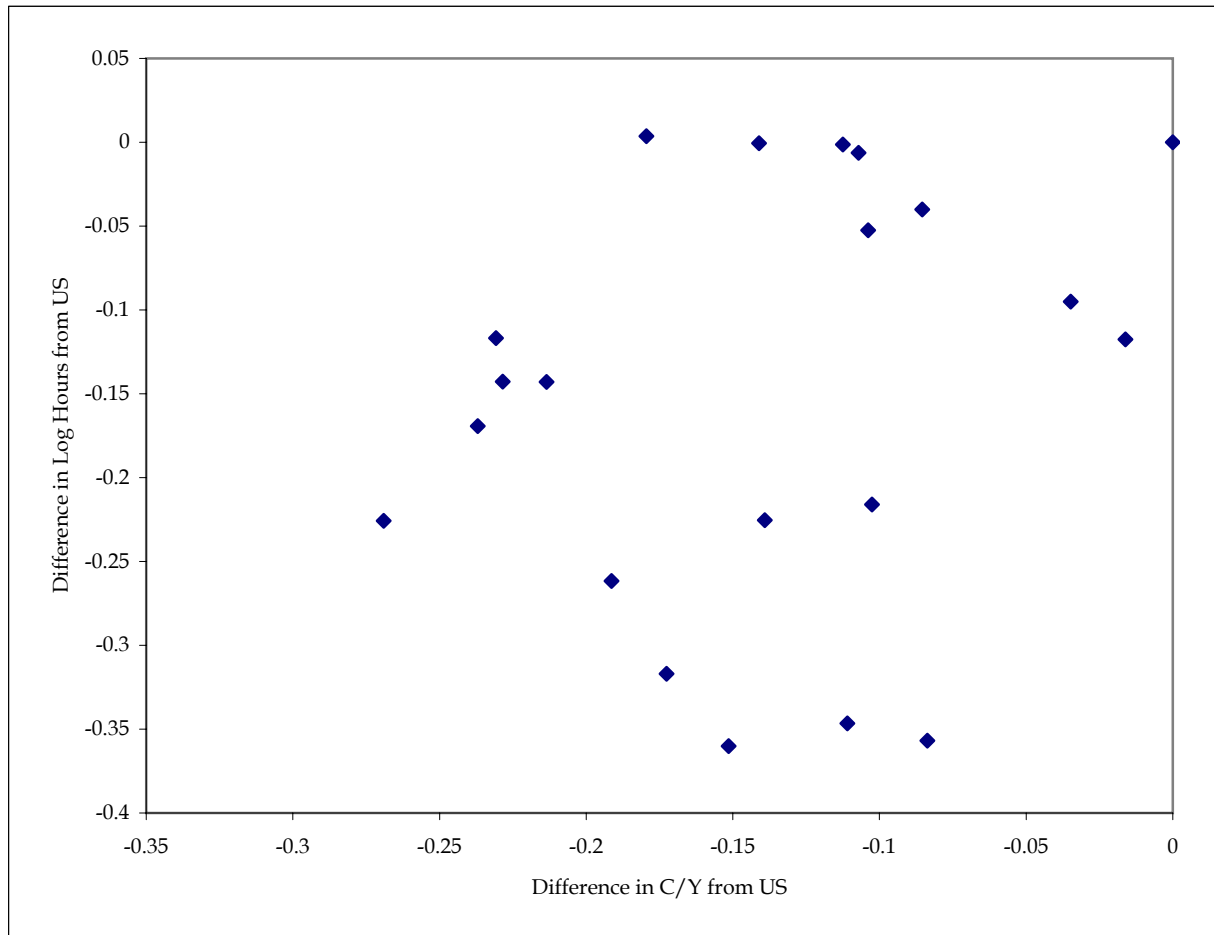


Figure 12. Wedge and Labor Supply Elasticity



Figure 13. Wedge and Government Consumption

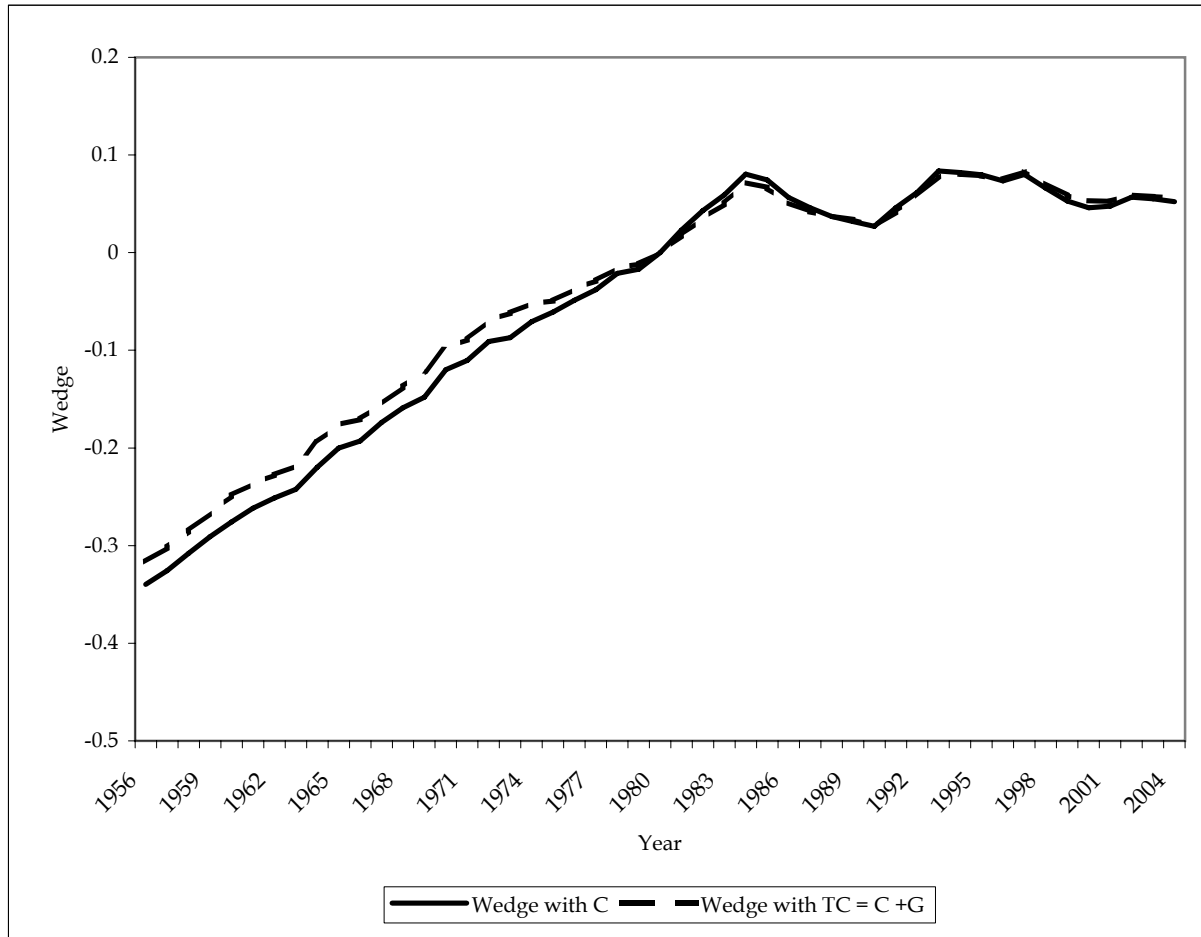


Figure 14. Effect of Subsistence on Wedge



Figure 15. Change in Hours and Wedges with Subsistence

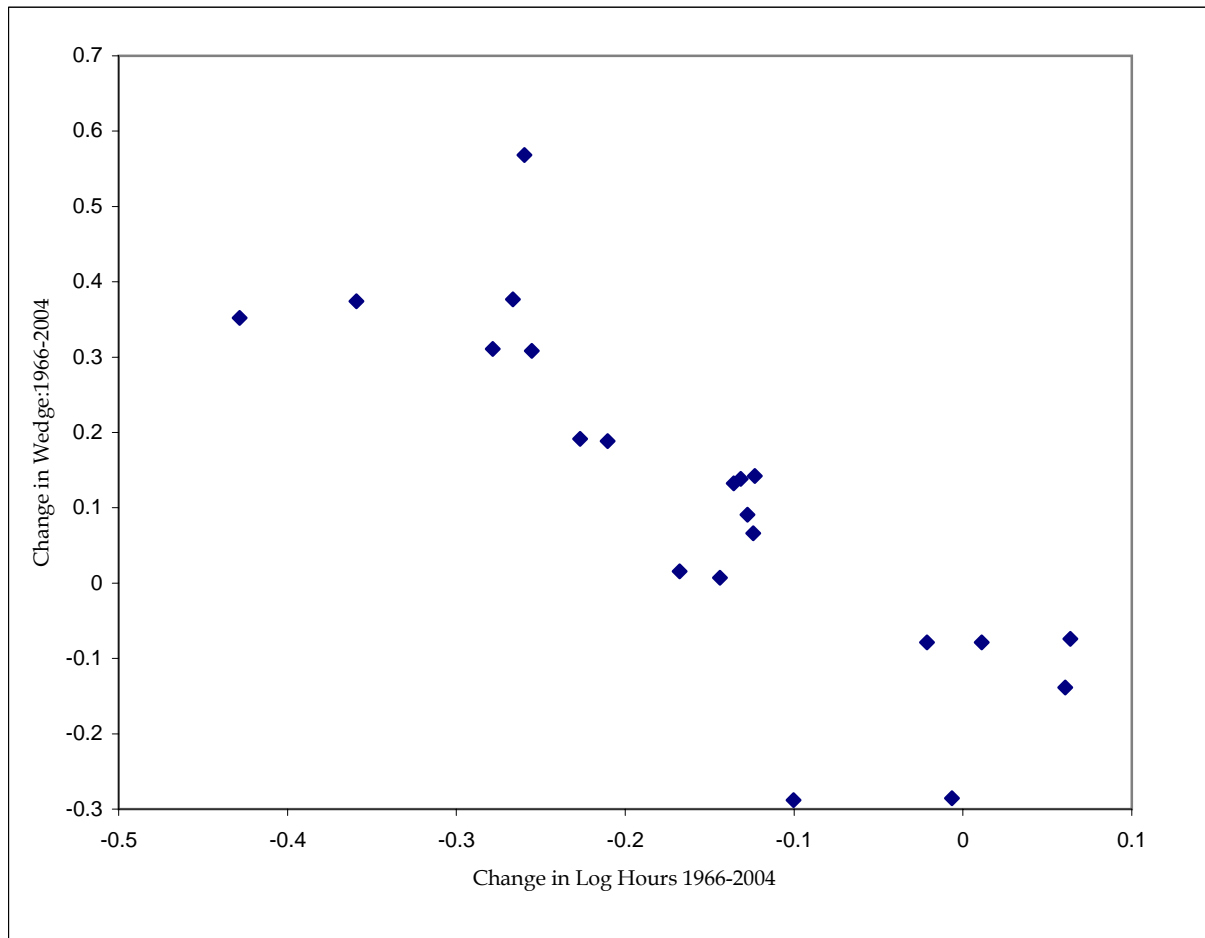


Figure 16a. Mean Wedge and Taxes



Figure 16b. Mean Wedge and Taxes (Stone Geary)

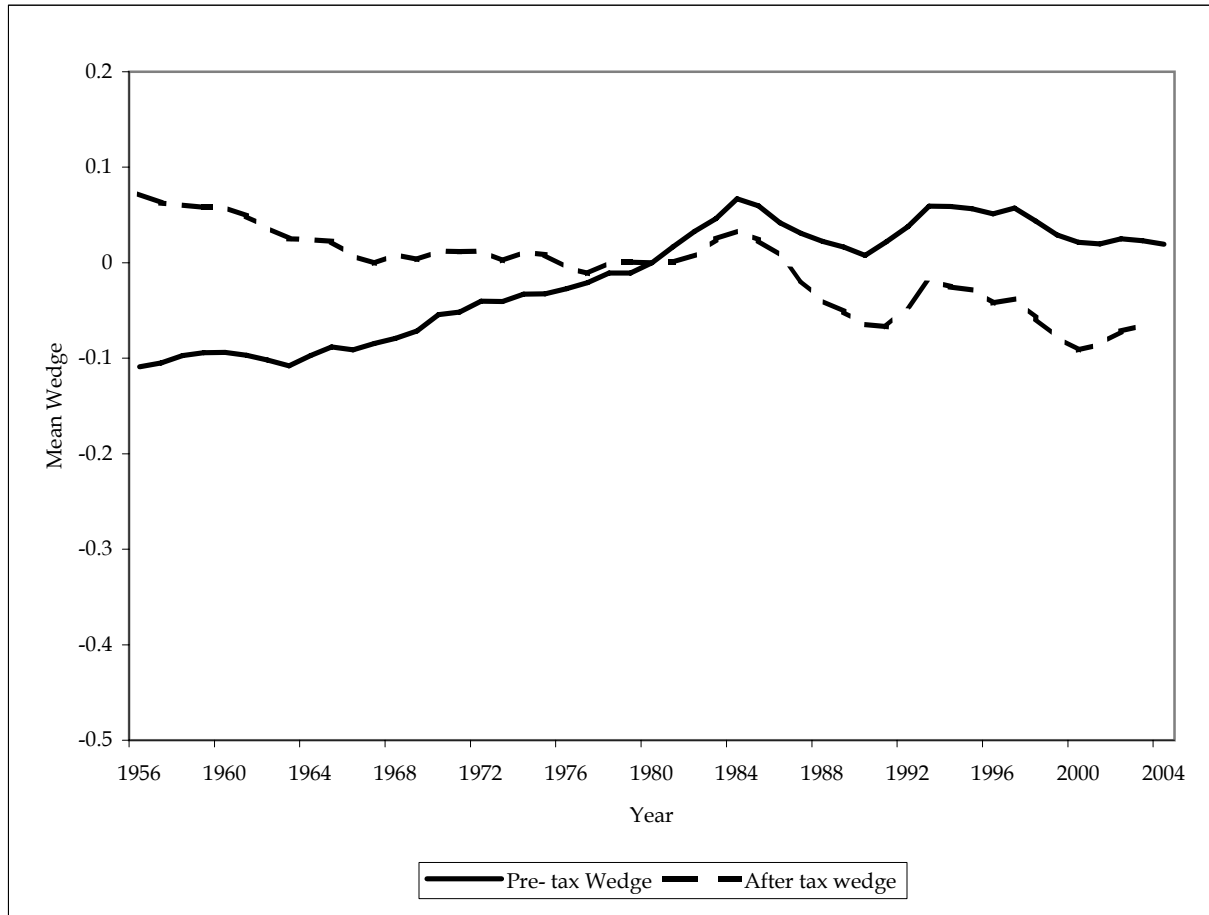


Figure 17a. After Tax Mean Wedges

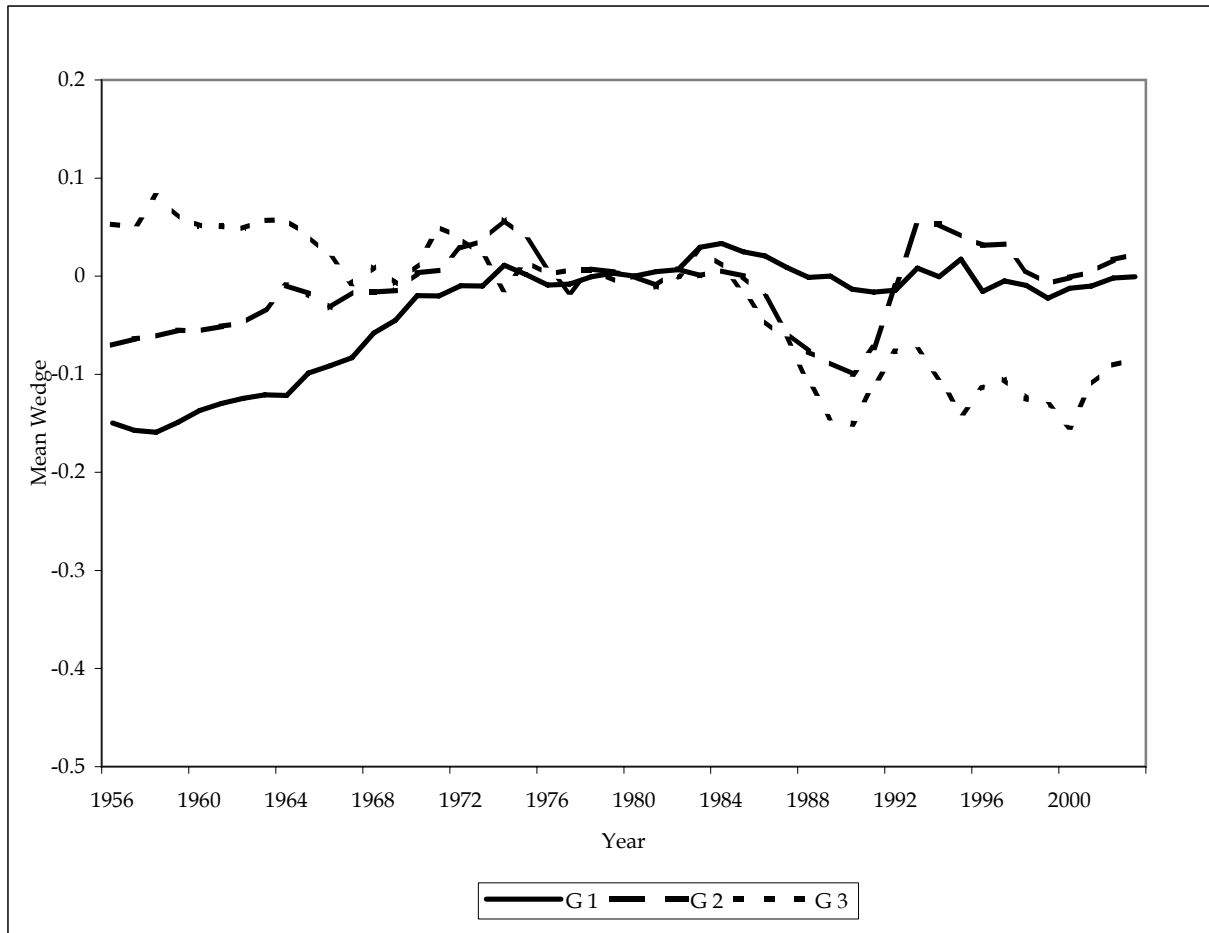


Figure 17b. After tax Mean Wedges (Stone-Geary)

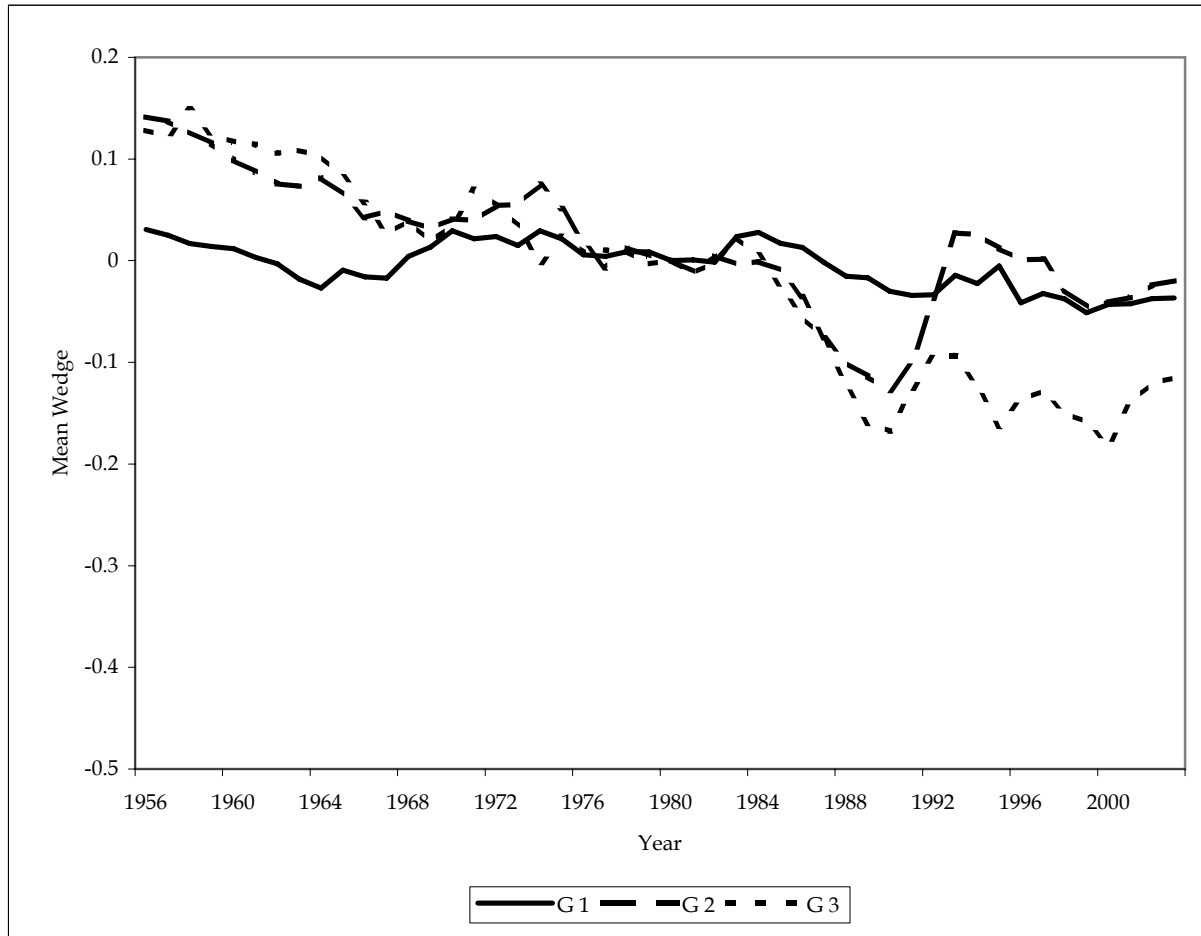


Figure 18. Actual and Predicted: Netherlands

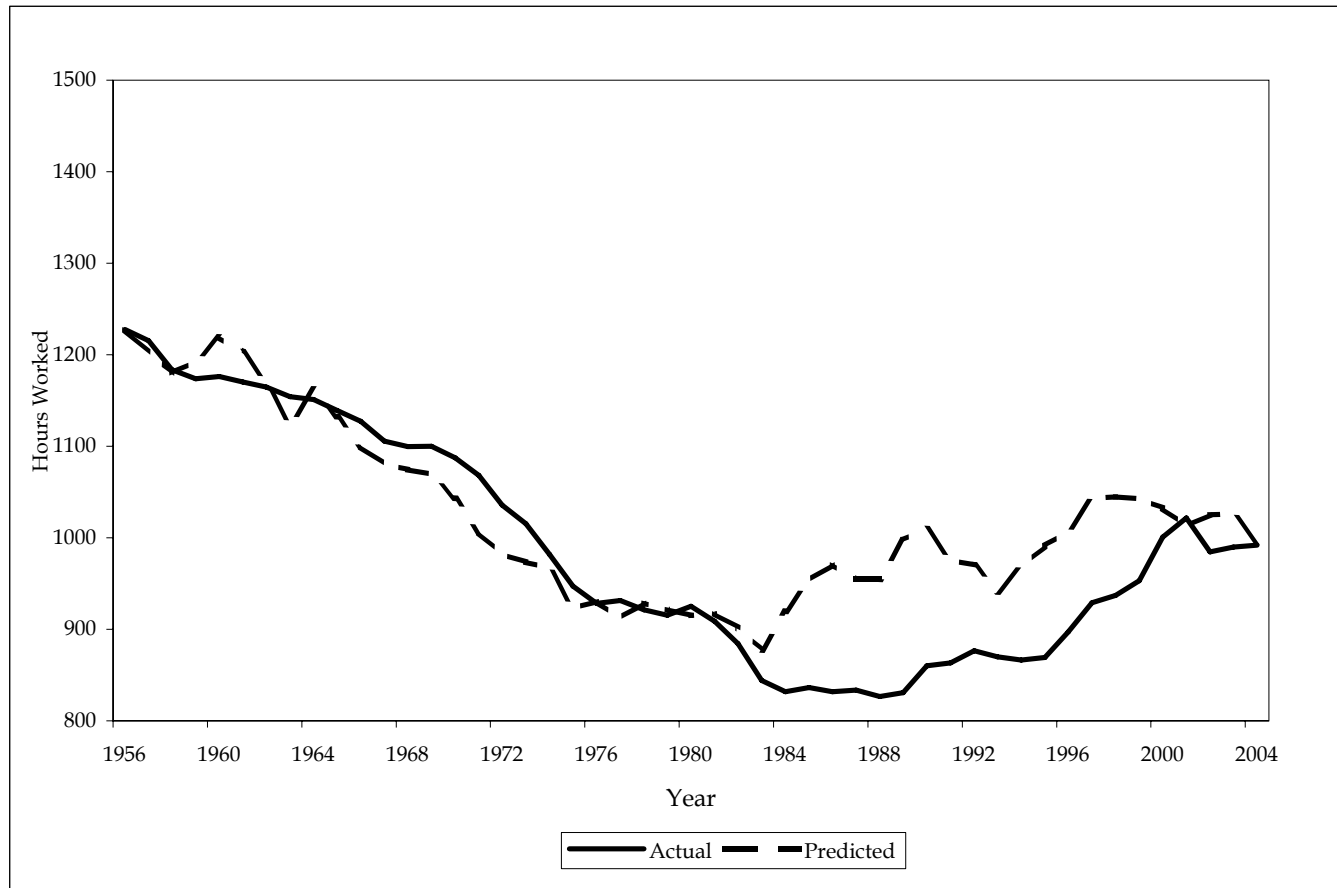


Figure 19. Actual and Predicted: Germany

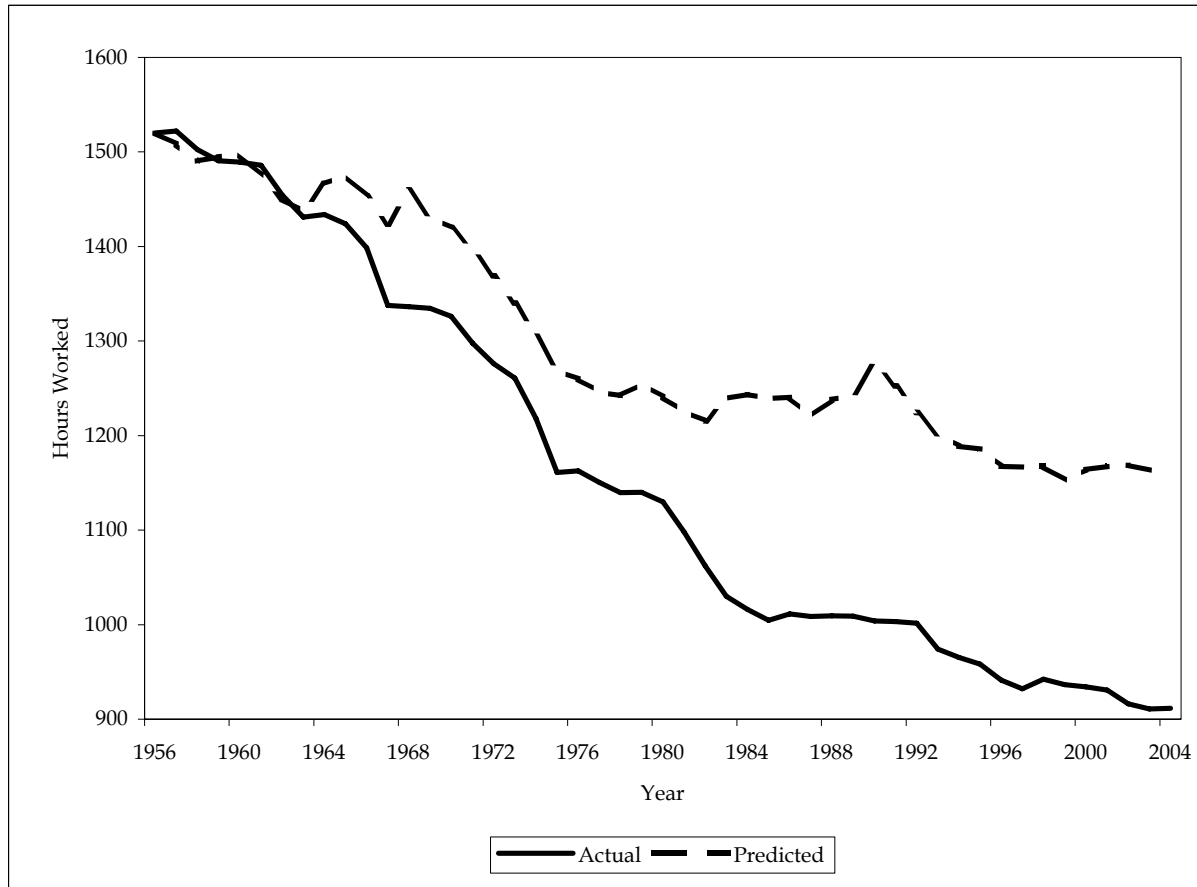


Figure 20. Actual and Predicted: Sweden

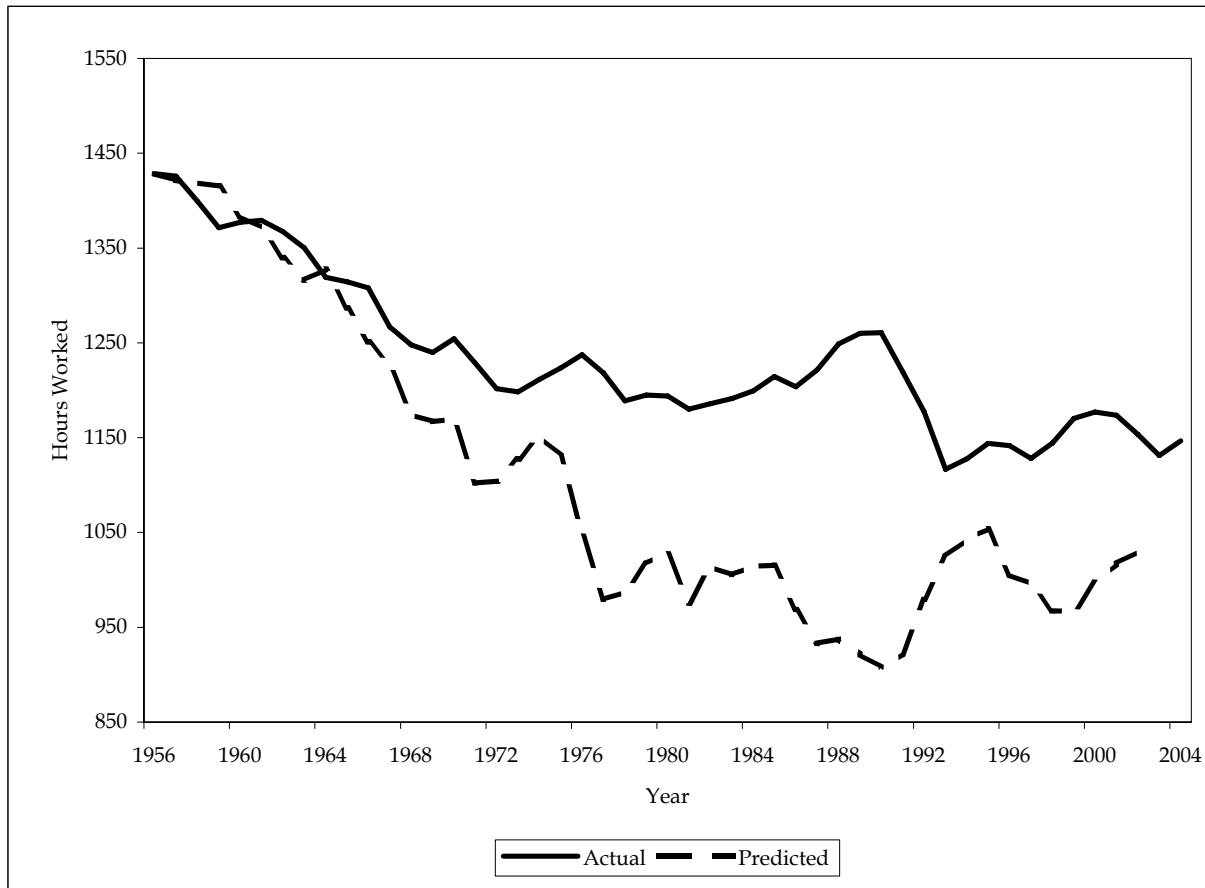


Figure 21. Actual and Predicted: US

