Male Labor Supply and Generational Fiscal Policy*

Christian vom Lehn† Aspen Gorry‡ Eric O'N. Fisher§

February 14, 2017

Abstract

Between 1948 and 2000, hours worked by men in the United States fell by twenty percent. Using a life cycle model of labor supply with intensive and extensive margins, we assess how much of this decline can be accounted for by changes in tax and transfer policies. We use policy measures from the generational accounting literature, capturing the lifetime fiscal burdens faced by each male birth-year cohort. Changes in age demographics and fiscal policy account for 44% of the decline in hours worked. Policy alone explains approximately a quarter of the decline, both in the aggregate and across age groups.

Keywords: Generational Accounts; Taxes; Male Labor Supply; Life Cycle
JEL codes: E24, E62, H24, H31, H6, J11, J22

*We are grateful for comments received from Alan Auerbach, Frank Caliendo, Lei Fang, Carlos Garriga, Victor Rios-Rull, and Johanna Wallenius. We also thank participants in seminars at the 2015 Western Economics Association Meetings, 2015 Fall BYU-USU Macro Workshop, Federal Reserve Bank at Philadelphia, Federal Reserve Bank at St. Louis, 2016 Spring Midwest Macro Meetings, 2016 North American Econometric Society Summer Meetings, 2016 Society for Economc Dynamics Meetings, 2016 National Tax Association Meetings, and the 2016 Southern Economic Association Meetings for their helpful comments.

†Corresponding author. cvomlehn@byu.edu, Brigham Young University, Department of Economics, 130 Faculty Office Building, Provo, UT 84602.

‡aspen.gorry@gmail.com, Department of Economics and Finance, Utah State University, 3565 Old Main Hill, Logan, UT 84322.

§efisher@calpoly.edu, Department of Economics, California Polytechnic State University, 03-444, 1 Grand Avenue, San Luis Obispo, CA 93407.

1 Introduction

Since World War II, the amount of work done by men in the United States has declined substantially. Figure 1 plots the annual total hours worked per person for men, showing that total male hours have declined by twenty percent from 1948 to 2000. Much of this decline has occurred along the extensive margin, as seen in Table 1. Employed men are not only working fewer hours, but fewer men are working. Table 1 also shows that the decline in hours worked has been particularly concentrated among younger and older men, when the decision of whether or not to work is most important.

In this paper, we ask how much of the decline in male work can be accounted for by changes in generational tax and transfer policies in the US. We combine a version of the life cycle labor supply model of Rogerson and Wallenius (2009), which features both intensive and extensive margins, with measures of lifetime tax and transfer rates from the generational accounting literature, originally pioneered by Auerbach, Gokhale, and Kotlikoff (1991). We use the model to simulate life cycle labor supply profiles for each birth cohort from 1900 to 1991 based on each cohort's lifetime tax and transfer rate. To compare our predicted labor supply profiles with the data, we aggregate labor supply across cohorts to construct total hours worked in each year. An advantage of using a quantitative life cycle model is that it also predicts how changes in taxes and transfers impact hours worked at different ages. Therefore, we also compare these predictions of hours worked by age group with the data.

This novel application of the generational accounts has several virtues. First, generational accounts policy measures include policies by gender, allowing us to use male-specific policy figures, which are more appropriate for the question we study. Second, they provide a transparent measure of how the burden of fiscal policy varies across individuals born in different years that is consistent with individual life cycle budget constraints. Moreover, these measures allow us to identify the degree to which tax revenues are transferred back to the household, an important ingredient of our model. Finally, as discussed in Gorry and Oberfield (2012), when individuals face a lifetime budget constraint, taxes faced at one age in the life cycle may influence labor supply at other ages. Thus, forward-looking lifetime tax rates may better account for the timing of how changes in fiscal policy impact aggregate male hours worked. In the discussion, we compare the empirical performance of these tax wedges to more traditional measures of annual tax wedges and find significant additional explanatory power coming from generational accounts measurements.

We find that the combination of changes in fiscal policy and the changes in the age distribution of the population account for 44% of the decline in hours worked per man from 1948 through 2000. Further, the timing of the decline generated in our results is consistent with observed declines, matching well trends for the entire time period, with the exception of the period of the 1970s and early 1980s. Holding fixed demographic changes, changes in policy alone account for roughly a quarter of the decline in hours worked.

Across age groups, changes in demographics and policy account for roughly 60% of the changes in hours worked by men under 16 and men over 55, compared to only 30% for prime aged males. Holding demographic changes fixed, we still find that policy generates roughly a quarter of the decline in hours worked for each age group. Incorporating the extensive margin of labor supply is particularly important for explaining changes in labor supply by the young and the old. Changes at the extensive margin generate roughly two thirds of the policy-driven decline for workers under 24 and workers over 55. These declines at the extensive margin for young and old workers together account for one third of the total decline in aggregate hours worked per man generated by the model. While our model omits trends in educational attainment, which may be relevant for the labor supply of young men, we also present empirical evidence on these trends and find that they are not large enough to render our predictions for young men implausible.

It is useful to compare our results to those of Ohanian, Raffo, and Rogerson (2008), who extend Prescott (2004) in using the first order conditions of a standard growth model to assess the contribution of tax changes to changes in hours worked across OECD countries.[1] By using annual calendar year tax measures, their predictions for hours worked over time look very similar to the decline in male hours worked. However, if demographic factors were added to their analysis, their results would overpredict declines in labor supply prior to 1980 and counterfactually predict increases thereafter. In contrast, while generational variation in tax rates has a much more gradual impact on labor supply, when coupled with demographic changes it generates dynamics of hours that are more consistent with the data. Further, relative to Ohanian, Raffo, and Rogerson (2008), we do not assume all tax revenues are rebated to the household, but use generational accounts data to pinpoint the fraction of income the household receives as transfers. The assumption of full rebates is important for how changes in tax rates influence labor supply, as pointed out by Rogerson (2006, 2007).[2] In the extreme case, where there are no transfers made, changes in taxes have no impact on hours worked. Thus, by using generational accounts policy figures, we are able to discipline the income effect of policy changes. However, the results generated by generational accounts policies look very similar to those in Ohanian, Raffo, and Rogerson (2008) when we assume all tax revenue is rebated.

In focusing on the decline in male labor supply, we abstract from the simultaneous increase in female labor supply that has occurred over this same time period. This decision is motivated by evidence from Juhn (1992), Juhn and Murphy (1997), and Juhn and Potter (2006), who argue that

[1]McDaniel (2011) also studies the effects of tax changes in a dynamic setting that includes productivity growth and home production.

[2]Recent studies, such as Ragan (2013) and Wallenius (2013), account for the role of transfers more carefully by explicitly modeling the benefits of particular transfer programs such as elderly care or social security.

the women who have increased their labor supply have not come from the same households as the men who have been decreasing their labor supply. We reproduce some of the evidence from Juhn (1992) in Appendix B and extend upon it by showing that the decline in hours worked for married and non-married men is extremely similar. While changes in female labor supply may have had an impact on individual household decisions, as in work such as Knowles (2013), these changes do not appear to be generating the observed decline in hours worked by men. Because we do not model labor demand, focusing only on labor supply, we feel this evidence justifies our focus on men.[3]

Our modeling strategy and exercises abstract from many other potential dimensions of heterogeneity in men, which may be interesting and relevant for this question. Most of this abstraction is generated by the choice of using generational accounts policy measures, which have not been constructed for more heterogeneous groups beyond men and women in each birth cohort. However, we find the comparative simplicity and transparency of our approach appealing relative to a more complicated framework with many more dimensions of heterogeneity. Given the magnitude of our results without accounting for these many additional dimensions, we see this as compelling evidence for the importance of fiscal policy and demographics in shaping the decline of hours worked by men.

We build on the literature studying labor supply decisions over the life cycle; see Erosa, Fuster, and Kambourov (2012), Wallenius (2013), Laun and Wallenius (2013a,b), and Alonso-Ortiz (2014). These papers generally focus on the role specific government programs play in explaining cross-country differences in the late-life labor supply patterns of men; our focus is on the time series changes in labor supply behavior by men at all ages in the United States. Our paper is also complementary to recent work by Bick and Fuchs-Schündeln (2014), who use a static household framework to study the labor supply behavior of married men and its interaction with the labor supply of married women, tax treatment of households, and cross-country variation in tax policies.

Elsby and Shapiro (2012) also study changes in male labor force patterns. They consider the extent to which slowing productivity growth in the 1970s can account for increasing nonemployment amongst men. We view our work as complementary to theirs, as our results are unable to explain the changes in hours worked over this time period. Other complementary explanations for why work done by men may have declined in the US over the period of the 1970s and 1980s include changing skill demands (Juhn, Murphy, and Topel (1991) and Juhn (1992)) increases in low-skilled immigration that would impact young workers (Smith (2012)), and trends in structural transformation away from male-dominated industries, such as manufacturing (Ngai and Petrongolo (2014)). Our work is also complementary to the empirical literature studying the effect of

[3]See Kaygusuz (2010), Guner, Kaygusuz, and Ventura (2012a,b), Bick and Fuchs-Schündeln (2014) for recent work on how tax policy changes may affect both female and household labor supply.

changes in social security and disability benefits on declines in male labor force participation (ex. Parsons (1980) and Autor and Duggan (2003)).

Finally, there has been a substantial amount of recent interest in male labor force participation since the financial crisis. We find that both secular trends in tax and transfer policies and demographic composition predict a substantial (and ongoing) decrease in hours worked by men since the year 2000. This is consistent with work by Aaronson et al. (2006) and Aaronson et al. (2014), who argue that signifcant portions of the recent decline in male labor force participation is attributable to long term trends, including changes in demographics.

2 Model

To model life cycle labor supply decisions, we construct a model closely following the setup of Rogerson and Wallenius (2009). Consider an individual who lives for one unit of time. The individual chooses a lifetime consumption plan, and whether and how much to work at each date. The preferences of individual j are represented by

$$\int_0^1 [u(c_j(a)) - v(h_j(a))]da,$$

where j denotes the individual's cohort, $c_j(a)$ is consumption at age a, and $h_j(a)$ is hours worked at age a.[4] The instantaneous utility from consumption $u(\cdot)$ is assumed to be twice continuously differentiable and strictly concave. The instantaneous disutility of labor is $v(\cdot)$. We assume that there is a fixed utility cost of working at each instant, χ. A possible interpretation of this is the fixed cost of commuting or getting ready for work.[5] Then $v(\cdot)$ is given by

$$v(h_j(a)) = \tilde{v}(h_j(a)) + \chi \mathbb{I}_{h_j(a)>0}.$$

Here, we assume that $\tilde{v}$ is twice continuously differentiable and strictly convex. $\mathbb{I}_{h_j(a)>0}$ is an indicator function that takes the value of one if the individual works a strictly positive number of hours at a given date.

The individual in cohort j faces a lifetime budget constraint,

$$\int_0^1 [(1-\tau_j)y_j(a) - c_j(a)]da + T_j \geq 0,$$

[4] We abstract from discounting and assume a zero interest rate for simplicity of exposition. We relax this assumption in calibrating the model, described in Section 4.

[5] Rogerson and Wallenius (2009) model the fixed costs as a time cost which maps more directly into commuting. Our setup follows Gorry and Oberfield (2012) where a fixed utility cost still generates a meaningful decision to work or not at the extensive margin.

where $y_j(a)$ is the flow of income at age a, τ_j is the proportional lifetime income tax rate faced by generation j, and T_j is the present value of lump-sum government transfers for this agent. When making labor supply decisions, the tax rate for the cohort and total transfers are taken as given.

We assume that

$$y_j(a) = w_j(a)h_j(a),$$

where $w_j(a)$ is the exogenously given wage for an individual of age a that varies deterministically over the life cycle.

In this formulation, the individual in cohort j chooses a consumption path $c_j(a)$ and working profile $h_j(a)$ for $a \in [0,1]$ to maximize

$$\int_0^1 [u(c_j(a)) - v(h_j(a))]da$$

subject to the lifetime budget constraint given above.

Since the discount rate and interest rate are equal, the individual chooses to perfectly smooth consumption. Hence, $c_j(a) = c_j$, a flat consumption profile. This smoothing also implies that differences in hours worked across different ages only depend on the wage that the individual faces and not on age itself. Moreover, given the fixed disutility of working at each date, the individual only works when his wage is above some reservation wage, w_j^*.

These properties allow us to reformulate the problem as one of choosing a consumption level and an hours profile that depends on wages rather than age. Let $F_j(w)$ be the cumulative distribution function of wages that the individual faces, with support of $[\underline{w}_j, \overline{w}_j]$. Assume that $F_j(\cdot)$ is continuously differentiable and that there is an interior solution for the reservation wage w_j^*. We can then rewrite the individual's problem as choosing a constant consumption profile c_j, a reservation wage w_j^*, and an hours worked function $h_j(w)$ for $w \in [w_j^*, \overline{w}_j]$ to maximize

$$u(c_j) - \int_{w_j^*}^{\overline{w}_j} v(h_j(w))dF_j,$$

subject to

$$c_j = (1-\tau_j)\int_{w_j^*}^{\overline{w}_j} wh_j(w)dF_j + T_j.$$

The first order necessary condition with respect to $h_j(w)$ is

$$\frac{v'(h_j(w))}{u'(c_j)} = (1-\tau_j)w,$$

and with respect to the reservation wage w_j^* is

$$\frac{v(h_j(w_j^*))}{u'(c_j)} = (1-\tau_j)w_j^* h_j(w_j^*).$$

These necessary conditions have standard interpretations. The first condition states that the after-tax real wage balances the marginal disutility of labor against the marginal utility of consumption. The second describes the extensive margin. At the reservation wage, one's entire take-home pay at balances the disutility of working $h(w_j^*)$ hours (the marginal disutility of choosing to work), with the marginal utility of consumption. The first equation gives the individual j's labor supply for $w \in [w_j^*, \overline{w}_j]$, and the second determines the reservation wage, w_j^*.

Individual j is the representative of a generation that faces an economic environment shaped by its lifetime tax rate on income τ_j, its lifetime lump-sum transfers T_j, and the distribution of wages during its life $F_j(w)$. It is convenient to express generational lump-sum transfers as a fraction θ_j of lifetime labor income:

$$T_j = \theta_j \int_{w_j^*}^{\overline{w}_j} w h_j(w) dF_j.$$

3 Measuring Taxes and Transfers

3.1 Generational Accounts

We draw from the literature on generational accounts to measure tax and transfer rates for men by cohort, τ_j and θ_j. Beginning with Auerbach, Gokhale, and Kotlikoff (1991), generational accounting is a proposed alternative to traditional deficit accounting in assessing the viability and sustainability of fiscal policy.[6] An important insight of this literature is that conventional cash-flow measures of policy variables do not adequately represent the dynamic nature of fiscal policy nor how their burdens are distributed across different generations, including those not yet born. Hence, it is important to keep track of how the benefits and costs of tax-transfer policies accrue to different agents in the economy.

This literature commonly measures how government consumption and transfers are financed

[6]See Conesa and Garriga (2015) for a recent discussion of generational accounting methodology and an analysis of how generational accounting measures relate to intergenerational welfare consequences of changing fiscal policies. Fisher (1995) shows that these ideas are also relevant in economies with incomplete financial markets, and Fisher and Kasa (1997) explores generational accounts in a model of the open economy where the crowding out of capital may also occur.

by different generations using the following intertemporal government budget constraint:

$$\sum_{j=t-D}^{t} N_{t,j} + \sum_{j=t+1}^{\infty} N_{t,j} = \sum_{s=t}^{\infty} G_{t,s} + B_t.$$

$G_{t,s}$ is the present value at time t of government consumption in year s and B_t is current government debt. This formulation assumes that individuals live for D years. $N_{t,j}$ is the present value at time t of net taxes that fall on cohort j, and is called the generational account for cohort j. Net tax burdens for future generations, $j > t$, are computed as a residual given measurement of other all other objects. While the generational accounting literature is commonly interested in the generational accounts for future generations, we are only interested in the generational accounts for past and current generations.

The generational accounts, $N_{t,j}$, are simply the difference of the present value of cohort $j's$ total tax payments at time t, $P_{t,j}$, and the present value of total transfers received by that cohort, $T_{t,j}$:

$$N_{t,j} = P_{t,j} - T_{t,j}. \tag{1}$$

Auerbach, Gokhale, and Kotlikoff (1993) show that generational accounts can be used to construct forward-looking lifetime tax and transfer rates by birth cohort and by gender. These rates are computed by dividing each generation's present value of tax payments and transfers at the time of birth, $P_{j,j}$ and $T_{j,j}$, by the present value of lifetime labor income at birth, $Y_{j,j}$. The lifetime gross tax rate τ_j and gross transfer rate θ_j for birth cohort j are:

$$\begin{aligned} \tau_j &= \frac{P_{j,j}}{Y_{j,j}} \quad \text{and} \\ \theta_j &= \frac{T_{j,j}}{Y_{j,j}}. \end{aligned}$$

Constructing these tax rates requires data on the taxes paid, transfers received, and labor income earned by both genders in each cohort at each point in time. Auerbach, Gokhale, and Kotlikoff (1993) and Auerbach, Gokhale, and Kotlikoff (1994) use micro data to assess the distribution of payments, transfers, and income in each year across men and women different birth cohorts. They then multiply these distributions by total tax receipts, transfer payments, and labor income in the relevant calendar year. Total tax payments include taxes on labor income, capital income, payroll taxes, excise taxes, and property taxes. Transfer payments include social security transfers, Medicaid and Medicare, and welfare payments. These tax and trasnfer rates are computed using data through the year 1991. For cohorts still alive in 1991, future taxes and transfers are calculated using projections of fiscal policy based upon Congressional Budget Office forecasts and payment distributions in the last year for which data are available coupled with population projections based on life expectancy tables from the Social Security Administration. Additional details are available in Auerbach, Gokhale, and Kotlikoff (1993).

Gross lifetime tax and transfer rates for each male cohort from 1900 to 1991 are plotted in Figure 2. The original figures in Auerbach, Gokhale, and Kotlikoff (1994) are only reported by decade starting in 1900; the annual measures shown in Figure 2 are obtained by polynomial interpolation. Both tax and transfer rates have been monotonically increasing over time, though growth in both slowed after the 1970 cohort. Notably, the transfer rate is substantially lower than the tax rate. This is in part because the some of the largest sources of government transfers are social security and Medicare, which are paid later in an individual's life and thus are heavily discounted. While some government spending is not included in the transfer rate, we follow Rogerson (2006, 2007) and argue that many of these expenditures, such as those for national defense, may have no impact on the marginal utility of private consumption. In Section 6, we consider how variation in the transfer rate affects our results by allowing for education expenditures to enter into the analysis.

3.2 Methodological Comparison with Other Tax Measures

The remainder of this section compares our tax and transfer measures with the measures of average annual tax rates typically used in the macro literature. The most comprehensive measure of aggregate tax rates comes from McDaniel (2007), who constructs the tax measures used in Ohanian, Raffo, and Rogerson (2008). Building on the approaches of Mendoza, Razin, and Tesar (1994) and Prescott (2004), McDaniel (2007) constructs an average labor tax rate by computing aggregate taxes paid in each calendar year and then dividing by aggregate labor income. The final tax wedges used in Prescott (2004) and Ohanian, Raffo, and Rogerson (2008) also depend on the consumption tax rates that are also computed by McDaniel (2007).

It is helpful to recall that in both our life cycle model and also the standard growth model used by Ohanian, Raffo, and Rogerson (2008), it is the *marginal* tax rate that determines hours worked per worker. However, in our model, the *average* tax rate is also important for hours worked decisions at the extensive margin. The tax rates computed by McDaniel (2007) and those computed from the generational accounts both measure *average* tax rates. In Ohanian, Raffo, and Rogerson (2008), the model has one person per year, so the average tax rate proxies for an aggregate marginal rate since there is no heterogeneity. We have a single lifetime rate faced by each cohort that proxies for both the marginal and average tax rates. Using this average lifetime tax rates fits with individual behavior on the extensive margin, where the bulk of the changes in hours worked occurs in the data.

Using generational accounts has three primary advantages. First, they provide a male-specific measure of taxes, which provides a better measure of policy for the question we study. In Appendix

C, we report results using combined policies for both men and women and find that this actually improves our results. Second, they are consistent with life cycle labor supply decisions. Traditional annual cash-flow measures of tax rates may not adequately capture the timing of the labor supply response, since labor supply decisions may depend on tax rates in other years than the current one. Third, transfers as well as taxes matter for labor supply decisions. With balanced-growth preferences, if taxes are collected and "thrown into the ocean" ($\theta_j = 0$), there is *no* effect on labor supply. Many prominent papers in the existing literature, such as Prescott (2004) and Ohanian, Raffo, and Rogerson (2008), go to the opposite extreme, assuming that *all* taxes collected are rebated lump sum to the household ($\theta_j = \tau_j$). Rogerson (2006, 2007) emphasizes that the theoretical implications of this assumption are significant. Using estimates of lifetime transfer rates provide us with empirical discipline on their level and thus a more precise estimate of their effect on labor supply.

A drawback of generational accounts relative to McDaniel (2007) is that it lumps all taxes, including capital and excise taxes, into one measure. This may misrepresent the relevant tax rate for labor supply decisions and changes in measured tax rates may not may not reflect actual changes in the labor tax wedge. However, McDaniel (2007) observes that, with the exception of payroll and labor income taxes, all other tax rates in the US have decreased or remained stable since the year 1950. Thus, although the exact level of taxes may be higher than the actual labor tax wedge, the inclusion of these other taxes may lead to a conservative estimate of the *increase* in taxes relevant for labor supply decisions.[7]

One final possible concern is applying the generational accounts tax measures in a life cycle model of labor supply misses progressivity in the US tax system, which may matter for labor supply decisions. If progressivity has been increasing over time, this could mean that increases in generational average taxes are unevenly distributed across ages. Thus, our estimates of labor supply changes at the extensive margin could be biased upward, as higher progressivity increases incentives to work at lower wages relative to higher wage times. As we discuss in the following section, part of our calibration strategy matches the life cycle hours worked profile in the model to the data, which thus implicitly accounts for the extent to which a constant level of progressivity shapes hours worked over the life cycle. Further, Piketty and Saez (2007) show that the progressivity of the US tax system has been declining at least since 1960, and possibly earlier. Thus, if anything, omitting progressivity from our analysis means our estimates of the labor supply response at the extensive margin are conservative.

In summary, we view the use of generational accounts as a useful complement to the existing literature and not a replacement for other measures of tax and transfer rates. In Section 6, we

[7]Notably, as discussed in Auerbach, Gokhale, and Kotlikoff (1994), inclusion of these other taxes has a significant impact on estimates of taxes for women, as some consumption taxes are allocated evenly across household members. While this concern may be less significant for men, we consider adjustments to generational taxes to remove these other taxes in Section 6.

explore the some of the quantitative implications of these measurement differences through several robustness checks, which make modifications to the rates computed from generational accounts, and discuss to what extent lifetime tax rates better account for the timing of the decline in hours.

4 Calibration

4.1 Aggregation over Time and Age

Since our model predicts a unique hours profile for each cohort, we want to compare the model predictions with detailed micro data on hours worked by each birth cohort in every year. Detailed annual data on hours worked by age, however, are available only since 1962. But annual data on hours worked per person in roughly 10 year age groups are available from the Bureau of Labor Statistics (BLS) since 1948.[8] Therefore, we focus our study on hours worked in these coarse age bins and aggregate labor supply for males 16 and older. This section describes how we aggregate the hours profiles by cohort from the model to be comparable with the data. For simplicity, although all of our calculations are for men and in per capita terms, we will henceforth refer to model and data results as simply hours worked.

4.1.1 Time Aggregation

Since the representative agent for each cohort j lives for one unit of continuous time in the model, we must map the interval $[0,1]$ into discrete ages. We assume that each cohort in the model lives from age 16 to 80, and integrate over discrete sets within the interval [0,1] to construct each cohort's hours worked at each age. Annual hours worked by cohort j at age $v = 16,...,80$ are

$$H_j(v) = \int_{a\in A_v} h_j(a)da,$$

where $A_v = \left[max\{\frac{v-16.5}{65},0\},min\{\frac{v-15.5}{65},1\}\right)$.[9] An agent of age v in calendar year t was born in year $j = t - v$, and he works

$$H(t,v) \equiv H_{t-v}(v)$$

[8]Because data on tax and transfer rates by cohort are available only for men born after 1899, the model cannot describe any policy-driven changes for workers over 55 until 1956 and for workers over 65 until 1966. Hence, in both the model and the data, hours worked by men age 55-64 are held constant at their 1960 level until 1960, and for men over age 65, hours worked are held constant until 1970. By these years, half of the relevant cohorts have potential variation in tax and transfer rates. As labor supply for young workers in near the end of the sample depends on cohorts born after the year 1991, we assume that these cohorts face the same policies as that born in 1991.

[9]We assume that individuals counted as, say age 20, are effectively agents between the ages of [19.5,20.5). This is consistent with taking monthly averages of BLS data to get annual hours worked by age group. Our results are robust to counting agents at 20 as being agents in ages [20,21).

hours in that year.

4.1.2 Age Aggregation

To aggregate model output for hours worked by each age in each year, we require data on the fraction of the population at each age, $\varphi(t,v)$, for ages $v = 16,...,80$ and for every year t. We obtain values of these population weights from the March Current Population Survey (CPS) and the Decennial Census; details are in Appendix A.

Given information on the fraction of the population at each age at time t, we aggregate the model's predictions for hours worked by age and year into broad aggregates. We consider five separate age bins, V_k, for k in the five age groups we consider: 16-24, 25-34, 35-44, 45-54, and 55+.

Hours worked per person for age bin V_k are:

$$H_k(t) = \frac{\sum_{v \in V_k} H(t,v)\varphi(t,v)}{\sum_{v \in V_k} \varphi(t,v)}.$$

Aggregrate hours worked per person at time t is then given by:

$$H(t) = \sum_k H_k(t)\varphi_k(t),$$

where $\varphi_k(t) = \sum_{v \in V_k} \varphi(t,v)$ is the population share in age group V_k.

4.2 Functional Forms and Parameter Values

We choose utility functions that are consistent with balanced growth preferences:

$$u(c_j) = \log c_j$$

and

$$v(h_j(w)) = \alpha \frac{h_j(w)^{1+\gamma}}{1+\gamma} + \chi \mathbb{I}_{h_j(w)>0}.$$

We also impose a subjective discount factor and a real interest rate that are both 3%. With these functional form choices, it is not difficult to show that the effective tax wedge in the model will be $\frac{(1-\tau_j)}{(1-\tau_j+\theta_j)}$. Again, if $\theta_j = 0$, changes in taxes have no effect on labor supply.

Calibration of the life cycle wage profile, $w_j(a)$, is important for the model's results. Rupert and Zanella (2015) show that standard life cycle models using actual wage profiles fail to reproduce profiles of hours worked that are consistent with the data, particularly for older and younger workers. We do not make progress in resolving this puzzle. Instead, we calibrate the wage profile for the 1940 cohort to match average labor supply behavior for men born from 1935 through 1945. We focus our calibration on the 1935-1945 cohorts as we have data on their hours worked for nearly their entire lives. In particular, we calibrate the profile of wages to match data on average hours per worker across the life cycle using the model's first order conditions for the intensive margin. We further assume that each cohort faces the same wage profile as the 1940 cohort, $w_j(a) = w_{1940}(a)$. Since income and substitution effects are off-setting with balanced growth preferences, this assumption is justified if the changes in wage profiles are proportional at each age in life. An advantage of calibrating the wage profile to match actual hours worked behavior for a particular cohort is that this implicitly accounts for any other constant non-wage factors that impact average hours worked over the life cycle, e.g. constant progressivity in the tax structure.

An alternative, isomorphic calibration for the wage profile would be to use observed wages from the data and then calibrate an age-varying disutility of work to match the life cycle labor supply profiles in the data. This calibration also lends itself naturally to allowing wage profiles to change over time. The primary issue with this calibration and time-varying wage profiles is accurate measurement of successive wage profiles, as reasonable age measures are only available in the CPS beginning in 1976. We further discuss measurement issues, what this alternate calibration would imply for parameters, and present tentative results using interpolated time-varying wage profiles in Appendix C.

The exact calibration of the wage profile depends on the parameterization of the agent's disutility of labor, particularly the parameter γ, which governs elasticities of labor supply. A wide range of values for this parameter have been considered in the micro and macro literatures and estimates of this parameter range widely depending on estimation procedures; see Keane and Rogerson (2012), Chetty et al. (2012), Fiorito and Zanella (2012) and Peterman (2016). We set $\gamma = 1$, as this is roughly the midpoint of the range of common value used or estimated, and this value is also consistent with micro estimates of the Hicksian elasticity of labor supply, which is the relevant elasticity for our cross-cohort comparisons.[10] However, Appendix C considers how results vary with different values of γ and we find that our results are robust to a wide range of values.

Figure 3 reports the calibrated wage profile, $w(a)$. The wage is normalized to 1 at age 16 and assumed to be zero above age 75.

We calibrate the two remaining parameters, α and χ, to normalize the maximum hours worked at a point in the life span to 1 and to match the fraction of life spent working for the 1940 cohort. We assume the representative agent in the 1940 cohort works for 44 years. This assumption is consistent with the data, which shows that median hours worked for individuals in the 1940 birth

[10]This value of γ is also consistent with micro estimates when the nonconvexity generating the extensive margin appears in hours worked, as in Rogerson and Wallenius (2009).

cohort are positive for 44 years. These calibration targets give values of $\alpha = 1.24$ and $\chi = 0.43$. Figure 3 also shows the reservation wage implied by these values for the 1940 cohort. This reservation wage implies that the 1940 cohort works from approximately age 19 to age 63. Appendix C shows that our results are robust to different values for χ.[11]

5 Results

5.1 Aggregate Results and Results by Age Group

Figure 4 shows the time series of hours worked from 1948 to 2014 in the model and the data; all series are normalized to 1 in the year 1948.[12] Prior to the cyclical downturns in the 2000s, the combination of policy changes and changes in the age distribution of the population accounts for a sizable fraction of the decline in the data. Further, the timing of this decline in our results closely matches what we observe in the data and our results even generate a meaningful decline in labor supply following the year 2000. The only exception is the period of the 1970s and the early 1980s, when the data shows a sizable decline in hours worked, but the model generates minimal decline.

Figure 5 shows the model's results for each decadal age group: 16-24, 25-34, 35-44, and 55+. The model best accounts for the decline in hours worked for men aged 16-24 and 55+, the ages that have experienced the greatest decline in the data. For individuals ages 25 through 54, the model generates much less decline. As with the aggregate results, results by age group match well the timing of labor supply changes in the data with the exception of the 1970s and early 1980s. As most of the decline in hours worked for prime-aged workers occurred in that decade, the model only accounts for a much smaller amount of the total decline observed in the data, compared to older and younger workers. The model generates larger declines for the youngest and oldest workers for two reasons: greater sensitivity to changes in demographics due to the steeper wage profile at these ages and importance of the extensive margin, which is operative only at these ages. Sections 5.2 and 5.3 disentangle each of these effects.

Table 2 summarizes these results for the period 1948 to 2000, showing the fraction of the decline in observed hours worked that the model explains for the aggregate and each age group. The data reported in this table is smoothed with an HP filter with parameter 6.25 to ensure that choosing an endpoint of 2000, when the business cycle was at a high point, does not generate artificially stronger results. In total, changes in age demographics and in tax and transfer policies explain 44%

[11] An alternative approach is to allow for some heterogeneity in χ, and then weight the labor responses of different χ values to smooth the age cutoffs across several different ages. We have experimented with this in a few simple ways and found similar results.

[12] As we report results both for the aggregate and for individual age groups, it is not possible to calibrate the model to match the initial levels of hours worked for the aggregate and each age group.

of the decline in male hours worked between 1948 and 2000. For older and younger workers, these two forces account for roughly 55-65%, whereas for prime-aged workers, the fraction explained is only between 20-30%.

5.2 Holding Demographic Change Fixed

Since the aggregate results from the model are generated using the age distribution of the population, it is important to understand how much of the predicted decline in hours is due to fiscal policy and how much is due to mechanical changes in demographics. We use the following standard decomposition:

$$\frac{\Delta H(t)}{H(t)} = \frac{\sum_k \overline{H}_k \Delta\varphi_k(t)}{H(t)} + \frac{\sum_k \Delta H_k(t)\overline{\varphi}_k}{H(t)}. \tag{2}$$

where $\overline{H}_k$ and $\overline{\varphi}_k$ are the averages – across the beginning and ending periods – of hours worked and the population share for age group k. The first term on the right shows the demographic effects alone, and the second term shows labor supply effects with fixed population weights. We use this decomposition initially to explore changes in the composition of the population across the broad age bins V_k. We initially focus on decadal age groups because we have a complete data time series of hours worked at this level of age aggregation.

Table 3 reports the results of this decomposition in the model and in the data. Since our results in Figure 4 suggested that the model explains less of the decline after 1970, we consider both the demographic decomposition for both the full sample from 1948 to 2000 and the period 1948 to 1970. For the decline between 1948 and 2000, changes in demographics contribute nothing to the decline in hours worked. However, a sizable part of the decline in both the data and the model between 1948 and 1970 comes from demographic changes. This is more clearly visible in Figure 6, which plots the hours worked generated from the model by either only allowing variation in the share of the population in each age bin or from holding this variation fixed. These changes in the age distribution contribute significantly to hours worked declines through the year 1970, but reverse and have minimal aggregate impact by the year 2000. However, demographics are prominent again in generating most of the predicted decline in labor supply after the year 2000.

While changes in the age composition of the population across these decadal age groups contribute minimally to the decline in hours worked through 2000, changes in the age composition at the individual age year level may yet be important. Figure 7 shows the counterfactual model hours worked series generated by allowing either only detailed population shares to or only policy to change over time. Demographic change generates most of the decline in hours worked through the year 1970, but generates an increase in hours worked throughout the 1980s, though there is still some impact of demographic change still present by the year 2000. In contrast, changes in policy generate a near linear trend decline in hours worked, and serve to offset the increase in hours worked coming from demopgrahics during the 1980s. Both changes in policy and changes in demographics contribute to the predicted decline in labor supply after the year 2000.

Although annual data on hours worked at each age are not available for the entire period we study, we can compute an estimate of the effect of demographic change in the data by combining microdata from the 1950 Census and the March CPS beginning in 1962.[13] With this estimate, we can determine how much of the decline in hours worked not explained by demographics can be accounted for by the changes in taxes and transfers of our model.

Table 4 reports the results of this more detailed decomposition for the years 1948-2000. After accounting for demographic changes, changes in policy alone account for between 20-25% of the total decline in both aggregate hours worked and hours worked holding demographics fixed. Within age groups, the model also accounts for 20-30% of the decline holding demographics fixed in each group. We also report the results of this more detailed decomposition for the years 1948-1970 in Table 5. As seen in Figure 7, demographic changes play an important role for this time period. However, holding demographic changes fixed, the model still explains roughly 20% of the decline in hours worked.

In summary, the model accounts for 44% of the decline in hours worked between 1948 and 2000. We find that about half of that decline comes from changing policies and half comes from demographic changes, though the contributions of these two forces have fluctuated significantly over time and both contribute to the decline in hours worked since 2000.[14]

5.3 Intensive versus Extensive Margins

To understand the relative importance of the intensive and extensive margins of labor supply, we decompose our results across these two margins. In year t, data measures of hours worked per person $H(t)$ can be written as

$$H(t) = \sum_{v} \frac{H(t,v)}{E(t,v)} E(t,v) \varphi(t,v),$$

where $\frac{H(t,v)}{E(t,v)}$ is hours per worker of age v (intensive margin), $E(t,v)$ is the fraction of the population working of age v (extensive margin), and $\varphi(t,v)$ is the corresponding share of the population. This decomposition allows us to compute the relative contribution of intensive and extensive margin

[13]Details of how this is computed are reported in the notes to Table 4.

[14]Notably, the model suggests a slightly larger role for demographic shifts than we estimate from the data over this time horizon. This may be partially due to the difficulties in carefully identifying annual labor supply at the individual age level prior to the 1960s. Another partial explanation is that the larger role of age demographics may be due to the stylized nature of the entry and exit decisions in the model.

changes, using the same approach as used for demographic change, shown in Equation 2.[15] In doing this decomposition, we hold age distributions fixed to eliminate confounding demographic changes.

The model permits a similar decomposition. For example, consider hours worked by a 40-year-old man in 1948 and those by his counterpart in 2000. In 2000, hours worked are determined by the labor supply rule $h_{1960}(w)$ and the cut-off wage w^*_{1960}. For a 40-year-old man in 1948, the analogous rules are $h_{1908}(w)$ and w^*_{1908}. We compute the changes in hours worked coming from the extensive margin by computing the hours worked by 40-year-old men in 2000 assuming they have the same intensive margin rule as 40-year-olds in 1948, $h_{1960}(w)$. With identical intensive margin behavior, differences in hours worked only come from the extensive margin, determined by w^*_j.[16] These changes are computed at each age, and aggregated holding population weights fixed. Changes coming from the intensive margin are counted as residual effects unexplained by the extensive margin.

Table 6 reports our decomposition. It is not surprising that there are no changes in the extensive margin for prime-aged males between 25 and 54, as the decisions to enter and leave the labor market occur around ages 19 and 63. This is generally consistent with the data, which also suggests a lesser role for changes in the extensive margin for these ages.

The results for young and old workers are more interesting. For younger workers, the model accounts for a sizable fraction of the decline on the extensive margin, but for less of the decline on the intensive margin. The model's inability to explain fewer hours worked per worker for young workers is perhaps unsurprising, as the model is silent regarding other changes affecting young men, such as trends in obtaining higher education.[17] For older workers, the model captures a modest fraction of the change along the extensive margin, where there has been the greatest total decline. The model also predicts that hours worked on the intensive margin have declined, but in the data these hours have actually increased.[18] However, Rupert and Zanella (2015) note that there has been a blurring of intensive and extensive margin adjustment among older workers in the data, which may complicate a detailed comparison.

At this point, one might be naturally concerned that the model generated decline in hours

[15]Due to data limitations, we make a small adjustment, relative to the decomposition shown in Equation 2. Details are described in the notes to Table 6.

[16]In empirical applications it is commonplace to compare hours worked averaging the intensive margin behavior at both the start and end dates, as in Equation 2. However, hours worked per worker is not well defined for those who are not currently employed. Hence, we fix hours worked per worker at its initial level. The same is done for the data decomposition; using averages in the data does not change the results significantly.

[17]Changes in higher education enrollment are likely to strongly impact the intensive margin of labor supply. In 1970, roughly 50% of college students ages 16-24 were employed during their studies, and that fraction has been steadily rising; see National Center for Education Statistics (2009).

[18]Note that this is only true if the hours worked by older workers is constrained to be unchanging prior to 1960 for 55-64 year old men and prior to 1970 for 65+ year old men.

worked for young workers is unusually high, given that the model does not account for the sizable trends occurring in educational attainment. To address this concern, we do a simple statistical decomposition of the contribution of schooling trends to the change in hours worked. Using data from the 1950 and 1980 Census on schooling patterns, we compute the fraction of the decline in hours worked for 16-24 year old men explained by only changing the fraction in those men in school. Changes in the fraction of men in school accounts for roughly 20% of the decline in hours worked for 16-24 males from 1950 to 1980. This figure is also likely biased upward, as demographic trends towards younger workers over this period naturally inflate the fraction of those in school; holding demographics fixed, schooling trends explain only 8% of the decline in hours worked.[19] Thus, while schooling is an important margin our model omits, it does not appear to be so quantitatively important so as to render our results implausibly large.

We raise two other important points in interpreting our results for young workers. First, with a typical concave wage profile, the model predicts that there will at most be two ages where the extensive margin is relevant. However, heterogeneity within cohorts regarding life cycle productivity, perhaps owing to education, could lead to extensive margin changes at other ages.[20] Second, the representative agent for each cohort has perfect foresight and is able to transfer resources across the life cycle. Loosening either of these assumptions would likely lead to smaller declines in labor supply early in life and greater declines later in life, as workers with additional savings are able to retire sooner.

In spite of these potential concerns, an important feature of introducing the extensive margin is the ability to account for differentially large labor supply changes for younger and older workers. Comparing the intensive margin changes across ages, we see that that the model produces very similar effects for each age group. Without an extensive margin, the model would grossly understate changes in labor supply for the young and old.

5.4 Changes in Taxes vs. Changes in Transfers

An important contribution of using generational accounts is that it gives separate measures of the taxes and transfers faced by each generation. This separate measure is an important contribution, as many prominent studies assume that all tax revenue is rebated back to the individual. To better understand the separate role of transfers in our model, we decompose our results into changes stemming from changes in taxes and changes in transfers.

[19]Further details of these exercises are available upon request from the authors.

[20]Juhn (1992) gives evidence that the decline in male labor force participation is linked to less educated workers. Our model treats every cohort as a representative agent who has average skills and faces one tax rate. High skilled workers would have different ages at which they entered and left employment. They would also face higher marginal tax rates. Without explicitly modeling this kind of heterogeneity, it may be difficult to make direct comparisons to the margins of adjustment in the data.

Table 7 reports the fixed demographic decline in hours worked accounted for by taxes and transfers separately. This is done by allowing one policy to vary while holding the other fixed at its value for the baseline 1940 cohort. The policy change with the greater impact is the change in the transfer rate, accounting for roughly four more of a labor supply decline than the change in the tax rate. To understand this result, it is helpful to recall the effective tax wedge in the model, $\frac{1-\tau_j}{1-\tau_j+\theta_j}$. From this expression, the level of transfers influences the effect of changes in taxes, and vice versa. In particular, if the transfer rate is zero, then wedge is identically equal to 1 for any tax rate, meaning that taxes would have no effect on labor supply. If on the other hand the transfer rate is identically equal to the tax rate, then the tax wedge is $1-\tau_j$. Given the transfer rate is only about 5% on average compared to a tax rate of around 30-35%, changes in taxes have a muted effect on the tax wedge and thus labor supply.

6 Discussion

What do we learn from applying generational accounting measures of fiscal policy in a macroeconomic model of life cycle labor supply? It is difficult to compare our results directly with the extant literature because our model and its application are idiosyncratic. Ohanian, Raffo, and Rogerson (2008) (hereafter ORR) is the paper closest in spirit to ours, although they use quite different methods. We consider three possible comparisons between our methods and results and theirs: we directly compare measurements of tax wedges from each paper, we consider modifications to the generational accounts tax wedge to resemble theirs more closely, and we consider an empirical exercise evaluating the value added from each measure. For a more direct comparison, we use total tax and transfer measures from the generational accounts instead of the male-specific ones, as ORR use aggregate tax rates. These rates are obtained from Gokhale, Page, and Sturrock (1999). As shown in Appendix C, using this measure of tax and transfer rates generates very similar results to those obtained using male-specific policies.

6.1 Aggregating Generational Accounts Tax Measures from Birth Cohorts to Calendar Years

One difficulty in comparing the generational accounts tax wedge with the one used in ORR, is that their tax wedge is defined with respect to calendar year, whereas ours is defined with respect to birth year. However, we can use population weights to aggregrate up generational accounts tax wedges by cohort to construct an approximate calendar year tax wedge. We write the tax wedges for the two measures such that higher tax wedges correspond to reductions in labor supply in each framework. The tax wedge for ORR is given by $1-\frac{1-\tau_l(t)}{1+\tau_c(t)}$, where $\tau_l(t)$ is the sum of payroll taxes and labor income taxes at time t, and $\tau_c(t)$ is the consumption tax at time t.[21] Our tax wedge from generational accounting, prior to aggregation, is $1-\frac{1-\tau_j}{1-\tau_j+\theta_j}$ for birth cohort j.

Figure 8 plots the average generational accounts tax wedge, gross tax rate and gross transfer rate against the tax wedge from ORR. The two policies are plotted on separate axes, scaled to emphasize the timing of policy changes between the two measures instead of their relative levels. We see that there is marked similarity between the behavior of these two tax wedges over time, as both have increased, with substantial increases prior to the year 1980. However, whereas the ORR tax wedge shows much less growth after 1980, the generational accounts tax wedge, gross tax rate, and gross transfer rate all continue to increase. The gross tax rate from generational accounts slows down eventually like the ORR tax wedge, but not until at least the 1990s. Given that the two tax trends are both driven by the same aggregate changes in labor income and payroll tax rates, the differences in timing come from the forward-looking nature of the generational accounts. When aggregated to a lifetime tax rate as in the generational accunts, the effects of the tax rate changes before the 1980s are still observed for at least the subsequent decade. And the increase in the transfer rate, not accounted for in the ORR tax wedge, further leads to increases in the total generational accounts tax wedge through the end of the sample.

Why do these timing differences matter? Both our results in our framework and the results of ORR in their framework generate relatively flat aggregate levels of hours worked after the 1980s.[22] However, our results in Section 5.2 showed that demographic trends from 1980 to 2000 would have led to increases in hours worked were it not for the continued downward pressure on labor supply coming from policy changes. Thus, we argue the continued increase in generational accounts policy measures after the year 1980 is necessary to generate a flat hours worked series for the period 1980 to 2000. Were demographic effects to be introduced to ORR, we expect this would lead to a counterfactual rise in hours worked occurring over this time period. Thus, in terms of generating the timing of labor supply changes from policy changes, generational accounts tax wedges offer improvement upon the existing approach.

6.2 Modifications to Generational Accounts Measurements

In terms of actual model predictions, ORR predict a 9.4% total decline in hours worked in the US between 1956 and 2003; for that period, using total taxes and transfers for men and women, our

[21] All results in this section are robust to using the tax wedge from ORR without the consumption tax component. ORR also have a subsistence consumption term in preferences, however, as they note, the quantitative impact of this is small after the year 1970.

[22] The result for ORR can be seen in their paper, but is also discernible from the time path of their tax wedge in Figure 8.

model generates a decline of 6.4%, of which 4.8% comes from policy changes alone.[23] As we have shown that there are clear differences in the tax wedges across these studies, we discuss how measurement adjustments affect our conclusions.[24]

Section 3 explained that generational accounts tax rates include not only labor taxes, but also those on capital and excise taxes, which differs from ORR, who only focus on labor taxes. Also, measured transfer rates do not include many categories of government expenditures, such as educational spending. As a robustness check, we consider three modifications to our tax and transfer rates: (1) adjusting the level of tax rates to represent just labor taxes; (2) adjusting the transfer rate to account for additional government spending in the form of educational expenditures; and (3) allowing for full rebate of tax receipts to the household ($\theta_j = \tau_j$).[25] Table 8 reports how these changes affect our results. We report both the decline from 1948 through 2000, as a comparison with our own benchmark, and the decline from 1956 through 2003 to compare directly to Ohanian, Raffo, and Rogerson (2008).

Table 8 shows that the first two modifications to tax and transfer rates generate only modest changes in the results. These changes do not reconcile the differences between our findings and those of Ohanian, Raffo, and Rogerson (2008). However, if transfers fully rebate all tax revenue and demographics are fixed, the model yields a decline of 9.5%, almost exactly the same as Ohanian, Raffo, and Rogerson (2008). We see this as additional evidence that measuring the transfer rate is important for understanding the effects of fiscal policy on labor supply. We prefer our baseline estimates, because Table 8 shows that the implied change in labor supply for older workers when taxes are fully rebated is 150% of what is observed in the data, a decline we find implausibly large.

6.3 Reduced Form Comparison of Tax Wedges

We conclude with an empirical exercise to assess the value added of the generational accounts tax wedge relative to the one used in ORR. We run regressions of male hours per person by age and calendar year on both our tax measures and the ORR tax wedge. By including both tax wedges simultaneously, the regression results show the the additional explanatory power in each tax wedge beyond any common component of the two wedges.[26]

[23] In their paper, this result runs contrary to the data, as their data includes both men and women. Due to the rise in female labor force participation, aggregate hours slightly rise over this time period.

[24] A natural concern is that the model's inability to produce changes in labor supply for cohorts prior to 1900 may partially account for the differences in results. However, if we assume the model could explain the same amount of the labor supply decline for these cohorts, roughly a quarter, then the most we would expect this adjustment to explain is one percentage point of the difference in the final results.

[25] Details of how these adjustments are made are available in the notes to Table 8.

[26] As it turns out, the two tax wedges are highly correlated for any given age: 0.73 on average, with values close to 1 for some ages.

The regressions we run take the form

$$log\left[H(t,v)\right] = \beta_v + \beta_{GA} log[\tau_{GA}(t-v)] + \beta_{ORR} log[\tau_{ORR}(t)] + \varepsilon_{vt},$$

where $H(t,v)$ is hours worked per person at time t by a man of age v, $\tau_{GA}(t-v)$ is the generational accounts tax wedge faced by a man born in year $j = t - v$, and $\tau_{ORR}(t)$ is the ORR wedge at time t.[27]

The regressions are run with age fixed effects because the generational accounts tax wedge varies across birth cohorts (and thus ages), and our focus is on explaining time series variation and not life cycle variation in hours worked. As there are more birth cohorts than calendar years represented in our regression, there is more variation in the generational accounts than in calendar year tax wedges. Thus, we also report regressions within narrower age groups to show that our results are not driven by this additional variation.

What do we expect to learn from these regressions? A negative and significant coefficient on the generational accounts tax measure would indicate that, after controlling for the common component of the two tax wedges, either the addition of transfers or the forward-looking nature of these tax rates has a meaningful impact on labor supply. Given that the average tax rate matters at ages where the extensive margin decision is prominent, we would also reasonably expect the coefficient to larger in magnitude for older and younger workers.

Table 9 reports the results of these regressions using hours worked per person from the 1950 and 1960 Census and 1962-2000 CPS, our generational tax wedges, and those from ORR. We explicitly exclude the post-2000 fluctuations in hours worked per person to avoid any unusual labor supply responses due to recent cyclical downturns; however, the results are similar if we include this time period.

Even controlling for the common variation observed between these two tax wedges, our tax wedge has sizable explanatory power, both for all ages and within narrow age groups. And as suggested above, the generational accounts tax wedge has a particularly large coefficient for older and younger workers. Overall, the empirical results highlight that tax wedges measured from the generational accounts provide additional value for explaining changes in labor supply over time.

7 Conclusion

In sum, we have found that a model with shifts in demographics and changing fiscal policy across birth cohorts can account for nearly half of the decline in hours worked by men between 1948 and

[27]Regression results are robust to a variety of specifications, though we prefer log-log because that is more consistent with the first order conditions of these models. Further, given that the ORR tax wedge is much more volatile, we have checked to see using if a smoothed version of their tax wedge changes the results. It does not.

2000, and that the timing of the decline generated by the model closely mirrors the data. Changes in generational policy alone account for nearly 25% of this decline, both in the aggregate and within age groups. We have further highlighted some of the value added from using the generational accounts tax wedges relative to the benchmark tax wedge in the existing literature.

References

AARONSON, S., T. CAJNER, B. FALLICK, F. GALBIS-REIG, C. SMITH, AND W. WASCHER (2014): "Labor force participation: recent developments and future prospects," *Brookings Papers on Economic Activity*, 2014, 197–275.

AARONSON, S., B. FALLICK, A. FIGURA, J. F. PINGLE, AND W. L. WASCHER (2006): "The recent decline in the labor force participation rate and its implications for potential labor supply," *Brookings Papers on Economic Activity*, 2006, 69–154.

ALONSO-ORTIZ, J. (2014): "Social security and retirement across the OECD," *Journal of Economic Dynamics and Control*, 47, 300–316.

AUERBACH, A. J., J. GOKHALE, AND L. J. KOTLIKOFF (1991): "Generational Accounts: A Meaningful Alternative to Deficit Accounting," *Tax Policy and the Economy*, 55–110.

——— (1993): "Generational Accounts and Lifetime Tax Rates," *Economic Review of Federal Reserve Bank of Cleveland*, 29.

——— (1994): "Generational accounting: a meaningful way to evaluate fiscal policy," *The Journal of Economic Perspectives*, 73–94.

AUTOR, D. H. AND M. G. DUGGAN (2003): "The Rise in the Disability Rolls and the Decline in Unemployment," *Quarterly Journal of Economics*, 118.

BICK, A. AND N. FUCHS-SCHÜNDELN (2014): "Taxation and labor supply of married couples across countries: a macroeconomic analysis," *Manuscript, Arizona State University*.

CHETTY, R., A. GUREN, D. MANOLI, AND A. WEBER (2012): "Does indivisible labor explain the difference between micro and macro elasticities? A meta-analysis of extensive margin elasticities," in *NBER Macroeconomics Annual 2012, Volume 27*, University of Chicago Press, 1–56.

CONESA, J. C. AND C. GARRIGA (2015): "Intergenerational policy and the measurement of tax incidence," *European Economic Review*.

ECKSTEIN, Z. AND O. LIFSHITZ (2011): “Dynamic female labor supply,” *Econometrica*, 79, 1675–1726.

ELSBY, M. W. AND M. D. SHAPIRO (2012): “Why does trend growth affect equilibrium employment? A new explanation of an old puzzle,” *The American Economic Review*, 102, 1378–1413.

EROSA, A., L. FUSTER, AND G. KAMBOUROV (2012): “Labor supply and government programs: A cross-country analysis,” *Journal of Monetary Economics*, 59, 84–107.

FIORITO, R. AND G. ZANELLA (2012): “The anatomy of the aggregate labor supply elasticity,” *Review of Economic Dynamics*, 15, 171–187.

FISHER, E. O. (1995): “A New Way of Thinking about the Current Account,” *International Economic Review*, 36, 555–568.

FISHER, E. O. AND K. KASA (1997): “Generational accounting in open economies,” *Economic Review-Federal Reserve Bank of San Francisco*, 3, 34–46.

FLOOD, S., M. KING, S. RUGGLES, AND J. R. WARREN (2015): “Integrated Public Use Microdata Series, Current Population Survey: Version 4.0,” [Machine-readable database]. Minneapolis: University of Minnesota.

GOKHALE, J., B. R. PAGE, AND J. STURROCK (1999): “Generational accounts for the United States: an update,” in *Generational accounting around the world*, University of Chicago Press, 489–518.

GORRY, A. AND E. OBERFIELD (2012): “Optimal taxation over the life cycle,” *Review of Economic Dynamics*, 15, 551–572.

GUNER, N., R. KAYGUSUZ, AND G. VENTURA (2012a): “Taxation and household labour supply,” *Review of Economic Studies*, 79, 1113–1149.

——— (2012b): “Taxing women: A macroeconomic analysis,” *Journal of Monetary Economics*, 59, 111–128.

JUHN, C. (1992): “Decline of male labor market participation: The role of declining market opportunities,” *The Quarterly Journal of Economics*, 79–121.

JUHN, C. AND K. M. MURPHY (1997): “Wage Inequality and Family Labor Supply,” *Journal of Labor Economics*, 15, 72–97.

JUHN, C., K. M. MURPHY, AND R. H. TOPEL (1991): “Why Has the Natural Rate of Unemployment Increased over Time?” *Brookings Papers on Economic Activity*, 22, 75–142.

JUHN, C. AND S. POTTER (2006): "Changes in labor force participation in the United States," *The Journal of Economic Perspectives*, 20, 27–46.

KAYGUSUZ, R. (2010): "Taxes and female labor supply," *Review of Economic Dynamics*, 13, 725–741.

KEANE, M. AND R. ROGERSON (2012): "Micro and macro labor supply elasticities: A reassessment of conventional wisdom," *Journal of Economic Literature*, 50, 464–476.

KNOWLES, J. A. (2013): "Why are married men working so much? An aggregate analysis of intra-household bargaining and labour supply," *The Review of Economic Studies*, 80, 1055–1085.

KONG, Y.-C., B. RAVIKUMAR, AND G. VANDENBROUCKE (2014): "Explaining Differences in Life Cycle Earnings Across Cohorts," .

LAUN, T. AND J. WALLENIUS (2013a): "A life cycle model of health and retirement: The case of Swedish pension reform," *Journal of Public Economics*.

——— (2013b): "Social Insurance and Retirement: A Cross-Country Perspective," .

MCDANIEL, C. (2007): "Average tax rates on consumption, investment, labor and capital in the OECD 1950-2003," *Manuscript, Arizona State University*, 1960–2004.

——— (2011): "Forces shaping hours worked in the OECD, 1960-2004," *American Economic Journal: Macroeconomics*, 27–52.

MCGRATTAN, E. R. AND R. ROGERSON (2004): "Changes in hours worked, 1950-2000," *Federal Reserve Bank of Minneapolis Quarterly Review*, 28, 14–33.

MENDOZA, E. G., A. RAZIN, AND L. L. 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.

NATIONAL CENTER FOR EDUCATION STATISTICS (2009): *The Condition of Education 2015*, US Dept. of Education, Office of Educational Research and Improvement, National Center for Education Statistics.

NGAI, R. AND B. PETRONGOLO (2014): "Gender gaps and the rise of the service economy," .

OHANIAN, L., A. RAFFO, AND R. ROGERSON (2008): "Long-term changes in labor supply and taxes: Evidence from OECD countries, 1956–2004," *Journal of Monetary Economics*, 55, 1353–1362.

PARSONS, D. O. (1980): "The decline in male labor force participation," *The Journal of Political Economy*, 117–134.

PETERMAN, W. B. (2016): "Reconciling micro and macro estimates of the Frisch labor supply elasticity," *Economic Inquiry*, 54, 100–120.

PIKETTY, T. AND E. SAEZ (2007): "How progressive is the US federal tax system? A historical and international perspective," *The Journal of Economic Perspectives*, 21, 3–24.

PRESCOTT, E. C. (2004): "Why Do Americans Work So Much More Than Europeans?" *Federal Reserve Bank of Minneapolis Quarterly Review*, 28, 2–13.

RAGAN, K. S. (2013): "Taxes and time use: Fiscal policy in a household production model," *American Economic Journal: Macroeconomics*, 5, 168–192.

ROGERSON, R. (2006): "Understanding differences in hours worked," *Review of Economic Dynamics*, 9, 365–409.

——— (2007): "Taxation and market work: is Scandinavia an outlier?" *Economic theory*, 32, 59–85.

ROGERSON, R. AND J. WALLENIUS (2009): "Micro and macro elasticities in a life cycle model with taxes," *Journal of Economic theory*, 144, 2277–2292.

RUGGLES, S., K. GENADEK, R. GOEKEN, J. GROVER, AND M. SOBEK (2015): "Integrated Public Use Microdata Series, Version 6.0," [Machine-readable database]. Minneapolis: University of Minnesota.

RUPERT, P. AND G. ZANELLA (2015): "Revisiting wage, earnings, and hours profiles," *Journal of Monetary Economics*, 72, 114–130.

SMITH, C. L. (2012): "The impact of low-skilled immigration on the youth labor market," *Journal of Labor Economics*, 30, 55–89.

WALLENIUS, J. (2013): "Social security and cross-country differences in hours: A general equilibrium analysis," *Journal of Economic Dynamics and Control*, 37, 2466–2482.

Table 1: Percentage Change in Hours Worked For Men, 1948-2000

	Aggregate	16-24	25-34	35-44	45-54	55-64	65+
Hours per Person	-20	-32	-10	-11	-12	-27	-72
Employment per Person	-14	-14	-3	-6	-7	-22	-63

Data are smoothed with HP filter (6.25) to avoid cyclical sensitivity in computing percentage declines. Appendix A describes how we measure hours and employment per person.

Table 2: Model Predictions of the Percentage Change in Hours Worked for Men, 1948-2000

	Aggregate	16-24	25-34	35-44	45-54	55+
Model	-7.8	-17.8	-3.0	-2.9	-2.6	-23.5
Data	-17.8	-31.5	-10.4	-10.9	-11.9	-35.2
Frac. Explained	**0.44**	**0.57**	**0.29**	**0.27**	**0.22**	**0.67**

Data are smoothed with HP filter (6.25) to avoid cyclical sensitivity in computing percentage declines. Because tax and transfer rates only begin in 1900, the aggregate hours for men 55 and over are constructed assuming that there are no changes in hours worked for men aged 55 through 64 until 1960 and no changes in the hours worked for men over 65 until 1970. The same assumption is applied to the model output for consistency.

Table 3: Contribution of Age Demographic Changes (Age Group Level) to Percentage Change in Hours Worked by Men

		Total	Demographic Changes	Constant Demographics
1948-2000	Model	-7.8	0.0	-7.8
	Data	-17.8	-0.1	-17.7
	Frac. Explained	**0.44**	**-**	**0.44**
1948-1970	Model	-8.3	-3.5	-4.8
	Data	-11.5	-3.0	-8.6
	Frac. Explained	**0.72**	**1.17**	**0.56**

Because of the small contributions of population shares to the labor supply decline from 1948-2000, the percent explained by population shares is not computed. Totals do not exactly sum because of rounding.

Table 4: Contribution of Age Demographic Changes (Individual Age Year Level) to Percentage Change in Hours Worked by Men, 1948-2000

		Aggregate	16-24	25-34	35-44	45-54	55+
Model	Total	-7.8	-17.8	-3.0	-2.9	-2.6	-23.5
	Demographics	-3.9	-9.3	0.1	0.0	0.0	-18.2
	Labor Supply	-3.9	-8.5	-3.0	-2.9	-2.6	-5.3
Data	Total	-17.8	-31.5	-10.4	-10.9	-11.9	-35.2
	Demographics	-2.1	-2.9	0.3	0.0	0.1	-10.9
	Labor Supply	-15.7	-28.6	-10.7	-10.9	-12.0	-24.3
Frac. Explained by Labor Supply:							
of Total		**0.22**	**0.27**	**0.29**	**0.27**	**0.22**	**0.15**
of Labor Supply		**0.25**	**0.30**	**0.28**	**0.27**	**0.22**	**0.22**

We use microdata from the 1950 Census and the March CPS for the years 1962-2000 to compute hours per population by age group. With this data, and the shares of the population for each age year, described in Appendix A, it is possible to obtain the contribution of changing population shares for the years 1950-1970 and 1950-2000 using equation 2. We use these as an approximation for the changes occurring between 1948 and 1970 and between 1948 and 2000, as there is not data for the year 1948. Comparisons to the model's results from 1950-1970 and 1950-2000 produce very similar results. Totals do not exactly sum because of rounding.

Table 5: Contribution of Age Demographic Changes (Individual Age Year Level) to Percentage Change in Hours Worked by Men, 1948-1970

		Aggregate	16-24	25-34	35-44	45-54	55+
Model	Total	-8.3	-19.3	-1.0	-0.9	-1.0	-9.0
	Demographics	-6.9	-15.6	0.0	0.0	0.0	-7.8
	Labor Supply	-1.4	-3.7	-1.0	-0.9	-1.0	-1.2
Data	Total	-11.5	-26.0	-3.1	-3.9	-5.8	-9.6
	Demographics	-4.8	-5.8	0.0	0.0	0.0	-4.1
	Labor Supply	-6.7	-21.2	-3.1	-3.9	-5.8	-5.5
Frac. Explained by Labor Supply:							
of Total		**0.12**	**0.14**	**0.32**	**0.23**	**0.17**	**0.13**
of Labor Supply		**0.21**	**0.17**	**0.32**	**0.23**	**0.17**	**0.22**

See notes to Table 4.

Table 6: Percentage Change in Intensive and Extensive Margins of Hours Worked by Men (Holding Demographics Fixed), 1948-2000

		Aggregate	16-24	25-34	35-44	45-54	55+
Model	Total	-3.9	-8.5	-3.0	-2.9	-2.6	-5.3
	Extensive	-1.3	-5.9	0.0	0.0	0.0	-3.7
	Intensive	-2.6	-2.6	-3.0	-2.9	-2.6	-1.6
Data	Total	-15.7	-28.6	-10.7	-10.9	-12.0	-24.3
	Extensive	-10.9	-12.1	-3.4	-5.6	-7.1	-28.0
	Intensive	-4.8	-16.5	-7.3	-5.3	-4.9	3.7
Frac. Explained:							
Total		**0.25**	**0.30**	**0.28**	**0.27**	**0.22**	**0.22**
Extensive Margin		**0.12**	**0.49**	**0.00**	**0.00**	**0.00**	**0.13**
Intensive Margin		**0.54**	**0.16**	**0.41**	**0.55**	**0.53**	**-0.43**

Data limitations complicate computing the contribution of population changes to the decline in extensive and intensive margins of labor supply. Thus, we assume that demographic change affects these two margins symmetrically. We compute the contribution of each margin to the total decline in hours worked and then subtract off half the contribution of demographic change from this figure. Attempts to more precisely pinpoint the separate impact of demographic change on each margin in micro data yield very similar results. Details of how the contribution of demographic change is computed are available in the notes to Table 4. Totals do not exactly sum because of rounding.

Table 7: Percentage Change in Hours Worked by Men (Holding Demographics Fixed) Coming From Taxes and Transfers Separately,1948-2000

		Aggregate	16-24	25-34	35-44	45-54	55+
Model	Total	-3.9	-8.5	-3.0	-2.9	-2.6	-5.3
	Tax Changes	-0.8	-1.5	-0.6	-0.6	-0.6	-1.4
	Transfer Changes	-3.3	-6.3	-2.4	-2.4	-2.3	-4.8
Data	Total	-15.7	-28.6	-10.7	-10.9	-12.0	-24.3
Frac. of Data Explained:							
Using Only Taxes		**0.05**	**0.05**	**0.06**	**0.06**	**0.05**	**0.06**
Using Only Transfers		**0.21**	**0.22**	**0.22**	**0.22**	**0.19**	**0.20**

We vary taxes and transfers separately by allowing one policy to change while holding the other fixed at its level for the 1940 cohort.

Table 8: Percentage Changes in Hours Worked from Model Simulations with Modified Tax/Transfer Rates, 1948-2000

	1948-2000						1956-2003	
		Fixed Demographics					Total	Fixed Dem.
	Agg.	16-24	25-34	35-44	45-54	55+	Agg.	Agg.
Baseline	-4.9	-9.3	-3.4	-3.5	-3.6	-7.9	-6.4	-4.8
Adj. Taxes	-3.8	-7.0	-2.7	-2.8	-2.9	-6.0	-5.4	-3.7
Adj. Trans.	-5.8	-23.2	-2.8	-3.3	-3.5	-9.4	-7.9	-5.4
Full Trans.	-10.4	-24.8	-2.9	-4.3	-5.7	-38.3	-12.3	-9.5

We modify total tax and transfer rates using the detailed composition of the generational accounts for the birth cohort born in 1995, reported in Gokhale, Page, and Sturrock (1999). We then assume that the fraction of taxes due to non-labor taxes (everything but payroll and labor income) is constant over time and subtract this value from the tax rate. For transfers, we assume that the fraction of transfers due to educational expenditures is also constant over time, and add this fraction to the transfer rate. For results in the full rebate transfer simulation, the underlying wage profile is smoothed with an HP filter (6.25) to avoid multiple entry/exit dates in late ages due to a slight non-monotonicity in the wage profile around age 58.

Table 9: Log-log Regressions of Male Hours per Person on Tax Wedges from Generational Accounts and Ohanian, Raffo, and Rogerson (2008)

Dep Var: Hours per Person	All Ages	16-24	25-34	35-44	45-54	55-64	65+
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Generational Accounts	-0.54***	-0.53***	-0.30***	-0.18***	-0.01	-1.06***	-0.87***
	(0.04)	(0.10)	(0.02)	(0.01)	(0.02)	(0.10)	(0.11)
Conventional Tax Wedge	0.25***	-0.01	0.10	-0.05**	-0.29***	0.77***	1.06**
	(0.06)	(0.17)	(0.02)	(0.02)	(0.04)	(0.17)	(0.49)
R^2	0.98	0.98	0.71	0.58	0.59	0.90	0.89
N	2444	369	410	410	406	393	456

Robust standard errors in parentheses.

Figure 1: Hours Worked Per Person for Men in the United States, 1948-2014

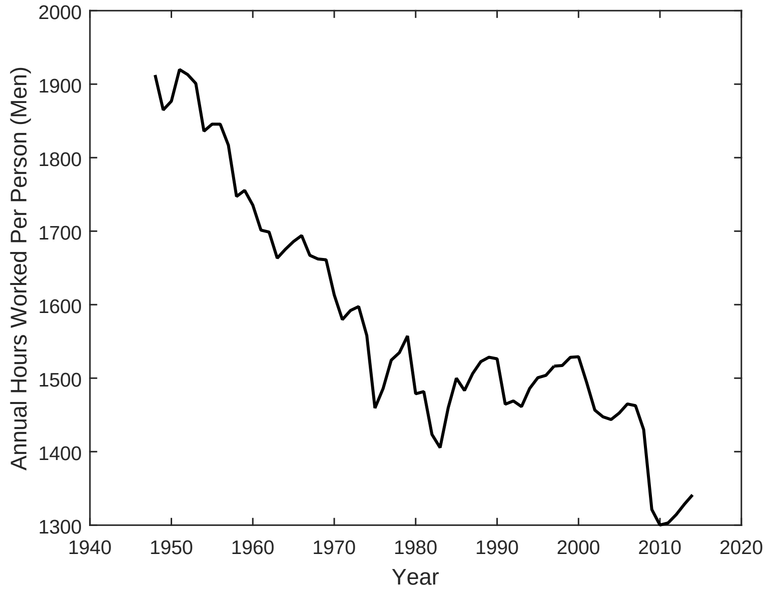

Annual hours worked are obtained based on estimates of weekly hours worked multiplied by 52. Further details regarding measurement are available in Appendix A.

Figure 2: Gross Lifetime Tax and Transfer Rates by Male Birth Cohort, 1900-1991

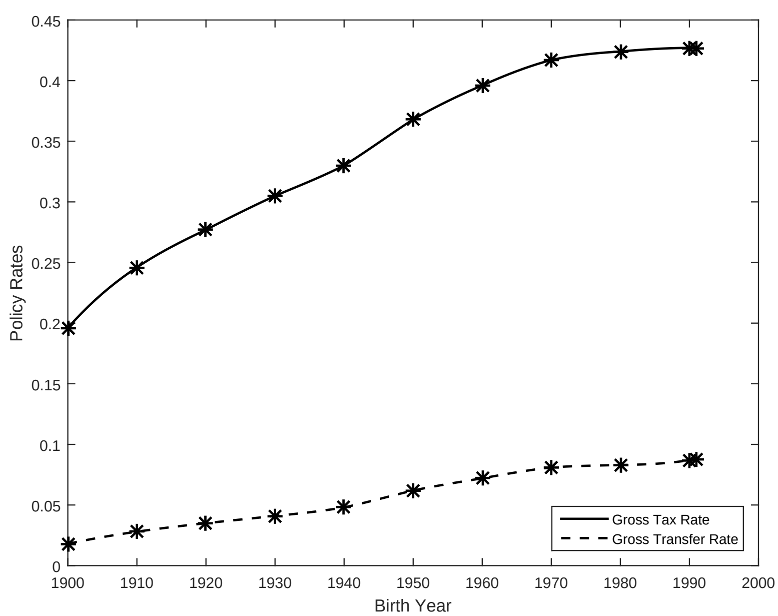

Gross tax and transfer rates are obtained by polynomial interpolation. Starred points represent the original tax and transfer rates in Auerbach, Gokhale, and Kotlikoff (1994).

Figure 3: Life Cycle Wage Profile

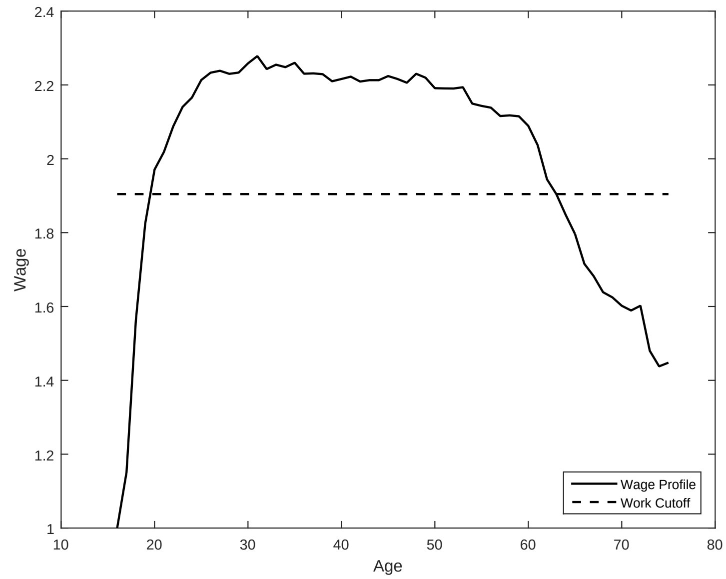

The solid line presents the life cycle wage profile calibrated to match the hours per worker data for the 1935-1945 cohorts. The wage at age 16 is normalized to 1. The dashed line is the reservation wage faced by the 1940 cohort.

Figure 4: Aggregate Hours Worked by Men, Model vs. Data

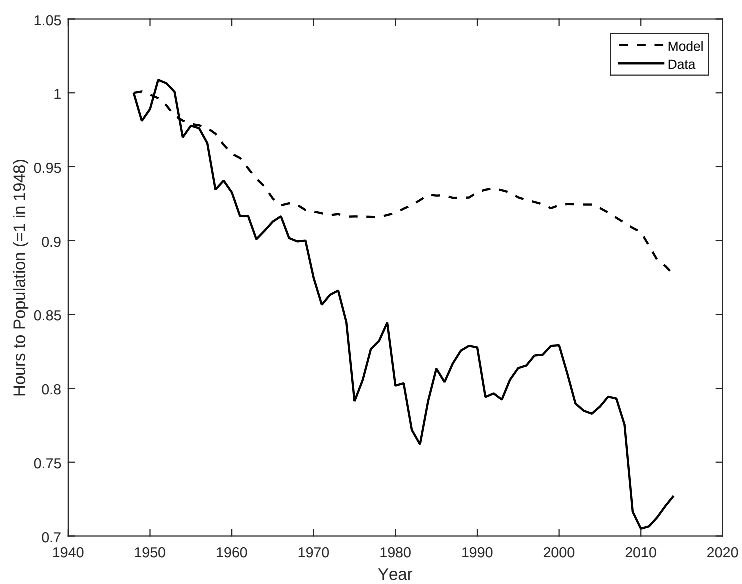

Data on taxes and transfers begin only in 1900. These series assume that there are no changes in hours worked for those aged 55 through 64 until 1960 and no such changes for men 65 and older until 1970.

Figure 5: Hours Worked by Men Across Age Groups

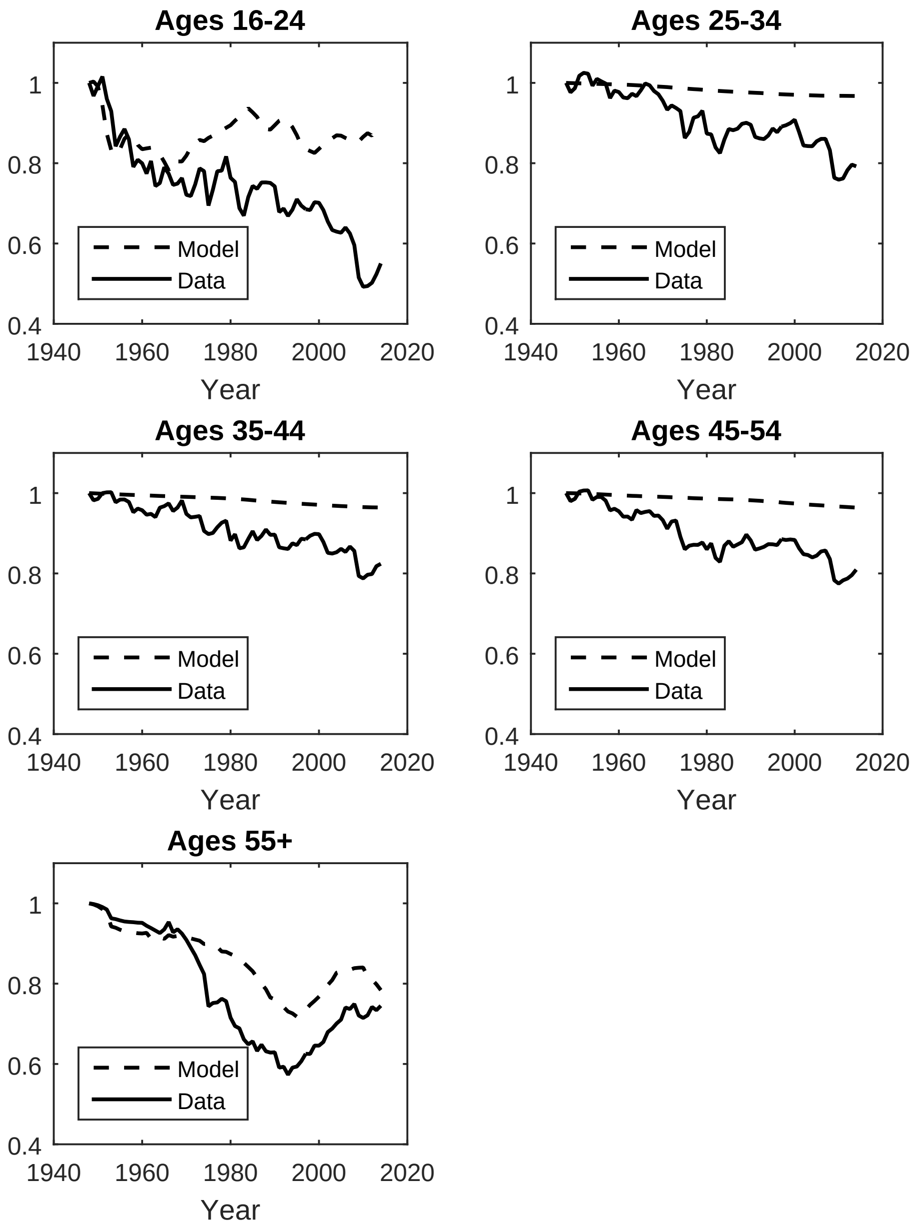

The y-axes denotes hours worked per person, where hours worked is normalized to 1 in the year 1948. See notes to Figure 4.

Figure 6: Contributions of Broad Demographic Changes to Hours Worked by Men

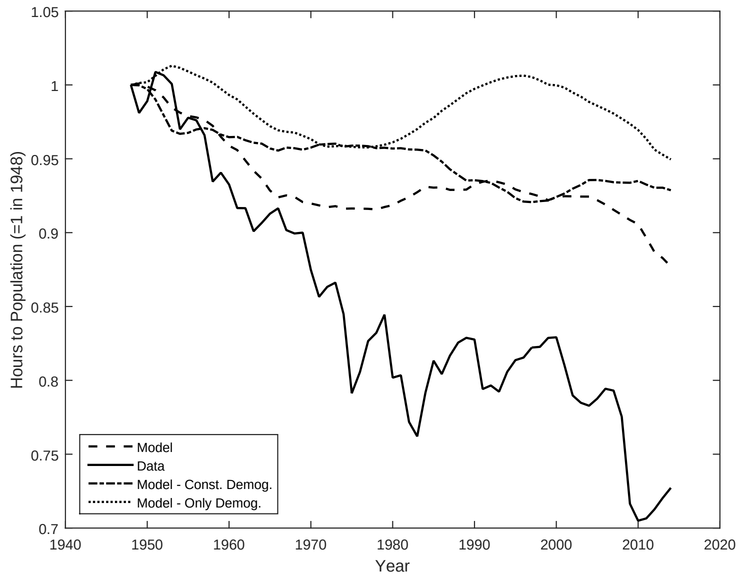

The dash-dotted line shows the evolution of hours worked from the model where the share of the population in each age bin is held fixed; the dotted line shows the evolution of hours worked from the model solely coming from changes in distribution across age bins. Details of this decomposition available in the text.

Figure 7: Contributions of Detailed Demographic Changes and Policy Changes to Hours Worked by Men

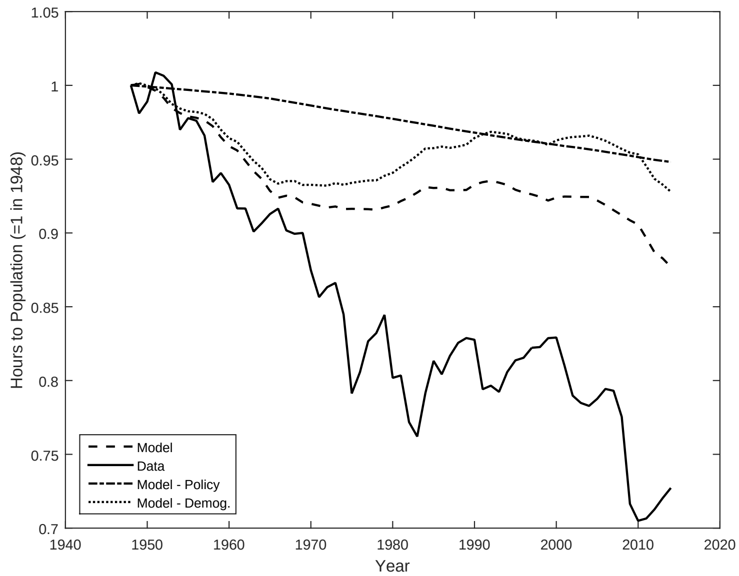

The dash-dotted line shows the evolution of hours worked from the model where the share of the population in each individual age is held fixed, and represents changes solely arising from policy; the dotted line shows the evolution of hours worked from the model solely coming from changes in the full age distribution of the population.

Figure 8: Comparison of Aggregated Generational Accounts Policy Measures with Ohanian, Raffo, and Rogerson (2008) Tax Wedge

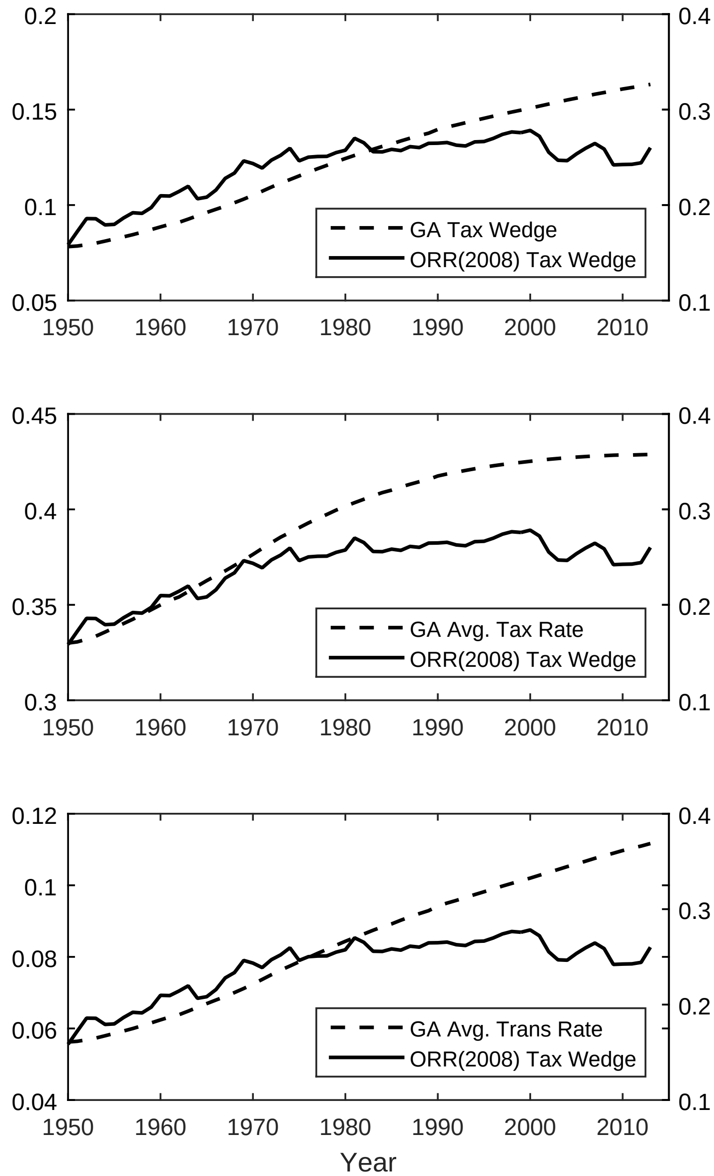

The Ohanian, Raffo, and Rogerson (2008) tax wedge is plotted on the right y-axis; all generational accounts policy figures are plotted on the left y-axis. The plot for the Ohanian, Raffo, and Rogerson (2008) tax wedge is identical in all cases. The top panel compares it to the total tax wedge from generational accounts, the middle panel to the average gross tax rate from generational accounts and the bottom panel to the average gross transfer rate from generational accounts.

Appendix A: Data Sources and Measurement

The data we use comes from the BLS compilation of monthly Current Population Survey data from 1948-2014, microdata from the decennial Census from 1940-1960, and the March Supplement to the CPS from 1962-2014, obtained from IPUMS (Flood et al. (2015), Ruggles et al. (2015)). All data is restricted to observations on men only.

The main series of interest is hours per population by detailed age group and year. There is no pre-existing source of data for these for the entire postwar period, so we construct them by combining data on employment per person and hours per employed person. For aggregation purposes, we are also interested in measuring the fraction of the male population in each group across time. Our methodology extends upon that of McGrattan and Rogerson (2004), who are interested in similar data, but only collect it by decade.

Measuring Employment Per Population

The Bureau of Labor Statistics compiles data from the monthly Current Population Survey on employment and population by gender for the age groupings 16-17, 18-19, 20-24, 25-34, 35-44, 45-54, 55-64, and 65+. This data is available from 1948 to the present. We aggregate the monthly data to an annual frequency by averaging. These data are used to construct the hours per population series in the aggregate and for the age groups, 16-24, 25-34, 35-44, 45-54, and 55+.

In Section 5.2, we decompose changes in hours worked into demographic changes and within age group changes at a much more detailed level than is available in the BLS data - individual age years. We use data from the 1950 Census and the CPS from 1962-2014 to construct employment per person for individual ages across time. We restrict the data to men 16 or older who are in the non-institutionalized civilian population and omit observations with imputed hours worked.

Measuring Hours Per Worker

To generate hours per population for age groups and in the aggregate, we need data on average hours per worker. In both the Census and the CPS, survey respondents report their hours worked for the prior week. We assume that this response provides a measure of the average hours per week in a given year. Since 1978, respondents in the CPS report their usual hours worked for the prior year which gives a potentially more accurate measure. However, since we combine this data with monthly employment status, using the past week's hours worked is more consistent with our measure of employment per population.

Hours per worker are thus measured from these data sources, with identical sample restrictions as those described above for employment per population. For the Census data in 1960, respondents only report their hours worked in the past week in a broad interval. We use detailed hours data from 1950 to impute the average hours worked within each interval to obtain a more accurate measure of hours worked in 1960. However, this still leaves gaps in the data for the years 1948-1949, 1951-1959, and 1961. We thus interpolate hours worked per worker based on data in 1940, 1950, 1960, and 1962-1965 at the age group level consistent with the data from the BLS using a third order polynomial to obtain measures of hours per worker for these missing years.[28]

Measurement of hours per worker at the individual age year level is also conducted in this way for the decomposition of intensive margin changes in labor supply in Section 5.3. We also use this data on hours per worker to construct the baseline wage profile and use data on the cohorts born between 1935-1945 to construct an approximate average profile centered around the 1940 cohort, for which we have the most data.

Measuring Hours Per Population

Hours per population for age groups and birth cohorts at an annual frequency are the product of employment per person and hours per employed person.

Measuring Population Shares

The share of the male population at each age in each year is needed for aggregating the model's output and for exercises decomposing the contribution of demographic change. Data on population weight by age group is available from the BLS data from 1948-2014, but only by decadal age group. Thus, we use the Census and CPS microdata to construct these population shares for the years 1950, 1960, and 1962-2014. We then linearly interpolate changes in shares while scaling them to match the shares observed by decadal age group in the BLS data. As top coding of ages has changed over time, we scale individual year population weights for ages 65-80 to match the total share of the population in the 65+ age group.

[28]For hours per worker used to construct age grouping totals and the aggregate, we also scale hours per worker from the Census to be consistent with CPS values based on the observed discrepancy in the 1970 Census and CPS results. This difference arises because of sampling and weighting schemes. This adjustment is not material for the results, nor is it sensitive to the Census or CPS year.

Appendix B: Evidence on Female Labor Supply and Male Hours Worked

Perhaps the decline in male labor supply is related to the simultaneous rise in female labor since 1948. Figure B.1 shows the changes in male and female labor force participation.While our paper focuses on hours worked, we begin by analyzing these simultaneous trends as (1) this is the outcome studied in Juhn (1992) and (2) while not reported, the entire decline in the extensive margin of male labor supply is driven by this decline in male labor force participation.

There might be a connection between male and female labor force trends, but Juhn (1992) casts doubts on the idea that this connectionis on the supply side. She argues that female employment in households with male non-participants has not risen at all over that period and also that earnings for those females have increased far less than for women at large. Using data from the March CPS, we corroborate these findings in Figures B.2 through B.4.[29] Figure B.2 shows the decline in male labor force participation opposite the change in hours worked by women in households with male non-participants. While there has been an increase in the hours worked by these women, it begins in the 1980s, well *after* most of the decline in male labor force participation. Eckstein and Lifshitz (2011) shows that the rise in female labor supply has been concentrated among married women. Hence, we analyze hours worked by married women in households with married male non-participants in Figure B.3. The timing of the rise in the hours worked by these women is even later, not beginning until nearly the 1990s. Finally, Figure B.4 compares the increase in earnings for all women with the increase in earnings for women in households with male non-participants. Although incomes have risen for women in households with male non-participants, the timing again does not align with the decline in male labor force participation.

We also present evidence on the decline in hours worked per person for married and non-married men from the March CPS. If there are changes in the costs and returns to market work for women, we might expect to see different trends in the hours worked by married and non-married men, as the labor supply of the latter would be less affected by these trends.[30] To the contrary, Figure B.5 shows the decline in hours worked from 1962-2014 for married and non-married men and finds that the declines are nearly identical. Incorporating observations from the 1950 Census, hours worked per person for married men fell by 12.1% and hours worked per person for non-married men fell by 14.7%. The declines are very similar for both groups.

While changing trends of women's participation in the workforce may have affected male labor

[29]Data is on the population 16 and older. Incomes are deflated using the price index for personal consumption expenditures.

[30]While it is possible that non-married men may be cohabitating with women and thus also subject to these changes, the prevalence of cohabitation is much more recent phenomenon and unlikely to affect more than a small group of men over the pre-1980s period, where most of the decline is observed.

supply decisions, all this evidence seems to to suggest that the decline in male hours worked since 1948 is not linked to the increase in women's supply of labor.

Figure B.1: Labor Force Participation Rates, Men and Women

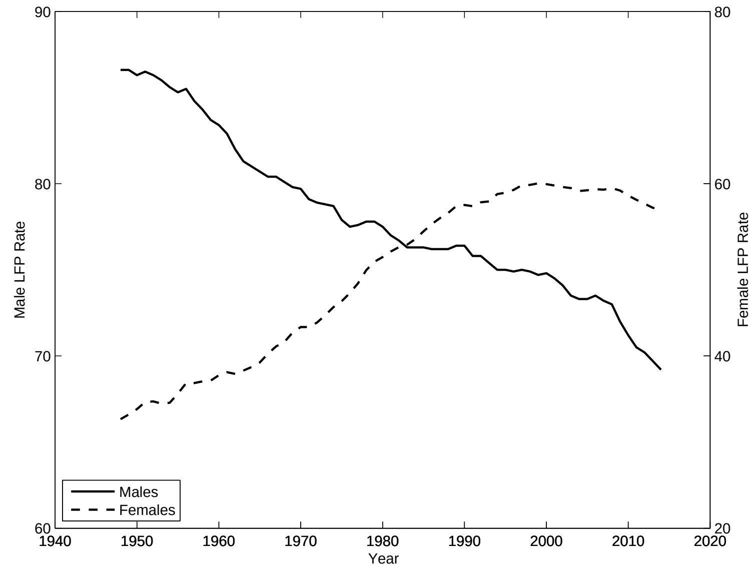

Data from BLS annual averages of monthly Employment and Earnings reports. Male labor force participation is the solid line (left axis); female labor force participation is the dashed line (right axis). Separate axes are used to emphasize the relative timing in labor force participation rates for men and women.

Figure B.2: Male Labor Force Participation and Hours per Woman in Households with Non-participating Males

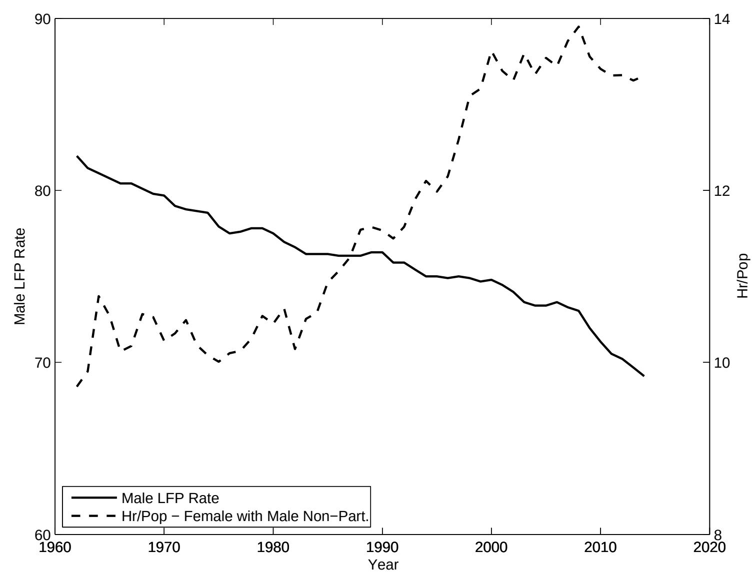

Data is from the March CPS.

Figure B.3: Male Labor Force Participation and Hours per Married Woman in Households with Non-participating Married Males

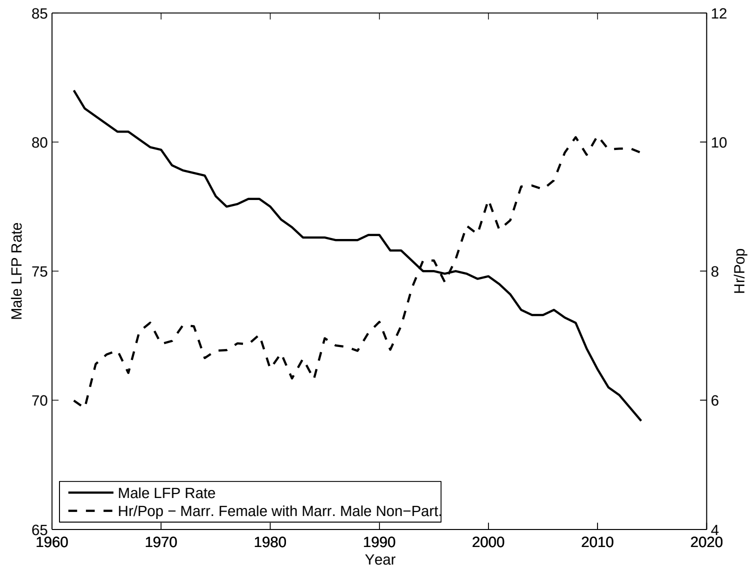

Data is from the March CPS.

Figure B.4: Average Real Incomes for All and Married Women, Total Average and for those in Household with Non-Participating Males

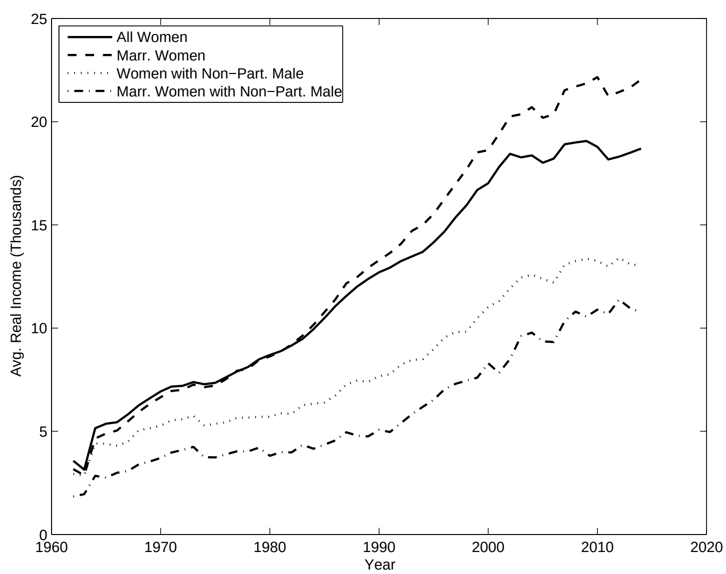

Data is from the March CPS. Real incomes are computed by dividing by the price index for personal consumption expenditures.

Figure B.5: Hours Per Person for Married and Non-Married Men

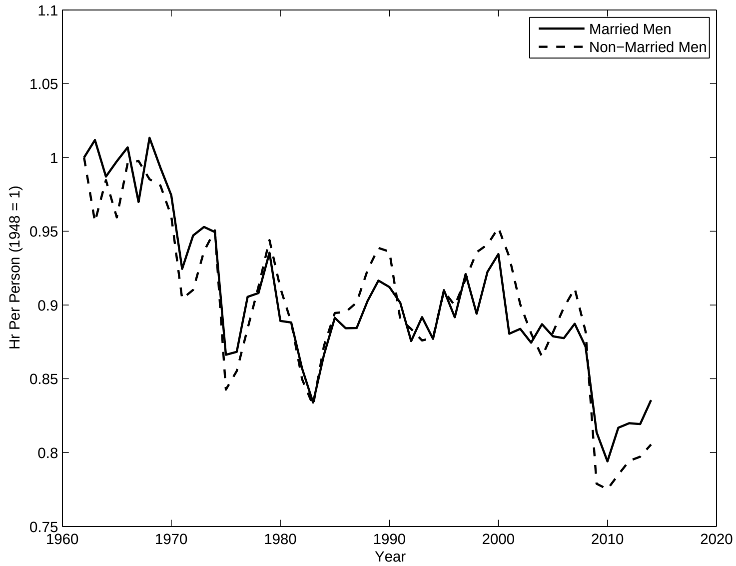

Data is from the March CPS. Hours per person (per week) for married men is 35.6 in 1962; for non-married men it is 26.8.

Appendix C: Robustness Checks

In this Appendix, we present several robustness checks on our primary results. We consider the following modifications: (1) an alternate calibration of the wage profile using measured wages, and potentially allowing those wages to vary over time; (2) using total generational tax and transfer rates instead of male-specific ones; and (3) variation in the parameters governing the disutility of labor, χ and γ.

C.1 Alternative Wage Calibration

In our baseline calibration, we calibrated the wage profile to match the life cycle hours worked behavior of the 1940 cohort. An alternative strategy would be to measure wages directly from the data and then allow the scale parameter governing the disutility of work, α, to vary over the life cycle to ensure the model generates a life cycle hours worked profile consistent with the data. In this case, the results will be identical, since this calibration strategy is isomorphic to our baseline case. However, it is now more natural to think about allowing wages to vary with successive birth cohorts, and we consider the implications of such variation for our results in this section.

The primary challenge with this alternative calibration strategy is simply measuring wages in the first place. A naive approach would be to construct a synthetic wage profile from a single cross-section of wages and age. For example, if working with the wage observed for a single cross-section of the population in the year 1980, the wage profile is constructed by assuming that the wages received by a 65 year old, born in 1915, will be the same as what a 16 year old in 1980, born in 1964, will face when he is 65 years old, in the year 2029. As discussed in Rupert and Zanella (2015), such wage profiles are likely to be biased because of variation in productivity growth over the lifetime.

A preferable approach for measuring wages is to either used repeated cross-sections or actual panel data to construct the wage profile for a given cohort. As discussed in Rupert and Zanella (2015), there are significant data limitations regarding the observation of wages for different cohorts. PSID data has reliable wage annual measures from 1968-1997 and biannually thereafter, and data from the CPS has reliable annual wage data from 1976 to the present. Thus, even using repeated cross-sections or panel data will still require substantial imputation of wages for ages where there is no data.

Using PSID and CPS data reported in Rupert and Zanella (2015), Figure C.1 reports the average wage profile for men born between the years 1937 and 1946, as well as a 2nd order fractional polynomial fitted on log wage data to project actual wages for ages where wages are not observed. Both wage profiles show minimal wage declines in old age, suggesting that some other factor must be changing late in the life cycle that induces lower hours worked. Figure C.2 reports how the scale parameter in labor disutility, α, must vary with age to generate life cycle hours worked behavior consistent with what is observed for the 1940 cohort given either of these wage profiles. Because the CPS data has no data for workers under 30, the fitted wage trend is much flatter early in the life cycle, meaning the implied α values for younger ages are much higher. But in both cases, the level of disutility of work increases throughout much of the lifetime, especially in the late stages of life.

Figure C.1: Average Wage Profiles for Cohorts Born Between 1937 and 1946

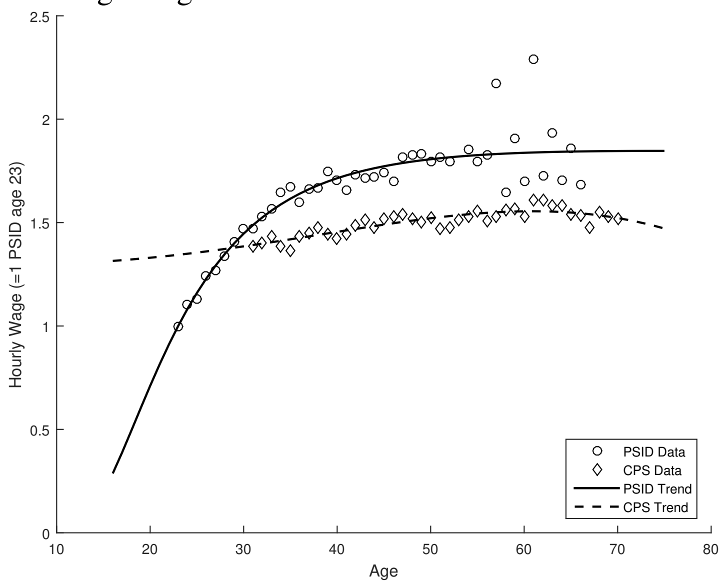

Data is from Rupert and Zanella (2015). Wages are normalized to 1 for 23 year olds in the PSID data.

Figure C.2: Implied Labor Disutility Parameters Over the Life Cycle Consistent with Hours Worked

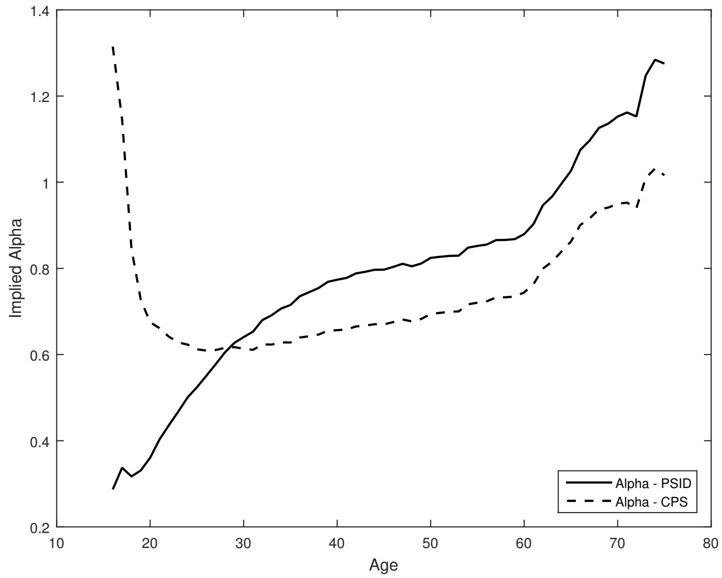

Implied α values are computed as measured wages at a given age divided by the observed hours per worker at that age.

Pursuing this alternative calibration strategy yields identical results for changes in hours worked, as this calibration is isomorphic to our baseline strategy. A natural question, however, is how our results would change if we allowed wage profiles to vary across different birth cohorts. Before discussing how this would impact our results, we emphasize that the degree of imputation required for measuring life cycle wage profiles for cohorts born from 1900 to 1990 is significant. For workers under the age of 24, there is no reliable data observed for cohorts born prior to the year 1944, meaning that nearly half of the cohorts studied will have imputed wage profiles for the young stages of life. We further emphasize that in our life cycle model, the decision regarding the extensive margin for work depends on wages faced at all ages of the life cycle, and thus will be highly sensitive to these imputed values.

To slightly improve the imputation problem, we construct crude measures of hourly earnings from the CPS and Census data by measuring hourly earnings as the ratio of average weekly earnings and average hours per week. This generates measures of wages for the years 1940, 1950, 1960 and 1963 to the present. Given these wages and a multivariate fractional polynomial imputation of all remaining unobserved data, we obtain wage profiles for each cohort born between 1900 and 1990. The growth in wages at each age between the 1900 and 1990 cohort is plotted in Figure C.3. Based on these measures of wages, wage growth has been fastest for younger workers, which is consistent with life cycle patterns in earnings growth discussed in Kong, Ravikumar, and Vandenbroucke (2014).

Figure C.3: Real Wage Growth Across the Life Cycle Between 1900 and 1990 Birth Cohorts

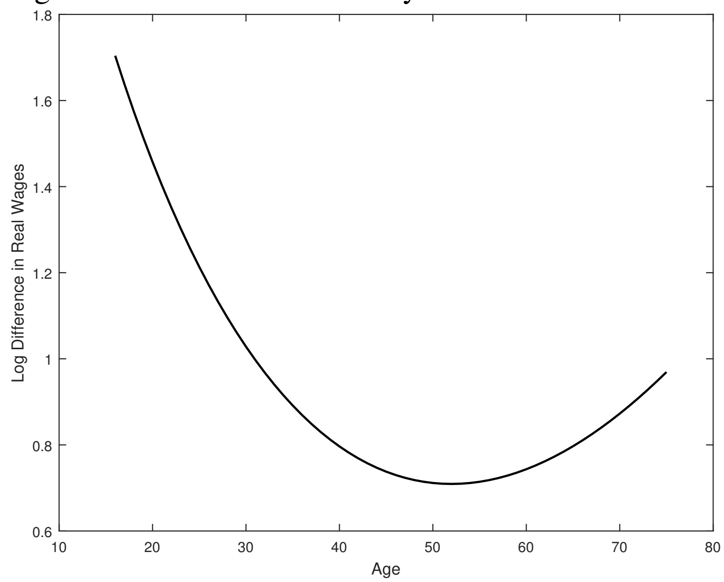

Table C.1 reports our results by age group where these wage profiles vary over the life cycle opposite our baseline results and the data. Allowing for time-varying wage profiles has minimal impact on our aggregate results, but dramatic impact on the results by age group. The model predicts a dramatic increase in the hours worked by young workers, contrary to what is observed in the data. With higher levels of productivity younger in life and slower productivity growth in much of the prime age years, declines in hours worked are greater for both prime-aged and older workers.

Table C.1: Model Predictions for Changes in Hours Worked with Changing Wage Profiles

%ΔHr/Pop	Total	16-24	25-34	35-44	45-54	55+
Model (Base)	-7.8	-17.8	-3.0	-2.9	-2.6	-23.5
Changing Wage	-8.9	61.5	-1.6	-13.6	-16.8	-39.4

We again emphasize that these results are heavily dependent on significant imputation of wage growth and thus highly sensitive to the nature of that imputation. Further, given that the exact mapping from wages to hours over the life cycle is yet a puzzle, more sophisticated theories of this relationship may generate very different results. There may also be selection concerns regarding which young workers choose to work that may cause wage measurement to be misleading. All in all, we find that the aggregate results are not substantially changed from allowing these wage profiles to vary, but view these results tentatively, given the above concerns.

C.2 Results Using Total Generational Tax and Transfer Rates

Our baseline results use the generational accounts lifetime tax and transfer rates for males. Here we show that similar results obtain when we use total lifetime tax rates computed from total income, taxes and transfers for each birth cohort. For these measures, we use the rates reported in Gokhale, Page, and Sturrock (1999), which have been updated through 1995.

Table C.2 shows the changes in hours worked generated by the model in the aggregate and for each age group when total tax and transfer rates are used. The results from using total measures of taxes and transfers are very similar to those from using male policy figures only, and actually show slightly larger declines in hours worked.

Table C.2: Model Predictions of the Percentage Change in Hours Worked for Men, Male Policies vs. Total Policies

	Aggregate	16-24	25-34	35-44	45-54	55+
Male Policies	-7.8	-17.8	-3.0	-2.9	-2.6	-23.5
Total Policies	-8.8	-18.6	-3.3	-3.5	-3.6	-25.9

C.3 Variation in Labor Disutility Parameters

Our baseline calibration assumes that the intensive margin labor disutility parameter, γ, is equal to 1. This parameter is critical for the elasticity of labor supply along the intensive margin, for which there are a wide range of estimates - see Keane and Rogerson (2012), Chetty et al. (2012), Fiorito and Zanella (2012) and Peterman (2016). Thus, in Table C.3, we present the model's results by age group for fixed demographic declines for different values of this parameter, covering a wide range of estimates in the literature. Variation in γ has minimal impact for the aggregate decline in hours worked, but significant impacts on which age groups experience reduce their hours the most. With a lower value of γ, agents have much higher elasticities of labor supply along the intensive margin, meaning changes in policy induce greater changes in hours worked for prime-aged males.

Our baseline calibration chose χ, the fixed cost of working, to generate a working life of 44 years for the 1940 cohort. Table C.4 shows robustness to varying the working life target used for calibrating χ. Variation in χ has minimal impact for the aggregate results. On the other hand, different values of χ have different implications for changes in hours worked across age groups. Part of the reason for this has to do with the steepness of the wage profile around that cutoff date, which varies with χ. The other reason, which primarily impacts young workers, is that with larger values of χ, the entry age draws closer to 24. The closer the entry age is to 24, the more the extensive margin will impact percent changes in hours worked for this young age group.

Table C.3: Robustness of Model Results from 1948-2000 to Variation in Intensive Margin Labor Disutility Parameter, γ

	Fraction Explained of Fixed Demographic (Age Year Level) Hours Decline					
	Aggregate	16-24	25-34	35-44	45-54	55+
$\gamma = 0.25$	0.36	0.26	0.56	0.52	0.42	0.18
$\gamma = 0.5$	0.31	0.28	0.43	0.40	0.32	0.20
$\gamma = 1$	0.25	0.30	0.28	0.27	0.22	0.22
$\gamma = 2$	0.21	0.31	0.17	0.16	0.13	0.23
$\gamma = 4$	0.18	0.33	0.09	0.09	0.07	0.25

See notes to Table 1 and 2.

Table C.4: Robustness of Model Results from 1948-2000 to Variation in Extensive Margin Labor Disutility Parameter, χ

	Fraction Explained of Fixed Demographic (Age Year Level) Hours Decline					
	Aggregate	16-24	25-34	35-44	45-54	55+
Work Life = 40y,$\chi = 0.56$	0.26	0.68	0.22	0.20	0.16	0.15
Work Life = 42y,$\chi = 0.50$	0.29	0.50	0.25	0.24	0.20	0.27
Work Life = 44y, $\chi = 0.45$	0.25	0.30	0.28	0.27	0.22	0.22
Work life = 46y, $\chi = 0.40$	0.24	0.25	0.31	0.29	0.24	0.15
Work life = 48y, $\chi = 0.35$	0.25	0.21	0.32	0.30	0.24	0.20

See notes to Table 1 and 2.