Male Labor Supply and Generational Fiscal Policy*

Christian vom Lehn† — Brigham Young University

Aspen Gorry‡ — Utah State University

Eric O'N. Fisher§ — California Polytechnic State University

April 15, 2016

Abstract

Between 1948 and 2000, hours worked per man 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 birth-year cohort. Changes in age demographics and fiscal policy together account for roughly half of the decline in hours worked. Policy alone explains approximately thirty percent, both in the aggregate and for different 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 Frank Caliendo and Victor Rios-Rull. We also thank participants in seminars at the Western Economics Association Meetings, BYU-USU Macro Workshop, and the Federal Reserve Bank at Philadelphia for their helpful comments.

†Department of Economics, 130 Faculty Office Building, Provo, UT 84602. cvomlehn@byu.edu.

‡Corresponding author. Department of Economics and Finance, 3565 Old Main Hill, Logan, UT 84322-3565. aspen.gorry@gmail.com.

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

1 Introduction

Hours worked per man in the United States declined by twenty percent between 1948 and 2000. As shown in Table 1, this decline has been concentrated among the young and the old, where the decision of whether or not to work is prominent. There have also been significant changes in tax and transfer policies over this period. The generational accounting literature, pioneered by Auerbach, Gokhale, and Kotlikoff (1991), provides estimates of how these policies have changed over time for different birth-year cohorts. In this paper, we assess how much of the decline in male labor supply can be explained by changes in generational fiscal policy.

We use a version of the life cycle labor supply model of Rogerson and Wallenius (2009), which features both intensive and extensive labor supply decisions, to construct life cycle labor supply for each birth cohort. Variation in a cohort's labor supply by age is determined by an exogenous age-dependent productivity profile that is calibrated to match a life cycle profile of hours worked. To compare our predicted labor supply profiles with the data, we aggregate labor supply across cohorts to construct total hours worked. 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.

Applying lifetime tax and transfer rates from the generational accounts provides a new perspective on how fiscal policy affects labor supply. These rates 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.[1] Moreover, generational accounts 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, using lifetime tax rates may better account for the timing of how changes in fiscal policy

[1] We contrast generational accounting measures of taxes and transfers with the average annual tax rates that are typically used in macro studies of taxes on labor supply in section 3. While standard models of labor supply imply that the marginal tax rate is important for individual labor supply decisions, our life cycle model implies that the average tax rate is also important for an individual's decision to work at each date.

impact aggregate hours worked.[2]

The combination of changes in fiscal policy and the changes in the age distribution of the population together account for roughly half of the decline in hours worked per man from 1948 through 2000. Changes in policy alone account for roughly 30% of the decline. Policy changes also explain about 30% of the decline in hours worked for each age group, including workers aged 16 through 24 and over 55, where the declines have been largest. Incorporating the extensive margin of labor supply is particularly important for explaining changes in labor supply by the young and the old. We find that the extensive margin generates roughly two thirds of the policy-driven decline for workers under 24 and half of that for workers over 55. These declines at the extensive margin for young and old workers together account for roughly 30% of the total decline in aggregate hours worked per man generated by the model.

For context, it is useful to compare our results to those of Prescott (2004) and Ohanian, Raffo, and Rogerson (2008), who use the first order conditions of a standard growth model to assess the contribution of tax changes to changes in hours worked across OECD countries.[3] One contribution of our approach is that generational accounts pinpoint the fraction of tax revenues rebated to the household, whereas previous studies have assumed that all tax revenues are rebated. As Rogerson (2006, 2007) points out, the extent to which tax revenues are rebated to the household has a strong influence on how taxes impact labor supply in growth models. In particular, with balanced growth preferences and zero rebates, taxes have no effect on labor supply as income and substitution effects exactly offset.[4] An additional benefit of our approach is that, as shown in Rogerson and Wallenius (2009), the inclusion of the extensive margin reconciles differing values of the Frisch elasticity that arise in micro and macro studies. Therefore, our results also contrast with the existing literature by not requiring a parameterization of labor disutility that implies a Frisch elasticity at odds with

[2]Auerbach, Kotlikoff, and Koehler (2016) have also recently argued that current year tax rates may not adequately capture effects that fiscal policy have on discouraging work.

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

[4]However, as we discuss in Section 6, if instead of measuring transfers from the generational accounts we assume that all tax revenues are transferred back to households as is done by Prescott (2004), our results are quantitatively similar to previous studies.

micro studies. In general, our results are not very sensitive to varying parameter values for the disutility of work.

In focusing on the decline in male labor supply, we abstract from the increase in female labor supply that has occurred since 1948. This decision is motivated by evidence from Juhn (1992), who shows that there has not been an increase in female employment in households with men outside the labor force, even as the number of these men has increased. We reproduce this evidence 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, these changes do not appear to be generating the observed decline in hours worked by men.[5] Because our model describes labor supply and not labor demand, we feel this evidence justifies our focus on men. Further, as our measures of policy are not gender-specific, we would expect similar outcomes for female labor supply based on our model. Thus, while female labor supply has increased, potential explanations for this pattern must account for even more than the increase as measured in the data, as female labor supply would have otherwise declined over this period.

To our knowledge, this is the first paper to study the decline in male labor supply in a quantitative life cycle model. 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. Earlier empirical work has also studied the decline in male labor force participation and its relation to changes in fiscal policy. Parsons (1980) shows a strong cross-sectional correlation between the value of potential government transfers and non-participation and that these estimates coupled with trends in the generosity of transfers can account for the decline in male labor supply. Juhn (1992) argues that a combination of real wage changes and transfer programs can explain much of

[5] A number of recent papers have studied the role of household decisions and taxes in determining male and female labor supply. See Kaygusuz (2010), Guner, Kaygusuz, and Ventura (2012a,b), and Bick and Fuchs-Schündeln (2014).

the decline, with changes in transfer programs showing more promise as an explanation prior to the early 1970s.

The remainder of the paper is laid out as follows. Section 2 extends the model of Rogerson and Wallenius (2009) to cohorts born in each year who face different lifetime tax and transfer rates. Section 3 describes how our estimates of lifetime tax and transfer policies are obtained from the generational accounting literature, and discusses the benefits and drawbacks of these measures. Section 4 describes the calibration strategy and Section 5 presents the results of our quantitative exercises. Section 6 discusses possible variations on the generational accounting tax and transfer rates, compares our work to the relevant literature, and concludes.

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.[6] 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

[6]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.

fixed cost of commuting or getting ready for work.[7] 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,

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

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.[8] 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

[7] 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.

[8] The assumption of a lifetime budget constraint implies the presence of complete securities markets, allowing for free transfer of resources across the agent's lifetime.

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 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.[9] The generational accounts are also useful for understanding how the burden of fiscal policy is distributed across birth cohorts. 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.[10]

This insight can be seen by writing down the government's intertemporal budget constraint at a point in time t, given by,

$$\sum_{j=t-D}^{t} N_j + \sum_{j=t+1}^{\infty} N_j = \sum_{s=t}^{\infty} G_s + B(t).$$

G_s is the present value at time t of government consumption in year s and $B(t)$ is current govern-

[9] 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.

[10] Fisher (1995) shows that these ideas are true also in economies with incomplete financial markets, and Fisher and Kasa (1997) extend them a model of the open eocnomy where the crowding out of capital may also occur.

ment debt. This formulation assumes that individuals live for D years. N_j is the present value at time t of net taxes that fall on cohort j. It is customary to differentiate between those currently alive and those not yet born with index $j > t$.

The primary objects of interest in generational accounting are the net payments of each generation to the government, N_j, which are often termed the generational accounts. The generational accounts are simply the difference of the present value of cohort $j's$ total tax payments, P_j, and the present value of total transfers received by that cohort, T_j:

$$N_j = P_j - T_j. \tag{1}$$

Auerbach, Gokhale, and Kotlikoff (1993) and Gokhale, Page, and Sturrock (1999) use generational accounts to construct lifetime tax and transfer rates by birth cohort.[11] These rates are computed by dividing each generation's present value of tax payments and transfers at the time of birth by the present value of lifetime labor income at birth, Y_j. The lifetime gross tax rate τ_j and gross transfer rate θ_j for birth cohort j are:

$$\begin{aligned} \tau_j &= \frac{P_j}{Y_j} \quad \text{and} \\ \theta_j &= \frac{T_j}{Y_j}. \end{aligned}$$

Constructing these tax rates requires data on the taxes paid, transfers received, and labor income earned by each cohort at each point in time. Gokhale, Page, and Sturrock (1999) use micro data to assess the distribution of payments, transfers, and income in each year across 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 primarily 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. For cohorts still alive at the

[11]Some generational accounts provide separate estimates of tax and transfer rates for males and females, as in Auerbach, Gokhale, and Kotlikoff (1993). However, interpretation of these gender-specific policy rates is difficult, as several components of taxes and transfers are only measured at the household level.

time these tax rates are computed, future taxes and transfers are calculated using 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, and projections of fiscal policy based upon Congressional Budget Office forecasts. Additional details are available in Auerbach, Gokhale, and Kotlikoff (1993).

Gross lifetime tax and transfer rates for each cohort from 1900 to 1995 are plotted in Figure 1. The original figures in Gokhale, Page, and Sturrock (1999) are only reported by decade starting in 1900; the annual measures shown in Figure 1 are obtained by polynomial interpolation. From 1900 to 1950, there were substantial increases in these rates. And although the gross tax rate has fallen some for more recent cohorts, the transfer rate has continued to increase. Notably, the transfer rate is substantially lower than the tax rate. This is in part because 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.[12]

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 (2011), who constructs the tax measures used in Ohanian, Raffo, and Rogerson (2008). Building on the approaches of Prescott (2004) and Mendoza, Razin, and Tesar (1994), McDaniel (2011) constructs an average labor tax rate by computing aggregate taxes paid in each calendar year and then dividing by aggregate labor income.[13]

It is helpful to recall that in both our life cycle model and also the standard growth model

[12]In Section 6, we consider how variation in the transfer rate affects our results by allowing for education expenditures to enter into the analysis.

[13]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 in McDaniel (2011).

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 because of the extensive margin. The tax rates computed by McDaniel (2011) 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 two obvious virtues. First, 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. Second, 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. Much of the existing literature, including Prescott (2004) and Ohanian, Raffo, and Rogerson (2008), goes to the opposite extreme, assuming that *all* taxes collected are rebated lump sum to the household ($\theta_j = \tau_j$). In practice, this assumption is not directly imposed on the quantitative work, but 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 (2011) 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. We take comfort, however, from the observations in McDaniel (2011), 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, the inclusion of these alternate tax rates means any measured increases in generational lifetime tax rates are likely to be understated.

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 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 would like ideally to compare the model 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.[14] 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

[14] 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 16-17 year old workers in 2013-2014 depends on cohorts born after the year 1995, we assume that these cohorts face the same policies as that born in 1995.

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.[15] 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 $A_v = \left[max\{\frac{v-16.5}{65}, 0\}, min\{\frac{v-15.5}{65}, 1\}\right)$.[16] An agent of age v in calendar year t was born in year $j = t - v$, and he works

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

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.[17]

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+.[18]

[15] Integrals are computed numerically using standard procedures. Our results are robust to reducing the life span to allow agents to live until only 70 or 75, or extending it for ages prior to 16. This robustness stems from our use of balanced growth preferences and the calibration of the life cycle wage profile, discussed in the next section.

[16] 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).

[17] Data on detailed population weights are also subject to the same limitations faced for hours worked. However, as the evolution of the age distribution of the population is far more predictable than hours worked and not our object of interest, we feel comfortable obtaining weights for missing years by interpolation. Further details are available in Appendix A.

[18] While there are slightly more detailed data available for workers under age 24 and over age 55, given the model's stylized hours worked decision at the extensive margin, we focus on these broader bins. Note that if the reservation wage is realized only at ages near the end of the interval in an age bin, this can dramatically misrepresent the hours worked changes for that bin. However, the next section shows that the exact age cutoffs for hours worked are nicely centered in the 16-24 and 55+ age bins.

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:[19]

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

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%.[20] With these functional form choices for preferences, 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

[19]We consider balanced growth preference to be a natural starting point. That said, a motivation for using balanced growth preferences is the constant level of hours observed in the US over time, although this is generated by a simultaneous decline in hours worked by men and a simultaneous rise in hours worked by women. As the evidence in Appendix B shows, there is limited evidence for the idea that these trends are interacting at the household level, and thus it may be of interest to consider departures from this assumption on preferences.

[20]The generational accounts literature typically assumes an interest rate of 6%. Our results are robust to using this value for the interest rate and discount rate.

to match average labor supply behavior for men born from 1935 through 1945. We focus 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.[21]

We further assume that each cohort faces the same wage profile $w(a) = w_j(a)$.[22] 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. We have also considered using the limited data available on wage growth across cohorts to compute wage profiles that vary by cohort; incorporating these effects does not change our aggregate results.[23] An advantage of calibrating the wage profile to match hours worked behavior for a particular cohort is that this effectively accounts for the potential impact of constant progressivity in the tax structure on the life cycle profile of hours worked. And while this calibration strategy matches life cycle hours worked for the target cohorts, holding the wage profile fixed while allowing policy to vary allows us to assess the impact of policy on hours of different cohorts.

The exact calibration of the wage profile depends on the parameterization of the agent's disutility of labor. Following Rogerson and Wallenius (2009), we choose $\gamma = 1$, as this choice is consistent with both micro and macro estimates of the Frisch elasticity of labor supply. Figure 2 reports the calibrated wage profile, $w(a)$. With $\gamma = 1$, wages roughly double over the life cycle, which is generally consistent with wage profile reported in Rupert and Zanella (2015). Appendix C shows that our results are robust to different values of γ. In constructing wage profiles for different values of γ, we find that the implied life cycle wage profile provides some support for $\gamma = 1$. Substantially larger or smaller values of γ imply vastly counterfactual predictions for the slope of

[21] We normalize the wage to 1 at age 16 and assume it to be zero above age 75. Hours per worker are measured from the March CPS and Census; details are in Appendix A. An equivalent alternative would be to feed in measured wage profiles and then allow for variation in preferences to match the labor supply behavior over the life cycle.

[22] Given the life cycle wage profile, $w(a)$, it is straightforward to construct the CDF over wages, $F(w)$, needed to solve the model.

[23] There are two problems with allowing for changing wage profiles across time. The first is measurement. Because of substantial data limitations, any attempt to construct wage profiles for each cohort from 1900 through 1995 will require substantial interpolation. Second, interpolated wage profiles imply the fastest wage growth among the young and the old as is Kambourov and Manovskii (2009) and Kong, Ravikumar, and Vandenbroucke (2014), ages at which the relationship between wages and labor supply is still not well understood. Thus, while allowing for these changing wage profiles does not change our aggregate results, we do not report them in the paper.

the wage profile.

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 cohort are positive for 44 years. These calibration targets give values of $\alpha = 1.24$ and $\chi = 0.43$. Figure 2 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.[24]

5 Results

5.1 Aggregate Results and Results by Age Group

Figure 3 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.[25] Prior to the cyclical downturns in the 2000s, the model accounts for approximately 50% of the decline in the data. This decline is generated from changes in both the fiscal burdens faced by individual cohorts and the share of the population at each age. The timing of the decline in the model is consistent with the data; there is an initial decline in the early years of the sample followed by a flat period and another drop near the end. The primary difference is that the model does not account for the decline observed in the 1970s.

Figures 4 through 8 show 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. In contrast, for individuals ages 25 through 54, the model generates much less decline. The model accounts for roughly 30%

[24]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.

[25]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 observed in the data, compared to the 60 to 75% explained for 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 the extensive margin is important at these ages. Sections 5.2 and 5.3 disentangle each of these effects.

Table 2 summarizes these results, showing the fraction of the decline in observed hours worked that the model explains for the aggregate and each age group. In total, changes in age demographics and in tax and transfer policies explain a substantial proportion of the decline in male hours worked between 1948 and 2000.

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. This decomposition only represents changes in the composition of the population across the broad age bins V_k. We initially focus on decomposing demographic changes across 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 3 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 partially explains why the model is more successful in accounting for the decline in hours worked in the earlier period.[26]

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. Although annual data on hours worked at each individual age are not available for the entire period, we combine microdata from the 1950 Census and the March CPS beginning in 1962 to compute an estimate of the effect of demographic change at this level.[27] With this estimate, we can then determine how much of the decline in hours worked can be accounted for solely by changes in taxes and transfers.

Table 4 reports the results of this more detailed decomposition for the years 1948-2000. While demographic change across decadal age groups had no impact on hours worked, changes in the age composition of each decadal bin has been important, particularly for the young and the old. After accounting for demographic changes, changes in policy alone account for roughly 30% of the total decline in both aggregate hours worked and hours worked holding demographics fixed. Within age groups, the model also accounts for roughly 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. What we find is that the model's higher explanatory power for this time period is largely driven by a greater importance of demographic changes. Holding demographic changes fixed, the model still explains roughly 30% of the decline in hours worked.[28]

In summary, the model accounts for 49% of the decline in hours worked between 1948 and 2000. We find that about 28% come from changing policies and the remaining 21% comes from demographic changes.[29]

---

[26]Though unreported, the sizable downturn in hours worked in the model in the late 2000s is primarily generated by demographic change as well.

[27]Details of how this is computed are reported in the notes to Table 5.

[28]The only significant remaining variation across time periods is observable for older workers, where policy explains much more of the decline through the year 1970 than it does through 2000. However, as described previously, hours worked by men over age 55 are held constant through the year 1960, and over age 65 through 1970, meaning that the data and the model capture different behavior in the post-1970 period for hours worked by men over age 55.

[29]Notably, the model suggests a slightly larger role for demographic shifts than does the data over this time horizon,

5.3 Intensive versus Extensive Margins

To understand the relative importance of the intensive and extensive margins of labor supply, we decompose the 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 changes, using the same approach as used for demographic change, shown in Equation 2.[30] 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.[31] 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

with only 8% coming from demographic shifts in the data. This may be partially due to the difficulties in carefully identifying annual labor supply at the individual age level. 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.

[30]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.

[31]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.

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.[32] 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.[33] 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.

We raise two points of caution in interpreting these results. 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.[34] 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.

[32] 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).

[33] 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.

[34] 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.

In spite of these cautions, 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 previous studies frequently 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.

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 more than three times as much 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 10% on average compared to a tax rate of around 40%, 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. Prescott (2004) and Ohanian, Raffo, and Rogerson (2008) are the papers closest in spirit to ours, although they use quite different methods. The model of Ohanian, Raffo, and Rogerson (2008) predicts a 9.4% total decline in hours worked in the US between 1956 and 2003; for that period, our model generates a decline of 6.4%, of which 4.8% comes from policy changes alone.[35] As there are clear differences in how taxes and transfers are measured in the two studies, we discuss how measurement adjustments affect our conclusions.[36]

Section 3 explained that generational accounts tax rates include not only labor taxes, but also those on capital and excise taxes, which differs from Ohanian, Raffo, and Rogerson (2008), 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 ($\theta_j = \tau_j$).[37] 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 our baseline results. These changes do not reconcile the differences between our find-

[35] As neither our nor their tax wedges have explicit reference to gender, we consider this to be a reasonable comparison. That said, 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.

[36] 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 30%, then the most we would expect this adjustment to explain is one percentage point of the difference in the final results.

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

ings 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 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.

We conclude with a direct comparison between generational accounts and the usual average tax wedge. A natural comparison would be to use rates like those in Ohanian, Raffo, and Rogerson (2008) fitted to our model framework. However, the basic assumptions of our model and theirs are just too different to allow a meaningful comparison. Instead, we conduct a simple empirical exercise to assess the additional value of using the generational accounting tax measures. We run regressions of male hours per person by birth year and also calendar year on both our tax measures and the tax wedge from Ohanian, Raffo, and Rogerson (2008). 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.[38]

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 lifetime tax wedge faced by a man born in year $j = t-v$, and $\tau_{ORR}(t)$ is the wedge from Ohanian, Raffo, and Rogerson (2008) at time t.[39]

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

[38]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.

[39]We use a log-log specification, as this is most consistent with the first order conditions of these models. A linear specification leads to very similar conclusions.

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.[40]

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 inclusion of tax rates at other dates in the life cycle, has a meaningful impact on labor supply. In particular, if agents choose their labor supply in a life cycle framework, we would also expect the coefficient to be particularly large and negative for older and younger workers, where the average of tax rates over the life cycle is particularly important.

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 Ohanian, Raffo, and Rogerson (2008).[41] 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 consistent with our theory described above, the generational accounts tax wedge has a particularly large coefficient for older and younger workers.[42] Overall, the empirical results highlight that tax wedges measured from the generational accounts provide additional value for explaining changes in labor supply over time.

In sum, we have found that a model with shifts in demographics and changing fiscal policy

[40] Though unreported, our conclusions hold for specific regressions on almost all individual ages as well.

[41] The tax wedge for Ohanian, Raffo, and Rogerson (2008) 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. Removing consumption taxes from their measure does not impact the results. Again, our tax wedge for generational j is $1-\frac{1-\tau_j}{1-\tau_j+\theta_j}$. Observations without data from the generational accounts are omitted.

[42] Since our tax wedges are different from those in the existing literature on several dimensions, we have also run regressions to disentangle which components generate the results in Table 9. Results from these regressions are available upon request.

across birth cohorts can account for nearly 50% of the decline in hours worked by men between 1948 and 2000. Changes in generational policy alone account for nearly 30% of this decline. The model consistently accounts for 30% of the decline within each age group, including the large declines in hours worked by the old and the young. The model is able to account for these larger declines in labor supply among the old and the young because of the inclusion of the extensive margin of labor supply. Finally, using tax wedges from the generational accounts complements past studies and provides a deeper understanding of the effects of fiscal policy on labor supply.

References

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.

AUERBACH, A. J., L. J. KOTLIKOFF, AND D. R. KOEHLER (2016): "U.S. Inequality, Fiscal Progressivity, and Work Disincentives: An Intragenerational Accounting," Working Paper 22032, National Bureau of Economic Research.

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

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.

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

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.

KAMBOUROV, G. AND I. MANOVSKII (2009): "Accounting for the changing life-cycle profile of earnings," *University of Pennsylvania. Mimeo*.

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

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. (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.

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.

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

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.

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	-8.8	-18.6	-3.3	-3.5	-3.6	-25.9
Data	-17.8	-31.5	-10.4	-10.9	-11.9	-35.2
Frac. Explained	**0.49**	**0.59**	**0.32**	**0.32**	**0.30**	**0.74**

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

		Total	Demographic Changes	Constant Demographics
1948-2000	Model	-8.8	0.0	-8.8
	Data	-17.8	-0.1	-17.7
	Frac. Explained	**0.49**	**-**	**0.50**
1948-1970	Model	-8.8	-3.4	-5.5
	Data	-11.5	-3.0	-8.6
	Frac. Explained	**0.77**	**1.13**	**0.64**

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, 1948-2000

		Aggregate	16-24	25-34	35-44	45-54	55+
Model	Total	-8.8	-18.6	-3.3	-3.5	-3.6	-25.9
	Demographics	-3.9	-9.3	0.1	0.0	0.0	-18.0
	Labor Supply	-4.9	-9.3	-3.4	-3.5	-3.6	-7.9
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		**0.28**	**0.30**	**0.33**	**0.32**	**0.30**	**0.22**
of Labor Supply		**0.31**	**0.33**	**0.32**	**0.32**	**0.30**	**0.33**

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, 1948-1970

		Aggregate	16-24	25-34	35-44	45-54	55+
Model	Total	-8.8	-20.1	-1.0	-1.1	-1.7	-10.9
	Demographics	-6.9	-15.6	0.0	0.0	0.0	-7.9
	Labor Supply	-1.9	-4.4	-1.0	-1.1	-1.7	-3.0
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		**0.17**	**0.17**	**0.32**	**0.28**	**0.29**	**0.31**
of Labor Supply		**0.28**	**0.21**	**0.32**	**0.28**	**0.29**	**0.55**

See notes to Table 4.

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

		Aggregate	16-24	25-34	35-44	45-54	55+
Model	Total	-4.9	-9.3	-3.4	-3.5	-3.6	-7.9
	Extensive	-1.7	-6.5	0.0	0.0	0.0	-5.5
	Intensive	-3.2	-2.8	-3.4	-3.5	-3.6	-2.3
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		**0.31**	**0.33**	**0.32**	**0.32**	**0.30**	**0.33**
Extensive Margin		**0.16**	**0.54**	**0.00**	**0.00**	**0.00**	**0.20**
Intensive Margin		**0.67**	**0.17**	**0.47**	**0.66**	**0.73**	**-0.62**

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 (Holding Demographics Fixed) Coming From Taxes and Transfers Separately,1948-2000

		Aggregate	16-24	25-34	35-44	45-54	55+
Model	Total	-4.9	-9.3	-3.4	-3.5	-3.6	-7.9
	Tax Changes	-1.1	-1.1	-0.6	-0.8	-1.0	-2.4
	Transfer Changes	-4.0	-8.0	-2.7	-2.8	-3.1	-6.6
Data	Total	-15.7	-28.6	-10.7	-10.9	-12.0	-24.3
Frac. of Data Explained:							
Using Only Taxes		**0.07**	**0.04**	**0.06**	**0.07**	**0.08**	**0.10**
Using Only Transfers		**0.26**	**0.28**	**0.25**	**0.26**	**0.26**	**0.27**

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 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: Gross Lifetime Tax and Transfer Rates by Birth Cohort, 1900-1995

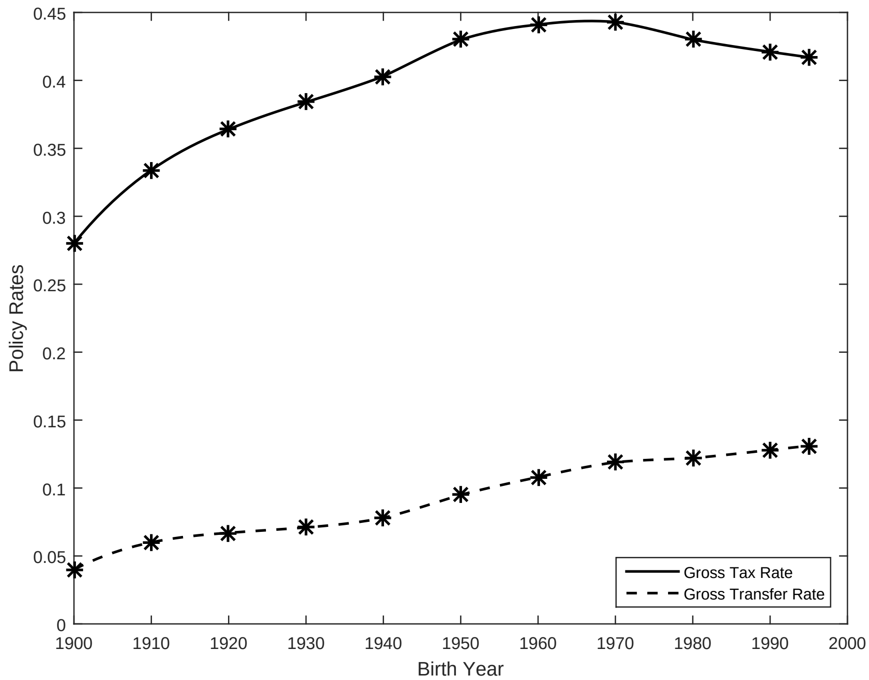

Gross tax and transfer rates are obtained by polynomial interpolation. Starred points represent the original tax and transfer rates in Gokhale, Page, and Sturrock (1999)

Figure 2: Life Cycle Wage Profile

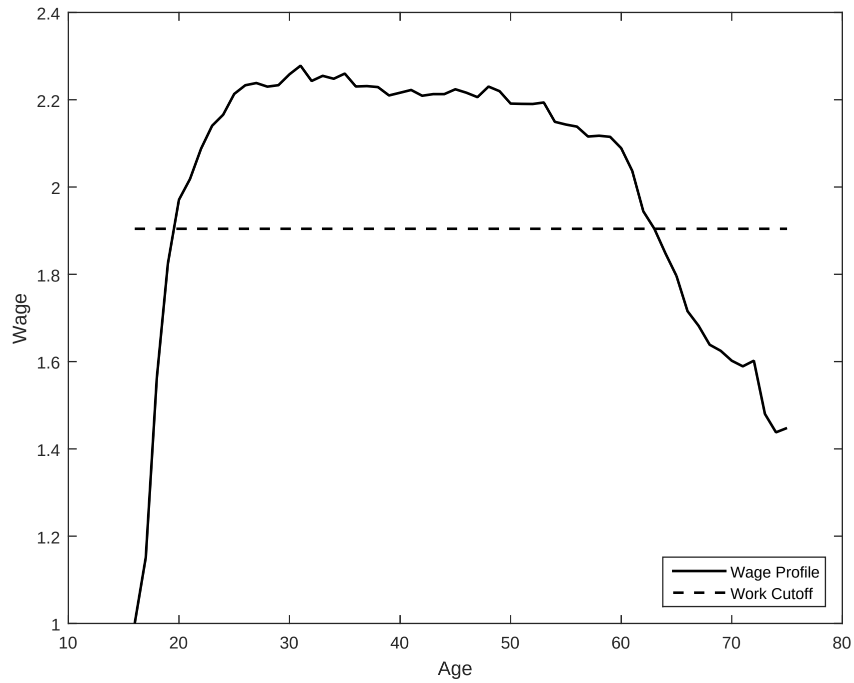

The solid line presents the life cycle wage profile calibrated to match the hours per worker data for the 1935-1945 cohorts. The wage at age 16 is normalized to 1. The dashed line is the reservation wage faced by the 1940 cohort.

Figure 3: Aggregate Hours Worked, Model vs. Data

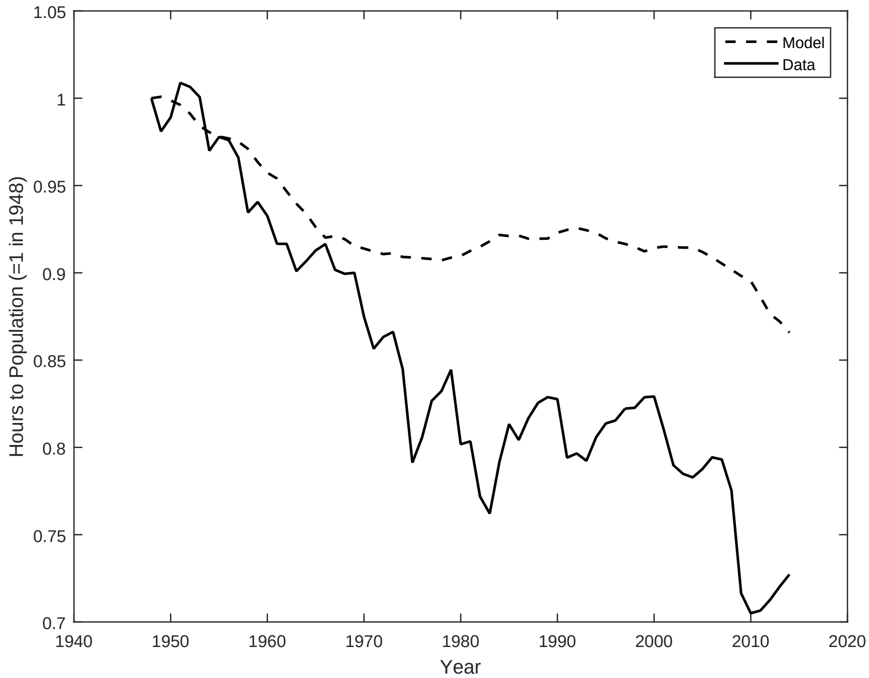

Data on taxes and transfers begin only in 1900. These series assume that there are no changes in hours worked for those aged 55 through 64 until 1960 and no such changes for men 65 and older until 1970.

Figure 4: Hours Worked, Age 16-24

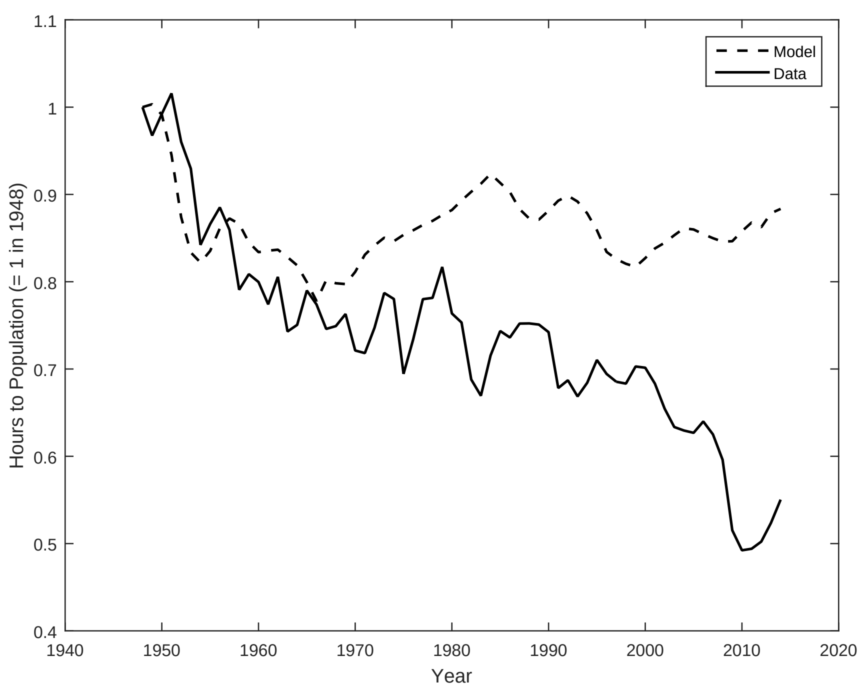

Figure 5: Hours Worked, Age 25-34

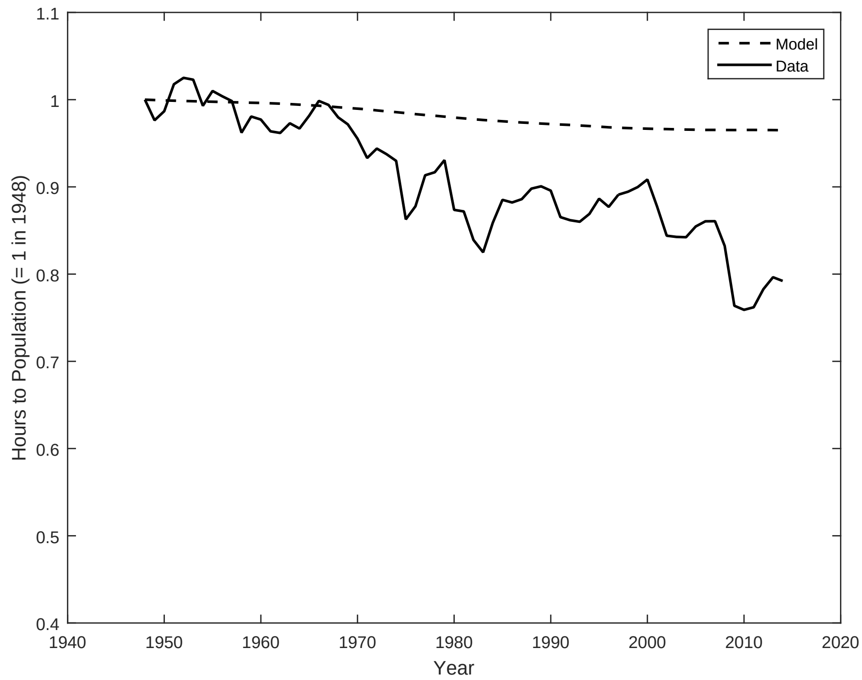

Figure 6: Hours Worked, Age 35-44

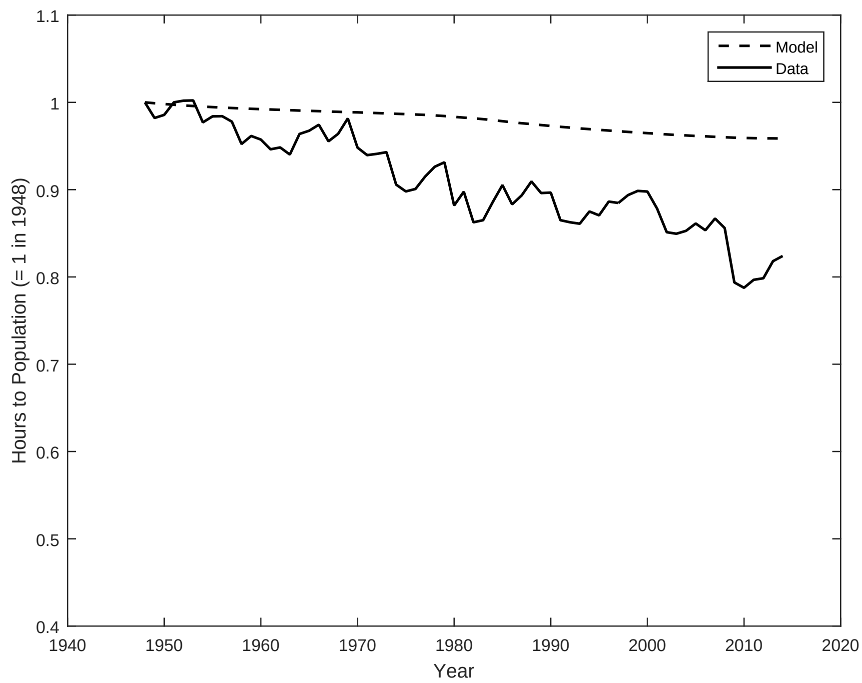

Figure 7: Hours Worked, Age 45-54

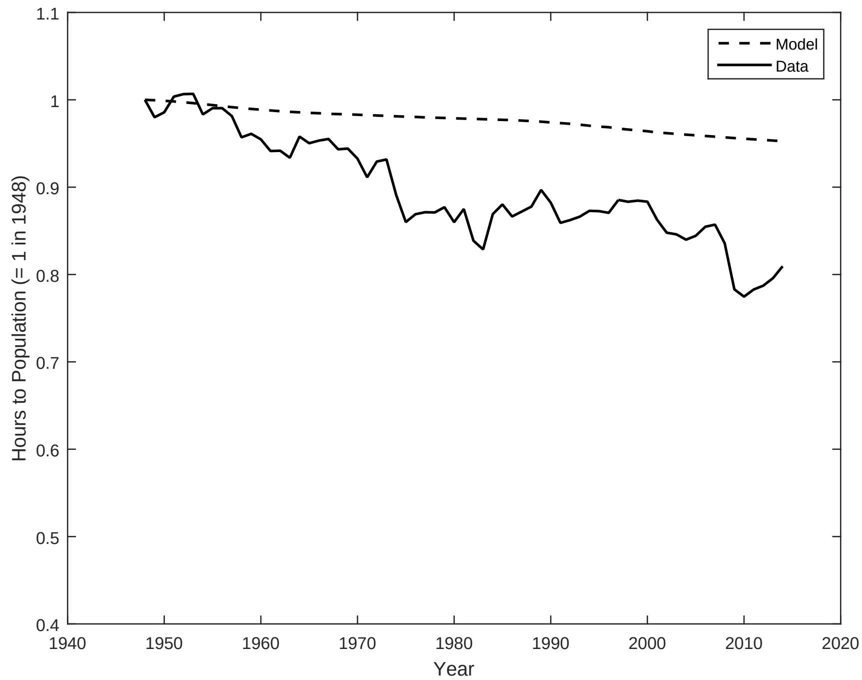

Figure 8: Hours Worked, Age 55+

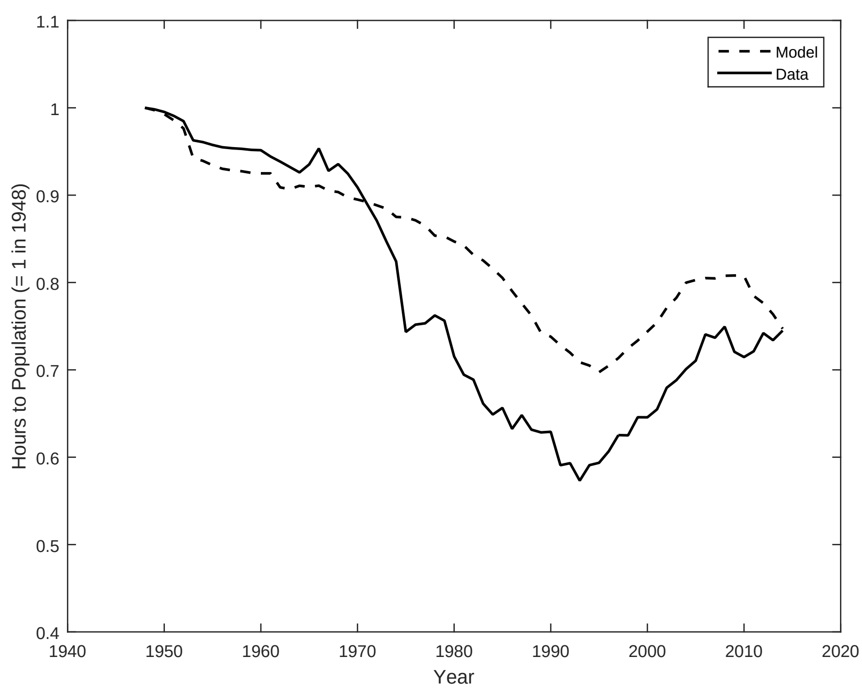

See notes to Figure 3.

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.[43]

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.

[43]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..

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.

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.[44] 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

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

men, as the labor supply of the latter would be less affected by these trends.[45] 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 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

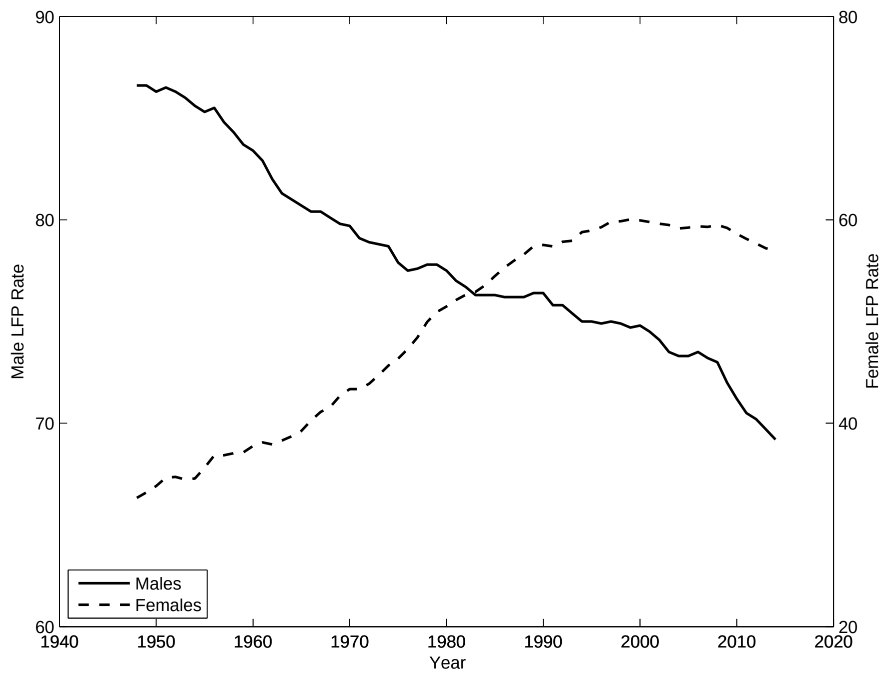

Data from BLS annual averages of monthly Employment and Earnings reports. Male labor force participation is the solid line (left axis); female labor force participation is the dashed line (right axis). Separate axes are used to emphasize the relative timing in labor force participation rates for men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

[45]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.

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

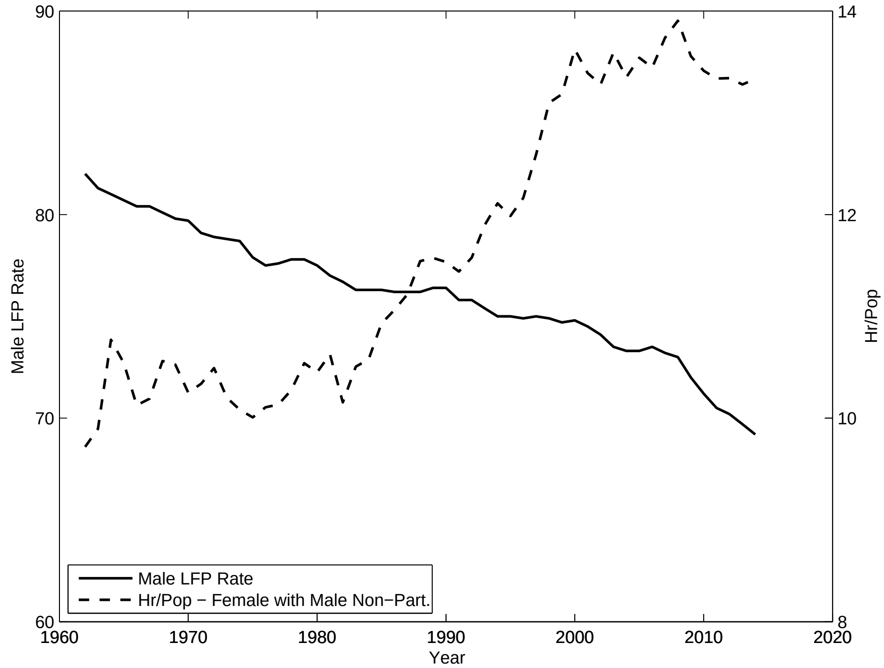

Data is from the March CPS.

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

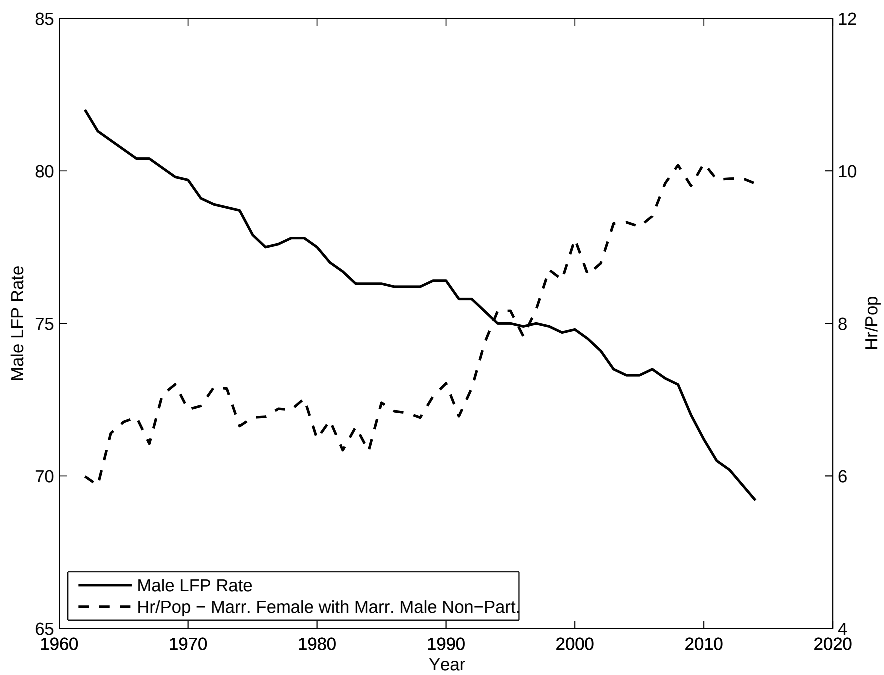

Data is from the March CPS.

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

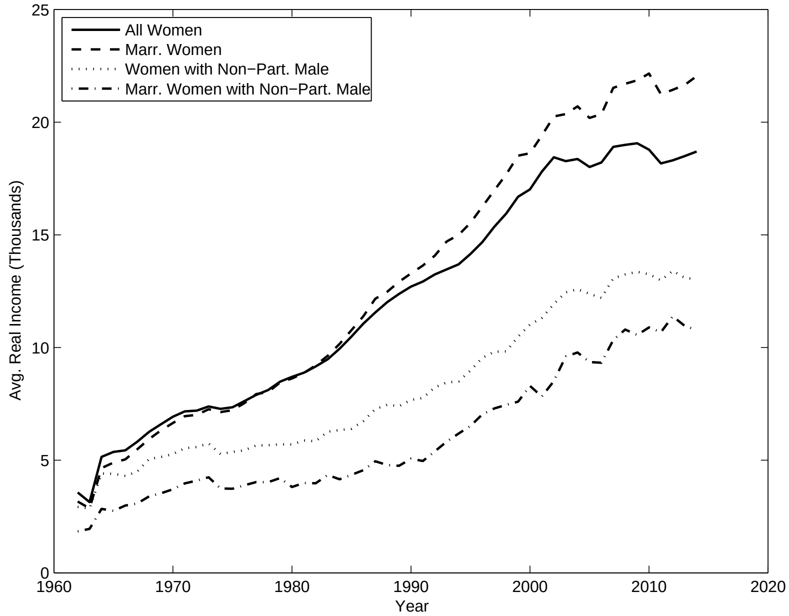

Data is from the March CPS. Real incomes are computed by dividing by the price index for personal consumption expenditures.

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

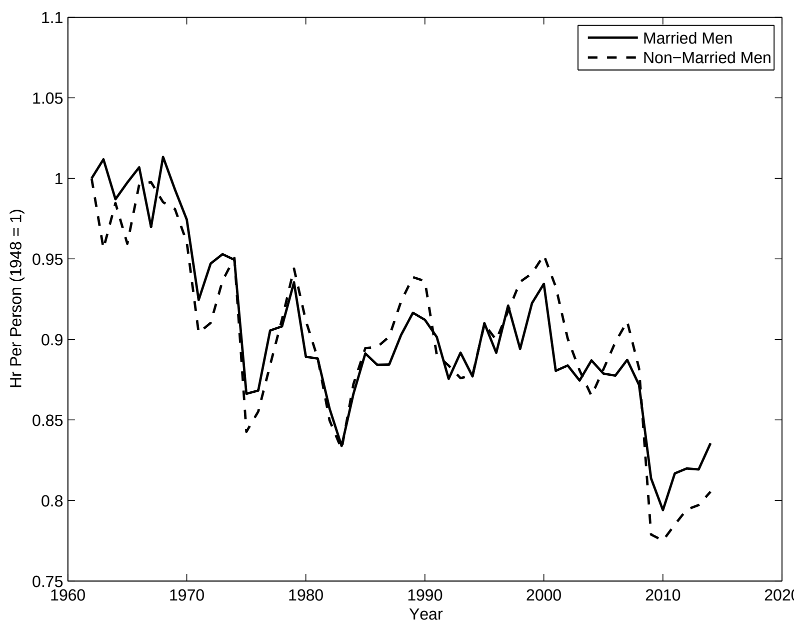

Data is from the March CPS. Hours per person (per week) for married men is 35.6 in 1962; for non-married men it is 26.8.

Appendix C: Labor Disutility Parameter Robustness

In this Appendix, we present several robustness checks on our primary results. We consider variation in the values for γ, which impacts intensive margin labor disutility, and χ, which impacts extensive margin labor disutility.

Our baseline calibration assumes that the intensive margin labor disutility parameter, γ, is equal to 1.[46] Table C.1 presents the model's results by age group for fixed demographic declines for different values of this parameter. The model's overall performance is largely unchanged for a reasonable range of values for γ, but declines by age group vary widely based on γ. With a lower values of γ, agents have a greater elasticity of labor supply when they do work. Hence, changes in taxes induce greater changes in hours worked for prime-aged males when γ is lower.

Our baseline calibration chose χ, the fixed cost of working, to generate a working life of 44 years for the 1940 cohort. Table C.2 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. The reason for this is clear when observing the wage profile shown in Figure 2. As χ varies, so do the cutoff years for entering and exiting the labor market. The model's results for the young and the old depend on the steepness of the wage profile around that cutoff date. If the wage profile is particularly steep (for example for ages 19-20, where the cutoff is in the baseline), then there will be greater changes in the extensive margin because the cost to adjusting the dates when the individual chooses to work relative to changing hours during other periods in life is lower.

[46] Note that unless the fixed cost, χ, equals zero, γ does not directly determine the Frisch elasticity in this model.

Table C.1: 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+
$\gamma = 0.25$	0.46	0.23	0.68	0.68	0.63	0.20
$\gamma = 0.5$	0.38	0.25	0.54	0.54	0.50	0.22
$\gamma = 1$	0.31	0.33	0.32	0.32	0.30	0.33
$\gamma = 2$	0.28	0.48	0.13	0.14	0.13	0.37
$\gamma = 4$	0.30	0.64	0.05	0.05	0.05	0.47

See notes to Table 1 and 2.

Table C.2: 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,$\chi = 0.53$	0.31	0.75	0.25	0.24	0.22	0.21
Work Life = 42y,$\chi = 0.47$	0.36	0.56	0.28	0.29	0.28	0.42
Work Life = 44y, $\chi = 0.43$	0.31	0.33	0.32	0.32	0.30	0.33
Work life = 46y, $\chi = 0.38$	0.29	0.27	0.35	0.35	0.33	0.20
Work life = 48y, $\chi = 0.33$	0.31	0.23	0.37	0.36	0.33	0.28

See notes to Table 1 and 2.