

**Opportunity &
Inclusive Growth
INSTITUTE**



FEDERAL RESERVE BANK OF MINNEAPOLIS



INSTITUTE WORKING PAPER
No. 41

Are Marriage-Related Taxes and Social Security Benefits Holding Back Female Labor Supply?

October 2020

Margherita Borella

*University of Torino and CeRP-Collegio
Carlo Alberto*

Mariacristina De Nardi

*University of Minnesota, Federal Reserve
Bank of Minneapolis, CEPR, and NBER*

Fang Yang

Louisiana State University

DOI: <https://doi.org/10.21034/iwp.41>

JEL classification: E21, H2, J22, J31

The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

Are marriage-related taxes and Social Security benefits holding back female labor supply?

Margherita Borella, Mariacristina De Nardi, and Fang Yang*

September 21, 2020

Abstract

In the United States, both taxes and old age Social Security benefits depend on one's marital status and tend to discourage the labor supply of the secondary earner. To what extent are these provisions holding back female labor supply? We estimate a rich life cycle model of labor supply and savings for couples and singles using the method of simulated moments (MSM) on the 1945 and 1955 birth-year cohorts and use it to evaluate what would happen without these provisions. Our model matches well the life cycle profiles of labor market participation, hours, and savings for married and single people and generates plausible elasticities of labor supply. Eliminating marriage-related provisions drastically increases the participation of married women over their entire life cycle, reduces the participation of married men after age 60, and increases the savings of couples in both cohorts, including the later one, which has similar participation to that of more recent generations. If the resulting government surplus were used to lower income taxation, there would be large welfare gains for the vast majority of the population.

*Margherita Borella: University of Torino and CeRP-Collegio Carlo Alberto, Italy. Mariacristina De Nardi: University of Minnesota, Federal Reserve Bank of Minneapolis, CEPR, and NBER. Fang Yang: Louisiana State University. De Nardi gratefully acknowledges support from the ERC, grant 614328 "Savings and Risks." Yang gratefully acknowledges MRRC grant number 08098401, pursuant to a grant from the U.S. Social Security Administration (SSA) funded as part of the Retirement Research Consortium through the University of Michigan Retirement Research Center Award RRC08098401. We thank Joe Altonji, Richard Blundell, Monica Costa Dias, Zvi Eckstein, Joan Gieseke, Rasmus Lenz, Derek Neal, and Jon Skinner for useful comments and suggestions. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research, MRRC, the SSA, the CEPR, any agency of the federal government, the Federal Reserve Bank of Minneapolis, or the Federal Reserve System.

1 Introduction

After increasing robustly from 1962 to the early 1990s, the labor force participation of women in the United States has been stagnating. Black, Schanzenbach, and Breitwieser (2017) write, “The U.S. economy will not operate at its full potential unless government and employers remove impediments to full participation by women in the labor market. The failure to address structural problems in labor markets—including tax and employment policy—does more than hold back women’s careers and aspirations for a better life. In fact, barriers to participation by women also act as brakes on the national economy, stifling the economy’s ability to fully apply the talents of 51 percent of the population.”

Barro and Redlick (2011) stress the importance of government policy and taxation in determining the aggregate performance of an economy and suggest an important role for the disincentives to female labor supply. More specifically, they study the effects of U.S. marginal tax rates over time and find that increases in average marginal tax rates have significant negative effects on aggregate output, with the notable exception of the 1948 tax cut, which was followed by the 1949 recession. Notably, this tax cut was in large part due to the introduction of joint filing for married couples and implied that the marginal tax rate for the secondary earner increased drastically.

In this paper, we ask, to what extent does the dependence of taxes and old-age Social Security benefits on marital status discourage female labor supply and affect welfare? The mechanisms are the following. First, since couples file taxes jointly, the secondary earner faces a higher marginal tax rate, which tends to discourage their labor supply. Second, married and widowed people can claim Social Security spousal and survivorship benefits under their spouses’ past contributions rather than their own. Hence, their reduced labor supply does not necessarily imply lower Social Security benefits. Since women have historically been the secondary earners, both provisions tend to discourage female labor supply, but to what extent are these disincentives holding it back? And to what extent do they affect cohorts with different female participation?

To answer this question, we develop and estimate a rich life cycle model with single and married people in which single people meet partners and married people might get divorced. Every working-age person experiences wage shocks and every retiree faces medical expenses and life span risk. People in couples face the risks of

both partners. Households can self-insure by saving and by choosing whether and how much to work (for both partners if in a couple). Consistent with the data, we allow for human capital to affect wages. We explicitly model Social Security with survival and spousal benefits, the differential tax treatment of married and single people, the progressivity of the tax system (including the earned income tax credit or EITC), and old-age means-tested transfer programs such as Medicaid and Supplemental Security Income (SSI). We also model the changes in the tax and Social Security systems that our two cohorts face over time.

We estimate our dynamic structural model using the method of simulated moments (MSM) and data from the Panel Study of Income Dynamics (PSID) and from the Health and Retirement Study (HRS) for the cohort born in the period 1941-1945 (referred to here as the “1945” cohort). That cohort has by now completed a large part of its life cycle and is covered by these two data sets, which provide excellent information over their working and retirement periods, respectively. Then, taking the estimated preference parameters from that cohort as given, we also estimate our model for the 1951-1955 cohort (referred to here as the “1955” cohort), which had much higher participation of married women (and closer to that of more recent cohorts) and for which policy and welfare implications might thus be very different.

Our estimated model matches the life cycle profiles of labor market participation, hours worked by the workers, and savings for married and single people for both cohorts very well. It also generates elasticities of labor supply by age, gender, and marital status that are consistent with those previously estimated by others. The latter provides an additional test of the reliability of our model and its policy implications.

For the 1945 cohort, we find that Social Security spousal and survivor benefits and the current structure of joint income taxation provide strong disincentives to work to married women and single women who expect to get married, and strong incentives to work for married men after age 60. For instance, the elimination of all of these marriage-based rules raises participation at age 25 by over 20 percentage points for married women and by 5 percentage points for single women. At age 45, participation for these groups is, respectively, still 15 and 3 percentage points higher without these marital benefits provisions. In addition, the elimination of marriage-based rules decreases the participation of married men starting at age 60, resulting in a participation rate that is 8 percentage points lower by age 65. Finally, for this

cohort, the elimination of marital provisions increases the savings of married couples by 20.3% at age 66.¹ In terms of welfare, abolishing these marital provisions would benefit most couples, all single men, and over one-third of single women and, thus, over 90% of the people in this cohort.

Given that the labor supply of married women has been increasing rapidly over time for cohorts born before the 1970s, a natural question that arises is whether the effects of these marital provisions are also large for more modern cohorts in which married women are more likely to work. To shed light on this question, we study a cohort that is 10 years younger than our reference cohort (that is, the 1955 cohort), for which we still have a completed labor market history and whose labor market behavior is close to that of more recent cohorts. By way of comparison, the labor market participation of married women at age 25 is just over 50% for our 1945 cohort, whereas it is over 60% for our 1955 cohort.

To estimate our model for the 1955 cohort, we assume that their preference parameters are the same as the ones we estimate for the 1945 cohort, but we give the 1955 cohort their observed marriage and divorce probabilities, number of children, initial conditions for wages and experience, and returns to working. We then estimate the child care costs, available time, and participation costs that reconcile their labor supply and saving behavior with the observed data. Finally, we run the policy experiment of eliminating the marriage-related provisions for both taxes and Social Security. We find that the effects for the 1955 cohort on participation, wages, earnings, and savings are large and similar to those in the 1945 cohort, thus indicating that the effects of marriage-related provisions are also large for cohorts in which the labor participation of married women is higher. We also find that abolishing these marriage-related provisions for this cohort at age 25 would also benefit most couples, all single men, and over two-thirds of single women. In addition, the welfare benefits to those gaining would be much higher, and the welfare costs of those losing would be very small, because the human capital of women in the cohort is already higher than that in the previous cohort at age 25.

Our paper provides several contributions. First, it is the first estimated structural model of couples and singles that allows for participation and hours decisions of both

¹While our model takes marriage and divorce behavior from the observed data from each cohort, we show in Section 8 that the empirical evidence finds small effects of these provisions on marriage and divorce and that our results are robust to large changes in marriage and divorce behavior.

men and women, including those in couples, in a framework with savings. Our results show that, in addition to lowering the participation of women, these marriage-related policies significantly reduce the savings of couples and increase the participation of married men later in life, and decrease welfare for the vast majority of the population. Second, it is the first paper to study all marriage-related taxes and benefits in a unified framework. Third, it does so by allowing for the large observed changes in the labor supply of married women over time by studying two different cohorts. Fourth, our framework is very rich along dimensions that are important in the study of our problem. For instance, allowing for labor market experience to affect wages (of both men and women) is important in that it captures the endogeneity of wages and their response to policy and marital status changes. Carefully modeling survival, health, and medical expenses in old age, and their heterogeneity by marital status and gender, is crucial to evaluate the effects on labor supply and savings of Social Security payments during old age and their interaction with taxation and old-age means-tested benefits such as Medicaid and SSI, which we also model. By modeling one-year periods, it gives people the ability to change their labor supply and savings in a more flexible and realistic way. Finally, our model fits the data for participation, hours worked, and savings, the estimated labor supply elasticities over the life cycle for single and married men and women, and thus provides a valid benchmark to evaluate the effects of the current marriage-related policies.

1.1 Related literature

We build on the literature on female labor supply over the life cycle. Within this literature, Attanasio, Low, and Sánchez-Marcos (2008) and Eckstein and Lifshitz (2011) point to the importance of changing wages and child care costs in explaining increases in female labor supply over time. Eckstein, Keane, and Lifshitz (2019) examine the changes over time in the selection and determinants of married women working. Hubener, Maurer, and Mitchell (2016) study the effects of exogenous family dynamics and endogenous labor supply on portfolio choice and retirement.

The structural papers in this branch of the literature typically assume that male labor supply is exogenously fixed and/or that the choice of hours of both partners is limited to full-time or full-time and part-time, and/or abstract from savings. We also add to this literature by quantifying the disincentive effects of the U.S. Social

Security and tax code on the labor supply of women.

We contribute to the small body of literature studying policy reforms in environments that include couples. Guner, Kaygusuz, and Ventura (2012a) study the switch to a proportional income tax and a reform in which married individuals can file taxes separately and find that these reforms substantially increase female labor participation. Nishiyama (2017), Kaygusuz (2015), and Groneck and Wallenius (2017) find that removing spousal and Social Security survivor benefits would increase female labor participation, female hours worked, and aggregate output. Bick and Fuchs-Schundeln (2018) focus on a simpler static model of married couples and find that income taxes are an important factor driving differences in the labor supply of married women across countries.

More generally, our paper differs from the previous literature in focus, methodology, and important model elements. In terms of focus, previous papers have only studied the effects of removing marriage-related rules that pertain to either Social Security or taxes and thus cannot answer the question as to what extent these provisions jointly hold back female labor supply, which is the focus of our paper. In terms of methodology, we not only estimate our model but also make sure that our model's inputs and outputs are consistent with the PSID and HRS data for the working and retirement periods, respectively. As a result, for instance, we estimate the accumulation of human capital on the job from the data and allow the tax structure to vary over time for each cohort (and estimate our tax functions from the PSID as a function of cohort, year, and marital status). Thus, we take this variation into account when we estimate our model. In terms of important model elements, none of the previous papers models health shocks and medical expenses in retirement, which are important to understand savings and the role of Social Security in insuring both mortality and medical expense risks, nor do they have flexible labor supply of both men and women, including in hours worked, over all of the working period. As we show, the labor supply of men also changes as a result of the reforms, thus allowing that of women to adjust differently than it would have if the labor supply and hours of men had been fixed.

2 Background on marriage and U.S. taxes and old-age Social Security benefits

Many countries tax the income of married people by making them file as if they were single (individual taxation). As a result, when the secondary earners in couples work, their marginal tax rate is based on their own income rather than on the sum of their partner's income and their own.

The United States, instead, taxes the income of married couples jointly (joint taxation) and uses a different tax schedule for married and single people. The combination of joint taxation and a progressive tax system typically implies that a married secondary earner faces a higher marginal tax rate than a single earner.

The question of when and why we ended up with such a system in the U.S. is an interesting one. Our reading of the literature is that joint taxation was implemented in 1948 with the goal of eliminating differences between community-property and common-law property states. In community property states, all income received by a married couple is considered jointly earned and owned. Thus, their residents felt legally entitled to pay taxes, including at the federal level, on the average income of each spouse. This was not possible for residents in common-law property states, where the couples with a main earner (most of them at the time) thus ended up facing a much higher average marginal tax rate. The 1948 reform was meant to eliminate this source of inequality and essentially imposed that all couples have to file jointly (under a different tax bracket system).

To illustrate the secondary earner's disincentive to work, we use the effective tax rates that we estimate from the PSID in 1988, a time period during which the earned income tax credit (EITC) program is already active and people in our 1945 cohort are still of working age (in 1988 the median woman in our 1945 cohort is 45 years old). The details of our tax computations are at the end of Appendix B.

The left panel of Figure 1 illustrates the incentives to work for these single and married women by plotting four marginal tax rates as a function of women's earnings: the marginal tax rates of single women and those of married women with husbands at three different percentiles of earnings. A single woman earning \$500 a year faces a marginal tax rate of -10%, while a married woman earning the same amount faces a marginal tax rate of 14%, 18%, and 21%, respectively, if she is married to a man in the 25th, 50th, and 75th income percentiles (which correspond to, respectively,




Figure 1: Left panel: 45-year-old women’s marginal tax rate when single (starred red line) or married to men at the 25th (dashed orange line), 50th (dotted orange line), and 75th (circled orange line) income percentiles, as a function of women’s earnings in 2016 dollars (minimum value \$500). Right panel: cumulative density function (cdf) of 45-year-old non-working wives’ marginal tax rates.

\$43,090, \$68,995, and \$113,288 in 2016 dollars). Our estimated negative tax rate at low income levels illustrates the impact of the EITC.

While this graph tells us that married women typically face a higher marginal tax rate than single women, it does not tell us the distribution of marginal tax rates for married women who are not working. Thus, the right panel of Figure 1 displays the distribution of marginal tax rates for 45-year-old men whose wives are not working; this marginal tax rate is also that of their wives, should they start working. Comparing the marginal tax rate of non-working wives with that of non-working single women reveals that single women who are starting to work face a -10% marginal tax rate, while 80% of married women face a marginal tax rate of 10% or higher, because of their husband’s earnings and joint taxation. These graphs thus suggest that making married people file as single rather than jointly could have large incentives for the labor market participation of married women.

Social Security for a single person is a function of one’s average lifetime earnings. Social Security for a married person is the higher between one’s own benefit entitlement and half of the spouse’s entitlement while the other spouse is alive (spousal benefit) and the higher between one’s own benefit entitlement and the deceased spouse’s after the spouse’s death (survival benefit).

We use data from the PSID for 66-year-old couples in our 1945 cohort and Social Security rules to generate Figure 2, which illustrates the magnitude of Social Security spousal benefits. The left panel of Figure 2 plots household Social Security benefits




Figure 2: Left panel: Average household Social Security benefits at age 66 by wife’s own Social Security benefit decile, with (circled blue line) or without (crossed red line) marital benefit, in 2016 dollars. Right panel: Average survivor benefit by wife’s own Social Security benefit decile, with (circled blue line) or without (crossed red line) marital benefits, 2016 dollars.

while the husband is alive. It takes married women at retirement age and, based on the deciles of their own Social Security entitlement, plots their average household yearly Social Security benefits with (circled line) and without (crossed line) marital benefits. For instance, the number 1 on the x -axis represents 66-year-old married women in our 1945 cohort that are in the lowest decile of their own Social Security contributions. At that decile, household Social Security benefits for those women and their husbands are \$32,000 under marital benefits and about \$22,000 without marital benefits. The comparison of the two lines in this picture reveals that about 50% of married households benefit from Social Security marital benefits while their husband is alive and that these benefits can be very large.

The right panel of Figure 2 takes the same married women and plots what their yearly Social Security benefits would be after their husband’s death with and without survivor’s benefits. For instance, once a widow, a 66-year old married woman at the lowest 10% of Social Security contributions would receive less than \$500 a month based on her own contributions only, whereas she would receive \$22,000 thanks to her husband’s contributions and survivorship benefits. The picture shows that because most women have lower potential wages than men’s, participate less, and work fewer hours, survivorship benefits are large for over 80% of married women in this cohort. This last set of graphs highlights that Social Security marital benefits are large and can also reduce married women’s incentives to work.

3 Life cycle patterns for single and married men and women in our cohorts

We pick the 1945 cohort because their entire adult life is first covered by the PSID, which starts in 1968 and has rich information for the working period, and then by the HRS, which starts covering people at age 50 in 1994 and has rich information for the retirement period, including on medical expenses and mortality. Thus, we have excellent data for this cohort over their entire life cycle. We pick our 1955 cohort to be as young as possible to maximize changes in their participation, conditional on having an almost complete working period for the same cohort.²

Figure 3 displays participation and average annual hours worked by workers. The top panels refer to the 1945 cohort.³ The top left panel shows that married men have the highest participation rate and only slowly decrease their participation starting from age 45, whereas single men decrease their participation much faster. The participation of single women starts about 10 percentage points lower than that of single men and gradually increases until age 50. Married women have the lowest participation rate. It starts around 50% at age 25, increases to 78% between ages 40 and 50, and gradually declines at a rate similar to that of the other three groups. The top right panel highlights that married men on average work more hours than everyone else. Women not only have a participation rate lower than men on average but also display lower average hours, even conditional on participation.

The middle panels display the analogous information for the 1955 cohort. Comparing the top and bottom panels shows a large increase in participation by married women across these two cohorts. Conditional on working, average annual hours have also increased for married women. Finally, annual hours worked by married men conditional on working are lower, which underscores the importance of modeling men's labor supply, in addition to that of women's.

Because the availability of asset data in the PSID is limited (available only every five years until 1999 and every other year afterward) and our 1955 cohort has not yet

²Appendix A provides details about our computations and also shows that the majority of men and women are married in both cohorts and that the fraction of married people goes down only slightly across these two cohorts. Appendix I validates our labor market outcomes from the PSID with those from the Current Population Survey (CPS) for both cohorts and shows that they are very similar.

³These profiles are obtained from the data by fitting a fourth-order polynomial in age fully interacted with marital status and cohort dummies, separately for each gender.




Figure 3: Life cycle profiles by gender and marital status for the 1945 cohort (top two graphs), 1955 cohort (middle two graphs), and both cohorts (bottom graph), PSID data

retired, we use the same asset profiles for both cohorts. The bottom panel in Figure 3 shows that average assets increase until age 70 for all groups, with single women accumulating the lowest amount and showing no sign of a slowdown in accumulation before age 75.

4 The model

Our model period is one year long and there are three stages in one's life: a working stage (ages 25 to 61), an early retirement stage (ages 62 to 65), and a retirement stage (age 66 to the maximum age of 99).

During the **working stage**, single and married people choose how much to work

and save and face wage shocks. Married people face divorce shocks, and single people might meet partners and get married.⁴

Wages are a function of one's human capital (which is endogenously accumulated while working) and are affected by shocks. We measure human capital at a point in time as a person's average accumulated earnings at that point in time. Thus, human capital is a function of one's past wages and labor supply (and of one's education, to the extent that education influences one's wages).

We model (and estimate) available time to be split between working and leisure and we allow it to depend on one's gender and marital status. We interpret it as net of home production, child care, and elderly care that one has to perform whether working or not (and that is not easy to outsource). All workers have to pay a fixed cost of working, which depends on their age, gender, and marital status. It represents the cost of commuting, getting ready for work, making arrangements for being able to go to work, and so on.

Single women and married people have children, and the number of their children depends on maternal age and marital status. We allow for both time costs and monetary costs of raising children. The time costs affect one's available time for working and enjoying leisure. The monetary costs enter our model in two ways. First, they affect consumption through an adult-equivalent scale family size. Second, working mothers have to pay child care costs that depend on the age and number of their children, and on their own earnings. We thus assume that child care costs are a normal good: women with higher earnings pay for more expensive child care.⁵

During the **early retirement stage**, people still experience wage shocks but single people don't get married anymore and couples no longer divorce.⁶ If they claim Social Security, they can no longer work. Couples claim Social Security at the same time.

During the first year of the **retirement stage**, those who have not already claimed

⁴For tractability, we assume that people survive to retirement for sure. Although the death of a spouse is a big shock for the households experiencing it (Fadlon and Nielsen, 2015), it is a low-probability event in the data.

⁵Introducing home production and child care choices is infeasible given the complexity of our framework. The main caveat with our assumptions is that we do not allow these choices to vary when policy changes.

⁶In the HRS data, we observe our 1941-1945 birth cohort between ages 62 and 72. Over that period, only 1% of couples get divorced and 4% of singles get married. Thus, the implied yearly probability of marriage and divorce is very small.

Social Security do so and stop working. People face out-of-pocket medical expenses and the risk of death. Thus, each married person faces the risk of his or her spouse dying, in addition to their own. Mortality risk and medical expenses depend on gender, age, health status, and marital status.

Given that we explicitly model labor participation and hours of husbands and wives, savings, and medical expenses in old age, our model is computationally very intensive (See Appendix D for more details.) For tractability, we make the following additional assumptions. First, people who are married to each other are the same age. Second, fertility is exogenous, and women have an age-varying number of children that depends on their age and marital status and that we estimate from the data. Lastly, we assume that marriage and divorce are exogenous processes that we also estimate from the data. Thus, our results should be interpreted as holding marriage and divorce patterns fixed at those historically observed for this cohort. We discuss the empirical literature on the responses of changes in marriage and divorce rates to policy changes and evaluate the robustness of our findings to this assumption in Section 8.

4.1 Preferences

Let t be age $\in \{t_0, t_1, \dots, \dots, t_d\}$, with $t_0 = 25$ and $t_d = 99$ being the maximum possible life span. For simplicity of notation, think of the model as being written for one cohort, so age t also indexes the passing of time for that cohort. We solve the model for the two cohorts separately and make sure that each cohort has the appropriate time- and age-varying inputs.

Households have time-separable preferences and discount the future at rate β . The superscript i denotes gender, with $i = 1, 2$ being a man or a woman, respectively. The superscript j denotes marital status, with $j = 1, 2$ being single or in a couple, respectively.

Each **single person** has preferences over consumption and leisure, and the period flow of utility is given by the standard CRRA utility function

$$v^i(c_t, l_t, \eta_t^{i,1}) = \frac{((c_t/\eta_t^{i,1})^\omega l_t^{1-\omega})^{1-\gamma} - 1}{1-\gamma},$$

where c_t is consumption and $\eta_t^{i,j}$ is the equivalent scale in consumption (which is a

function of family size, including children) and $\eta_t^{i,1}$ corresponds to that for singles.

The term $l_t^{i,j}$ is leisure, which is given by

$$l_t^{i,j} = L^{i,j} - n_t^i - \Phi_t^{i,j} I_{n_t^i}, \quad (1)$$

where $L^{i,j}$ is available time endowment, which can be different for single and married men and women and should be interpreted as available time net of home production. It is a convenient way to represent activities that require time and cannot easily be outsourced. Leisure equals available time endowment less n_t^i , hours worked on the labor market and the fixed time cost of working. That is, the term $I_{n_t^i}$ is an indicator function that equals 1 when hours worked are positive and zero otherwise, while the term $\Phi_t^{i,j}$ represents the fixed time cost of working.

The fixed cost of working should be interpreted as including commuting time, time spent getting ready for work, and so on. We allow it to depend on gender, marital status, and age because working at different ages might imply different time costs for married and single men and women. We assume the following functional form, whose three parameters we estimate using our structural model:

$$\Phi_t^{i,j} = \frac{\exp(\phi_0^{i,j} + \phi_1^{i,j}t + \phi_2^{i,j}t^2)}{1 + \exp(\phi_0^{i,j} + \phi_1^{i,j}t + \phi_2^{i,j}t^2)}.$$

We assume that **couples** maximize their joint utility function

$$w(c_t, l_t^1, l_t^2, \eta_t^{i,j}) = \frac{((c_t/\eta_t^{i,j})^\omega (l_t^1)^{1-\omega})^{1-\gamma} - 1}{1-\gamma} + \frac{((c_t/\eta_t^{i,j})^\omega (l_t^2)^{1-\omega})^{1-\gamma} - 1}{1-\gamma}.$$

Note that for couples, $\eta_t^{i,j}$ does not depend on gender and that $j = 2$.

4.2 Environment

People can hold assets a_t at a rate of return r . The timing is as follows.

At the beginning of each working period, each single person observes his/her current idiosyncratic wage shock, age, assets, and accumulated earnings. Each married person also observes their partner's wage shock and accumulated earnings. At the beginning of each early retirement period, single people may claim Social Security benefits and couples may claim retirement benefits jointly.

At the beginning of each retirement period, each single person observes his/her current age, assets, health, and accumulated earnings. Each married person also observes their partner's health and accumulated earnings.

Decisions are made after everything has been observed and new shocks hit at the end of the period after decisions have been made.

4.2.1 Human capital and wages

We define human capital, \bar{y}_t^i , as one's average past earnings at each age. Thus, our definition of human capital implies that it is a function of one's initial wages and schooling and subsequent labor market experience and wages.⁷

Wages have two components. The first is a deterministic function of age, gender, and human capital: $e_t^i(\bar{y}_t^i)$. The second component is a persistent earnings shock ϵ_t^i that evolves as follows:

$$\ln \epsilon_{t+1}^i = \rho_\epsilon^i \ln \epsilon_t^i + v_t^i, \quad v_t^i \sim N(0, (\sigma_v^i)^2).$$

The product of $e_t^i(\cdot)$ and ϵ_t^i determines an agent's units of effective wage per hour worked during a period.

4.2.2 Marriage and divorce

During the working period, a single person gets married with an exogenous probability that depends on his/her age, gender, and wage shock. The probability of getting married at the beginning of next period is

$$\nu_{t+1}(\cdot) = \nu_{t+1}(i, \epsilon_t^i).$$

Conditional on meeting a partner, the probability of meeting with a partner p with wage shock ϵ_{t+1}^p is

$$\xi_{t+1}(\cdot) = \xi_{t+1}(\epsilon_{t+1}^p | \epsilon_{t+1}^i, i). \quad (2)$$

Allowing this probability to depend on the wage shock of both partners generates assortative mating. We assume random matching over assets a_{t+1} and average accumulated earnings of the partner \bar{y}_{t+1}^p , conditional on the partner's wage shock. Thus,

⁷It also has the important benefit of allowing us to have only one state variable keeping track of human capital and Social Security contributions.

we have

$$\theta_{t+1}(\cdot) = \theta_{t+1}(a_{t+1}^p, \bar{y}_{t+1}^p | \epsilon_{t+1}^p). \quad (3)$$

A working-age couple can be hit by a divorce shock at the end of the period that depends on age and the wage shock of both partners,

$$\zeta_{t+1}(\cdot) = \zeta_{t+1}(\epsilon_t^1, \epsilon_t^2).$$

If the couple divorces, they split the assets equally, and each of the ex-spouses becomes single and moves on with half of the assets, their own wage shock, and own Social Security contributions. Since we do not distinguish the previously divorced from the singles, these two groups have the same number of children. We also abstract from alimony in the case of divorce.

4.2.3 Costs of raising children and running a household

We keep track of the total number of children and children's age as a function of mothers' age and marital status. The total number of children by one's age affects the economies of scale of single women and couples.

The number of children between ages 0 to 5 and 6 to 11 determines the child care costs of working mothers ($i = 2$). The term $\tau_c^{0,5}$ is the child care cost for each child ages 0 to 5, where that number of children is $f^{0,5}(i, j, t)$, while $\tau_c^{6,11}$ is the child care cost for each child ages 6 to 11, which are $f^{6,11}(i, j, t)$. We use our structural model to estimate these costs.

4.2.4 Medical expenses and death

At age 66, we endow people with a distribution of health that depends on their marital status and gender. After that, they face survival, medical expenses, and health shocks. Survival $s_t^{i,j}$ depends on one's age, gender, and marital status. Health status ψ_t^i can be either good or bad and evolves according to a Markov process $\pi_t^{i,j}(\psi_t^i)$ that depends on age, gender, and marital status. Medical expenses $m_t^{i,j}(\psi_t^i)$ are a function of age, gender, marital status, and health status.

4.2.5 Initial conditions

We take the fraction of single and married people at age 25 and their distribution over the relevant state variables from the PSID (that is, assets, human capital, and wage shocks, with the latter two being for each of the spouses in the case of a married couple) for each of our two cohorts. We define notation for all of our state variables in Section 4.4.

4.3 Government

Each cohort in our model faces the effective time-varying tax rates that it experienced in the data. As in Benabou (2002) and Heathcote et al. (2014), we adopt a functional form that allows for negative tax rates (and thus incorporates the EITC), and we allow our effective tax rates to depend on marital status, gender, and age for each cohort (and thus time). Taxes paid are thus given by

$$T(Y, i, j, t) = (1 - \lambda_t^{i,j} Y^{-\tau_t^{i,j}}) Y. \quad (4)$$

We estimate these functions using the PSID.

The government also uses a proportional payroll tax τ_t^{SS} on labor income, up to a Social Security cap \tilde{y}_t , to help finance old-age Social Security benefits. We also allow the payroll tax and the Social Security cap to change over time for each cohort as in the data. We thus assume that the tax changes were anticipated by the households.

We use human capital \bar{y}_t^i (computed as an individual's average earnings at age t , up to the cap \tilde{y}_t) to determine both wages and old-age Social Security payments.

The insurance provided by Medicaid and SSI in old age is represented by a means-tested consumption floor, $\underline{c}(j)$, as in Hubbard, Skinner, and Zeldes (1995).⁸

4.4 Recursive formulation

We define and compute nine sets of value functions: the value function of working-age singles, the value function of singles during the early retirement stage, the value function of retired singles, the value function of working-age couples, the value function of couples during the early retirement stage, the value function of retired couples,

⁸Borella, De Nardi, and French (2018) discuss Medicaid rules and observed outcomes after retirement.

the value function of an individual who is of working age and in a couple, the value function of an individual who is at early retirement stage and in a couple, and the value function of an individual who is retired and in a couple.

4.4.1 The value function of working-age singles

The state variables for a single person during one's working period are age t , gender i , assets a_t^i , the persistent earnings shock ϵ_t^i , and human capital \bar{y}_t^i . The corresponding value function is

$$W^s(t, i, a_t^i, \epsilon_t^i, \bar{y}_t^i) = \max_{c_t, a_{t+1}, n_t^i} \left(v^i(c_t, l_t^{i,j}) + \beta(1 - \nu_{t+1}(i))E_t W^s(t+1, i, a_{t+1}^i, \epsilon_{t+1}^i, \bar{y}_{t+1}^i) + \beta \nu_{t+1}(i)E_t \left[\hat{W}^c(t+1, i, a_{t+1}^i + a_{t+1}^p, \epsilon_{t+1}^i, \epsilon_{t+1}^p, \bar{y}_{t+1}^i, \bar{y}_{t+1}^p) \right] \right), \quad (5)$$

subject to equation (1) and

$$Y_t^i = e_t^{i,j}(\bar{y}_t^i)\epsilon_t^i n_t^i, \quad (6)$$

$$\tau_c(i, j, t) = \tau_c^{0,5} f^{0,5}(i, j, t) + \tau_c^{6,11} f^{6,11}(i, j, t), \quad (7)$$

$$T(\cdot) = T(ra_t + Y_t, i, j, t), \quad (8)$$

$$c_t + a_{t+1} = (1+r)a_t^i + Y_t^i(1 - \tau_c(i, j, t)) - \tau_t^{SS} \min(Y_t^i, \tilde{y}_t) - T(\cdot), \quad (9)$$

$$\bar{y}_{t+1}^i = (\bar{y}_t^i(t - t_0) + \min(Y_t^i, \tilde{y}_t)) / (t + 1 - t_0), \quad (10)$$

$$a_{t+1} \geq 0, \quad (11)$$

$$n_t^i \geq 0. \quad (12)$$

The expectation of the value function next period if one remains single integrates over one's wage shock next period. When one gets married, we not only take a similar expectation but also integrate over the distribution of the state variables of one's partner: $(\xi_{t+1}(\epsilon_{t+1}^p | \epsilon_{t+1}^i, i))$ is the distribution of the partner's wage shock defined in equation (2), and $\theta_{t+1}(\cdot)$ is the distribution of the partner's assets and human capital defined in equation (3).

The value function \hat{W}^c is the discounted present value of the utility for the same individual, once he or she is in a married relationship with someone with given state

variables, not the value function of the married couple, which counts the utility of both individuals in the relationship. We discuss the computation of the value function of an individual in a marriage later in this section.

Equation (10) describes the evolution of human capital, which we measure as average accumulated earnings (up to the Social Security earnings cap \tilde{y}_t) and which we use as a determinant of future wages and Social Security payments after retirement.

4.4.2 The value function of singles during the early retirement stage

Let tr denote the age at which someone first claimed Social Security. The recursive problem for an individual that has claimed Social Security at age tr can be written as

$$S^s(t, i, a_t^i, \bar{y}_r^i, tr) = \max_{c_t, a_{t+1}} \left(v^i(c_t, L^{i,j}) + \beta E_t S^s(t+1, i, a_{t+1}^i, \bar{y}_r^i, tr) \right), \quad (13)$$

subject to equations (8), (11), and

$$Y_t = SS(\bar{y}_r^i, tr) \quad (14)$$

$$c_t + a_{t+1} = (1+r)a_t + Y_t - T(\cdot). \quad (15)$$

The term $SS(\bar{y}_r^i, tr)$ is a function of the income that the single person earned during his or her working life, \bar{y}_r^i , and claiming age tr .

Let $N^s(t, i, a_t^i, \epsilon_t^i, \bar{y}_t^i)$ denote the value function of a person during the early retirement period who has not yet claimed benefits,

$$N^s(t, i, a_t^i, \epsilon_t^i, \bar{y}_t^i) = \max_{c_t, a_{t+1}, n_t^i} \left(v^i(c_t, l_t^{i,j}) + \beta E_t V^s(t+1, i, a_{t+1}^i, \epsilon_{t+1}^i, \bar{y}_{t+1}^i) \right), \quad (16)$$

subject to equations (1), (6), (8), (10), (11), (12), and

$$c_t + a_{t+1} = (1+r)a_t^i + Y_t^i - \tau_t^{SS} \min(Y_t, \tilde{y}_t) - T(\cdot). \quad (17)$$

Let $V^s(t, i, a_t^i, \epsilon_t^i, \bar{y}_t^i)$ denote the value function for a person during the early retirement stage who has not yet retired. At the beginning of each period, that person chooses whether to claim Social Security benefits, and D_t^i is an indicator function for that decision which maximizes

$$V^s(t, i, a_t^i, \epsilon_t^i, \bar{y}_t^i) = \max_{D_t^i} \left((1 - D_t^i) N^s(t, i, a_t^i, \epsilon_t^i, \bar{y}_t^i) + D_t^i S^s(t, i, a_t^i, \bar{y}_t^i, t) \right). \quad (18)$$

4.4.3 The value function of retired singles

The state variables for a retired single are age t , gender i , assets a_t^i , health ψ_t^i , average realized lifetime earnings \bar{y}_r^i , and Social Security claiming age tr . His or her recursive problem can be written as

$$R^s(t, i, a_t, \psi_t^i, \bar{y}_r^i, tr) = \max_{c_t, a_{t+1}} \left(v^i(c_t, L^{i,j}) + \beta s_t^{i,j}(\psi_t^i) E_t R^s(t+1, i, a_{t+1}, \psi_{t+1}^i, \bar{y}_r^i, tr) \right), \quad (19)$$

subject to equations (8), (11), (14), and

$$B(a_t, Y_t, \psi_t^i, \underline{c}(j)) = \max \left\{ 0, \underline{c}(j) - [(1+r)a_t + Y_t - m_t^{i,j}(\psi_t^i) - T(\cdot)] \right\} \quad (20)$$

$$c_t + a_{t+1} = (1+r)a_t + Y_t + B(a_t, Y_t, \psi_t^i, \underline{c}(j)) - m_t^{i,j}(\psi_t^i) - T(\cdot) \quad (21)$$

$$a_{t+1} = 0, \quad \text{if } B(\cdot) > 0. \quad (22)$$

The term $s_t^{i,j}(\psi_t^i)$ is the survival probability as a function of age, gender, marital status, and health status. The function $B(a_t, Y_t^i, \psi_t^i, \underline{c}(j))$ represents old-age means-tested government transfers (such as Medicaid and SSI) that ensure a minimum consumption floor $\underline{c}(j)$.

4.4.4 The value function of couples during the working period

The state variables for a married couple during the working stage are $(t, a_t, \epsilon_t^1, \epsilon_t^2, \bar{y}_t^1, \bar{y}_t^2)$, where 1 and 2 refer to gender, and the recursive problem for the married couple ($j = 2$) can be written as

$$\begin{aligned} W^c(t, a_t, \epsilon_t^1, \epsilon_t^2, \bar{y}_t^1, \bar{y}_t^2) = & \max_{c_t, a_{t+1}, n_t^1, n_t^2} \left(w(c_t, l_t^{1,j}, l_t^{2,j}) \right. \\ & + (1 - \zeta_{t+1}(\cdot)) \beta E_t W^c(t+1, a_{t+1}, \epsilon_{t+1}^1, \epsilon_{t+1}^2, \bar{y}_{t+1}^1, \bar{y}_{t+1}^2) \\ & \left. + \zeta_{t+1}(\cdot) \beta \sum_{i=1}^2 \left(E_t W^s(t+1, i, a_{t+1}/2, \epsilon_{t+1}^i, \bar{y}_{t+1}^i) \right) \right), \quad (23) \end{aligned}$$

subject to equations (1), (6), (7), (10), and

$$T(\cdot) = T(ra_t + Y_t^1 + Y_t^2, i, j, t) \quad (24)$$

$$c_t + a_{t+1} = (1+r)a_t + Y_t^1 + Y_t^2(1 - \tau_c(2, 2, t)) - \tau_t^{SS}(\min(Y_t^1, \tilde{y}_t) + \min(Y_t^2, \tilde{y}_t)) - T(\cdot) \quad (25)$$

$$a_t \geq 0, \quad n_t^1, n_t^2 \geq 0. \quad (26)$$

The expected value of the couple's value function is taken with respect to the conditional probabilities of ϵ_{t+1} given the current value of the ϵ_t for each of the spouses (we assume independent draws). The term $\zeta_{t+1}(\cdot)$ represents the probability of divorce that we defined in Section 4.2.2. The expected values for the newly divorced people are taken using the appropriate conditional distribution for their own wage shocks.

4.4.5 The value function of couples during the early retirement period

For tractability, we assume that during the early retirement stage, couples can no longer divorce. The recursive problem for couples that have claimed Social Security at age tr can be written as

$$S^c(t, a_t, \bar{y}_r^1, \bar{y}_r^2, tr) = \max_{c_t, a_{t+1}} \left(w(c_t, L^{1,j}, L^{2,j}) + \beta E_t S^c(t+1, a_{t+1}, \bar{y}_r^1, \bar{y}_r^2, tr) \right), \quad (27)$$

subject to equations (8), (15), (11), and

$$Y_t = \max \left\{ (SS(\bar{y}_r^1, tr) + SS(\bar{y}_r^2, tr)), \frac{3}{2} \max(SS(\bar{y}_r^1, tr), SS(\bar{y}_r^2, tr)) \right\} \quad (28)$$

In equation (28), the variable Y_t represents the spousal benefit from Social Security, which gives a married person the right to collect the higher amount between one's own benefit and half of their spouse's benefit.

Let $N^c(t, a_t, \epsilon_t^1, \epsilon_t^2, \bar{y}_t^1, \bar{y}_t^2)$ denote the value function of a couple that has not yet claimed benefits,

$$N^c(t, a_t, \epsilon_t^1, \epsilon_t^2, \bar{y}_t^1, \bar{y}_t^2) = \max_{c_t, a_{t+1}, n_t^1, n_t^2} \left(w(c_t, l_t^{1,j}, l_t^{2,j}) + \beta E_t V^c(t+1, a_{t+1}, \epsilon_{t+1}^1, \epsilon_{t+1}^2, \bar{y}_{t+1}^1, \bar{y}_{t+1}^2) \right), \quad (29)$$

subject to equations (1), (6), (10), (24), (26), and

$$c_t + a_{t+1} = (1 + r)a_t + Y_t^1 + Y_t^2 - \tau_t^{SS}(\min(Y_t^1, \tilde{y}_t) + \min(Y_t^2, \tilde{y}_t)) - T(\cdot). \quad (30)$$

Let $V^c(t, a_t, \epsilon_t^1, \epsilon_t^2, \bar{y}_t^1, \bar{y}_t^2)$ denote the value function for a married couple during the early retirement stage that has not yet claimed Social Security benefits. At the beginning of each period, a couple chooses whether to claim Social Security benefits, that is $D_t = 1$. The early claiming decision maximizes

$$V^c(t, a_t, \epsilon_t^1, \epsilon_t^2, \bar{y}_t^1, \bar{y}_t^2) = \max_{D_t} \left((1 - D_t)N^c(t, a_t, \epsilon_t^1, \epsilon_t^2, \bar{y}_t^1, \bar{y}_t^2) + D_t S^c(t, a_t, \bar{y}_t^1, \bar{y}_t^2, t) \right). \quad (31)$$

4.4.6 The value function of couples during retirement

During retirement, each of the spouses faces health shocks ψ_t^i and survival shocks $s_t^{i,2}(\psi_t^i)$. We assume that the health shocks of each spouse are independent of each other and that the death shocks of each spouse are as well. During each period, the married couple's recursive problem ($j = 2$) can be written as

$$\begin{aligned} R^c(t, a_t, \psi_t^1, \psi_t^2, \bar{y}_r^1, \bar{y}_r^2, tr) = \max_{c_t, a_{t+1}} & \left(w(c_t, L^{1,j}, L^{2,j}) + \right. \\ & \beta s_t^{1,j}(\psi_t^1) s_t^{2,j}(\psi_t^2) E_t R^c(t + 1, a_{t+1}, \psi_{t+1}^1, \psi_{t+1}^2, \bar{y}_r^1, \bar{y}_r^2, tr) + \\ & \beta s_t^{1,j}(\psi_t^1) (1 - s_t^{2,j}(\psi_t^2)) E_t R^s(t + 1, 1, a_{t+1}, \psi_{t+1}^1, \bar{y}_r^1, tr) + \\ & \left. \beta s_t^{2,j}(\psi_t^2) (1 - s_t^{1,j}(\psi_t^1)) E_t R^s(t + 1, 2, a_{t+1}, \psi_{t+1}^2, \bar{y}_r^2, tr) \right), \end{aligned} \quad (32)$$

subject to equations (8), (11), (22), (28), and

$$\bar{y}_r^i = \max(\bar{y}_r^1, \bar{y}_r^2), \quad (33)$$

$$B(a_t, Y_t, \psi_t^1, \psi_t^2, \underline{c}(j)) = \max \left\{ 0, \underline{c}(j) - [(1 + r)a_t + Y_t - m_t^{1,j}(\psi_t^1) - m_t^{2,j}(\psi_t^2) - T(\cdot)] \right\} \quad (34)$$

$$c_t + a_{t+1} = (1 + r)a_t + Y_t + B(a_t, Y_t, \psi_t^1, \psi_t^2, \underline{c}(j)) - m_t^{1,j}(\psi_t^1) - m_t^{2,j}(\psi_t^2) - T(\cdot) \quad (35)$$

In equation (33), the variables $\bar{y}_r^i, i = 1, 2$ represent that the survivor collects

benefits based on the higher amount between their own contributions and those of their deceased spouse.

4.4.7 The value functions of individuals in couples during working age and retirement

We have to compute the joint value function of the couple to appropriately compute joint labor supply and savings under the married couples' available resources. However, when computing the value of getting married for a single person, the relevant object for that person is his or her discounted present value of utility in the marriage. We thus compute this object for person of gender i who is married with a specific partner.

Let $\hat{c}_t(\cdot)$, $\hat{l}_t^{i,j}(\cdot)$, $\hat{a}_{t+1}(\cdot)$, and $\hat{D}_t(\cdot)$ denote, respectively, the optimal consumption, leisure, saving, and claiming decision for an individual of gender i in a couple with a given set of state variables. During the working period, we have

$$\begin{aligned} \hat{W}^c(t, i, a_t, \epsilon_t^1, \epsilon_t^2, \bar{y}_t^1, \bar{y}_t^2) &= v^i(\hat{c}_t(\cdot), \hat{l}_t^{i,j}) + \\ &\beta(1 - \zeta(\cdot))E_t\hat{W}^c(t+1, i, \hat{a}_{t+1}(\cdot), \epsilon_{t+1}^1, \epsilon_{t+1}^2, \bar{y}_{t+1}^1, \bar{y}_{t+1}^2) + \\ &\beta\zeta(\cdot)E_tW^s(t+1, i, \hat{a}_{t+1}(\cdot)/2, \epsilon_{t+1}^i, \bar{y}_{t+1}^i). \end{aligned} \quad (36)$$

During the early retirement period, we have

$$\begin{aligned} \hat{N}^c(t, i, a_t, \epsilon_t^1, \epsilon_t^2, \bar{y}_t^1, \bar{y}_t^2) &= v^i(\hat{c}_t(\cdot), \hat{l}_t^{i,j}) \\ &+ \beta E_t\hat{V}^c(t+1, i, \hat{a}_{t+1}(\cdot), \epsilon_{t+1}^1, \epsilon_{t+1}^2, \bar{y}_{t+1}^1, \bar{y}_{t+1}^2) \end{aligned} \quad (37)$$

$$\hat{S}^c(t, i, a_t, \bar{y}_r^1, \bar{y}_r^2, tr) = v^i(\hat{c}_t(\cdot), L^{i,j}) + \beta E_tS^c(t+1, i, \hat{a}_{t+1}(\cdot), \bar{y}_r^1, \bar{y}_r^2, tr) \quad (38)$$

$$\hat{V}^c(t, i, a_t, \epsilon_t^1, \epsilon_t^2, \bar{y}_t^1, \bar{y}_t^2) = (1 - \hat{D}_t(\cdot))\hat{N}^c(t, i, a_t, \epsilon_t^1, \epsilon_t^2, \bar{y}_t^1, \bar{y}_t^2) + \hat{D}_t(\cdot)\hat{S}^c(t, i, a_t, \bar{y}_r^1, \bar{y}_r^2, t). \quad (39)$$

During the retirement period, we have

$$\begin{aligned} \hat{R}^c(t, i, a_t, \psi_t^1, \psi_t^2, \bar{y}_r^1, \bar{y}_r^2, tr) &= v^i(\hat{c}_t(\cdot), L^{i,j}) + \\ &\beta s_t^{i,j}(\psi_t^i) s_t^{p,j}(\psi_t^p) E_t\hat{R}^c(t+1, i, \hat{a}_{t+1}(\cdot), \psi_{t+1}^1, \psi_{t+1}^2, \bar{y}_r^1, \bar{y}_r^2, tr) + \\ &\beta s_t^{i,j}(\psi_t^i) (1 - s_t^{p,j}(\psi_t^p)) E_tR^s(t+1, i, \hat{a}_{t+1}(\cdot), \psi_{t+1}^i, \bar{y}_r^i, tr), \end{aligned} \quad (40)$$

where $s_t^{p,j}(\psi_t^p)$ is the survival probability of the partner of the person of gender i .

5 Estimation

We estimate our model on our two birth cohorts separately. For each cohort, we adopt a two-step estimation strategy, as done by Gourinchas and Parker (2002) and De Nardi, French, and Jones (2010 and 2016). We extend their approach to match the life cycle profiles of labor market participation and hours (in addition to savings).

In the first step, for each cohort, we use data on the initial distributions at age 25 for our model’s state variables, and we estimate or calibrate those parameters that can be cleanly identified outside our model. For example, we directly estimate from the data the probabilities of marriage, divorce, and death, as well as the wage processes while working and medical expenses during retirement.

In the second step, we use the method of simulated moments to estimate the remaining model parameters. For the 1945 cohort, we estimate 19 model parameters $(\beta, \omega, (\phi_0^{i,j}, \phi_1^{i,j}, \phi_2^{i,j}), (\tau_c^{0,5}, \tau_c^{6,11}), L^{i,j})$.⁹ For the 1955 cohort, we assume that the households of the 1955 cohort have the same discount factor β and weight on consumption ω as the 1945 cohort, and we estimate the remaining 17 parameters.

To perform the estimation, for each cohort, we use the model to simulate a representative population of people as they age and die, and we find the parameter values that allow simulated life cycle decision profiles to “best match” (as measured by a GMM criterion function) the data profiles for that cohort. The data that inform the estimation of the parameters of our model are composed of the following 448 moments for each cohort:

1. To better evaluate the determinants of labor market participation and their responses to changes in taxes and transfers, we match the labor market participation of four demographic groups (married and single men and women) starting at age 25 and up to age 65 (41 time periods for each group).
2. To better evaluate the determinants of hours worked and their responses to changes in taxes and transfers, we match hours worked conditional on participation for four demographic groups (married and single men and women) starting at age 25 and up to age 65 (41 time periods for each group).

⁹We normalize the time endowment for single men.

3. Because net worth, together with labor supply, is essential to smooth resources during the working period and finance retirement, we match net worth for three groups (couples and single men and women) starting at age 26 and up to age 65 (40 time periods for each group).¹⁰ Because people save to self-insure against shocks and for retirement, matching assets by age is essential to evaluate the effects of policy instruments and other forces, not only on saving but also on participation and hours.

The mechanics of our MSM approach draw heavily from De Nardi, French, and Jones (2010 and 2016) and are as follows. We discretize the asset grid, and, using value function iteration, we solve the model numerically (see Appendix D for details). This yields a set of decision rules that allows us to simulate life cycle histories for assets, participation, and hours. We keep track of a large number of artificial individuals, which are initially endowed with a value of the state vector drawn from the data distribution for each cohort at age 25, generate their histories, and use them to construct moment conditions and evaluate the match using our GMM criterion. We search over the parameter space for the values that minimize the criterion. We repeat the estimation procedure for each cohort.

Appendix E contains a detailed description of our moment conditions, the weighting matrix in our GMM criterion function, and the asymptotic distribution of our parameter estimates.

5.1 First-step estimation

Table 12 (in Appendix B) and Table 20 (in Appendix C) summarize our first-step estimated or calibrated model inputs. The procedures for estimating wages as a function of age and previous experience and earnings are new, as are the estimates of the probability of marriage and divorce by age, gender, and wage shocks. Appendix B details all of these inputs and reports additional first-steps inputs for both of each cohorts.

5.1.1 Wages

We assume that wages are composed of a persistent stochastic shock and a component that is a function of age, gender, and human capital. We measure human

¹⁰Net worth at age 25 is an initial condition.

capital at a given point in time as one’s average realized earnings up to that time. Thus, we allow past wages (and education, to the extent that it affects wages) and labor market experience to affect one’s wage today. We estimate this relationship from the PSID data.¹¹




Figure 4: Wage profile for single and married men and women at the average level of human capital by age and subgroup. Left panel: 1945 cohort. Right panel, 1955 cohort. PSID data

Figure 4 displays the average age-efficiency profiles computed from the estimated wage process that we estimate for men and women, evaluated at the average values of human capital or average accumulated earnings at each age, (\bar{y}_t) . It shows that, consistent with the evidence on the marriage premium, the wages of married men are higher than those of single men. In contrast, the wages of married women are lower than those of single women in our 1945 cohort, but this gap shrinks for our 1955 cohort because the average wage of married women has increased, while the average wage for single women has stagnated. This is due to a combination of both different

¹¹Human capital, measured as average past earnings, soaks up more heterogeneity in wages than education. Hence, we do not miss much by ignoring education when we take human capital into account. To see this, consider the following. For our baseline specification, we estimate a fixed effect regression of potential wage on age and human capital and their interactions with cohort and gender. As shown in Table 13 in Appendix B, it yields an R-square of 0.103. We have also run an alternative specification in which we run fixed effects regressions of potential wage on a polynomial in age, interacted with gender and education. The resulting R-square is 0.067. Thus, the variability in the wage data as measured by the R-square indicates that our measure of human capital explains more of the variability in the data than a typical measure of education. The economic intuition is that, conditional on years of education, types of major and quality of college imply much more variation in wages than the variation that is implied by our measure of human capital. In addition, we have also estimated a fixed effect regression which adds interactions with education for all of the variables already included in our baseline specification (human capital, age, cohort, and gender). This specification delivers an R-square of 0.116, which is only slightly higher than the one for our base case.

returns to human capital and accumulated human capital levels. The stagnation of men’s wages that we observe for our two cohorts is consistent with findings on wages over time reported by Acemoglu and Autor (2011) and Roys and Taber (2017).

Table 15 in Appendix B reports our estimates for the earnings shock processes. They imply that men and women face a similar persistence and earnings shock variance and that the initial variance upon labor market entry for men is a bit larger than that for women.

5.1.2 Marriage, divorce, spousal wage shocks, spousal assets, and Social Security benefits

We use the PSID to estimate the probabilities of marriage and divorce. Figure 14 in Appendix B displays our estimated benchmark probabilities of marriage for both cohorts. Men with higher wage shocks are more likely to get married, but this gap shrinks with age. In contrast, the probability of marriage for women displays little dependence on their wage shocks. The comparison with the 1955 cohort shows that the probability of getting married is smaller for the 1955 cohort, for both men and women. Figure 15 in Appendix B reports results for our benchmark estimation of divorce probabilities and shows that married men with lower wage shocks are more likely to get divorced. The probability of divorce decreases with age, as does the gap in the probabilities of divorce as a function of wage shocks. The probability of divorce for women displays less dependence on the wage shock. The comparison with the 1955 cohort shows that divorce rates are a bit lower in our more recent cohort, once we condition on age and wage shocks.

We also estimate the joint distribution of (the logarithm of) the wage shocks of new husbands and wives¹² by age and we assume it is lognormal. We find that the correlation of the logarithm of initial wage shocks between spouses is 0.27 in the 25-34 age group, 0.39 in the 35-44 group, and 0.45 after age 45. Because of these initial correlations and the high persistence of shocks that we estimate at the individual level, partners tend to have positively correlated shocks even after getting married.

Appendix B reports spousal assets and spousal Social Security earnings by spousal wage shocks in case of marriage next period for both of our cohorts.

¹²We assume it to be the same for both cohorts because the number of new marriages after age 25 is small during this time period.

5.1.3 Children

Figure 16 in Appendix B displays the average total number of children and average number of children in the 0-5 and 6-11 age groups by parental age. It shows that the number of children has decreased for married women and, to a smaller extent, for single women in the 1955 cohort compared to the 1945 cohort.

5.1.4 Health, mortality, and medical expenses

Health, survival, and medical expenses in old age interact in an important way to determine old-age longevity and medical expense risks. These risks, in turn, are affected by the structure of taxation and Social Security rules. For these reasons, it is important to capture the key aspects of health, mortality, and medical expenses to evaluate the effects of these programs.

We take these data from the HRS, and because we have no data after age 65 for the 1955 cohort, we assume that the 1955 cohort faces the same risks as the 1945 cohort in terms of health, mortality, and medical expenses.

Based on self-reported health status, we assume that health takes on two values, good and bad. Figure 17 in Appendix B reports our estimated health transition matrices by gender, age, and marital and health status. Women, married people, and healthy people have longer life expectancies (Figure 18 in Appendix B displays the survival probabilities by gender and marital and health status). Figure 19 in Appendix B displays the importance of medical expenditures after retirement. Average medical expenses climb fast past age 85 and are highest for single and unhealthy people.

5.2 Second-step estimation

Table 1 presents our estimated preference parameters for both cohorts.¹³ For the 1945 cohort, our estimated discount factor is 0.990, the same value estimated by De Nardi, French, and Jones (2016) on a sample of elderly retirees, and our estimated weight on consumption is 0.406. We assume that the 1955 cohort shares these preference parameters. While we normalize the total weekly time endowment of single men to 5,840 hours a year, and thus 112 hours a week, for our 1945 cohort, we

¹³Table 21 in Appendix F reports all of our estimated parameters for both cohorts and their standard errors.

Estimated parameters	1945 cohort	1955 cohort
β : Discount factor	0.990	Same as 1945 cohort
ω : Consumption weight	0.406	Same as 1945 cohort
$L^{2,1}$: Time endowment (weekly hours), single women	107	112
$L^{1,2}$: Time endowment (weekly hours), married men	107	101
$L^{2,2}$: Time endowment (weekly hours), married women	88	88
$\tau_c^{0,5}$: Prop. child care cost for children ages 0-5	30%	25%
$\tau_c^{6,11}$: Prop. child care cost for children ages 6-11	7%	19%
$\Phi_t^{i,j}$: Partic. cost	Fig. 23	Fig. 23

Table 1: Second-step estimated model parameters

estimate that single women have a total weekly time endowment of 107 hours a week. We interpret this as single women having to spend five more hours a week managing their household and rearing children (they have fewer children than married women but still more than single men) or taking care of elderly parents. The corresponding time endowments for married men and women are, respectively, 107 and 88 hours. This implies that people in the latter two groups spend 5 and 24 hours a week, respectively, running households, raising children, and taking care of aging parents.

Our estimates of non-market work time are remarkably similar to those reported by Aguiar and Hurst (2007), who find that, in the 1985 American Time Use Survey (ATUS) data set (when our 1945 cohort was 42 years old), men and women spent 14 and 27 hours a week, respectively, engaging in non-market work. Using more recent data, Dotsey, Li, and Yang (2014) find that, similarly to Aguiar and Hurst (2007), people spend 17 hours per week on average on activities related to home production. It should be noted that, even for a working woman, 24 hours can amount to, for example, spending seven hours each day on Saturday and Sunday and two hours a day the other five days by parenting, cooking, doing laundry, cleaning, organizing one's house, and taking care of one's parents. Thus, the data and model estimates are very consistent in the way households spend time running their households and providing care.

Our estimates for the 1945 cohort imply that the per-child child care cost of having a child ages 0-5 and 6-11 are, respectively, 30% and 7% of a woman's wage. In the PSID data, child care costs are not broken down by age of the child, but per-child

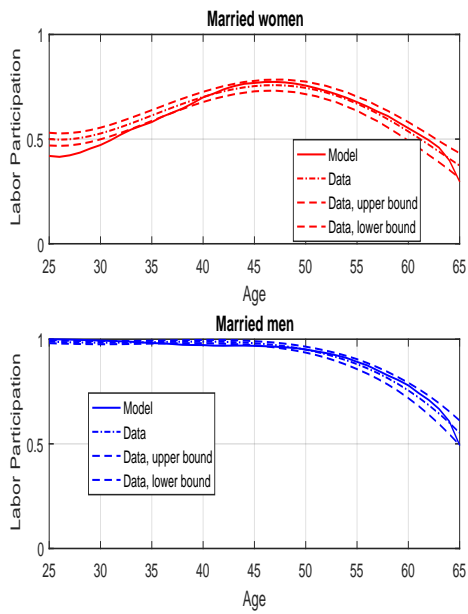
child care costs (for all children in the age range 0-11) of a married woman are 31% and 20% of her earnings at ages 25 and 30, respectively. Computing our model's implications, we find our corresponding numbers for a married woman are 23% and 18% of her earnings, respectively, at ages 25 and 30. Thus, our model infers child care costs that are similar to those in the PSID data.

For the 1955 cohort, we notice two main changes compared with the 1945 cohort. First, to help reconcile the lower hours worked by married men in this cohort, the model estimates that their available time to work and enjoy leisure decreases by six hours a week. Second, to help reconcile the slopes of hours and participation over the life cycle by married women in the presence of fewer children, the model estimates that the per-child child care costs of having younger children goes up, while that of having older children goes down. While decomposing the effects of changing labor supply between the two cohorts is very interesting (see, for instance, Attanasio, Low, and Sánchez-Marcos (2008) and Eckstein and Lifshitz (2011)), we abstract from analyzing it further because of space constraints.

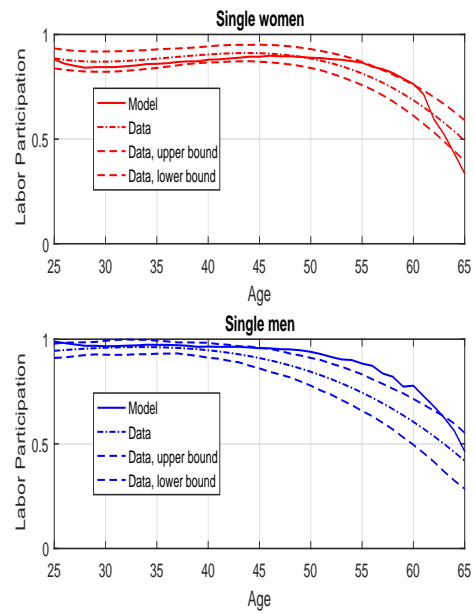
Figure 23 in Appendix F reports the age-varying time costs of working by age expressed as fraction of the time endowment of single men that are necessary to reconcile the labor market participation of our four groups of people in each cohort. Our estimated participation costs are relatively high when people are younger and, with the exception of single men, increase again after 45. The time costs of going to work might include factors other than commuting time. For instance, they might be higher when children are youngest because, for instance, during that period parents might need additional time to get their children back and forth from day care. They also show that, conditional on all aspects of our environment, the participation costs of married women are the lowest ones because married women face lower wages, have a smaller time endowment (because of time spent engaging in home production and child care), and tend to have higher-wage husbands who work.

5.3 Model fit

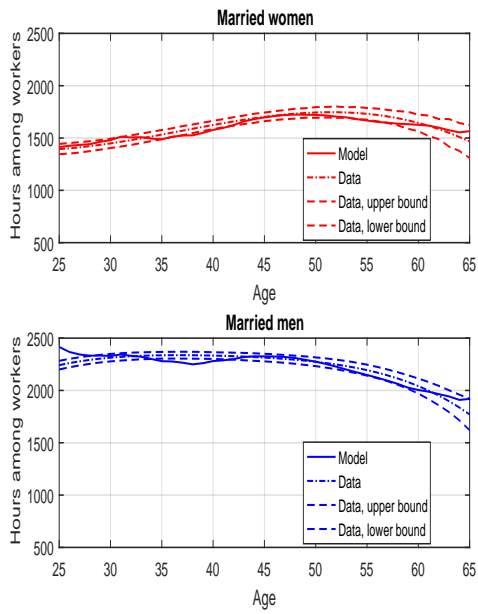
Figures 5 and 6 report our model-implied moments, as well as the moments and 95% confidence intervals from the PSID data for our 1945 cohort for the moments that we target in our estimation procedure. They show that our model matches participation, hours conditional on participation, and asset accumulation for all of




(a) Participation, couples



(b) Participation, singles



(c) Hours for workers, couples



(d) Hours for workers, singles

Figure 5: Model fit for participation (top graphs) and hours (bottom graphs) and average and 95% confidence intervals from the PSID data

our demographic groups.

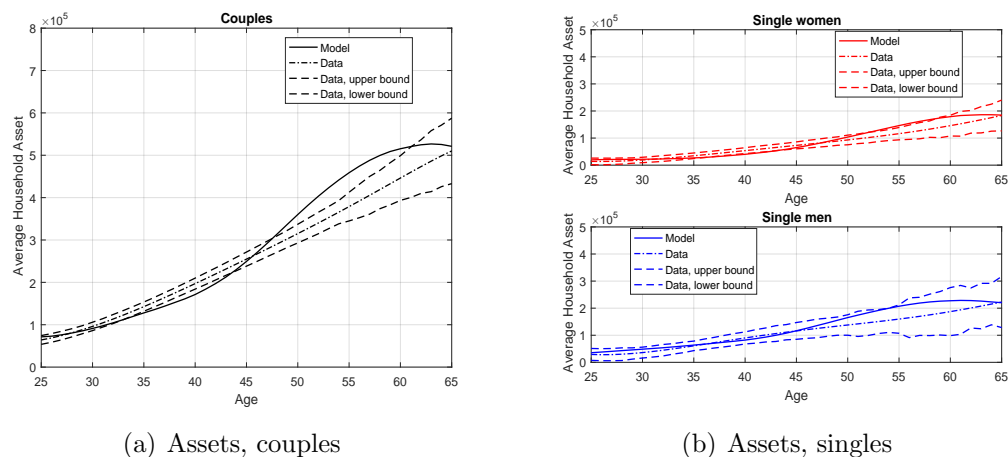


Figure 6: Model fit for assets and average and 95% confidence intervals from the PSID data

Figure 25 in Appendix G compares additional model implications for couples with those in the data for our 1945 cohort for moments that we do not target in estimation. They show that our estimated model also matches the fraction of couples with two workers, with only the husband working, with only the wife working, or with none working by age. They also display that our model produces reasonable implications for the hours worked over the life cycle for each type of couple. Our model fits the data well for the 1955 cohort too (to save on space, we show the graphs for the 1955 cohort in Appendix G).

Thus, our parsimoniously parameterized model reproduces all of these features of the data well, including those that are not matched by construction, which is remarkable given that it is tightly parameterized. In fact, we estimate 19 parameters and 448 targets for the 1945 cohort and 17 parameters and 448 targets for the 1955 cohort.

5.4 Identification

The fixed cost of participation by age and subgroup ($\Phi_t^{i,j}$) especially affects participation by subgroup over the life cycle. The available time endowment ($L^{i,j}$) has first-order effects on hours worked by workers. Child care costs have a larger effect on hours than participation and especially affect hours worked by women when they

have young children. This effect is especially large for married women, as they have more children than single women.

The discount factor (β) has large effects on savings. The weight on consumption (ω) affects the intratemporal substitution between consumption and leisure and thus affects hours worked at all ages. Because the wage is increasing with human capital (and past hours worked), a high ω increases the value of consumption at all ages but has a larger impact on the hours of older workers relative to younger workers.

6 Model validation in terms of elasticities

To help build confidence in our model's responses to policy changes, we report its labor supply elasticities. Table 2 shows the (compensated) elasticities of participation and hours among workers with respect to an anticipated change to their own wage.¹⁴ It shows that, first, the elasticity of participation of women is larger than that of men, for both married and singles. Second, it shows that married men have the lowest elasticity of participation. Third, it shows that the elasticity of participation for all groups is largest around retirement age, a finding that confirms that of French (2005) for men. Fourth, our elasticities are consistent with those in Liebman, Luttmer, and Seif (2009), which uses HRS data for people over age 50 and variation stemming from Social Security rules. Their results imply that the yearly elasticity at the extensive margin is 0.7 for the sample of men and women, 1.1 for women, and 0.2 (but not statistically significant) for men. At the intensive margin, their elasticity is 0.4 for men and women, 0.7 for men, and -0.3 (but not statistically significant) for women. Thus, their estimated labor supply elasticities at the extensive and intensive margins are consistent with those in our 50 and 60 age groups. More generally, our model's implied elasticities at all ages are in line with those in the vast existing literature, as surveyed in Blundell and MaCurdy (1999) and more recently estimated by Attanasio et al. (2018).

While important to compare with the empirical estimates, the compensated wage elasticities are not necessarily indicative of how participation and hours would change as a result of a wage change that permanently affects all of the population and which is more similar to that implied by a permanent tax change at all income levels.

¹⁴For this computation, we temporarily increase the wage for only one age and one group at a time (married men, married women, single men, or single women) by 5%.

	Participation				Hours among workers			
	Married		Single		Married		Single	
	W	M	W	M	W	M	W	M
30	1.0	0.0	0.5	0.2	0.2	0.3	0.4	0.3
40	0.7	0.1	0.4	0.2	0.3	0.5	0.5	0.5
50	0.6	0.2	0.4	0.5	0.5	0.5	0.8	0.5
60	1.1	0.8	1.4	2.0	0.4	0.2	0.5	0.3

Table 2: Model-implied elasticities of labor supply

To help shed light on what we should expect from our policy experiments, here we report the effects of a permanent increase of 5% in the wage schedule of married women when the wage structure of the three other demographic groups remains the same. The left panel of Figure 7 shows that a permanent wage increase for married women implies a much higher, and U-shaped, elasticity of participation for married women, which peaks at 2.5 at age 25. It also reports the cross-elasticities of the other groups to changes in the wages of married women. The right panel highlights that permanent wage changes can lead to high increases in married women’s participation, with participation being 4-7 percentage points higher over all of their life cycle. It also shows that the participation of single women rises because they expect to get married and obtain higher wages (and higher returns to their accumulated human capital) upon marriage. There is little response in the participation of single men. In contrast, married men’s participation after age 40 decreases when women’s wage schedule increases. This shows that modeling men’s labor supply is important to assess the effects of reforms affecting the wages of married women in a long-lasting way.

7 Policy experiments: eliminating marital Social Security benefits and joint taxation

We now turn to evaluating the effects of various policy reforms. We first show the labor supply and savings responses resulting from the elimination of various marital policies, and we then evaluate their welfare implications.

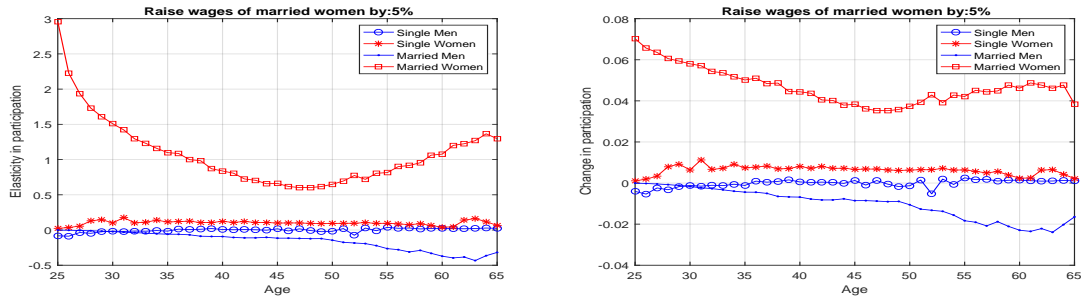


Figure 7: Elasticity of participation (left graph) and change in participation (right graph) for a 5% permanent increase in the wage schedule for married women. Effect on all four demographic groups. Model implications.

7.1 Outcomes

For each policy counterfactual, we compute two sets of results. The first one balances the government budget constraint by adjusting the proportional component of the income tax, while the second one keeps the government budget constraint unbalanced. Because of space constraints, we report the effects of the latter set of experiments in Appendix H.

7.1.1 Eliminating spousal Social Security benefits, 1945 cohort

According to the current Social Security rules, one’s spouse can receive half of his or her partner’s contribution while their partner is alive and all of the benefits of their deceased spouse. This provision potentially has three effects. First, it discourages the labor supply of the secondary earner, given that he or she can benefit from spousal benefits. Second, it encourages the labor supply of the main earner, who is also working to provide Social Security benefits to the secondary earner spouse. Third, it reduces retirement savings because it raises the annuitized income flow of the secondary earner or non-participant.

When eliminating both spousal Social Security benefits, the government runs a budget surplus and can cut the proportional component of the income tax from 4.0% to 1.8%. The left panel of Figure 8 shows that the participation of married women is, respectively, 10, 11, and 4 percentage points higher at ages 25, 55-60, and 65 without spousal Social Security benefits. In contrast, married men decrease their participation starting at age 55, and their participation is 6 percentage points lower by age 65. A

model in which married men cannot change their participation or can do it only after a certain age would miss this effect. The participation of single women at ages 25-30 increases (by 3 percentage points) because, should they marry, they now expect no Social Security benefits coming from their spouse’s labor supply. As they age, the probability that they marry becomes negligible, and the effect on their participation of the elimination of spousal benefits fades.




Figure 8: Changes in participation (left panel) and labor income (right panel) after the elimination of all spousal Social Security benefits when the income tax is reduced to balance the government budget.

An important reason why these reforms have such large effects on the labor supply of married women resides in the initial distribution of potential wages of men and women at age 25. Table 3 shows that, in the 1945 cohort, 60% of women and only 20% of men belong to the bottom two quintiles of wages at age 25. Thus, most women have low wages and tend to be secondary earners in this cohort. For this reason, they react strongly to the elimination of spousal benefits.

	Wage quintile				
	1	2	3	4	5
Men	7.9 %	12.4 %	21.3 %	28.2 %	30.2 %
Women	32.3 %	27.8 %	18.7 %	11.6 %	9.6 %

Table 3: Distribution of men and women across potential wage quintiles at age 25, 1945 cohort, PSID data.

Groneck and Wallenius (2017) and Kaygusuz (2015) study the effects of marital Social Security benefits in simpler models than ours in which, for instance, men

cannot change their labor supply and women can do so to a limited extent.¹⁵ They report that, over all of the working period, their model implies an increase in the participation by married women of 6.4 and 6.1 percentage points, respectively. Because we also allow men to adjust their labor supply, and they choose to reduce it in older ages, and because women (as in the data) have more flexibility in their hours worked, we find effects that are a bit larger but in the same ballpark.

The right panel of Figure 8 reports changes in labor income for our four demographics groups. Married women work more, accumulate more human capital, and earn more as a result of the reform. Married women’s labor income is about, respectively, 18%, 12%, and 11% higher at ages 25, 55-60, and 65. The labor income of married men drops by about 13% by age 65.

	Couples	Single men	Single women
Savings, balanced government budget	14.9%	7.8%	11.2%

Table 4: Change in savings at age 66, in percentages, as a result of removing spousal Social Security benefits when the income tax is reduced to balance the government budget.

Table 4 shows the resulting changes in assets at retirement time. The reform increases savings by reducing government payments to spouses and widows during retirement, and assets at retirement go up by 14.9%, 7.8%, and 11.2% for couples, single men, and single women, respectively.

7.1.2 Eliminating joint income taxation, 1945 cohort

Figure 9 displays the effects on participation of having everyone file as singles (the married men file as single men and the married women as single women) and reducing the proportional component of the income tax to balance the government budget (from 4.0% to 3.5%). As a result of this policy, the participation of married women increases by more than 20 percentage points until age 35 and by 10 percentage points between ages 45 and 60. The participation of single women also increases slightly until age 60. The right panel of Figure 1 provides the key intuition for this result:

¹⁵Their models are also less rich along other important dimensions and are calibrated rather than estimated.




Figure 9: Changes in participation after the elimination of joint income taxation when the income tax is reduced to balance the government budget.

the marginal tax rates for married women working are much lower when they do not file jointly with their husbands.

Guner, Kaygusuz, and Ventura (2012a) study the switch from current U.S. taxation to single filer taxation in a calibrated model of a steady state and find that the labor supply of married women goes up by 10-20 percentage points. Our effects on the labor supply of married women are thus close to theirs.

7.1.3 Eliminating spousal Social Security benefits and joint income taxation, 1945 cohort

This policy change implies a reduction in the proportional component of the income tax from 4.0% to 2.0%. Figure 10 displays the participation profiles in our

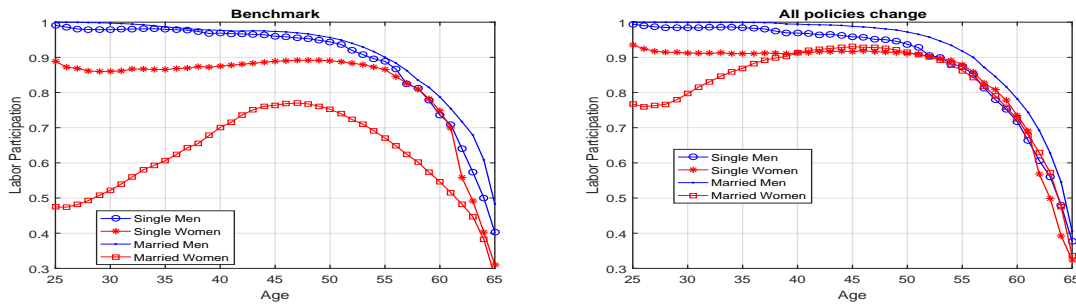


Figure 10: 1945 cohort: Participation in the benchmark (left panel) and reformed economy (right panel) after the elimination of all spousal Social Security benefits and joint income taxation when the income tax is reduced to balance the government budget.

benchmark economy and this counterfactual economy. Eliminating spousal Social

Security benefits and joint income taxation has a large effect on the participation of married and single women. To make magnitudes clearer, the left panel of Figure 11 plots the differences in participation between the benchmark and this counterfactual for each group of people by age. It shows that the participation of married women is 16-30 percentage points higher until age 62 in the no-marital-provisions economy. The participation of single women is about 5 percentage points higher until age 40. The participation of married men is higher in their middle age, reaching a peak of 2 percentage points higher than in the benchmark, but is 8 percentage points lower than in the benchmark at age 65. Thus, the timing of their participation changes over their life cycle. This highlights the importance of also modeling their labor supply behavior over their life cycle, in addition to that of their wives when we change provisions that affect both members in the household.




Figure 11: Changes in participation (left panel) and labor income (right panel) after the elimination of all spousal Social Security benefits and joint income taxation when the income tax is reduced to balance the government budget.

	Couples	Single men	Single women
Savings, balanced government budget	20.3%	8.8%	14.8%

Table 5: Change in assets at age 66, in percentages, as a result of removing spousal Social Security benefits and joint income taxation when the income tax is reduced to balance the government budget.

Table 5 displays the effects on savings at retirement time. Couples now save 20.3% more for retirement, while single men and women save, respectively, 8.8% and 14.8% more.

7.1.4 Eliminating Marital Social Security benefits and joint taxation for the 1955 cohort

We now turn to studying the effects of marriage-related taxes and Social Security benefits for the 1955 cohort. In the interest of space, we report results only for the case in which we eliminate all three marriage-related provisions at the same time.

The left panel of Figure 12 displays the difference in the participation profile. This graph shows that eliminating all marital-related provisions also has large effects for the 1955 cohort, in which labor supply participation is much higher to start with. Thus, the effects of these policies on a relatively younger cohort with a much higher participation of married women continue to be very large.

The effects of increased labor market experience on wages are similar to those in the 1945 cohort, and, as for the 1945 cohort, increased wages and participation (hours increase little for the workers) imply higher average earnings of \$5,000-6,000 per year for married women and \$3,000 for single women for most of their life cycle. Average earnings of married men start dropping earlier for this cohort, that is, at age 50, compared with age 55 for the 1945 cohort, but their drop is smaller by age 65 (see right panel in Figure 12).




Figure 12: 1955 cohorts: Changes in participation (left panel) and labor income (right panel) after the elimination of all spousal Social Security benefits and joint income taxation when the income tax is reduced to balance the government budget.

Table 6 displays the effects on savings at retirement time. Couples now save 19.7% more for retirement, while single men and women save, respectively, 8.4% and 14.9% more.

Comparing Tables 3 and 7 highlights that the fraction of women in the lowest wage quintile has decreased and the fraction of women in the highest one has increased from

	Couples	Single men	Single women
Savings, balanced government budget	19.7%	8.4%	14.9%

Table 6: Change in assets at age 66, in percentages, as a result of removing spousal Social Security benefits and joint income taxation when the income tax is reduced to balance the government budget.

	Wage quintile				
	1	2	3	4	5
Men	10.4 %	14.0 %	19.1 %	26.0 %	30.4 %
Women	28.1 %	25.1 %	20.6 %	15.1 %	11.1 %

Table 7: Distribution of men and women across potential wage quintiles at age 25, 1955 cohort, PSID data.

the 1945 to the 1955 cohort, but it is still the case that, even in the 1955 cohort, most women tend to have lower wages and thus be secondary earners in this cohort, and therefore respond strongly to the elimination of marital provisions.

7.2 Welfare

To evaluate welfare changes, we calculate the asset compensation required for each household at age 25 to stay in the benchmark economy and report it as a fraction of average income in the benchmark economy. Thus, negative asset compensations mean that households are better off in the benchmark economy. The first three columns in Table 8 report the average welfare gains or losses conditional on one’s marital status at age 25. We also report the fraction of households gaining and losing and the average gains and losses among each of these groups.

The top panel of counterfactuals refers to the 1945 birth cohort. The first set of results in the top panel compares our benchmark economy with one in which there are no marital Social Security benefits and taxes remain unchanged despite the resulting government surplus. On average, couples would need to be compensated by a onetime asset transfer at age 25 that is equivalent to 0.25 of average earnings in the economy, while single women would require 0.23 average earnings, as they expect to marry and potentially lose these benefits after marriage. While sizable, these welfare costs are not very large because, as of age 25, people already know that these benefit changes

	All			Winners			Losers		
	Couples	SW	SM	Couples	SW	SM	Couples	SW	SM
1945 cohort									
(1). Remove Social Security spousal benefits, unbalanced budget									
Average	-0.25	-0.23	0.31	0.00	0.00	0.31	-0.25	-0.23	-0.02
Percentage				0.0	0.0	100.0	100.0	100.0	0.0
Percentage (No marriage prob.)				0.0	0.0	0.0	100.0	0.0	0.0
(2). Remove Social Security spousal benefits, balanced budget									
Average	0.71	0.20	1.30	0.71	0.22	1.30	0.00	-0.04	0.00
Percentage				100.0	93.4	100.0	0.0	6.6	0.0
(3). Remove joint income taxation, unbalanced budget									
Average	0.13	-0.19	1.04	0.35	0.07	1.04	-0.18	-0.20	0.00
Percentage				58.4	4.5	100.0	41.6	95.5	0.0
(4). Remove joint income taxation, balanced budget									
Average	0.33	-0.10	1.25	0.45	0.11	1.25	-0.09	-0.15	0.00
Percentage				78.5	17.9	100.0	21.5	82.1	0.0
(5). Remove all marital related policies, balanced budget									
Average	0.83	0.03	2.24	0.84	0.31	2.24	-0.04	-0.13	0.00
Percentage				98.9	35.8	100.0	1.1	64.2	0.0
1955 cohort									
(6). Remove all marital related policies, balanced budget									
Average	0.75	0.21	1.31	0.77	0.31	1.31	-0.05	-0.05	-0.02
Percentage				97.2	70.9	100.0	2.8	29.1	0.0

Table 8: Asset compensation required for staying in the benchmark economy, normalized as a fraction of average income in the benchmark economy. SM: single men, SW: single women. Top line for each experiment: average welfare gain or loss. Bottom line for each experiment: fraction in that group gaining or losing welfare.

will take place at retirement time and they have many years to work and save to make up for the potential loss in benefits. In contrast, single men benefit from this policy change because their wives will work more, earn more, and accumulate more human capital after they marry, and single men do not take into account their future wife’s disutility from working more. The remaining columns in the table distinguish the effects for winners and losers. The three “winners” columns show that all couples and single women, on average, lose from this policy, whereas all single men gain from it. The third row in this panel reports the percentage of people winning and losing for the same policy experiment but in an economy in which people no longer marry after

age 25. It clarifies the role of marriage expectations in driving our welfare calculations and highlights that this benefit removal would have no effects on single people, who don't expect to receive this benefit anyway if they do not expect to get married.

The second set of results removes marital Social Security provisions and balances the government budget by reducing the proportional component of the income tax from 4.0% to 1.8%. The first three columns display large welfare gains: couples would be willing to pay, on average, an asset amount that corresponds to 0.7 times average income in welfare terms, single women 0.2 times, and single men 1.3 times. The second set of three columns shows that all couples, 93.4% of single women, and all single men would benefit from these changes. The last set of three columns shows that the 6.6% of the single women who lose would face very small losses. These are women whose initial human capital is very low and are heavily relying on marital benefits. Thus, this counterfactual suggests that eliminating these benefits while reducing the income tax would benefit the vast majority of the young population and would have only small welfare costs for a small fraction of single women.

The third set of results makes everyone file as single people. The first line in this panel does not balance the government budget constraint and shows that the willingness to pay for this policy, measured as a onetime asset amount as a fraction of average income, is equal to 0.13 and 1.04 for couples and single men, respectively. In contrast, single women would lose and require an asset compensation of 0.19 average income. This happens because they know they will be working more and enjoying less leisure in the future, and especially so after marriage. The winners and losers columns reveal that 58.4% of couples, 4.5% of single women, and 100% of single men would favor this policy and that both gains for the winners and losses for the losers would be sizable.

The fourth set of results balances the government budget constraint by reducing the income tax (from 4.0% to 3.5%) and generates more winners with larger welfare gains and fewer losers, who also experience smaller welfare losses than in the previous experiment. For instance, 78.5% of the couples would be willing to give up an asset amount corresponding to 0.84 average income to live under this policy, while the compensation for the remaining 21.5% would amount to assets equal to 0.09 of average income.

The fifth set of results for the 1945 cohort eliminates all of the marriage-related policies that we consider and balances the government budget constraint by reducing

the income tax, which goes down from 4.0% to 2.0%. This policy change generates the largest aggregate welfare gains among the set that we consider for the 1945 cohort: 0.83, 0.03, and 2.24 times average income for couples, single women, and single men, respectively. Among couples, 98.9% would gain, compared with 35.8% of single women and 100% of single men. The bigger losers coming out of this policy are 64.2% of single women, who lose, on average, 0.13 of average income.

The results in the last panel refer to the 1955 cohort and show that there would also be large aggregate gains from removing marriage-related provisions and reducing the income tax and that single women in this cohort would be less disadvantaged by this policy than single women in the 1945 cohort: only 29.1% of them would lose, compared with 64.2%. In addition, their loss would be much smaller (0.05 average income, compared with 0.13 in the 1945 cohort). In both cohorts, only a minority of couples would lose and would experience a small welfare loss.

Overall, our policy experiments thus indicate that removing marriage-related taxes and Social Security benefits would increase female labor supply and the welfare of the majority of the populations, whereas the rest would only bear small welfare costs.

8 Changes in marriage and divorce patterns in response to policy

Because we study labor supply and savings responses to the elimination of two important marriage-based policies, the question of the robustness of our findings to changes in marriage and divorce patterns naturally arises. To address this question, we first turn to summarizing previous empirical findings on the effects of changes in Social Security rules and income taxes on marriage and divorce patterns. Then, to evaluate the robustness of our results, we perform policy experiments in which marriage and divorce exogenously change at the same time as we eliminate marriage-based taxes and Social Security benefits by more than what has been found in the empirical literature.

8.1 The effects of Social Security and income taxes on marriage and divorce in the empirical literature

Before 1977, Social Security spousal benefits were available to the secondary earner in case of divorce after 20 years of marriage. After that date, the threshold for eligibility became 10 years of marriage before divorce.¹⁶ Dickert-Conlin and Meghea (2004) examine the 1977 U.S. policy switch (using data from the U.S. National Vital Statistics System, the 1980 census, and the Current Population Survey) and conclude that it had no effect on divorces and remarriages. Goda, Shoven, and Slavov (2007) also find no impact of the 10-year eligibility discontinuity on divorces (using data from the PSID Marital History File). Dillender (2016) confirms that these rules have no effects overall and small ones on a very small number of people: those who married late. Thus, previous literature indicates that the effects of Social Security benefits kinks are negligible in the United States.

Turning to the effect of income taxes on marriages and divorces, Alm and Whittington (1995) use time series data from 1947 to 1988 and argue that “the magnitude of this impact is quite small. This result suggests that some individuals respond to tax incentives in their marriage choices, but that for many individuals taxes do not affect these decisions.” Alm and Whittington (1997) use data from the PSID and estimate a discrete-time hazard model of the probability of divorce from the first marriage. They conclude that “couples respond to tax incentives in their decision to divorce, although these responses are typically small.” Alm and Whittington (1999) utilize the same data to estimate a discrete-time hazard model of the time to first marriage from 1968 to 1992, and uncover that the income tax has no effect on the marriage decisions of men and only a small effect on the marriage decisions of women. They thus conclude that, in the context of the United States, “In general, the impacts of the income tax variables, even when statistically significant, are small.”

Thus, for the United States, previous empirical studies on the impact of income taxation and Social Security benefits on marriage and divorce find either no significant effects or very small effects that apply to tiny groups of people.

Looking into welfare programs, Low et al. (2018) study the U.S. subpopulation of low-education mothers on welfare and the 1996 welfare reform, which was meant to

¹⁶We do not model this part of the benefits because the fraction of people divorcing after 10 years of marriage is small, and this addition would add great computational complexity to our framework.

encourage labor supply by welfare recipients and reduce marital disincentives. They use the Survey of Income and Program Participation (SIPP) data and document that the reform greatly reduced welfare recipience and increased labor participation of mothers, but had no effects on marriage and fertility. They do find an effect on divorce rates, which declined from 0.9% before the reform to 0.7% after the reform. This is a non-trivial drop as a fraction of divorces, but in absolute terms, the reform reduced a very small number to a tiny one and refers to a small population.

More broadly on the effects of welfare programs, a survey by Moffitt (1998) concludes, “Most find that the majority of studies show either no significant effects of AFDC and other welfare programs, effects that are statistically significant but small in magnitude, a set of mixed effects indicating some that are favorable and some unfavorable, or effects that occur only for some specific types of programs. Although the research reviewed in these chapters does not support a finding of no effect whatsoever of welfare programs on demographic behavior, it would be difficult to argue that the research often indicates very sizable or stable effects.”

Persson (2017) studies the elimination of marital survivorship benefits that took place in 1989 in Sweden and infers larger effects than those found for the United States. In terms of comparison with our work, her main finding is that the divorce rate increased by 10% as a result of the elimination of marital survivorship benefits. Although this effect is sizeable in percentage terms, it is a small change from the standpoint of the overall population because the divorce rate is small.

In comparing findings for the United States and Sweden, it is also important to keep in mind that cultural and religious factors are important reasons why people marry and stay married, and that marriage is much more widespread in the United States than in Sweden. For instance, United Nations (Department of Economic and Social Affairs, 2012) reports that in 1985, a time period preceding the 1989 Swedish marital benefits reform, only 35.8% of the 25- to 29-year-old Swedish women were married, compared with 64.3% in the United States. In addition, in the 1980s unmarried women in the United States accounted for 18% of live births, compared to 40% in Sweden (Sorrentino, 1990).

8.2 Robustness of policy results to large changes in marriage and divorces

In this subsection, we compare the effects of a policy experiment in which we eliminate joint income taxation of couples and Social Security marital benefits (for the 1945 cohort) for given marriage and divorce patterns with the effects of the same policy when there are also two possible alternative changes in marriage and divorce patterns. In the first robustness exercise, the policy decreases marriage rates by 20% and increases divorce rates by 20%. Alternatively, in the second robustness exercise, the policy increases marriage rates by 20% and decreases divorce rates by 20%. In both, we also balanced the government budget.




Figure 13: Differences in participation after the elimination of all spousal Social Security benefits and joint income taxation for the 1945 cohort. Left panel: benchmark economy with unchanging marriage and divorce after the policy change. Middle panel: 20% lower marriage probability and 20% higher divorce rate after the policy change. Right panel: 20% higher marriage probability and 20% lower divorce rate after the policy change.

Figure 13 highlights several important findings. First, all changes in participation of the four groups are very similar whether marriage and divorce patterns change or not. Second, comparing the left panel (no marriage and divorce changes) and the middle panel (decreased marriage and increased divorce probability) shows that a reform that lowers the probability of marriage and raises the probability of divorce makes women more self-reliant on their own labor supply and human capital. Married women work more (accumulating more human capital) to edge against divorce risk. Single women are less likely to marry and also work more (accumulating more human capital). Comparing the left panel (no marriage and divorce changes) and the right panel (increased marriage and decreased divorce probability) highlights that increasing the marriage rate and lowering the divorce rate has the opposite effect, but that

this effect is small and does not change the conclusions reached in our benchmark policy experiment.

	Couples	Single men	Single women
Benchmark	20.3%	8.8%	14.8%
Low marriage, high divorce	20.3%	7.6%	14.7 %
High marriage, low divorce	21.1%	11.7%	15.9 %

Table 9: Change in assets at age 66, in percentages, as a result of removing spousal Social Security benefits and joint income taxation.

Table 9 displays the effects on savings at retirement time for the three experiments. The second and third rows show that the effects on savings of couples, who make up the vast majority of the population, are also very robust to changes in expected marital patterns. Our results are thus robust to large changes in marriage and divorce rates.

9 Conclusions

We estimate a model of labor supply and savings for single and married people that allows for a rich representation of the risks that people face over their entire life cycle and for the important provisions of taxes and Social Security for singles and couples. We do so for both the 1945 and the 1955 birth cohorts, and we show that our model fits the data very well, including along important dimensions that it was not meant to match by construction, such as the elasticities of labor supply. We find that the fact that young women entering the labor market have much lower wages than those of men and the time and monetary costs that children imply are important determinants in the labor supply of single and married men and women.

We use our model to evaluate the effect of marriage-based Social Security benefits and the marriage tax bonus and penalty. We find that these marriage-based provisions have a strong disincentive effect on the labor supply not only of married women, but also of single young women who expect to get married. This lower participation reduces their labor market experience, which, in turn, reduces their wages over their life cycle. These provisions also induce married men to work longer at their careers and depress the savings of couples. Our findings are robust to changes in marriage

and divorce patterns. These effects are very similar for the 1945 and 1955 birth cohorts, even though the labor market participation of young married women in the 1955 cohort is over 10 percentage points higher than that of the 1945 cohort. We also show that, if the government surplus resulting from the elimination of marriage-related provisions were used to lower income taxation, there would be large welfare gains for the vast majority of the population and the few losing would experience small welfare losses.

Our paper provides several contributions. First, it is the first estimated structural model of couples and singles that allows for participation and hours decisions of both men and women, including those in couples, in a framework with savings. Second, it is the first paper to study all marriage-related taxes and benefits in a unified framework. Third, it does so by allowing for the large observed changes in the labor supply of married women over time by studying two different cohorts. Fourth, our framework is very rich along dimensions that are important in the study of our problem, including labor market experience affecting wages and carefully modeling survival, health, and medical expenses in old age, and their heterogeneity by marital status and gender.

References

- [1] Acemoglu, Daron, and David Autor. 2011. “Skills, Tasks and Technologies: Implications for Employment and Earnings.” in D. Card and O. Ashenfelter, ed., *Handbook of Labor Economics, volume 4*, (North Holland), 1043-1171.
- [2] Aguiar, Mark, and Erik Hurst. 2007. “Measuring Trends in Leisure: The Allocation of Time Over Five Decades.” *Quarterly Journal of Economics* 122(3), 969-1006.
- [3] Alm, James, and Leslie A. Whittington. 1995. “Income Taxes and the Marriage Decision.” *Applied Economics* 27(1), 25-31.
- [4] Alm, James, and Leslie A. Whittington. 1997. “Til Death or Taxes Do Us Part: The Effect of Income Taxation on Divorce.” *The Journal of Human Resources* 32(2), 388-412.
- [5] Alm, James, and Leslie A. Whittington. 1999. “For Love or Money? The Impact of Income Taxes on Marriage.” *Economica* 66, 297-316.
- [6] Altonji, Joseph G., and Lewis M. Segal. 1996. “Small-Sample Bias in GMM Estimation of Covariance Structures.” *Journal of Business and Economic Statistics*, 14(3), 353–366.
- [7] Attanasio, Orazio, Hamish Low, and Virginia Sánchez-Marcos. 2008. “Explaining Changes in Female Labor Supply in a Life-Cycle Model.” *American Economic Review*, 98(4), 1517-1552.
- [8] Attanasio, Orazio, Peter Levell, Hamish Low, and Virginia Sánchez-Marcos. 2018. “Aggregating Elasticities: Intensive and Extensive Margins of Women’s Labor Supply.” *Econometrica*, 86(6), pp. 2049-2082.
- [9] Barro, Robert J., and Charles J. Redlick, 2011. “Macroeconomic Effects of Government Purchases and Taxes.” *Quarterly Journal of Economics*, February 2011, 51–102.
- [10] Benabou, Roland. 2002. “Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?” *Econometrica*, 70(2), 481–517.

- [11] Bick, Alexander, and Nicola Fuchs-Schundeln. 2018. "Taxation and Labour Supply of Married Couples across Countries: A Macroeconomic Analysis." *Review of Economic Studies*, 85(3), 1543-1576.
- [12] Black, Sandra E., Diane Whitmore Schanzenbach, and Audrey Breitwieser. 2017. "The Recent Decline in Women's Labor Force Participation" in *The 51%: Driving Growth through Women's Economic Participation*, October. Brookings.
- [13] Blundell, Richard, and Tom MaCurdy. 1999. "Labour Supply: A Review of Alternative Approaches" in *Handbook of Labor Economics*, O.Ashenfelter and D. Card (eds.), North-Holland, December 1559-1695.
- [14] Blundell, Richard, Jack Britton, Monica Costa Dias, and Eric French. 2017. "The impact of health on labour supply near retirement." Mimeo.
- [15] Borella, Margherita, Mariacristina De Nardi, and Eric French. 2017. "Who Receives Medicaid in Old Age? Rules and Reality." *Fiscal Studies*, forthcoming.
- [16] Browning, Martin, and Pierre-André Chiappori. 1998. "Efficient Intra-Household Allocations: A General Characterization and Empirical Tests." *Econometrica*, 66(6), 1241-78.
- [17] Casanova, Maria. 2012. "Happy Together: A Structural Model of Couples' Joint Retirement Choices." Mimeo.
- [18] Chiappori, Pierre-André. 1988. "Rational Household Labor Supply." *Econometrica*, 56, 63-89.
- [19] Chiappori, Pierre-André. 1992. "Collective Labor Supply and Welfare." *Journal of Political Economy*, 100(3), 437-67.
- [20] De Nardi, Mariacristina, Eric French, and John B. Jones. 2010. "Why Do the Elderly Save? The Role of Medical Expenses," *Journal of Political Economy*, 118, 39-75.
- [21] De Nardi, Mariacristina, Eric French, and John B. Jones. 2016. "Medicaid Insurance in Old Age," *American Economic Review*, 106(11), 3480-3520.
- [22] De Nardi, Mariacristina, Svetlana Pashchenko, and Ponpoje Porapakarm. 2017. "The Lifetime Costs of Bad Health" NBER working paper 23963.

- [23] Dickert-Conlin, Stacy and Christian Meghea. 2004. “The Effect of Social Security on Divorce and Remarriage Behavior,” Working paper no. 9, Center for Retirement Research.
- [24] Dillender, Marcus. 2016. “Social Security and Divorce.” *B.E. Journal of Economic Analysis and Policy*, 16(2), 931-971.
- [25] Dotsey, Michael, Wenli Li, and Fang Yang, 2014, “Consumption and Time Use over the Life Cycle.” *International Economic Review*, 55(3), 665-692.
- [26] Duffie, Darrell, and Kenneth J. Singleton. 1993. “Simulated Moments Estimation of Markov Models of Asset Prices.” *Econometrica*, 61(4), 929–952.
- [27] Eckstein, Zvi, and Osnat Liftshitz. 2011. “Dynamic Female Labor Supply.” *Econometrica* 79(6), 1675-1726.
- [28] Eckstein, Zvi, Michael Keane, and Osnat Liftshitz. 2019. “Career and Family Decisions: Cohorts Born 1935–1975.” *Econometrica*, 87(1), 217-253.
- [29] Fadlon, Itzik, and Torben H. Nielsen. 2015. “Family Labor Supply Responses to Severe Health Shocks.” NBER Working Paper 21352.
- [30] Fitzgerald, John, Peter Gottschalk, and Robert Moffitt. 1998. “An Analysis of Sample Attrition in Panel Data: The PSID.” *Journal of Human Resources*, 33(2), 251-299.
- [31] French, Eric. 2005. “The Effects of Health, Wealth, and Wages on Labor Supply and Retirement Behavior.” *Review of Economic Studies*, 72(2), 395-427.
- [32] Gourinchas, Pierre-Olivier, and Jonathan A. Parker. 2002. “Consumption Over the Life Cycle.” *Econometrica* 70(1), 47–89.
- [33] Goda, Gopi Shah, John Shoven and Sita Nataraj Slavov. 2007. “Social Security and the Timing of Divorce,” NBER Working Paper 13382.
- [34] Groneck, Max, and Johanna Wallenius. 2017. “It Sucks to Be Single! Marital Status and Redistribution of Social Security. Working Paper.
- [35] Guner, Nezih, Remzi Kaygusuz and Gustavo Ventura. 2012a. “Taxation and Household Labour Supply.” *Review of Economic Studies*, 79, 1113-1149.

- [36] Guner, Nezih, Remzi Kaygusuz and Gustavo Ventura. 2012b. “Income Taxation of U.S. Households: Basic Facts.” CEPR Discussion Paper 9078.
- [37] Heathcote, Jonathan, Kjetil Storesletten and Giovanni L. Violante. (2014). “Consumption and labor supply with partial insurance: An analytical framework.” *The American Economic Review*, 104(7), 2075-2126.
- [38] Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes, 1995. “Precautionary Saving and Social Insurance.” *Journal of Political Economy*, 103(2), 360-399.
- [39] Hubener, Andreas, Raimond Maurer, and Olivia S. Mitchell. 2016. “How Family Status and Social Security Claiming Options Shape Optimal Life Cycle Portfolios.” *Review of Financial Studies*, 29(4), 937-978.
- [40] Kaygusuz, Remzi. 2015. “Social Security and Two-Earner Households.” *Journal of Economic Dynamics and Control*, 59, 163-178
- [41] Liebman, Jeffrey B., Erzo F.P. Luttmer, and David G. Seif. 2009. “Labor Supply Responses to Marginal Social Security Benefits: Evidence from Discontinuity.” *Journal of Public Economics*, 93(11-12), 1208-1223.
- [42] Low, Hamish, Costas Meghir, Luigi Pistaferri, and Alessandra Voena. 2018. “Marriage, Labor Supply and the Social Safety Net.” Mimeo.
- [43] Moffitt, Robert. 1998. *Welfare, The Family, And Reproductive Behavior, Research Perspectives*, ed. by Robert A. Moffitt. National Academies Press. Washington D.C.
- [44] Moffitt, Robert, and Sisi Zhang. 2018. “Income Volatility and the PSID: Past Research and New Results.” *AEA Papers and Proceedings*, 108, 277-280.
- [45] National Research Council, 1995, *Measuring poverty: a new approach*, Washington, DC: The National Academies Press.
- [46] Nishiyama, Shinichi. 2017. “The Joint Labor Supply Decision of Married Couples and the Social Security Pension System.” *Review of Economic Dynamics*, 31, 277-304.

- [47] Pakes, Ariel, and David Pollard. 1989. "Simulation and the Asymptotics of Optimization Estimators." *Econometrica*, 57(5), 1027–1057.
- [48] Persson, Petra. 2017. "Social Insurance and the Marriage Market." Mimeo.
- [49] Roys, Nicolas and Christopher Taber. 2017. "Skills Prices, Occupations and Changes in the Wage Structure for Low Skilled Men." Mimeo.
- [50] Sorrentino, Constance. 1990. "The Changing Family in International Perspective." *Monthly Labor Review*, 113(3), 41-58.
- [51] United Nations, Department of Economic and Social Affairs Population Division, Fertility and Family Planning Section. 2012. "World Marriage Data," <http://www.un.org/esa/population/publications/WMD2012/MainFrame.html>.

Appendix A. Data: The PSID and the HRS

We use the Panel Study of Income Dynamics (PSID) to estimate the wage process, the marriage and divorce probabilities, the initial distribution of couples and singles over state variables, taxes, and the sample moments that we match using our structural model.

The PSID is a longitudinal study of a representative sample of the U.S. population. The original 1968 PSID sample was drawn from a nationally representative sample of 2,930 families designed by the Survey Research Center at the University of Michigan (referred to here as the SRC sample), and from an over-sample of 1,872 low-income families from the Survey of Economic Opportunity (referred to here as the SEO sample). Individuals have been followed over time to maintain a representative sample of families.

We study the two cohorts born in the periods 1941-1945 and 1951-1955. More specifically, we select all individuals in the SRC sample who are interviewed at least twice in the sample years 1968-2013, select only heads and their wives, if present, and keep individuals born between 1931 and 1955. The resulting sample includes 5,129 individuals ages 20 to 70, for a total of 103,420 observations. In general, to gather the information we need, we control for birth cohort effects in our estimates (we use 5-year-of-birth windows) and use the results relative to the cohorts of interest. Table 10 details our PSID sample selection.

Selection	Individuals	Observations
Initial sample (observed at least twice)	30,587	893,420
Heads and wives (if present)	18,304	247,203
Born between 1931 and 1955	5,153	105,985
Age between 20 and 70	5,129	103,420

Table 10: Sample selection in the PSID.

Table 11 shows that the majority of men and women are married and that the fraction of married people goes down only slightly across these cohorts.

We use the Health and Retirement Study (HRS) to compute inputs for the retirement period because this data set contains a large number of observations and high-quality data for this stage of the life cycle. In fact, the HRS is a longitudinal data set collecting information on people ages 50 and older, including a wide range

	Born in 1941-1945			Born in 1951-1955		
Gender	Age 25	Age 40	Age 55	Age 25	Age 40	Age 55
Men	0.87	0.90	0.88	0.82	0.86	0.84
Women	0.86	0.84	0.79	0.83	0.81	0.76

Table 11: Fraction of married men and women by age and cohort, PSID data.

of demographic, economic, and social characteristics, as well as physical and mental health, and cognitive functioning.

The HRS started collecting information in 1992 on individuals born between 1931 and 1941, the so-called initial HRS cohort, which was then reinterviewed every two years. Other cohorts were introduced over the years. Our data set is based on the RAND HRS files for the period 1995-2012, to which we add the EXIT files to include information on the wave right after death. Our sample selection is as follows. Of the 37,317 individuals initially present, we drop individuals for whom marital status is not observed (1,548 individuals). This yields 35,769 individuals and 185,255 observations. We then select individuals in the age range 66-100 born in 1900 to 1945, obtaining a sample of 16,118 individuals and 81,246 observations. As we cannot observe individuals born after 1945 and older than age 66 in the HRS, for the 1955 cohort we use the same estimates obtained for the 1945 one.

Appendix B. First step estimation, methodology

Table 12 summarizes our estimated model inputs.

Wages

Because we allow our initial conditions, assortative matching in marriage, and marriage and divorce probabilities to depend on the realized values of wage shocks, we need to estimate not only wages as a function of human capital, age, and gender, and the stochastic process for the wage shocks, but also the realized wage shocks for all men and women of working age in our sample (whether they are working or not).

To do so, we proceed as follows. First, we impute potential wages for individuals who are not working, so that we are able to construct potential wages as actual wages for participants and potential wages for non-participants. Second, we estimate potential wages as a function of age, gender, and human capital. Third, we estimate

Estimated processes		Source
Wages		
$e_t^{i,j}(\cdot)$	Endogenous age-efficiency profiles	PSID
ϵ_t^i	Wage shocks	PSID
Demographics		
$s_t^{i,j}(\psi_t^i)$	Survival probability	HRS
$\zeta_t(\cdot)$	Divorce probability	PSID
$\nu_t(\cdot)$	Probability of getting married	PSID
$\xi_t(\cdot)$	Matching probability	PSID
$\theta_t(\cdot)$	Partner's assets and earnings	PSID
$f^{0,5}(i, j, t)$	Number of children ages 0-5	PSID
$f^{6,11}(i, j, t)$	Number of children ages 6-11	PSID
Health shock		
$m_t^{i,j}(\psi_t^i)$	Medical expenses	HRS
$\pi_t^{i,j}(\psi_t^i)$	Transition matrix for health status	HRS
Government policy		
λ_t^j, τ_t^j	Income tax	See text

Table 12: First-step estimated inputs summary.

the persistence and variance of its unobserved component and the realized wage shocks using Kalman filtering.

Missing wages imputation. The observed wage rate is computed as annual earnings divided by annual hours worked. Gross annual earnings are defined as labor income during the previous year. Annual hours are given by annual hours spent working for pay during the previous year.¹⁷

We impute missing wages by using coefficients from fixed effects regressions that we run separately for men and women. To avoid endpoint problems with the polynomials in age, we include individuals ages 22 to 70 in the sample. Define the observed wage for labor market participants as $\ln wage_{kt} = I_{kt}^n \widetilde{\ln wage}_{kt}$, where k denotes an individual and t is age. The term I_{kt}^n is an indicator for participation (which is equal to 1 if the individual participates in the labor market and has no missing hours or earnings) and $\widetilde{\ln wage}$ is the potential wage that we wish to estimate and do not observe. We estimate the unobserved potential wage by running the following regression

¹⁷Wages may be missing both because an individual has not been active in the labor market and because (s)he may have been active, but earnings or hours (or both) are missing. In addition, because estimated variances are very sensitive to outliers, we set to missing observations with an hourly wage rate below half the minimum wage and above \$368 (in 2016 values). We use the same imputation procedure for all these cases.

on observables as $\ln wage_{kt} = Z'_{kt}\beta_z + f_k + \varsigma_{kt}$, where the dependent variable is the logarithm of the observed hourly wage rate, f_k is an individual-specific fixed effect, and ς_{kt} is an error term. We include a rich set of explanatory variables in Z_{kt} : a fifth-order polynomial in age, a third-order polynomial in experience (measured in years of labor market participation), marital status (a dummy for being single), family size (dummies for each value), number of children (dummies for each value), age of youngest child, and an indicator of partner working if married. As an indicator of health, we use a variable recording whether bad health limits the capacity of working (this is the only health indicator available in the PSID for all years). Because this health indicator is not collected for wives, we do not include it in the regression for married women. Both regressions also include interaction terms between the explanatory variables. Variables that do not vary over time are captured by the individual effect f_k .

Using the estimated coefficients, we take the predicted value of the wage to be the potential wage for observations with missing wages. Hence, we define potential wage as

$$\ln \overline{wage}_{kt} = \begin{cases} \ln wage_{kt} & \text{if } I_{kt}^n = 1 \\ Z'_{kt}\hat{\beta}_z + \hat{f}_k & \text{if } I_{kt}^n = 0 \end{cases}$$

Wage function estimation. We model wages as a function of human capital, age, and gender, and we measure human capital as average realized earnings accrued up to the beginning of age t (\bar{y}_t).

To estimate the wage profiles, we proceed in two stages. First, we run the following fixed effect regression for the logarithm of potential wages

$$\ln \overline{wage}_{kt} = d_k + f^i(t) + \sum_{g=1}^G \beta_g D_g \ln(\bar{y}_{kt} + \delta_y) + u_{kt}, \quad (41)$$

on a gender-specific fifth-order polynomial in age $f^i(t)$, gender-cohort cells g , and gender-cohort dummies D_g .¹⁸ The shifter δ_y is set equal to \$5,000 to avoid taking the logarithm of values that are too small.¹⁹ We also experimented by adding marital

¹⁸Instead of following our general methodology of defining 5-year-of-birth cohorts, to estimate the cohort-specific effect of human capital on wages in equation (41), we take two broader windows: the 1940s cohort includes the generations born in 1931-1945, while the 1950s cohort includes those born in 1946-1955. We do so because we do not observe the complete age profile for the wages of the 1955 cohort.

¹⁹While we use earnings subject to the Social Security cap to compute average earnings (the state

status dummies to capture the effect of changing marital status on wages, but they did not turn out to be statistically different from zero, conditional on average earnings. Second, to fix the constant of the wage profile for our cohorts of interest, we regress the sum of the residuals and fixed effects $d_k + u_{kt+1} \equiv w_{kt+1}$ on cohort dummies to compute the average effects for the cohorts born in 1941-1945 and in 1951-1955, respectively. Table 13 reports the coefficients of the estimated equation from the first stage, the fixed effect regression, while Table 14 reports those from the second stage.

	Coefficient	Standard Error
$\ln(\bar{y}_t + \delta_y)$	0.307***	(0.0216)
$\ln(\bar{y}_t + \delta_y)$ *female	0.0419	(0.0277)
$\ln(\bar{y}_t + \delta_y)$ *born in 1950s	0.118***	(0.0265)
$\ln(\bar{y}_t + \delta_y)$ *born in 1950s*female	-0.0398	(0.0334)
Age	-0.567***	(0.177)
$Age^2/(10^2)$	2.679***	(0.861)
$Age^3/(10^4)$	-6.135***	(2.029)
$Age^4/(10^6)$	6.908***	(2.321)
$Age^5/(10^8)$	-3.095***	(1.033)
Age*female	0.343	(0.220)
$Age^2/(10^2)$ *female	-1.695	(1.070)
$Age^3/(10^4)$ *female	3.955	(2.526)
$Age^4/(10^6)$ *female	-4.433	(2.895)
$Age^5/(10^8)$ *female	1.947	(1.291)
Constant	2.044**	(0.851)
N	93363	
R-sq	0.103	

Table 13: Coefficients from fixed effects estimates. Dependent variable: logarithm of the potential wage. PSID data. Robust standard errors in parentheses, clustered at the individual level. * p<0.10, ** p<0.05, *** p<0.01

The estimated potential wage profiles, computed at average values of $\ln(\bar{y}_t)$, are shown in the main text. The shock in log wage is modeled as the sum of a persistent component plus white noise, which we assume captures measurement error:

$$\tilde{w}_{kt+1} = \ln \epsilon_{kt+1} + \xi_{kt+1} \quad (42)$$

$$\ln \epsilon_{kt+1} = \rho_\epsilon \ln \epsilon_{kt} + v_{kt+1}, \quad (43)$$

where \tilde{w}_{kt+1} are the residuals from the second stage, and ξ_{kt+1} and v_{kt+1} are independent white-noise processes with zero mean and variances equal to σ_ξ^2 and σ_v^2 ,

variable in our model), estimating this wage regression by using uncapped previous average earnings yields very similar estimates.

	Men	Women
Born in 1931-35	0.0178 (0.0406)	-0.0537* (0.0318)
Born in 1936-40	-0.00663 (0.0385)	-0.0537* (0.0319)
Born in 1946-50	-1.265*** (0.0277)	-0.776*** (0.0239)
Born in 1951-55	-1.314*** (0.0273)	-0.795*** (0.0227)
Constant	2.227*** (0.0231)	-0.953*** (0.0192)
N	45366	47996
R-sq	0.584	0.382

Table 14: Second stage: coefficients from OLS estimates. Dependent variable: residuals from fixed effects estimates. PSID data. Robust standard errors in parentheses, clustered at the individual level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

respectively. We estimate these process separately for each gender.²⁰

To estimate the realized value of the wage shocks, we estimate the system composed by equations (42) and (43) by maximum likelihood, which can be constructed assuming that the initial state of the system and the shocks are Gaussian, and using standard Kalman filter recursions. With that, we can estimate both the parameters in (42) and (43) and the entire state, that is, $\ln \epsilon_{kt}, t = 1, \dots, T$.

Table 15 reports our estimates for the AR component of earnings.

Parameter	Men	Women
Persistence	0.947	0.945
Variance prod. shock	0.026	0.016
Initial variance	0.112	0.098

Table 15: Estimated processes for the wage shocks for men and women, PSID data.

Average realized earnings and accumulated Social Security contributions

In the model we keep track of average accumulated earnings for a person (\bar{y}_{kt}) subject to Social Security cap that is applied to yearly earnings and is time varying. To

²⁰For this, we limit the age range between 25 and 65 and, because we rely on residuals also taken from imputed wages, we drop the highest 0.5% residuals for both men and women. This avoids large outliers to inflate the estimated variances (however, the effect of this drop on our estimates is negligible).

do so, we assume that individuals start working at age 22, and we compute individual-level capped average earnings. This computation requires taking a stand on people who appear in our data after age 22. Some individuals (5%) enter the sample after turning 22 either because in 1968, the first year the PSID was collected, they were older or because they entered as spouses or descendants and might thus be older than 22. Among people in this group, 46 enter the sample before turning 27: for those individuals we assume average accumulated earnings at entry are equal to zero. For the remaining 189 individuals, we use an imputation procedure to recover average realized earnings at entry and then update the value following each individual over time. We run a regression of capped earnings on a fourth-order polynomial in age fully interacted with gender, education dummies, interactions of education and gender, marital status, and race dummies also interacted with gender. Cohort dummies are also included. We use the predicted values of this regression as the entry value for individuals entering the sample after turning 27. Average earnings are then updated for each individual following his/her observed earnings history (as done in the model).

For the purposes of imputing missing values of wealth, we also compute uncapped average realized earnings using the same methodology for missing values of accumulated earnings at entry as above.

Wealth

We define wealth as total assets (defined as all asset types available in the PSID) plus home equity. Wealth in the PSID is only recorded in 1984, 1989, 1994, and then in each (biennial) wave from 1999 onward. We rely on an imputation procedure to compute wealth in the missing years, starting in 1968. This imputation is based on the following fixed effect regression:

$$\ln(a_{kt} + \delta_a) = Z'_{kt}\beta_z + da_k + wa_{kt}, \quad (44)$$

where k denotes the individual and t is age. The parameter δ_a is a shifter for assets to have only positive values and to be able to take logs, and the variables Z includes polynomials in age, also interacted with health status, and with average earnings (uncapped), family size, and a dummy for health status. The term da_k is the individual fixed effect and wa_{kt} is a white-noise error term. Equation (44) is estimated separately for single men, single women, and couples, as wealth is measured at the

household level, on an enlarged sample of individuals born between 1931 and 1965.

We then use the imputed as well as the actual observations to estimate the wealth profiles used as target moments and to parameterize the joint distribution of initial assets, average realized earnings, and wage shocks for single men, single women, and couples.

Distributions upon entering the model and for prospective spouses

For single men and women, separately, we parameterize the joint distribution of initial assets, average realized earnings, and wage shocks at each age as a joint lognormal distribution:

$$\begin{pmatrix} \ln(a_t^i + \delta_a^i) \\ \ln(\bar{y}_t^i) \\ \ln \epsilon_t^i \end{pmatrix} \sim N \begin{pmatrix} \mu_{at}^i + \delta_a^i \\ \mu_{\bar{y}t}^i \\ \mu_{\epsilon t}^i \end{pmatrix}, \quad \Sigma_{st}^i, \quad (45)$$

where Σ_s is a 3x3 covariance matrix. We estimate its mean and variance as a function of age t . For the mean, we regress the logarithm of assets plus shift parameter, average earnings, and productivity shock $\ln \hat{\epsilon}_t^i$ on a third-order polynomial in age and cohort dummies. The predicted age profile, relative to cohorts born in 1945 and 1955, is the age-specific estimate of the mean of the lognormal distribution. Taking residuals from the above estimates, we can estimate the elements of the variance-covariance matrix by computing the relevant squares or cross-products. We regress the squares or the cross-products of the residuals on a third-order polynomial in age to obtain, element by element, a smooth estimate of the variance-covariance matrix at each age.

For couples, we compute the initial joint distribution at age 25 of the following variables:

$$\begin{pmatrix} \ln(a + \delta_a) \\ \ln(\bar{y}^1) \\ \ln(\bar{y}^2) \\ \ln(\epsilon^1) \\ \ln(\epsilon^2) \end{pmatrix} \sim N \begin{pmatrix} \mu_a + \delta_a \\ \mu_{\bar{y}1} \\ \mu_{\bar{y}2} \\ \mu_{\epsilon 1} \\ \mu_{\epsilon 2} \end{pmatrix}, \quad \Sigma_c, \quad (46)$$

where Σ_c is a 5x5 covariance matrix computed on the data for married or cohabiting

couples.

Marriage and divorce probabilities

We model the probability of getting married, ν_{t+1} , as a function of gender, age and the wage shock and perform the estimation separately for men and women using PSID data. Our estimated equation is

$$\nu_{t+1}^i = \text{Prob}(\text{Married}_{t+1} = 1 | \text{Married}_t = 0, Z_t) = F(Z_t' \beta_m),$$

where F denotes the standard logistic distribution and Z_t includes a polynomial in age, cohort dummies, the logarithm of the wage shock, and the after-1997 dummy²¹. Using the estimated coefficients on the cohort dummies, we then adjust the probability for the 1945 and 1955 cohorts, respectively.

Similarly, we estimate the probability of divorce as

$$\zeta_t = \text{Prob}(\text{Divorced}_{t+1} = 1 | \text{Married}_t = 1, Z_t) = F(Z_t' \beta_d),$$

where F denotes the standard logistic distribution and Z_t includes a polynomial in age, husband's wage shock, wife's wage shock, cohort dummies, and an indicator for biennial waves.




Figure 14: Marriage probabilities by gender, age, and one's wage shock for the 1945 cohort (left panel) and 1955 cohort (right panel), PSID data.

²¹The PSID goes from a yearly to a biennial frequency in 1997. To take this into account, we include an indicator variable taking a value of one from 1997 on in the regression, which we then abstract from when constructing the yearly probabilities.

	Single Men Marriage	Single Women Marriage	Couples Divorce
<i>Age</i>	0.0386 (0.0310)	-0.0424 (0.0324)	0.0469 (0.0294)
<i>Age</i> ² /10 ²	-0.109*** (0.0404)	-0.0324 (0.0415)	-0.114*** (0.0368)
$\ln \epsilon_{kt}$	0.308** (0.131)	0.0578 (0.143)	-0.399*** (0.101)
Spouse's $\ln \epsilon_{kt}$			-0.116 (0.120)
<i>I</i> (<i>year</i> > 1997)	0.588*** (0.208)	0.378* (0.208)	0.505*** (0.167)
Born in 1931-35	0.144 (0.269)	-0.329 (0.229)	-0.259 (0.173)
Born in 1936-40	-0.0219 (0.238)	-0.586** (0.238)	-0.115 (0.141)
Born in 1946-50	-0.198 (0.181)	-0.292* (0.160)	-0.122 (0.111)
Born in 1951-55	-0.364** (0.174)	-0.352** (0.156)	-0.0762 (0.118)
Constant	-1.545** (0.606)	-0.135 (0.634)	-3.537*** (0.574)
N	5064	8860	32071
pseudo- <i>R</i> ²	0.026	0.060	0.023

Table 16: Estimated coefficients from logistic regressions. Column 1: Marriage of single men; column 2: marriage of single women; column 3: divorce of couples. PSID data. Robust standard errors in parentheses, clustered at the individual level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$




Figure 15: Divorce probabilities by gender, age, and one's wage shock for the 1945 cohort (left panel) and 1955 cohort (right panel), PSID data.

Figures 14 and 15 report the resulting marriage and divorce probabilities for both cohorts.

Conditional on meeting a partner, the probability of meeting a partner p with wage shock ϵ_{t+1}^p is $\xi_{t+1}(\cdot) = \xi_{t+1}(\epsilon_{t+1}^p | \epsilon_{t+1}^i, i)$. Using our estimated wage shocks and partitioning households in age groups (25-35, 35-45, 45-65), we compute the variance-covariance matrix of newly matched partners' wage shocks by age groups. We then

derive the conditional distribution of meeting a partner assuming lognormality. In the whole sample we observe only 750 new marriages in the age range 25-65, therefore we do not allow this probability to depend on cohort.

Number of children

To compute the average number of children by age group, we use the individual information in the PSID and classify as children of the family in the following categories: sons or daughters of the head, stepsons or stepdaughters of the head, sons or daughters of the cohabiting partner but not of the head, foster sons or foster daughters (not legally adopted), and children of the first-year cohabitor but not of the head. Having done that, we add up the number of children in each age category (0 to 5, 6 to 11, or 0 to 17 for the total number of children) and run a regression on a fifth-order polynomial in age of the mother, interacted with marital status, and cohort dummies to construct the average age profile of children in each age group for single and married women. We use the profiles for the cohorts of mothers born in 1941-1945 and in 1951-1955.

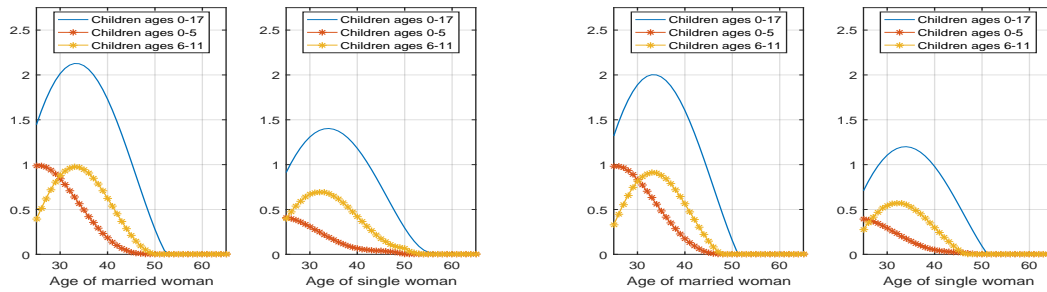


Figure 16: Number of children for married and single women for the 1945 cohort (left panel) and 1955 cohort (right panel), PSID data.

Health status at retirement

We define health status on the basis of self-reported health. In the HRS, this variable can take five possible values (excellent, very good, good, fair, poor). As standard, we take health to be a dichotomous variable equal to 1 if self-reported

health is fair or poor and 0 otherwise.²² We estimate the probability of being in bad health at age 66, using the observed frequencies for the 1941-1945 cohort, which is the youngest cohort that we can observe at age 66+ in the HRS data. All the inputs estimated from the HRS correspond to the 1941-1945 cohort. For lack of better data, we also use them for our 1951-1955 cohort. For singles, we compute the sample fraction of single men and single women in bad health in the age range 65-67, which ensures that the sample size is big enough. For couples, we define the first member in the couple as the husband and the second as the wife, and compute the sample frequencies for the four possible health states in the couple as (good, good), (good, bad), (bad, good), and (bad, bad).

Health dynamics after retirement

As before, we use the HRS data, and we define the health status variable ψ equal to 1 if self-reported health at time t is equal to “fair” or “bad” and 0 otherwise. We model the probability of being in bad health during retirement as a logit function:

$$\pi_{\psi t} = Prob(\psi_t = 1 \mid X_t^\psi) = \frac{\exp(X_t^{\psi'} \beta^\psi)}{1 + \exp(X_t^{\psi'} \beta^\psi)},$$

which we then use to construct the transition matrix at each age, gender, and marital status. The set of explanatory variables X_t^ψ includes cohort dummies, a second-order polynomial in age, previous health status, gender, marital status, and interactions between these variables when they are statistically different from zero. As the HRS data are collected every two years, we obtain two-year probabilities and convert them into one-year probabilities. Table 17 reports our estimated coefficients, while Figure 17 displays the health transition matrix by gender, age, marital status, and health status that we estimated.

²²Looking at labor supply behavior around retirement time, Blundell et al. (2017) show that this measure of self-reported health captures health well and about as well as more involved measures such as using large numbers of objective measures to predict health.




Figure 17: Health transition probabilities for singles and couples by age. HRS data.

	Coefficient	SE
Age	0.0934***	(0.0305)
$Age^2/10^2$	-0.0463**	(0.0190)
$Health_{t-1}$	6.754***	(0.265)
$Health_{t-1} * Age$	-0.0547***	(0.00335)
Male	0.476**	(0.242)
Male*Age	-0.00551*	(0.00313)
Age*Married	-0.0262***	(0.00329)
$Age^2/10^2 * Married$	0.0307***	(0.00420)
Born in 1936-40	0.166***	(0.0610)
Born in 1931-35	0.221***	(0.0609)
Born in 1926-30	0.349***	(0.0698)
Born in 1921-25	0.392***	(0.0692)
Born in 1916-20	0.519***	(0.0720)
Born in 1900-15	0.677***	(0.0788)
Constant	-6.200***	(1.215)
N	58547	
Pseudo- R^2	0.236	

Table 17: Health dynamics over two-year periods. Logistic regression coefficients, dependent variable: health status. HRS data. Robust standard errors in parentheses, clustered at the individual level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Survival probabilities

We model the probability of being alive at time t as a logit function:

$$s_t = Prob(Alive_t = 1 | X_t^s) = \frac{\exp(X_t^{s'} \beta^s)}{1 + \exp(X_t^{s'} \beta^s)}$$

which we estimate using the HRS data. Among the explanatory variables, we include a fourth-order polynomial in age, gender, marital status, and health status in the previous wave, as well as interactions between these variables and age, whenever they are statistically different from zero. As the HRS is collected every two years, we transform the biennial probability of surviving into an annual probability by taking

the square root of the biennial probability. Table 18 reports estimated coefficients, and Figure 18 displays the implied survival probability by age, gender, and marital and health status.

	Coefficient	SE
Age	-19.03***	(6.826)
$Age^2/10^2$	34.66***	(12.53)
$Age^3/10^4$	-27.96***	(10.17)
$Age^4/10^6$	8.377***	(3.081)
$Health_{t-1}$	-3.816***	(0.307)
$Health_{t-1} * age$	0.0313***	(0.00370)
Male	-1.213***	(0.332)
Male*Age	0.00836**	(0.00405)
Married	1.302***	(0.375)
Married*Age	-0.0128***	(0.00465)
Born in 1936-40	0.161	(0.129)
Born in 1931-35	0.0817	(0.128)
Born in 1926-30	-0.00885	(0.138)
Born in 1921-25	0.0434	(0.139)
Born in 1916-20	0.0363	(0.142)
Born in 1900-15	0.0644	(0.145)
Constant	395.6***	(138.8)
N	63746	
Pseudo- R^2	0.171	

Table 18: Logistic regression coefficients, dependent variable: survival over a two-year period. HRS data. Robust standard errors in parentheses, clustered at the individual level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$




Figure 18: Survival probability by age, gender, and marital and health status, both cohorts. HRS data.

Out-of-pocket medical expenditures

Out-of-pocket medical expenses are defined as the total amount that the individual spends out of pocket in hospital and nursing home stays, doctor visits, dental costs,

outpatient surgery, average monthly prescription drug costs, home health care, and special facilities charges. They also include medical expenses in the last year of life, as recorded in the exit interviews. In contrast, expenses covered by public or private insurance are not included in our measure, as they are not directly incurred by the individual. The estimated equation is

$$\ln(m_{kt}) = X_{kt}^{m'} \beta^m + \alpha_k^m + u_{kt}^m,$$

where explanatory variables include a fourth-order polynomial in age fully interacted with gender and current health status, and we include these interactions whenever they are statistically different from zero. We estimate the equation on the HRS data using a fixed effects estimator, which takes into account all unmeasured fixed-over-time characteristics that may bias the age profile, such as differential mortality (as discussed in De Nardi, French, and Jones (2010)). Marital status (also interacted with other variables) does not turn out to be significantly different from zero in the first step. We then regress the residuals and fixed effects from this equation on cohort, gender, and marital status dummies to compute the average effect for each group of interest. Table 19 reports estimated coefficients, while Figure 19 displays medical expenditure by age, gender, and marital and health status.

Finally, we model the variance of the shocks by regressing the squared residuals from the regression in logs on a third-order polynomial in age fully interacted with gender and current health status, and on cohort, gender, and marital status dummies and use it to construct average medical expenses as a function of age by adding half of the variance to the average in logs before exponentiating.




Figure 19: Medical expenditure by age, gender, and marital and health status. HRS data.

	Coefficient	SE
Age	9.770***	(2.416)
Age ² /10 ²	-18.63***	(4.612)
Age ³ /10 ⁴	15.68***	(3.893)
Age ⁴ /10 ⁶	-4.901***	(1.226)
Bad health	3.819***	(1.012)
Bad health*Age	-0.0961***	(0.0263)
Bad health*Age ² /10 ²	0.0624***	(0.0169)
Male*Age	-9.160***	(3.793)
Male*Age ² /10 ²	17.76***	(7.261)
Male*Age ³ /10 ⁴	-15.14***	(6.147)
Male*Age ⁴ /10 ⁶	4.792***	(1.942)
Constant	-109.9***	(36.32)
Second stage		
Male	174.9***	(0.0263)
Married	0.330***	(0.0195)
Male*Married	-0.0469**	(0.0316)
Born in 1936-40	-0.000573	(0.0302)
Born in 1931-35	-0.0534**	(0.0296)
Born in 1926-30	-0.118***	(0.0358)
Born in 1921-25	-0.0954***	(0.0338)
Born in 1916-20	-0.143***	(0.0344)
Born in 1900-15	-0.309***	(0.0378)
Constant	-73.77***	(0.0287)
N	65917	
R ² first stage	0.017	
R ² second stage	0.99	

Table 19: Estimates for the logarithm of medical expenses, first stage (fixed effects) and second stage (OLS). HRS data. Robust standard errors in parentheses, clustered at the individual level. * p<0.1, ** p<0.05, *** p<0.01

Spousal assets and Social Security benefits

We assume random matching over asset and lifetime income of the partner conditional on partner's wage shock. Thus, we compute $\theta_{t+1}(\cdot) = \theta_{t+1}(a_{t+1}^p, \bar{y}_{t+1}^p | \epsilon_{t+1}^p)$ using sample values of assets, average capped earnings, and wage shocks. More specifically, we assume θ_{t+1} is lognormally distributed at each age with mean and variance computed from sample values. Assets include a shifter as described for the computation of the joint distribution at age 25 (see Wealth subsection in this Appendix).

Figure 20 reports spousal assets by spousal wage shocks in case of marriage next period. Both panels show that both women and men marrying early on in life expect their partner to have relatively low assets on average, even conditional on the various wage shocks. In contrast, those who marry later experience a much larger variation in their partner's assets conditional on their partner's wage shocks. The gradient in average assets by wage shocks increases especially fast for male partners and thus exposes women to much more variability in their partner's resources as they marry




Figure 20: Spousal assets by spousal wage shocks in case of marriage next period for the 1945 cohort (left panel) and 1955 cohort (right panel), PSID data.

later and later. The patterns are very close for the two cohorts.

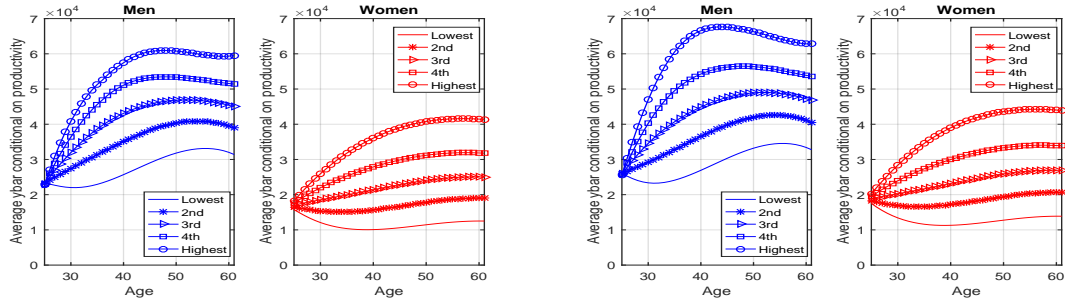


Figure 21: Spousal Social Security earnings by spousal wage shocks in case of marriage next period for the 1945 cohort (left panel) and 1955 cohort (right panel), PSID data.

Figure 21 reports spousal Social Security earnings by spousal wage shocks in case of marriage next period. Given that male wage shocks are higher on average, Social Security earnings for men are higher than those for women at all levels of the wage shocks.

Taxes

We model taxes T on total income Y as $T(Y) = Y - \lambda Y^{1-\tau}$, where τ captures the degree of progressivity and λ captures the average level of taxation in the system. Since this specification implies $(Y - T(Y)) = \lambda Y^{1-\tau}$ and $\ln(Y - T(Y)) = \ln(\lambda) + (1 - \tau)\ln(Y)$, we estimate τ and λ by regressing the logarithm of after-tax household income on a constant and on the logarithm of pre-tax household income by cohort,

year, and household type (single man, single woman, couple).

We use PSID data from 1968 to 2015 (tax years 1967-2014) to estimate cohort- and time-specific tax functions. Information about federal taxes paid is provided directly by the PSID up to 1991. After that year it is gathered using TAXSIM, the NBER simulation program computing taxes. In particular, we build on and extend the program written by Kimberlin et al. (2015) to prepare the input needed by TAXSIM.²³

Before-tax household income is defined as the sum of all money income received by the spouses (or by the individual if single) in a given tax year. It therefore includes income of the head and the wife (if present), that is, labor income, the asset part of income from farm, business, roomers, and so on, plus income from rent, interest dividends, and so on, and wife's income from assets, plus transfer income, that is, Social Security, pension, annuities, other retirement income, welfare, aid to dependent children, unemployment or worker's compensation, help from relatives, alimony, or child support. After-tax household income is defined as before-tax income minus the federal income tax liability (including capital gains rates, surtaxes, AMT, and refundable and non-refundable credits, as computed by TAXSIM).

To keep the number of observations large while keeping different tax regimes separate, we follow a slightly different procedure for couples and singles. For couples, we define two five-year cohorts (one born in 1941-1945, one in 1951-1955) and estimate the tax functions over two- or three-year intervals. For single men and women, the 1945 cohort includes individuals born in 1938-1947, while the 1955 one includes those born in 1948-1957. Then, we estimate yearly tax functions, using data relative to a moving five-year window for each function, to have enough observations and to capture relevant changes in the legislation.

All the inputs needed by TAXSIM are gathered directly from the PSID, for the sample years 1992-2015. However, for years prior to 1999, medical expenses and charitable contributions are not available and need to be imputed, as they may be deducted from gross income (if the household chooses to itemize). Hence, we impute

²³The program by Kimberlin et al. (2015) prepares the input for TAXSIM for the PSID years 1999-2011 following Butrica and Burkhauser (1997). It differs from more simplified PSID TAXSIM interface approaches in that multiple tax units are identified within each PSID family unit; thus, cohabiting couples are treated as two separate tax units, with children assigned to the appropriate tax unit (head or cohabitor) using relationship codes. We extend their program to include all years between 1992 and 2015.

them by regressing the sum of the two items for pooled years 1999-2015 and predicting the value using the estimated parameters (out-of-sample prediction). The included explanatory variables are demographic and income variables, such as family size, employment status of the head and the spouse if present, state of residence, wages, pensions, other incomes, education, number of children, age, and marital status. Then, we add an error term to that prediction to tackle the attenuation in the variance of the distribution of the imputed values, following the procedure in David et al. (1986), and French and Jones (2011). More precisely, the procedure is as follows. First, we regress the sum of the two items on the vector of observables for the sample of heads who choose to itemize, $deduc_i = z_i\beta + \epsilon_i$. Second, for each household i for which $deduc$ is observed, we calculate the predicted value $\widehat{deduc}_i = z_i\hat{\beta}$ and the residual $\hat{\epsilon}_i = deduc_i - \widehat{deduc}_i$. Third, we sort the predicted value \widehat{deduc}_i into deciles and keep track of all values of $\hat{\epsilon}_i$ within each decile. Next, for every individual j with missing $deduc$ we impute $\widehat{deduc}_j = z_j\hat{\beta}$. Then we impute $\hat{\epsilon}_j$ for households with missing $deduc$ by finding a random individual i in the non-missing sample with a value of \widehat{deduc}_i in the same decile as \widehat{deduc}_j and set $\hat{\epsilon}_j = \hat{\epsilon}_i$. The imputed value of $deduc$ is $\widehat{deduc}_j + \hat{\epsilon}_j$.

Appendix C. Calibrated model parameters

Table 20 summarizes our first-step calibrated model inputs. We set the interest rate r to 4% and the utility curvature parameter, γ , to 2.5. The equivalence scales are set to $\eta_t^{i,j} = (j + 0.7 * f_t^{i,j})^{0.7}$, as estimated by National Research Council (1995). With the equivalence scale, \$1 spending in a household of a couple with two kids gives each household member a consumption of 0.42 cents. The term $f_t^{i,j}$ is the average total number of children for single and married men and women by age.

The most recent paper estimating the consumption floor during retirement is the one estimated by De Nardi, French, and Jones (2016) in a rich model of retirement savings with endogenous medical expenses. In their framework, they estimate a utility floor that corresponds to consuming \$4,600 a year when healthy. However, they note that SSI recipients are guaranteed a minimum income of \$6,670. As a compromise, we use \$5,900 as our consumption floor for elderly singles, which is \$8,687 in 2016 dollars, and the one for couples to be 1.5 the amount for singles, which is the statutory ratio between benefits of couples to singles. The retirement benefit at age 66 is calculated to mimic the Old Age and Survivor Insurance component of the Social Security system.

Calibrated parameters		Source
Preferences and returns		
r	Interest rate	4% De Nardi, French, and Jones (2016)
γ	Utility curvature parameter	2.5 see text
η_t	Equivalence scales	PSID
Government policy		
$SS(\bar{y}_r^i)$	Social Security benefit	See text
τ_t^{SS}	Social Security tax rate	See text
\tilde{y}_t	Social Security cap	See text
$\underline{c}(1)$	Minimum consumption, singles	\$8,687, De Nardi, French, and Jones (2016)
$\underline{c}(2)$	Minimum consumption, couples	\$8,687*1.5 Social Security rules

Table 20: First-step calibrated inputs summary

Social Security benefits

The Social Security benefit at age 66 is calculated to mimic the Old Age and Survivor Insurance component of the Social Security system:

$$SS(\bar{y}_r) = \begin{cases} 0.9\bar{y}_r, & \bar{y}_r < 0.1115; \\ 0.1004 + 0.32(\bar{y}_r - 0.1115), & 0.1115 \leq \bar{y}_r < 0.6725; \\ 0.2799 + 0.15(\bar{y}_r - 0.6725), & 0.6725 \leq \bar{y}_r < y_t^{cap} \end{cases}$$

The marginal rates and bend points, expressed as fractions of average household income, come from the Social Security Administration.²⁴ The Social Security tax and Social Security cap shown in Figure 22 have been changing over time. We also allow them to change over time for the households in our cohorts.




Figure 22: Social Security tax and Social Security cap over time (expressed in 2016 dollars)

²⁴Social Security Administration, “Benefit Formula Bend Points,” <https://www.ssa.gov/oact/cola/bendpoints.html>. We use their values for 2009.

Appendix D. The solution algorithm

This appendix describes the solution algorithm. We first solve the value functions and policy functions. Then we simulate our model economy using the inputs and estimate parameters following the procedures that we describe in the next section.

We optimize over six value functions over multiple time periods, compute three more value functions, and have six continuous state variables. In addition, there can be kinks in the value functions because both husbands and wives choose their participation. Thus, to have reliable solutions, we compute them brute force on a grid. To get a sense of dimensionality, the value function for working couples has the following dimensions in terms of state variables: age (41 periods, as we have yearly periods), assets, earnings shocks for each spouse, and human capital for each spouse. Over these grids, we evaluate choices for consumption, savings, and labor supply of both household members and compute all of the relevant expected values at each and marital status for each of the value functions.

Even parallelizing our model in C on high-end workstations, the model requires 37 minutes for each set of parameter values to be solved. Estimating the model for one cohort implies solving it thousands of times, which thus requires at least three or four weeks each time. We reestimate our model for each cohort many times to check for local minima, robustness, and so on. The computation time required is substantial.

During the retirement stage, single people do not get married anymore; hence, their value function can be computed independently of the other value functions. The value function of couples depends on their own future continuation value and the one of the singles, in case of death of a spouse. Then there is the value function of the single person being married in a couple, which depends on the optimal policy function of the couple, taking the appropriate expected values. We compute them as follows:

1. Compute the value function of the retired single person for all time periods after retirement by backward iteration starting from the last period.
2. Compute the value function of the retired couple for all time periods after retirement, which uses the value function for the retired single person in case of death of one of the spouses by backward induction starting from the last period.
3. Compute the value function of the single person in a marriage for all time periods after retirement.

During the early retirement stage, single people do not get married, and married individuals do not divorce or die; hence, the value function of the single person and that of the couple can be computed independently. We compute them as follows:

1. Compute the value function of the single person for all time periods by backward iteration starting from the last period in the early retirement stage.
2. Compute the value function of the couple for all time periods by backward iteration starting from the last period in the early retirement stage.
3. Compute the value function of the single person in a marriage for all time periods in the early retirement stage.

During the working age, the value functions are interconnected; hence, we solve each of them at time t , working backward over the life cycle, at each period:

1. Take as given the value of being a single person in a married couple for next period and the value function of being single next period, which have been previously computed, and compute the value function of being single this period.
2. Given the value function of being single, compute the value function of the couple for the same age.
3. Given the optimal policy function of the couple, use the implied policy functions to compute the value function for a person in a couple.
4. Keep going back in time until the first period.

Appendix E. Moment Conditions and Asymptotic Distribution of Parameter Estimates

In this appendix, we review the two-step estimation strategy, the moment conditions, and the asymptotic distribution of our estimation. To simplify notation, we do not include a separate indicator for each of the two cohorts.

In the first step, we estimate the vector χ , the set of parameters than can be estimated without explicitly using our model. In the second step, we use the method of simulated moments (MSM) to estimate the remaining parameters, which are contained in the $M \times 1$ vector Δ . For the 1945 cohort, the elements of Δ are the 19 model parameters $(\beta, \omega, (\phi_0^{i,j}, \phi_1^{i,j}, \phi_2^{i,j}), (\tau_c^{0,5}, \tau_c^{6,11}), L^{i,j})$.²⁵ For the 1955 cohort, we assume

²⁵We normalize the time endowment of single men.

that the households have the same β and ω as the 1945 cohort, and we thus estimate the remaining 17 parameters. Our estimate, $\hat{\Delta}$, of the “true” parameter vector Δ_0 is the value of Δ that minimizes the (weighted) distance between the lifecycle profiles found in the data and the simulated profiles generated by the model.

From ages 25 to 65, we match average assets for single men, single women, and couples, as well as working hours and participation for single men, single women, married men, and married women. For the generic variable z equal to hours (H), participation (In), and assets (a), we denote $z_{k,t}^{i,j}$ the sample observation relative to person k , of gender i , marital status j , and age t . Denoting $z_t^{i,j}(\Delta, \chi)$ the model-predicted expected value of z for age i , gender i , and marital status j , where χ is the vector of parameters estimated in the first step, we write the moment conditions as

$$E[a_{k,t}^{i,j} - a_t^{i,j}(\Delta_0, \chi_0)] = 0, \quad \forall t = 2, \dots, 41 \quad (47)$$

$$E[H_{k,t}^{i,j} - H_t^{i,j}(\Delta_0, \chi_0)] = 0, \quad \forall t = 1, \dots, 41 \quad (48)$$

$$E[In_{k,t}^{i,j} - In_t^{i,j}(\Delta_0, \chi_0)] = 0, \quad \forall t = 1, \dots, 41. \quad (49)$$

Note that assets for couples, $a_{k,t}^{i,j}$, do not depend on gender when marital status is $j = 2$. Also, as assets at age 25 ($t = 1$) is an initial condition, it is matched by construction. Thus, we have a total of $J = 448$ moment conditions. In practice, we compute the sample expectations in equations (47), (48), and (49) conditional on a flexible polynomial in age. More specifically, we regress each variable z on a fourth-order polynomial in age and on a set cohort of dummies, fully interacted with marital status and separately for each gender. We then compute the conditional expectations for each cohort in turn using the estimated marital- and gender-specific polynomial in age as well as coefficients relative to that cohort. These average age profiles, conditional on gender, marital status, and cohort, are those shown in the figures in the main text.

Suppose we have a data set of K persons that are each observed at up to T separate calendar years. Let $\varphi(\Delta; \chi_0)$ denote the J -element vector of moment conditions described immediately above, and let $\hat{\varphi}_K(\cdot)$ denote its sample analog.

Letting $\widehat{\mathbf{W}}_K$ denote a $J \times J$ positive definite weighting matrix, the MSM estimator $\hat{\Delta}$ is given by

$$\underset{\Delta}{\operatorname{argmin}} \hat{\varphi}_K(\Delta; \chi_0)' \widehat{\mathbf{W}}_K \hat{\varphi}_K(\Delta; \chi_0). \quad (50)$$

Note that we also estimate χ_0 . For tractability reasons, and following much of the literature, we treat it as known.

Under the regularity conditions stated in Pakes and Pollard (1989) and Duffie and Singleton (1993), the MSM estimator $\hat{\Delta}$ is both consistent and asymptotically normally distributed:

$$\sqrt{K} \left(\hat{\Delta} - \Delta_0 \right) \rightsquigarrow N(0, \mathbf{V}), \quad (51)$$

with the variance-covariance matrix \mathbf{V} given by

$$\mathbf{V} = (\mathbf{D}'\mathbf{W}\mathbf{D})^{-1}\mathbf{D}'\mathbf{W}\mathbf{S}\mathbf{W}\mathbf{D}(\mathbf{D}'\mathbf{W}\mathbf{D})^{-1}, \quad (52)$$

where \mathbf{S} is the variance-covariance matrix of the data;

$$\mathbf{D} = \left. \frac{\partial \varphi(\Delta; \chi_0)}{\partial \Delta'} \right|_{\Delta=\Delta_0} \quad (53)$$

is the $J \times M$ gradient matrix of the population moment vector; and $\mathbf{W} = \text{plim}_{K \rightarrow \infty} \{\widehat{\mathbf{W}}_K\}$. When $\mathbf{W} = \mathbf{S}^{-1}$, \mathbf{V} simplifies to $(\mathbf{D}'\mathbf{S}^{-1}\mathbf{D})^{-1}$.

The asymptotically efficient weighting matrix arises when $\widehat{\mathbf{W}}_K$ converges to \mathbf{S}^{-1} , the inverse of the variance-covariance matrix of the data. However, as Altonji and Segal (1996) point out, the optimal weighting matrix is likely to suffer from small sample bias. We thus use a diagonal weighting matrix that is the same as \mathbf{S} along the diagonal and has zeros off the diagonal of the matrix. We estimate \mathbf{D} and \mathbf{W} with their sample analogs.

Appendix F. Parameter Estimates

Figure 23 reports the age-varying time costs of working by age expressed as a fraction of the time endowment of single men that are necessary to reconcile the labor market participation of our four groups of people in each cohort.

Appendix G. Model fit, additional information

We do not match savings after age 66 because the asset data become very noisy after that. However, the model does fit them well. Figure 24 shows the full profile of assets generated by the model and those in the data for the 1945 cohort.

Figure 25 compares additional model implications with those in the data for couples. The top panels display participation patterns within married couples from the model and the PSID data. While we match the participation of married men and

	Cohort 1945		Cohort 1955	
	Coeff.	Std. Err.	Coeff.	Std. Err.
β	0.9898	(0.00025)		
ω	0.4057	(0.00121)		
<i>Participation costs:</i>				
$\phi_2^{1,1}$	0.0001	(0.00003)	0.0008	(0.00003)
$\phi_1^{1,1}$	-0.0044	(0.00054)	-0.0312	(0.00149)
$\phi_0^{1,1}$	-1.3769	(0.02518)	-0.9631	(0.02210)
$\phi_2^{2,1}$	0.0006	(0.00003)	-0.0001	(0.00001)
$\phi_1^{2,1}$	-0.0170	(0.00095)	0.0024	(0.00029)
$\phi_0^{2,1}$	-1.4778	(0.03245)	-0.8859	(0.01233)
$\phi_2^{1,2}$	0.0007	(0.00002)	0.0016	(0.00004)
$\phi_1^{1,2}$	-0.0187	(0.00087)	-0.0655	(0.00164)
$\phi_0^{1,2}$	-1.5621	(0.03095)	-1.4657	(0.02473)
$\phi_2^{2,2}$	0.0033	(0.00008)	0.0064	(0.00012)
$\phi_1^{2,2}$	-0.1345	(0.00394)	-0.2723	(0.00706)
$\phi_0^{2,2}$	-2.1433	(0.04556)	-1.8571	(0.04541)
<i>Time endowments:</i>				
$FL^{2,1}$	-3.0381	(0.08251)	-5.5347	(0.40068)
$FL^{1,2}$	-3.0197	(0.07796)	-2.1473	(0.03101)
$FL^{2,2}$	-1.2955	(0.01272)	-1.2654	(0.00768)
<i>Childcare costs:</i>				
$\tau_c^{0,5}$	0.3002	(0.01316)	0.2502	(0.01527)
$\tau_c^{6,11}$	0.0683	(0.00718)	0.1885	(0.01007)

Table 21: Estimates of parameters. Standard errors in parentheses. We estimate $FL^{i,j}$ and time endowment in the model is given by $L^{i,j} = \frac{L}{1+\exp(FL^{i,j})}$, where we normalize L to 112 hours a week.

women by estimation, we do not match the fraction of couples with no earners or with only women earners. The bottom panels show hours worked by husbands whose wife is not working, husbands whose wife is working, and wives whose husband is working. We report them by age over the whole working period. This, also, is not a target that our estimation procedure seeks to match. Both sets of graphs reveal that the model also reproduces these aspects of the data well. This is remarkable given that our model is tightly parameterized compared with the number of targets that it matches.

Figures 26 and 27 report our model-implied moments as well as target moments





Figure 23: Estimated life cycle labor participation costs expressed as a fraction of the time endowment of single men. SM: single men; SW: single women; MM: married men; MW: married women. Left panel: 1945 cohort. Right panel: 1955 cohort. Model estimates




(a) Assets, couples


(b) Assets, singles

Figure 24: 1945 cohort. Model fit for assets and average and 95% confidence intervals from the PSID data


and 95% confidence intervals from the PSID data for our 1955 cohort. They show that our parsimoniously parameterized model also fits the data for the 1955 cohort well.




(a) Participation, model



(b) Participation, data




(c) Hours for workers, model




(d) Hours for workers, data


Figure 25: 1945 cohort. Participation and worker's hours patterns for people in couples. Model and PSID data comparison




(a) Participation, couples



(b) Participation, singles




(c) Hours for workers, couples




(d) Hours for workers, singles

Figure 26: 1955 cohort. Model fit for participation (top graphs) and hours (bottom graphs) and average and 95% confidence intervals from the PSID data



(a) Assets, couples



(b) Assets, singles

Figure 27: 1955 cohort. Model fit for assets and average and 95% confidence intervals from the PSID data

Appendix H. Policy experiments results without balancing government budget for both cohorts

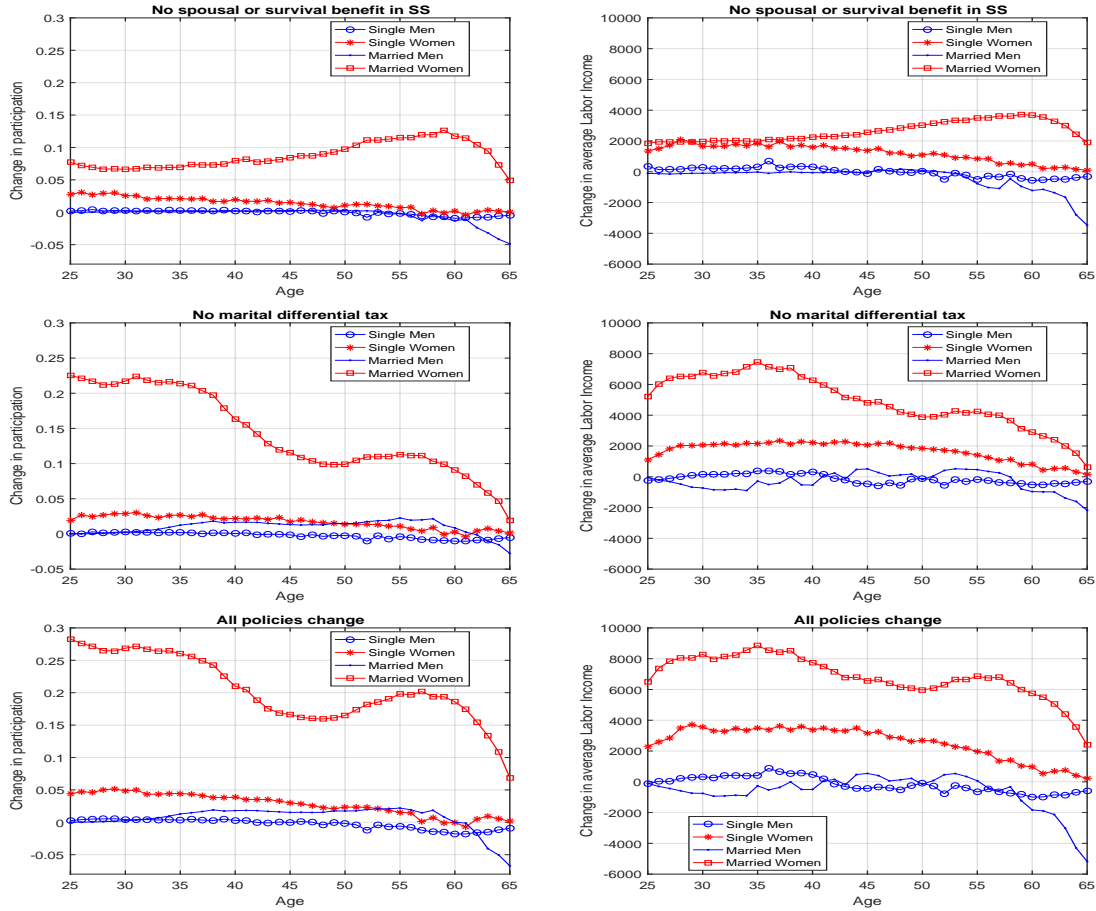


Figure 28: 1945 cohorts: Changes in participation (left panels) and labor income (right panels), unbalanced government budget. Top panels: after the elimination of all the spousal Social Security benefits; middle panels: after the elimination of joint income taxation; bottom panels: after the elimination of all marital-related policies.

	Couples	Single men	Single women
Removing spousal Social Security benefits	9.7%	1.7%	4.3%
Removing all marital-related policies	15.3%	3.3%	8.5%

Table 22: 1945 cohorts: Change in assets at age 66, in percentages, unbalanced government budget.




Figure 29: 1955 cohorts: Changes in participation (left panel) and labor income (right panel) after the elimination of all spousal Social Security benefits and joint income taxation. Unbalanced government budget

	Couples	Single men	Single women
Savings, unbalanced government budget	15.6%	3.8%	9.5%

Table 23: 1955 cohorts: Change in assets at age 66, in percentages, as a result of removing spousal Social Security benefits and joint income taxation, unbalanced government budget

Appendix I. Comparing PSID and CPS data

Starting in 1968, the PSID has excellent data for the cohort of people we want to study. Its design allows the sample to remain representative of the US population. Despite attrition, it has maintained its cross-sectional validity, as discussed by Fitzgerald et al., 1998, and Moffitt and Zhang, 2018. Nonetheless, in this appendix, we compare the key moments from the PSID with the corresponding ones that we compute from the Current Population Survey (CPS) which does not have a panel dimension, and hence does not allow us to compute many of the inputs that we need, but has a relatively larger sample size.

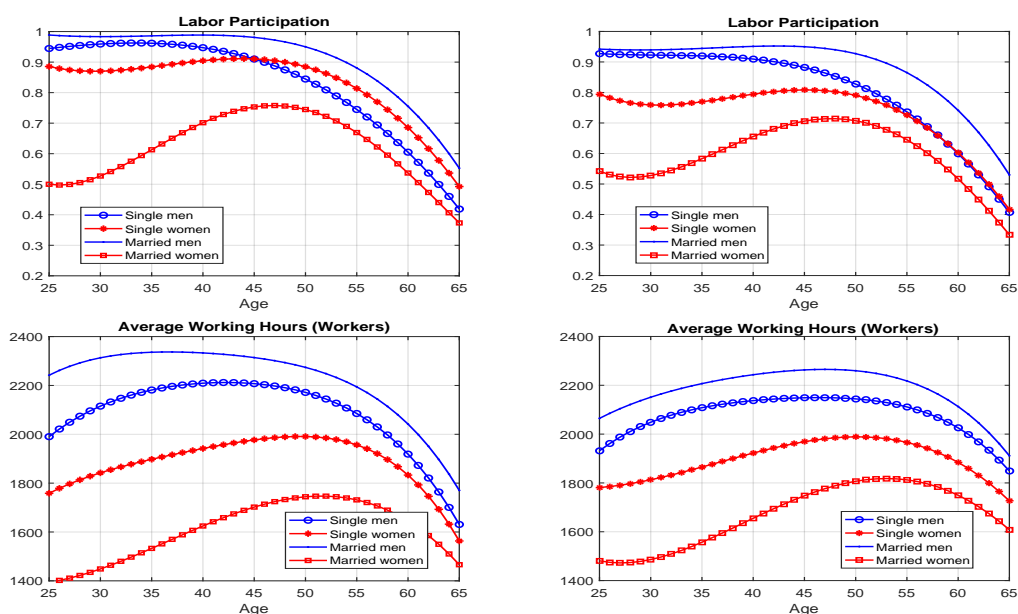


Figure 30: Life-cycle profiles by gender and marital status for the 1945 cohort in the PSID (left-hand-side panel) and CPS (right-hand-side panel) data

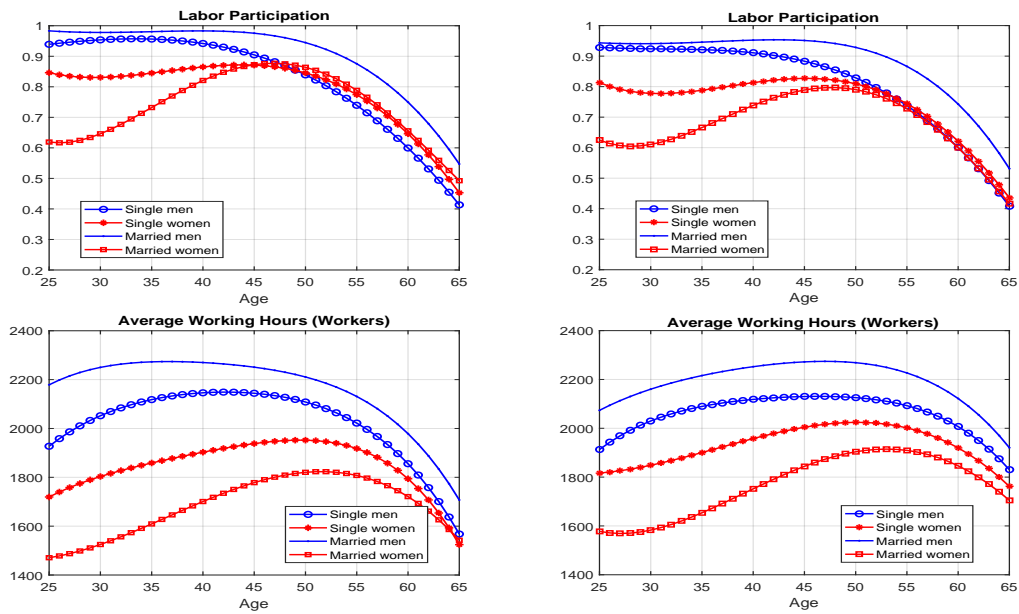


Figure 31: Life-cycle profiles by gender and marital status for the 1955 cohort in the PSID (left-hand-side panel) and CPS (right-hand-side panel) data