

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

Margherita Borella, Mariacristina De Nardi, and Fang Yang*

February 15, 2019

Abstract

In the U.S., 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 we 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 55, and increases the savings of couples in both cohorts, including in the later one, which has similar participation to that of more recent generations.

*Margherita Borella: University of Torino and CeRP-Collegio Carlo Alberto, Italy. Mariacristina De Nardi: UCL, Federal Reserve Bank of Chicago, IFS, 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, 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 Chicago, or the IFS.

1 Introduction

After increasing robustly from 1962 to the early 1990s, the labor force participation of prime-age women in the U.S. 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.”

In this paper, we ask to what extent the fact that both taxes and old age Social Security benefits depend on one’s marital status discourages female labor supply. The mechanisms are the following. First, as couples file taxes jointly, the secondary earner in the married couple 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 back women from working?

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 to work and how many hours 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 and pension payments with survival and spousal benefits, the differential tax treatment of married and single people, the progressivity of the tax system, and old-age means-tested transfer programs such as Medicaid and Supplemental Social Insurance (SSI). We also model the changes of the tax and Social Security system that our two cohorts faced over time.

We estimate our dynamic structural model using the Method of Simulated mo-

ments (MSM) and data from the Panel Study of Income Dynamics (PSID) and from the Health and Retirement Study (HRS) for the cohort born in 1941-1945 (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 period and retirement period, respectively. Then, taking the estimated preference parameters from that cohort as given, we also estimate our model for the 1951-1955 cohort (the “1955” cohort), which had much higher participation of married women (and closer to that of even more recent cohorts) and for which policy 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, but also to single women who expect to get married, and to married men after age 55. 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 five percentage points for single women. At age 45, participation for these groups is, respectively, still 15 and three percentage point higher without these marital benefits provisions. In addition, these marriage-based rules reduce the participation of married men starting at age 55, resulting in a participation that is three percentage points lower by age 65. Finally, for these cohorts, these marital provisions decrease savings of married couples by \$40,000 at age 70, and wages for married women by about 10%, due to the experience effect on wages. While our model takes marriage and divorce behavior from the observed data from each cohort, we show that our results are robust to changes in marriage and divorce behavior in Section 10.1.

Given that the labor supply of married women has been increasing fast 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 much more likely to work. To shed light on this question, we study a cohort that is ten years younger than our reference cohort, for which we still have a

completed labor market history, and whose labor market behavior is close to that of more recent cohorts, that is cohort born in 1955. By way of comparison, the labor market participation of married women at age 25 is just over 50% for our 1945 cohort, while 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 to the observed data. Finally, we run the policy experiment of eliminating the marriage-related provisions of 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 large also for cohorts in which the labor participation of married women is higher.

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. Our results show that, in addition to lowering the participation of women, these marriage-related policies also significantly reduce the savings of couples and the participation of married men starting in their middle age. Second, it is the first paper that studies 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 to study 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. Finally, our model fits the data for participation, hours worked, and savings, as well as 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.

2 Related literature

We build on the literature on the determinants of 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 to explain increases in female labor supply over time, while Eckstein, Keane, and Lifshitz (2016) examine the changes over time in the selection of married women working and find that it accounts for 75% of the observed increase over time in the marriage-wage premium (the differential in salary for married versus single women). Bronson (2015) uses a dynamic structural model of marriage, education choices, and labor supply to explain gender gaps in college attendance and choice of major. Her paper does not model marriage-related policies, nor their effects. Altonji and Vidangos (2017) empirically characterize the dynamics of marriage, divorce, children, on labor-market outcomes of married households. 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 also contribute to the literature on tax and benefit reforms. The vast majority of the existing literature on social insurance program reform adopts the paradigm of a one-person household, or a household in which only one person supplies labor, and abstracts from many important risks that households face over the life cycle and for which progressive taxation and social insurance provide valuable protection.

There is a smaller literature studying policy reforms in environments that includes 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 (2015), 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-Shundeln (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.

In the context of understanding portfolio choice, Love (2010) proposes the first model to evaluate the effects of family and marital status risk on savings and portfolio choice. Hubener, Maurer, and Mitchell (2016) study the effects of exogenous family dynamics and endogenous labor supply on portfolio choice and retirement behavior.

Compared with these papers, we propose a richer model, which allows for intensive and extensive labor supply decisions for both men and women, savings, medical expenses during retirement, marriage and divorce, and life expectancy heterogeneity as a function of marital status. In addition, we estimate our model by using detailed data from the PSID and the HRS for the 1945 and 1955 cohorts. Finally, our paper is the first one to examine the role of U.S. taxes and Social Security transfers, jointly and for two cohorts, for which we observe large changes in taxes and Social Security policies and economic behavior across cohorts and over time.

3 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). This implies that, when the secondary earners in couples work, their marginal tax rate is based on their own income only rather than on the sum of their partner's income and their own, and is thus typically lower under progressive taxation. The U.S., instead, taxes the income of married couples jointly (joint taxation) and uses a different tax schedule for married and single individuals. 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. This feature of the tax system tends to be larger the larger the husband's earnings.

To quantify this disincentive to work for the secondary earner, we use data from the PSID for 25-year-old single and married women in our 1945 cohort and our estimated effective tax functions. The left hand side panel of Figure 1 highlights the incentives to work for 25 year old single and married women. It plots four marginal tax rates as a function of women's earnings: the marginal tax rate for a single woman and the marginal tax rates of married women with husbands at the 25th, 50th, and 75th




Figure 1: Left panel: 25-year-old women’s marginal tax rate when single (red-starred line) or married to a men at the 25th (orange-dashed line), 50th (orange-dotted line), and 75th earnings percentile (orange-circled line) as a function of women’s earnings in 1998 dollars. Right panel: cumulative density function (cdf) of 25-year-old non-working wives’ marginal tax rates.

earnings percentiles. A 25-year-old single woman born in the 1945 cohort pays a zero percent marginal tax rate to start working, while a married woman faces a marginal tax rate of 14%, 16%, and 17%, respectively, if she is married to a 25-year old men in the 25th, 50th, and 75th earnings percentile. The corresponding earnings levels for these men in 1998 dollars are, respectively, \$21,552, \$32,729, and \$42,167.

While this graph tells us that married women typically pay a larger marginal tax rate than single women, it does not tell us what is the distribution of marginal tax rates for married women to start working. This is what the right hand side panel of Figure 1 is meant to illustrate. It displays the distribution of the marginal tax rates for young 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 face a 0% marginal tax rate, while 90% of married women face a marginal tax rate of 10% or higher, due to 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 on the labor market participation of married women.

Social Security for a single individual is a function of one’s average lifetime earnings. Social Security for a married person is the highest of one’s own benefit enti-

tlement and half of the spouse's entitlement while the other spouse is alive (spousal benefit) and the highest of own benefit entitlement and the deceased spouse's after the spouse's death (survivor 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 hand panel of Figure 2 plots household Social Security




Figure 2: Left panel: Average household Social Security benefits at age 66 by wife's own Social Security benefit decile, with (blue-dotted line) or without (red-starred line) marital benefit, in 1998 dollars. Right panel: Average survivor benefit by wife's own Social Security benefit decile, with (blue-dotted line) or without (red-starred line) marital benefits, 1998 dollars.

benefits while the husband is alive. It takes married women at retirement age and, based on the deciles of their own Social Security's entitlement, plots their average household yearly Social Security benefits with (circled line) and without (crossed line) marital benefits. For instance, the number one on the x axis represents married women aged 66 in our 1945 cohort that are in the lowest decile of their own Social Security contributions. At that decile, household Social Security benefit for those women and their husbands are \$22,000 under marital benefits and about \$15,000 without marital benefits. The comparison of the two lines in this picture reveals that about 50% of married women benefit from Social Security marital benefits while their husband is alive and that these benefits can be very large.

The righthand 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, a 66-year old married women at the lowest 10% of Social Security contributions, once a widow, would receive less than \$500 dollars a month based on her own contributions only, while she receives \$15,000 thanks to her husbands’ 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.

While survival benefits are larger than marital benefits, they are typically received only for a few years after one’s husband’s death. In contrast, marital benefits are received over a longer period. This last set of graphs highlights that Social Security marital benefits are large and can also reduce married women’s incentives to work.

4 Data

We use the PSID and the HRS data. 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, this is a cohort for which we have excellent data 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.

We use the PSID to estimate all of the data we need for the working period and the HRS to compute inputs for the retirement period. In Appendix A we discuss these data sets and provide details about our computations.

4.1 The fraction of married people and life-cycle patterns for single and married men and women in our cohorts

Table 1 shows that the majority of men and women are married and that the fraction of married people goes down only slightly across these cohorts.

Figure 3 displays participation and average annual hours worked for workers. The top panel refers to the 1945 cohort¹. In that cohort, the top left panel shows

¹The smoothed profiles are obtained by fitting a fourth-order polynomial in age fully interacted

Gender	Born in 1941 – 1945			Born in 1951 – 1955		
	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 1: Fraction of married men and women by age and cohort, PSID data

that married men have the highest participation rate and only slowly decrease their participation starting from age 45, while single men decrease their participation much faster. The participation of single women starts about 10 percentage points lower than that of men, but it gradually increases until age 50. Married women have the lowest participation. It starts around 50% at age 25, increases to 78% between age 40 and 50, and gradually declines at a similar rate as 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 bottom panel displays the same data for the 1955 cohort. Comparing the top and bottom panels shows a large increase in participation by married women across these two cohorts and, to a much smaller extent, by single women. 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.

Due to the limited availability of asset data in the PSID (it is available only every 5 years until 1999 and every other year afterwards) and to the fact that our 1955 cohort has not yet retired, we use the same asset profiles for both cohorts. Figure 3 displays average assets increase until age 70 for all groups, with women accumulating the lowest amount and showing no sign of a slowdown in accumulation before age 75.

5 The model

Our model period is one year long. People start their economic life at age 25, stop working at age 66 at the latest, and live up to the maximum age of 99.

During the working stage, people choose how much to save and how much to
with marital status and cohort dummies




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

work, face wage shocks and, if they are married, divorce shocks. Single people meet partners. Given that we explicitly model labor participation and hours of husbands and wives, savings and medical expenses in old age, the model is computationally very intensive.² For tractability, we assume exogenous marriage and divorce probabilities and we estimate them 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

²As we show later, we optimize over four value functions, compute other two, and have six continuous state variables. Even parallelizing on C on high-end computers, the model requires 37 minutes to be solved. Estimating the model implies solving it thousands of times, that is 3 or 4 weeks at least.

assumption in section 10. Also for tractability, we assume exogenous fertility and that women have an age-varying number of children that depends on their age and marital status. We take the number of children from the data.

During the retirement stage, people face out-of-pocket medical expenses (which are net of Medicare payments) and are partly covered by Medicaid and by Social Security payments. Married retired couples also face the risk of one of the spouses dying. Single retired people face the risk of their own death. We allow mortality risk and medical expenses to depend on gender, age, health status, and marital status.

We allow for both time costs and monetary costs of raising children and running households. They enter our problem in the following way. We allow for available time to be split between work and leisure to depend on gender and marital status. We interpret this available time to be net of home production and child care that one has to perform whether working or not. We estimate available time using our model. We will then compare our model's implications on time use to those from the PSID data and the literature.

All workers have to pay a fixed cost of working which, for women, depends on their age (which also maps in her number of children). Finally, when women work, they also have to pay a child care cost that depends on her number and age of children, and her earnings. That is, child care costs are a normal good: women with higher earnings pay for higher-quality (and more expensive) child care. We estimate their size using our model.³

5.1 Preferences

Let t be age $\in \{t_0, t_1, \dots, t_r, \dots, t_d\}$, with $t_0 = 25$, $t_r = 66$ being retirement time and $t_d = 99$ being the maximum possible lifespan. 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 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,

³Introducing 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.

respectively.

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

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

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.

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

$$l_t^{i,j} = L^{i,j} - n_t - \Phi_t^{i,j} I_{n_t},$$

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. This is a convenient way to represent activities that require time and cannot easily be outsourced.

Leisure equals available time endowment less n_t , hours worked on the labor market, less the fixed time cost of working. That is, the term I_{n_t} is an indicator function which 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 function⁴

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

⁴This a generalization of the functional form in Casanova (2012). An alternative is to use the collective model and solve for intra-household allocation as in Chiappori (1988, 1992), and Browning and Chiappori (1998)). We abstract from that for tractability.

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

5.2 The 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 individual observes his/her current idiosyncratic wage shock, age, assets, and accumulated earnings. Each married person also observes their partner's labor wage shock and accumulated earnings.

At the beginning of each retirement period, each single individual 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.

5.2.1 Human capital and wages

There are two components to wages. The first component is human capital, which is a function of one's initial conditions, individual's labor market experience, past earnings, age, gender, and marital status and that we denote $e_t^i(\cdot)$ ⁵.

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.

5.2.2 Marriage and divorce

During the working period, a single person gets married with an exogenous probability which depends on his/her age, gender, and wage shock. To simplify our computations, we assume that people who are married to each other have the same age.

The probability of getting married at the beginning of next period is

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

⁵We make this relationship explicit when describing our value functions for the working age and we define all of our state variables.

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

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

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 partner's wage shock. Thus, we have

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

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 moves on with those assets and their own wage shock and Social Security contributions.

After retirement, single individuals don't get married anymore while people in couples no longer divorce and only lose their spouse because of death. In the data, these events happen very infrequently in our cohort.⁶

5.2.3 The 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 determine the child care costs of working mothers ($i = 2$). The term $\tau_c^{0,5}$ is the child care cost for each child age 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 age 6 to 11, which are $f^{6,11}(i, j, t)$. We use our structural model to estimate these costs.

⁶In the HRS data, we observe our 1941-1945 birth cohort between the age of 66 and 72. Over that six-year period, only 1% of couples get divorced and 4% of singles get married. Thus, the implied yearly probability of marriage and divorce is tiny.

5.2.4 Medical expenses and death

After retirement, surviving people face medical expenses and health and death shocks. At age 66, we endow people with a distribution of health that depends on their marital status and gender. 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.

5.2.5 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 data for each of our two cohorts.

5.3 The Government

We model taxes on total income Y as Gouveia and Strauss (1994) and we allow them to depend on marital status as follows

$$\tau(Y, j) = (b^j - b^j(s^j Y + 1)^{-\frac{1}{p^j}})Y.$$

The government also uses a proportional payroll tax on labor income τ_t^{SS} , up to a Social Security cap $\tilde{y}_t$, to help finance old-age Social Security benefits. We use $\bar{y}_t$ to denote an individual's average earnings at age t , which we use to determine old age Social Security and defined benefit pensions.

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

We allow both the payroll tax and the Social Security cap to change over time for each cohort, as in the data. We do not require the government to balance its budget, as it is not done cohort by cohort or for a couple of cohorts.

The insurance provided by Medicaid and SSI in old age is represented by a means-tested consumption floor, $\underline{c}(j)$. Borella, De Nardi and French (2017) discuss Medicaid rules and observed outcomes after retirement.

5.4 Recursive formulation

We define and compute six sets of value functions: the value function of working age singles, the value function of retired singles, the value function of working age couples, the value function of retired couples, the value function of an individual who is of working age and in a couple, the value function of an individual who is retired and in a couple.

5.4.1 The singles: working age and retirement

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

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

$$l_t^{i,j} = L^{i,j} - n_t^i - \Phi_t^{i,j} I_{n_t^i}, \quad (2)$$

$$Y_t = e_t^i(\bar{y}_t^i) \epsilon_t^i n_t^i, \quad (3)$$

$$\tau_c(i, j, t) = \tau_c^{0,5} f^{0,5}(i, j, t) + \tau_c^{6,11} f^{6,11}(i, j, t), \quad (4)$$

$$T(\cdot) = \tau(ra_t + Y_t, j), \quad (5)$$

$$c_t + a_{t+1} = (1 + r)a_t^i + Y_t(1 - \tau_c(i, j, t)) - \tau_t^{SS} \min(Y_t, \tilde{y}_t) - T(\cdot), \quad (6)$$

$$\bar{y}_{t+1}^i = \bar{y}_t^i + \min(Y_t, \tilde{y}_t)/(t_r - t_0), \quad (7)$$

$$a_t \geq 0, \quad n_t \geq 0, \quad \forall t. \quad (8)$$

Equation 3 shows that the deterministic component of wages is a function of age, gender, and labor market experience through $\bar{y}_t^i$, which, being the individual's average labor income so far, includes the effects of both previous labor supply and wages, and thus human capital. This formulation allows us to capture the important aspects of labor market experience and previous wages on current wages but does not force us

to keep track of additional state variables.

The expectation operator is taken with respect of the distribution of ϵ_{t+1}^i conditional on ϵ_t^i and with respect to the probability distribution of the partner's characteristics for people getting married $\xi_{t+1}(\cdot)$ and $\theta_{t+1}(\cdot)$. 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.

The state variables for a retired single individual are age t , gender i , assets a_t^i , health ψ_t^i , and average realized lifetime earnings $\bar{y}_r^i$. Because we assume that the retired individual can no longer get married, his or her recursive problem can be written as

$$R^s(t, i, a_t, \psi_t^i, \bar{y}_r^i) = \max_{c_t, a_{t+1}} \left(v(c_t, L^{i,j}) + \beta s_t^{i,j}(\psi_t^i) E_t R^s(t+1, i, a_{t+1}, \psi_{t+1}^i, \bar{y}_r^i) \right) \quad (9)$$

$$Y_t = SS(\bar{y}_r^i) \quad (10)$$

$$T(\cdot) = \tau \left(Y_t + r a_t, j \right) \quad (11)$$

$$B(a_t, Y_t, \psi_t^i, \underline{c}(j)) = \max \left\{ 0, \underline{c}(j) - \left\{ (1+r)a_t + Y_t - m_t^{i,j}(\psi_t^i) - T(\cdot) \right\} \right\} \quad (12)$$

$$c_t + a_{t+1} = (1+r)a_t + Y_t + B(a_t, Y_t, \psi_t^i, \underline{c}(j)) - m_t^{i,j}(\psi_t^i) - T(\cdot) \quad (13)$$

$$a_{t+1} \geq 0, \quad \forall t \quad (14)$$

$$a_{t+1} = 0, \quad \text{if } B(\cdot) > 0 \quad (15)$$

The term $SS(\bar{y}_r^i)$ includes Social Security and defined benefit plans, which for the single individual is a function of the income earned during their work life, $\bar{y}_r^i$, while $s_t^{i,1}(\psi_t^i)$ is the survival probability as a function of age, gender, marital 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, which ensure a minimum consumption floor $\underline{c}(j)$.

5.4.2 The couples: working age and retirement

The state variables for a married couple in 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 ($j = 2$)

before t_r can be written as:

$$\begin{aligned}
W^c(t, a_t, \epsilon_t^1, \epsilon_t^2, \bar{y}_t^1, \bar{y}_t^2) = & \max_{c_t, a_{t+1}, n_t^1, n_t^2} \left(w(c_t, l_t^{1,j}, l_t^{2,j}) \right. \\
& + (1 - \zeta_{t+1}(\cdot)) \beta E_t W^c(t+1, a_{t+1}, \epsilon_{t+1}^1, \epsilon_{t+1}^2, \bar{y}_{t+1}^1, \bar{y}_{t+1}^2) \\
& \left. + \zeta_{t+1}(\cdot) \beta \sum_{i=1}^2 \left(E_t W^s(t+1, i, a_{t+1}/2, \epsilon_{t+1}^i, \bar{y}_{t+1}^i) \right) \right) \quad (16)
\end{aligned}$$

$$l_t^{i,j} = L^{i,j} - n_t^i - \Phi_t^{i,j} I_{n_t^i}, \quad (17)$$

$$Y_t^i = e_t^i(\bar{y}_t^i) \epsilon_t^i n_t^i, \quad (18)$$

$$\tau_c(i, j, t) = \tau_c^{0,5} f^{0,5}(i, j, t) + \tau_c^{6,11} f^{6,11}(i, j, t), \quad (19)$$

$$T(\cdot) = \tau(r a_t + Y_t^1 + Y_t^2, j) \quad (20)$$

$$c_t + a_{t+1} = (1+r)a_t + Y_t^1 + Y_t^2 (1 - \tau_c(2, 2, t)) - \tau_t^{SS}(\min(Y_t^1, \tilde{y}_t) + \min(Y_t^2, \tilde{y}_t)) - T(\cdot) \quad (21)$$

$$\bar{y}_{t+1}^i = \bar{y}_t^i + (\min(Y_t^i, \tilde{y}_t)) / (t_r - t_0), \quad (22)$$

$$a_t \geq 0, \quad n_t^1, n_t^2 \geq 0, \quad \forall t \quad (23)$$

The expected value of the couple's value function is taken with respect to the conditional probabilities of the two ϵ_{t+1} s given the current values of the ϵ_t s for each of the spouses (we assume independent draws). The term $\zeta_{t+1}(\cdot) = \zeta_{t+1}(\epsilon_t^1, \epsilon_t^2)$ represents the probability of divorce for a couple at age $t+1$ with wage shocks ϵ_t^1 and ϵ_t^2 . The expected values for the newly divorced people are taken using the appropriate conditional distribution for their own labor wage shocks.

During retirement, that is from age t_r on, each of the spouses is hit with a health shock ψ_t^i and a realization of the survival shock $s_t^{i,2}(\psi_t^i)$. Symmetrically with the other shocks, $s_t^{1,2}(\psi_t^1)$ is the after retirement survival probability of husband, while $s_t^{2,2}(\psi_t^2)$ is the survival probability of the wife. We assume that the deaths of the each spouse are independent of each other.

In each period, the married couple's ($j = 2$) recursive problem can be written as

$$R^c(t, a_t, \psi_t^1, \psi_t^2, \bar{y}_r^1, \bar{y}_r^2) = \max_{c_t, a_{t+1}} \left(w(c_t, L^{1,j}, L^{2,j}) + \right. \\ \beta s_t^{1,j}(\psi_t^1) s_t^{2,j}(\psi_t^2) E_t R^c(t+1, a_{t+1}, \psi_{t+1}^1, \psi_{t+1}^2, \bar{y}_r^1, \bar{y}_r^2) + \\ \beta s_t^{1,j}(\psi_t^1) (1 - s_t^{2,j}(\psi_t^2)) E_t R^s(t+1, 1, a_{t+1}, \psi_{t+1}^1, \bar{y}_r^1) + \\ \left. \beta s_t^{2,j}(\psi_t^2) (1 - s_t^{1,j}(\psi_t^1)) E_t R^s(t+1, 2, a_{t+1}, \psi_{t+1}^2, \bar{y}_r^2) \right) \quad (24)$$

$$Y_t = \max \left\{ (SS(\bar{y}_r^1) + SS(\bar{y}_r^2), \frac{3}{2} \max(SS(\bar{y}_r^1), SS(\bar{y}_r^2))) \right\} \quad (25)$$

$$\bar{\bar{y}}_r^i = \max(\bar{y}_r^1, \bar{y}_r^2), \quad (26)$$

$$T(\cdot) = \tau \left(Y_t + r a_t, j \right) \quad (27)$$

$$B(a_t, Y_t, \psi_t^1, \psi_t^2, \underline{c}(j)) = \max \left\{ 0, \underline{c}(j) - \left[(1+r)a_t + Y_t - m_t^{1,j}(\psi_t^1) - m_t^{2,j}(\psi_t^2) - T(\cdot) \right] \right\} \quad (28)$$

$$c_t + a_{t+1} = (1+r)a_t + Y_t + B(a_t, Y_t, \psi_t^1, \psi_t^2, \underline{c}(j)) - m_t^{1,j}(\psi_t^1) - m_t^{2,j}(\psi_t^2) - T(\cdot) \quad (29)$$

$$a_{t+1} \geq 0, \quad \forall t \quad (30)$$

$$a_{t+1} = 0, \quad \text{if } B(\cdot) > 0. \quad (31)$$

In equation (25), the evolution of variable Y_t mimics the spousal benefit from Social Security and pension which gives a married person the right to collect the higher of own benefit entitlement and half of the spouse's entitlement. In equation (26), the evolution of variables $\bar{\bar{y}}_r^i, i = 1, 2$ represents survivorship benefits from Social Security and pension in case of death of one of the spouses. The survivor has the right to collect the higher of own benefit entitlement and the deceased spouse's entitlement.

5.4.3 The individuals in couples: 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, while when computing the value of getting married for a single person, the

relevant object for that person is his or her the 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

$$\begin{aligned}\hat{W}^c(t, i, a_t, \epsilon_t^1, \epsilon_t^2, \bar{y}_t^1, \bar{y}_t^2) = & v(\hat{c}_t(\cdot)/\eta_t, \hat{l}_t^{i,j}) + \\ & \beta(1 - \zeta(\cdot))E_t\hat{W}^c(t+1, i, \hat{a}_{t+1}(\cdot), \epsilon_{t+1}^1, \epsilon_{t+1}^2, \bar{y}_{t+1}^1, \bar{y}_{t+1}^2) + \\ & \beta\zeta(\cdot)E_tW^s(t+1, i, \hat{a}_{t+1}(\cdot)/2, \epsilon_{t+1}^i, \bar{y}_{t+1}^i)\end{aligned}\quad (32)$$

where $\hat{c}_t(\cdot)$, $\hat{l}_t^{i,j}(\cdot)$, and $\hat{a}_{t+1}(\cdot)$ are, respectively, optimal consumption, leisure, and saving for an individual of gender i in a couple with the given state variables.

During the retirement period, we have

$$\begin{aligned}\hat{R}^c(t, i, a_t, \psi_t^1, \psi_t^2, \bar{y}_r^1, \bar{y}_r^2) = & v(\hat{c}_t(\cdot), \eta_t, L^{i,j}) + \beta s_t^{i,j}(\psi_t^i) s_t^{p,j}(\psi_t^p) E_t\hat{R}^c(t+1, i, \hat{a}_{t+1}(\cdot), \psi_{t+1}^1, \psi_{t+1}^2, \bar{y}_r^1, \bar{y}_r^2) + \\ & \beta s_t^{i,j}(\psi_t^i) (1 - s_t^{p,j}(\psi_t^p)) E_tR^s(t+1, i, \hat{a}_{t+1}(\cdot), \psi_{t+1}^i, \bar{y}_r^i).\end{aligned}\quad (33)$$

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

where $s_t^{p,j}(\psi_t^p)$ is the survival probability of the partner of the person of gender i and the term $\bar{\bar{y}}_r^i$ represents the Social Security survivor benefits.

6 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 estimate or calibrate those parameters that can be cleanly identified outside our model. For example, we estimate the probabilities of marriage, divorce, and death, as well as wage processes while working and medical expenses during retirement, directly from the data for that cohort, and we calibrate the interest rate and a few other model parameters.

In the second step, we use the method of simulated moments. For the 1945

cohort, we estimate 19 model parameters $(\beta, \omega, (\phi_0^{i,j}, \phi_1^{i,j}, \phi_2^{i,j}), (\tau_c^{0,5}, \tau_c^{6,11}), L^{i,j})$.⁷ 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 until 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 until age 65 (41 time periods for each group).
3. Because net worth, together with labor supply, is an essential to smooth resources during the working period and to finance retirement we match net worth for three groups (couples and single men and women) starting at age 26 and until age 65 (40 time periods for each group).⁸ 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 which allows us to simulate life-cycle histories for asset, participation, and hours. We simulate a large number of artificial individuals, that are

⁷We normalize the time endowment of single men.

⁸Net worth at age 25 is an initial condition.

initially endowed with a value of the state vector drawn from the data distribution for each cohort at age 25 (that is, assets, accumulated Social Security, and wage shocks for singles and the same variables for each of the partners for a couple), 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.

6.1 First-step estimation for our two cohorts

Table 2 summarizes 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 and so are the estimates of the probability of marriage and divorce by age, gender, and wage shocks. Appendix B details all of these inputs, while Appendix C reports additional first-steps inputs for both of each cohorts.

6.1.1 Marriage, divorce, spousal assets and Social Security benefits, and wages

We use the PSID to estimate the probabilities of marriage and divorce. Figure 4 displays our estimated 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 less 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 5 shows that married men with lower wage shocks are more likely to get divorced. The probability of divorce decreases with age, and so 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 smaller in our more recent cohort once we condition on age and wage shocks.

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

Estimated processes		Source
Wages		
$e_t^i(\cdot)$	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 age 0-5	PSID
$f^{6,11}(i, j, t)$	Number of children age 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
Calibrated parameters		Source
Preferences and returns		
r	Interest rate	4% De Nardi, French, and Jones (2016)
η_t	Equivalence scales	PSID
γ	Utility curvature parameter	2.5 see text
Government policy		
b^j, s^j, p^j	Income tax	Guner et al. (2012b)
$SS(\bar{y}_r^i)$	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	\$6,950, De Nardi et al. (2016)
$\underline{c}(2)$	Minimum consumption, couples	\$6,950*1.5 Social Security rules

Table 2: First-step inputs summary

We also estimate the joint distribution of (the logarithm of) the wage shocks of new husbands and new wives⁹ 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. Due to these initial

⁹We assume it to be the same for both cohorts because the number of new marriages after age 25 is small during this time period.




Figure 4: Marriage probabilities by gender, age and one's wage shock for the 1945 cohort (left panel) and 1955 cohort (right panel), PSID data




Figure 5: Divorce probabilities by gender, age and one's wage shock for the 1945 cohort (left panel) and 1955 cohort (right panel), PSID data

correlations and the high persistence of shocks for an individual that we estimate, partners tend to have positively correlated shocks even after getting married.

6.1.2 Children

Figure 6 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.

We use the average total number of children for single and married women by age to compute equivalence scales (as $\eta(f_t) = (j + 0.7 * f_t)^{0.7}$, as estimated by Citro and Michael (1995), with j being equal to 1 or 2 depending on marital status). We also use the number of children in those two different age groups to compute child care costs.




Figure 6: Number of Children for married and single women for the 1945 cohort (left panel) and 1955 cohort (right panel), PSID data

6.1.3 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 capital at a given point in time as one's average realized earnings up to that time. Thus, we allow past wages and labor market experience to affect one's wage today. We estimate this relationship from the PSID data.¹⁰




Figure 7: 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 7 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

¹⁰Since we already keep track of average realized earnings to compute Social Security benefits, this formulation does not require us to add state variables to our already computationally intensive model.

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. The marriage premium has decreased from the 1945 to the 1955 cohort because the average wage of married women has increased, while the average wage for men has stagnated. This is due to a combination of both different 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).

To show the effect of human capital on wages, Figure 8 displays the implied profiles of wages for men and women that we estimate, evaluated at different percentiles of human capital ($\bar{y}_t$) for each cohort in our PSID data. The bottom line refers to the the lowest level of human capital (which is zero), while the top one corresponds to the top 1% of each cohort. The left and right panels respectively refer to the 1945 and 1955 cohorts.

The left panel of Figure 8 shows that a man of the 1945 cohort entering the labor market at age 25 with no accumulated human capital earns an hourly wage of about \$10, while a woman of the same cohort, human capital, and age would earn less than \$9. The wage at entry, when human capital is equal to zero, is slightly increasing for both men and women up to age 50, and then slightly decreasing after age 55. At




Figure 8: The effect of human capital on men's and women's wage, 1945 cohort. Average wage profiles for constant values of human capital at 0\$, and at the 25th, 50th, 75th and 99th percentiles. Left panel: 1945 cohort. Right panel, 1955 cohort. PSID data

every age, the hourly wage rate increases with human capital, peaking at \$28 for men in the top 1%. The profile for women peaks later, at age 54, and women in the top 1% of human capital earn \$24 an hour.

The right panel of the figure refers to the 1955 cohort, for which average wage at a given human capital level is slightly lower than that for the 1945 cohort. Thus, consistent with in Figure 7, wages for men have fallen from the 1945 cohort to the 1955 one. Those for women, for given level of human capital, have dropped by less and women’s human capital has increased, which explains why the wages of married women have increased on average, as shown in Figure 7.

The shock in log wages is modeled as the sum of a persistent component and a white noise, which we assume captures measurement error, and thus we do not include in our structural model. We assume that this shock processes are cohort-independent. Table 3 reports our estimates for the AR component of earnings. They imply that men and women face 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.

Parameter	Men	Women
Persistence	0.941	0.946
Variance prod. shock	0.026	0.015
Initial variance	0.114	0.095

Table 3: Estimated processes for the wage shocks for men and women, PSID data

6.1.4 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 this 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. The left panel of Figure 9 displays the survival probabilities by gender and marital and health status. Women, married people, and healthy people have longer life expectancies. In Borella, De Nardi, and Yang (2017), we have shown that

our estimated mortality rates line up very well with the life tables.




Figure 9: Left panel: Survival probability by age, gender, and marital and health status, both cohorts. Right panel: Medical expenditure by age, gender, and marital and health status, both cohorts. HRS data

The left panel of Figure 9 displays the importance of medical expenditures after retirement. Average medical expenses climb fast past age 85 and are highest for single and unhealthy people. Figure 24 in Appendix B reports our estimated health transition matrices by gender, age, and marital and health status.

6.1.5 Calibrated parameters

We set the interest rate r to 4% and the utility curvature parameter, γ , to 2.5. We use the tax function for married and single people estimated by Guner et al. (2012b). We set the minimum consumption for the elderly singles at \$6,950 in 1998 dollars, as in De Nardi, et al. (2016) 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. Appendix B reports the Social Security benefit function and the parameters in the tax function.

6.2 Second step estimates

Table 4 presents our estimated preference parameters for both cohorts.¹¹

For the 1945 cohort, our estimated discount factor is .994, the same value estimated by De Nardi et al. (2016) on a sample of elderly retirees and our estimated

¹¹Appendix F reports all of our estimated parameters for both cohorts and their standard errors.

Estimated parameters	1945 cohort	1955 cohort
β : Discount factor	0.994	0.994
ω : Consumption weight	0.403	0.403
$L^{2,1}$: Time endowment (weekly hours), single women	107	105
$L^{1,2}$: Time endowment (weekly hours), married men	102	98
$L^{2,2}$: Time endowment (weekly hours), married women	85	84
$\tau_c^{0,5}$: Prop. child care cost for children age 0-5	34%	56%
$\tau_c^{6,11}$: Prop. child care cost for children age 6-11	18%	14%
$\Phi_t^{i,j}$: Partic. cost	Fig. 10	Fig. 10

Table 4: Second step estimated model parameters

weight on consumption is 0.4. We assume that the 1955 cohort shares these preference parameters.

While we normalize total weekly time endowment of single men to 5840 hours a year, and thus 112 hours a week, for our 1945 cohort, we 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, 102 and 85 hours. This implies that people in the latter two groups spend 10 and 27 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). They find that, in the 1985 American Time Use Survey (ATUS) dataset (when our 1945 cohort was age 42), 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 activities related to home production on average. It should be noted that, even for a working woman, 28 hours can amount to, for example, spending nine 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 with 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 age 0-5 and 6-11 are, respectively, 34% and 18% of a woman's wage. The PSID only reports child care costs for all children in the age range 0-11 and implies that the per-child child care costs of a married woman vary with age and they are for example 29% and 19% of her wages at ages 25 and 30 respectively. To compare our results with the PSID data, we aggregate the costs obtained in estimation and find that the per child care costs (for all children in the age range 0-11) of a married woman are 30% and 25% of her wages, respectively, at ages 25 and 30. Thus, our model infers child care cost for women that are very similar to those in the PSID data.

For the 1955 cohort, we notice two main changes compared to 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 here due to space constraints.




Figure 10: Estimated lifecycle labor participation costs expressed as fraction of the time endowment of a single men. SM: single men; SW: single women; MM: married men; MW: married women. Left panel: 1945 cohort. Right panel: 1955 cohort. Model estimates

Figure 10 reports the age-varying time costs of working by age expressed as frac-

tion of the time endowment of a single men that are necessary to reconcile the labor market participation of our four groups of people in each cohort. Our estimated participation cost 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 other factors 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 daycare. They also show that, conditional on all aspects of our environment, the participation costs of married women are the lowest ones. This is because married women face lower wages, have a smaller time endowment (due to the time spent engaging in home production and childcare), and tend to have higher-wage husbands who work.


To reconcile the higher participation of married women at younger ages in the 1955 cohort, our model estimates a lower time cost of working for married women when younger. Coupled with our previous finding that that married men have less available time to work and enjoy leisure, and that per-child child care costs are higher for this cohort, this might indicate that married women in the 1955 cohort are both better able to get their husbands to help with childcare and to purchase more childcare on the market, compared to the married women in the 1945 cohort.

6.3 Model fit


Figures 11 and 12 report our model-implied moments, as well as the moments and 95% confidence intervals from the PSID data for our 1945 cohort. They show that our parsimoniously parameterized model (19 parameters and 448 targets) fits the data very well. It reproduces the important patterns of participation, hours conditional on participation, and asset accumulation for all of our demographic groups. 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).

6.4 Identification


The fixed cost of participation by age and subgroup ($\Phi_t^{i,j}$) especially impacts participation by subgroup and over the lifecycle. 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




(a) Participation, couples



(b) Participation, singles



(c) Hours for workers, couples




(d) Hours for workers, singles


Figure 11: Model fit for participation (top graphs) and hours (bottom graphs) and average and 95% confidence intervals from the PSID data

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 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 than for younger workers.



(a) Assets, couples



(b) Assets, singles

Figure 12: Model fit for assets and average and 95% confidence intervals from the PSID data

7 Model validation and implications

The previous section shows that our model fits the data targets by age for assets, participation, and hours worked. In this section, we show that our model also matches labor income over the life cycle, Social Security benefits for single and married men and women (which are endogenous to experience formation and labor supply in our framework) at age 70, and labor supply elasticities at the extensive and intensive margin for the 1945 cohort. To save space, we do not evaluate all of these implications for the 1955 cohort.

Figure 13 shows that our model reproduces the main features of labor income observed in the data for married and single men and women. This is not an obvious implication because, while we match participation and hours conditional on participation, our wages are endogenous to labor supply decisions and could thus generate earning profiles that do not resemble those in the data.

Table 5 reports the Social Security benefits at age 70 implied by our model and in the HRS data and shows that these numbers closely match. The fact that the model's implications for both labor income and Social Security benefits over the life cycle are




Figure 13: Labor income, data (left graph) and model implications (right graph).

consistent with the data is important because these implications have to do with the available resources that the households have to consume, save, and pay taxes, first during their working period, and then during their retirement.

	Married men	Married women	Single men	Single women
Model	14,690	10,560	14,460	11,100
Data	14,290	9,160	13,450	12,800

Table 5: Average Social Security benefits at age 70, model and HRS data

To compare the model's implications for labor supply elasticity with those estimated by many others, we compute a version of compensated labor supply elasticity. To do so, we temporarily increase the wage for only one age and one group (either married men, or married women, or single men, or single women) at a time by 5%. This change is anticipated. Table 6 shows the (compensated) elasticities of participation and hours among workers with respect to a change to their own wage. It shows several interesting features. First, the elasticity of participation of women is larger than that of men, both for married and singles. Second, married men display small participation elasticities to temporary and small wage changes, and have the lowest elasticity of participation among the four groups that we consider. Third, the elasticity of participation for all groups is largest around retirement age, a finding that confirms that of French (2005) for men. Our elasticities are in line with those estimated by the existing literature.

Our model was not required to match any of these additional important aspects of the data. The fact that it does provides further support for the validity of the model's prediction in response to policy changes.

	Participation				Hours among workers			
	Married		Single		Married		Single	
	W	M	W	M	W	M	W	M
30	0.7	0.0	0.5	0.1	0.1	0.2	0.1	0.3
40	0.5	0.1	0.4	0.2	0.3	0.4	0.4	0.2
50	0.6	0.2	0.4	0.6	0.5	0.4	0.5	0.3
60	1.0	0.8	1.7	1.1	0.3	0.3	0.4	0.5

Table 6: Model-implied elasticities of labor supply

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 permanent wages change such as those implied by a permanent tax change. To help shed light on what we should expect from our policy experiments, we report here the effects of a permanent increase of 5% in the wage schedule of married women when the wage structure of the other three demographic groups remains the same.

Panel (a) of Figure 14 shows that a permanent wage increase for married women implies a much larger, 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. Panel (b) 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.

Turning to the elasticity of married women’s behavior to changes in their husband’s wages, Figure 15 highlights a negative cross-elasticity of participation for married women which is as high as -1.8 at age 25, decreases to -0.7 between ages 43 and 50, and then increases again to about -1.5 after age 60. This change also increases the labor supply of married men after age 40.

We do not report the effects of permanently changes the wages of singles, as their




Figure 14: 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




Figure 15: Change in participation for a 5% permanent increase in the wage schedule of married men. Effect on all four demographic groups. Model implications

effects are very small, nor the effect of wages on hours, which are also small.

8 Eliminating Marital Social Security benefits and joint taxation for the 1945 cohort

We now turn to using our model to evaluate the effect of various policy reforms. For each of them, we leave everything else unchanged, including taxes. This has the benefit of focusing on the policy change at hand and does not require the government to balance the budget for one cohort.

8.1 Eliminating all 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 the partner is alive and all of the benefits of the 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 all spousal Social Security benefits, both while the spouse is alive and after his or her death, the government budget surplus increases by 24% and the top left panel of Figure 16 shows that the participation of married women is, respectively, eight, twelve, and six percentage points higher at ages 25, 55-60, and 65. In contrast, men decrease their participation starting at age 55 and their participation is three 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. Even the participation of single women increases (by four percentage points) because, should they get married, they now expect no Social Security benefits coming from their spouse's labor supply. As they age, the probability that they get married becomes negligible and the effect on spousal benefits elimination on their participation fades.

Groneck and Wallenius (2017) focus on the redistribution of Social Security related to marital benefits in a simpler model than ours, for instance with three-year periods, much less flexibility in adjusting hours and calibrating all of their model parameters, but also report that, over all of the working period, their model implies an increase in the participation by married women of 6.4 percentage points. Kaygusuz (2015) finds an increase in the participation by married women of 6.1 percentage points overall. Their findings are a bit smaller but in the ballpark of ours.

The top right panel of Figure 16 reports changes in labor income for our four demographics groups. As a result of their increased labor supply and labor market experience, married women's labor income is about, respectively, \$1,200, \$2,500 and \$1,500 higher ages 25, 55-60 and 65. The labor income of married men drops by




Figure 16: Changes in participation (top left panel), labor income (top right panel) and savings (bottom panel) after the elimination of all the spousal Social Security benefits

about \$1,500 by age 65.

The bottom panel of Figure 16 displays changes in savings over the life cycle for married people. The effect of this reform is large and peaks at an accumulation of additional \$35,000 between the ages of 75 and 80 compared to our benchmark economy with Social Security marital benefits.

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

Government surplus increases by 25% as a result of this reform. The left sub-panel of Figure 17 displays the participation profiles in our benchmark economy and the economy in which every one files as an individual and there are no Social Security spousal benefits. The sub-panel on the left plots the difference in the participation profile in each group. These graphs highlight several important findings. First, eliminating all marriage-related taxes and transfers have a large effect of the participation of married and single women. The participation of married women is 15-22 percentage points higher until age 60 in the no-marital provisions economy. The one for single women is about five percentage points higher until age 40. The participation

of married men is higher in their middle age, reaching a peak of three percentage points higher than in the benchmark, but is three percentage points lower than in the benchmark after age 60. 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 17: 1945 cohort: Differences in participation (left panel) and participation in the benchmark and reformed economy (right panel) after the elimination of all the spousal Social Security benefits and of joint income taxation

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. They do not study the effect of marital Social Security benefits. When we perform the experiment of switching to single filer taxation in isolation, we find that our effect on the labor supply of married women is closer to the lower bound of those found by Guner, Kaygusuz, and Ventura (2012a).

As a result of increased labor market experience, the average wages of both married and single women increase; this increase peaks at \$0.8-0.9 between age 50 and 60 (left panel in Figure 18). Increased wages and participation (hours increase little for the workers) imply higher average earnings of \$3,000 per year over for married women and \$2,000 for single women for most of their life cycle. In contrast, average earnings of married men start dropping at age 62 and are \$1,500 a year lower by age 65. Looking at savings, couples accumulate less assets before age 40 because both partners work more and consume less leisure and more consumption goods. However, they then save much more rapidly and end up accumulating over \$40,000 more by age 75.




Figure 18: 1945 cohort: Changes in wages (left panel), labor income (right panel), and assets for couples (bottom panel) after the elimination of all the spousal Social Security benefits and with joint income taxation

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 when they enter the labor market at age 25. Table 7 shows that, in the 1945 cohort, 57% of women and only 23% 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 marital tax benefits.

	Wage quintile				
	1	2	3	4	5
Men	7.3%	16.3%	18.8%	24.9%	32.7%
Women	39.3%	17.8%	21.4%	15.0%	6.6%

Table 7: Distribution of men and women across potential wage quintiles at age 25, 1945 cohort, PSID data

9 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 benefits for the 1955 cohort. In the interest of space, we only report results for the case in which we eliminate all three marriage-related provisions at the same time.

The right sub-panel of Figure 19 displays the participation profiles in our benchmark economy and the economy in which every one files as an individual and there are no Social Security spousal benefits. The sub-panel on the right plots the difference in the participation profile in each group. These graphs show 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 relative younger cohort with much higher participation of married women continues to be very large.




Figure 19: 1955 cohort: Participation in the benchmark and reformed economy (left panel), differences in participation (right panel) after the elimination of all the spousal Social Security benefits and with joint income taxation

The effects on increased labor market experience on wages are similar to those in the 1945 cohort, (left panel in Figure 20) and, as for the 1945 cohort, increased wages and participation (hours increase little for the workers) imply higher average earnings of \$3-4,000 per year over for married women and \$2,000 for single women for most of their life cycle. Average earnings of married men start dropping earlier for this cohort, that is at age 55, compared to age 62 for the 1945 cohort, but their drop is smaller by age 65. Changes in savings for married people are very close to those for the older cohort and also peak at over \$40,000 more by age 75.




Figure 20: 1955 cohort: Changes in wages (left panel), labor income (right panel), and assets for couples (bottom panel) after the elimination of all the spousal Social Security benefits and with joint income taxation

	Wage quintile				
	1	2	3	4	5
Men	8.6%	16.9%	19.5%	28.1%	27.4%
Women	32.9%	26.9%	16.6%	12.3%	11.3%

Table 8: Distribution of men and women across potential wage quintiles at age 25, 1955 cohort, PSID data

Comparing Tables 7 and 8 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 the 1945 to the 1955 cohort but that it is still the case that, even in the 1955 cohort, most women tend to have lower wages and thus to be secondary earners in this cohort, and thus to respond strongly to the elimination of marital provisions.

10 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, 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, to evaluate the robustness of our results.

10.1 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 twenty years of marriage. After 1977, the threshold for eligibility became ten years of marriage before divorce.¹²

Dickert-Conlin and Meghea (2004) examine the 1977 U.S. policy switch (using data from the National Vital Statistics, the 1980 Census, and The Current Population Survey) and conclude that it had no effects on divorces and remarriages. Goda et al. (2007) also find no impact of the 10-year eligibility discontinuity (using data from the PSID Marital History File) on divorces. Dillender (2016) confirms that these rules have no effects overall and small ones on a very small number of people, those who married late. This group of people accounts for a very small fraction of marriages and divorces in the aggregate. Thus, previous literature indicates that while 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

¹²We 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.

affect these decisions.” Alm and Whittington (1997) use data from the Panel Study on Income Dynamics 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 U.S., “In general, the impacts of the income tax variables, even when statistically significant, are small.”

Thus, for the U.S., 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. The small effects apply to tiny groups of people which have very little weight on the aggregates.

Looking into welfare programs, Low, Meghir, Pistaferri, and Voena (2018) study the U.S. subpopulation of mothers on welfare and the 1996 welfare reform, which was meant to 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 immediately reduced the welfare recipience and increased the labor participation of mothers, but had no effects on marriage and fertility. They do highlight an effect on divorces, 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.

More broadly on the effects of welfare programs, a previous survey by Moffitt (1998) concludes that “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 U.S. In terms of comparison with our work, her main finding is that divorce rate increased

by 10% as a result of the elimination of marital survivorship benefits. Although this is a sizeable effect 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, The United Nations (Department of Economic and Social Affairs) reports that in 1985, a time period proceeding the 1989 Swedish marital benefits reform, only 35.8% of the 25-29 year old Swedish women were married, compared to 64.3% in the U.S.. In addition, in the 1980s 18% of live births were born to unmarried women in the United States, compared to 40% in Sweden (Sorrentino, 1990).

10.2 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 to the effects of the same policy when there are also two alternative possible changes in marriage and divorce patterns. In the first robustness exercise, the policy also decreases marriage rates by 20% and increases divorce rates by 20%. Alternatively, in the second robustness exercise, the policy also increases marriage rates by 20% and decreases divorces rates by 20%.

Figure 21 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) shows that a reform that lowers the probability of marriage and raises that of divorce makes women more self-reliant on their own labor supply and human capital. Married women work more (and accumulate more human capital) to edge against divorce risk. Single women are less likely to get married and also work more (and also accumulate more human capital). Comparing the left panel (no marriage and divorce changes) and the right panel (increased marriage and decreased divorce) highlights that increasing marriage rates




Figure 21: Differences in participation after the elimination of all the 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.

and lowering divorce rates has the opposite effects, but that these effects are small and do not change the conclusions that we reach in our original policy experiment.

Changes in hours worked are also very similar to each other across these three experiments and we do not report them due to space constraints.




Figure 22: Changes in assets for couples after the elimination of all the spousal Social Security benefits and with 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 22 compares savings for the three experiments. The middle and right panel show that the effects on savings of couples are also very robust to changes in expected marital patterns.

Our results are thus robust to large changes in marriage and divorce rates.

11 Conclusions

We estimate a model of labor supply and savings of single and married people, which allows for a rich representation of the risks that people face over their entire life cycle and that allows for the important provisions of taxes and Social Security for singles and couples.

We estimate our model to 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 of 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 of married women, but also on that 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 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 the 1955 birth cohorts, despite the fact that the labor market participation of young married women in the 1955 cohort is over ten percentage points higher than that of the 1945 cohort.

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 that studies 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 to study 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* 13(3), 193-204.
- [4] Alm, James and Leslie A. Whittington. 1997. “’Til Death or Taxes Do Us Part: The effects 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 Ivan Vidangos. 2017. “Marriage Dynamics, Earnings Dynamics, and Lifetime Family Income. ” Mimeo.
- [7] 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.
- [8] Attanasio, Orazio, Hamish Low, and Virginia Sánchez-Marcos. 2008. “Explaining Changes in Female Labor Supply in a Life-Cycle Model.” *The American Economic Review*, 98(4), 1517-1552.
- [9] Bick, Alexander, and Nicola Fuchs-Schündeln. 2018. “Taxation and Labor Supply of Married Couples across Countries: A Macroeconomic Analysis. ” *The Review of Economic Studies*, 85(3), 1543-1576.
- [10] Blau, David M. and Donna Gilleskie. 2006. “Health Insurance and Retirement of Married Couples.” *Journal of Applied Econometrics*, 21(7), 935-953.

- [11] Blundell, Richard and Luigi Pistaferri and Itay Saporta-Eksten. 2016. "Consumption Inequality and Family Labor Supply." *American Economic Review*, 106(2), 387–435.
- [12] Blundell, Richard and Costa Dias, Monica and Meghir, Costas and Shaw, Jonathan. 2016. "Female labor supply, human capital, and welfare reform." *Econometrica*, 84(5), 1705-1753.
- [13] Blundell, Richard and Jack Britton and Monica Costa Dias and Eric French. 2017. "The impact of health on labour supply near retirement." Mimeo.
- [14] Bronson, Mary Ann. 2015. "Degrees Are Forever: Marriage, Educational Investment, and Lifecycle Labor Decisions of Men and Women." Mimeo.
- [15] Borella, Margherita and Mariacristina De Nardi and Eric French. 2017. "Who received Medicaid in old age? Rules and reality." *Fiscal Studies*, forthcoming.
- [16] Borella, Margherita and Mariacristina De Nardi and Fang Yang. 2017. "The Aggregate Implications of Gender and Marriage." *The Economics of Ageing*, May 2018, 6-26.
- [17] Browning, Martin, and Pierre-André Chiappori. 1998. "Efficient intrahousehold allocation: a characterisation and tests. " *Econometrica*, 66(6), 1241-78.
- [18] Casanova, Maria. 2012. "Happy Together: A Structural Model of Couples' Joint Retirement Choices." Mimeo.
- [19] Chiappori, Pierre-André. 1988. "Rational Household Labor Supply. " *Econometrica*, 56, 63-89.
- [20] Chiappori, Pierre-André. 1992. "Collective labor supply and welfare. " *Journal of Political Economy*, 100(3), 437-67.
- [21] Citro, Constance Forbes, and Robert T. Michael, 1995, Measuring poverty: a new approach, National Academy Press.
- [22] De Nardi, Mariacristina and Eric French, and John Jones. 2010. "Why do the Elderly Save? The Role of Medical Expenses," *Journal of Political Economy*, 118, 39-75.

- [23] De Nardi, Mariacristina and Eric French, and John Jones. 2016. "Medicaid Insurance in Old Age," *The American Economic Review*, 106(11), 3480-3520.
- [24] Dickert-Conlin, Stacey and Meghea, Christian. 2004. "The Effect of Social Security on Divorce and Remarriage Behavior," Working paper no. 9, Center for Retirement Research.
- [25] Dillender, Marcus. 2016. "Social Security and Divorce. " *The B.E. Journal of Economic Analysis and Policy*, 16(2), 931-971.
- [26] Dotsey, Michael, Wenli Li, and Fang Yang, 2014, "Consumption and Time Use over the Life Cycle." *International Economic Review*, 55 (3), 665-692.
- [27] Duffie, Darrell, and Kenneth J. Singleton. 1993. "Simulated Moments Estimation of Markov Models of Asset Prices. " *Econometrica*, 61(4), 929-952.
- [28] Eckstein, Zvi and Osnat Liftshitz. 2011. "Dynamic Female Labor Supply." *Econometrica* 79(6), 1675-1726.
- [29] Eckstein, Zvi, and Michael Keane, and Osnat Liftshitz. 2016. "Sources of Change in the Life-Cycle Decisions of American Men and Woman: 1962-2014. " Working paper.
- [30] French, Eric. 2005. "The Effects of Health, Wealth, and Wages on Labor Supply and Retirement Behavior." *Review of Economic Studies*, 72(2), 395-427.
- [31] Gouveia, Miguel and Robert P. Strauss. 1994. "Effective Federal Individual Income Tax Functions: An Exploratory Empirical Analysis." *National Tax Journal*, 47(2), 317-39.
- [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 Papers 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] Hubener, Andreas, Raimond Maurre, and Olivia Mitchell. 2016. "How Family Status and Social Security Caliming Options Shape Optimal Life Cycle Portfolios. " *The Review of Financial Studies*, 29(4), 937-978.
- [38] Kaygusuz, Remzi. 2015. "Social Security and Two-Earner Households." Mimeo.
- [39] Love, David. 2010. "The Effects of Marital Status and Children on Savings and Portfolio Choice. " *Review of Financial Studies*, 23, 328-355.
- [40] Low, Hamish, Costas Meghir, Luigi Pistaferri, and Alessandra Voena. 2018. "Marriage, Labor Supply and the Social Safety Net." Mimeo.
- [41] Robert Moffit. 1998. "Welfare, The Family, And Reproductive Behavior, Research Perspectives. " Robert A. Moffitt, Editor. National Academy Press. Washington D.C.
- [42] Persson, Petra. 2017. "Social Insurance and the Marriage Market." Mimeo.
- [43] Nishiyama, Shinichi. 2015. "The Joint Labor Supply Decision of Married Couples and the Social Security Pension System." Mimeo.
- [44] Pakes, Ariel, and David Pollard. 1989. "Simulation and the Aysmptotics of Optimization Estimators. " *Econometrica*, 57(5), 1027-1057.
- [45] Roys, Nicolas and Christopher Taber. 2017. "Skills Prices, Occupations and Changes in the Wage Structure for Low Skilled Men. " Mimeo.
- [46] Sorrentino, Constance. 1990. "The Changing Family in International Perspective." *Monthly Labor Review*, 113(3), 41-58.
- [47] 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>
" Web page.

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, and to compute 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 (the “SRC sample”), and from an over-sample of 1,872 low-income families from the Survey of Economic Opportunity (the “SEO sample”). Individuals have been followed over time to maintain a representative sample of families.

We study the two cohorts born in 1941-45 and in 1951-55. 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 aged 20 to 70, for a total of 103,420 observations. In general, to gather the information we need, we control for cohort effects in our estimates, and use the results relative to the cohorts of interest. Table 9 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	5,137	105,381
Age between 20 and 70	5,129	103,420

Table 9: Sample Selection in the PSID

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 aged 50 or older, including a wide range 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 re-interviewed every two years. Other cohorts were introduced over the years, the AHEAD (Assets and Health

Dynamics Among the Oldest Old) cohort, born before 1924, was first interviewed in 1993, while the CODA (Children of the Depression) cohort, born 1924 to 1930, and the WB (War-Baby) cohort, which includes individuals born 1942 to 1947, were introduced in 1998 and subsequently interviewed every two years. Younger cohorts, the EBB (Early Baby Boomer), born 1948 to 1953, and the MBB (Mid Baby Boomer), born 1954 to 1959, were first interviewed in 2004 and 2010 respectively.

Our data set is based on the RAND HRS files and 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 (2,275 individuals) because marital status is crucial information in our analysis. This sample consists of 35,042 individuals and 176,698 observations. We then select individuals in the age range 66-100 born in 1900 to 1945, obtaining a sample of 15,072 individuals and 67,744 observations. As we cannot observe individuals born after 1945 and aged at least 66 in the HRS, for the 1955 cohort we use the same estimates obtained for the 1945 one.

Appendix B. First step estimation, methodology

Wages

In this section we describe how we estimate the deterministic wage function and the distribution of the wage shocks. We proceed in two steps, because in the PSID data the wage is observed only for those who work. In the first step, we obtain an estimate of latent wages for those who are not working by imputation. In the second step, we estimate the average latent wage profile by age and experience, and the persistence and variance of its unobserved component.

Imputation. The observed wage rate is defined as annual earnings divided by annual hours worked. Gross annual earnings are defined as previous year’s income from labor, while annual hours are previous year’s annual hours spent working for pay.

Wages 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 \$250 (in 1998 values). We use the same imputation procedure for all these cases.

More specifically, we impute wage values using coefficients from fixed effects regressions run separately for men and women. To avoid endpoint problems with the polynomials in age, we include individuals aged 22 to 70 in the sample. Define:

$$\ln wage_{kt}^i = I_{ikt}^n \widehat{\ln wage}_{kt}^i.$$

where i denotes gender, k denotes an individual, and t is age. I_{ikt}^n is an indicator of participation (equal to 1 if the individual participates in the labor market and has no missing hours or wage) and $\widehat{\ln wage}$ is latent or potential wage, that we wish to estimate. We estimate:

$$\ln wage_{kt}^i = Z'_{kit} \beta_z^i + f_k^i + \varsigma_{kt}^i,$$

where the index i emphasizes the fact the we run separate regressions for each gender. The dependent variable is the logarithm of the observed hourly wage rate, f_k^i is an individual-specific fixed effect and ς_{kt}^i is an error term. We include a rich set of explanatory variables: a fifth-order polynomial in age, a third-order polynomial in experience (measured in years), 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. We also add interactions between these variables. As an indicator of health, we use a variable recording whether bad health limits the capacity of working, as this is the only health indicator available for all years (self-reported health starts in 1984 and is not asked before). However, as this health indicator is not collected for wives, we do not include it in the regression for married women. All regressions also include interaction terms between the explanatory variables. Using the estimated coefficients, the predicted value of the (logarithm of the) wage is taken as a measure of the potential wage for observations with a missing wage. When the wage is observed, we use the actual wage. Our estimated latent wage is then:

$$\begin{aligned} \widehat{\ln wage}_{kt}^i &= Z'_{kit} \hat{\beta}_z^i + \hat{f}_k^i \quad \text{if } I_{ikt}^n = 0 \\ \widehat{\ln wage}_{kt}^i &= \ln wage_{kt}^i \quad \text{if } I_{ikt}^n = 1. \end{aligned}$$

Wage function. Having estimated the (logarithm of) latent wage for all obser-

variations in the sample, we can estimate the deterministic wage function e_t^i .

To take into account the effect of human capital on the current wage in a way that does not add state variables to the structural model, we condition the age-efficiency profile on age and average realized earnings accrued up to the beginning of age t ($\bar{y}_t^i$), which are a measure of accumulated human capital.

We describe the computation of average realized earnings later in this Appendix. As it turns out, conditional on an individual's previous average earnings, differences due to marital status are not statistically significant, and we do not include them.

To estimate the wage profile, we run a fixed-effect regression of the logarithm of the latent wage rate on a fourth-order age polynomial, fully interacted with gender, and include average realized earnings as a regressor. We then regress the residuals from this regression on cohort dummies to compute the average effect for our two cohorts of interest.

More specifically, we estimate:

$$\widehat{\ln wage}_{kt}^i = X'_{kit}\beta_x + \ln u_{kt}^i, \text{ with } \ln u_{kt}^i = d_k^i + w_{kt}^i$$

where the dependent variable is imputed wage as computed in the previous step, and the explanatory variables include age, gender, average earnings $\ln(\bar{y}_{kt})$ and interaction terms. To fix the constant of the wage profile for our cohorts of interest, we then regress the residuals on cohort dummies to compute the average effects for the cohorts born in 1941-45 and in 1951-55 respectively.

The estimated potential wage profiles, computed at average values of $\ln(\bar{y}_t)$, are shown in the main text. We also use these residuals to estimate an AR(1) plus white noise process. In this case, we limit the age range between 25 and 65. As we rely on residuals also taken from imputed wages, we drop the highest 0.5% residuals both for men and women, in order to avoid large outliers to inflate the estimated variances (however, the effect of this drop is negligible on our estimates).

The shock in log wage is modeled as the sum of a persistent component plus white noise, which we assume captures measurement error:

$$\ln u_{k,t+1}^i = \ln \epsilon_{k,t+1}^i + \xi_{k,t+1}^i \quad (34)$$

$$\ln \epsilon_{k,t+1}^i = \rho_\epsilon^i \ln \epsilon_{k,t}^i + v_{k,t+1}^i, \quad (35)$$

where i indicates gender, and $\xi_{k,t+1}^i$ and $v_{k,t+1}^i$ are independent white-noise processes with zero mean and variances equal to σ_{ξ}^2 and σ_{iv}^2 , respectively. This last variance together with the persistent parameter ρ_{ϵ}^i characterize the AR process in the model. Estimation is carried out on residuals $\ln \hat{u}_{kt}^i$.

Because initial conditions and marriage and divorce are functions of one's wage shocks, we need the value of those wage shocks for each person of working age over time. To obtain them, we notice that the system formed by (34) and (35) can be estimated by Maximum Likelihood, which can be constructed assuming the initial state of the system and the shocks are Gaussian, and using the standard Kalman Filter recursions. We then also use this estimates of the productivity shocks as an explanatory variable when computing marriage/divorce probabilities, as well as the initial distribution.

Wealth

We define wealth as total assets defined as the sum of all assets types available in the PSID, net of debt and plus the value of home equity. We apply an imputation procedure to gather information on assets in missing years, as described below. All monetary values are expressed in 1998 prices.

Wealth in the PSID is only recorded in 1984, 1989, 1994, and then in each (biennial) wave from 1999 onwards. We rely on an imputation procedure to compute wealth in the missing years, from 1968; this procedure is based on the estimation of a fixed-effect regression which allows to estimate each individuals' fixed effect. In order to get a reliable estimate of the age profile at young ages, however, we include also younger cohorts. The cohorts used in all previous steps were born between 1931 and 1955: this implies that in 1984 the youngest individuals were aged 29, and if we restrict the sample as usual to those cohorts we cannot estimate the age profile at the beginning of the life cycle. We therefore include, in the imputation regression only, individuals born between 1956 and 1965, in order to have enough observations observed at young ages to rely on when estimating the age profile to be used to impute wealth. The equation we estimate is:

$$\ln(a_{kt}^i + \delta_a^i) = Z_{kit}'\beta_z^i + da_k^i + wa_{kt}^i, \quad (36)$$

where, as before, k denotes the individual, i gender, and t is age. δ_a^i is a shift

parameter for assets to have only positive values and to be able to take logs, and the variables Z include polynomials in age, average earnings (uncapped), hours of work, family size, age of the youngest child, dummies for marital and health status, and interactions among the variables. da_k^i is the individual fixed effect, invariant regardless the individual is single, married or divorced, and wa_{kt}^i is a white-noise error term. Equation (36) is estimated separately for men and women, on an enlarged sample of individuals born between 1931 and 1965.

We then use the estimated coefficients and individual fixed effects to impute missing observations (that is, wealth in missing years) using:

$$\ln(\widehat{a_{kt}^i + \delta_a^i}) = Z'_{kit} \hat{\beta}_z^i + \hat{da}_k^i$$

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

$$\begin{pmatrix} \ln(a_t^i + \delta_a^i) \\ \ln(\bar{y}_t^i) \\ \ln \epsilon_t^i \end{pmatrix} \sim N \begin{pmatrix} \mu_{at}^i + \delta_a^i \\ \mu_{yt}^i \\ \mu_{\epsilon t}^i \end{pmatrix}, \Sigma_{st}^i, \quad (37)$$

where Σ_s is a 3x3 covariance matrix. We completely characterize the distribution by estimating its mean and its variance, both depending on age t . To get an estimate of 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 in turn to the cohort born in 1945 and in 1955, is the age-specific estimate of the mean of the log-normal distribution. Taking residuals from the above estimates, we can obtain an estimate of the elements of the variance-covariance matrix, by taking 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

$$\begin{pmatrix} \ln(a + \delta_a) \\ \ln(\bar{y}^1) \\ \ln(\bar{y}^2) \\ \ln(\epsilon^1) \\ \ln(\epsilon^2) \end{pmatrix} \sim N \begin{pmatrix} \mu_a + \delta_a \\ \mu_{\bar{y}1} \\ \mu_{\bar{y}2} \\ \mu_{\epsilon1} \\ \mu_{\epsilon2} \end{pmatrix}, \quad \Sigma_c, \quad (38)$$

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

We use wealth with imputed missing observations when computing the initial joint distribution (equations (37) and (38)) and the descriptive wealth age profiles.

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 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 (4.86 per cent, that is 5,153 out of 105,985) 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, 1,969 enter the sample before turning 27: for those individuals we assume average accumulated earnings at entry is equal to zero. For the remaining 3,184 we use an imputation procedure to recover average realized earnings at entry and then we 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 entry value for individuals entering the sample after turning 27. Average earnings is then updated for each individual following his/her observed earnings history.

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.

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.¹³

The Social Security tax and Social Security cap shown in Figure 23 have been changing over time. We also allow them to change over time for the households in our cohorts.




Figure 23: Social Security tax and Social Security cap over time (expressed in 1998\$) for the 1945 cohort (graphs on the left) and 1955 cohort (graphs on the right)

Taxes

Guner et al. (2012b) estimate the tax function by marital status. The resulting values for a married couple are: $b^2 = 0.2338$, $s^2 = 0.0032$, $p^2 = 1.493$; Those for singles are: $b^1 = 0.2462$, $s^1 = 0.0311$, $p^1 = 0.8969$.

¹³<https://www.ssa.gov/oact/cola/bendpoints.html>. We use value in 2009 year and convert them into 1998 dollar.

Marriage and divorce probabilities

We also use the PSID to estimate the probabilities of marriage and divorce. More specifically, we model the probability of getting married, $\nu_{t+1}(\cdot)$, as a function of gender, age and the wage shock. We perform the estimation separately for men and women. Our estimated equation is

$$\Pr(Married_{t+1} = 1 | Married_t = 0, Z_{t+1}^i) = F(\beta_m^i Z_{t+1}^i),$$

where the variables Z include age, the logarithm of the wage shock, squares of all variables, interactions of levels and squares, and cohort dummies. Using the estimated coefficients on the cohort dummies, we then adjust the probability for the 1945 and the 1955 cohort respectively. F is the standard logistic distribution.

Similarly, we estimate the probability of divorce, $\xi_{t+1}(\cdot)$, as

$$\Pr(Divorced_{t+1} = 1 | Married_t = 1, X_{t+1}^i) = F(\beta_d^i X_{t+1}^i)$$

where the explanatory variables are: age of the husband, age squared, husband's wage shock, wife's wage shock, and interactions between age and wage shocks. Also in this case, we add cohort dummies and use the estimated coefficients to adjust the probability of divorce for the 1945 and 1955 cohort. F is the standard logistic distribution.

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

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

We compute the above probability using our estimates of the wage shocks, by partitioning households in age groups (25-35; 35-45; 45-65) and computing the variance-covariance matrix of newly matched partners' wage shocks in each age group. We then assume that the joint distribution is lognormal. As in the whole sample we observe 750 new marriages in the age range 25-65, we do not allow this probability to depend on cohort.

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 log-normally distributed at each age with mean and variance computed from sample values. Assets include a shifter as described for the computation of joint the distribution at age 25 (see Wealth subsection in this Appendix).

Number of children

To compute the average number of children by age group we use the PSID, where information on the total number of children and on the age of the youngest child is available, while information on the age of each child is missing. Using the panel structure of the data to update the age of existing children and to infer the arrival of new children in the family, we are able to construct the age structure of children in each family. More specifically, of the 2,543 women in our dataset, 71% enter the sample with no children or one child. As information about the age of the youngest child is available, following those families through time enables us to update the age of the initial child. In addition, we can assign a year of birth and follow through time any child born during the sample period. Having done that, we add up the number of children in each age category (0 to 5, 6 to 11) and run a regression on a fifth-order polynomial in age, 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 relative to the cohorts of mothers born in 1941-45 and in 1951-55.

Similarly, we compute the average age profile of the total number of children, which is available in the PSID, running a regression on a polynomial in age, interacted with marital status, and cohort dummies to construct the average age profile of total number of children for single and married women.

Health status at retirement

We use the HRS data and define health status on the basis of self-reported health, a variable that can take five possible values (excellent, very good, good, fair, poor). As standard, bad health status is defined as a dichotomous variable equal to 1 if self-reported health is fair or poor and 0 otherwise.¹⁴

¹⁴Looking at labor supply behavior about retirement time, Blundell, Britton, Costa Diaz and French (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.

We estimate the probability of being in bad health at the start of retirement, that is at age 66, using the observed frequencies for the 1941-1945 cohort, which is the youngest cohort that we can observe aged 66+ in the HRS data. All the inputs estimated from the HRS correspond to the 1941-45 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, in the HRS, 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:

$$\pi_{\psi t} = Prob(\psi_t = 1 \mid X_t^\psi) = \frac{\exp(X_t^\psi)}{1 + \exp(X_t^\psi)}.$$

The set of explanatory variables X_t^ψ includes a third-order polynomial in age, previous health status, gender, marital status, and interactions between these variables. The logit is estimated on all individuals born in 1900 to 1945, and includes cohort dummies among the regressors; we then use coefficients relative to the 1941-1945 cohort as inputs in our model for both of our cohorts. As the HRS data are collected every two years, we obtain two-year probabilities, and convert them into one-year probabilities. Figure 24 displays the health transition matrix by gender, age, marital status, and health status that we estimated.

Survival probabilities

We assume that people are alive until retirement age for sure. After that, we model the probability of being alive at time t as a logit function

$$\pi_{st} = Prob(S_t = 1 \mid X_t^s) = \frac{\exp(X_t^s)}{1 + \exp(X_t^s)}.$$




Figure 24: Health transition probabilities for singles and couples by age. HRS data

that we estimate using the HRS data. Among the explanatory variables, we include a fourth-order polynomial in age, gender, marital status (single or married) and health status in the previous period, as well as interactions between these variables. We also include cohort dummies and use coefficients relative to the 1941-1945 cohort for both cohorts. 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.

Out-of-pocket medical expenditures

Out-of-pocket (oop) medical expenditures are defined as the sum the individual spends out of pocket in hospital costs, nursing home costs, doctor visits costs, dentist costs, outpatient surgery costs, average monthly prescription drug costs, home health care and special facilities cost. They include medical expenditures in the last year of life, as recorded in the exit interviews. On the other hand, expenses covered by public or private insurance are not included in our measure. The estimated equation is:

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

where explanatory variables include a third-order polynomial in age, current health status, gender, marital status, and interactions between these variables. 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).

Appendix C. First step estimation, additional results

Spousal assets and Social Security benefits




Figure 25: Spousal assets by spousal wages shocks in case of marriage next period for the 1945 cohort (left panel) and 1955 cohort (right panel) panel, PSID data

Figure 25 reports spousal assets by spousal wage shocks in case of marriage next period. Both panels show that both women and men getting married 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 get married later experience much larger variation in partner's assets conditional on 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 get married later and later. The patterns are very close for the two cohorts.




Figure 26: 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) panel, PSID data

Figure 26 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. Their shape by age reflects the different participation and hours decisions made by women and men over their life cycle. Here, too, the graphs for the two cohorts are similar, but reveal a bit more Social Security benefits by spousal wage shocks in the 1955 cohort, and especially so when marrying middle-age men.

Appendix D. The solution algorithm

We compute three sets of value functions: the value function of being single, the joint value function of the two people in the married couple, and the value function of each individual in a couple. During the retirement period, single people do not get married anymore, hence their value function can be computed independently of the other value functions. The value function of the 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. This is thus how the value functions are computed during retirement:

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 individual in case of death of one of the spouses, doing the usual backward iteration starting from the last period.
3. Compute the value function of the single individual in a marriage for all time periods after retirement.

During the working age, the value functions are interconnected, hence we solve each of them at a given time t . Working backwards over the life cycle

1. For any given time period, 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 we then 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 an individual of being in a couple.
4. Keep going back in time until the first period.

After we solve the value functions and policy functions, we simulate our model economy using the inputs and the procedures that we describe in the next section.

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 keep notation light, we do not include a separate indicator for each of the two cohorts that we estimate the model for.

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 $(\beta, \omega, (\phi_0^{i,j}, \phi_1^{i,j}, \phi_2^{i,j}), (\tau_c^{0,5}, \tau_c^{6,11}), L^{i,j})$.¹⁵ For the 1955 cohort, we assume that the households have the same β and ω as the 1945 cohort and we 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 life-cycle profiles found in the data and the simulated profiles generated by the model.

For each working age from 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 individual k , of gender i , marital status j , and age t . Denoting $z_t^{i,j}(\Delta, \chi)$ the model predicted expected value of the generic variable z at 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:

$$E[a_{k,t}^{i,j} - a_t^{i,j}(\Delta_0, \chi_0)] = 0, \quad \forall t = 2, \dots, 41 \quad (39)$$

¹⁵We normalize the time endowment of single men.

$$E[H_{k,t}^{i,j} - H_t^{i,j}(\Delta_0, \chi_0)] = 0, \quad \forall t = 1, \dots, 41 \quad (40)$$

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

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. In the end, we have a total of $J = 448$ moment conditions.

In practice, we compute the sample expectations in equations (39), (40) and (41) conditional on a flexible polynomial in age. More in detail, 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 dataset of K individuals 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

$$\underset{\Delta}{\operatorname{argmin}} \hat{\varphi}_K(\Delta; \chi_0)' \widehat{\mathbf{W}}_K \hat{\varphi}_K(\Delta; \chi_0). \quad (42)$$

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

Under 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:

$$\sqrt{K} \left(\hat{\Delta} - \Delta_0 \right) \rightsquigarrow N(0, \mathbf{V}), \quad (43)$$

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

$$\mathbf{V} = (\mathbf{D}'\mathbf{W}\mathbf{D})^{-1} \mathbf{D}'\mathbf{W}\mathbf{S}\mathbf{W}\mathbf{D}(\mathbf{D}'\mathbf{W}\mathbf{D})^{-1}, \quad (44)$$

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

$$\mathbf{D} = \frac{\partial \varphi(\Delta; \chi_0)}{\partial \Delta'} \Big|_{\Delta=\Delta_0} \quad (45)$$

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) pointed 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


	Cohort 1945		Cohort 1955	
	Coeff.	Std. Err.	Coeff.	Std. Err.
β	0.9942	(0.00032)		
ω	0.4032	(0.00176)		
<i>Participation costs:</i>				
$\phi_2^{1,1}$	0.0002	(0.00002)	0.0003	(0.00003)
$\phi_1^{1,1}$	-0.0126	(0.00086)	-0.0260	(0.00110)
$\phi_0^{1,1}$	-1.0997	(0.02180)	-0.7464	(0.01712)
$\phi_2^{2,1}$	0.0005	(0.00002)	0.0006	(0.00002)
$\phi_1^{2,1}$	-0.0198	(0.00094)	-0.0266	(0.00103)
$\phi_0^{2,1}$	-1.1557	(0.02927)	-1.0316	(0.02444)
$\phi_2^{1,2}$	0.0011	(0.00003)	0.0012	(0.00004)
$\phi_1^{1,2}$	-0.0479	(0.00135)	-0.0604	(0.00137)
$\phi_0^{1,2}$	-1.3315	(0.02468)	-1.3661	(0.02541)
$\phi_2^{2,2}$	0.0043	(0.00010)	0.0041	(0.00014)
$\phi_1^{2,2}$	-0.1700	(0.00488)	-0.1507	(0.00836)
$\phi_0^{2,2}$	-2.3437	(0.05299)	-3.5088	(0.10222)
<i>Time endowments:</i>				
$FL^{2,1}$	-2.9758	(0.10054)	-2.6273	(0.06388)
$FL^{1,2}$	-2.3407	(0.03637)	-1.9071	(0.02466)
$FL^{2,2}$	-1.1139	(0.01369)	-1.0902	(0.00644)
<i>Childcare costs:</i>				
$\tau_c^{0,5}$	0.3422	(0.01232)	0.5576	(0.00756)
$\tau_c^{6,11}$	0.1820	(0.00728)	0.1442	(0.00596)

Table 10: Estimates of parameters. Standard Errors in parenthesis. 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.


Appendix G. Model fit for the 1955 cohort




(a) Participation, couples



(b) Participation, singles




(c) Hours for workers, couples




(d) Hours for workers, singles

Figure 27: Model fit for participation (top graphs) and hours (bottom graphs) and average and 95% confidence intervals from the PSID data

Figures 27 and 28 report our model implied moments as well as moments 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) Assets, couples



(b) Assets, singles

Figure 28: Model fit for assets and average and 95% confidence intervals from the PSID data