

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

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

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1853902> since 2023-01-27T09:09:04Z

Published version:

DOI:10.1093/restud/rdac018

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

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

Margherita Borella, Mariacristina De Nardi, and Fang Yang*

November 17, 2021

Abstract

In the United States, both taxes and old-age Social Security benefits depend on one's marital status and tend to reduce the labor supply of the secondary earner. To what extent are these provisions holding back female labor supply? We estimate a rich dynamic life-cycle model of labor supply and savings for couples and singles using the Method of Simulated Moments for the 1945 and 1955 birth cohorts. 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. It implies that eliminating these marriage-related provisions would drastically increase the participation of married women over their entire life cycle, reduce the participation of married men after age 60, and increase savings. If the resulting government surplus were used to lower income taxation, there would be large welfare gains for the vast majority of the population. These results hold for both cohorts, including the later one, which has participation similar to that of more recent generations.

*Borella: University of Torino and CeRP-Collegio Carlo Alberto. De Nardi: University of Minnesota, Federal Reserve Bank of Minneapolis, CEPR, and NBER. Yang: Louisiana State University. De Nardi gratefully acknowledges support from the ERC, grant 614328, "Savings and Risks." Yang gratefully acknowledges MRRC grant number 08098401 and hospitality from the Opportunity and Inclusive Growth Institute at the Federal Reserve Bank of Minneapolis. We thank Veronica Guerrieri and four anonymous referees, Joe Altonji, Richard Blundell, Monica Costa Dias, Zvi Eckstein, Joan Gieseke, James Holt, Rasmus Lenz, Derek Neal, Aviv Nevo, and Jon Skinner for useful comments and suggestions. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research, MRRC, the SSA, the CEPR, any agency of the federal government, the Federal Reserve Bank of Minneapolis, or the Federal Reserve System.

1 Introduction

After increasing robustly from the 1960s to the early 1990s, the labor force participation of women in the United States has been stagnating. In this paper, we ask to what extent the dependence of taxes and old-age Social Security benefits on marital status has held back female labor supply and affected the welfare of cohorts with different degrees of female labor market participation, including more recent ones.

The mechanisms through which these marital provisions work are the following. First, since couples file taxes jointly, the secondary earner faces a higher marginal tax rate, which tends to lower the secondary earner's labor supply. Second, married and widowed people can claim Social Security spousal and survivor 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 secondary earners, both provisions tend to hinder female labor supply. But to what extent are these disincentives holding it back?

To answer this question, we develop a rich dynamic life-cycle model with single and married people and estimate it for two cohorts using the Method of Simulated Moments (MSM) and data from the Panel Study of Income Dynamics (PSID) and the Health and Retirement Study (HRS). Our first cohort is the one born in 1941-1945 (the "1945" cohort), which has by now completed a large part of its life cycle and is covered by these two data sets (which provide excellent information on their working and retirement periods, respectively). Our second cohort is the 1951-1955 one (the "1955" cohort), which has by now completed a large part of its working period and has much higher levels of participation of married women (which are closer to those of more recent cohorts). Hence, for this cohort, policy and welfare implications might be quite different.

In our model, single people meet partners, and married people might get divorced. Every working-age person experiences wage shocks, and every retiree faces medical expenses and lifespan risk. Couples face the risks of both partners. Households can self-insure by saving and by choosing how much to work (or, in the case of couples, how much each partner works) and when to retire. We allow for human capital to affect wages, as is consistent with the data. We explicitly model Social Security with spousal and survival benefits, the differential tax treatment of married and single people, the progressivity of the tax system (including the Earned Income Tax

Credit (EITC)), and old-age, means-tested transfer programs such as Medicaid and Supplemental Security Income (SSI). We also model the changes in the tax and Social Security systems over time.

For both cohorts, our estimated model does a good job of matching the life-cycle profiles of labor market participation, hours worked by the workers, and savings for married and single people. It also generates elasticities of labor supply by age, gender, and marital status that are consistent with those estimated by others.

For the 1945 cohort, we find that the elimination of Social Security spousal and survivor benefits and joint income taxation has large effects on participation and savings. In particular, it raises the participation of married women at age 25 by over 20 percentage points and that of single women by 5 percentage points. Participation for these groups at age 45 is still 15 and 3 percentage points higher, respectively. In addition, the participation of married men decreases starting at age 60, resulting in a participation rate that is 7 percentage points lower by age 65. It also increases the savings of married couples by 20% at age 66. In terms of welfare, abolishing these marital provisions would benefit most couples, all single men, and over one-third of single women—thus benefitting over 90% of the people in this cohort.

The effects of these marital provisions on the participation, wages, earnings, and savings of the 1955 cohort are also large, thus indicating that they continue to affect more recent cohorts. In terms of welfare, abolishing these marital provisions at age 25 for this cohort would benefit most couples, all single men, and over three-fourths of single women. In addition, compared with the older cohort, the welfare benefits to those gaining would be much higher and the costs of those losing would be very small. This is because at age 25, the human capital of women in the 1955 cohort is higher than that in the previous cohort.

Our paper provides several contributions. First, it is the first paper to study all marriage-related taxes and benefits in a unified framework. Second, it does so by allowing for the large observed changes in the labor supply of married women over time by studying two different cohorts. Third, it is the first dynamic estimated structural model of couples and singles that allows for participation and hours decisions of both men and women in a framework with savings. Fourth, our framework is very rich along many other dimensions that are important for studying our problem. For instance, allowing for labor market experience to affect wages (of both men and women) is important for capturing their endogeneity and response to policy and marital status

changes. Allowing the tax structure to vary over time for each cohort (we estimate our tax functions from the PSID as a function of cohort, year, and marital status) takes this important variation into account when we estimate our model. Carefully modeling survival, health, and medical expenses in old age, and their heterogeneity by marital status and gender, is crucial to evaluate the effects of policy reforms on labor supply and savings. Modeling one-year periods gives people the ability to change their labor supply and savings in a more flexible and realistic way. Fifth, the model fits the data well, including in terms of elasticities that we do not match by construction, and thus provides a valid benchmark to evaluate policy reforms.

We build on the literature on female labor supply over the life cycle. Attanasio, Low, and Sánchez-Marcos (2008), Eckstein and Lifshitz (2011), and Eckstein, Keane, and Lifshitz (2019) examine the determinants of married women’s participation over time. Hubener, Maurer, and Mitchell (2016) study the effects of family dynamics and labor supply on portfolio choice and retirement. In addition, we contribute to the small body of literature studying policy reforms in environments that include couples. Guner, Kaygusuz, and Ventura (2012) study the switch to a proportional income tax and a reform in which married individuals can file taxes separately. They find that these reforms substantially increase female labor participation. Kaygusuz (2015), Nishiyama (2019), and Groneck and Wallenius (2020) find that removing Social Security spousal and survivor benefits would increase female labor participation, female hours worked, and aggregate output. Bick and Fuchs-Schündeln (2018) focus on a simpler static model of married couples and find that income taxes are an important factor driving differences across countries in the labor supply of married women. Finally, our work is also related to the literature on optimal income taxation of couples. Kleven, Kreiner, and Saez (2009) and Gayle and Shephard (2019) use static models of the couple. Kleven, Kreiner, and Saez (2009) show that dual-earner couples can arise either because the secondary earner has low costs of participating in the labor market or because he or she has low ability in home production. As the first case is the most empirically relevant, optimal tax rates display positive tax rates on secondary earnings along with negative jointness, whereby the tax rate of one person decreases with the earnings of the spouse. Gayle and Shephard (2019) empirically explore taxation design in a collective model of the household with a marriage market. They find that the optimal tax system for couples is characterized by negative jointness, but the welfare gains from jointness are modest.

2 Marriage, taxes, and Social Security benefits

Many countries tax the income of married people as if they were single (individual taxation). As a result, when the secondary earners in couples work, their marginal tax rate is based on their own income rather than on the sum of their partners' income and their own. The United States instead taxes the income of married couples jointly (joint taxation) and uses a different tax schedule for married and single people. The combination of joint taxation and a progressive tax system typically implies that a married secondary earner faces a higher marginal tax rate than a single earner.




Figure 1: Women's marginal tax rates as a function of earnings (2016 dollars). Single (starred) and married to men at different earnings percentiles: 25th (dashed), 50th (dotted), and 75th (circled).

To illustrate the secondary earner's disincentives to work, in Figure 1 we display the effective tax rates that we estimate from the PSID in 1988, when the EITC program was already active and people in our 1945 cohort were 45 years old.¹ It illustrates that, for instance, a single woman earning \$500 a year faces a marginal tax rate of -10%, while a married woman earning the same amount faces marginal tax rates of 14%, 18%, and 21%, if she is married to a man in the 25th, 50th, and 75th income percentiles, respectively (which correspond to, \$43,090, \$68,995, and \$113,288 in 2016 dollars, respectively). Our estimated negative tax rate at low income levels is due to the EITC. This graph shows that married women tend to face a higher marginal tax rate than single women, thus suggesting that making married people file as single rather than jointly could have large incentives for the labor market participation of married women.

¹The details of our tax computations are at the end of Appendix B.

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

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




Figure 2: Average household Social Security benefits (left) and average survivor benefit (right), both by wife’s own Social Security benefit decile at age 66, with marital benefits (circled) or without (crossed). All in 2016 dollars.

the husband is alive. It takes married women at retirement age and, based on the deciles of their own Social Security entitlement, plots their average household yearly Social Security benefits with (circled line) and without (crossed line) marital benefits. For instance, the number 1 on the x -axis represents 66-year-old married women who are in the lowest decile of their own Social Security contributions. At that decile, household Social Security benefits for those women and their husbands are \$32,000 under marital benefits and about \$22,000 without marital benefits. The comparison of the two lines in this picture reveals that about 50% of married women benefit from Social Security marital benefits while their husband is alive and that these benefits can be very large. The right graph of Figure 2 takes the same married women and plots what their yearly Social Security benefits would be with and without survivor’s benefits after their husband’s death. For instance, a 66-year old widow at the lowest 10% of Social Security contributions would receive less than \$500 a month based on her own contributions only, while she would receive \$22,000 with survivorship benefits. Because in this cohort most women have lower wages than men, participate

less, and work fewer hours, survivorship benefits are large for over 80% of married women. This last set of graphs highlights that Social Security marital benefits are large and can also reduce married women’s incentives to work.

3 Life-cycle patterns for single and married people

We pick the 1945 cohort because its members’ entire adult life is covered first by the PSID, which starts in 1968 and has rich information for the working period, and then by the HRS, which starts with people age 50 and older in 1994 and has rich information for the retirement period. 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.²

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

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

The bottom panel in Figure 3 shows average wealth (or net worth), which we assume to be the same for both cohorts because wealth data in the PSID are limited. Despite the fact that couples enjoy economies of scale, the wealth of couples is over

²Appendix A details our computations and shows that the majority of men and women in both cohorts are married. Appendix M validates our labor market outcomes from the PSID with those from the Current Population Survey (CPS).

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




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

twice that of singles at all ages.

4 The model

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

During the **working stage**, single and married individuals choose how much to work and save and face wage shocks. Married people face divorce shocks, and single people might meet partners and get married. Wages are a function of one's human capital (which is endogenously accumulated while working) and are affected by shocks.

We model (and estimate) available time to be split between working and leisure

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

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

During the **early retirement stage**, people still experience wage shocks, but single people don't get married anymore and couples no longer divorce.⁴ If they claim Social Security, they can no longer work. Couples claim Social Security jointly.

During the first year of the **retirement stage**, those who have not already claimed Social Security do so and stop working. People face health shocks, out-of-pocket medical expenses, and mortality shocks. Thus, each married person faces the risk of his or her spouses's death, in addition to his or her own. Mortality risk and medical expenses depend on gender, age, health status, and marital status.

Given that we explicitly model labor participation and working hours of husbands and wives, savings, and medical expenses in old age, our model is computationally intensive (see Appendix E for more details). For tractability, we assume the following. First, marriage and divorce are exogenous processes that we estimate from the data. Second, the amount of time spent in home production is fixed. Third, fertility is exogenous, women have an age-varying number of children that depends on their age and marital status (which we estimate from the data), and we do not model child quality. Fourth, people who are married to each other are of the same age. Lastly, divorced and never-married people face the same problem (we thus abstract from Social Security benefits entitlement for divorcees). We discuss our model's limitations

⁴Only 1% of couples get divorced and 4% of singles get married between ages 62 and 72 in the HRS data for our 1945 cohort.

and perform robustness checks in Section 8.

4.1 Preferences

Let t be age, with people entering at age 25 and dying by age 99. For simplicity of notation, we write our model for one cohort; hence, t indexes both age and time for that cohort. We solve the model separately for the two cohorts and make sure that each has the appropriate time- and age-varying inputs.

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

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

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

where c_t is consumption; $\eta_t^{i,j}$ is its equivalent scale for couples and $\eta_t^{i,1}$ is the one for singles; and $l_t^{i,j}$ is leisure, which is given by

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

where $L^{i,j}$ is available time, net of home production, which can be different for single and married men and women. The functional form we use for it is

$$L^{i,j} = \frac{L}{1 + \exp(FL^{i,j})}, \quad (2)$$

where we normalize L to 112 hours a week and estimate $FL^{i,j}$ using our structural model. The term n_t^i is hours worked, and $I_{n_t^i}$ is an indicator function that equals 1 when hours worked are positive.

The term $\Phi_t^{i,j}$ is the fixed time cost of working, which depends on gender, marital status, and age. It assumes the following functional form, whose parameters we estimate using our structural model,

$$\Phi_t^{i,j} = \frac{\exp(\phi_0^{i,j} + \phi_1^{i,j}t + \phi_2^{i,j}t^2)}{1 + \exp(\phi_0^{i,j} + \phi_1^{i,j}t + \phi_2^{i,j}t^2)}. \quad (3)$$

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

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

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

4.2 Human capital and wages

We define human capital, \bar{y}_t^i , as one's average past earnings at each age (see Equation (12) for a formal definition). It is therefore a function of one's initial wages (and schooling, to the extent that it is reflected in one's wages) and subsequent labor market experience and wages, not just of experience measured as the amount of time one has previously worked. Our definition has two important benefits. First, it respects previous findings that the returns to experience depend on one's education and thus human capital and earnings (Blundell, Dias, Meghir, and Shaw (2016) and Costa Dias, Joyce, and Parodi (2018)). Second, it allows us to use only one state variable to keep track of both human capital and Social Security contributions, which keeps our framework manageable.

Wages have two components, a deterministic function of age, gender, and human capital, $e_t^i(\bar{y}_t^i)$, and a persistent 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 one's effective hourly wage per hour.

4.3 Marriage and divorce

During the working period, the probability that a single person gets married at the beginning of next period depends on age, gender, and wage shock: $\nu_{t+1}(\cdot) = \nu_{t+1}(i, \epsilon_t^i)$.

To allow for assortative mating, conditional on meeting a partner, the probability of meeting with a partner p with wage shock ϵ_{t+1}^p is

$$\xi_{t+1}(\cdot) = \xi_{t+1}(\epsilon_{t+1}^p | \epsilon_{t+1}^i, i). \quad (4)$$

We assume random matching over wealth a_{t+1} and average accumulated earnings of

the partner \bar{y}_{t+1}^p , conditional on the partner's wage shock. Thus, we have

$$\theta_{t+1}(\cdot) = \theta_{t+1}(a_{t+1}^p, \bar{y}_{t+1}^p | \epsilon_{t+1}^p). \quad (5)$$

A working-age couple can be hit by a divorce shock 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 their wealth equally (we experimented with different asset splits with very similar results). We abstract from alimony.

4.4 Costs of raising children and running a household

We keep track of both the total number of children and their age as a function of mothers' age and marital status. The term $f^{0,5}(i, j, t)$ is the number of children age 0-5, and $\tau_c^{0,5}$ is the child care cost for each child in that age group. Similarly, $f^{6,11}(i, j, t)$ is the number of children age 6-11, and $f^{0,5}(i, j, t)$ is the corresponding child care cost for each child. We use our structural model to estimate these costs.

4.5 Medical expenses and death

At age 66, people are endowed with a distribution of health that depends on their marital status and gender. After that, they face survival probabilities, medical expenses, and health shocks. 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 also 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. Survival probabilities $s_t^{i,1}(\psi_t^i)$ are a function of age, gender, marital status, and health.

4.6 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 (wealth, human capital, and wage shocks; in the case of couples we take the latter two for each of the spouses) for each of our two cohorts.

4.7 Government

Each cohort in our model faces the effective time-varying tax rates that it experienced in the data and that we estimate from the PSID. Like Benabou (2002), we adopt a functional form that allows for negative tax rates and thus incorporates the EITC. We allow our effective tax rates to depend on marital status, gender, and age for each cohort (and thus on time). Taxes paid are a function of total income Y :

$$T(Y, i, j, t) = (1 - \lambda_t^{i,j} Y^{-\tau_t^{i,j}}) Y. \quad (6)$$

The government uses a proportional payroll tax τ_t^{SS} , up to a Social Security cap \tilde{y}_t , to help finance old-age Social Security benefits, which are a function of average past earnings (or human capital, as discussed in Section 4.2). We also allow the payroll tax and the Social Security cap to change over time for each cohort as in the data. We thus assume that the tax changes were anticipated by the households. The insurance provided by Medicaid and SSI in old age is represented by a means-tested consumption floor, $\underline{c}(j)$, as in Hubbard, Skinner, and Zeldes (1995).⁵

4.8 Recursive formulation

We compute nine value functions for the following groups and stages of life.

4.8.1 The value function of working-age singles

The value function of a working-age single depends on his or her age t , gender i , wealth a_t^i , persistent earnings shock ϵ_t^i , and human capital \bar{y}_t^i :

$$W^s(t, i, a_t^i, \epsilon_t^i, \bar{y}_t^i) = \max_{c_t, a_{t+1}, n_t^i} \left(v^i(c_t, l_t, \eta_t^{i,1}) + \beta(1 - \nu_{t+1}(i)) E_t W^s(t+1, i, a_{t+1}^i, \epsilon_{t+1}^i, \bar{y}_{t+1}^i) + \beta \nu_{t+1}(i) E_t \left[\hat{W}^c(t+1, i, a_{t+1}^i + a_{t+1}^p, \epsilon_{t+1}^i, \epsilon_{t+1}^p, \bar{y}_{t+1}^i, \bar{y}_{t+1}^p) \right] \right), \quad (7)$$

subject to Equation (1) and

$$Y_t^i = e_t^{i,j} (\bar{y}_t^i) \epsilon_t^i n_t^i, \quad (8)$$

$$\tau_c(i, j, t) = \tau_c^{0,5} f^{0,5}(i, j, t) + \tau_c^{6,11} f^{6,11}(i, j, t), \quad (9)$$

⁵Borella, Nardi, and French (2018) discuss Medicaid rules and observed outcomes after retirement.

$$T(\cdot) = T(ra_t + Y_t, i, j, t), \quad (10)$$

$$c_t + a_{t+1} = (1 + r)a_t^i + Y_t^i(1 - \tau_c(i, j, t)) - \tau_t^{SS} \min(Y_t^i, \tilde{y}_t) - T(\cdot), \quad (11)$$

$$\bar{y}_{t+1}^i = (\bar{y}_t^i(t - t_0) + (\min(Y_t^i, \tilde{y}_t)))/(t + 1 - t_0), \quad (12)$$

$$a_{t+1} \geq 0, \quad (13)$$

$$n_t^i \geq 0. \quad (14)$$

If one remains single, the expectation of the value function for next period integrates over one's wage shock next period. If one gets married, it also integrates over the distribution of the partner's state variables. The value function \hat{W}^c is the person's discounted present value of utility once he or she is in a married relationship with someone with given state variables (see Appendix C). Equation (12) describes the evolution of human capital, measured as average accumulated earnings (up to the Social Security earnings cap \tilde{y}_t) and in which $t_0 = 25$.

4.8.2 The value function of singles during the early retirement stage

The recursive problem for someone who has claimed Social Security at age tr is

$$S^s(t, i, a_t^i, \bar{y}_r^i, tr) = \max_{c_t, a_{t+1}} \left(v^i(c_t, L^{i,j}, \eta_t^{i,1}) + \beta E_t S^s(t + 1, i, a_{t+1}^i, \bar{y}_r^i, tr) \right), \quad (15)$$

subject to equations (10), (13), and

$$Y_t = SS(\bar{y}_r^i, tr) \quad (16)$$

$$c_t + a_{t+1} = (1 + r)a_t + Y_t - T(\cdot). \quad (17)$$

The term $SS(\bar{y}_r^i, tr)$ is a function of the income that the single person earned during his or her working life, \bar{y}_r^i , and claiming age tr . Let $N^s(t, i, a_t^i, \epsilon_t^i, \bar{y}_t^i)$ denote the value function of a person during the early retirement period who has not yet claimed benefits.

$$N^s(t, i, a_t^i, \epsilon_t^i, \bar{y}_t^i) = \max_{c_t, a_{t+1}, n_t^i} \left(v^i(c_t, l_t^{i,j}, \eta_t^{i,1}) + \beta E_t V^s(t + 1, i, a_{t+1}^i, \epsilon_{t+1}^i, \bar{y}_{t+1}^i) \right), \quad (18)$$

subject to equations (1), (8), (10), (12), (13), (14), and

$$c_t + a_{t+1} = (1 + r)a_t^i + Y_t^i - \tau_t^{SS} \min(Y_t, \tilde{y}_t) - T(\cdot). \quad (19)$$

Let $V^s(t, i, a_t^i, \epsilon_t^i, \bar{y}_t^i)$ denote the value function for a person during the early retirement stage who has not yet claimed and who, at the beginning of each period, chooses whether to do so, where D_t^i is an indicator function for claiming.

$$V^s(t, i, a_t^i, \epsilon_t^i, \bar{y}_t^i) = \max_{D_t^i} \left((1 - D_t^i)N^s(t, i, a_t^i, \epsilon_t^i, \bar{y}_t^i) + D_t^i S^s(t, i, a_t^i, \bar{y}_t^i, t) \right). \quad (20)$$

4.8.3 The value function of retired singles

The value function of a retired single with health ψ_t^i , average realized lifetime earnings \bar{y}_r^i , and Social Security claiming age tr is

$$R^s(t, i, a_t, \psi_t^i, \bar{y}_r^i, tr) = \max_{c_t, a_{t+1}} \left(v^i(c_t, L^{i,j}, \eta_t^{i,1}) + \beta s_t^{i,j}(\psi_t^i) E_t R^s(t+1, i, a_{t+1}, \psi_{t+1}^i, \bar{y}_r^i, tr) \right), \quad (21)$$

subject to equations (10), (13), (16), and

$$B(a_t, Y_t, \psi_t^i, \underline{c}(j)) = \max \left\{ 0, \underline{c}(j) - [(1 + r)a_t + Y_t - m_t^{i,j}(\psi_t^i) - T(\cdot)] \right\} \quad (22)$$

$$c_t + a_{t+1} = (1 + r)a_t + Y_t + B(a_t, Y_t, \psi_t^i, \underline{c}(j)) - m_t^{i,j}(\psi_t^i) - T(\cdot) \quad (23)$$

$$a_{t+1} = 0, \quad \text{if } B(\cdot) > 0. \quad (24)$$

The term $B(a_t, Y_t, \psi_t^i, \underline{c}(j))$ represents old-age, means-tested government transfers (such as Medicaid and SSI), which ensure a minimum consumption floor $\underline{c}(j)$.

4.8.4 The value function of couples during the working period

The value function of a married couple at this stage depends on both partners' state variables, where 1 and 2 refer to gender and $j = 2$:

$$\begin{aligned}
W^c(t, a_t, \epsilon_t^1, \epsilon_t^2, \bar{y}_t^1, \bar{y}_t^2) = & \max_{c_t, a_{t+1}, n_t^1, n_t^2} \left(w(c_t, l_t^{1,j}, l_t^{2,j}, \eta_t^{i,j}) \right. \\
& + (1 - \zeta_{t+1}(\cdot)) \beta E_t W^c(t+1, a_{t+1}, \epsilon_{t+1}^1, \epsilon_{t+1}^2, \bar{y}_{t+1}^1, \bar{y}_{t+1}^2) \quad (25) \\
& \left. + \zeta_{t+1}(\cdot) \beta \sum_{i=1}^2 \left(E_t W^s(t+1, i, a_{t+1}/2, \epsilon_{t+1}^i, \bar{y}_{t+1}^i) \right) \right),
\end{aligned}$$

subject to equations (1), (8), (9), (12), and

$$T(\cdot) = T(ra_t + Y_t^1 + Y_t^2, i, j, t) \quad (26)$$

$$c_t + a_{t+1} = (1+r)a_t + Y_t^1 + Y_t^2(1 - \tau_c(2, 2, t)) - \tau_t^{SS}(\min(Y_t^1, \tilde{y}_t) + \min(Y_t^2, \tilde{y}_t)) - T(\cdot) \quad (27)$$

$$a_t \geq 0, \quad n_t^1, n_t^2 \geq 0. \quad (28)$$

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

4.8.5 The value function of couples during the early retirement period

The recursive problem for couples that have claimed Social Security at age tr is

$$S^c(t, a_t, \bar{y}_r^1, \bar{y}_r^2, tr) = \max_{c_t, a_{t+1}} \left(w(c_t, L^{1,j}, L^{2,j}, \eta_t^{i,j}) + \beta E_t S^c(t+1, a_{t+1}, \bar{y}_r^1, \bar{y}_r^2, tr) \right), \quad (29)$$

subject to equations (10), (17), (13), and

$$Y_t = \max \left\{ (SS(\bar{y}_r^1, tr) + SS(\bar{y}_r^2, tr)), \frac{3}{2} \max(SS(\bar{y}_r^1, tr), SS(\bar{y}_r^2, tr)) \right\} \quad (30)$$

The variable Y_t represents the Social Security spousal benefit: married people receive either their own benefit or half of their spouse's, depending on which amount is higher.

The value function of a couple that has not yet claimed benefits is

$$N^c(t, a_t, \epsilon_t^1, \epsilon_t^2, \bar{y}_t^1, \bar{y}_t^2) = \max_{c_t, a_{t+1}, n_t^1, n_t^2} \left(w(c_t, l_t^{1,j}, l_t^{2,j}, \eta_t^{i,j}) + \beta E_t V^c(t+1, a_{t+1}, \epsilon_{t+1}^1, \epsilon_{t+1}^2, \bar{y}_{t+1}^1, \bar{y}_{t+1}^2) \right), \quad (31)$$

subject to equations (1), (8), (12), (26), (28), and

$$c_t + a_{t+1} = (1+r)a_t + Y_t^1 + Y_t^2 - \tau_t^{SS}(\min(Y_t^1, \tilde{y}_t) + \min(Y_t^2, \tilde{y}_t)) - T(\cdot). \quad (32)$$

The value function of a married couple during the early retirement stage that has not yet claimed Social Security benefits is

$$V^c(t, a_t, \epsilon_t^1, \epsilon_t^2, \bar{y}_t^1, \bar{y}_t^2) = \max_{D_t} \left((1 - D_t) N^c(t, a_t, \epsilon_t^1, \epsilon_t^2, \bar{y}_t^1, \bar{y}_t^2) + D_t S^c(t, a_t, \bar{y}_t^1, \bar{y}_t^2, t) \right). \quad (33)$$

4.8.6 The value function of couples during retirement

During this stage, the married couple's recursive problem ($j = 2$) depends on each spouse's health shocks ψ_t^i , and there are survival shocks $s_t^{i,2}(\psi_t^i)$. We assume that the health shocks of each spouse are independent of each other and that the death shocks of each spouse are also independent of each other. The value function is

$$R^c(t, a_t, \psi_t^1, \psi_t^2, \bar{y}_r^1, \bar{y}_r^2, tr) = \max_{c_t, a_{t+1}} \left(w(c_t, L^{1,j}, L^{2,j}, \eta_t^{i,j}) + \beta s_t^{1,j}(\psi_t^1) s_t^{2,j}(\psi_t^2) E_t R^c(t+1, a_{t+1}, \psi_{t+1}^1, \psi_{t+1}^2, \bar{y}_r^1, \bar{y}_r^2, tr) + \beta s_t^{1,j}(\psi_t^1) (1 - s_t^{2,j}(\psi_t^2)) E_t R^s(t+1, 1, a_{t+1}, \psi_{t+1}^1, \bar{y}_r^1, tr) + \beta s_t^{2,j}(\psi_t^2) (1 - s_t^{1,j}(\psi_t^1)) E_t R^s(t+1, 2, a_{t+1}, \psi_{t+1}^2, \bar{y}_r^2, tr) \right), \quad (34)$$

subject to equations (10), (13), (24), (30), and

$$\bar{y}_r^i = \max(\bar{y}_r^1, \bar{y}_r^2), \quad (35)$$

$$B(a_t, Y_t, \psi_t^1, \psi_t^2, \underline{c}(j)) = \max \left\{ 0, \underline{c}(j) - [(1+r)a_t + Y_t - m_t^{1,j}(\psi_t^1) - m_t^{2,j}(\psi_t^2) - T(\cdot)] \right\}, \quad (36)$$

$$c_t + a_{t+1} = (1+r)a_t + Y_t + B(a_t, Y_t, \psi_t^1, \psi_t^2, \underline{c}(j)) - m_t^{1,j}(\psi_t^1) - m_t^{2,j}(\psi_t^2) - T(\cdot). \quad (37)$$

Survivors collect benefits based on the higher amount between their own contributions and those of their deceased spouse (Equation (35)).

4.8.7 The value functions of individuals in couples

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

5 Estimation

We estimate our model on our two birth cohorts separately by adopting a two-step estimation strategy (like that of Gourinchas and Parker (2002)). In the first step, 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.

In the second step, we use the MSM. For the 1945 cohort, we normalize the time endowment for single men and estimate 19 model parameters $(\beta, \omega, (\phi_0^{i,j}, \phi_1^{i,j}, \phi_2^{i,j}), (\tau_c^{0,5}, \tau_c^{6,11}), L^{i,j})$. For the 1955 cohort, we assume that it has the same discount factor β and weight on consumption ω as the 1945 cohort, and thus we estimate 17 parameters. The data that inform the estimation of the parameters of our model are composed of the following 448 moments for each cohort:

1. Labor market participation of married and single men and women age 25-65
2. Hours worked, conditional on working, for married and single men and women age 25-65
3. Wealth for couples and single men and women age 26-65

Appendixes B, D, and F discuss our first-step and second-step inputs in detail.

5.1 First-step estimation

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

Table 1: First-step estimated inputs summary.

Table 1 summarizes our first-step estimated model inputs.




Figure 4: Average wage. Left: 1945 cohort. Right, 1955 cohort. PSID.

Wages. Figure 4 displays the average wage profiles implied by our estimated processes. Consistent with the evidence on the marriage premium, the wages of married men are higher than those of single men. In contrast, the wages of married women are lower than those of single women in our 1945 cohort. Across the two cohorts, the gender gap shrinks because the average wage of married women increases.

This is due to a combination of different returns to human capital and human capital accumulation. The stagnation of men’s wages across cohorts is consistent with the findings by Acemoglu and Autor (2011) and Roys and Taber (2017) on wages over time. Table 3 in Appendix B reports our estimates for the earnings shock processes.⁶

Marriage, divorce, spousal wage shocks, spousal wealth, and Social Security benefits. Men with higher wage shocks are more likely marry and less likely to divorce. In contrast, these shocks have little effect on the marriage and divorce probabilities of women. In addition, both the probability of marrying and that of divorcing are smaller for our 1955 cohort (Figures 1 and 2 in Appendix B). Our estimated correlation of the logarithm of initial wage shock between spouses is 0.22 at age 25-34, 0.36 at age 35-44, and 0.42 after age 45.⁷ Because of these initial correlations and the high wage shock persistence that we estimate, spouses also have positively correlated shocks after getting married. Appendix B reports spousal wealth and Social Security earnings.

Children. The number of children decreases for married women and, to a smaller extent, for single women from the 1945 to the 1955 cohort (Figure 5 in Appendix B).

Health, mortality, and medical expenses. Because of a lack of old-age data for the 1955 cohort, we assume that it and the 1945 one face the same risks after age 65. We find that women, married people, and healthy people have longer life expectancies (Figure 7 in Appendix B) and that medical expenses climb fast past age 85 and are highest for single and unhealthy people (Figure 8 in Appendix B).

5.2 Second-step estimation

Table 2 presents our estimated preference parameters for both cohorts (see Appendix G for more details). For the 1945 cohort, our estimated discount factor is 0.990, the same value estimated by De Nardi, French, and Jones (2016) for elderly retirees, and our estimated weight on consumption is 0.406. We assume that the 1955 cohort shares these preference parameters.

We normalize the time endowment of single men to 5,840 hours a year and 112 hours a week. For our 1945 cohort, single women have a weekly time endowment of 107 hours a week. We interpret this as their having to spend five more hours a week

⁶We discuss the role of education in Appendix B.

⁷We assume this correlation is the same for both cohorts, because the number of new marriages after age 25 is small during this time period.

Estimated parameters	1945 cohort	1955 cohort
β : Discount factor	0.990	As 1945 cohort
ω : Consumption weight	0.407	As 1945 cohort
$L^{2,1}$: Time endowment (weekly hours), single women	107	112
$L^{1,2}$: Time endowment (weekly hours), married men	107	100
$L^{2,2}$: Time endowment (weekly hours), married women	88	88
$\tau_c^{0,5}$: Prop. child care cost for children age 0-5	28%	22%
$\tau_c^{6,11}$: Prop. child care cost for children age 6-11	7%	19%
$\Phi_t^{i,j}$: Partic. cost	Fig. 10 App. G	Fig. 10 App. G

Table 2: Second-step estimated model parameters.

managing their household, rearing children, and taking care of elderly parents. The time endowments for married men and women are 107 and 88 hours. This implies that people in the latter two groups spend five hours and 24 hours a week in home production activities, respectively. These estimates are remarkably similar to those from time diary data in Aguiar and Hurst (2007) and Dotsey, Li, and Yang (2014).

Our estimates for the 1945 cohort imply that child care costs for each child age 0-5 and 6-11 are 28% and 7% of a woman’s earnings, respectively. The PSID data report child care costs only for all children in the 0-11 age range. They correspond to 31% and 20% of married women’s earnings at ages 25 and 30, respectively. These numbers line up well with those from our model, which are 22% and 17%, respectively.

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

Figure 10 in Appendix G reports the age-varying time costs of working by age, expressed as fraction of the time endowment of single men. Our estimated participation costs are relatively high when people are younger and increase again after 45, except for those for single men. The time costs of going to work might include factors

other than commuting time. For instance, they might be higher when children are youngest, because during that period, parents might need additional time to transport their children back and forth from day care. They also show that conditional on all aspects of our environment, the participation costs of married women are the lowest ones because married women face lower wages, have a smaller time endowment (because of time spent engaging in home production and child care), and tend to have higher-wage husbands who work. Table 11 in Appendix G reports all of our structural model estimates and their standard errors.

5.3 Model fit

Figures 5 and 6 report our model-implied moments and the data moments and 95% confidence intervals from the PSID for our 1945 cohort (for brevity, we report the fit for the 1955 cohort in Appendix H). Our model fits the targeted data well for both cohorts, which is remarkable given that it is tightly parameterized: we have 448 targets for each cohort, and we estimate 19 parameters for the 1945 cohort and 17 for the 1955 one.

5.4 Identification





In this section, we briefly discuss why our target moments allow us to identify the 19 model parameters that we estimate (see Table 2) for our 1945 cohort.⁸

The discount factor β is the only parameter that we estimate that has a large effect on the savings of both couples and singles over much of their life cycle (compare Figure 14, in which we change it, with all of the other figures in Appendix I, which change the other parameters). In contrast, it has relatively small effects on both participation and hours over the life cycle for our four subgroups.

The weight on consumption ω (Figure 15) is the only parameter that we estimate that affects the participation and hours, conditional on participation, for all four of our subgroups (single and married men and women). It also has overall minor effects on savings.





Our estimated available time for married men and women and for single women (we normalize available time for single men) mostly affect the group they apply to.

⁸Given that we estimate only 17 parameters for the 1955 cohort and that we match the same targets, the parameters for the 1955 cohort are identified as well.



(a) Participation, couples

(b) Participation, singles



(c) Hours for workers, couples

(d) Hours for workers, singles

Figure 5: Model-implied participation and hours, and average and 95% confidence intervals from the PSID.




Figure 6: Model-implied wealth, and average and 95% confidence intervals from the PSID.

More specifically, changing the available time of single women mostly affects their hours, but it does not have much effect on their participation (Figure 16) and has no discernible effects on savings. Reducing the available time of married men (Figure 17), who are the breadwinners for these cohorts, marginally decreases their participation later in life, decreases their hours, and reduces couples' savings, but not the savings of singles. Interestingly, it also increases the participation of married women over all of their life cycle but has no effects on the hours worked by married women, conditional on participation. Reducing instead the available time of married women (Figure 18), reduces their participation and hours and the participation of married men in middle and late age.

Child care costs mostly affect the participation of the married women, who have more of them. As younger women tend to have younger children, child care costs for children ages 0-5 mostly affect the participation of very young married women (until age 35, Figure 19). In contrast, child care costs for children ages 6-11 have relatively more of an effect on older married women (Figure 20). Both types of child care costs have little effect on the rest of our target moments.

The remaining set of parameters that we estimate governs a quadratic function affecting the fixed time cost of working (as a function of age) for each of our four groups (see Equation 3). Thus, it comprises 12 additional parameters. We also plot the implied fixed cost of working for each parameter change.

Decreasing the fixed-time costs of working for single men (Figure 21) increases their participation after age 40 and hours worked throughout their working period,

especially at younger ages. Doing the same for single women goes in the same direction, but single women start increasing their participation at younger ages and adjust by less at each age. For single women, the comparison between Figure 16 and Figure 22 reveals that while the fixed cost of working affects both participation and hours, available time mainly affects their hours worked.

If we turn to married men, Figure 23 shows that decreasing their fixed-time cost of participation increases their participation after age 40 and their hours worked over all of their life cycle. In contrast, changing available time for married men affects their hours much more than their participation. Decreasing the fixed time cost of participation for married women (Figure 24) increases their labor market participation at all ages but has no effects on hours worked conditional on participation. The comparison with Figure 18 reveals that, in contrast, changing the available time of married women changes both their participation and hours worked, conditional on participation.

The remaining set of figures changes the linear term or the quadratic term fixed time cost of working with age. Consistent with our previous set of results, these changes mostly affect participation rather than hours. In addition, the linear term has a relatively stronger effect in middle age than the constant term, while the quadratic term has the strongest effect on labor market participation closer to retirement age.

Table 13 in Appendix I reports our GMM criterion function disaggregated by group and outcome for our benchmark model and for the experiments that change one parameter at a time. It shows that the GMM criterion rises along all of these deviations. This and many other checks that we have run during the estimation stage tell us that our model's parameters are well identified given the target moments that we have chosen.

6 Model validation

We start by analyzing the joint distribution of married couples' wages, because within-couple relative wages play an important role in generating responses to policy changes. Table 3 shows that the correlation of the wages of working husbands and wives by age group for the 1945 cohort generated by the model broadly replicates that found in the data, with the correlation in the model actually being higher than that in the data. Table 12 in Appendix H shows that our model also matches well

the joint distribution of husbands' and wives' wages by wage tercile.

Age group	Model	Data
25 – 34	0.415	0.338
35 – 44	0.376	0.314
45 – 54	0.385	0.300

Table 3: Correlation of husbands and wives' wages at various ages, workers, model and data, 1945 cohort.




Figure 7: Participation (left) and working hours (right) for people in couples. Model (solid) and PSID (dashed). 1945 cohort.

Figure 7 compares additional model implications for couples with those in the data for our 1945 cohort. It shows that our estimated model also matches the fraction by age of couples with two workers, with only the husband working, with only the wife working, or with none working. It also highlights that our model produces reasonable implications for the hours worked over the life cycle for each of these type of couples.

To further build confidence in our model's responses to policy changes, we also evaluate its labor supply elasticities. Table 4 shows the elasticities of participation and hours among workers with respect to an anticipated change to their own wage.⁹ It shows that the elasticity of participation of women is larger than that of men,

⁹For this computation, we temporarily increase the wage by 5% for only one age and one group (married men, married women, single men, or single women) at a time. While we do not compensate this wage change, a temporary change in wage of this size is very small compared with a family's lifetime earnings and consumption. For instance, at age 30, it corresponds to 0.04% of average discounted family income. The corresponding numbers for temporary wage increases at ages 40, 50, and 60, are 0.04, 0.03, and 0.02%, respectively.

that married men have the lowest elasticity of participation, and that the elasticity of participation for all groups is largest around retirement age (a finding that confirms that of French (2000) for men). Importantly, our elasticities are consistent with those in Liebman, Luttmer, and Seif (2009), Blundell and Macurdy (1999), and Attanasio, Levell, Low, and Sánchez-Marcos (2018).

	Participation				Hours among workers			
	Married		Single		Married		Single	
	W	M	W	M	W	M	W	M
30	1.0	0.0	0.5	0.2	0.2	0.3	0.4	0.4
40	0.7	0.1	0.4	0.2	0.4	0.5	0.4	0.5
50	0.6	0.2	0.4	0.5	0.4	0.5	0.8	0.5
60	1.1	0.8	1.8	1.4	0.3	0.3	0.5	0.4

Table 4: Model-implied elasticities of labor supply.

However, these wage elasticities to a small and transitory wage change are not necessarily indicative of how participation and hours would change as a result of a permanent wage change, which is more relevant given our policy experiments. The left graph of Figure 8 shows that a permanent 5% increase in the wage of married women leads to large increases in married women’s participation, with participation being 3-7 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 human capital) upon marriage. Married men’s participation after age 40 decreases when women’s wages increase, a finding that confirms the importance of modeling the labor supply of both spouses when studying couples. The right panel translates these changes in participation to elasticities by taking the ratio of the percentage change of labor force participation at each age divided by percentage change of the wage at the same age. We find an implied elasticity of participation for married women that is U-shaped and peaks at 2.8 at age 25. This panel also reports the cross-elasticities of the other groups to changes in the wages of married women.




Figure 8: Change in participation (left) and implied participation elasticities (right) for a 5% permanent increase in the wage schedule of married women. Model’s implications.

7 Eliminating marital provisions

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

7.1 Outcomes

For each policy counterfactual, we compute two sets of results. The first one balances the government budget constraint by adjusting the proportional component of the income tax, while the second one keeps it unbalanced. The comparison of the first and second set of results (we report the latter in Appendix K) shows that most of the changes in labor supply and savings that take place as a result of a given reform do not come from our adjustment of the proportional component of the labor tax.

7.1.1 Eliminating spousal Social Security benefits, 1945 cohort

According to the current Social Security rules, a spouse can receive half of his or her partner’s contribution while the partner is alive and all of the benefits after the spouse’s death. This provision holds back the labor supply of the secondary earner, given that he or she can benefit from spousal benefits. It also encourages the labor supply of the main earner, who also works to provide Social Security benefits to the secondary earner. In addition, it reduces retirement savings by raising the annuitized income flow of the secondary earner or non-participant.

When eliminating both spousal Social Security benefits, the government runs a budget surplus and can cut the proportional component of the income tax from 4.0% to 1.8%. The top left panel of Figure 9 shows that without spousal Social Security benefits, the participation of married women is, respectively, 9, 11, and 4 percentage points higher at ages 25, 55-60, and 65. In contrast, married men decrease their participation starting at age 55, and their participation is 6 percentage points lower by age 65. The participation of single women at ages 25-30 increases by 3 percentage points, because should they marry, they now expect no Social Security benefits coming from their spouse's labor supply. As single women age, the probability that they marry becomes negligible, and the effect of the elimination of spousal benefits on their participation fades.




Figure 9: Changes in participation as a result of the elimination of various marriage-related policies under government budget balance.

An important reason why this reform has such large effects on the labor supply of married women lies in the initial distribution of potential wages of men and women at age 25. The top panel of Table 5 shows that in the 1945 cohort, 60% of women and only 20% of men belong to the bottom two quintiles of wages at age 25. Thus, most women have low wages and tend to be secondary earners in this cohort. For

this reason, they react strongly to the elimination of spousal benefits.

	Wage quintile				
	1	2	3	4	5
1945 cohort					
Men	7.9 %	12.3 %	21.7 %	29.1 %	29.1 %
Women	32.5 %	27.9 %	18.3 %	10.7 %	10.7 %
1955 cohort					
Men	10.1 %	14.3 %	20.0 %	25.5 %	30.1 %
Women	28.5 %	24.9 %	20.2 %	15.1 %	11.4 %

Table 5: Distribution of men and women at age 25 across wage quintiles, 1945 (top) and 1955 cohort (bottom). PSID data.

Groneck and Wallenius (2020) and Kaygusuz (2015) study the effects of marital Social Security benefits in simpler models than ours, in which, for instance, men cannot change their labor supply and women can do so to a limited extent.¹⁰ These two papers report that over all of the working period, their model implies an increase in the participation by married women of 12.2 and 6.1 percentage points, respectively. Hence, we find effects in the same ballpark.

Married women work more and accumulate more human capital. As a result, their labor income is over 10% higher over all of their working period. Row 1 in Table 6 shows the resulting changes in wealth at retirement time. The reform increases savings by reducing government payments to spouses and widows during retirement, and wealth at retirement goes up by 14.6%, 8.7%, and 11.2% for couples, single men, and single women, respectively.

	Couples	Single men	Single women
1945 cohort, no spousal Social Security benefits	14.6%	8.7%	11.2%
1945 cohort, no marital-related policies	19.5%	10.1%	15.4%
1955 cohort, no marital-related policies	17.9%	10.0%	13.7%

Table 6: Change in wealth at age 66, in percentages, as a result of each policy change when the income tax is reduced to balance the government budget.

¹⁰Their models are also less rich along other important dimensions and are calibrated rather than estimated.

7.1.2 Eliminating joint income taxation, 1945 cohort

The top right panel of Figure 9 displays the effects on participation of having everyone file as singles (the married men file as single men and the married women as single women) and reducing the proportional component of the income tax to balance the government budget (from 4.0% to 3.4%). As a result of this policy, the participation of married women increases by more than 20 percentage points until age 35 and by 10 percentage points between ages 45 and 60. The participation of single women also increases slightly until age 60. Figure 1 provides the key intuition for this result: the marginal tax rates for working married women are much lower when they do not file jointly with their husbands. Guner, Kaygusuz, and Ventura (2012) study the switch from current U.S. taxation to single filer taxation in a calibrated model of a steady state and find that the labor supply of married women goes up by 10-20 percentage points. Our effects on the labor supply of married women are thus close to theirs.

7.1.3 Eliminating spousal benefits and joint taxation, 1945 cohort

This policy change implies a reduction in the proportional component of the income tax from 4.0% to 1.8%. The bottom left panel of Figure 9 shows that as a result of this reform, the participation of married women is 15-28 percentage points higher until age 62 and that of single women is about 5 percentage points higher until age 40. The participation of married men is higher in middle age, reaching a peak of 2 percentage points higher than in the benchmark, but at age 65 is 7 percentage points lower than in the benchmark. Thus, the timing of their participation changes over their life cycle. The effects of increased participation and human capital accumulation imply that average yearly earnings at age 45 are \$6,000 and \$3,000 higher for married and single women, respectively. Average earnings of married men are similar until about age 55 and then drop sharply (see left panel of Figure 33 in Appendix J). Row 2 in Table 6 displays the effects on savings at retirement time: couples now save 19.5% more for retirement, while single men and women save 10.1% and 15.4% more.

7.1.4 Eliminating spousal benefits and joint taxation, 1955 cohort

For brevity, we report results only for the 1955 cohort for the case in which we eliminate all three marriage-related provisions at the same time. The bottom right

panel of Figure 9 shows that the effects of these policies on a relatively younger cohort with a much higher degree of participation of married women continue to be very large. The effects of increased labor market experience and participation on income are a bit smaller but overall similar to those in the 1945 cohort (left panel of Figure 33 in Appendix J). Row 3 in Table 6 displays the effects on savings at retirement time. Couples now save 17.9% more for retirement, while single men and women save, respectively, 10.0% and 13.7% more. Comparing the top and bottom panel of Table 5 highlights that in the 1955 cohort, the fraction of women in the lowest wage quintile has decreased, and the fraction of women in the highest wage quintile has increased. However, it is still the case that in the 1955 cohort, most women tend to have lower wages and be secondary earners, and hence they respond strongly to the elimination of marital provisions.

7.2 Welfare

To evaluate welfare changes, we calculate the one-time asset compensation required for each household at age 25 to stay in the benchmark economy and report it as a fraction of average income. Thus, negative asset compensations mean that households are better off in the benchmark economy. Table 7 reports the average welfare gains or losses conditional on one’s marital status at age 25, the fraction of households gaining and losing, and average gains and losses in each of these groups.

The top panel refers to the 1945 cohort. The first set of results refers to the economy in which there are no marital Social Security benefits and taxes are unchanged. On average, couples need to be compensated by an asset transfer that is equivalent to 0.24 times average income in the economy, while single women require 0.20 times average income (they expect to marry and lose these benefits after marriage). These welfare costs are not very large, because these reforms are anticipated as of age 25 and people have many years to work and save to make up for the loss in benefits. Single men benefit from this policy because their future wives work harder and earn more, and they do not take into account their future wife’s disutility of doing so. The remaining columns in the table distinguish the effects for winners and losers. The three “winners” columns show that all couples and single women lose from this policy, while all single men gain from it.

The second set of results removes marital Social Security provisions and balances

	All			Winners			Losers		
	Couples	SW	SM	Couples	SW	SM	Couples	SW	SM
1945 cohort									
(1). Remove Social Security spousal benefits, unbalanced budget									
Average	-0.24	-0.20	0.25	0.00	0.00	0.25	-0.24	-0.20	-0.02
Percentage				0.0	0.0	100.0	100.0	100.0	0.0
(2). Remove Social Security spousal benefits, balanced budget									
Average	0.66	0.19	1.15	0.66	0.20	1.15	0.00	-0.03	0.00
Percentage				100.0	92.5	100.0	0.0	7.5	0.0
(3). Remove joint income taxation, unbalanced budget									
Average	0.06	-0.18	0.81	0.29	0.06	0.81	-0.19	-0.19	0.00
Percentage				52.8	4.9	100.0	47.2	95.1	0.0
(4). Remove joint income taxation, balanced budget									
Average	0.31	-0.08	1.06	0.42	0.12	1.06	-0.09	-0.13	0.00
Percentage				78.8	20.6	100.0	21.2	79.4	0.0
(5). Remove all marital related policies, balanced budget									
Average	0.80	0.05	1.97	0.80	0.33	1.97	-0.03	-0.12	0.00
Percentage				98.8	37.4	100.0	1.2	62.6	0.0
1955 cohort									
(6). Remove all marital related policies, balanced budget									
Average	0.73	0.21	1.14	0.74	0.30	1.14	-0.04	-0.04	-0.03
Percentage				98.2	74.1	100.0	1.8	25.9	0.0

Table 7: Asset compensation required for staying in the benchmark economy, normalized as a fraction of average income. SM: single men, SW: single women. Top line for each experiment: average welfare gain or loss. Bottom line for each experiment: fraction in that group gaining or losing welfare.

the government budget by reducing the proportional component of the income tax from 4.0% to 1.8%. The first three columns display large welfare gains: couples are willing to pay, on average, an asset amount that corresponds to 0.66 times average income, single women 0.19 times, and single men 1.15 times. The second set of three columns shows that all couples, 92.5% of single women, and all single men benefit from these changes. The last set of three columns shows that 7.5% of the single women who lose face very small losses. These are women whose initial human capital is very low and who rely on marital benefits. Thus, this counterfactual suggests that eliminating these benefits while reducing the income tax benefits the vast majority of the young population and has small welfare costs for a minority of single women.

The third set of results makes everyone file as single people without balancing the government budget constraint and shows that the willingness to pay for this policy equals 0.06 and 0.81 times average income for couples and single men, respectively. In contrast, single women require an asset compensation of 0.18 times average income. This happens because they will be working more in the future, especially after marriage. The winners and losers columns reveal that 52.8% of couples, 4.9% of single women, and 100% of single men favor this policy and that both gains for the winners and losses for the losers are sizable.

The fourth set of results balances the government budget constraint by reducing the income tax (from 4.0% to 3.4%) and generates more winners with larger welfare gains and fewer losers, with smaller welfare losses than in the previous experiment. For instance, 78.8% of the couples are willing to give up an asset amount corresponding to 0.42 times average income to live under this policy, while the compensation for the remaining 21.2% amounts to 0.09 times average income.

The fifth set of results for the 1945 cohort eliminates all marriage-related policies and balances the government budget constraint by reducing the income tax from 4.0% to 1.8%. This policy generates the largest aggregate welfare gains among the set that we consider for the 1945 cohort: 0.80, 0.05, and 1.97 times average income for couples, single women, and single men, respectively. Among couples, 98.8% gain, compared with 37.4% of single women and 100% of single men. The bigger losers coming out of this policy are 62.6% of single women, who lose, on average, 0.12 times average income.

The results in the bottom panel refer to the 1955 cohort. They show that not only there are large aggregate gains from removing marriage-related provisions and reducing the income tax but also that single women in this cohort are less hurt by this policy than those in the 1945 cohort. In fact, only 25.9% of them lose, compared with 62.6% in the older cohort. In addition, their loss is much smaller (0.04 times average income, compared with 0.12 in the 1945 cohort). In both cohorts, only a minority of couples lose, and those who do experience a small welfare loss.

Overall, our policy experiments indicate that removing marriage-related taxes and Social Security benefits would increase female labor supply and the welfare of the majority of the population, whereas the rest would bear only small welfare costs. On this point, it is important to note that welfare calculations are usually sensitive to all aspects of the environment. In our specific case, they could be especially sensitive

to modeling the couple as a collective, rather than unitary, decision making unit. We leave this very interesting extension for future work.

8 Limitations, robustness, and future research

To keep our model tractable, we have made several important assumptions. Marriage, divorce, and home production patterns are estimated from the data and do not change with policy. Fertility is exogenous, and we do not model children's human capital accumulation. Couples behave according to the unitary model. Health, life expectancy, and medical expenses do not depend on one's earnings.

Before we start discussing our findings' robustness to changes in marriage and divorce, it is worth pointing out that the empirical literature for the United States shows that the effects of marital provisions on marriage and divorce behavior coming from Social Security (Dickert-Conlin and Meghea (2004), Goda, Shoven, and Slavov (2007), and Dillender (2016)) and income taxes (Alm and Whittington (1995, 1997, 1999)) are tiny. Turning to welfare programs, Low, Meghir, Pistaferri, and Voena (2018) document no effects on marriage and fertility and a decline in divorce rates, from 0.9% to 0.7%, as a result of the 1996 reform. Persson (2020) finds that in Sweden, the divorce rate increased by 10% as a result of the elimination of marital survivorship benefits.

We now turn to showing that changes of this magnitude in marriage and divorce patterns have little effect on our results. To do so, we compare the effects of eliminating marital benefits for the 1945 cohort in our benchmark economy with those of two alternative scenarios. In the first one, the policy decreases marriage rates by 20% and increases divorce rates by 20%. In the second one, the policy increases marriage rates by 20% and decreases divorce rates by 20% (in both, we also balance the government budget). These changes are large compared with those that have been documented in the empirical literature. Figure 10 shows that changes in participation are similar whether marriage and divorce patterns change or not. Comparing the middle panel (no marriage and divorce changes) and the left panel highlights that a reform that lowers the probability of marriage and raises the probability of divorce makes women more self-reliant on their own labor supply. Married women work more to hedge against divorce risk. Single women are less likely to marry, and they also work more. Both groups accumulate more human capital as a result. Comparing the




Figure 10: Changes in participation as a result of the elimination of all marriage-related policies when marriage and divorce patterns also change with policy. Government budget balance, 1945 cohort.

middle and right panel shows that the opposite changes (increasing the marriage rate and lowering the divorce rate) have the opposite effect on participation, but that this effect is small. The effects on savings at retirement time are also robust to changes in expected marital patterns.

Because we fix the amount of time that single and married men and women have to split between work and leisure, we also implicitly fix the amount of time that they spend in home production. To evaluate the effects of this assumption on our results, in addition to changing policy, we exogenously change available time. More specifically, we assume that as women’s incentives to work more on the labor market increase, women make more time available for labor and leisure. In this robustness check, as all policies change, the available time of married women increases by three hours a week. Hence, if they decide to work longer hours, they also have to pay for more child-care. Figure 11 shows that the increase in female labor supply would be even larger in this case. Table 8 shows that savings by couples, which make up

	Couples	Single men	Single women
Benchmark	19.5%	10.1%	15.4%
Reduce home production	21.5%	10.8%	16.8%

Table 8: Change in wealth at age 66 after removing all marriage-related policies when available time increases by 3 hours each week as policy changes. Government budget balance, 1945 cohort.

the vast majority of the population, also go up in this robustness check, and that they do so by a little more than in our benchmark policy experiment. While this




Figure 11: Changes in participation as a result of the elimination of all marriage-related policies when available time increases by 3 hours each week as policy changes. Government budget balance, 1945 cohort.

is a very stylized and suggestive experiment, if we were to add a cost for women to outsource some of their home production (for instance to pay a cleaning person, in addition to paying for additional child care), those who choose not to pay this cost would exhibit an increase in labor market participation very similar to that in our benchmark, because their available time would not change. Those who choose to pay this cost would have a negative wealth effect as a result and would increase their labor at least as much as they do in our robustness exercise. Thus, the increase in female labor market participation would be at least as large as in our benchmark policy experiment.

Because the importance of Social Security depends on one’s life expectancy, we also perform a robustness check, in which we allow health, survival, and medical expenses to depend on one’s average accumulated earnings (or human capital) at retirement time. That is, we re-estimate our model inputs and structural model for this case and re-run our policy experiments. Appendix L shows that the results are very similar in this case too.

Our behavioral responses to marital policy changes have thus proven to be robust to all of the checks that we have performed. Turning to their welfare implications, if we were to give households more choices—including, for instance, time spent in home production—we would provide more levers for households to optimize over. Hence, it should be the case that in the presence of home production, the welfare benefits from policy reforms should be no smaller (and possibly larger) in the presence of more choices than in our benchmark economy.

While our results are robust to marriage and divorce patterns and changes in home production activities, exploring marriage and divorce decisions is an important avenue for understanding the decisions of couples more generally. Furthermore, we do not model investment in children’s human capital, marital happiness, and the collective model of the households, which are also very important and promising avenues that we leave for future research.

9 Conclusions

We estimate a model of labor supply and savings for single and married people. It allows for a rich representation of the risks that people face over their entire life cycle and for the important provisions of taxes and Social Security benefits for couples and singles. We do so for both the 1945 and the 1955 birth cohorts, and we show that our model fits the data very well, including along important dimensions that we do not match by construction.

We use our model to evaluate the effects of marriage-based Social Security benefits and taxation. We find that these marriage-based provisions act as a strong disincentive on the labor supply of both married women and single women. This lower participation also reduces their labor market experience, which, in turn, reduces their wages over all of their life cycle. These provisions also induce married men to work longer, and they depress the savings of couples.

Our findings are very similar for the 1945 and 1955 birth cohorts, even though the labor market participation of young married women in the 1955 cohort is over 10 percentage points higher than that of the 1945 cohort. They thus suggest that these marital provisions have large negative effects on female labor supply, including in recent cohorts. If the government surplus resulting from the elimination of marriage-related provisions were used to lower income taxation, there would be large welfare gains for the vast majority of the population, and the few losing would experience small welfare losses.

Our paper provides several contributions. First, it is the first paper to study all marriage-related taxes and benefits in a unified framework. Second, it does so by allowing for the large observed changes in the labor supply of married women over time by studying two different cohorts. Third, 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 dynamic framework with savings. Fourth, our framework is very rich along other dimensions that are important in the study of our problem, including looking at how labor market experience affects wages and carefully modeling survival, health, and medical expenses in old age, and their heterogeneity by marital status and gender.

References

- Acemoglu, D. and D. Autor (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. In O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics*, Volume 4, pp. 1043–1171. Elsevier.
- Aguiar, M. and E. Hurst (2007). Measuring Trends in Leisure: The Allocation of Time Over Five Decades. *Quarterly Journal of Economics* 122(3), 969–1006.
- Alm, J. and L. A. Whittington (1995). Income Taxes and the Marriage Decision. *Applied Economics* 27(1), 25–31.
- Alm, J. and L. A. Whittington (1997). 'Til Death or Taxes Do Us Part: The Effect of Income Taxation on Divorce. *Journal of Human Resources* 32(2), 388–412.
- Alm, J. and L. A. Whittington (1999). For Love or Money? The Impact of Income Taxes on Marriage. *Economica* 66, 297–316.
- Attanasio, O., P. Levell, H. Low, and V. Sánchez-Marcos (2018). Aggregating Elasticities: Intensive and Extensive Margins of Women's Labor Supply. *Econometrica* 86(6), 2049–2082.
- Attanasio, O., H. Low, and V. Sánchez-Marcos (2008). Explaining Changes in Female Labor Supply in a Life-Cycle Model. *American Economic Review* 98(4), 1517–1552.
- Benabou, R. (2002). Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency? *Econometrica* 70(2), 481–517.
- Bick, A. and N. Fuchs-Schündeln (2018). Taxation and Labour Supply of Married Couples across Countries: A Macroeconomic Analysis. *Review of Economic Studies* 85(3), 1543–1576.
- Blundell, R., M. C. Dias, C. Meghir, and J. Shaw (2016, September). Female Labor Supply, Human Capital, and Welfare Reform. *Econometrica* 84(5), 1705–1753.
- Blundell, R. and T. Macurdy (1999). Labor supply: A review of alternative approaches. In O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics*, Volume 3, pp. 1559–1695. Elsevier.

- Borella, M., M. D. Nardi, and E. French (2018, March). Who Receives Medicaid in Old Age? Rules and Reality. *Fiscal Studies* 39(1), 65–93.
- Costa Dias, M., R. Joyce, and F. Parodi (2018). The Gender Pay Gap in the UK: Children and Experience in Work. IFS Working Paper W18/02.
- De Nardi, M., E. French, and J. B. Jones (2016). Medicaid Insurance in Old Age. *American Economic Review* 106(11), 3480–3520.
- Dickert-Conlin, S. and C. Meghea (2004). The Effect of Social Security on Divorce and Remarriage Behavior. Working paper no. 9, Center for Retirement Research, Boston College.
- Dillender, M. (2016). Social Security and Divorce. *B. E. Journal of Economic Analysis and Policy* 16(2), 931–971.
- Dotsey, M., W. Li, and F. Yang (2014). Consumption and Time Use Over the Life Cycle. *International Economic Review* 55(3), 665–692.
- Eckstein, Z., M. Keane, and O. Lifshitz (2019). Career and Family Decisions: Cohorts Born 1935–1975. *Econometrica* 87(1), 217–253.
- Eckstein, Z. and O. Lifshitz (2011). Dynamic Female Labor Supply. *Econometrica* 79(6), 1675–1726.
- French, E. (2000). The Effects of Health, Wealth, and Wages on Labor Supply and Retirement Behavior. Working Paper, Federal Reserve Bank of Chicago.
- Gayle, G.-L. and A. Shephard (2019). Optimal Taxation, Marriage, Home Production, and Family Labor Supply. *Econometrica* 87(1), 291–326.
- Goda, G. S., J. B. Shoven, and S. N. Slavov (2007). Social Security and the Timing of Divorce. NBER Working Paper 13382.
- Gourinchas, P.-O. and J. A. Parker (2002). Consumption over the Life Cycle. *Econometrica* 70(1), 47–89.
- Groneck, M. and J. Wallenius (2020). It Sucks to Be Single! Marital Status and Redistribution of Social Security. *Economic Journal* 131(633), 327–371.

- Guner, N., R. Kaygusuz, and G. Ventura (2012). Taxation and Household Labour Supply. *Review of Economic Studies* 79(3), 1113–1149.
- Hubbard, R. G., J. Skinner, and S. P. Zeldes (1995). Precautionary Saving and Social Insurance. *Journal of Political Economy* 103(2), 360–399.
- Hubener, A., R. Maurer, and O. S. Mitchell (2016). How Family Status and Social Security Claiming Options Shape Optimal Life Cycle Portfolios. *Review of Financial Studies* 29(4), 937–978.
- Kaygusuz, R. (2015). Social Security and Two-Earner Households. *Journal of Economic Dynamics and Control* 59, 163–178.
- Kleven, H. J., C. T. Kreiner, and E. Saez (2009). The Optimal Income Taxation of Couples. *Econometrica* 77(2), 537–560.
- Liebman, J. B., E. F. Luttmer, and D. G. Seif (2009). Labor Supply Responses to Marginal Social Security Benefits: Evidence from Discontinuity. *Journal of Public Economics* 93(11-12), 1208–1223.
- Low, H., C. Meghir, L. Pistaferri, and A. Voena (2018). Marriage, Labor Supply and the Social Safety Net. Working Paper.
- Nishiyama, S. (2019). The Joint Labor Supply Decision of Married Couples and the U.S. Social Security Pension System. *Review of Economic Dynamics* 31, 277–304.
- Persson, P. (2020). Social Insurance and the Marriage Market. *Journal of Political Economy* 128(1), 252–300.
- Roys, N. and C. Taber (2017). Skills Prices, Occupations and Changes in the Wage Structure for Low Skilled Men. Working Paper.