Why Do Households Save and Work?

Margherita Borella, Mariacristina De Nardi, Fang Yang, and Johanna Torres-Chain*

January 22, 2026

Abstract

This paper quantifies why households save and work using a life-cycle model that incorporates wage risk, endogenous labor supply of both spouses, marital transitions, health, medical expenses, mortality, and bequest motives at the death of the first and last household member. We estimate it using PSID and HRS data and conduct counterfactuals to assess the quantitative role of individual mechanisms. Precautionary saving against wage risk is smaller than in models that abstract from labor supply and within-household insurance. Bequest motives and medical expenses remain important drivers of wealth, while marriage and divorce generate large but offsetting effects across household types.

*Borella: University of Torino and CeRP-Collegio Carlo Alberto. De Nardi: University of Minnesota, Federal Reserve Bank of Minneapolis, CEPR, and NBER. Yang: Federal Reserve Bank of Dallas. Torres Chain: University of Minnesota. De Nardi gratefully acknowledges support from NSF grant SES-2044748. Borella gratefully aknowledges funding by the European Union - Next Generation EU, Mission 4 Component 2, within the NRRP (National Recovery and Resilience Plan) project "Age-It - Ageing well in an ageing society" (PE0000015 - CUP H73C22000900006). Torres-Chain gratefully acknowledges support from the Alfred P. Sloan Foundation, Grant Number G-2019-12400. We thank Adrien Auclert, Tim de Silva, Chris Tonetti, and seminar participants at many institutions 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, the CEPR, any agency of the federal government, the Federal Reserve Bank of Minneapolis, the Federal Reserve Bank of Dallas, the Federal Reserve System, the European Union, or the European Commission.

1 Introduction

Households face a rich set of risks over the life cycle. Earnings fluctuate, individuals marry and divorce, and health and mortality shocks arrive. At the same time, households have several effective margins of adjustment. They choose how much to work, how much to save, and whether to leave bequests. In the data, these risks and margins operate jointly over the full life cycle.

Much of the existing literature studies these forces in isolation or restricts attention to particular stages of the life cycle. Models that abstract from family structure, labor supply responses, or late-life risks inevitably attribute observed behavior to the remaining margins. As a result, conclusions about saving and labor supply depend critically on which risks are modeled, which adjustment margins are available, and which parts of the life cycle are taken into account.

We develop and estimate a life-cycle model in which single and married individuals jointly choose labor supply and savings from labor market entry until death. The model incorporates endogenous labor supply of both spouses, marital transitions, retirement decisions, health and longevity risk, medical expenses in old age, and bequest motives at the death of both the first and the last household member.

Two features of the framework are central. First, key adjustment margins operate at the household level. Marriage pools income, generates economies of scale in consumption, and provides insurance through two labor income draws and two labor supply margins. In retirement, couples face joint health and survival risk, which further shapes saving behavior. Second, decisions earlier in life respond to anticipated late-life risks. Labor supply and saving during working years reflect not only earnings risk but also expected medical expenses, survival, and bequests.

We estimate the model using PSID and HRS data for the 1941–1945 birth cohort, matching life-cycle profiles of labor force participation, hours worked, and wealth by marital status. Despite its scope, the model is parsimonious: a limited number of parameters is disciplined by a large set of moments, forcing the model to match interactions across risks and behaviors rather than fitting individual mechanisms in isolation.

Our model fits the targeted profiles well and reproduces systematic differences in wealth and labor supply across household types. Although labor supply elasticities are not targeted, our estimated model implies elasticities that are consistent with existing evidence.

We use the estimated model to evaluate the quantitative role of different risks

and behavioral margins through counterfactual experiments. These counterfactuals isolate the contribution of a given mechanism in an environment where households can adjust optimally along all relevant margins.

We begin by summarizing the main qualitative conclusions of the analysis. Four findings stand out. First, when labor supply provides an adjustment margin through both one's own labor supply and a spouse's labor supply, precautionary saving accounts for a much smaller share of wealth accumulation in response to earnings risk than in working-life models with fixed labor supply. As in earlier work, precautionary responses to earnings risk operate primarily earlier in life, but once labor supply and within-household insurance are available, they translate far less into precautionary wealth accumulation. Second, bequest motives and medical expense risk remain important drivers of aggregate wealth, with bequest motives playing a larger role, particularly at older ages, consistent with the retirement literature (e.g., Lockwood 2012, 2018; Ameriks et al. 2020a; De Nardi et al. 2025). Their effects arise well before retirement, as households adjust labor supply and saving in anticipation of late-life needs. Third, late-life saving motives interact strongly: removing bequest motives and medical expense risk jointly reduces aggregate wealth by far more than the sum of their individual effects. This non-additivity mirrors interactions documented in retirement-only frameworks (e.g., De Nardi et al. 2025), but in our model it reflects behavior shaped over the full life cycle. Fourth, marriage and divorce have modest effects on aggregate wealth but generate large and offsetting responses across household types, with substantial implications for both saving and labor supply over the life cycle.

Taken together, these qualitative findings show that several central messages emphasized in earlier work are preserved once risks and adjustment margins are modeled jointly. In particular, precautionary responses to earnings risk operate primarily earlier in life, while bequest motives and late-life medical risk remain key drivers of wealth accumulation, especially at older ages, consistent with the retirement literature (e.g., Lockwood 2012, 2018; Ameriks et al. 2020a; De Nardi et al. 2025). What changes sharply is their quantitative interpretation. Allowing labor supply, family structure, and late-life risks to interact within a unified life-cycle framework substantially reduces the contribution of precautionary saving against wage risk to aggregate wealth accumulation, while bequest motives and medical expense risk dominate. Moreover, the forces that shape aggregate wealth differ sharply from those that shape aggregate earnings.

To make these mechanisms concrete, we now describe the quantitative responses

underlying these conclusions. Eliminating wage risk reduces aggregate wealth by 10.4% and raises aggregate earnings by 2.3%. In the absence of wage risk, households, especially younger ones, supply less labor and accumulate less wealth early in life, when precautionary motives are strongest in models with wage uncertainty. Later in life, lower accumulated wealth and lower income from wealth induce households to work more in order to rebuild savings. As a result, precautionary responses to earnings risk translate less into wealth accumulation and more into labor supply adjustments. Quantitatively, the contribution of wage risk to aggregate wealth is much smaller than in working-life life-cycle models that abstract from labor supply responses and from other saving motives, such as Gourinchas and Parker (2002) and Cagetti (2003), which attribute 50–70% of wealth accumulation to precautionary saving against earnings risk. In our unified framework, wage risk affects both saving and labor supply decisions over the full life cycle, and allowing labor supply and within-household insurance substantially reduces the share of aggregate wealth attributed to precautionary saving, consistent with evidence that labor supply provides an important source of insurance against earnings risk (e.g., Low (2005); Pijoan-Mas (2006); Blundell, Pistaferri, and Saporta-Eksten (2016b)).

Bequest motives remain an important driver of wealth accumulation even in the presence of endogenous labor supply and other risks. Eliminating voluntary bequest motives lowers aggregate wealth by 23.8% and aggregate earnings by 1.2%, with effects that arise well before retirement. Households adjust labor supply and saving during working years in anticipation of future bequests, and accounting for bequests at the death of the first spouse, in addition to those at the death of the last household member, is important for understanding saving behavior within couples and the wealth profiles of singles following widowhood. The magnitude of this effect lies well below the upper-bound transfer-wealth shares emphasized in early accounting exercises, such as Kotlikoff and Summers (1981), who attribute roughly 80% of aggregate wealth to intergenerational transfers, and is closer to the lower estimates stressed by Modigliani (1988), who argues for a contribution of about 20%, with intermediate values reported by Gale and Scholz (1994). Unlike these accounting decompositions, which attribute observed wealth to transfers broadly construed, our counterfactual isolates the behavioral role of an intentional bequest motive in a unified life-cycle model in which households can adjust labor supply and saving over the full life cycle and face medical-expense and longevity risk.

Medical expenses during retirement also play an important role. Removing medical expenses lowers aggregate wealth by 13.1% and aggregate earnings by 0.7%. These

effects are substantially larger than in retirement-focused frameworks that condition on wealth at retirement and abstract from working-age adjustments. For example, De Nardi et al. (2025) find that eliminating medical expenses reduces average wealth by about 3.1%, with effects concentrated late in life. In contrast, by modeling the full adult life cycle, our framework allows anticipated medical expense risk to affect saving and labor supply decisions well before retirement, generating a larger impact on aggregate wealth. Our estimate is close to those obtained in life-cycle models that allow saving to adjust over the full life span but abstract from labor supply and family structure, such as Kopecky and Koreshkova (2014), who estimate that out-of-pocket medical expenses account for about 13.5% of aggregate wealth.

Late-life saving motives also interact strongly. Removing bequest motives and medical expense risk jointly reduces aggregate wealth by substantially more than the sum of their individual effects. This non-additivity is consistent with interactions emphasized in retirement-focused analyses (e.g., De Nardi et al. 2025; Ameriks et al. 2020a), but in our framework it reflects behavior shaped by anticipation and adjustment over the full life cycle rather than responses confined to old age.

Finally, marriage and divorce have modest effects on aggregate wealth but generate large and offsetting responses across household types. Eliminating marriage and divorce dynamics changes aggregate wealth by only 0.7%, but this masks substantial heterogeneity. In the absence of divorce risk, couples save and work less, reflecting reduced precautionary motives, while eliminating marriage prospects leads singles to save and work substantially more. As a result, marriage and divorce have much larger effects on labor market outcomes than on aggregate wealth.

Removing all saving motives other than retirement needs and uncertain and heterogeneous lifespans reduces lifetime wealth accumulation by 56.9% and aggregate earnings by 2.7%. This counterfactual highlights how the quantitative contribution of any single motive depends on the broader risk environment in which households operate.

Our paper makes three main contributions. First, we develop and estimate a unified life-cycle model that jointly captures saving and labor supply decisions of singles and couples over the full adult life cycle, allowing for multiple sources of risk, endogenous labor supply of both spouses, and selection into marriage, health, and survival. Second, we estimate the model over the full adult life cycle, matching detailed profiles of participation, hours, wealth, and marital status across groups and ages, and accounting for survival and health selection effects. Third, we show how the quantitative conclusions of the life-cycle literature change once households face multiple

risks and adjustment margins jointly, clarifying the implications of abstracting from particular risks, margins, or stages of the life cycle when studying saving and labor supply behavior.

2 Related literature

A large literature studies saving and labor supply over the life cycle, emphasizing different sources of risk and adjustment margins. Our contribution is not to challenge the importance of these mechanisms, but to reassess their relative magnitude and interactions once households face multiple risks and adjustment margins jointly over the adult life cycle.

Early work on precautionary saving considers limited forms of earnings risk and abstracts from endogenous labor supply. Seminal contributions such as Aiyagari (1994) study unemployment risk faced by infinitely lived agents and find relatively modest effects on aggregate wealth. Life-cycle models that focus on the working period, including Gourinchas and Parker (2002) and Cagetti (2003), introduce finite lives and richer earnings dynamics and conclude that precautionary saving against earnings risk accounts for a large share of wealth accumulation. A closely related literature shows, however, that labor supply itself provides an important margin of insurance. Allowing hours to respond to wage shocks substantially reduces the need for precautionary saving Low (2005); Pijoan-Mas (2006). This adjustment margin remains relevant later in life, including through retirement behavior (e.g., Ameriks et al. 2020b). At the household level, spousal labor supply provides even stronger insurance than individual labor supply alone, as documented by Blundell, Pistaferri, and Saporta-Eksten (2016b). Taken together, these findings highlight that estimates of precautionary saving depend critically on whether labor supply and within-household insurance are available margins of adjustment.

A separate and long-standing literature emphasizes bequest motives as a key determinant of wealth accumulation. The classic debate between Kotlikoff and Summers (1981) and Modigliani (1988) illustrates sharply different views on the contribution of bequests to aggregate wealth. Subsequent work, including Gale and Scholz (1994), Lockwood (2012, 2018), and Ameriks, Briggs, Caplin, Shapiro, and Tonetti (2020a), produces intermediate estimates. Structural retirement-focused models, such as Ameriks et al. (2020a), show that late-life saving reflects the interaction of medical expense risk, longevity risk, and utility from holding wealth late in life, rather than any single motive in isolation. Much of this literature abstracts from labor supply

and family structure and typically focuses on behavior after retirement. More recent work documents that bequests may also occur at the death of the first spouse, with important implications for saving behavior within couples (De Nardi, French, Jones, and McGee (2025)).

Another strand of the literature studies late-life risks, including medical expenses and longevity risk. This work shows that uncertainty in old age plays an important role for saving behavior but typically conditions on wealth at retirement or later ages (De Nardi, French, Jones, and McGee (2025); Ameriks, Briggs, Caplin, Shapiro, and Tonetti (2020a)). By construction, retirement-focused frameworks limit the scope for late-life risks to affect labor supply and saving decisions earlier in life. Our approach complements this literature by modeling the full adult life cycle, allowing anticipated late-life risks to influence working-age behavior and thereby shape wealth accumulation at older ages.

Marriage and divorce are another important source of economic risk. Cubeddu and Ríos-Rull (2003) emphasize marital risk by modeling marriage formation and dissolution as sources of uncertainty faced by households and show that the timing of marital transitions has first-order implications for saving behavior. Subsequent work often abstracts from marital transitions or studies singles and couples separately, limiting its ability to capture within-household insurance, selection into marriage and divorce, and composition effects over the life cycle. Our framework allows marital transitions to interact with labor supply, saving motives, and other risks, which is essential for understanding both group-specific behavior and aggregate outcomes.

Our paper is closely related to Borella, De Nardi, and Yang (2023b) and De Nardi, French, Jones, and McGee (2025). Relative to Borella, De Nardi, and Yang (2023b), which focuses on saving and labor supply during working life, we show how retirement risks and bequest motives feed back into behavior earlier in life and alter the quantitative role of working-age risks. Relative to De Nardi, French, Jones, and McGee (2025), which studies saving behavior in retirement conditional on observed wealth at older ages, we endogenize labor supply, marriage, health, and survival earlier in the life cycle and show how these choices shape wealth and labor outcomes later on. Together, these comparisons underscore that conclusions about the importance of particular risks or motives depend on which stages of the life cycle and which adjustment margins are modeled.

3 Savings and labor supply over the life cycle

We use data from the PSID, supplemented with the HRS to improve coverage at older ages, focusing on the 1941–1945 birth cohort (henceforth, the 1945 cohort). Although this cohort is older, it offers a key advantage: we observe both their working years and most of their retirement period. This allows us to estimate retirement risks and bequest motives, which are central to understanding saving at older ages.

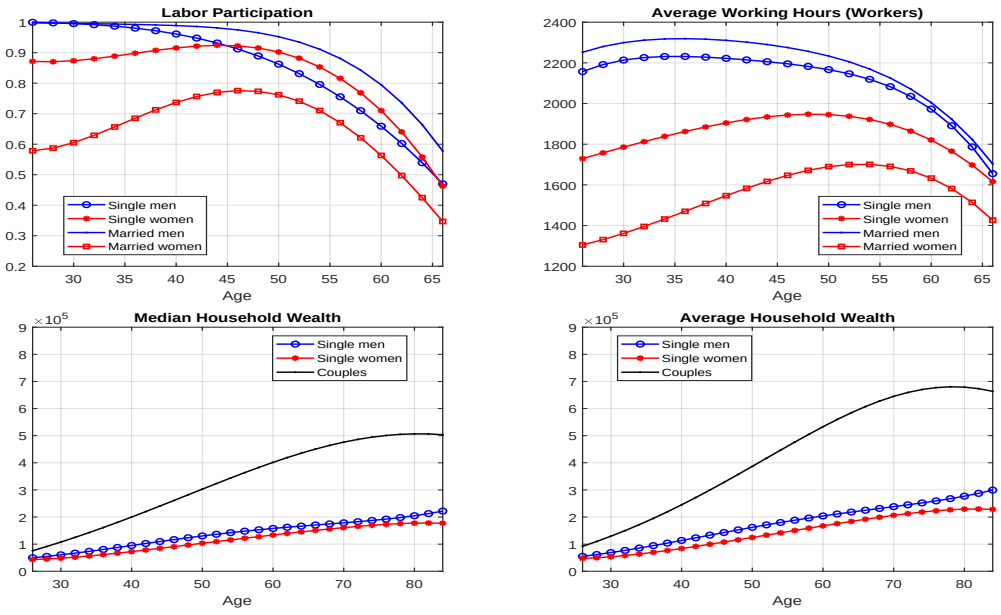


Figure 1: Life-cycle profiles by gender and marital status, PSID.

The top panels of Figure 1 show average labor force participation and hours worked for single and married men and women.[1] The data for couples include only households in which both spouses are alive and remain married. The data for singles include individuals who are single at age 26 as well as those who are single at a given age due to divorce or widowhood. The left-hand-side panel shows that married men have the highest participation rate, with only a gradual decline after age 45, while single men experience a much steeper drop. Single women begin about 10 percentage points lower than single men and increase participation until peaking around age 50. Married women have the lowest participation rates: around 57% at age 26, rising to 78% between ages 40 and 50, then declining similarly to the other groups. The top-

[1]We include cohabiting couples among married couples. In the PSID, it is not possible to distinguish between legal marriage and cohabitation prior to 1983. In that year, only 2% of couples in our cohort were not legally married. The profiles are estimated by fitting a fourth-order polynomial in age, fully interacted with marital status and cohort dummies, separately for each gender. We express all monetary values in 2018 dollars.

right panel shows that married men work more hours, on average, and that women work fewer hours than men, even conditional on participation.

These labor supply differences reflect both behavioral and compositional factors: individuals with higher wages are more likely to marry and remain married. Our analysis accounts for these compositional differences both at model entry (age 26) and through subsequent marital transitions, which shape the evolving composition of married and single individuals.

The bottom panels of Figure 1 show median and average net worth by age and marital status. Two features stand out. First, despite economies of scale, couples hold more than twice the wealth of singles at all ages. Second, wealth declines only modestly after retirement, contrary to the predictions of a basic life-cycle model.

This limited drawdown reflects both saving motives and mortality bias. Saving for medical expenses and bequest motives can lead households to hold more wealth at older ages. Mortality bias also matters: wealthier individuals live longer, even conditional on gender and marital status. This has two important implications. Because wealthier individuals survive longer, average wealth among the surviving population flattens or even rises after age 66. In addition, since couples are wealthier than singles, those who become single through widowhood tend to be wealthier than those who were always single, further stabilizing singles' wealth profiles after retirement.

Our analysis extends previous literature by examining how medical expenses, bequest motives, and selection effects shape wealth and labor supply decisions across the full life cycle. It also accounts for mortality bias and for marital transitions earlier in life, both of which are salient in the data.

4 The model

Life is divided into three stages.[2] The first is a working stage (ages 26 to 61), followed by an early retirement stage (ages 62 to 65), and finally a retirement stage (ages 66 to the maximum age of 99). Our model period is two years long, consistent with the frequency of our data.

During the **working stage**, single and married individuals decide how much to

[2]This model builds on Borella, De Nardi, and Yang (2023b), particularly in its structure and exposition of the working-life stage. We retain many modeling features from that earlier framework and follow its exposition where content overlaps. However, this paper introduces several key innovations, including a richer treatment of the retirement period, endogenous life expectancy, and new estimation moments. These improvements significantly extend the earlier model and are essential to addressing the questions posed in this paper.

work and save. Wages are subject to shocks and depend on one's human capital endogenously accumulated through work. Married people face divorce shocks, and single people may meet partners and get married.

We model (and estimate) how available time is split between working and leisure, allowing it to depend on gender and marital status. We interpret available time as net of home production, child care, and elderly care—tasks that are difficult to outsource.

All workers incur fixed costs of working, which vary by age, gender, and marital status. These costs represent commuting, preparing for work, and making arrangements to be able to work.

Single and married women have children, with the number of children depending on maternal age and marital status. Raising children imposes both time and monetary costs. Time costs reduce the available hours for work and leisure. Monetary costs arise from increased consumption needs, modeled through equivalence scales, and child care expenses. For working mothers, child care costs vary based on the number and age of their children, as well as their earnings. Consequently, 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 no longer marry, and couples no longer divorce. These assumptions are data-driven[3] and help keep our model tractable. For tractability, if they claim Social Security, they can no longer work, and 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, medical expenses, and mortality shocks. Thus, each married person faces the risk of his or her spouse dying, in addition to their own. Mortality and health risks, as well as medical expenses, depend on age, gender, health status, human capital, and marital status. Because human capital is, at least to some extent, endogenous, so are lifespan, healthspan, and medical expenses.

Because we explicitly model labor force participation and hours for both spouses, savings, bequest motives at the death of the first and last spouse, and medical expenses in old age, the model is computationally intensive (see Appendix E for details). For tractability, we also 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

[3]Only 1% of couples divorce and 4% of singles marry between ages 62 and 72 in the HRS data for our 1945 cohort.

from the data). 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).

4.1 Preferences

Let t be age, with people entering at age 26 and dying by the end of age 99. Households discount the future at a 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.

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

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

c_t is consumption, $\eta_t^{i,1}$ is the equivalence scale for singles and $\eta_t^{i,2}$ is the one for couples, $l_t^{i,j}$ is leisure, and $b \geq 0$ is the value of life.[4] Leisure is given by

$$l_t^{i,j} = L^{i,j} - n_t^i - P_t^{i,j} I_{n_t^i}, \tag{1}$$

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

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

where we normalize L to 112 hours a week. The parameter $FL^{i,j}$ captures heterogeneity in time endowments across gender and marital status and is estimated using our structural model. The term n_t^i, is hours worked, $I_{n_t^i}$ is an indicator function that equals 1 when hours worked are positive.

The term $P_t^{i,j}$ is the fixed time cost of work, which depends on gender, marital status, and age. It assumes the following functional form, whose parameters we

[4] Given that, to some extent, our life expectancy is endogenous, we follow the approach in Hall and Jones (2007) in adding a parameter that ensures that people are happy to be alive.

estimate using our structural model

$$P_t^{i,j} = \frac{\exp(p_0^{i,j} + p_1^{i,j}t + p_2^{i,j}t^2)}{1 + \exp(p_0^{i,j} + p_1^{i,j}t + p_2^{i,j}t^2)}. \tag{3}$$

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

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

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

Utility from bequests. We model utility from two types of bequests: those distributed to heirs other than the spouse at the death of the first household member and those at the death of the last. While the latter has been extensively studied, the former was recently introduced and estimated by De Nardi, French, Jones, and McGee (2025) in a model of couples and singles over age 70. Their analysis shows that, upon a spousal death, over 30% of cases involve substantial "side" bequests to heirs other than the surviving spouse, which significantly influence household saving and bequest behavior. Households may allocate side bequests at this time because a spousal death resolves uncertainty about the deceased spouse's remaining lifespan and potential medical expenses.

The household derives utility $\theta_j(e)$ from leaving a bequest $e \geq 0$ to heirs other than the spouse. Specifically, $\theta_1(e)$ represents the utility of bequests when there is one person in the household (and thus no surviving spouse), and $\theta_2(e)$ captures the utility from bequests to non-spousal heirs when there are two spouses and one survives. The bequest function is defined as

$$\theta_j(e) = \phi_j \frac{(e + k_j)^{1-\gamma}}{1-\gamma}$$

where k_j determines the curvature of the bequest function, and ϕ_j governs its intensity.

De Nardi (2004) employed this functional form to explain life-cycle savings behavior. It accommodates several interpretations of the "bequest motive," including dynastic or "warm glow" altruism (Becker and Tomes, 1986; Abel and Warshawsky, 1988; Andreoni, 1989), strategic motives (Bernheim, Shleifer, and Summers, 1985; Brown, 2003), and utility derived directly from wealth (Carroll, 1998; Hurd, 1989).

4.2 Human capital and wages

We define human capital, $\bar{y}_t^i$, as the average past earnings of an individual at each age (see Equation (12) for a formal definition). This measure depends on initial wages, which reflect schooling, as well as on subsequent labor market experience and earnings, rather than solely on cumulative time worked.

This definition offers two key advantages. First, it aligns with empirical findings that individuals with higher education levels enjoy greater returns on their labor market experience (Blundell, Costa Dias, Meghir, and Shaw, 2016a; Costa Dias, Joyce, and Parodi, 2020). Second, it enables the use of a single state variable to track both human capital and Social Security contributions, ensuring computational tractability within our framework.[5]

In the model, human capital, combined with stochastic wage shocks, determines effective wages. Wages are composed of two components: a deterministic function of age, gender, and human capital, $z_t^i(\bar{y}_t^i)$, and a persistent stochastic shock, ϵ_t^i, which evolves as

$$\ln \epsilon_{t+1}^i = \rho_\epsilon^i \ln \epsilon_t^i + \upsilon_t^i, \qquad \upsilon_t^i \sim N(0, (\sigma_\upsilon^i)^2).$$

The effective hourly wage is given by the product of $z_t^i(\bar{y}_t^i)$ and ϵ_t^i. This formulation captures both systematic wage variation, influenced by demographic and human capital factors, and idiosyncratic wage shocks.

4.3 Marriage and divorce

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

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

We assume random matching over wealth a_{t+1} and partner's human capital $\bar{y}_{t+1}^p$, conditional on the partner's wage shock. Thus, we have

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

[5] In estimating the wage function, we use human capital defined in Equation (12) to maintain consistency with the model. See Appendix C.1 for details.

For working-age couples, divorce can occur as a result of a shock that depends on the age and wage shocks of both partners. This probability is denoted by $\zeta_{t+1}(\cdot) = \zeta_{t+1}(\epsilon_t^1, \epsilon_t^2)$. If a couple divorces, their wealth is split equally. For simplicity, we abstract from alimony payments.

4.4 Costs of raising children and running a household

We keep track of the total number of children and their ages as a function of the mothers' age and marital status. The number of children aged 0-5 is denoted by $f_t^{0,5}(i,j)$, and the corresponding childcare cost for each child is $\tau_c^{0,5}$. Similarly, the number of children age 6-11 is $f_t^{6,11}(i,j)$ and the corresponding child care cost per child is $\tau_c^{6,11}$.

4.5 Health, medical expenses, and death

At age 66, individuals are endowed with a health state h_{66}, which can be good, bad, or in a nursing home (disabled). The distribution of this initial health state depends on gender, marital status, and one's human capital, allowing the model to capture observed patterns where women, married individuals, and those with higher incomes typically enjoy better health outcomes.

After age 66, health evolves according to a Markov process. For singles, the transition probability is a function of age, gender, marital status, current health, and human capital and can be written as $\pi_t^{i,1}\left(h_t^i, \bar{y}_r^i\right)$. For individuals in couples, we also allow the transition probability to depend on the human capital of both partners and can thus be expressed as $\pi_t^{i,2}\left(h_t^i, \bar{y}_r^i, \bar{y}_r^p\right)$.

We allow the medical expenses of singles to be a function of age, gender, marital status, current health, and human capital: $m_t^{i,1}\left(h_t^i, \bar{y}_r^i\right)$, while the medical expenses of couples can depend on both partners' human capital $m_t^{i,2}\left(h_t^i, \bar{y}_r^i, \bar{y}_r^p\right)$.

To capture the sharp rise in medical expenses during the period preceding death and the additional costs associated with funerals and burials, the model includes an additional expense shock at death. For singles it is given by $d_t^{i,1}\left(h_t^i, \bar{y}_r^i\right)$ and for couples, by $d_t^{i,2}\left(h_t^i, \bar{y}_r^i, \bar{y}_r^p\right)$.

Survival probabilities $s_t^{i,j}(h_t^i, \bar{y}_r^i)$ depend on age, gender, marital status, current health, and human capital.

4.6 Initial conditions

We take the fraction of single and married people at age 26 and their distribution over the relevant state variables (wealth, human capital, and wage shocks, with the latter two being for each of the spouses in the case of couples) from the PSID.

4.7 Government

Our cohort faces the effective time-varying tax rates that it experienced in the data and that we estimate from the PSID. As 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 (thus time). Taxes paid are a function of total income Y

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

The government imposes a proportional payroll tax τ_t^{SS}, applied to earnings up to a Social Security cap $\widetilde{y}_t$. These revenues help finance old-age Social Security benefits, which depend on average past earnings (human capital, as discussed in Section 4.2). We also allow the payroll tax and the Social Security cap to change over time, as in the data. Households are assumed to have anticipated these changes.

The insurance provided by the government, such as Medicaid and SSI, is captured by a means-tested consumption floor, $\underline{c}(j)$, as in Hubbard, Skinner, and Zeldes (1995).[6]

4.8 Recursive formulation

We compute nine value functions for different groups and stages of life. For brevity, we report in the main text only those that are most informative about our modeling choices and economic mechanisms. The remaining ones pertain to early retirement for couples and singles, and those of individuals in couples, and in Appendix B.

4.8.1 The value function of working-age singles

The value function of a working-age single person ($j = 1$) depends on age t, gender i, wealth a_t^i, persistent earnings shock ϵ_t^i, and human capital $\bar{y}_t^i$. Net worth a_t earns

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

a rate of return r.

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

subject to Equation (1) and

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

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

$$T(\cdot) = T(ra_t + Y_t, i, j, t), \tag{10}$$

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

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

$$a_{t+1} \geq 0, \tag{13}$$

$$n_t^i \geq 0. \tag{14}$$

The expectation of the value function for the next period if one remains single 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. Equation (12) describes the evolution of human capital, measured as average accumulated earnings up to the Social Security earnings cap $\widetilde{y}_t$, in which $t_0 = 26$. This cap reflects the fact that earnings above $\widetilde{y}_t$ do not affect payroll taxes or future Social Security benefits, ensuring consistency between human capital accumulation, taxation, and benefit determination.

We assume a constant rate of return, abstracting from time-varying and stochastic returns. This choice streamlines the computational structure and keeps the focus on life-cycle labor supply, marital risk, health risk, and bequest motives, which are the central mechanisms of the paper.

4.8.2 The value function of retired singles

The value function of a retired single ($j = 1$) with health h_t^i, human capital $\bar{y}_r^i$, and Social Security claiming age tr is

$$R^s(t, i, a_t, h_t^i, \bar{y}_r^i, tr) = \max_{c_t, a_{t+1}} \Bigg(v^i(c_t, L^{i,j}, \eta_t^{i,j}) + \beta s_t^{i,j}(\cdot) E_t R^s(t+j, i, a_{t+1}^i, h_{t+1}^i, \bar{y}_r^i, tr) + \beta(1 - s_t^{i,j}(\cdot))\theta_0[\max(0, (a_{t+1}^i - d_t^{i,j}\left(h_t^i, \bar{y}_r^i\right))] \Bigg), \tag{15}$$

with $\theta_0(\cdot)$ denoting utility from terminal bequests made at the death of the last household member. The problem is subject to equations (10) and (13) in the main text and (29) in Appendix B, and there is a consumption floor determined by

$$B(\cdot) = \max\Big\{0, \underline{c}(j) - \big[(1+r)a_t^i + Y_t - m_t^{i,1}\left(h_t^i, \bar{y}_r^i\right) - T(\cdot)\big]\Big\}, \tag{16}$$

where $m_t^{i,1}\left(h_t^i, \bar{y}_r^i\right)$ are medical expenses (out-of-pocket plus Medicaid payments) and $d_t^{i,j}\left(h_t^i, \bar{y}_r^i\right)$ are additional expenses at death. Also,

$$a_{t+1}^i = 0 \quad \text{if} \quad B(\cdot) > 0, \tag{17}$$

and the budget constraint is given by

$$c_t + a_{t+1}^i = (1+r)a_t^i + Y_t + B(\cdot) - m_t^{i,1}\left(h_t^i, \bar{y}_r^i\right) - T(\cdot). \tag{18}$$

4.8.3 The value function of working-age couples

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

$$W^c(t, a_t, \epsilon_t^1, \epsilon_t^2, \bar{y}_t^1, \bar{y}_t^2) = \max_{c_t, a_{t+1}, n_t^1, n_t^2} \Bigg(w(c_t, l_t^{1,j}, l_t^{2,j}, \eta_t^{i,j}) + (1 - \zeta_{t+1}(\cdot))\beta E_t W^c(t+1, a_{t+1}, \epsilon_{t+1}^1, \epsilon_{t+1}^2, \bar{y}_{t+1}^1, \bar{y}_{t+1}^2) + \zeta_{t+1}(\cdot)\beta \sum_{i=1}^{2} \Big(E_t W^s(t+1, i, a_{t+1}/2, \epsilon_{t+1}^i, \bar{y}_{t+1}^i) \Big) \Bigg), \tag{19}$$

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

$$T(\cdot) = T(ra_t + Y_t^1 + Y_t^2, i, j, t), \tag{20}$$

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

$$a_t \geq 0, \qquad n_t^1, n_t^2 \geq 0. \tag{22}$$

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.

4.8.4 The value function of retired couples

During this stage, the married couple's recursive problem ($j = 2$) depends on each of the spouses health shocks h_t^i and there are survival shocks $s_t^{i,j}(\cdot)$ and medical expenses $m_t^{i,j}(\cdot)$ of each spouse. 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 couple also decides, in the case that each spouse dies, the distribution of bequests towards other heirs and the survivor.

$$\begin{aligned} R^c(t, a_t, h_t^1, h_t^2, \bar{y}_r^1, \bar{y}_r^2, tr) = \max_{c_t, a_{t+1}, e_t^1, e_t^2} \Bigg(& w(c_t, L^{1,j}, L^{2,j}, \eta_t^{i,j}) + \\ & \beta s_t^{1,j}(\cdot) s_t^{2,j}(\cdot) E_t R^c(t+1, a_{t+1}, h_{t+1}^1, h_{t+1}^2, \bar{y}_r^1, \bar{y}_r^2, tr) + \\ & \beta s_t^{1,j}(\cdot)(1 - s_t^{2,j}(\cdot)) \Big(\theta_1(e_t^2) + E_t R^s(t+1, 1, a_{t+1} - e_t^2 - d_t^{2,j}\left(h_t^i, \bar{y}_r^i, \bar{y}_r^p\right), h_{t+1}^1, \bar{\bar{y}}_r^1, tr) \Big) + \\ & \beta s_t^{2,j}(\cdot)(1 - s_t^{1,j}(\cdot)) \Big(\theta_1(e_t^1) + E_t R^s(t+1, 2, a_{t+1} - e_t^1 - d_t^{1,j}\left(h_t^i, \bar{y}_r^i, \bar{y}_r^p\right), h_{t+1}^2, \bar{\bar{y}}_r^2, tr) \Big) + \\ & 2\beta(1 - s_t^{1,j}(\cdot))(1 - s_t^{2,j}(\cdot))\theta_0(\max(0, (a_{t+1} - d_t^{1,j}\left(h_t^i, \bar{y}_r^i, \bar{y}_r^p\right) - d_t^{2,j}\left(h_t^i, \bar{y}_r^i, \bar{y}_r^p\right))/2)) \Bigg), \end{aligned} \tag{23}$$

subject to equations (10), (13), (17) in the main text, and (35) in Appendix B, and

$$\bar{\bar{y}}_r^i = \max(\bar{y}_r^1, \bar{y}_r^2), \tag{24}$$

$$B(\cdot) = \max\Big\{0, \underline{c}(j) - \left[(1+r)a_t + Y_t - m_t^{1,j}\left(h_t^i, \bar{y}_r^i, \bar{y}_r^p\right) - m_t^{2,j}\left(h_t^i, \bar{y}_r^i, \bar{y}_r^p\right) - T(\cdot)\right]\Big\} \tag{25}$$

$$c_t + a_{t+1} = (1+r)a_t + Y_t + B(\cdot) - m_t^{1,j}\left(h_t^i, \bar{y}_r^i, \bar{y}_r^p\right) - m_t^{2,j}\left(h_t^i, \bar{y}_r^i, \bar{y}_r^p\right) - T(\cdot). \tag{26}$$

$$e_t^i \in [0, a_{t+1}] \tag{27}$$

A survivor collects benefits based on the higher amount between their own contributions and those of their deceased spouse (Equation (24)).

5 Estimation

We estimate our model by adopting a two-step estimation strategy (as Gourinchas and Parker (2002)). In the first step, we calibrate some parameters (see Appendix D) and estimate others (see next subsection and Appendix C). In the second step, we use the MSM (See Subsection 5.2 and Appendices F and G). We normalize the time endowment for single men and estimate 21 model parameters, including the discount factor (β), the leisure weight in utility (ω), the parameters governing labor participation costs ($p_0^{i,j}, p_1^{i,j}, p_2^{i,j}$), the time endowment by gender and marital status ($L^{i,j}$), and the parameters related to the curvature and intensity of the bequest functions (ϕ_1, k_1, ϕ_0, k_0).

We estimate the model using 334 moments that capture key features of the data. These include labor market participation and hours worked (conditional on working) for married and single men and women ages 26 to 64, as well as average and median wealth profiles for couples and single men and women from ages 28 to 84. We stop matching wealth moments after age 84 due to the limited sample size at older ages.

5.1 First-step calibration and estimation

For brevity in this subsection, we only discuss the first-step inputs that we change in our counterfactuals. Table D.1 in Appendix D reports the parameters that we calibrate. Appendix C discussed all of our first-step estimation in detail.

Wage Risks. The persistence of wage shocks is 0.9360 for men and 0.9484 for women, while the variance of shocks is larger for men (0.0312) than for women (0.0206).

Marriage and Divorce Risks. Figure 2 shows marriage and divorce probabilities by gender, age, and wage shocks up to age 62, when individuals can still experience marital transitions in the model. At age 26, marriage probabilities are about 35% for men and 30% for women, declining to 5% and 3%, respectively, by age 60. Divorce probabilities start at around 8% for both genders early in life and fall to about 1% later on. There is notable dispersion in these probabilities by wage shocks, particularly for men: higher-wage individuals are generally more likely to marry and

less likely to divorce.

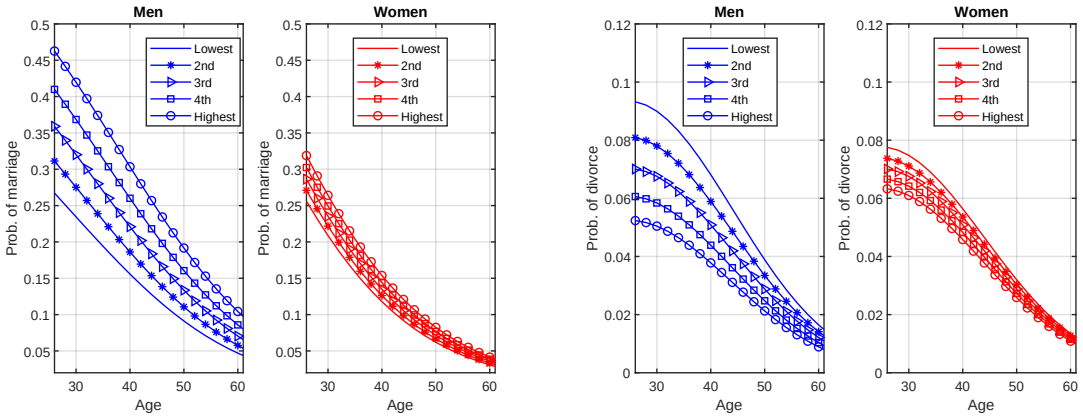


Figure 2: Marriage (on the left) and divorce (on the right) probabilities by gender, age, and one's wage shock. PSID data.

Medical expenses. Figure 3 shows medical expenses (out-of-pocket plus Medi-

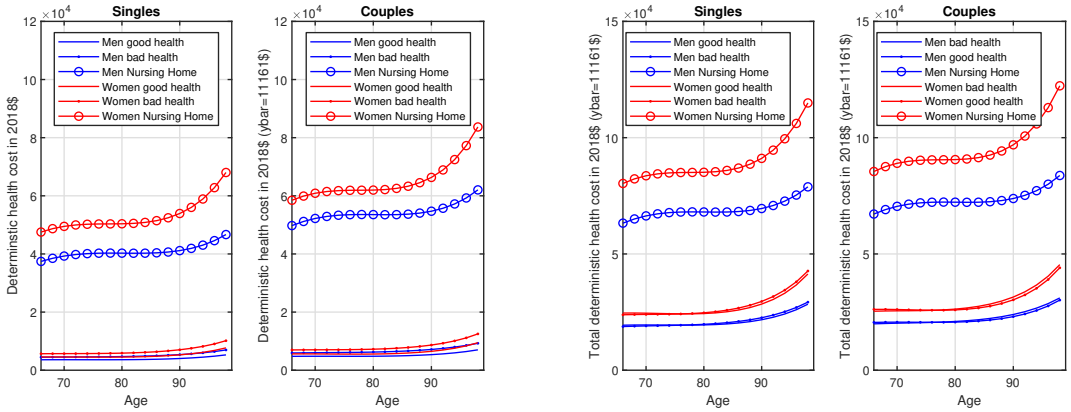


Figure 3: Medical expenses (out-of-pocket plus Medicaid payments) by age, gender, ybar, and marital and health status. Left panel: conditional on survival during the period. Right panel: conditional on dying during the period. HRS data.

caid) by age, gender, and marital and health status. The left panel reports average expenses for individuals who survive the period. The right panel shows medical expenses incurred by individuals who die during the period, which reflect the sharp rise in healthcare use typically associated with the end of life—such as hospitalizations and nursing home care. Among survivors, average medical expenses remain relatively low even at older ages, except for those in nursing homes. In contrast, medical expenses in the period of death are substantially higher across all health states.

5.2 Second-step estimation

Table 1 presents our estimated preference parameters (Appendix G reports all of our parameter estimates and their standard errors). Our estimated discount factor

Estimated parameters	
β: Discount factor	0.9958
ω: Consumption weight	0.4478
$L^{2,1}$: Time endowment (weekly hours), single women	104
$L^{1,2}$: Time endowment (weekly hours), married men	108
$L^{2,2}$: Time endowment (weekly hours), married women	74
$P_t^{i,j}$: Partic. cost	Fig. 10 App. G
ϕ_0: Terminal bequest, strength	1,902,590
k_0: Terminal bequest, shifter	975,581
ϕ_1: Side bequest, strength	38,703,874
k_1: Side bequest, shifter	2,826,257

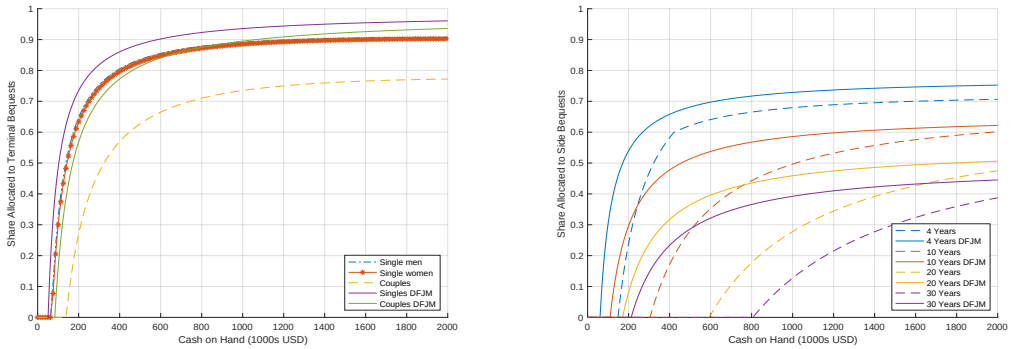
Table 1: Second-step estimated model parameters

is very similar to that estimated by De Nardi, French, and Jones (2016) for elderly retirees, and our estimated weight on consumption is 0.4497. We normalize the time endowment of single men to 112 hours a week (and 5,840 hours a year). Single women have a time endowment of 104 hours a week. We interpret this as their having to spend eight more hours a week managing their household, rearing children, and taking care of elderly parents. The time endowments for married men and women are 108 and 74 hours. This implies that people in the latter two groups spend 4 and 38 hours a week on 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).

Figure G.1 in Appendix G reports the age-varying time costs of working by age, expressed as a fraction of the time endowment of single men. Our estimated participation costs are relatively high when people are younger and increase again after 45. The time costs of going to work might include factors other than commuting time. For instance, they might be higher when children are youngest, because 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.

The bequest parameters are difficult to interpret. To clarify their implications,

we follow De Nardi, French, Jones, and McGee (2025) and plot the share of resources allocated to bequests that they imply. Figure 4, Panel (a), shows the share of terminal bequests as a function of total resources for singles and couples facing certain death in the next period. Bequests behave as a luxury good for singles, and even more so for couples. Panel (b) displays side bequests for couples when one spouse dies, as a function of total resources and the expected remaining lifespan of the surviving spouse. Consistent with De Nardi, French, Jones, and McGee (2025), we find that the longer the survivor's planning horizon, the smaller the share allocated to side bequests, however, our model implies that both terminal and side bequests exhibit stronger luxury behavior than in their framework.



(a) Estimates terminal bequests

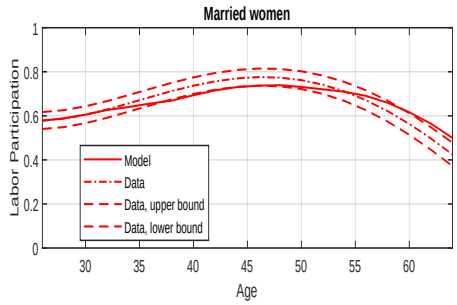
(b) Estimates side bequests

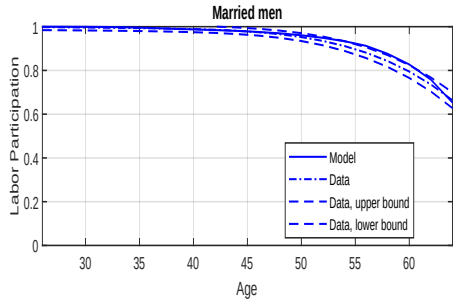
Figure 4: Implications of our estimated bequest parameters. Panel (a): Share of total resources allocated to terminal bequests for couples and single men and women, who face certain death next period. Panel (b): Share allocated to side bequests by couples, upon death of ones spouse, under varying survivor lifespans.

5.3 Model fit

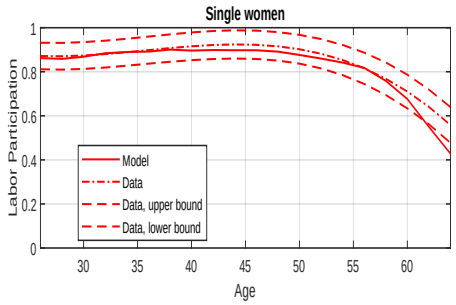
Figures 5 and 6 report the moments from the data and their 95% confidence intervals, and the moments implied by our model. Our model fits the targeted data well, which is remarkable given that it is tightly parameterized: we have 334 targets, and we estimate 21 parameters.[7]

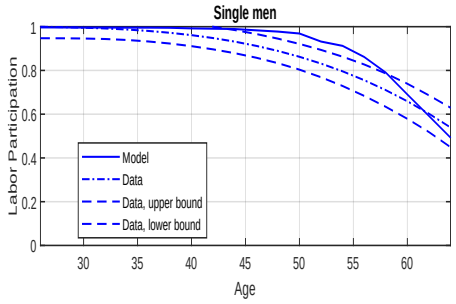
[7]Solving the model for a given parameter vector takes approximately two minutes on a standard high-performance workstation. Estimation requires solving the model several thousand times using the method of simulated moments and typically takes one to two weeks, depending on starting values and the degree of parallelization. We re-estimate the model from multiple starting values to assess convergence and reduce sensitivity to local minima.



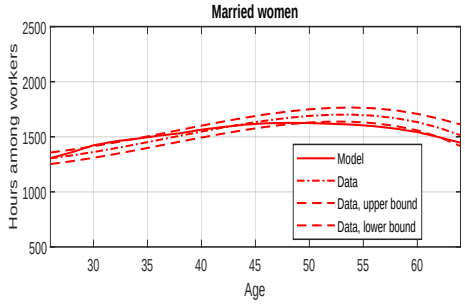


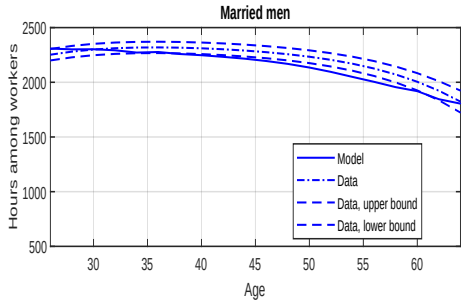
(a) Participation, couples



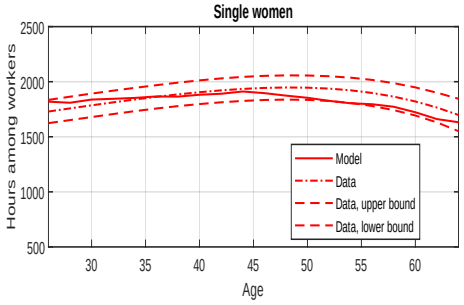


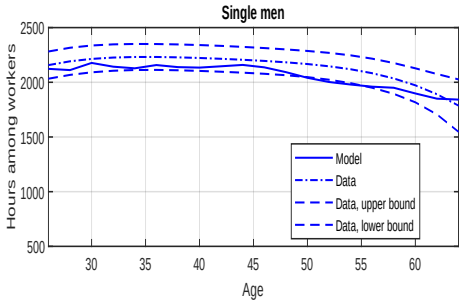
(b) Participation, singles





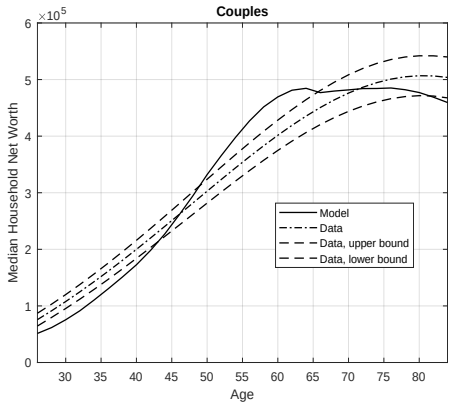
(c) Hours for workers, couples



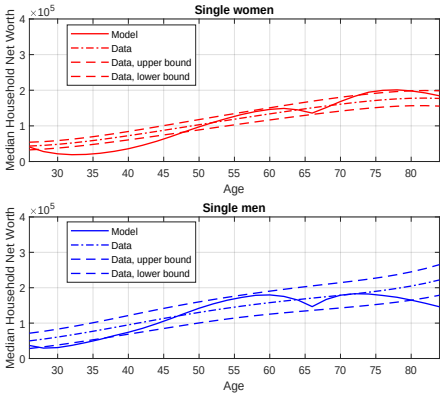


(d) Hours for workers, singles

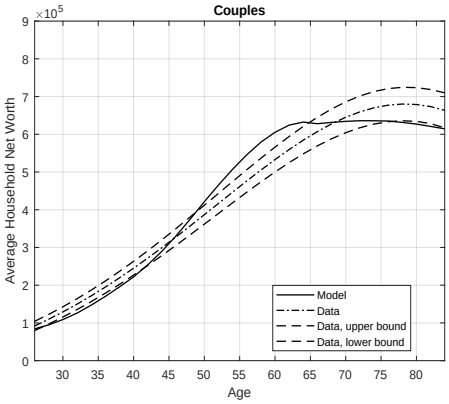
Figure 5: Model-implied participation and hours, and average and 95% confidence intervals from the PSID.



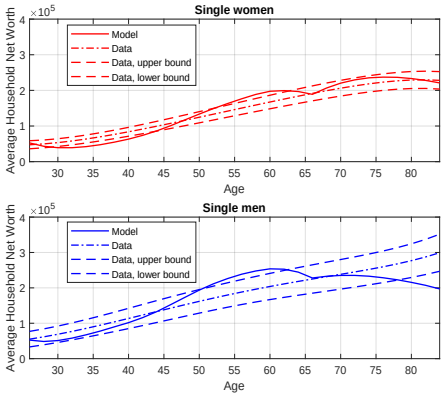
(a) Median wealth, couples



(b) Median wealth, singles



(c) Average wealth, couples



(d) Average wealth, singles

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

5.4 Model-implied labor supply elasticities

Table 2 reports the model-implied elasticities of labor force participation and hours worked (conditional on working) in response to a temporary, anticipated increase in own wages.[8] We find that women's participation elasticities are larger than men's, and that married men are the least responsive group. For all groups, participation elasticities increase sharply near retirement age—a pattern consistent with earlier findings by French (2000) for men.

Our elasticity estimates align well with those in Liebman, Luttmer, and Seif

[8]Each elasticity is computed by raising the wage of a given group (married men, married women, single men, or single women) by 5% at a single age, holding all else constant. The change is not compensated elsewhere, but its lifetime impact is negligible. For example, at age 40, the present value of this temporary wage change amounts to just 0.04% of average discounted family income. The analogous figures at ages 50 and 60 are 0.03% and 0.02%, respectively.

(2009), Blundell and Macurdy (1999), and Attanasio, Levell, Low, and Sánchez-Marcos (2018) and reflect the role of endogenous family structure, participation costs, and time constraints. Married women's high elasticities stem from joint decision-making, marital risk, and child-rearing responsibilities—channels absent in models with fixed households or exogenous labor supply. These mechanisms amplify labor supply responses and highlight its role as a key margin of insurance against income and family shocks.

	Participation				Hours among workers			
	Married		Single		Married		Single	
	W	M	W	M	W	M	W	M
40	0.6	0.1	0.4	0.1	0.2	0.1	0.4	0.3
50	0.5	0.2	0.5	0.3	0.2	0.1	0.5	0.3
60	1.0	0.5	1.6	1.7	0.2	0.0	0.3	0.1

Table 2: Model-implied elasticities of labor supply.

5.5 The economic forces driving identification

In this subsection, we provide a broad overview of the economic forces driving our model parameters' identification. We provide a detailed exploration of how each parameter affects our target moments individually in Appendix H.

Although every parameter affects more than one set of our target moments, there are broad economic forces driving the identification of our model's parameters. The discount factor and bequest motives both tend to increase savings across all groups of households, but in distinct ways. The discount factor affects wealth profiles starting at younger ages, whereas bequest motives tend to raise wealth later in life, typically after age 40.

Terminal and side bequests differ in their effects across groups. Reducing the importance of terminal bequests lowers the wealth of both couples and singles during retirement. In contrast, reducing the desire to leave side bequests increases the wealth of newly singles, who inherit more from their deceased spouse, and thus raises the retirement wealth of all singles.

The strength and luxuriousness of bequest motives further separate their effects. Stronger bequest motives influence a broader set of households and wealth earlier during retirement. Luxury-good bequest motives, instead, mainly affect richer households, who are more likely to survive to older ages. As a result, making bequests more

of a luxury good increases savings more at older ages than at younger ones during the retirement period.

The weight on consumption in the utility function has a large effect on the desire to work for all household groups. In contrast, available time and participation costs mainly affect the labor market behavior of the group they are calibrated for. For couples, there is some spillover from the parameters for husbands to wives and vice versa, but this separation broadly holds.

While available time tends to have the largest effect on hours worked, the level parameter in participation costs has a larger effect on participation. These effects are intuitive for the following reasons. First, an upward shift in the time cost of participating makes participation less attractive. Among those who still participate, there is less time available for work and leisure, so hours worked also decline. Second, a reduction in available time applies to everyone—regardless of participation status and at all ages. As a result, it implies a larger reduction in leisure over the life cycle (until consumption adjusts) than an equivalent increase in the cost of participating. Reducing available time thus leads to a larger decline in hours worked than raising participation costs, especially later in life. Finally, because available time affects both participants and non-participants, it has a smaller effect on the participation margin.

The linear and quadratic terms in participation costs shape the slopes of both participation and hours worked over the life cycle: the linear term matters more early in life, while the quadratic term becomes more important later in the working period.

In sum, the identification of our model's key parameters relies on intuitive economic forces: the life-cycle evolution of wealth and labor supply, heterogeneity across marital groups, and the dynamic responses to risks and incentives. The combination of dynamic optimization, selection into marriage and health, and rich heterogeneity in observed profiles provides substantial leverage for estimating the parameters governing saving, labor supply, and bequest motives. Our approach ensures that key margins are pinned down by meaningful and interpretable patterns in the data.

6 What drives savings and earnings?

We address this question using a sequence of decomposition exercises that alter one feature of the environment at a time, holding all other model inputs and parameters fixed at their benchmark values. This section reports how each change affects life-cycle profiles of labor supply and savings across household groups. The next section summarizes the implications for aggregate savings and earnings over the cohort's life

cycle. Appendix I and Appendix J report additional counterfactuals.

No wage risk. We eliminate wage risk by assuming that individuals retain their initial wage shock from age 26 until retirement and adjusting the deterministic portion of wages by age and gender to keep the exogenous portion of wages unchanged.[9] This preserves initial heterogeneity and the exogenous portion of wages but removes wage risk after entering the economy.

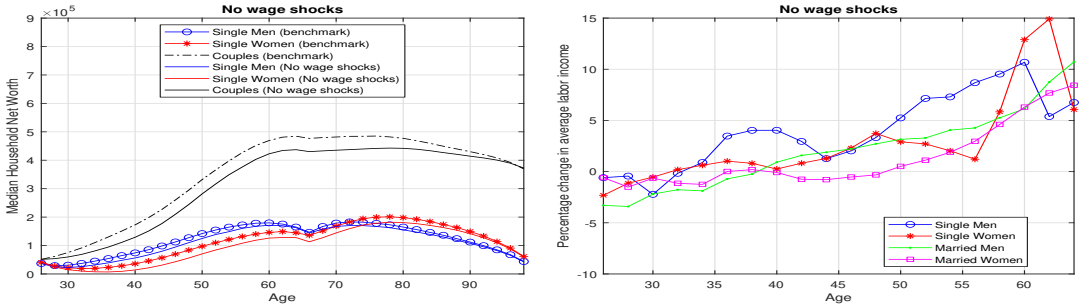


Figure 7: The effects of wage shocks

The left panel of Figure 7 shows that eliminating wage risk reduces median savings for all three groups over most of the life cycle. The largest declines occur around age 38 for single women (68%), age 35 for single men (30%), and age 36 for couples (27%). These patterns are consistent with evidence that precautionary saving against earnings risk is particularly important early in life (e.g., Gourinchas and Parker 2002), but in our framework the effect persists well beyond early working ages because wage risk interacts with other saving motives present later in life.

Eliminating wage risk also affects labor market behavior. The right panel of Figure 7 shows that labor earnings before ages 30–35 are lower for both married and single women, reflecting reduced labor supply when precautionary saving needs decline. At later ages, labor earnings increase as households work more to compensate for lower accumulated income and savings. Thus, wage risk primarily affects the timing of labor supply over the life cycle rather than total lifetime earnings, with lower labor supply early on and higher labor supply later in life. As a result, models that shut down labor supply responses necessarily attribute a much larger share of wealth accumulation to precautionary saving than is warranted once these adjustment margins are available.

[9]For couples, this means that both spouses keep their initial wage shock for all of the working period, including if they divorce. For singles, their own wage shock is fixed. If they marry, they face the same uncertainty about their spouse's wage shock upon marriage, but after that, their spouse's shock will also be fixed during all of the working period.

No bequest motives. We eliminate all voluntary bequest motives by removing utility from bequests at both the death of the first spouse (side bequests) and the death of the last household member (terminal bequests).

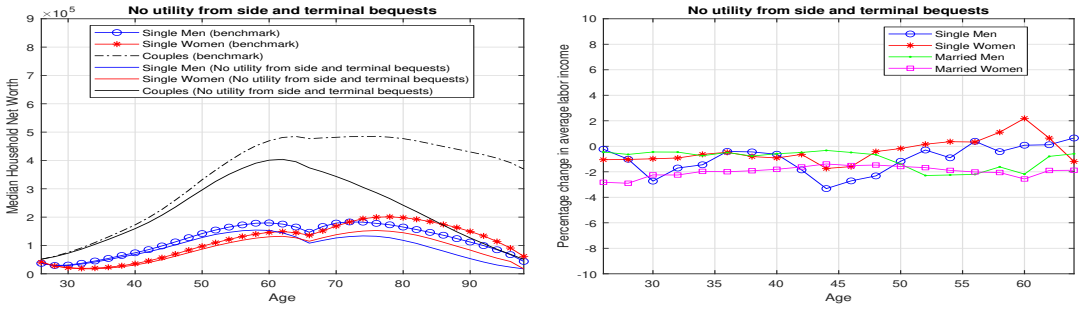


Figure 8: The effects of bequests

The left panel of Figure 8 shows that wealth declines significantly across all household groups. By age 50, median wealth is about 10% lower for all three groups, and by age 80 it falls by roughly 50% for couples, over 25% for single men, and about 20% for single women. Married women's labor income is also lower throughout working life, by nearly 2% relative to the benchmark. These effects indicate that bequest motives influence not only late-life asset decumulation but also labor supply and saving decisions well before retirement.

Notably, when all bequest motives are removed, the model generates a standard hump-shaped wealth profile that is closer to the predictions of a canonical life-cycle model and markedly different from observed saving behavior in retirement. This contrast highlights the role of bequests in sustaining high wealth levels at advanced ages, even in the presence of medical expense and longevity risk.

Appendix I decomposes these effects into side and terminal bequests. The decomposition shows that side bequests are the main driver of wealth accumulation and labor supply responses for couples, while terminal bequests play a relatively larger role for singles, particularly during retirement. Removing both motives jointly therefore eliminates an important force sustaining wealth late in life.

No medical expenses. To study the role of medical expenses as a saving motive in a way that parallels the removal of bequest motives, we eliminate out-of-pocket medical expenditures after age 66. This counterfactual removes both the need to save for medical expenses and the associated risk during retirement.

Medical expenses generate a "disaster risk" of living long and incurring high medical expenses. This low-probability but high-cost risk is difficult to insure through

precautionary saving and increases households' reliance on the consumption floor. Removing medical expenses reduces the size of the negative shocks in case of a disaster, makes self-insurance less costly, and reduces reliance on the consumption floor. As a result, the disincentive effects of the consumption floor on saving, highlighted by Hubbard et al. (1994b,a), are weaker. Consequently, eliminating medical expenses can raise saving incentives for some households, even as it reduces overall precautionary motives.

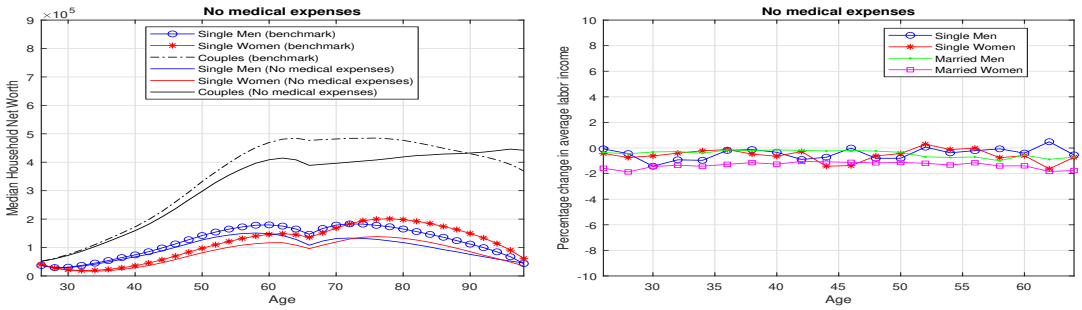


Figure 9: The effects of medical expenses during retirement

The left panel of Figure 9 shows that eliminating medical expenses after age 66 reduces median wealth at all ages for single men and women. For couples, median wealth is lower over most of the life cycle, but becomes higher at very advanced ages, after about age 90, in the absence of medical expenses. This pattern reflects the removal of the implicit saving penalty created by the consumption floor once medical expenses are eliminated.

The right panel displays average labor income. Married women's earnings decline, with the largest gap at age 65, when they earn 2% less than in the benchmark case. Married men earn slightly less over all of their working lives, while the earnings of single men and single women remain largely unchanged.

Appendix I reports the counterfactual that eliminates medical expenses at death and shows that this change has little effect on labor income and wealth prior to old age. This reflects the fact that medical expenses at death typically occur very late in life and therefore have limited effects on behavior before retirement.

No bequest motives and no medical expenses. To evaluate the interaction between saving motives and assess whether their effects are additive, we also consider counterfactuals that remove more than one key element of the model. Here, we focus on the joint removal of bequest motives and medical expenses.

The left panel of Figure 10 reports median wealth in the benchmark economy

and in this counterfactual. Under the joint removal of bequest motives and medical expenses, median household wealth at age 55 is about 25% lower for all household types. By age 70, median wealth is about 45% lower for single women and roughly 55% lower for couples and single men. These results indicate that bequest motives and medical expenses are important drivers of wealth accumulation throughout the life cycle.

Comparing across counterfactuals reveals that these effects are not additive. For couples, eliminating bequest motives alone leads to a large decline in median wealth at age 90, while eliminating medical expenses alone has little effect at that age. When both motives are removed, median wealth for couples at age 90 falls to near zero. Thus, the combined effect cannot be inferred from single-motive experiments and reflects a strong interaction between late-life saving motives. We quantify these interactions further in the next section.

Even in the absence of bequest motives and medical expenses, households continue to hold substantial wealth at older ages. This contrasts with a standard life-cycle model without lifespan uncertainty and reflects continued saving for uncertain longevity and health transitions, which remain operative in this counterfactual.

The right panel of Figure 10 shows that removing bequest motives and medical expenses also affects labor market behavior. Labor income declines for married men and women throughout their working lives and for single men over most of their careers, with smaller effects for single women. As a result, consumption increases in this counterfactual, particularly for couples and single men, while single women allocate relatively more resources to leisure.

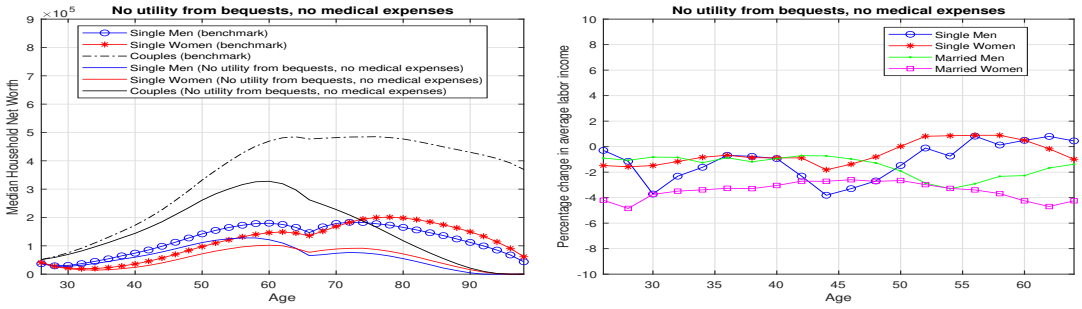


Figure 10: The effects of bequest motives and medical expenses during retirement

No marriage and divorce risks. To assess the role of marital dynamics, we eliminate both marriage formation and divorce after age 26. This counterfactual shuts down changes in marital status over the life cycle while preserving the initial

marital composition.

The top left panel of Figure 11 shows that eliminating marriage and divorce leaves the fraction of married men largely unchanged, while the fraction of married women before retirement rises from 80% to 87%, matching the share of married women at age 26. The bottom left panel shows that married households hold less wealth throughout

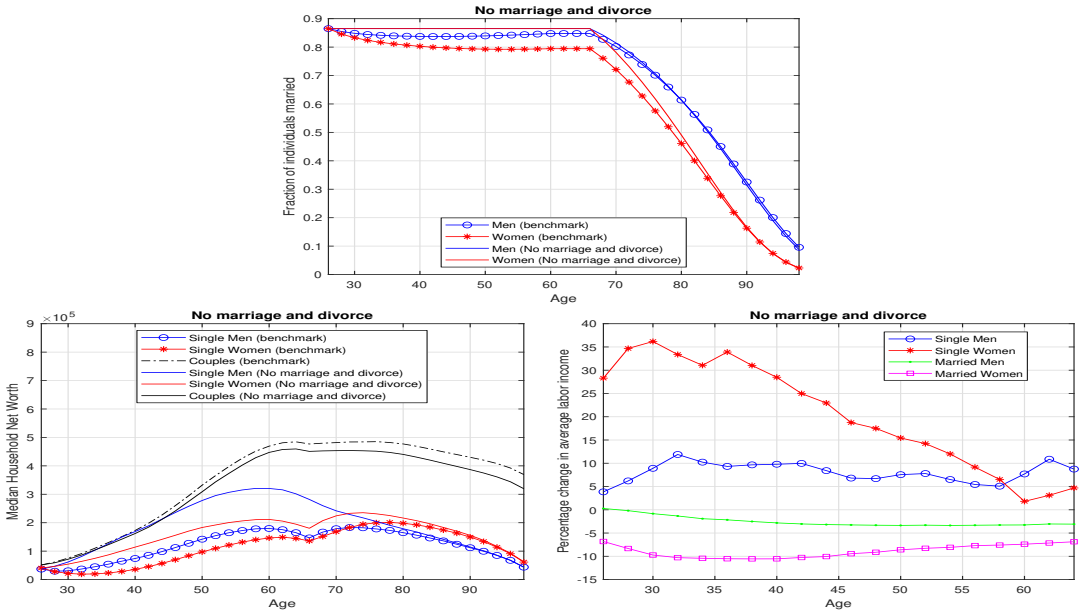


Figure 11: The effects of no new marriages and divorces

the life cycle, reflecting the removal of divorce risk and the associated precautionary saving motive. Single men and women instead hold more wealth at all ages, as they must rely solely on their own resources. After retirement, this pattern also reflects that wealthier widows and widowers enter the group of singles.

The bottom right panel shows the corresponding labor supply responses. Single men and women finance their higher wealth accumulation through increased labor supply, particularly early in life. Married men and women reduce labor supply once marital risk is removed. These offsetting group-specific responses explain why marital transitions have large behavioral effects but small aggregate implications.

Appendix I separates marriage formation and divorce. Removing marriage formation primarily affects behavior early in life through expectations about future household formation and composition effects among singles, while removing divorce mainly affects saving and labor supply later in life by reducing precautionary motives within couples.

None of the previous saving motives. We eliminate utility from bequests, medical expenses, wage risk, and marriage and divorce dynamics. The remaining motives are to finance retirement consumption and to self-insure against longevity risk. Figure 12 shows that for couples, wealth accumulation falls sharply: median wealth is about 50% lower than in the baseline between ages 50 and 60 and nearly zero by age 90. For singles, the absence of marriage prospects dominates early behavior, leading to substantially higher wealth accumulation at young ages. After retirement, however, wealth declines sharply for singles as well. Labor market responses are large and heterogeneous. Married women's earnings fall by 15–20% throughout working life. Married men's earnings are initially lower but rise above the benchmark by retirement. Single women's earnings are more than 25% higher early in life, while single men's earnings change little. This counterfactual provides a lower bound on wealth accumulation once all major saving motives are removed, clarifying the role of retirement consumption and longevity risk alone.

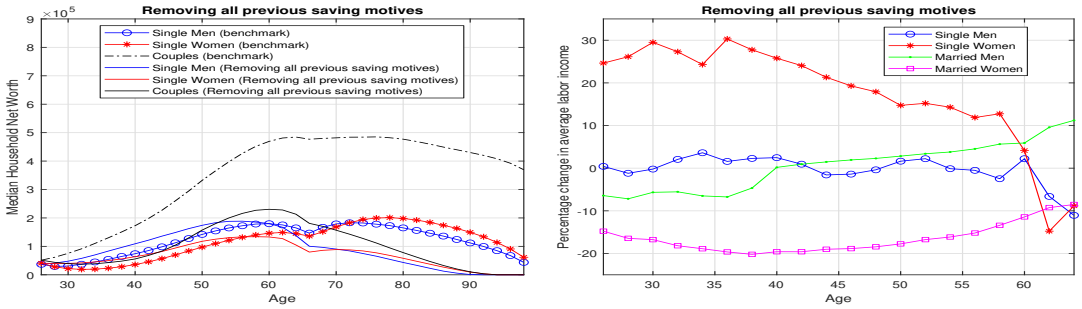


Figure 12: The effects of removing all of the previously considered saving motives: utility from leaving both side bequests and terminal bequests, medical expenses, wage risks, and marriage and divorce dynamics.

7 What drives aggregate savings and earnings?

In this section, we aggregate the life-cycle results discussed above, first by household type and then overall, and compare average outcomes in the benchmark and counterfactual economies. If demographics were stationary, the reported changes in average wealth and earnings over the life cycle would correspond to changes in aggregate outcomes. Appendix J reports aggregate results for additional counterfactuals.

Table 3 reports how various risks and saving motives account for average wealth accumulation over the life cycle, while Table 4 reports how the same forces affect average lifetime earnings. Together, these tables summarize the aggregate implications

of the life-cycle mechanisms studied in Section 6 and clarify how their quantitative importance changes once risks and adjustment margins are modeled jointly.

Counterfactual	Couples	SM	SW	All
No wage shocks	-10.3	-9.4	-11.8	-10.4
No utility from bequests	-25.0	-21.0	-18.6	-23.8
No medical expenses	-11.3	-16.0	-21.9	-13.1
No utility from bequests + no med. exp.	-39.8	-42.6	-46.1	-40.9
No marriage and divorce	-7.2	41.4	24.1	0.7
Removing all previous saving motives	-60.6	-35.0	-48.3	-56.9

Table 3: Changes in average wealth per person (in % from our benchmark economy). SM: single men, SW: single women

Aggregate wealth. Table 3 shows that several forces contribute to average wealth accumulation over the life cycle, with substantial differences in their quantitative importance. Removing wage risk lowers average wealth by 10.4% overall, with similar declines across couples (10.3%), single men (9.4%), and single women (11.8%). Wage risk therefore accounts for a non-trivial share of aggregate savings. This estimate is close to the upper bound of the precautionary wealth effects found in stylized incomplete-markets models without life-cycle structure, which find effects on the order of 5–14% (e.g., Aiyagari 1994), but much smaller than estimates from working-life life-cycle models that incorporate richer earnings risk while abstracting from endogenous labor supply and other saving motives. For example, Gourinchas and Parker (2002) and Cagetti (2003) attribute roughly 50–70% of wealth accumulation to precautionary saving against earnings risk. The difference reflects a structural distinction: in unified life-cycle models with endogenous labor supply and family insurance, savings are no longer the residual adjustment margin.

This comparison matters because it changes the interpretation of precautionary saving as a dominant driver of wealth accumulation. In our unified framework, households can adjust labor supply over the full life cycle and face competing late-life saving motives, including bequests and medical expense risk. Once these margins are available, precautionary saving against wage risk no longer operates in isolation and accounts for a smaller share of aggregate wealth accumulation.

Bequest motives are the single most important contributor to aggregate wealth. Eliminating the desire to leave bequests lowers average wealth by 23.8%, with declines of 25.0% for couples, 21.0% for single men, and 18.6% for single women. Bequest

motives thus account for nearly one-quarter of average wealth accumulated over the life cycle.

This estimate is well below the classic accounting results of Kotlikoff and Summers (1981), who attribute roughly 80% of aggregate wealth to intergenerational transfers, but is close to the lower estimates emphasized by Modigliani (1988), who argues that bequests account for about 20% of aggregate wealth. Subsequent empirical and structural work has produced intermediate estimates. For example, Gale and Scholz (1994) estimate that bequests account for roughly 30-55% of aggregate wealth, while more recent structural models of retirement behavior typically find bequest-related wealth shares in the range of 15–35% (e.g., Lockwood 2012, 2018; Ameriks et al. 2020a). Our estimate of 23.8% therefore lies within the range of modern quantitative estimates, while remaining well below the upper bounds implied by early accounting exercises.

A key distinction relative to retirement-focused analyses is that we allow bequest motives to affect saving and labor supply decisions over the full life cycle. In De Nardi et al. (2025), eliminating bequest motives reduces average wealth by 14.8%, with the effect concentrated late in life. In that framework, households enter retirement with given wealth and cannot adjust labor supply or saving behavior earlier in life. In contrast, in our framework households work less and save less earlier in life when bequest motives are removed. As we also account for bequests made at the death of the first spouse and of the last household member, our estimates capture both late-life and working-life channels through which bequest motives operate, leading to a larger life-cycle average effect.

Medical expenses are the second most important force for wealth accumulation. Eliminating them lowers average wealth by 13.1% overall, with declines of 11.3% for couples, 16.0% for single men, and 21.9% for single women.

This estimate is substantially larger than those obtained in retirement-focused frameworks that condition on wealth at retirement and abstract from working-age adjustments. For example, De Nardi et al. (2025) find that eliminating medical expenses lowers average wealth by about 3.1%, with the effect concentrated late in life. Life-cycle models that allow saving to adjust over the full life span but hold labor supply exogenous and abstract from marital dynamics find larger effects. In particular, Kopecky and Koreshkova (2014) estimate that out-of-pocket medical expenses account for about 13.5% of aggregate wealth. Our estimate of 13.1% is close to theirs, but arises in a framework that additionally allows labor supply and family structure to change over the life cycle.

The joint removal of bequest motives and medical expenses reduces average wealth by 40.9%, substantially more than either force alone. De Nardi et al. 2025 find a very similar joint effect in a retirement-focused framework, with average wealth falling by about 42.5%. This similarity indicates that strong interactions between bequest motives and medical expense risk arise primarily from late-life saving behavior and do not rely on working-age labor supply adjustments. In our model, allowing households to adjust saving and labor supply earlier in life changes the quantitative importance of the individual motives, but does not eliminate the powerful interaction between them.

Marriage and divorce dynamics have a small effect on average wealth in the aggregate: eliminating marriage and divorce lowers average wealth by only 0.7%. As shown in Section 6, however, this small aggregate effect reflects large but offsetting responses across household types, and therefore depends critically on demographic composition. To our knowledge, existing quantitative life-cycle models have not previously assessed the aggregate wealth implications of marital transitions in an environment with endogenous labor supply and multiple saving motives.

Removing all previously considered saving motives—bequests, medical expenses, wage risk, and marriage and divorce dynamics—reduces average wealth by 56.9% overall, with declines of 60.6% for couples, 35.0% for single men, and 48.3% for single women. Even in this extreme counterfactual, households continue to hold positive wealth, reflecting saving for retirement consumption under uncertain lifespans.

Taken together, Table 3 shows that bequest motives are the largest force shaping aggregate wealth accumulation, followed by medical expenses, wage risk, and marriage and divorce dynamics, and that the quantitative role of each motive depends on the presence of other risks and adjustment margins.

Comparison with existing quantitative conclusions about aggregate wealth. Taken together, these results show that the quantitative importance of individual saving motives depends critically on which risks and adjustment margins are modeled jointly. Precautionary saving against wage risk accounts for a non-trivial but much smaller share of aggregate wealth than in working-life models with fixed labor supply, while bequest motives and medical expense risk are dominant drivers of wealth accumulation. At the same time, strong interactions among late-life saving motives are present even in retirement-only frameworks, as shown by the similarity of the joint bequest-and-medical-expense counterfactual to De Nardi et al. 2025. Allowing households to adjust labor supply and saving earlier in life primarily reallocates

the quantitative importance across individual motives, rather than eliminating their interaction, and changes how saving behavior responds over the full life cycle.

Counterfactual	MM	MW	SM	SW	All
No wage shocks	2.3	1.5	4.2	3.0	2.3
No utility from bequests	-1.0	-1.9	-1.1	-0.3	-1.2
No medical expenses	-0.4	-1.3	-0.4	-0.6	-0.7
No utility from bequests + no med. exp.	-1.6	-3.4	-1.3	-0.4	-1.9
No marriage and divorce	-2.9	-10.1	7.4	19.3	-2.0
Removing all previous saving motives	0.2	-17.6	-0.6	16.9	-2.7

Table 4: Changes in average labor income. MM: married men, MW: married women, SM: single men, SW: single women

Aggregate earnings. Table 4 shows a different ranking of forces affecting average lifetime earnings. Wage risk is the dominant force shaping aggregate earnings. Eliminating wage risk raises average earnings by 2.3% overall, with increases of 2.3% for married men, 1.5% for married women, 4.2% for single men, and 3.0% for single women. As discussed in Section 6, this positive aggregate response reflects offsetting life-cycle adjustments: households work less earlier in life when wage risk is eliminated, but increase labor supply later on to offset lower accumulated savings and income from assets. As a result, changes in wage risk mainly affect the timing of work rather than total lifetime earnings, consistent with structural life-cycle models with endogenous labor supply Low, Meghir, and Pistaferri (2010); Blundell, Pistaferri, and Saporta-Eksten (2016b).

The modest magnitude of this aggregate earnings response is consistent with the labor-supply-with-risk literature, which finds that allowing households to adjust labor supply substantially attenuates the impact of wage risk on lifetime earnings, even when labor supply responses themselves are sizable (e.g., Low 2005; Pijoan-Mas 2006; Blundell et al. 2016b).

Marriage and divorce dynamics are the second most important force for earnings. Eliminating marriage and divorce lowers average earnings by 2.0% overall, with particularly large declines for married women (10.1%) and increases for singles (7.4% for single men and 19.3% for single women). Divorce risk alone accounts for 2.2% of average earnings over the life cycle, reflecting its effect on labor supply incentives within couples. Existing quantitative work on family labor supply emphasizes large hours and participation responses to household composition and wage shocks, particularly

for women (e.g., Blundell et al. 2016b; Attanasio et al. 2018), but does not aggregate these responses into lifetime earnings effects. To our knowledge, the aggregate earnings implications of marital transitions have not previously been quantified in a life-cycle model with endogenous labor supply.

Bequest motives and medical expenses play a smaller role for aggregate earnings. Eliminating bequest motives lowers average earnings by 1.2%, while eliminating medical expenses lowers earnings by 0.7%.

The joint removal of bequest motives and medical expenses lowers average earnings by 1.9%. Removing all saving motives lowers average earnings by 2.7% overall, with heterogeneous responses across household types.

Although the aggregate earnings effects are smaller in magnitude than those for wealth, Table 4 shows that earnings responses are also not additive. The combined effect of removing multiple motives differs from the sum of the corresponding single-motive effects, reflecting interacting labor supply adjustments over the life cycle when multiple risks and incentives change simultaneously. This implies that extrapolating aggregate earnings effects from single-risk elasticities can be misleading.

Putting the results together. Taken together, Tables 3 and 4 show that different forces govern aggregate wealth and aggregate earnings. Wealth accumulation is driven primarily by late-life saving motives, such as bequests and medical expenses, while earnings are shaped mainly by working-life risks and family structure, particularly wage risk and marriage and divorce dynamics. Although some aggregate effects, especially for earnings, are modest in size, they arise from large and offsetting group-specific responses that would be missed in models that abstract from household structure or endogenous labor supply.

Finally, removing all saving motives except retirement consumption and longevity risk lowers average wealth by 56.9% and average earnings by 2.7%. This result underscores that saving for an uncertain retirement length remains a sizable determinant of household behavior even after accounting for all additional saving motives considered in the model.

8 Conclusions and directions for future work

This paper develops and estimates a dynamic life-cycle model that captures how wage risk, marital dynamics, health and longevity risks, bequest motives, and medical expenses shape savings and labor supply. The model jointly tracks the behavior of

single and married individuals from labor market entry until death, accounts for selection into marriage and divorce, and allows health and survival to depend on income and prior choices, consistent with the data.

We show that these risks and motives play distinct roles across groups and stages of the life cycle. Bequest motives, medical expenses, and wage risk each account for substantial shares of aggregate wealth, with bequest motives alone responsible for 23.8%. Medical expenses account for a further 13.1%, and wage risk for 10.4%. While the effects of marital dynamics on aggregate wealth are modest in the aggregate, they generate large heterogeneity by group. Wage risk, in particular, plays a dominant role in early-life behavior, substantially increasing savings at younger ages. On the labor margin, marital dynamics and wage risk significantly alter earnings profiles, with aggregate earnings falling when marriage and divorce transitions are removed, and rising when wage risk is eliminated due to increased late-career labor supply. Removing all saving motives other than retirement needs and uncertain lifetimes reduces lifetime wealth accumulation by 56.9% and aggregate labor earnings by 2.7%.

A central contribution of the paper is to clarify how conclusions drawn from models that isolate individual risks change once households face multiple risks and adjustment margins jointly. Relative to working-life life-cycle models with fixed labor supply, we find that precautionary saving against wage risk accounts for a much smaller share of aggregate wealth once labor supply and within-household insurance are available. Relative to retirement-only frameworks, we find that bequest motives and medical expense risk have larger life-cycle effects because households adjust saving and labor supply earlier in life in anticipation of late-life risks. At the same time, strong interactions between bequest motives and medical expenses persist even in retirement-only settings, indicating that some interactions arise primarily from late-life behavior rather than working-age adjustments.

These findings underscore the importance of jointly modeling wage, marital, and health risks, bequest motives, and medical expenses to understand the distribution and dynamics of savings and labor supply. Even when aggregate effects appear modest, the underlying group-specific responses are large and policy-relevant, and models that abstract from key risks or adjustment margins may misstate both the magnitude and the channels through which policies affect household behavior. By integrating these dimensions into a unified life-cycle framework, our paper provides a more complete understanding of household behavior and the channels through which risk and policy interact.

Our model makes several simplifying assumptions. First, we hold rates of re-

turn constant. This abstraction allows us to focus on the effects of wage, marital, health, and longevity risks without introducing additional intertemporal price variation. That said, incorporating trend variation or return risk is a natural extension and appears computationally feasible within our framework—at least as a robustness exercise. Second, we treat marriage and divorce as exogenous. Although empirical work suggests limited responsiveness to economic incentives (Dickert-Conlin and Meghea, 2004; Goda, Shoven, and Slavov, 2007; Dillender, 2016; Alm and Whittington, 1995, 1997, 1999), larger shifts in policy or norms could alter marriage patterns in ways that warrant endogenous modeling. Third, we fix home production and do not allow time use to adjust in response to counterfactuals. Given the large labor supply responses we estimate—especially for women—allowing home production to adjust would be a valuable extension. Finally, we treat fertility as exogenous and simplify the distribution of children by maternal age and marital status. While relaxing this additional set of assumptions would enhance realism, it would also introduce substantial computational complexity.

References

Abel, Andrew B. and Mark Warshawsky (1988), "Specification of the joy of giving: Insights from altruism." *The Review of Economics and Statistics*, 70, 145–149.

Aguiar, Mark and Erik Hurst (2007), "Measuring trends in leisure: The allocation of time over five decades." *Quarterly Journal of Economics*, 122, 969–1006.

Aiyagari, S. Rao (1994), "Uninsured idiosyncratic risk and aggregate saving." *The Quarterly Journal of Economics*, 109, 659–684.

Alm, James and Leslie A. Whittington (1995), "Income taxes and the marriage decision." *Applied Economics*, 27, 25–31.

Alm, James and Leslie A. Whittington (1997), "'til death or taxes do us part: The effect of income taxation on divorce." *Journal of Human Resources*, 32, 388–412.

Alm, James and Leslie A. Whittington (1999), "For love or money? the impact of income taxes on marriage." *Economica*, 66, 297–316.

Altonji, Joseph G and Lewis M Segal (1996), "Small-sample bias in gmm estimation of covariance structures." *Journal of Business & Economic Statistics*, 14, 353–366.

Ameriks, John, Joseph Briggs, Andrew Caplin, Matthew D. Shapiro, and Chris Tonetti (2020a), "Long-term care utility and late in life saving." *Journal of Political Economy*, 128, 2375–2451.

Ameriks, John, Andrew Caplin, Minjoon Lee, Matthew D. Shapiro, and Chris Tonetti (2020b), "Older americans would work longer if jobs were flexible." *American Economic Journal: Macroeconomics*, 12, 174–211.

Andreoni, James (1989), "Giving with impure altruism: Applications to charity and ricardian equivalence." *Journal of Political Economy*, 97, 1447–1458.

Attanasio, Orazio, Peter Levell, Hamish Low, and Virginia Sánchez-Marcos (2018), "Aggregating elasticities: Intensive and extensive margins of women's labor supply." *Econometrica*, 86, 2049–2082.

Becker, Gary S. and Nigel Tomes (1986), "Human capital and the rise and fall of families." *Journal of Labor Economics*, 4, 1–39.

Benabou, Roland (2002), "Tax and education policy in a heterogeneous-agent economy: What levels of redistribution maximize growth and efficiency?" *Econometrica*, 70, 481–517.

Bernheim, B. Douglas, Andrei Shleifer, and Lawrence H. Summers (1985), "The strategic bequest motive." *Journal of Political Economy*, 93, 1045–1076.

Blundell, Richard, Monica Costa Dias, Costas Meghir, and Jonathan Shaw (2016a), "Female labor supply, human capital, and welfare reform." *Econometrica*, 84, 1705–1753.

Blundell, Richard and Thomas Macurdy (1999), "Labor supply: A review of alternative approaches." In *Handbook of Labor Economics* (O. Ashenfelter and D. Card, eds.), volume 3, 1559–1695, Elsevier.

Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten (2016b), "Consumption inequality and family labor supply." *American Economic Review*, 106, 387–435.

Borella, Margherita, Mariacristina De Nardi, and Eric French (2018), "Who receives medicaid in old age? rules and reality." *Fiscal Studies*, 39, 65–93.

Borella, Margherita, Mariacristina De Nardi, Michael Pak, Nicolo Russo, and Fang Yang (2023a), "Heterogeneous agents, taxation, and social security." *Journal of the European Economic Association*, 21, 129–169.

Borella, Margherita, Mariacristina De Nardi, and Fang Yang (2023b), "Are marriage-related taxes and social security benefits holding back female labour supply?" *The Review of Economic Studies*, 90, 102–131.

Brown, Meta (2003), "End-of-life transfers and the decision to care for a parent." Mimeo.

Butrica, Barbara and Richard Burkhauser (1997), "Estimating federal income tax burdens for panel study of income dynamics (PSID) families using the national bureau of economic research taxsim model." Working Paper, Maxwell Center for Demography and Economics of Aging, Syracuse University.

Cagetti, Marco (2003), "Wealth accumulation over the life cycle and precautionary savings." *Journal of Business and Economic Statistics*, 21, 339–353.

Carroll, Christopher D. (1998), "Why do the rich save so much?" NBER Working Paper 6549.

Citro, Constance F. and Robert T. Michael (1995), "Adjusting poverty thresholds." In *Measuring Poverty: A New Approach* (Constance F. Citro and Robert T. Michael, eds.), 159–201, The National Academies Press, Washington, DC.

Costa Dias, Monica, Robert Joyce, and Francesca Parodi (2020), "The gender pay gap in the uk: children and experience in work." *Oxford Review of Economic Policy*, 36, 855–881.

Cubeddu, Luis and José-Víctor Ríos-Rull (2003), "Families as shocks." *Journal of the European Economic Association*, 1, 671–682.

David, Martin, Roderick J. A. Little, Michael E. Samuhel, and Robert K. Triest (1986), "Alternative methods for cps income imputation." *Journal of the American Statistical Association*, 81, 29–41.

De Nardi, Mariacristina (2004), "Wealth inequality and intergenerational links." *Review of Economic Studies*, 71, 743–768.

De Nardi, Mariacristina, Eric French, John Bailey Jones, and Rory McGee (2025), "Why do couples and singles save during retirement? household heterogeneity and its aggregate implications." *Journal of Political Economy*.

Dickert-Conlin, Stacy and Cristian Meghea (2004), "The effect of social security on divorce and remarriage behavior." *Boston College, Center for Retirement Research Working Paper*.

Dillender, Marcus (2016), "Social security and divorce." *B. E. Journal of Economic Analysis and Policy*, 16, 931–971.

Dotsey, Michael, Wenli Li, and Fang Yang (2014), "Consumption and time use over the life cycle." *International Economic Review*, 55, 665–692.

De Nardi, Mariacristina, Eric French, and John Bailey Jones (2010), "Why do the elderly save? the role of medical expenses." *Journal of Political Economy*, 118, 39–75.

De Nardi, Mariacristina, Eric French, and John Bailey Jones (2016), "Medicaid insurance in old age." *American Economic Review*, 106, 3480–3520.

Duffie, Darrell and Kenneth J. Singleton (1993), "Simulated moments estimation of markov models of asset prices." *Econometrica*, 61, 929–952.

French, Eric (2000), "The effects of health, wealth, and wages on labor supply and retirement behavior." Working Paper, Federal Reserve Bank of Chicago.

French, Eric and John Bailey Jones (2011), "The effects of health insurance and self-insurance on retirement behavior." *Econometrica*, 79, 693–732.

Gale, William G. and John Karl Scholz (1994), "Intergenerational transfers and the accumulation of wealth." *Journal of Economic Perspectives*, 8, 145–160.

Goda, Gopi Shah, John B. Shoven, and Sita Nataraj Slavov (2007), "Social security and the timing of divorce." NBER Working Paper 13382.

Gourinchas, Pierre-Olivier and Jonathan A. Parker (2002), "Consumption over the life cycle." *Econometrica*, 70, 47–89.

Hall, Robert E. and Charles I. Jones (2007), "The value of life and the rise in health spending." *The Quarterly Journal of Economics*, 122, 39–72.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes (1994a), "Expanding the life-cycle model: Precautionary saving and public policy." *American Economic Review*, 84, 174–179.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes (1994b), "The importance of precautionary motives in explaining individual and aggregate saving." *Carnegie Rochester Series on Public Policy*, 59–125.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes (1995), "Precautionary saving and social insurance." *Journal of Political Economy*, 103, 360–399.

Hurd, Michael D. (1989), "Mortality risk and bequests." *Econometrica*, 57, 779–813.

Kimberlin, Sara, Jiyoon Kim, and Luke Shaefer (2015), "An updated method for calculating income and payroll taxes from PSID data using the NBER's TAXSIM, for PSID survey years 1999 through 2011." University of Michigan manuscript.

Kopecky, Karen and Tatyana Koreshkova (2014), "The impact of medical and nursing home expenses and social insurance policies on savings and inequality." *American Economic Journal: Macroeconomics*, 6, 29–72.

Kotlikoff, Laurence J. and Lawrence H. Summers (1981), "The role of intergenerational transfers in aggregate capital accumulation." *Journal of Political Economy*, 89, 706–732.

Liebman, Jeffrey B., Erzo F.P. Luttmer, and David G. Seif (2009), "Labor supply responses to marginal social security benefits: Evidence from discontinuity." *Journal of Public Economics*, 93, 1208–1223.

Lockwood, Lee M. (2012), "Bequest motives and the annuity puzzle." *Review of Economic Dynamics*, 15, 226–243.

Lockwood, Lee M. (2018), "Incidental bequests and the choice to self-insure late-life risks." *American Economic Review*, 108, 2513–50.

Low, Hamish, Costas Meghir, and Luigi Pistaferri (2010), "Wage risk and employment risk over the life cycle." *American Economic Review*, 100, 1432–1467.

Low, Hamish W. (2005), "Self-insurance in a life-cycle model of labour supply and savings." *Review of Economic Dynamics*, 8, 945–975.

Modigliani, Franco (1988), "Measuring the constribution of intergenerational transfers to total wealth: Conceptual issues and empirical findings." In *Modeling the Accumulation and Distribution of Wealth* (Denis Kessler and Andre Masson, eds.), Clarendon Press, Oxford.

Pakes, Ariel and David Pollard (1989), "Simulation and the aysmptotics of optimization estimators." *Econometrica*, 57, 1027–1057.

Pijoan-Mas, Josep (2006), "Precautionary savings or working longer hours?" *Review of Economic Dynamics*, 9, 326–352.

Online appendix

The appendices are organized as follows. Appendix A describes the PSID and HRS data and sample construction. Appendix B presents additional value functions and recursive problems omitted from the main text. Appendix C details the first-step estimation of exogenous processes, including wages, marriage and divorce, health dynamics, medical expenses, survival, and taxes. Appendix E describes the numerical solution algorithm. Appendix F outlines the estimation strategy, while Appendix H discusses parameter identification. The remaining appendices report parameter estimates, counterfactual experiments, and additional results.

A Data: The PSID and the HRS

The PSID is a longitudinal study of a representative sample of the U.S. population that started in 1968. We select all individuals in the main sample who were interviewed at least twice during the years 1968-2019 (34,856 individuals); only heads and their wives (including cohabiting partners), if present (20,759); people born between 1941 and 1965 (7,581); and those between the ages of 23 and 70 (6,721). The resulting sample includes 121,603 observations.

The HRS is a biennial survey that started collecting information on people born between 1931 and 1941 in 1992. Other cohorts were introduced over the years. Our data set is based on the RAND HRS files for the period 1995-2018 (waves 3 to 14), to which we add the Exit files to include information on the wave right after death. Our sample selection is as follows. Of the 38,915 people initially present, we drop people for whom marital status is not observed (11 people). We then select people born between 1911 to 1965 and drop people older than 100, obtaining a sample of 37,287 people and 225,073 observations.

All monetary values reported throughout the paper are in 2018 U.S. dollars.

B Additional value functions

B.1 Singles during the early retirement stage

The recursive problem for singles ($j = 1$) who has claimed Social Security at age tr but is younger than age 66 is

$$S^s(t, i, a_t^i, \bar{y}_r^i, tr) = \max_{c_t, a_{t+1}} \left(v^i(c_t, L^{i,j}, \eta_t^{i,j}) + \beta E_t S^s(t+1, i, a_{t+1}^i, \bar{y}_r^i, tr) \right), \tag{28}$$

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

$$Y_t = SS(\bar{y}_r^i, tr), \tag{29}$$

$$c_t + a_{t+1}^i = (1+r)a_t^i + Y_t - T(\cdot). \tag{30}$$

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 and decide not to claim benefits this period.

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

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

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

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 claim or not, where D_t^i is an indicator function for claiming

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

B.2 Couples during the early retirement period

The recursive problem for couples ($j = 2$) that have claimed Social Security at age tr but are younger than 66 is

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

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

$$Y_t = \max\left\{ (SS(\bar{y}_r^1, tr) + SS(\bar{y}_r^2, tr), \frac{3}{2}\max(SS(\bar{y}_r^1, tr), SS(\bar{y}_r^2, tr)) \right\}. \quad (35)$$

The variable Y_t represents the Social Security spousal benefit: a married person receives the highest amount between one's own benefit and half of their spouse's benefit.

The value function of a couple that has not yet claimed benefits and decides not to claim benefits this period is

$$\begin{aligned} N^c(t, a_t, \epsilon_t^1, \epsilon_t^2, \bar{y}_t^1, \bar{y}_t^2) = \max_{c_t, a_{t+1}, n_t^1, n_t^2} \Bigg(& w(c_t, l_t^{1,j}, l_t^{2,j}, \eta_t^{i,j}) \\ & + \beta E_t V^c(t+1, a_{t+1}, \epsilon_{t+1}^1, \epsilon_{t+1}^2, \bar{y}_{t+1}^1, \bar{y}_{t+1}^2) \Bigg), \end{aligned} \quad (36)$$

subject to equations (1), (8), (12), (20), (22), and

$$c_t + a_{t+1} = (1+r)a_t + Y_t^1 + Y_t^2 - \tau_t^{SS}(\min(Y_t^1, \widetilde{y}_t) + \min(Y_t^2, \widetilde{y}_t)) - T(\cdot). \quad (37)$$

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

$$V^c(t, a_t, \epsilon_t^1, \epsilon_t^2, \bar{y}_t^1, \bar{y}_t^2) = \max_{D_t} \Big((1 - D_t) N^c(t, a_t, \epsilon_t^1, \epsilon_t^2, \bar{y}_t^1, \bar{y}_t^2) + D_t S^c(t, a_t, \bar{y}_t^1, \bar{y}_t^2, t) \Big). \quad (38)$$

B.3 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 marrying for a single person, the relevant object for that person is his or her discounted present value of utility in the marriage.

We compute this object for a person of gender i who is married to a specific partner by using the policy functions of the corresponding couple and the present and future utility of the person for whom we are computing the value function.

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

$$\begin{aligned}\hat{W}^c(t,i,a_t,\epsilon_t^1,\epsilon_t^2,\bar{y}_t^1,\bar{y}_t^2) = v^i(\hat{c}_t(\cdot),\hat{l}_t^{i,j},\eta_t^{i,j})+ \\ \beta(1-\zeta(\cdot))E_t\hat{W}^c(t+1,i,\hat{a}_{t+1}(\cdot),\epsilon_{t+1}^1,\epsilon_{t+1}^2,\bar{y}_{t+1}^1,\bar{y}_{t+1}^2)+ \\ \beta\zeta(\cdot)E_tW^s(t+1,i,\hat{a}_{t+1}(\cdot)/2,\epsilon_{t+1}^i,\bar{y}_{t+1}^i).\end{aligned} \tag{39}$$

During the early retirement period, we have

$$\begin{aligned}\hat{N}^c(t,i,a_t,\epsilon_t^1,\epsilon_t^2,\bar{y}_t^1,\bar{y}_t^2) = v^i(\hat{c}_t(\cdot),\hat{l}_t^{i,j},\eta_t^{i,j}) \\ +\beta E_t\hat{V}^c(t+1,i,\hat{a}_{t+1}(\cdot),\epsilon_{t+1}^1,\epsilon_{t+1}^2,\bar{y}_{t+1}^1,\bar{y}_{t+1}^2)\end{aligned} \tag{40}$$

$$\hat{S}^c(t,i,a_t,\bar{y}_r^1,\bar{y}_r^2,tr) = v^i(\hat{c}_t(\cdot),L^{i,j},\eta_t^{i,j}) + \beta E_tS^c(t+1,i,\hat{a}_{t+1}(\cdot),\bar{y}_r^1,\bar{y}_r^2,tr) \tag{41}$$

$$\hat{V}^c(t,i,a_t,\epsilon_t^1,\epsilon_t^2,\bar{y}_t^1,\bar{y}_t^2) = (1-\hat{D}_t(\cdot))\hat{N}^c(t,i,a_t,\epsilon_t^1,\epsilon_t^2,\bar{y}_t^1,\bar{y}_t^2) + \hat{D}_t(\cdot)\hat{S}^c(t,i,a_t,\bar{y}_r^1,\bar{y}_r^2,t). \tag{42}$$

During the retirement period, we have

$$\begin{aligned}&\hat{R}^c(t,i,a_t,h_t^1,h_t^2,\bar{y}_r^2,\bar{y}_r^1,\bar{y}_r^2,tr) = v^i(\hat{c}_t(\cdot),L^{i,j},\eta_t^{i,j})+ \\ &\beta s_t^{i,j}(\cdot)s_t^{p,j}(\cdot)E_t\hat{R}^c(t+1,i,\hat{a}_{t+1}(\cdot),h_{t+1}^1,h_{t+1}^2,\bar{y}_r^1,\bar{y}_r^2,tr)+ \\ &\beta s_t^{i,j}(\cdot)(1-s_t^{p,j}(\cdot))E_tR^s(t+1,i,\hat{a}_{t+1}(\cdot)-\hat{e}_t^p(\cdot)-d_t^{p,j}\left(h_t^i,\bar{y}_r^i,\bar{y}_r^p\right),h_{t+1}^i,\bar{\bar{y}}_r^i,tr)+ \\ &\beta(1-s_t^{i,j}(\cdot))s_t^{p,j}(\cdot)\theta_1(\hat{e}_t^i(\cdot))+ \\ &\beta(1-s_t^{i,j}(\cdot))(1-s_t^{p,j}(\cdot))\theta_0(\max(0,(\hat{a}_{t+1}-d_t^{i,j}\left(h_t^i,\bar{y}_r^i,\bar{y}_r^p\right)-d_t^{p,j}\left(h_t^i,\bar{y}_r^i,\bar{y}_r^p\right))/2)),\end{aligned} \tag{43}$$

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

Estimated processes		Source
Working period		
$z_t^{i,j}(\cdot)$	Endogenous age-efficiency profiles	PSID, See Appendix C.1 and C.2
ϵ_t^i	Wage shocks	PSID, See Appendix C.1
$\zeta_t(\cdot)$	Divorce probability	PSID, See Appendix C.5
$\nu_t(\cdot)$	Probability of marrying	PSID, See Appendix C.5
$\xi_t(\cdot)$	Matching probability	PSID, See Appendix C.5
$\iota_t(\cdot)$	Partner's wealth and earnings	PSID, See Appendix C.6
$f_t^{0,5}(i,j)$	Number of children ages 0-5	PSID, See Appendix C.7
$f_t^{6,11}(i,j)$	Number of children ages 6-11	PSID, See Appendix C.7
Retirement period		
$\pi_t^{i,j}(\cdot)$	Transition matrix for health status	HRS, See Appendix C.9
$m_t^{i,j}(\cdot)$	Medical expenses	HRS, See Appendix C.11
$s_t^{i,j}(h_t^i, \bar{y}^i)$	Survival probability	HRS, See Appendix C.10
$d_t^{i,j}(\cdot)$	End of life medical expenses	HRS, See Appendix C.11
b	Value of life	See Appendix C.12
Government policy		
λ_t^j, τ_t^j	Income tax	See Appendix C.13

Table C.1: First-step estimated inputs

C First step estimation

C.1 Wages

Our initial conditions, assortative matching in marriage, and marriage and divorce probabilities depend on the realized values of wage shocks. We estimate these values as follows. First, we impute potential wages for non-workers using fixed effects regressions. (As shown in Borella, De Nardi, and Yang (2023b), our results are robust to using a Heckman selection correction instead.) Second, we use both observed and imputed wages to estimate wage functions by age, gender, and human capital. Third, we estimate the persistence and variance of the unobserved wage component. Finally, we recover the realized wage shocks using Kalman smoothing.

Missing wages imputation. The observed wage rate is defined as annual earnings divided by annual hours worked. Gross annual earnings refer to labor income during the previous year; annual hours reflect total hours spent working for pay over the same period. We set to missing any observations with hourly wage rates below half the minimum wage ($7.5) , or above $350, both expressed in 2018 dollars.

We impute missing wages using fixed effects regressions estimated separately by

gender:

$$\ln wage_{kt} = f_k + Z'_{kt}\beta_z + \varsigma_{kt}.$$

The dependent variable is the log hourly wage; f_k denotes individual fixed effects, and ς_{kt} is an idiosyncratic error term. The vector Z_{kt} includes a rich set of explanatory variables: a fifth-order polynomial in age; a third-order polynomial in labor market experience (measured as years of participation); marital status (indicator for being single); family size (categorical dummies); number of children; age of youngest child; and an indicator for whether the partner works (if married). To capture health, we use an indicator for whether bad health limits work capacity–the only health variable consistently available in the PSID. Since this variable is not collected for wives, we omit it from the regression for married women. Both regressions also include interaction terms among regressors. Time-invariant variables are absorbed by the individual fixed effect f_k. To avoid endpoint issues in the age polynomials, we include individuals ages 24 to 70 in the sample.

We define potential wages $\ln \overline{wage}_{kt}$ as observed wages when available and imputed wages otherwise.

Wage function estimation. The wage function in our structural model depends on age, gender, and human capital, measured as average realized earnings accrued up to the beginning of age t, $\bar{y}_t$ [10].

We estimate it in two stages. In the first stage, we run a fixed effects regression for the logarithm of potential wages:

$$\ln \overline{wage}_{kt} = d_k + f^i(t) + \sum_{g=1}^{G} \beta_g D_g \ln(\bar{y}_{kt} + \delta_y) + u_{kt}, \tag{44}$$

where d_k is an individual fixed effect, $f^i(t)$ is a fifth-order polynomial in age specific to each marital status–gender group, and D_g is a set of dummies for these groups. The earnings shifter δ_y = \$5,000 avoids taking logs of small values.[11] We also experimented with including marital status dummies to capture short-term effects of marital transitions, but they were not statistically significant once average earnings were controlled for.

Second, to fix the constant of the wage profile, we regress the sum of the residuals

[10] Appendix C.2 provides details on the computation of $\bar{y}_t$.

[11] While we compute average earnings using income subject to the Social Security cap, using uncapped earnings in the regression yields very similar results.

and fixed effects,

$$w_{kt+1} = d_k + u_{kt+1}, \tag{45}$$

on cohort dummies to compute the average effects for the cohorts born in 1941-1945.[12]

Table C.2 reports the coefficients from the first-stage regression, and Table C.3 reports the second-stage cohort adjustments. We use these coefficients to compute the deterministic average logarithm of wages at age t for married and single men and women, $z_t^{i,j}(\bar{y}_t^i)$.

	Coefficient	Standard Error
$\ln(\bar{y}_t + \delta_y)$	0.501***	(0.0101)
$\ln(\bar{y}_t + \delta_y)$*female	-0.188***	(0.0115)
$\ln(\bar{y}_t + \delta_y)$*single	-0.0532***	(0.0148)
$\ln(\bar{y}_t + \delta_y)$*single*female	0.0598***	(0.0169)
Age	0.170***	(0.0260)
$Age^2/(10^2)$	-0.533***	(0.0932)
$Age^3/(10^4)$	0.813***	(0.143)
$Age^4/(10^6)$	-0.492***	(0.0799)
*Age*female	-0.174***	(0.0349)
$Age^2/(10^2)$*female	0.556***	(0.126)
$Age^3/(10^4)$*female	-0.781***	(0.194)
$Age^4/(10^6)$*female	0.413***	(0.109)
Single	-1.877***	(0.470)
*Age*single	0.230***	(0.0480)
$Age^2/(10^2)$*single	-0.785***	(0.175)
$Age^3/(10^4)$*single	1.143***	(0.276)
$Age^4/(10^6)$*single	-0.601***	(0.158)
*Age*single*female	-0.0706***	(0.0189)
$Age^2/(10^2)$*single*female	0.283***	(0.0804)
$Age^3/(10^4)$*single*female	-0.464***	(0.149)
$Age^4/(10^6)$*single*female	0.265***	(0.0979)
Constant	-2.343***	(0.201)
N	113163	
R-sq	0.130	

Table C.2: Dependent variable: $\ln \overline{wage}_{kt}$. Coefficients from fixed effects regressions. PSID data. Robust standard errors in parentheses, clustered at the individual level.
* p<0.10, ** p<0.05, *** p<0.01.

Figure C.1 shows the average wage profiles implied by our estimated processes. Consistent with the evidence on the marriage premium, married men earn more than single men. In contrast, single women earn more than married women throughout the working life of this cohort.

The shock in log wages is modeled separately by gender as the sum of a persistent component and a white noise term, which we interpret as measurement error:

$$\tilde{w}_{kt+1} = \ln \epsilon_{kt+1} + \xi_{kt+1} \tag{46}$$

[12]Our sample includes individuals born in more cohorts than those directly analyzed in our study in order to ensure an adequate sample size, which tends to be quite small–especially for single men.

	Men	Women
Born in 1946-50	0.0269***	0.0681***
	(0.00720)	(0.00679)
Born in 1951-55	-0.00703	0.0554***
	(0.00726)	(0.00666)
Born in 1956-60	-0.00708	0.0359***
	(0.00751)	(0.00696)
Born in 1961-65	-0.0451***	0.0414***
	(0.00818)	(0.00733)
Constant	-2.000***	1.851***
	(0.00566)	(0.00536)
N	55016	58147
R-sq	0.002	0.002

Table C.3: Dependent variable: w_{kt+1}. Coefficients from fixed effects regressions. PSID data. Robust standard errors in parentheses, clustered at the individual level. * p<0.10, ** p<0.05, *** p<0.01.

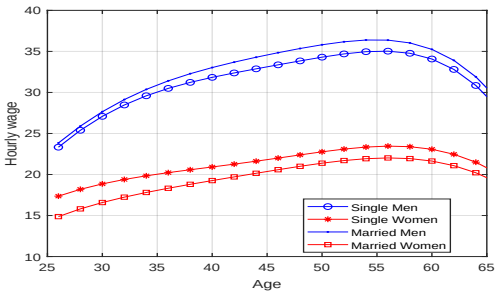


Figure C.1: Average wage over the life cycle for the 1945 birth cohort, PSID data

$$\ln \epsilon_{kt+1} = \rho_\epsilon \ln \epsilon_{kt} + v_{kt+1}, \tag{47}$$

where $\tilde{w}_{kt+1}$ denotes the residual from the second-stage wage regression, and ξ_{kt+1} and v_{kt+1} are independent white-noise processes with zero mean and variances σ_ξ^2 and σ_v^2, respectively. We estimate this system using equations (46) and (47) to recover the persistence and variances of the underlying shocks, as well as the realized wage shocks for each individual via Kalman smoothing.

Specifically, we first estimate the persistence and variances by maximum likelihood, assuming that the initial state and the shocks are Gaussian, and using standard Kalman filter recursions.[13] Given these estimates, we then obtain smoothed values of the latent state—that is, $\ln \epsilon_{kt}$ for $t = 1, \ldots, T$.

[13] We limit the age range to 26–66 and drop the largest 0.5% of residuals for both men and women. This trimming avoids undue influence from outliers on the estimated variances and has negligible effects on our results.

C.2 Average realized earnings

Our model tracks each individual's average past accumulated earnings ($\bar{y}_{kt}$) over time, subject to a time-varying Social Security cap on yearly earnings. To initialize this variable, we require data on average accumulated earnings at age 26, when individuals enter the model. Using the PSID data, we assume individuals begin working at age 22. Those who enter our data before age 26 are assigned zero average accumulated earnings until they are first observed. For those entering our data after age 26, we impute their initial average earnings using a regression-based approach. Specifically, we estimate a regression of capped earnings on a fourth-order polynomial in age, fully interacted with gender. The regression also includes education dummies, their interactions with gender, marital status, and race dummies interacted with gender, as well as cohort dummies. The predicted values from this regression determine the initial average earnings for entrants above age 26. After entry, the model updates each individual's average accumulated earnings based on their observed earnings history.

In our model, average past accumulated earnings at the time of retirement ($\bar{y}_{kr}$) influence post-retirement health, medical expenses, and survival probabilities. These relationships are estimated using HRS data, in which average past accumulated earnings are not directly observable. We therefore impute this variable using the following procedure. We start by estimating the relationship between average past accumulated earnings at retirement ($\bar{y}_{kr}$) and retirement income using PSID data. Specifically, we regress $\bar{y}_{kr}$ on a flexible specification that includes retirement income, education, gender, and marital status. Then, we use the estimated coefficients from this regression to impute $\bar{y}_{kr}$ for each individual in the HRS.

C.3 Wealth

We define wealth, or net worth, as all assets less all liabilities. Wealth in the PSID is recorded only in 1984, 1989, and 1994, and then in each (biennial) wave from 1999 onward. We rely on an imputation procedure to compute wealth in the missing years, starting in 1968. This imputation is based on the following fixed-effect regression:

$$\ln(a_{kt} + \delta_a) = Z'_{kt}\beta_z + da_k + wa_{kt}, \tag{48}$$

where k denotes the individual and t is age. The parameter δ_a is a shifter for wealth to have only positive values and to be able to take logs. The variables Z include polynomials in age, also interacted with health status and with average earnings (uncapped),

family size, and a dummy for health status. The term da_k is the individual fixed effect, and wa_{kt} is a white-noise error term. Equation (48) is estimated separately for single men, single women, and couples, as wealth is measured at the household level.

To estimate the wealth profile up to older ages, we merge individuals from the PSID—where we observe the active part of the life cycle well—and from the HRS sample, which provides information on the retirement years. When combining the samples, we account for two key issues. First, the sample designs in the PSID and HRS differ: in particular, the HRS oversamples non-white individuals. Second, the definitions of wealth in the two surveys are not exactly the same.

We address the first issue by using sample weights when estimating the profiles. Specifically, we assign a weight of 1 to individuals in the PSID and white individuals in the HRS sample, and a weight of 0.5 to non-white individuals in the HRS sample.[14]

To address the second issue—the differences in wealth definitions between the two surveys—we include a dummy variable indicating whether the individual comes from the HRS or PSID sample, using the PSID level as the benchmark in our profiles.

We then use the estimated wealth profiles as target moments and to parameterize the joint distribution of initial wealth, average realized earnings, and wage shocks for single men, single women, and couples.

C.4 Initial distributions

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

$$\begin{pmatrix} \ln(a_t^i + \delta_a^i) \\ \ln(\bar{y}_t^i) \\ \ln \epsilon_t^i \end{pmatrix} \sim N \begin{pmatrix} \mu_{at}^i + \delta_a^i \\ \mu_{\bar{y}t}^i \\ \mu_{\epsilon t}^i \end{pmatrix}, \Sigma_{st}^i \Bigg), \tag{49}$$

where $\Sigma^i st$ is a 3×3 covariance matrix. We estimate its mean and variance as functions of age t. For the mean, we regress the logarithm of wealth (plus a shift parameter), average earnings, and the wage shock $\ln \hat{\epsilon}_t^i$ on a third-order polynomial in age and cohort dummies. The predicted age profiles for the 1945 cohort provide age-specific estimates of the mean of the lognormal distribution. Using the residuals from these regressions, we estimate the elements of the variance-covariance matrix

[14] We implement this simple strategy because the sample weights provided by the two surveys follow different methodologies and are not comparable.

by computing the relevant squares and cross-products. We then regress each squared residual or cross-product on a third-order polynomial in age to obtain smooth, age-specific estimates of each element of the variance-covariance matrix.

For couples, we compute the age-26 initial joint distribution of the following variables

$$\begin{pmatrix} \ln(a+\delta_a) \\ \ln(\bar{y}^1) \\ \ln(\bar{y}^2) \\ \ln(\epsilon^1) \\ \ln(\epsilon^2) \end{pmatrix} \sim N \begin{pmatrix} \mu_a + \delta_a \\ \mu_{\bar{y}1} \\ \mu_{\bar{y}2} \\ \mu_{\epsilon 1} \\ \mu_{\epsilon 2} \end{pmatrix}, \Sigma_c \Bigg), \tag{50}$$

where Σ_c is a 5×5 covariance matrix estimated using data for married or cohabiting couples.

C.5 Marriage and divorce probabilities

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

$$\nu^i_{t+1} = Prob(Married_{t+1} = 1 | Married_t = 0, Z_t) = F(Z_t' \beta_m),$$

where F denotes the standard logistic function and Z_t includes a polynomial in age, cohort dummies, and the logarithm of the wage shock.[15] Using the estimated coefficients on the cohort dummies, we then adjust the marriage probability for the 1945 cohort.

We estimate the probability of divorcing as

$$\zeta_t = Prob(Divorced_{t+1} = 1 | Married_t = 1, Z_t) = F(Z_t' \beta_d),$$

where Z_t includes a polynomial in age, the husband's wage shock, the wife's wage shock, and cohort dummies.

Table C.4 reports our estimated coefficients from the marriage and divorce regressions.

Conditional on meeting a partner, the probability of meeting a partner p with wage shock ϵ^p_{t+1} is $\xi_{t+1}(\cdot) = \xi_{t+1}(\epsilon^p_{t+1} | \epsilon^i_{t+1}, i)$. Using our estimated wage shocks and

[15] The PSID shifted from a yearly to a biennial frequency in 1997. To estimate the biennial probabilities required by our model, we use only the waves collected in odd years.

	Single Men Marriage	Single Women Marriage	Couples Divorce
Age	-0.0181	-0.0585	0.0722**
	(0.0415)	(0.0409)	(0.0338)
$Age^2/10^2$	-0.0479	-0.0134	-0.146***
	(0.0493)	(0.0485)	(0.0388)
$\ln \epsilon_{kt}$	0.602***	0.247*	-0.449***
	(0.147)	(0.132)	(0.0982)
Spouse's $\ln \epsilon_{kt}$			-0.185*
			(0.105)
Born in 1946-50	-0.115	-0.297	-0.0201
	(0.211)	(0.189)	(0.129)
Born in 1951-55	-0.303	-0.262	-0.00545
	(0.203)	(0.179)	(0.127)
Born in 1956-60	-0.307	-0.134	-0.0928
	(0.203)	(0.177)	(0.129)
Born in 1961-65	-0.460**	-0.205	-0.0547
	(0.204)	(0.177)	(0.140)
Constant	0.196	0.695	-3.477***
	(0.857)	(0.830)	(0.717)
N	4001	6196	22234
Pseudo R-sq	0.059	0.066	0.034

Table C.4: Column 1: Marriage of single men; column 2: marriage of single women; column 3: divorce of couples. Estimated coefficients from logistic regressions. PSID data. Robust standard errors in parentheses, clustered at the individual level. * p<0.10, ** p<0.05, *** p<0.01.

partitioning households into three age groups (26–35, 36–45, and 46–65), we compute the variance-covariance matrix of newly matched partners' wage shocks within each age group. We then derive the conditional distribution of partner characteristics under the assumption of lognormality.

C.6 Spousal wealth and Social Security benefits

We assume random matching over a partner's wealth and lifetime income, conditional on the partner's wage shock. Accordingly, we compute $\iota_{t+1}(\cdot) = \iota_{t+1}(a^p_{t+1}, \bar{y}^p_{t+1}|\epsilon^p_{t+1})$ using sample values of wealth, average capped earnings, and wage shocks. Specifically, we assume that ι_{t+1} follows a lognormal distribution at each age, with its mean and variance estimated from the data. Wealth includes a shift term, as described in the construction of the joint distribution at age 26 (see the "Initial Distributions" subsection of this Appendix).

Figure C.2 presents average spousal wealth (left panel) and Social Security–covered earnings (right panel) by the spouse's wage shock. The dispersion of these resources increases rapidly with age, indicating that individuals who marry later face a more dispersed distribution of spousal resources.

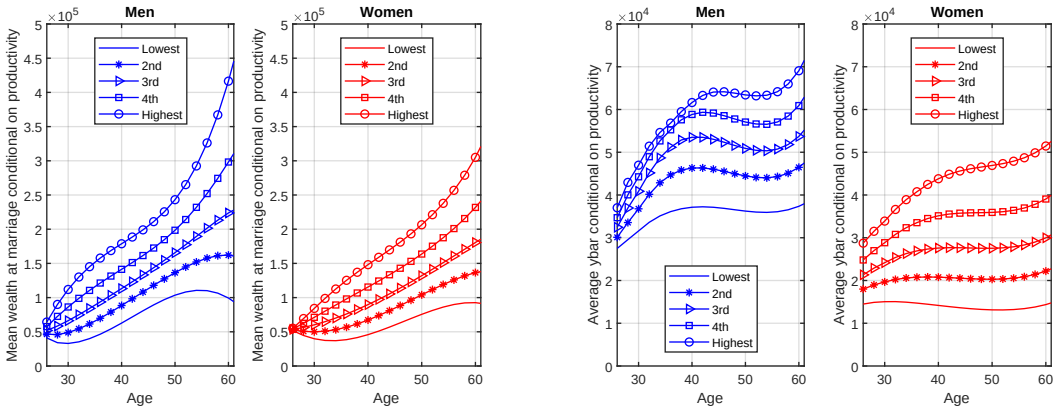


Figure C.2: Average spousal wealth (on the left) and Social Security Earnings (on the right) by spousal wage shocks in case of marriage. PSID data.

C.7 Number of children

To compute the average number of children by age group, we use individual-level information from the PSID and classify as children of the family those in the following categories: sons or daughters of the head, stepsons or stepdaughters of the head, sons or daughters of the cohabiting partner but not of the head, foster sons or daughters (not legally adopted), and children of the cohabitant but not of the head. We then sum the number of children in each age category (0–5, 6–11, and 0–17 for the total number) and regress the resulting count on a fifth-order polynomial in the mother's age, interacted with marital status and cohort dummies. This yields the average age profile of children in each group for single and married women. Figure C.3 displays these profiles.

C.8 Health Status at Retirement

We define health status as a categorical variable with three possible states: "good," "bad," or "nursing home." We assign individuals to the nursing "home state" if they meet one of the following criteria: (1) they spend at least 120 days in a nursing home between interviews (on average, 60 days per year), or (2) they spend at least 60 days in a nursing home before the next scheduled interview and die before that interview. For those not in a nursing home, we classify their health as "good" if they report excellent, very good, or good health, and as "bad" if they report fair or poor health.

We estimate the probability of each health state at age 66 as a function of average past earnings, using a multinomial logit model. We run separate estimations for

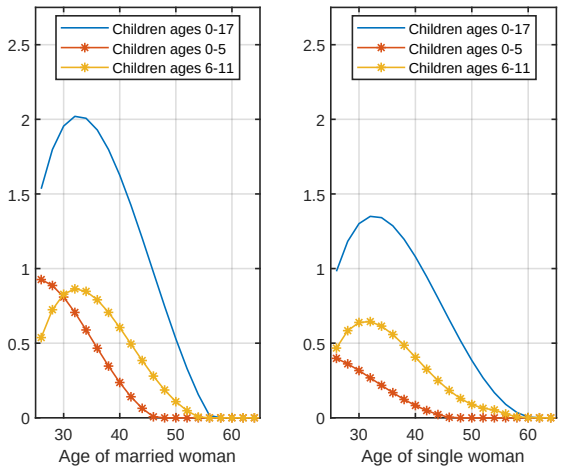


Figure C.3: Number of children for married (on the left) and single women (on the right) for the 1945 cohort, PSID data.

single men and single women, using individuals aged 65–67 and controlling for cohort dummies and average earnings. For couples, we estimate a multinomial logit model for the six possible joint health states (e.g., both partners in good health; one partner in good health and the other in bad health, etc.), using the same explanatory variables and restricting the husband's age to 65–67.

C.9 Health dynamics after retirement

We model transitions across the three health states–good, bad, and nursing home–using a multinomial logit model. The probability of being in health state i at time t equals

$$\pi_{ht}^{i} = \text{Prob}\left(h_t = i \mid X_t^h\right) = \frac{\exp\left(X_t^{h\prime}\beta_i^h\right)}{\sum_{j=0}^{2}\exp\left(X_t^{h\prime}\beta_j^h\right)},$$

where i indexes the health state and X_t^h includes age, gender, average earnings $\bar{y}_t$, marital status, previous-period health status, and interactions between these variables.

Table C.5 presents the estimated coefficients, and Figure C.4 shows the transition probabilities by gender, age, marital status, and prior health status.

Figure C.4 shows health transition probabilities, conditional on survival, for singles and couples by age, gender, and initial health status and $\bar{y}$ (which we set to $17,439$). These transitions vary substantially, particularly in the likelihood of being in good or bad health in the next period, conditional on current health.

	Coefficient	SE
Health state: Bad Health		
Age	-0.0246**	(0.0125)
Age^2	0.0301***	(0.0087)
$health_{t-1}$ = Bad	-1.1151*	(0.6128)
$health_{t-1}$ = Nursing Home	-2.8882	(5.0913)
$health_{t-1}$ = Bad × Age	0.1257***	(0.0175)
$health_{t-1}$ = Nursing Home × Age	0.1849	(0.1331)
$health_{t-1}$ = Bad × Age^2	-0.1022***	(0.0122)
$health_{t-1}$ = Nursing Home × Age^2	-0.1548*	(0.0860)
$\bar{y}_r$	-0.1725***	(0.0038)
Male	0.2928***	(0.0244)
Male × Married	0.0008	(0.0306)
Spouse $\bar{y}_r$ × Married	-0.0816***	(0.0041)
Married	0.1131***	(0.0249)
Born 1911-1915	0.2066***	(0.0529)
Born 1916-1920	0.2287***	(0.0411)
Born 1921-1925	0.1845***	(0.0352)
Born 1926-1930	0.1019***	(0.0317)
Born 1931-1935	0.0814***	(0.0269)
Born 1936-1940	0.0614**	(0.0245)
Born 1946-1950	0.0322	(0.0275)
Born 1951-1955	0.0480	(0.0297)
Born 1956-1960	0.0110	(0.0350)
Born 1961-1965	0.0467	(0.0544)
Constant	-1.0906**	(0.4434)
Health state: Nursing Home		
Age	0.1070*	(0.0560)
Age^2	0.0164	(0.0347)
$health_{t-1}$ = Bad	1.5597	(2.6158)
$health_{t-1}$ = Nursing Home	10.0765**	(4.8246)
$health_{t-1}$ = Bad × Age	0.0704	(0.0670)
$health_{t-1}$ = Nursing Home × Age	0.0412	(0.1240)
$health_{t-1}$ = Bad × Age^2	-0.0790*	(0.0426)
$health_{t-1}$ = Nursing Home × Age^2	-0.1156	(0.0789)
$\bar{y}_r$	-0.1294***	(0.0136)
Male	0.3516***	(0.0632)
Male × Married	-0.3454***	(0.1058)
Spouse $\bar{y}_r$ × Married	-0.0087	(0.0199)
Married	-0.6671***	(0.1078)
Born 1911-1915	0.4521***	(0.1383)
Born 1916-1920	0.4343***	(0.1296)
Born 1921-1925	0.3116**	(0.1253)
Born 1926-1930	0.1816	(0.1238)
Born 1931-1935	0.2052*	(0.1163)
Born 1936-1940	-0.0587	(0.1175)
Born 1946-1950	-0.1752	(0.1582)
Born 1951-1955	-0.2920	(0.2006)
Born 1956-1960	-0.0385	(0.2445)
Born 1961-1965	0.3877	(0.3811)
Constant	-13.1377***	(2.2299)
N	175103	
Pseudo-R^2	0.291	

Table C.5: Health dynamics over two-year periods. Multinomial logit coefficients for health state transitions. HRS data. Robust standard errors in parentheses clustered at the individual level. * p<0.1, ** p<0.05, *** p<0.01.

The top panel illustrates that, at age 70, individuals in good health face an approximately 80% chance of remaining in good health. Those in bad health recover with a probability of around 20–25%, while individuals in nursing homes have less than a 10% chance of returning to good health. Across all ages and initial health states, women remain in good health more often than men.

The middle panel shows the probability of being in bad health in the next period. At age 70, individuals already in bad health face a roughly 75% chance of remaining in that state. The probability of transitioning into bad health is about 15–20% for those currently in good health or in a nursing home. As individuals age, the likelihood of entering a nursing home in the next period rises for those in good health but declines for those already in a nursing home.

The bottom panel shows the probability of being in a nursing home in the next period. This probability remains highest—around 70–80%–for individuals already in a nursing home and stays relatively flat across ages. In contrast, those in good or bad health at age 70 face less than a 5% chance of entering a nursing home.

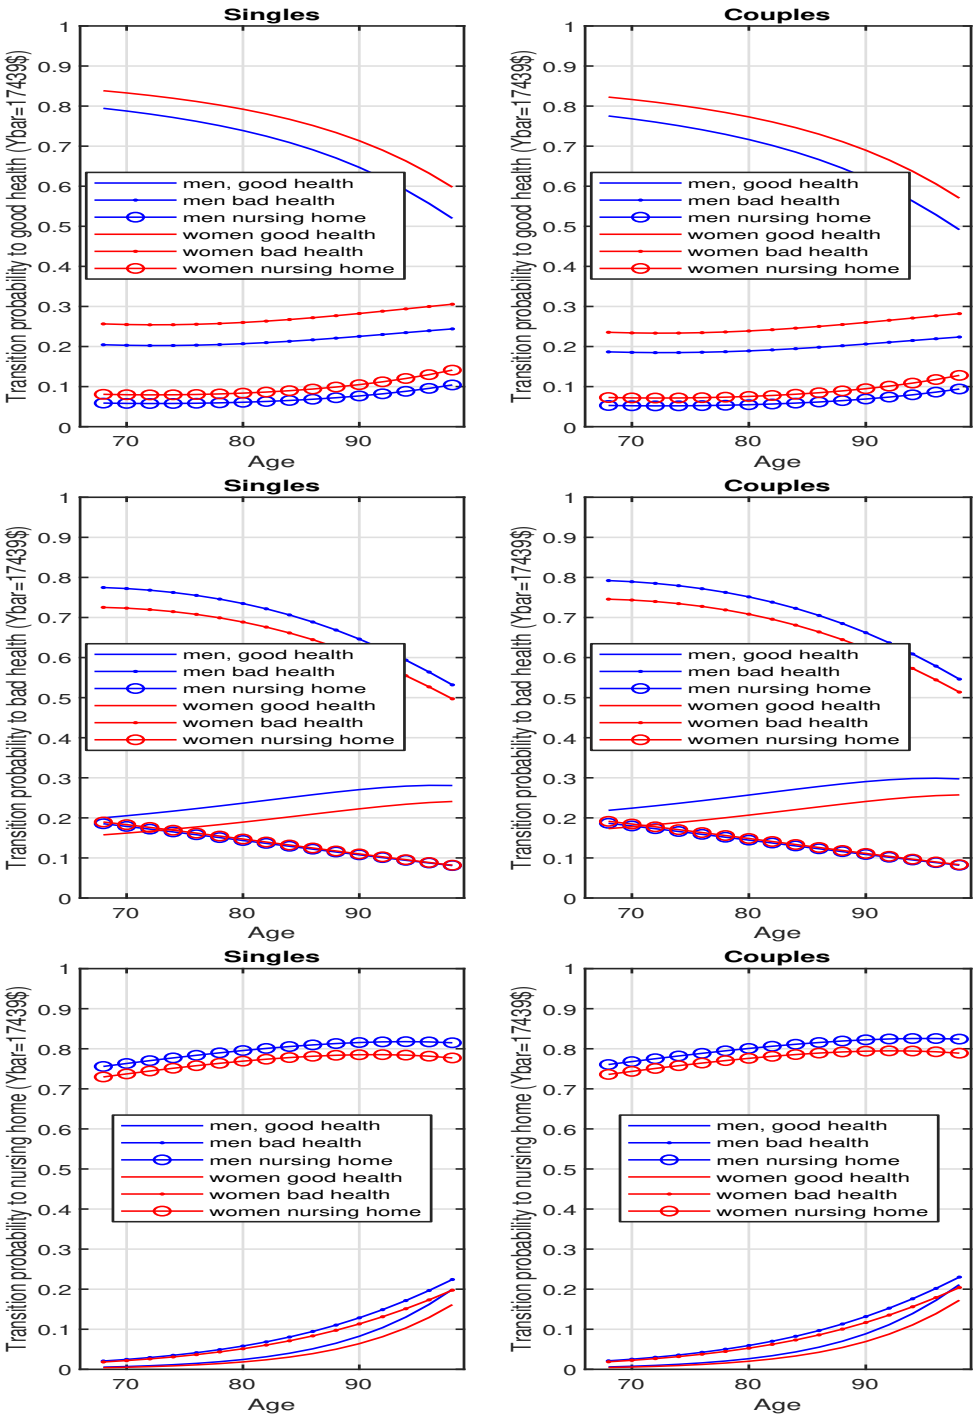


Figure C.4: Health transition probabilities, conditional on surviving, for singles and couples by age, HRS data.

C.10 Survival probabilities

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

$$s_t = \text{Prob}\left(\text{Alive}_t = 1 \mid X_t^s\right) = F\left(X_t^{s\prime}\beta^s\right),$$

where X_t^s is the vector of explanatory variables at time t, and β^s denotes the estimated coefficients. The covariates include a third-order polynomial in age, gender, marital status, human capital $\bar{y}_r$, and previous-period health status, along with interaction terms included when statistically significant.

Table C.6 reports the estimated coefficients from the logistic regression.

	Coefficient	SE
Age	-0.3078**	(0.1345)
Age^2	0.3667**	(0.1773)
Age^3	-0.1971**	(0.0771)
$\bar{y}_r$	0.0723***	(0.0064)
$health_{t-1}$ = Bad	-3.4617***	(0.1667)
$health_{t-1}$ = Nursing Home	-7.6339***	(0.3571)
$health_{t-1}$ = Bad × Age	0.0295***	(0.0021)
$health_{t-1}$ = Nursing Home × Age	0.0637***	(0.0043)
Male	-0.6434***	(0.0255)
Married	0.0892***	(0.0225)
Born 1911-1915	-0.1151*	(0.0646)
Born 1916-1920	-0.1500**	(0.0587)
Born 1921-1925	-0.1637***	(0.0553)
Born 1926-1930	-0.1559***	(0.0531)
Born 1931-1935	-0.0876*	(0.0490)
Born 1936-1940	-0.0308	(0.0474)
Born 1946-1950	0.1886***	(0.0617)
Born 1951-1955	0.2147***	(0.0705)
Born 1956-1960	0.3441***	(0.0946)
Born 1961-1965	0.3448**	(0.1614)
Constant	14.1139***	(3.3715)
N	187853	
Pseudo-R^2	0.198	

Table C.6: Survival over a two-year period. Logistic regression coefficients. HRS data. Robust standard errors in parentheses clustered at the individual level. * p<0.1, ** p<0.05, *** p<0.01.

Figure C.5 reports survival probabilities. Survival is more likely for women, individuals in better health, younger individuals, and those who are married.

C.11 Medical expenses

We model and estimate medical expenses using a comprehensive approach that captures spending patterns throughout life, including end-of-life expenditures.

Because our model explicitly accounts for out-of-pocket medical expenses and treats Medicaid payments as part of the consumption floor, the relevant measure of

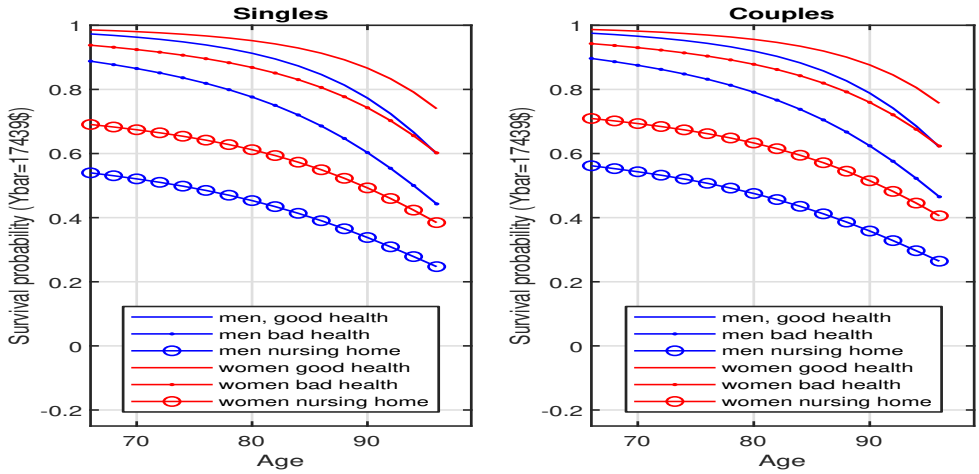


Figure C.5: Survival probability by age, gender, and marital and health status, HRS data.

medical expenses must include both components. While the HRS provides data on out-of-pocket medical expenditures, it lacks direct information on Medicaid spending. To fill this gap, we impute Medicaid payments by combining HRS data with administrative records from the Medicare Current Beneficiary Survey (MCBS), following the method described in De Nardi, French, Jones, and McGee (2025).

We define out-of-pocket medical expenses as total individual spending on hospital and nursing home stays, doctor visits, dental care, outpatient surgery, monthly prescription drugs, home health care, special facilities charges, and insurance premia for private plans and Medicare. These payments are made directly by individuals and exclude amounts covered by insurance. For deceased individuals, we also include death-related expenses–such as funeral costs, unpaid insurance premia, and legal fees–using information from HRS exit interviews.

Because our model assumes that medical expenses are deterministic conditional on observable states–age, gender, marital status, and health–we do not include stochastic shocks to medical spending. Instead, we estimate a flexible specification of log medical expenses as a function of these state variables, along with interactions that capture end-of-life spending. To recover the average level of medical spending–not just its logarithm–we assume that residuals are normally distributed (i.e., medical expenses are log-normally distributed). We therefore adjust the conditional means of $\ln m_{kt}$ by adding half the estimated residual variance before exponentiating. This yields smooth, state-dependent profiles of expected medical spending, consistent with the model's structure and inputs.

The estimation equation for out-of-pocket medical expenses plus Medicaid pay-

ments is:

$$\ln(m_{kt}) = \alpha_k^m + X_{kt}^{m\prime}\beta^m + \delta_{\text{death}} \cdot D_{\text{death},t} + (X_{kt}^m \times D_{\text{death},t})'\gamma^m + u_{kt}^m,$$

where m_{kt} denotes out-of-pocket medical expenses plus Medicaid payments for individual k at time t, including both during life and in the period of death. The vector X_{kt}^m includes age (up to a third-order polynomial), gender, health status, marital status, and relevant interactions (e.g., age–gender). The indicator $D_{\text{death},t}$ equals one if the individual dies during the period, and its interaction with covariates, $(X_{kt}^m \times D_{\text{death},t})'\gamma^m$, allows health and demographic characteristics to affect medical expenses differently in the period of death. This captures heterogeneity in death-period costs by, for example, chronic conditions, nursing home status, and gender. The term α_k^m denotes individual fixed effects. We estimate the specification using a fixed effects model to control for persistent unobserved heterogeneity—such as differential mortality—that could otherwise bias the estimated age profiles, as discussed in De Nardi, French, and Jones (2010).

Then, we regress the sum of residuals and fixed effects from this equation on cohort, gender, average past earnings at retirement, and marital status dummies to estimate average medical expenses for each group of interest. Table C.7 reports the estimated coefficients, and Figure 3 displays medical expenses by age, gender, and marital and health status.

Finally, we model the variance of medical expense shocks by regressing the squared residuals (in logs) on a second-order polynomial in age, fully interacted with current health status, as well as on a death-period indicator and cohort, gender, and marital status dummies. We use these estimates to construct average medical expenses as a function of age by adding half the estimated variance to the predicted mean in logs prior to exponentiation.

C.12 The parameter b and the value of life

The parameter b implies a value of statistical life (VSL) for each person. The VSL is defined as the compensation that people require to bear an increase in their probability of death, expressed as "dollars per death." For example, suppose that people are willing to tolerate an additional fatality risk of 1/10,000 during a given period for a compensation of \$500 per person. Among 10,000 people, there will be one death, and it will cost the society 10,000 times \$500 = \$5 million, which is the implied VSL. We set b so that the implied VSL is to \$2.4 million, which is in the

First Stage: Fixed Effects Regression	Coefficient	SE
Age	0.541***	(0.0451)
Age^2	-0.743***	(0.0644)
Age^3	0.337***	(0.0302)
$Health = Bad$	2.811*	(1.529)
$Health = NursingHome$	-4.432	(9.973)
$Health = Bad \times Age$	-0.106	(0.0659)
$Health = NursingHome \times Age$	0.162	(0.390)
$Health = Bad \times Age^2$	0.137	(0.0934)
$Health = NursingHome \times Age^2$	-0.114	(0.502)
$Health = Bad \times Age^3$	-0.0569	(0.0436)
$Health = NursingHome \times Age^3$	0.0150	(0.213)
Male $\times Age$	-0.106	(0.0671)
Male $\times Age^2$	0.157	(0.0957)
Male $\times Age^3$	-0.0767*	(0.0449)
Married	0.0693***	(0.0114)
Dead	1.651***	(0.0260)
$Health = Bad\times$ Dead	-0.196***	(0.0285)
$Health = NursingHome\times$ Dead	-0.998***	(0.0511)
Dead $\times$ Married	-0.145***	(0.0268)
Constant	-3.978***	(0.864)
Second Stage: Residual Regression	Coefficient	SE
$\bar{y}_r$	0.0517***	(0.000986)
Spouse $\bar{y}_r$ $\times$ Married	0.0176***	(0.00116)
Male	2.083***	(0.00898)
Married	0.0478***	(0.00830)
Male $\times$ Married	0.0774***	(0.0111)
Dead	0.0894***	(0.0142)
Male $\times$ Married $\times$ Dead	-0.0795***	(0.0248)
Born 1911-1915	0.0821***	(0.0152)
Born 1916-1920	0.0560***	(0.0122)
Born 1921-1925	0.0345***	(0.0108)
Born 1926-1930	-0.00959	(0.0102)
Born 1931-1935	-0.0317***	(0.00915)
Born 1936-1940	-0.0467***	(0.00874)
Born 1946-1950	0.0705***	(0.00977)
Born 1951-1955	0.0773***	(0.0101)
Born 1956-1960	0.0886***	(0.0110)
Born 1961-1965	0.0757***	(0.0148)
Constant	-1.178***	(0.00831)
Observations	219670	
R^2 First Stage	0.0249	
R^2 Second Stage	0.5132	

Table C.7: Estimates for the logarithm of medical expenses. First stage (fixed effects) and second stage (residual regression). Robust standard errors in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

range of those estimated by the literature. Our results are not very sensitive to the value of this parameter.

C.13 Taxes

We model taxes T on total income Y as

$$T(Y) = Y - \lambda Y^{1-\tau},$$

where τ governs the degree of progressivity and λ determines the average level of taxation. This implies

$$Y - T(Y) = \lambda Y^{1-\tau} \quad \text{and} \quad \ln(Y - T(Y)) = \ln(\lambda) + (1-\tau)\ln(Y).$$

We estimate τ and λ by regressing the logarithm of after-tax household income on a constant and the logarithm of pre-tax household income, separately by cohort, year, and household type (single men, single women, and couples).

We use PSID data from 1968 to 2017 to estimate cohort– and time-specific tax functions. Up to 1991, the PSID provides reported information on federal taxes paid. For years after 1991, we compute taxes using TAXSIM, the NBER simulation program. Specifically, we use the version developed by Borella, De Nardi, Pak, Russo, and Yang (2023a), which extends the earlier work of Kimberlin, Kim, and Shaefer (2015) to prepare PSID input files for TAXSIM.[16]

We define before-tax household income as the sum of all income received by the spouses (or by the individual if single) in a given tax year. It includes the labor income of the head and wife (if present), asset income from farms, businesses, roomers, rent, interest, and dividends, and the wife's income from wealth. It also includes transfer income such as Social Security, pensions, annuities, other retirement income, welfare, aid to dependent children, unemployment or workers' compensation, help from relatives, alimony, and child support. We define after-tax household income as before-tax income minus the federal income tax liability. The latter includes capital gains taxes, surtaxes, the alternative minimum tax, and refundable and non-refundable credits, as computed using TAXSIM.

[16] Kimberlin, Kim, and Shaefer (2015) developed code to prepare TAXSIM input files for PSID data from 1999 to 2011, based on the specifications in Butrica and Burkhauser (1997). Their approach differs from simplified PSID-TAXSIM interfaces in that it identifies multiple tax units within a PSID family. Cohabiting couples are treated as two separate tax units, and children are assigned to the appropriate unit (head or cohabitant) based on relationship codes. We extended their program to cover the full range from 1992 to 2017.

All inputs required by TAXSIM come from the PSID for the years 1992–2017. However, medical expenses and charitable contributions are not available prior to 1999 and must be imputed. We do so by pooling data from 1999–2015, regressing the sum of these two items on a set of observables, and predicting their values for earlier years using the estimated coefficients. The explanatory variables include demographic and income characteristics: family size, employment status of the head and spouse, state of residence, wages, pensions, other income, education, number of children, age, and marital status.

To preserve the variance of the imputed values, we follow the procedure in David, Little, Samuhel, and Triest (1986) and French and Jones (2011). Specifically, we:

1. Regress the sum of medical expenses and charitable contributions on observables for the sample of heads who itemize:

$$deduc_i = x_i\beta + \epsilon_i.$$

2. For each household i with non-missing $deduc$, compute the predicted value $\widehat{deduc}_i = x_i\hat{\beta}$ and residual $\hat{e}_i = deduc_i - \widehat{deduc}_i$.

3. Sort $\widehat{deduc}_i$ into deciles and retain all corresponding $\hat{e}_i$ values within each decile.

4. For households j with missing $deduc$, compute $\widehat{deduc}_j = x_j\hat{\beta}$, assign it to a decile, and draw a random residual $\hat{e}_i$ from the same decile in the non-missing sample.

5. Impute $deduc_j$ as $\widehat{deduc}_j + \hat{e}_i$.

This procedure ensures that the imputed distribution preserves both the conditional mean and variance structure of the observed data.

D Calibrated model parameters

Table D.1 summarizes our first-step calibrated model inputs. We set the interest rate r to 4%, the curvature of the utility function, γ, to 2.5, and the equivalence scales to $\eta_t^{i,j} = (j + 0.7 * f_t^{i,j})^{0.7}$ (as estimated by Citro and Michael (1995)). The term $f_t^{i,j}$ refers to the average total number of children for single and married men and women by age. Our equivalence scale implies that \$1 spent by a household in a couple with two children gives each household member a consumption of 0.42 cents.

While we do not have data on childcare costs by age of children as a fraction of maternal earnings, the PSID data report both childcare costs for all children aged 0-11 in a household and maternal earnings. We thus calibrate our per-child child care costs to match the observed data on total childcare costs as a fraction of earnings for working women with children age 0-11. The range of per-woman total childcare costs as a fraction of her earnings implied by our model ranges from about 5% to 10% of mothers' salaries before age 38, depending on age and mother's marital status and matches the range that we estimate from the PSID data.

De Nardi, French, Jones, and McGee (2025) estimate a model of Medicaid recipiency in old age for single and married people and obtain a utility floor that corresponds to consuming $4,101 (in 2017 dollars, which is $4,347 in 2018 dollars) a year for singles. Consistent with the statutory rules, we assume that the one for couples is 1.5 that for singles.

Calibrated parameters		Source
Preferences and returns		
r	Interest rate	4% De Nardi, French, and Jones (2016)
γ	Utility curvature parameter	2.5 see text
η_t	Equivalence scales	Citro and Michael (1995)
Proportional child care cost		
$\tau_c^{0,5}$	For each child age 0-5	9.00%
$\tau_c^{6,11}$	For each child age 6-11	7.50%
Government policy		
$SS(\bar{y}_r^i)$	Social Security benefit	See text
τ_t^{SS}	Social Security tax rate	See text
$\widetilde{y}_t$	Social Security cap	See text
$\underline{c}(1)$	Minimum consumption, singles	$4,347, See text
$\underline{c}(2)$	Minimum consumption, couples	$4,347*1.5, See text

Table D.1: First-step calibrated inputs summary.

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

$$SS(\bar{y}_r) = \left\{ \begin{array}{ll} 0.9\bar{y}_r, & \bar{y}_r < 11,112 \\ 10,001 + 0.32(\bar{y}_r - 11,112), & 11,112 \leq \bar{y}_r < 66,996 \\ 27,885 + 0.15(\bar{y}_r - 0.6725), & 66,996 \leq \bar{y}_r < y_t^{cap} \end{array} \right\},$$

where the marginal rates and bend points come from the Social Security Administration.[17] The Social Security tax and Social Security cap, shown in Figure D.1, have

[17] Social Security Administration "Benefit Formula Bend Points":

been changing over time.

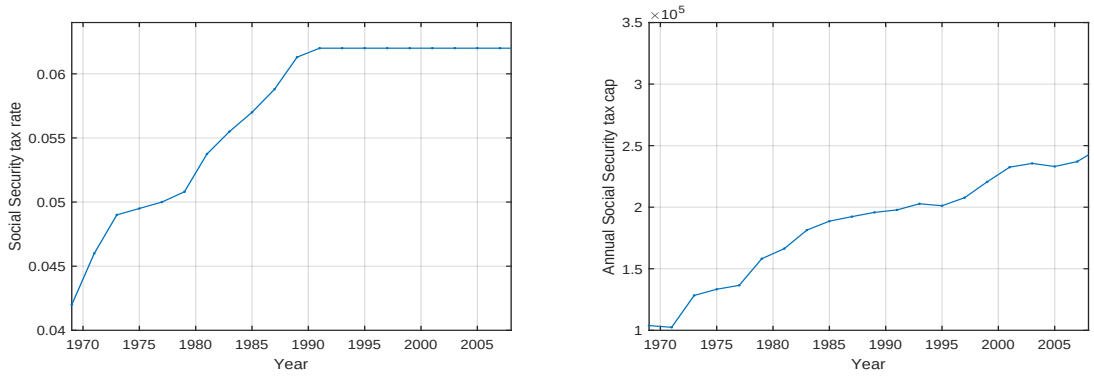


Figure D.1: Social Security tax (left) and cap (right) over time, for the 1945 cohort in 2018 dollars.

E Solution algorithm

We solve for the value functions and policy functions, and then we simulate our model economy. In the latter step, we use our estimated marriage and divorce probabilities both to compute the dynamic programming problem and to generate the distribution of people by state variables that is generated by our policy functions. Finally, we estimate our model parameters following the procedures that we describe in Appendix F.

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

During the retirement stage, single people do not get married anymore; hence, their value function can be computed independently of the other value functions. The value function of couples depends on their own future continuation value and

https://www.ssa.gov/oact/cola/bendpoints.html. We use values for 2019.

that of singles, in case of death of a spouse. Then there is the value function of each person in a couple, which depends on the optimal policy function of the couple, taking the appropriate expected values. We compute the value functions as follows:

1. Compute the value function of a retired single person for all time periods after retirement by backward iteration starting from the last period.

2. Compute the value function of a retired couple for all time periods after retirement by backward induction starting from the last period. This also uses the value function for the retired single person in case of death of one of the spouses.

3. Compute the value function of the single person in a marriage for all time periods after retirement.

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

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

2. Compute the value function of the couple for all time periods by backward iteration starting from the last period in the early retirement stage.

3. Compute the value function of each individual in a marriage for all time periods in the early retirement stage.

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

1. Take as given the value of being a single person in a married couple for next period and the value function of being single next period, which have been previously computed, and compute the value function of being single this period.

2. Given the value function of being single, compute the value function of the couple for the same age.

3. Given the optimal policy function of the couple, use the implied policy functions to compute the value function for a person in a couple.

4. Keep going back in time until the first period.

Computational considerations. Even with parallelization in C on high-performance workstations, solving the model for a given parameter vector takes approximately two minutes. Given the need to solve the model repeatedly over a large state space, the computational burden of the solution algorithm is substantial.

F Estimation strategy

In this appendix, we review the two-step estimation strategy, the moment conditions, and the asymptotic distribution of our estimator.

In the first step, we estimate the vector χ, which consists of the set of parameters that 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 21 model parameters (β, ω, $(\phi_0^{i,j}, \phi_1^{i,j}, \phi_2^{i,j})$, $L^{i,j}$).[18] Our estimate, $\hat{\Delta}$, of the "true" parameter vector Δ_0 is the value of Δ that minimizes the (weighted) distance between the lifecycle profiles found in the data and the simulated profiles generated by the model.

More specifically, we match median and average wealth for single men, single women, and couples from ages 28 to 84, and working hours and participation for single men, single women, married men, and married women from ages 26 to 64. For the generic variable z equal to hours (H), participation (In), and wealth (a), $x_{k,t}^{i,j}$ denotes the sample observation relative to person k, of gender i, marital status j, and age t. Using $x_t^{i,j}(\Delta, \chi)$ to denote the model-predicted expected value of x for age i, gender i, and marital status j, where χ is the vector of parameters estimated in the first step, we write the moment conditions as

$$E[a_{k,t}^{i,j} - a_t^{i,j}(\Delta_0, \chi_0)] = 0, \quad \forall t = 2, ..., 41 \tag{51}$$

$$E[H_{k,t}^{i,j} - H_t^{i,j}(\Delta_0, \chi_0)] = 0, \quad \forall t = 1, ..., 41 \tag{52}$$

$$E[In_{k,t}^{i,j} - In_t^{i,j}(\Delta_0, \chi_0)] = 0, \quad \forall t = 1, ..., 41. \tag{53}$$

Note that wealth for couples, $a_{k,t}^{i,j}$, does not depend on gender when marital status is $j = 2$. Also, as wealth at age 26 ($t = 1$) is an initial condition, it is matched by construction. In practice, we compute the sample expectations in equations (51), (52), and (53) conditional on a flexible polynomial in age. More specifically, we regress each

[18] We normalize the time endowment of single men.

variable x 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 using the estimated marital- and gender-specific polynomial in age and coefficients relative to that cohort. These average age profiles, conditional on gender, marital status, and cohort, are those shown in the figures in the main text.

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

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

$$\underset{\Delta}{\operatorname{argmin}}\ \hat{\varphi}_K(\Delta;\chi_0)'\widehat{\mathbf{W}}_K\hat{\varphi}_K(\Delta;\chi_0). \tag{54}$$

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

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

$$\sqrt{K}\left(\hat{\Delta}-\Delta_0\right) \rightsquigarrow N(0,\mathbf{V}), \tag{55}$$

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

$$\mathbf{V}=(\mathbf{D}'\mathbf{W}\mathbf{D})^{-1}\mathbf{D}'\mathbf{W}\mathbf{S}\mathbf{W}\mathbf{D}(\mathbf{D}'\mathbf{W}\mathbf{D})^{-1}, \tag{56}$$

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

$$\mathbf{D}=\frac{\partial\varphi(\Delta;\chi_0)}{\partial\Delta'}\Big|_{\Delta=\Delta_0} \tag{57}$$

is the $J\times M$ gradient matrix of the population moment vector; and $\mathbf{W}=\operatorname{plim}_{K\to\infty}\{\widehat{\mathbf{W}}_K\}$. 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 $\mathbf{S}^{-1}$. However, as Altonji and Segal (1996) point out, the optimal weighting matrix is likely to suffer from small sample bias. We thus use a diagonal weighting matrix that is the same as $\mathbf{S}$ along the diagonal and has zeros off the diagonal of the matrix. We estimate $\mathbf{D}$ and $\mathbf{W}$ with their sample analogs.

G Parameter estimates

	Coeff.	Std. Err.
β: Discount factor	0.9958	(0.00036)
ω: Consumption weight	0.4478	(0.00184)
ϕ_0: Terminal bequest, strength	1379.3442	(96.64710)
k_0: Terminal bequest, shifter	987.7149	(42.31610)
ϕ_1: Side bequest, strength	6221.2437	(230.86900)
k_1: Side bequest, shifter	1681.1474	(29.12160)
Time endowment parameters:		
$FL^{2,1}$: single women	-1.6060	(0.05834)
$FL^{1,2}$: married men	-2.5873	(0.10860)
$FL^{2,2}$: married women	-3.3447	(0.26815)
Participation cost parameters:		
$p_2^{1,1}$: Single men: quadratic term	0.0005	(0.00003)
$p_1^{1,1}$: Single men: linear term	-0.0093	(0.00141)
$p_0^{1,1}$: Single men: constant term	-0.6683	(0.02469)
$p_2^{2,1}$: Single women: quadratic term	0.0004	(0.00001)
$p_1^{2,1}$: Single women: linear term	-0.0044	(0.00084)
$p_0^{2,1}$: Single women: constant term	-0.8599	(0.03491)
$p_2^{1,2}$: Married men: quadratic term	0.0004	(0.00002)
$p_1^{1,2}$: Married men: linear term	-0.0048	(0.00087)
$p_0^{1,2}$: Married men: constant term	-0.7686	(0.01991)
$p_2^{2,2}$: Married women: quadratic term	0.0015	(0.00011)
$p_1^{2,2}$: Married women: linear term	-0.0725	(0.00441)
$p_0^{2,2}$: Married women: constant term	-1.6060	(0.05834)

Table G.1: Estimated parameters from the structural model. Standard errors in parentheses.

Table G.1 reports all of our structural model estimates and their standard errors. Figure G.1 displays the age-varying time costs of working by age.

H What drives our estimates and identification

This appendix studies how the model's target moments change when we vary one of our 21 estimated parameters at a time (Table 1). These changes help us understand how our model's parameters are identified by the target moments we use.

The **discount factor** β is the only parameter with a large effect on the savings of both couples and singles across much of the life cycle (Figure H.1). Decreasing

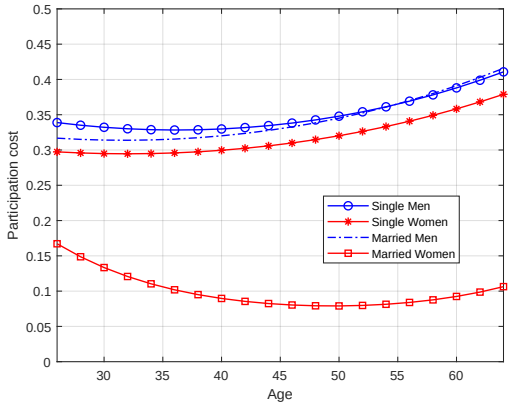


Figure G.1: Participation cost

patience lowers the participation of married women at all ages and slightly tilts the participation profile of single women, reducing it before age 40 and raising it after age 50. Hours change little.

The **weight on consumption** ω (Figure H.2) is the only parameter that shifts both participation and hours worked (conditional on participation) for all four subgroups (single and married men and women). As expected, reducing ω weakens incentives to work. It also lowers savings for couples and singles after age 60.

Turning to the effects of **bequest motives**, all of the parameters affecting terminal and side bequests have much larger effects on savings than on labor market outcomes. For the changes in parameter values that we examine here, large changes in bequest motives have large effects on savings, but little to no effect on participation and hours worked. It is worth noting that in the paper we show that eliminating bequests entirely can have substantial effects on labor supply. This is because this is a much larger change in parameters. Turning to the effects on wealth, terminal and side bequest motives have very different implications for couples and singles. Reducing the desire to leave terminal bequests lowers the wealth held by both couples and singles after age 60. In contrast, reducing the desire to leave side bequests reduces the wealth held by couples but tends to increase that of singles because the new singles turn out to be wealthier. Comparing the effects of the intensity and the shifter of the bequest parameters, we find that the shifter tilts the wealth-age profile more sharply at older ages, when those who survive tend to be those with higher lifetime income.

We next examine how **available time and participation costs** affect labor market outcomes. Reducing married women's available time by three hours per week (Figure H.7) lowers their participation by more than ten percentage points at age 26. It also reduces hours worked across the life cycle, although by less than the

mechanical weekly three-hours reduction, implying a decrease in their leisure. This change induces higher participation among husbands at older ages without affecting their hours worked, and reduces couples' wealth after age 55.

Participation costs vary across groups and are governed by a quadratic function of age (Equation 3). For each group, we estimate three parameters and plot the implied fixed costs for each parameter change.

Raising the constant term in married women's participation cost by three hours (Figure H.8) lowers participation among younger married women even further and steepens their participation profile, with a much smaller effect on hours worked. The effects on married men's participation and couples' wealth are in the same direction but smaller. Increasing the linear term in the participation cost (Figure H.9) raises costs with age, driving married women's participation to zero at older ages and reducing their hours worked, especially later in life. It also raises married men's participation after age 40 and reduces the participation of single women by expanding the pool of divorced women with lower human capital. The resulting decline in couples' retirement wealth is large. Adjusting the quadratic term in the participation cost (Figure H.10) steepens them later on, shifting the timing of participation.

We then study changes in available time and participation costs for married men. Reducing their available time by three hours per week (Figure H.11) has little effect at younger ages but leads to faster declines in participation and hours worked later in life. Married women respond by increasing participation, although their hours remain stable. Couples' wealth holdings decline. Raising married men's participation cost reduces participation at older ages and decreases hours worked more sharply than reducing available time. Changes in the linear and quadratic terms of the participation cost function similarly alter the slopes of participation and hours over the life cycle.

We next turn to single women. Figures H.15 to H.18 show that reducing available time lowers both participation and hours worked, with a larger effect on participation. Lowering the constant in their participation cost reduces participation even more and has a smaller effect on hours. Changes in the linear and quadratic terms affect the age profiles of participation and hours differently, shifting both timing and magnitude.

Finally, we examine single men. Figures H.19 to H.21 show the effects of adjusting the constant, linear, and quadratic terms of their participation costs (with a normalized time endowment). Changes in these terms affect participation at different points in the life cycle—earlier for the constant term, and later for the linear and quadratic terms—and similarly influence hours worked conditional on participation.

Our model's identification thus rests on the distinct ways structural parameters map into savings, participation, hours worked, and wealth profiles across groups and over the life cycle. Preference parameters such as β and ω influence savings and labor supply differently across subgroups, allowing their identification. Bequest motives affect savings but leave labor supply largely unchanged for the range of parameter changes considered, and their differential effects on couples versus singles allow us to separately identify side and terminal bequests. Participation costs and time endowments are identified through their differential impacts on participation rates versus hours worked, exploiting variation over age, gender, and marital status. The rich structure of life-cycle behavior across demographic groups provides the variation necessary to pin down each parameter.

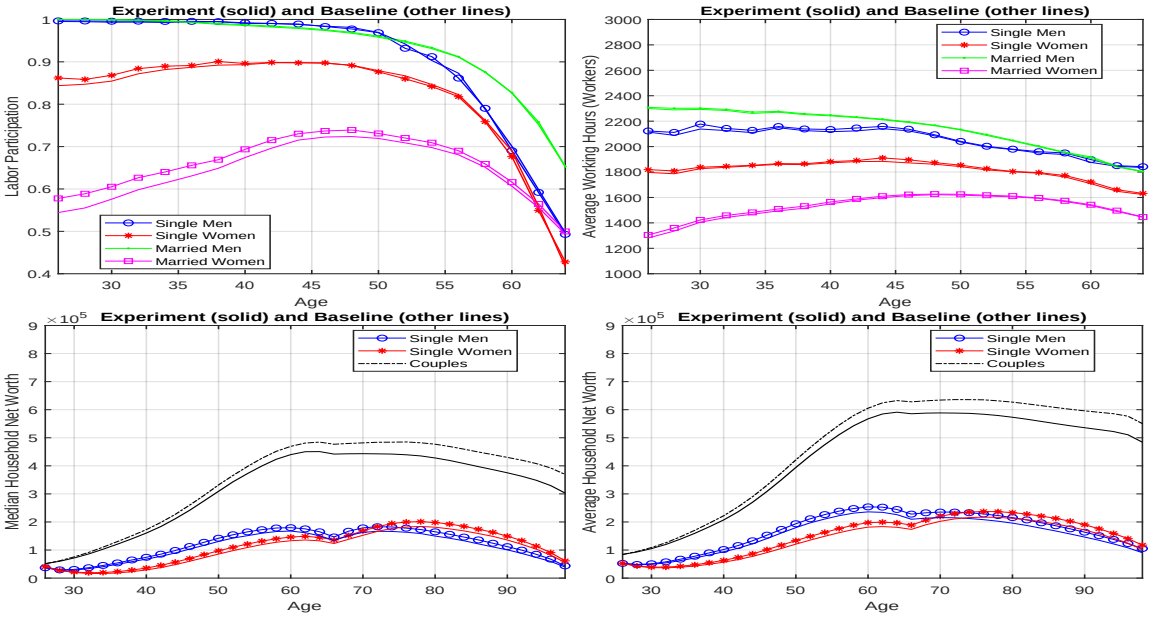


Figure H.1: Discount factor decreased by 0.2% (solid), benchmark (other lines)

Table H.1 in Appendix H reports the GMM criterion disaggregated by group and outcome for the benchmark model and for experiments varying one parameter at a time. In all cases, the GMM criterion rises, confirming that the model's parameters are well identified given the target moments we have chosen.

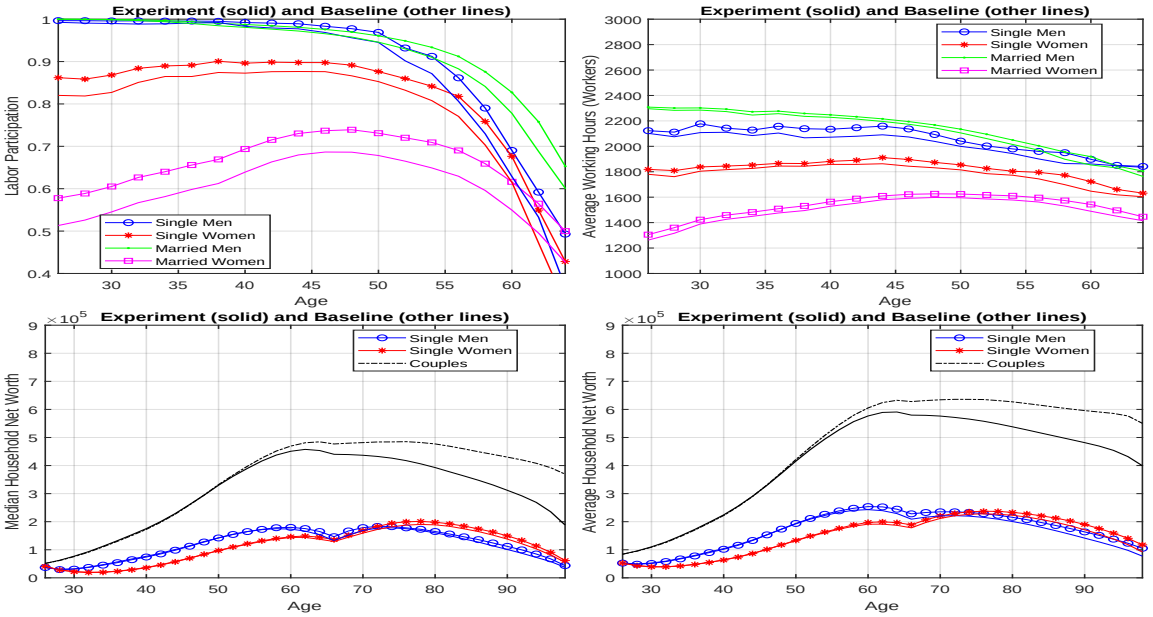


Figure H.2: Weight on consumption decreased by 2.0% (solid), benchmark (other lines)

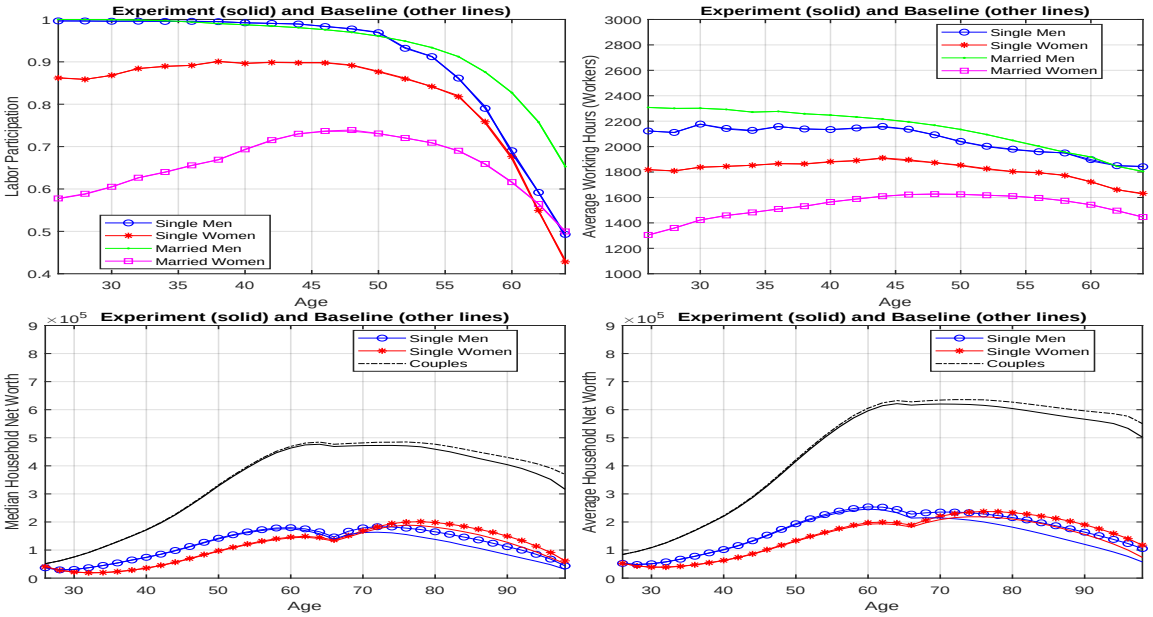


Figure H.3: Terminal bequest strength decreased by 20% (solid), benchmark (other lines)

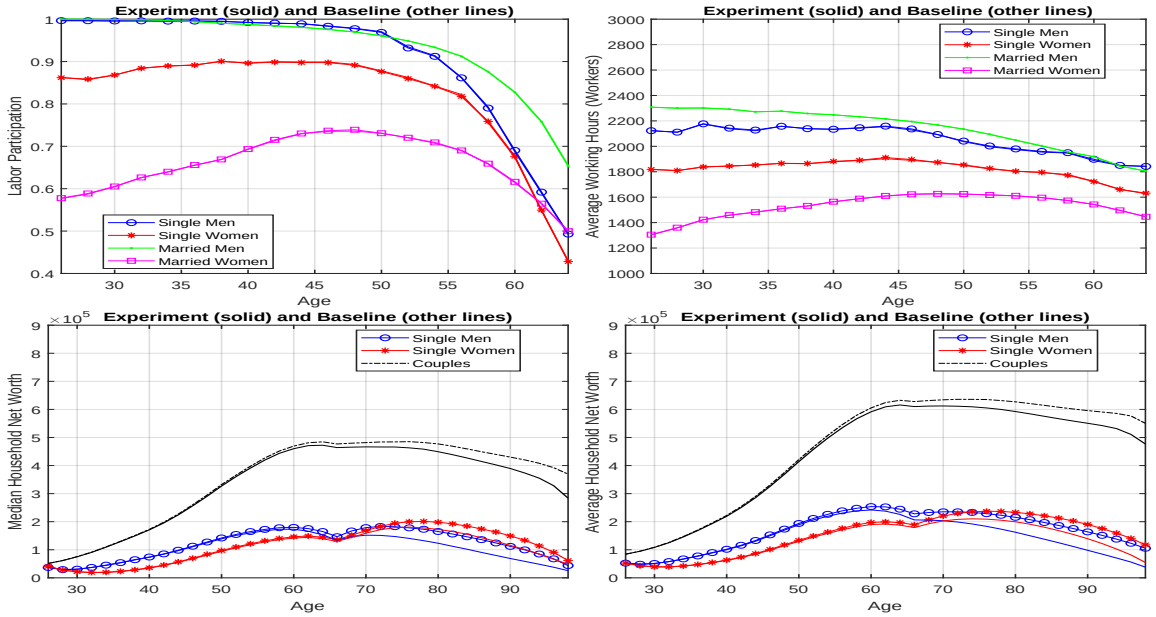


Figure H.4: Terminal bequest shifter increased by 20% (solid), benchmark (other lines)

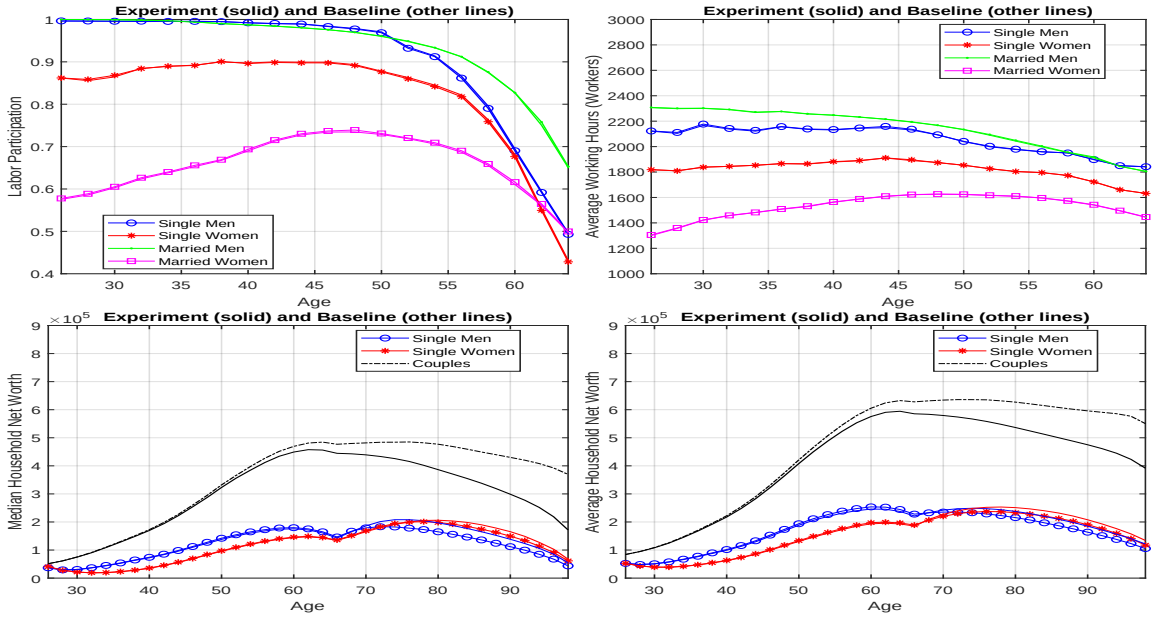


Figure H.5: Side bequest strength decreased by 20% (solid), benchmark (other lines)

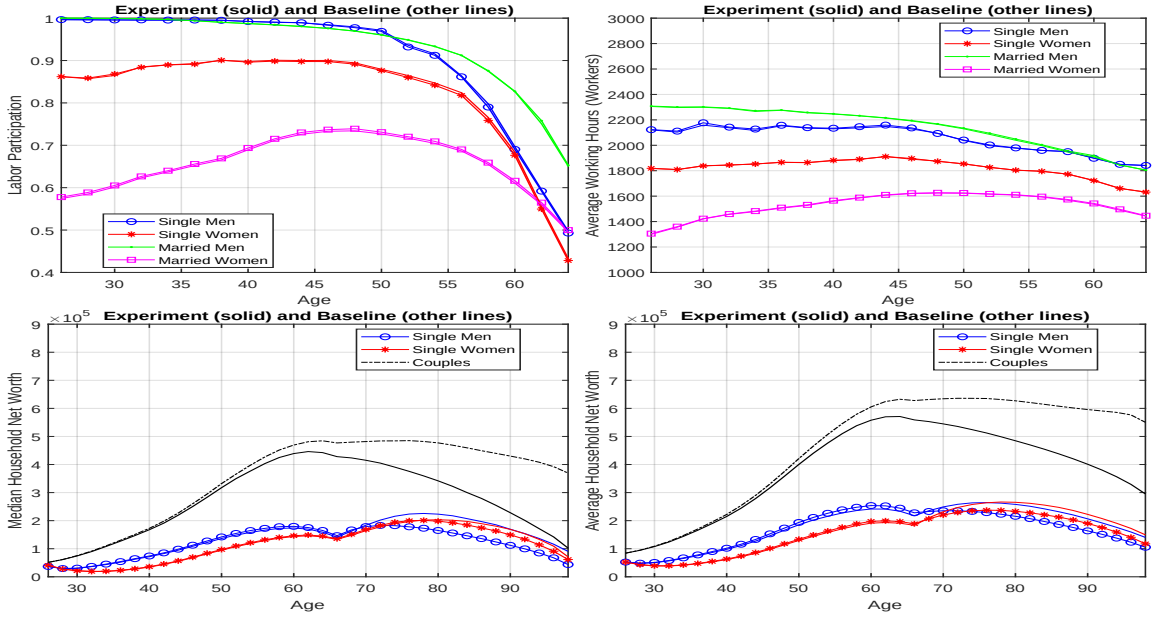


Figure H.6: Side bequest shifter increased by 20% (solid), benchmark (other lines)

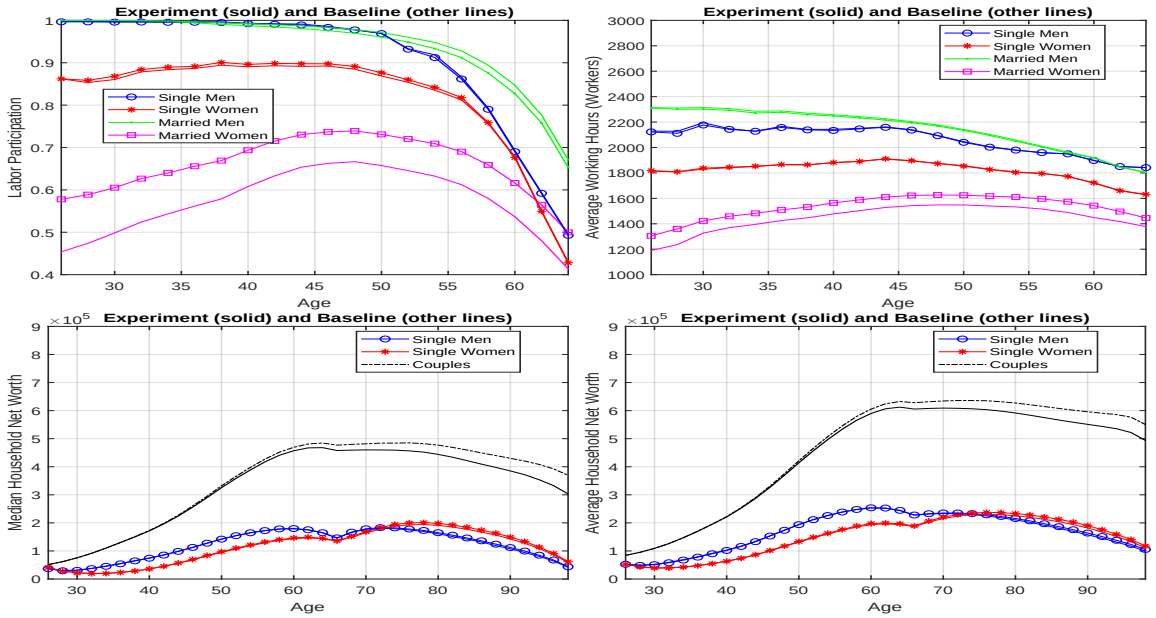


Figure H.7: Married women's available time decreased by 3 hours a week (solid), benchmark (other lines)

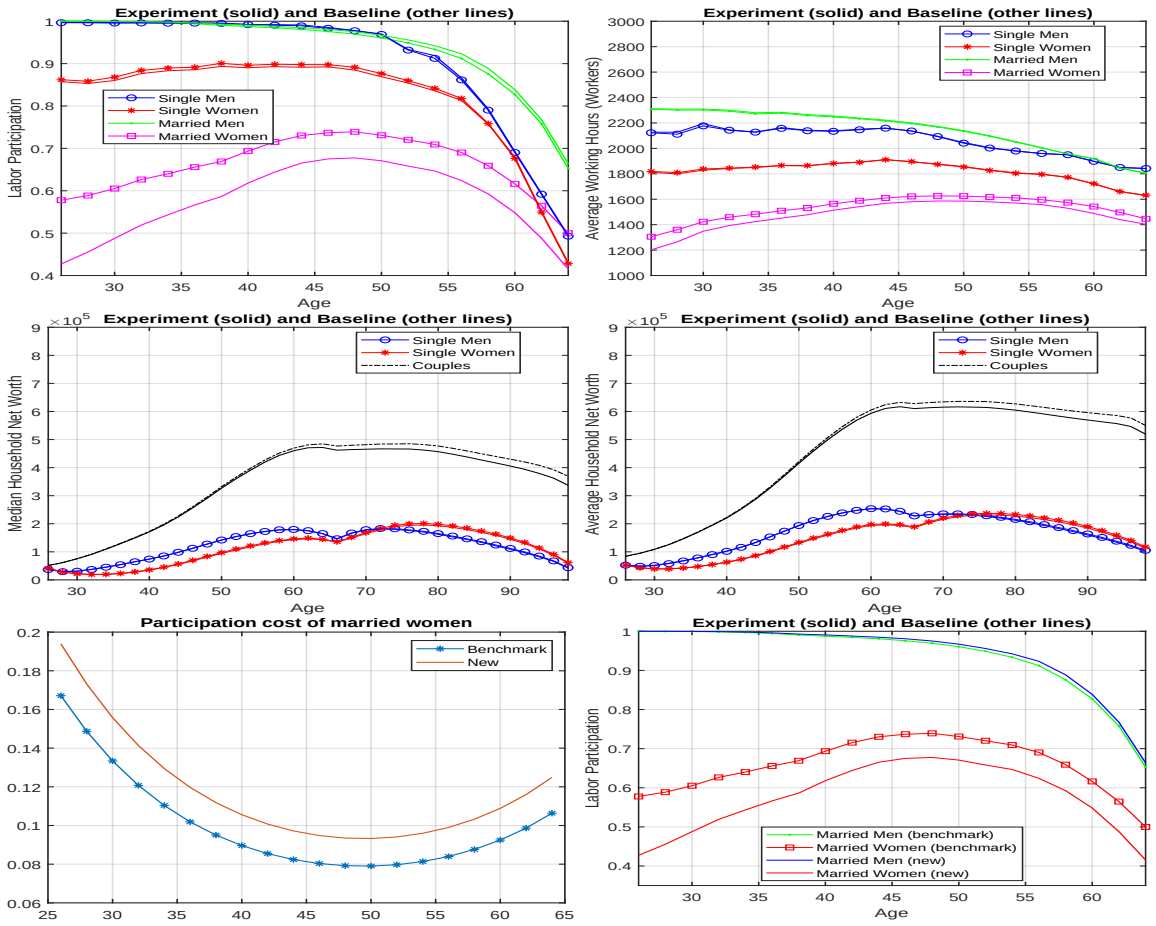


Figure H.8: Married women: constant in the participation cost increased by 3 hours per week (solid), benchmark (other lines)

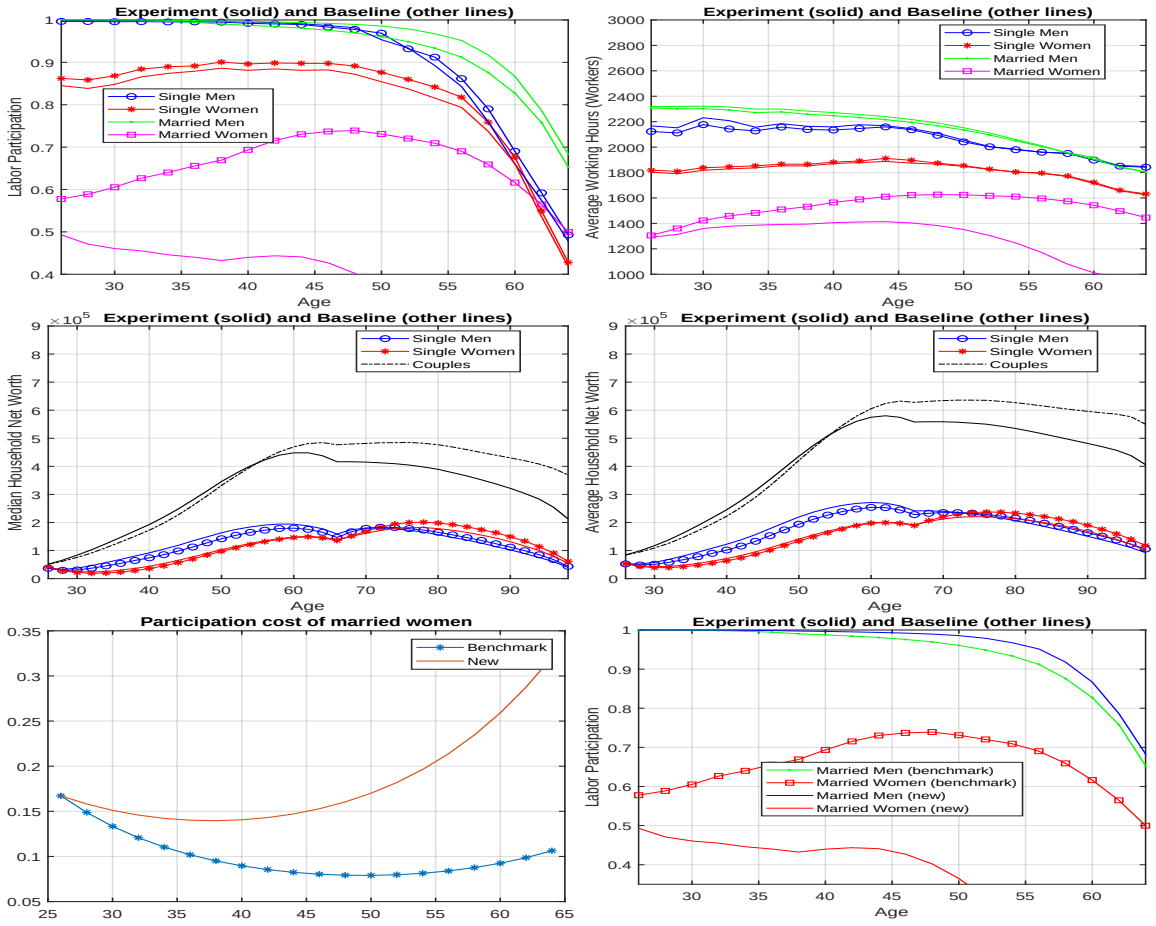


Figure H.9: Married women: linear term in age in the participation cost ($p_1^{2,2}$) increased by 50% (solid), benchmark (other lines)

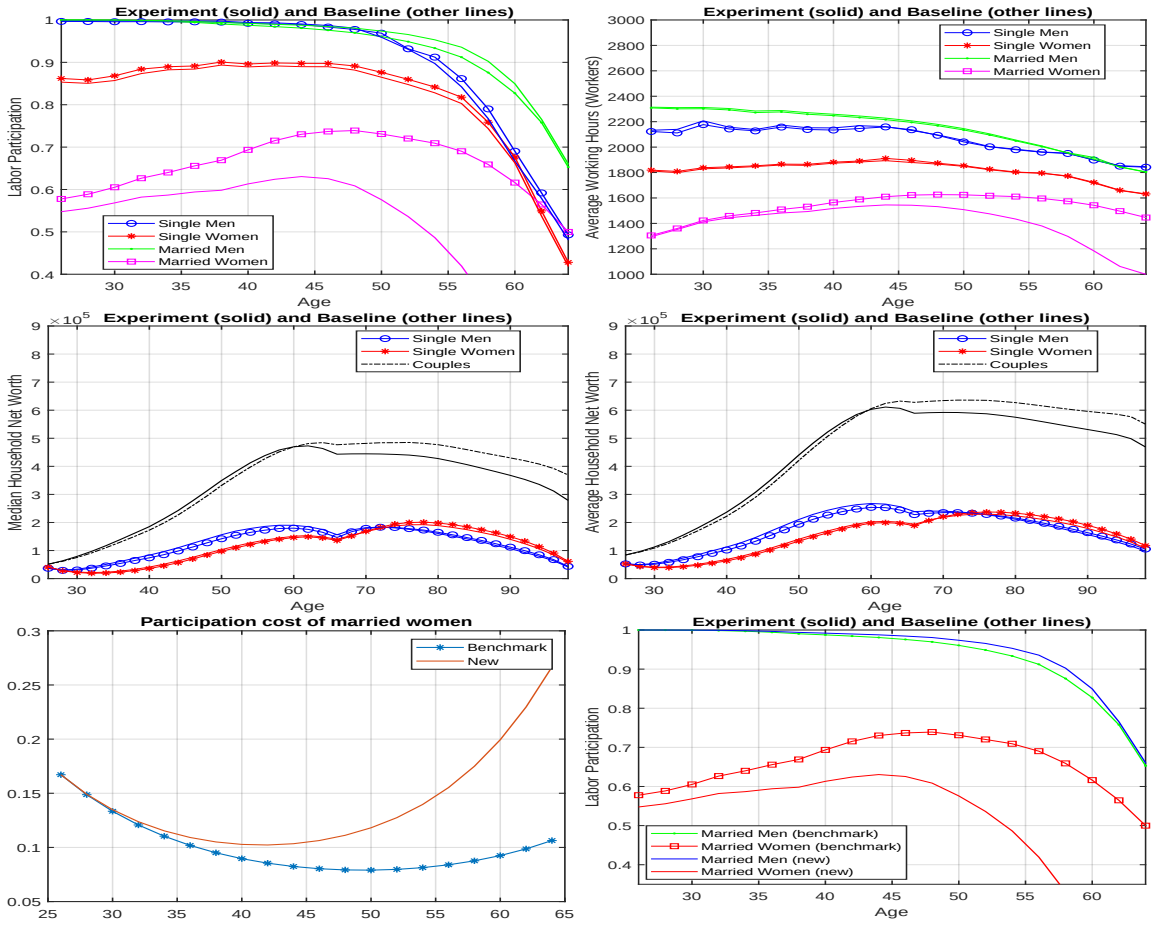


Figure H.10: Married women: quadratic term in age in the participation cost ($p_2^{2,2}$) increased by 50% (solid), benchmark (other lines)

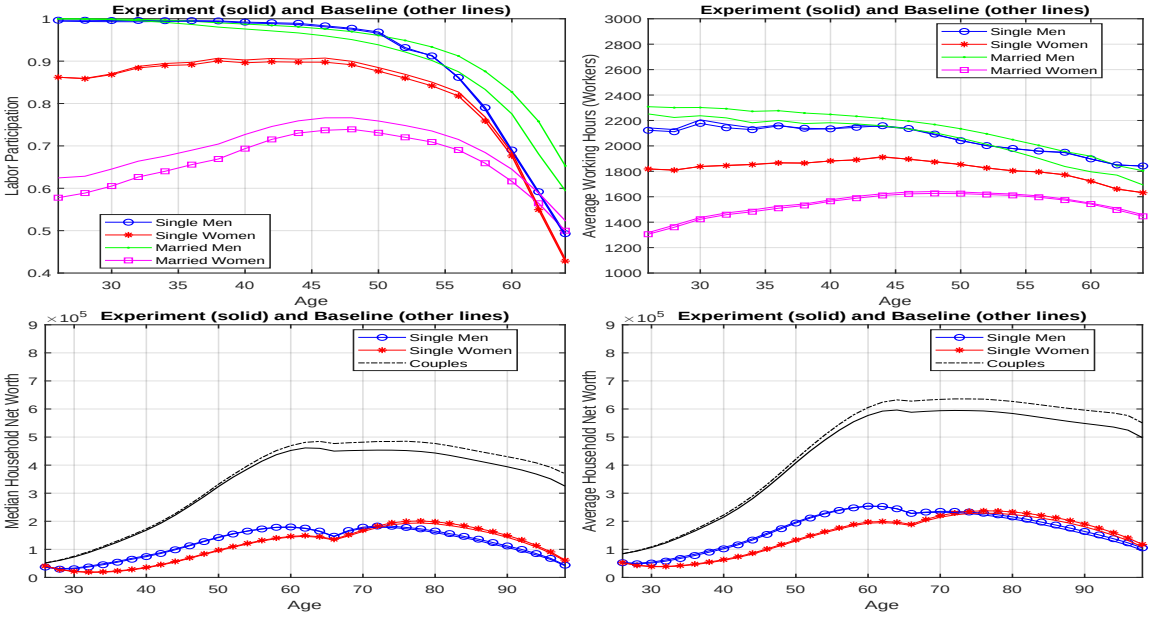


Figure H.11: Married men's available time decreased by 5 hours per week (solid), benchmark (other lines)

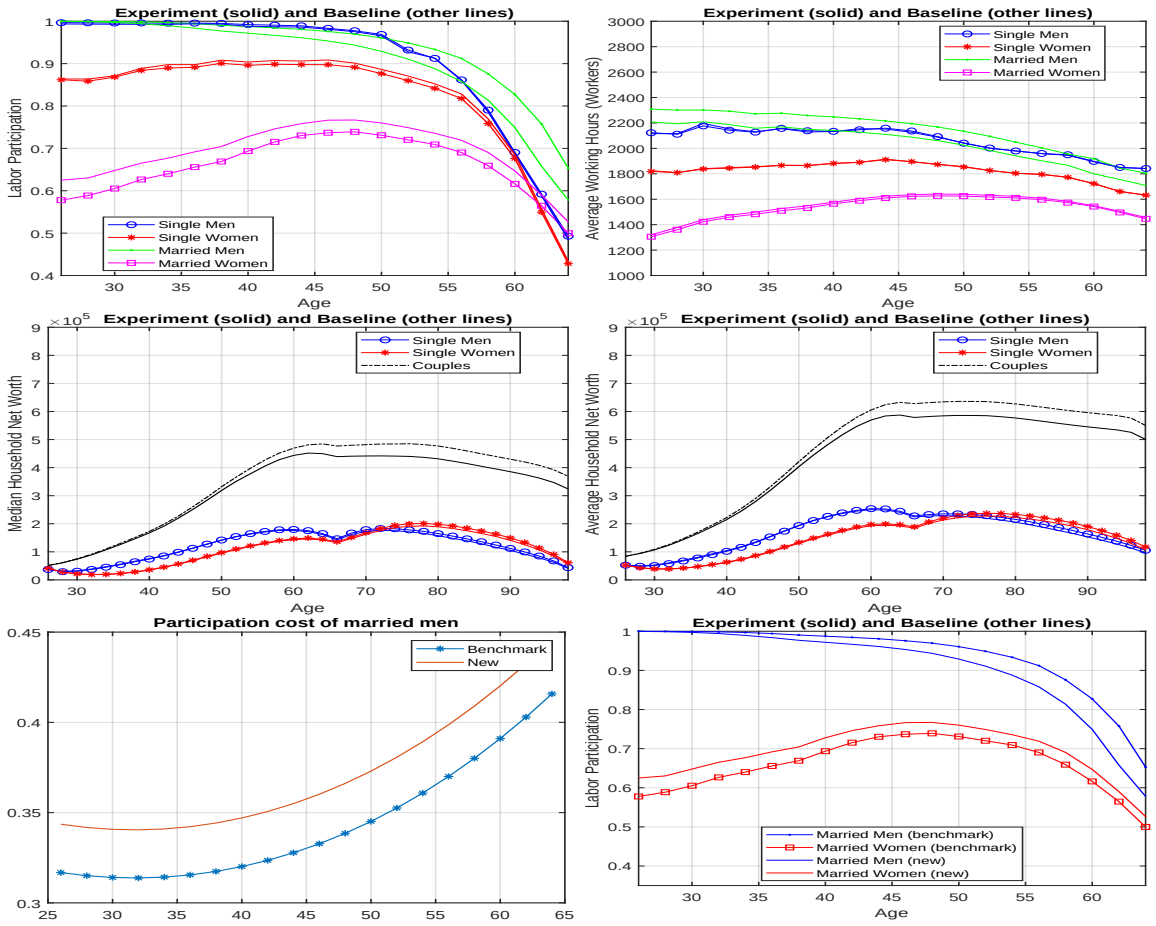


Figure H.12: Married men: constant in the participation cost increased by 3 hours per week (solid), benchmark (other lines)

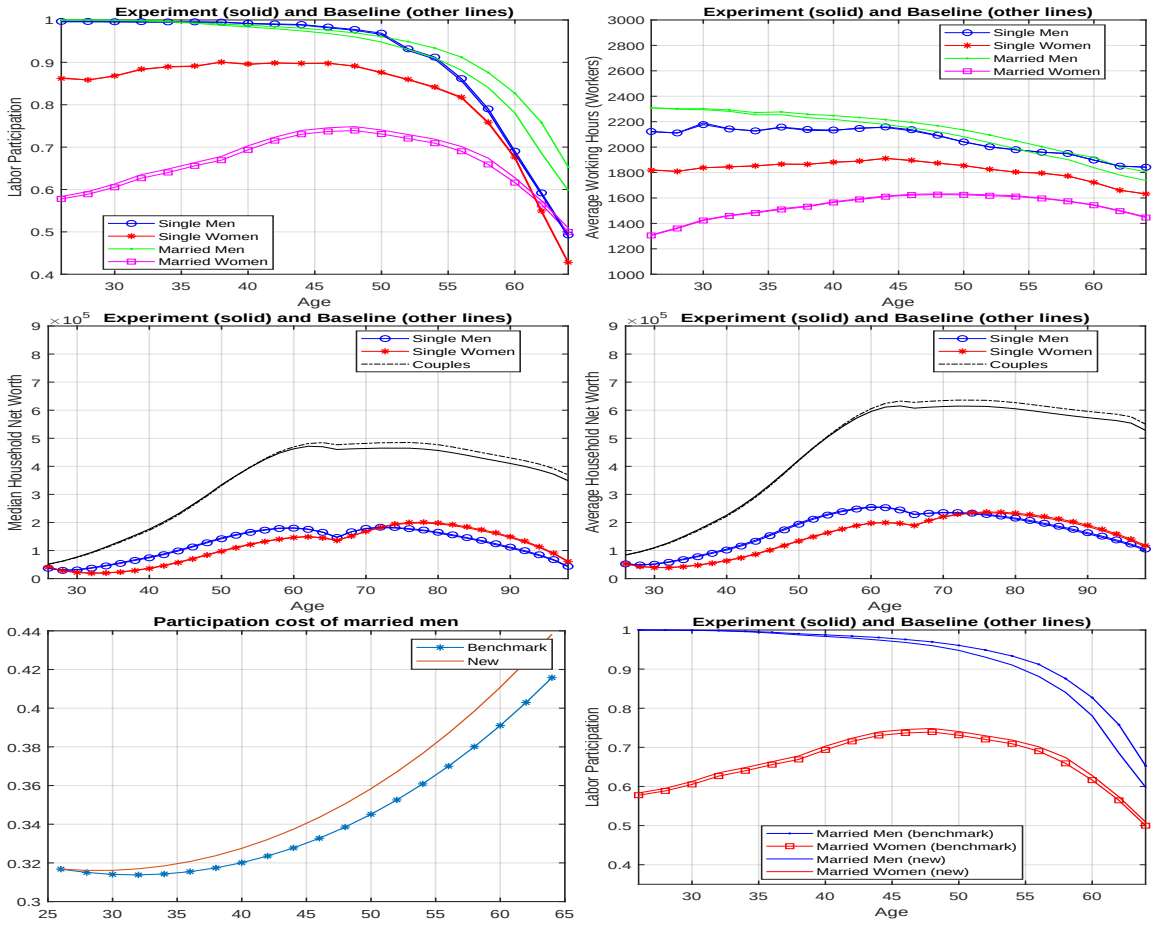


Figure H.13: Married men: linear term in age in the participation cost ($p_1^{1,2}$) increased by 50% (solid), benchmark (other lines)

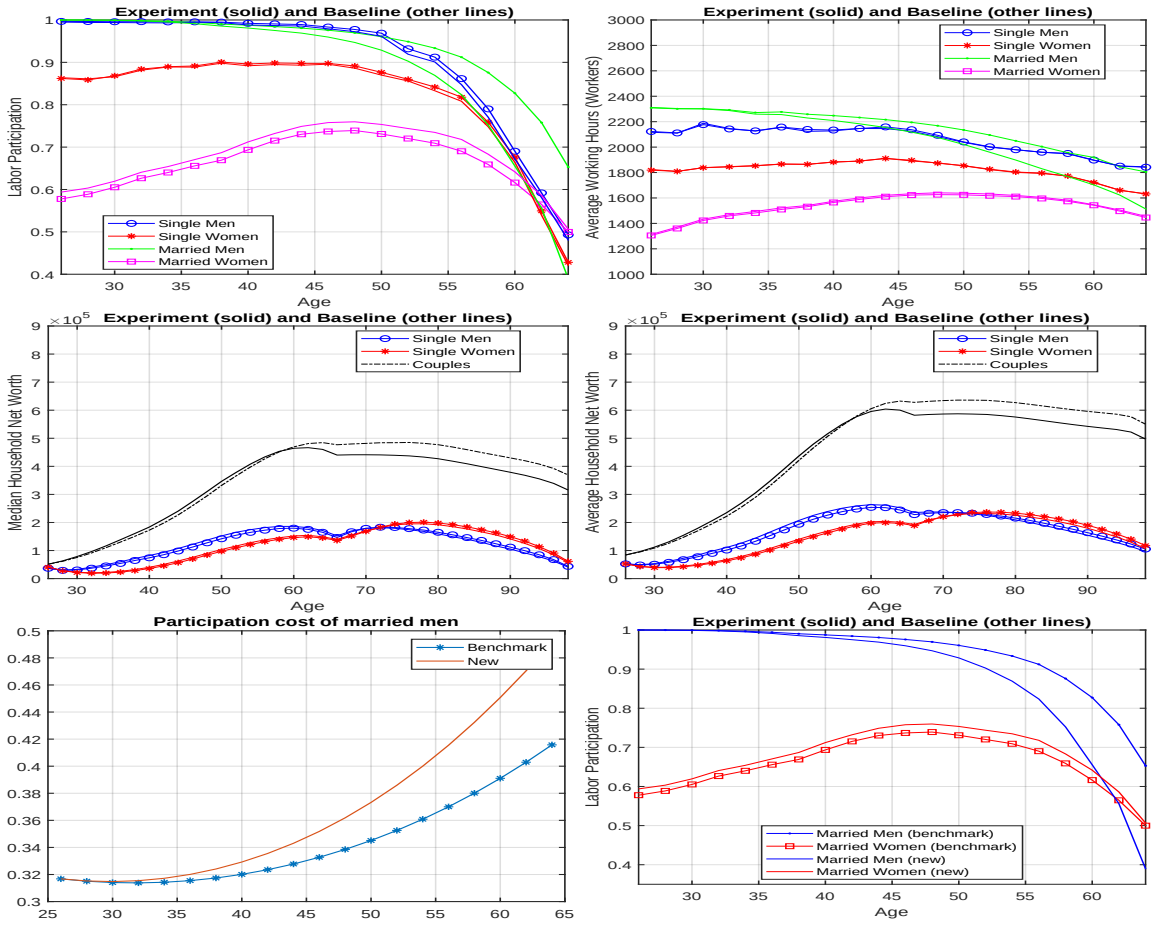


Figure H.14: Married men: quadratic term in age in the participation cost ($p_2^{1,2}$) increased by 50% (solid), benchmark (other lines)

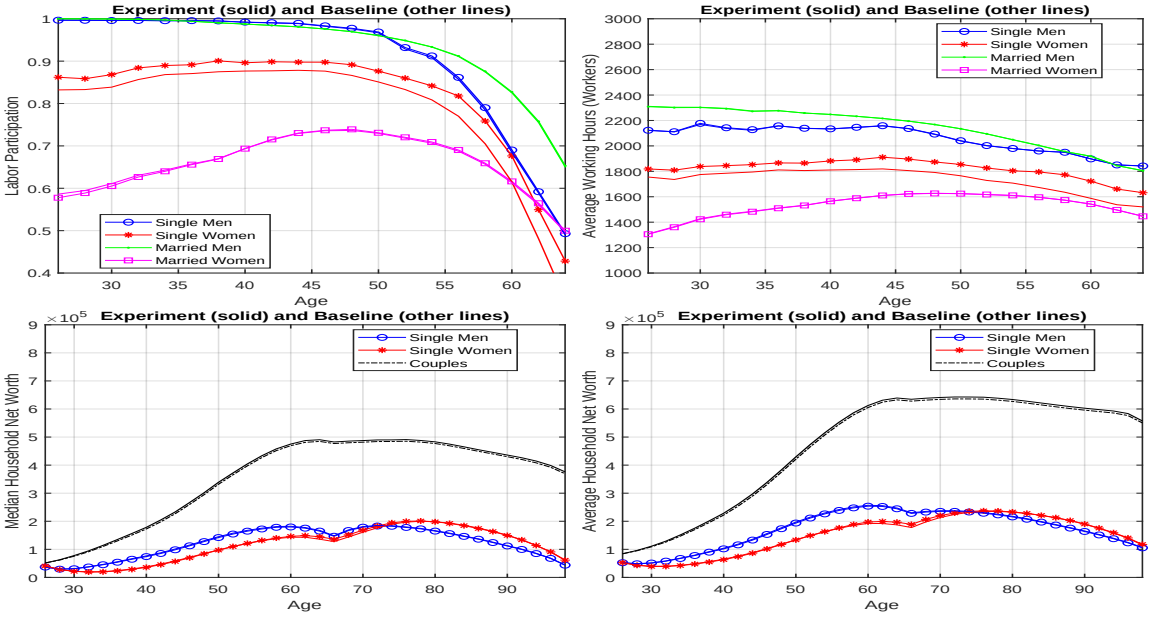


Figure H.15: Single women: available time decreased by 3 hours per week (solid), benchmark (other lines)

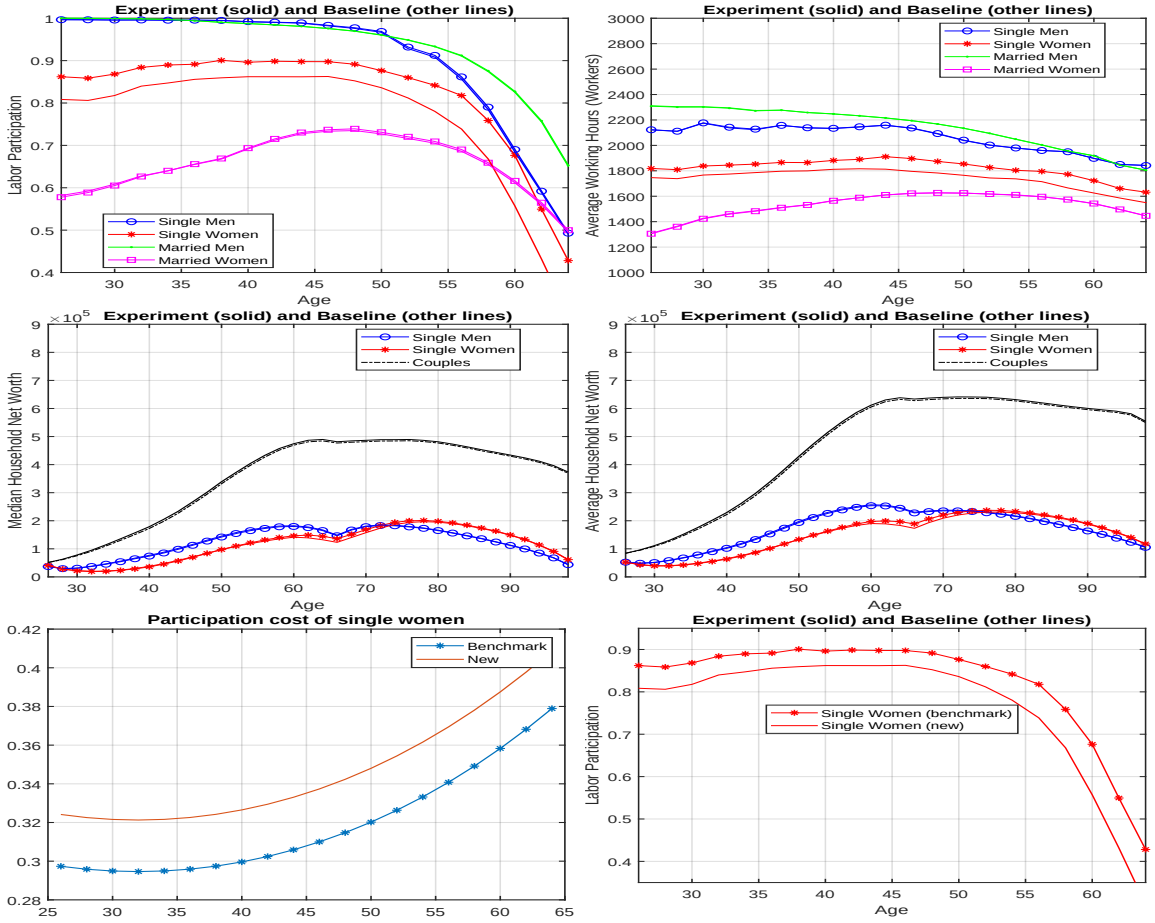


Figure H.16: Single women: constant in the participation cost increased by 3 hours per week (solid), benchmark (other lines)

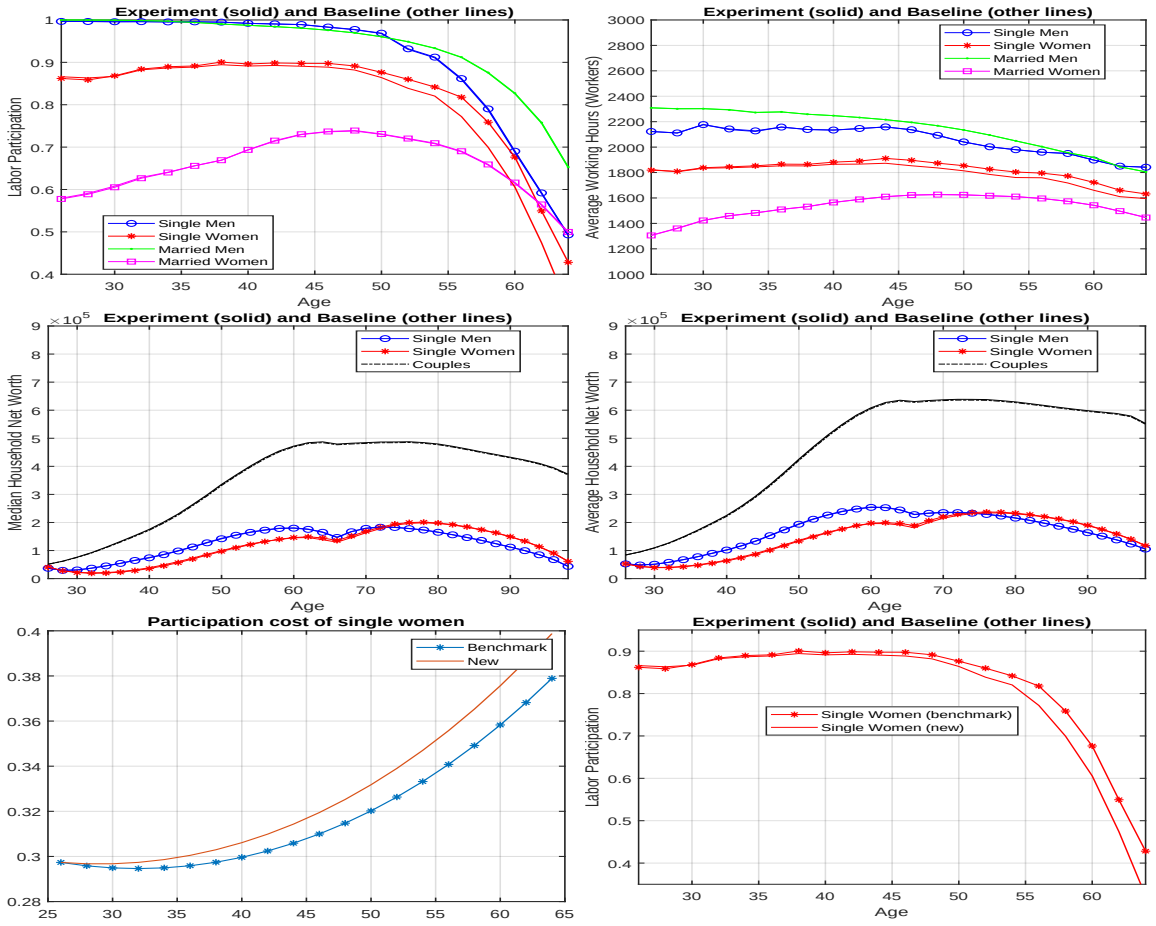


Figure H.17: Single women: linear term in age in participation cost ($p_1^{2,1}$) increased by 50% (solid), benchmark (other lines)

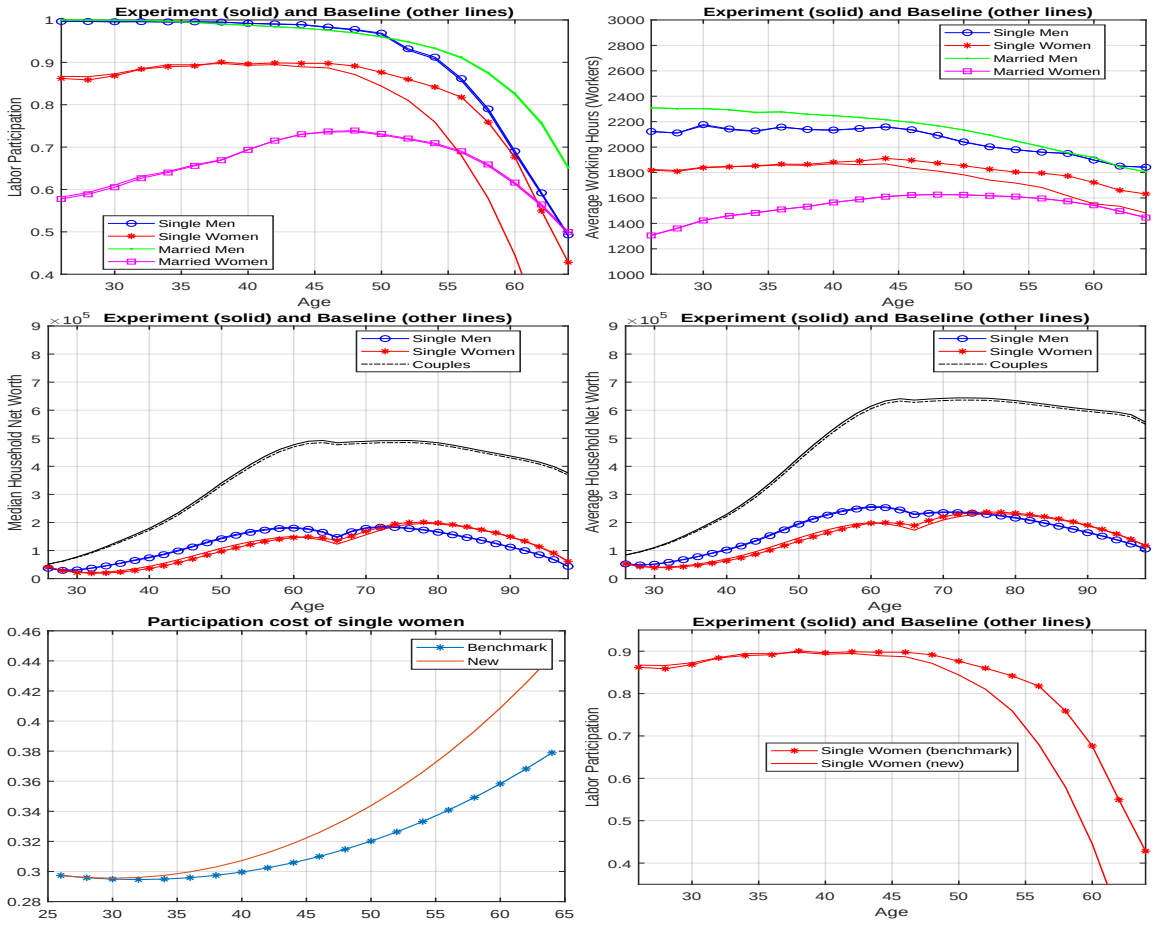


Figure H.18: Single women: quadratic term in age in the participation cost ($p_2^{2,1}$) increased by 50% (solid), benchmark (other lines)

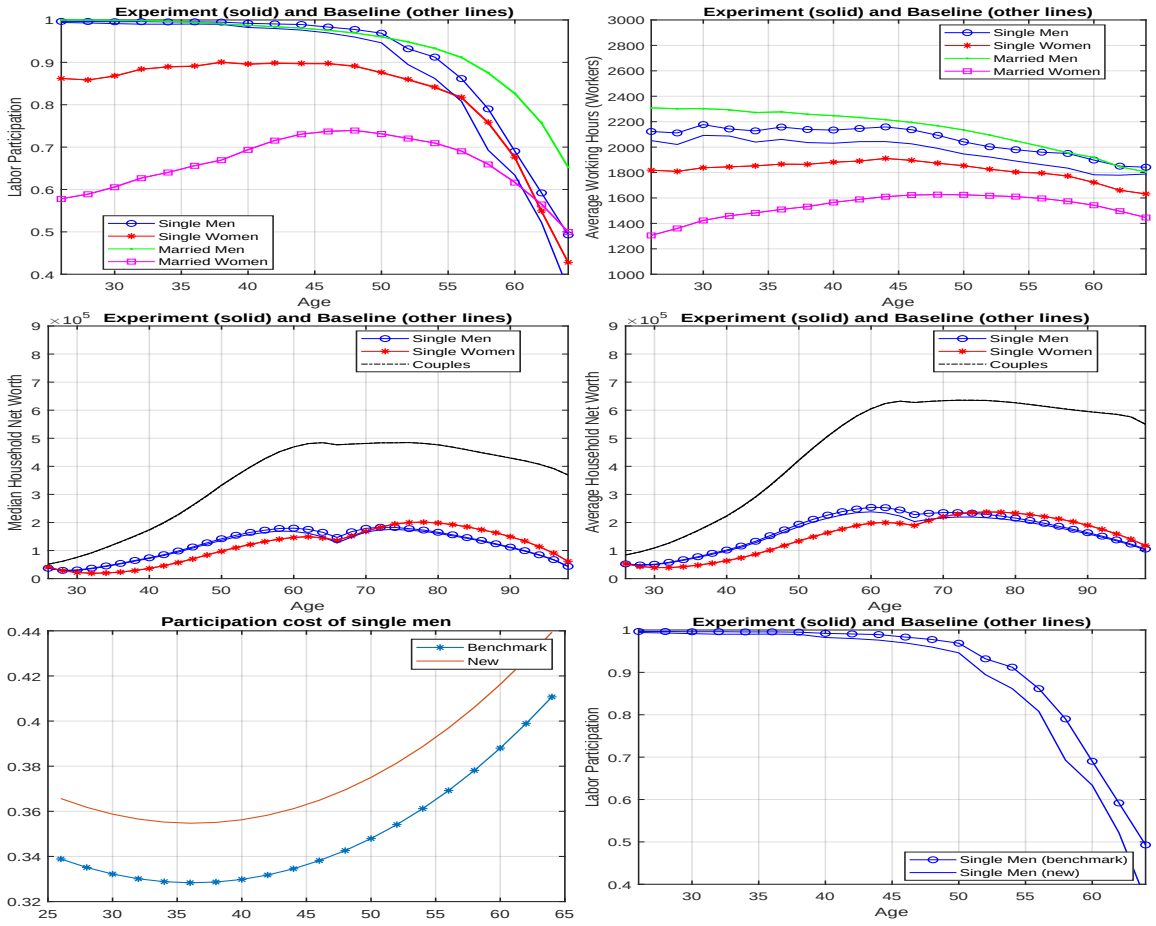


Figure H.19: Single men: constant term in the participation cost ($p_0^{1,1}$) increased by 50% (solid), benchmark (other lines)

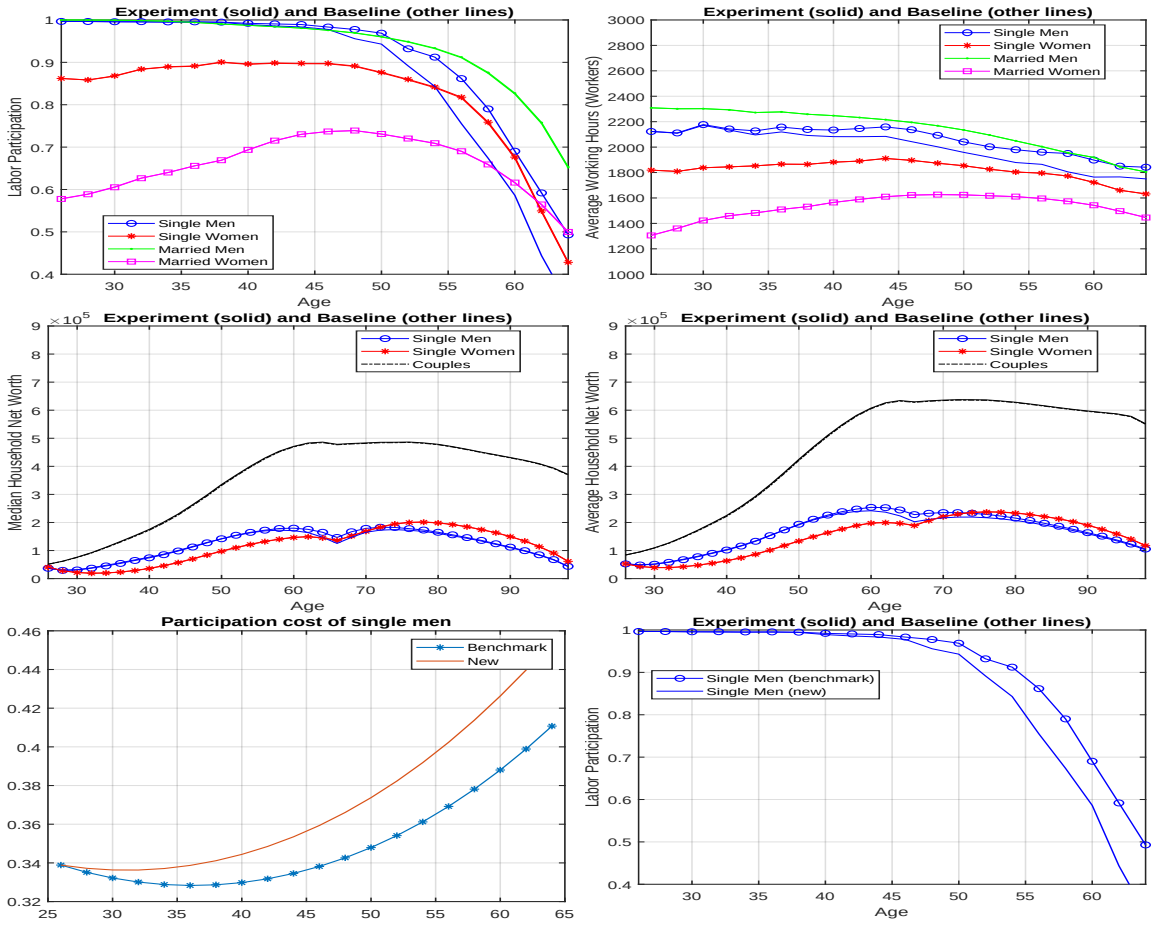


Figure H.20: Single men: linear term in age in the participation cost ($p_1^{1,1}$) increased by 50% (solid), benchmark (other lines)

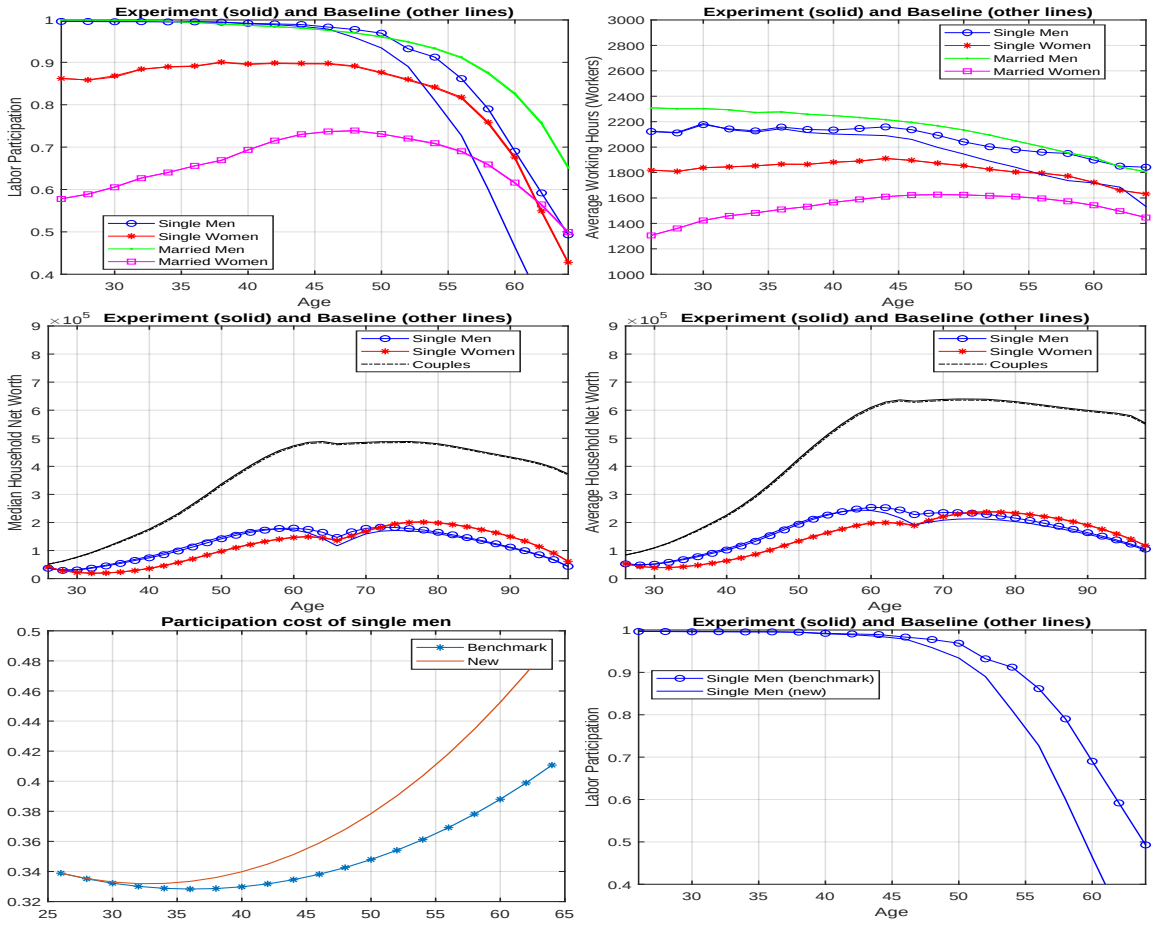


Figure H.21: Single men: quadratic term in age in the participation cost ($p_2^{1,1}$) increased by 50% (solid), benchmark (other lines).

	Mean wealth			Median wealth			Hours among workers				Participation				
Specification	SM	SW	C	SM	SW	C	SM	SW	MM	MW	SM	SW	MM	MW	Total
Benchmark	96	210	138	74	322	340	75	20	27	54	38	29	119	67	1608
$\beta \downarrow 0.2\%$	11%	39%	14%	38%	25%	28%	4%	9%	-13%	83%	30%	14%	12%	-4%	24.1%
$\omega \downarrow 2.0\%$	8%	65%	-16%	6%	43%	-10%	-30%	332%	13%	372%	138%	102%	77%	68%	43.5%
$\phi_0 \downarrow 20\%$	44%	4%	-1%	52%	0%	-3%	2%	-2%	1%	2%	5%	4%	1%	-1%	5.3%
$k_0 \uparrow 20\%$	83%	10%	12%	98%	3%	-1%	2%	-1%	1%	2%	8%	3%	1%	-1%	12.7%
$\phi_1 \downarrow 20\%$	-33%	72%	14%	-22%	55%	6%	4%	-7%	-8%	5%	7%	1%	6%	1%	20.8%
$k_1 \uparrow 20\%$	-49%	174%	40%	-21%	129%	6%	5%	-12%	-8%	8%	10%	0%	11%	3%	50.7%
$p_2^{1,1} \uparrow 50\%$	19%	4%	2%	21%	2%	-2%	86%	7%	-15%	-1%	277%	1%	1%	3%	13.2%
$p_1^{1,1} \uparrow 50\%$	-0%	1%	1%	7%	-1%	-2%	-6%	5%	-10%	-0%	205%	1%	1%	2%	4.5%
$p_0^{1,1} \uparrow 18\%$	91	207	138	85	315	331	52	21	24	54	155	29	120	68	1689
$p_2^{2,1} \uparrow 50\%$	4%	9%	-2%	-2%	5%	-24%	-5%	1288%	-14%	-3%	3%	299%	2%	5%	18.3%
$p_1^{2,1} \uparrow 50\%$	2%	2%	-12%	-1%	0%	-12%	-2%	251%	-6%	-0%	2%	101%	1%	2%	1.7%
$p_0^{2,1} \uparrow 15\%$	100	216	109	72	318	313	70	142	24	56	39	103	121	70	1755
$p_2^{1,2} \uparrow 50\%$	35%	43%	3%	8%	23%	-16%	-18%	24%	1924%	-2%	8%	4%	415%	-2%	71.6%
$p_1^{1,2} \uparrow 50\%$	9%	0%	-4%	-0%	-5%	-7%	-7%	4%	4%	-6%	3%	1%	134%	-2%	7.3%
$p_0^{1,2} \uparrow 16\%$	102	258	130	77	354	335	70	16	117	73	39	29	633	70	2303
$p_2^{2,2} \uparrow 50\%$	45%	53%	-1%	14%	28%	-22%	-10%	45%	238%	2715%	-17%	7%	-15%	1077%	149.9%
$p_1^{2,2} \uparrow 50\%$	74%	86%	-17%	25%	47%	-37%	-11%	85%	548%	9435%	-30%	6%	-26%	2718%	454.4%
$p_0^{2,2} \uparrow 11\%$	98	211	126	76	316	339	79	23	53	488	36	27	109	145	2124
$L^{2,1} \downarrow 3\%$	4%	4%	-16%	-3%	0%	-7%	-5%	287%	-7%	-2%	4%	322%	1%	5%	7.2%
$L^{1,2} \downarrow 3\%$	9%	19%	0%	-1%	8%	4%	-5%	-20%	101%	26%	-7%	-0%	284%	2%	28.6%
$L^{2,2} \downarrow 4\%$	4%	6%	-13%	2%	1%	-6%	6%	14%	177%	843%	-8%	-4%	-17%	300%	41.6%

Table H.1: Effects of parameter changes on GMM criteria. SM: single men; SW: single women; C: couples; MM: married men; MW: married women.

I Additional counterfactuals

I.1 Bequests

No side bequest motives. We begin by removing the utility from leaving side bequests, that is, bequests to heirs other than the spouse when the first member of a couple dies. Since side bequests are never accidental, this change also eliminates all bequests upon the first spouse's death.

The top row of Figure I.1 displays the life cycle profiles for median wealth and the percentage change in average labor income for our model economy with and without side bequests. Starting from couples, the left-hand-side graph shows that eliminating the desire to leave side bequests substantially reduces couples' wealth, with larger effects at older ages. By age 80, median wealth declines by more than 35%. The right-hand-side graph illustrates labor income changes: removing side bequests lowers earnings for both married men and women, with a larger reduction for women–about 2% over much of their working lives, compared to about 1% for men past age 55. These findings indicate that side bequests have large effects on couples' savings and non-negligible impacts on their labor supply. Both effects are already present by age 50, well before retirement.

The effects on singles are also noteworthy. Without side bequest motives, singles save slightly less until age 66, anticipating that after marriage, they would inherit more wealth in case their spouse dies. After age 66, the saving patterns of singles reflect both the earlier decrease in retirement savings and a composition effect. In this counterfactual, couples save less and leave no side bequests, resulting in new widows and widowers inheriting more wealth than existing singles. As more individuals lose their spouse, the median wealth among singles rises, generating the bump in percentage differences in wealth accumulation seen in the top right panel. Side bequests have small effects on the labor market behavior of singles. Young singles save less in the absence of side bequests and expect larger inheritances if they marry, leading them to work slightly less until age 50. As they age and the probability of marriage declines, they eventually work more to compensate for lower savings accumulated earlier. Thus, while side bequest motives affect singles, the largest impact arises from differences in the wealth of new widow(er)s and other single individuals.

No terminal bequest motives. The graphs in the bottom row in Figure I.1 examine the effects of removing the utility from terminal bequests, for both couples and singles. In our model, households do not fully annuitize their wealth, and uncertain lifespans generate accidental bequests. Eliminating terminal bequest motives

removes only voluntary bequests at the death of the last household member, but not side bequests and accidental bequests.

The bottom-left graph shows that removing terminal bequests reduces couples' wealth throughout the life cycle, though the effect of terminal bequests is smaller than that of side bequests. At age 80, median wealth is about 15% lower without terminal bequests but over 35% lower without side bequests. The additional wealth accumulation due to terminal bequests is financed by increased labor supply from married women during all of the working period and a slight reduction in consumption over much of married life.

For singles, eliminating terminal bequest motives leads to significantly lower median wealth at all ages, with a larger percentage change effect than for couples. This occurs because couples still retain side bequest motives, while singles never have them. Singles save less before age 66 because they no longer wish to leave a terminal bequest, whether they remain single or marry. After 66, a composition effect further lowers their wealth. Since couples are poorer without terminal bequests but still leave side bequests, widows and widowers inherit less, narrowing the wealth gap between new and existing singles. Consequently, median wealth among singles during retirement is much lower than in the benchmark: at age 80, it is more than 50% lower for men and about 30% lower for women. In the absence of terminal bequest motives, single men work slightly less, and both single men and women consume slightly more.

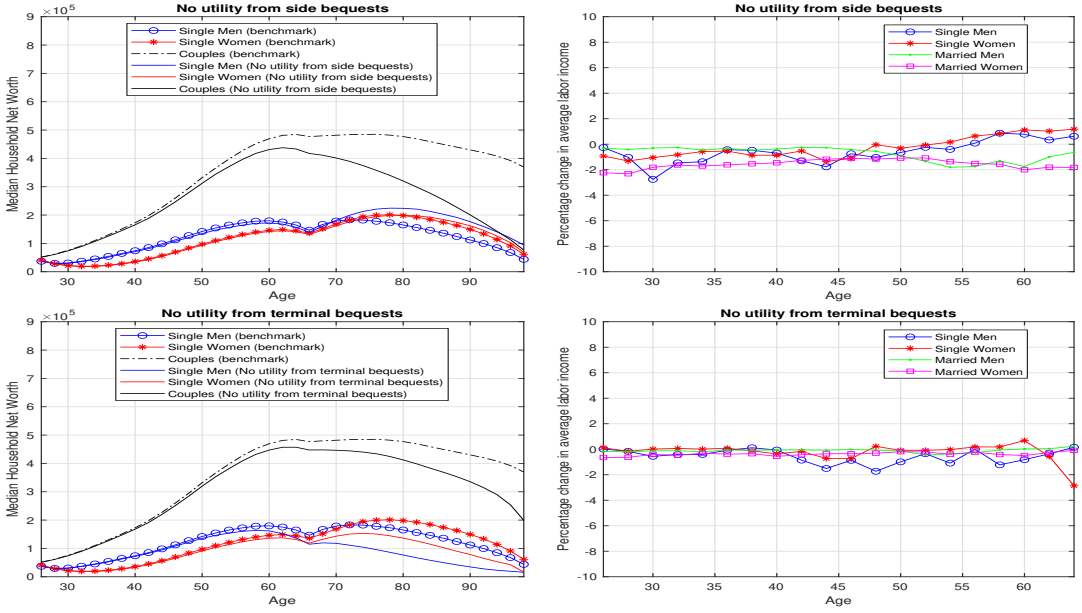


Figure I.1: The effects of bequests

I.2 Medical expenses

No additional medical expenses at death. The panels in Figure I.2 show changes in median wealth and average labor income when we eliminate the spike in medical expenses at death while maintaining regular medical expenditure risk in all periods. The left panel shows that wealth holdings remain nearly unchanged until age 60 for all three groups of households. After that age, saving behavior becomes more heterogeneous: couples hold slightly less wealth, while singles' wealth stagnates but then rises fast after age 90. To understand these patterns is important to keep in mind that the wealth of singles after age 66 is also affected by a composition effect: widows and widowers, who inherit wealth from their spouses and tend to be richer than previously singles individuals, enter the singles population after that age. Thus, their wealth is affected by both this composition effect and by the fact that the consumption floor provides less of an implicit tax on savings in the absence of medical expenses. The right panel indicates that this change has little effect on labor income.

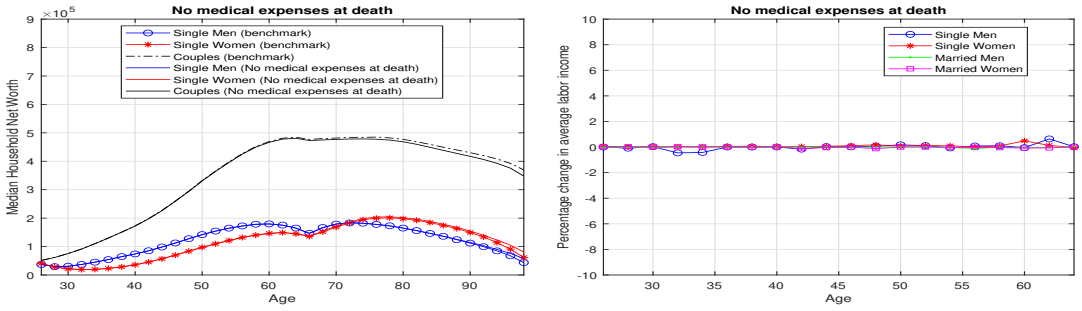


Figure I.2: The effects of medical expenses during retirement

I.3 Marriage and divorce risks

No new marriages. When marriage is no longer an option, two main forces shape outcomes. First, individuals adjust their expectations about the future and modify their behavior accordingly. Second, a composition effect emerges, changing who remains single and who marries after entering our model. In our benchmark economy, as in the data, men and women who earn higher wages are more likely to marry. As a result, during the working years, the group of single individuals in the benchmark economy tends to include more lower-wage individuals than in the no-new-marriages counterfactual.

The top left panel of Figure I.3 compares the fraction of married men and women

in our benchmark and counterfactual economies. At age 26, about 87% of both men and women in our cohort are married, and this fraction declines only gradually until age 66, reaching 80% for women and 85% for men. After that age, mortality drives further declines in marriage. The counterfactual with no new marriages reveals that, given historically realized divorce rates, new marriages play a crucial role in sustaining the married population. Without them, the fraction of married men and women falls below 40% by age 66 and continues to decline rapidly thereafter. This highlights the potentially large impact of the composition effect on household dynamics.

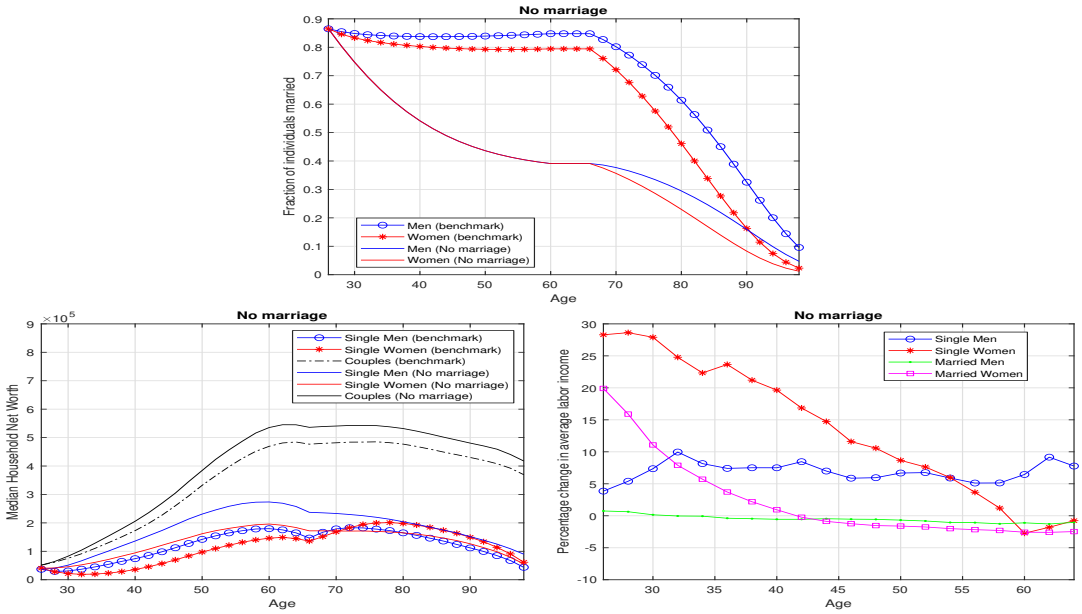


Figure I.3: The effects of not having any new marriages

The bottom left panel of Figure I.3 shows that until retirement, both single men and single women hold more wealth in the counterfactual economy with no new marriages. For instance, at age 38, single women and men hold over 200% more and 80% more wealth, respectively. This increase partly occurs because, without new marriages, individuals can no longer rely on finding a spouse and benefiting from the economies of scale in marriage. And partly, because single individuals now tend to have higher wages due to a composition effect of higher-wages singles not marrying. Later on in life, the savings patterns of single women and single men differ. Single women's wealth decreases fast after age 66. As their wealth at retirement is lower, they find it more advantageous to save less and rely on the consumption floor to insure against longevity risk and medical expenses. In contrast, single men, who have higher wages during their working ages, hold more wealth at retirement and depend less on social insurance, hence, they self-insure through savings and stay wealthier

even after that age. Couples also hold more wealth–up to 20% more at age 40—since they know that if they divorce, they will not remarry and must rely solely on their own resources. This effect is partly counteracted by the fact that low-wage couples, who are poorer, are now staying married. What we measure is the net effect of these two forces.

The bottom right panel shows that these changes in wealth holdings are accompanied by a large increase in labor income for all groups, except married men. Single women's labor income is about 28% higher around age 30 and remains elevated for much of their working life. Married women's labor income also rises–about 20% at age 26–as they anticipate needing financial independence if they divorce and can no longer remarry. As their probability of divorce decreases, their labor income gradually converges to that in the benchmark, around age 45, and falls slightly below it thereafter. This rise in married women's labor income helps finance couples' increased savings. Single men's labor income also increases. In contrast, married men experience a small decline in labor income, which becomes more pronounced after age 35. The large increases in wealth that we observe in the absence of new marriages turn out to be financed not only by higher earnings, but also by reduced consumption until age 40 for both single men and women, and couples. After making the sacrifice of working more (and thus accumulating more human capital) and consuming less when young, they enjoy more consumption through much of their remaining lives.

No divorce. When couples no longer divorce, almost all men and women in our cohort end up married by age 66 (see top left panel of figure I.4). The bottom right panel displays median wealth for couples and singles. In the absence of divorce, couples hold less wealth, partly because the pool of couples now includes individuals who would have divorced in the benchmark economy—-who tend to have lower wealth and human capital, and partly because they no longer need to save for when they divorce. Because single women expect to marry men with higher wages than theirs, the absence of divorce makes marriage an even better deal. Thus, they save much less when they are young and still have a high probability of marrying. Single men also expect to marry, but typically to women with lower wages than theirs, hence the effect on their wealth is smaller and actually positive early on, when their probability of marriage is highest. A notable pattern is the wealth jump at age 66 for singles. Without divorce, no newly divorced individuals–who tend to be wealthier–join the singles group. However, at age 66, some individuals lose their spouses and enter the singles pool with significantly higher wealth. This effect is particularly strong for single women, because new widows tend to be much wealthier than lifelong single

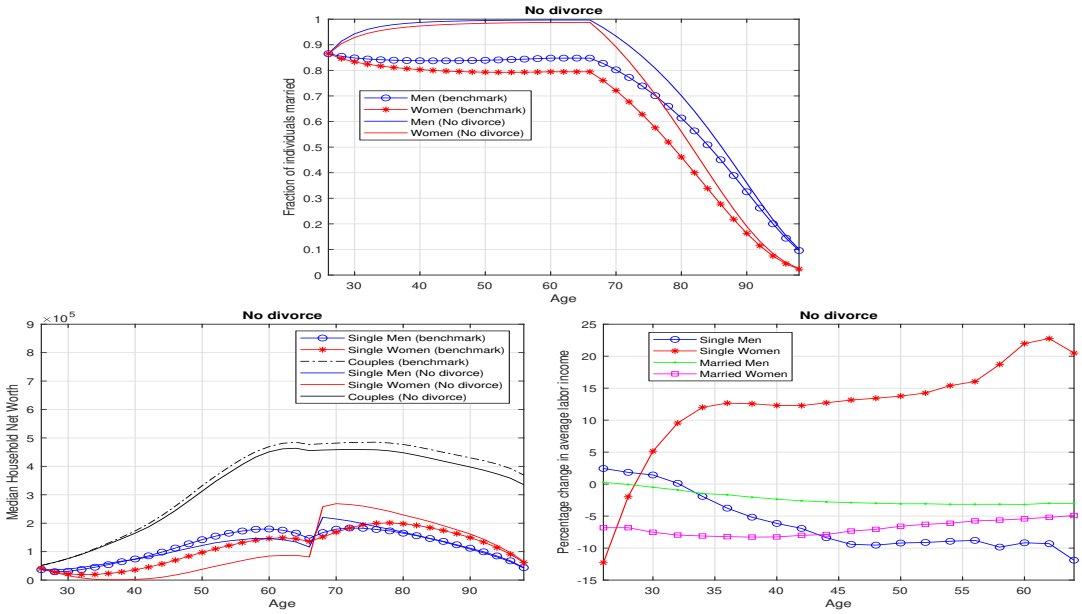


Figure I.4: The effects of not divorcing

women.

The bottom right panel of Figure I.4 shows that the effects of eliminating divorce on labor income are highly heterogeneous. Young single men maintain stable earnings until age 30 but then work less, as they are less likely to marry and support a lower-wage spouse. Consequently, their earnings decline by about 10% by retirement time. Young single women, who have a high probability of marrying and staying married–typically to higher-wage men–work less before age 30. However, as their likelihood of marriage declines after that age, they increase their labor supply to save for retirement. For married women, labor income drops by over 5% during all of the working period. For married men, we observe a steady decline, reaching over 3% by age 66. This decline in labor income during working years is also reflected in decreased consumption.

J Additional results on aggregate savings and earnings

Table J.1 shows that, as we have seen in I, **removing the desire to save to leave side bequests** reduces the wealth holding of couples (by 16.5% on average over the life cycle), but increases those of single men and women (by 15.9% and 14.4%, respectively). In the aggregate, couples are much richer than singles. Hence, overall,

Counterfactual	Couples	SM	SW	All
No utility from side bequests	-16.5	15.9	14.4	-9.6
No utility from terminal bequests	-6.6	-32.5	-28.5	-11.7
No medical expenses at death	-0.8	-0.6	0.4	-0.6
No marriage	-5.7	48.9	14.7	6.6
No divorce	-4.3	-4.2	21.7	-0.1

Table J.1: Changes in average wealth per person (in % from our benchmark economy). SM: single men, SW: single women

Counterfactual	MM	MW	SM	SW	All
No utility from side bequests	-0.8	-1.5	-0.6	-0.2	-0.9
No utility from terminal bequests	-0.1	-0.4	-0.6	-0.2	-0.2
No medical expenses at death	-0.0	-0.0	0.0	0.1	-0.0
No marriage	-2.2	-3.8	6.9	11.5	0.8
No divorce	-2.3	-7.1	-8.4	-0.1	-2.2

Table J.2: Changes in average labor income. MM: married men, MW: married women, SM: single men, SW: single women

the average aggregate wealth of our cohort drops by 9.6%. This saving motive also has non-trivial, but heterogeneous effects on average labor earnings which drop for all groups (table J.2). The largest drop happens for married women (1.5%), whose labor supply is more elastic, followed by married men (0.8%), single men (0.6%), and single women (0.2%). Overall, earnings drop by 0.9%.

Removing the desire to leave terminal bequests, instead, reduces the savings of all three of our groups, and by more for singles (32.5% for men and 28.5% for women) than for couples (6.6%). Overall, this change in saving motives reduces average aggregate wealth by 11.7%. It also reduces labor earnings for all groups (and here slightly more so for single men than for married women), resulting in a decrease in average labor earnings by 0.2%.

The next two lines of tables J.1 and J.2 turn to studying the role of medical expenses at death. As the jump in medical expenses in the period preceding death is large, **eliminating medical expenses at death** has non-trivial negative effects on wealth: declines of 0.8% for couples, 0.6% for single men, and an increase in 0.4% for single women. The aggregate effect on average wealth is a decline of 0.9%. Table J.2 shows that they, instead, have little effect on labor earnings.

removing the possibility of new marriages has large and heterogeneous effects on wealth. Its aggregate implications are that, without new marriages, wealth would be 6.6% lower and earnings 0.8% lower. **Without divorce**, average wealth

would be 0.1% lower and average earnings 2.2% lower. In terms of labor income, divorce risk alone accounts for 2.2% of average earnings over the life cycle.