End-of-Life Liquidity*

Neha Bairoliya[†]
Neha.Bairoliya@marshall.usc.edu

Giovanni Gallipoli[‡]
giovanni.gallipoli@ubc.ca

Kathleen McKiernan[§]
kathleen.mckiernan@vanderbilt.edu

October 14, 2024

Abstract

Uncertainty about one's lifespan induces a preference for end-of-life liquidity (Yaari, 1965). Such preference, which can be characterized as a warm-glow motive but need not be interpreted that way, interacts with institutional constraints to shape life-cycle behaviors. Using data on wills, life insurance, and designated beneficiaries we document the presence of end-of-life motives across different demographic groups, regardless of marital status and children. Then, we illustrate their quantitative importance in a model of consumption, labor supply, and retirement decisions. Our analysis highlights the life-cycle distortions that annuity plans, such as the U.S. Social Security, impose on life-cycle decisions due to the illiquid nature of entitlements. A marginal reduction in retirement annuities, in exchange for a guaranteed amount upon death, induces large effects on life-cycle allocations and raises welfare, especially among non-married individuals with low education. These individuals are often ignored in the debate on end-of-life provisions.

Keywords: liquidity, consumption, life-cycle, labor supply, self-insurance, annuity, health, marriage, social security, entitlements, survivor benefits, retirement.

JEL Classification Numbers: J01, J26, E21, H55

*We are grateful to several seminar and workshop participants. In particular, we thank Andy Abel, V. V. Chari, Mariacristina De Nardi, Jonathan Heathcote, Selahattin Imrohoroglu, Pat Kehoe, and Jane Ryngaert for their comments and suggestions.

[†]University of Southern California, Department of Finance and Business Economics, Los Angeles, CA 90089-1422, United States.

[‡]University of British Columbia, Vancouver School of Economics, Vancouver, BC V6T 1L4, Canada.

[§]Vanderbilt University, Department of Economics, Nashville, TN 37235, United States.

1 Introduction

Consumption plans are made under uncertainty about the length of one's life so that consumers may value liquidity at the time of death. Such end-of-life motive is not unlike a warm-glow; however, it need not be interpreted that way as it generally arises from the objectives and constraints of expected utility maximization over an uncertain lifespan and may apply to individuals with no spouses or dependants (Yaari, 1965). Several studies have examined the empirical content of end-of-life motives and have shown that they account for a variety of life-cycle observations (Dynan et al., 2002; French et al., 2006; De Nardi et al., 2016; Jones et al., 2020; De Nardi et al., 2021; Lockwood, 2012, 2018).

In this study, we examine how end-of-life motives interact with institutional arrangements imposed by the U.S. Social Security system, and result in life-cycle distortions. We do so in an environment where all households (couples and singles) are allowed to have an end-of-life liquidity preference that is consistent with their life-cycle behavior. The role of Social Security in providing annuitized consumption to poorer households has been studied before in settings where households have no concerns over others (Abel, 1986; Hubbard and Judd, 1987; İmrohoroglu et al., 1995; Conesa and Krueger, 1999) or where such concerns apply only to multiperson households (Hong and Ríos-Rull, 2007; Hong and Rios-Rull, 2012).

To motivate our work, we begin by documenting that end-of-life motives are prevalent across all demographic groups, including less educated, non-married individuals with no children. Using information about wills, designated beneficiaries, and life insurance, we show evidence of behaviors consistent with pervasive preferences for late-life liquidity. The evidence also suggests that these preferences can be partly, but not completely, satisfied through private life insurance contracts; consistent with previous work, we use the model to show that poorer and more constrained households may be unable to de-annuitize their lifetime wealth through private life insurance contracts.

By allowing for end-of-life motives across the whole range of households, our analysis brings to the fore severe distortions that result from the lack of end-of-life provisions in the U.S. Social Security system. These costs accrue mostly to unmarried individuals with little wealth. We estimate how much different individuals value access to liquidity at the end of life and, in counterfactual experiments, we illustrate that the impact of the poor design of the S.S. system can be mitigated through a minor reform that marginally reduces the S.S. annuity value in exchange for a small end-of-life payment.

The interest in end-of-life motives is related to a lively debate about the value of annuitizing lifetime wealth. To understand the nature of this debate it is useful to follow Yaari's original arguments and define an "actuarial note" that can be bought or sold by a consumer. Such a contract

remains enforceable as long as the buyer is alive and becomes null and void thereafter. The buyer of the contract acquires a promise to receive a regular allowance while alive (that is, an annuity). The allowance is determined by a rate of interest higher than the market interest on one-period debt contracts (the higher interest compensates the consumer for the possibility they will not survive to the next period, at which point the insurance company that sells the annuity is held free of any obligation). A consumer can go short and sell the actuarial note: in this case, the consumer acquires life insurance and commits to paying a rate of interest higher than the market as long as they live, in exchange for a lump-sum payment that becomes part of their estate when they pass away (the lump-sum payment is left behind with the rest of their wealth upon death, with no further obligations on anyone).

Having defined the actuarial contract, it is straightforward to show that all wealth should be annuitized unless preferences feature an end-of-life motive. If an end-of-life motive is present, the life-cycle consumption problem entails a portfolio choice where the agent has to decide how much to consume and how much to leave behind upon death (as Yaari puts it, the consumer can use the actuarial note to separate the consumption decision from the end-of-life decision).

This preamble helps clarify why the literature on household consumption is preoccupied with the observation that wealth is rarely annuitized while many people hold life insurance (LI). This stylized fact holds true also among single individuals with no dependants, which is hard to reconcile with explanations based exclusively on within-family linkages and bequest motives. Theoretical work (e.g. Davidoff et al., 2005) shows that less than full annuitization may be optimal when markets are incomplete even in the absence of end-of-life liquidity motive. However, in this case, large shares of wealth should still be annuitized. In reality, over half of unattached individuals, including many with no offspring, hold sizable amounts of LI and little annuitized wealth: the lack of voluntary annuitization among people with different marital arrangements and fertility suggests that end-of-life motives may be more prevalent than previously thought.

Given these observations, a body of work has developed computational models of consumption behaviors under different assumptions about family structure and institutional constraints. For example, Hong and Ríos-Rull (2007) examine the relative value of annuities and life insurance within an overlapping generations incomplete market setting. Allowing for multi-person households and family linkages, they show that the absence of life insurance contracts has significant welfare costs whereas access to annuities has little or no effect. This explains why the value of social security as a surrogate for missing markets is small: its focus on annuities limits its value.

We document that a majority of people hold life insurance (LI), including singles with no dependants. Existing research suggests that LI holdings are the outcome of voluntary choices (Hong and Rios-Rull, 2012) and would be even larger in the absence of financial constraints (Chambers et al., 2011).

The literature highlights that a mismatch exists between access to life insurance and the financial vulnerabilities of households (Auerbach and Kotlikoff, 1985, 1991; Bernheim et al., 2003). Social Security does provide a modicum of life insurance to married households through survivor benefits. However, since married households tend to be richer and better insured, the current S.S. system does a poor job of helping the most financially vulnerable (for example, singles with low assets), who are the very people it was designed to support. However, a small, incremental reform can unlock liquidity at the end of life: specifically, we examine a policy that allows all individuals to claim the equivalent of roughly five years' worth of social security annuity benefits upon death; this comes at the cost of lower per-year benefits (a 4 percent drop in yearly payments) so that aggregate government outlays are constant. For singles, the policy provides access to a guaranteed minimum claim in exchange for a marginally lower annuity. Since the end-of-life entitlement is valid upon death even if it occurs early, married individuals value its flexibility as it can be claimed at any time.

Our analysis yields several insights. (1) Poor single households benefit the most from the introduction of end-of-life entitlements; this comes with a large change in their life-cycle labor supply behavior. (2) The small policy departure equalizes the value of the program across participants. Under the current US Social Security system, FICA taxes are levied on all workers and used to finance survivor benefits, which are available only to married individuals. Unmarried workers pay for benefits they will never receive unless they enter formal marital relationships. To the extent that all individuals have some end-of-life motive, the reform guarantees a more even treatment of beneficiaries. (3) The life-cycle behavior of married households also responds strongly to the reform but they experience small welfare changes. More generally, the labor supply of all agents is sufficiently strong to pay for the reform in a revenue-neutral setting. (4) The consumption equivalent gains of the reform are maximized when the end-of-life entitlement is between five and six years' worth of the annuity; anything larger still results in positive but less sizeable gains. (5) A reform that only applies to single households (that is, an even smaller departure from the *status quo* with no changes for the married) can still generate welfare gains and large behavioral changes among the targeted households.

Our analysis delivers three main insights. First, we show that end-of-life motives induce different responses in the life-cycle labor supply and consumption behaviors of households. The needs for liquidity extend up to the time of death but shape behaviors much earlier in the life cycle (Dynan et al., 2002; French et al., 2006; Kopczuk and Lupton, 2007; De Nardi et al., 2016; Jones et al., 2020; De Nardi et al., 2021). Second, we document the significant welfare implications of marginal reforms to large-scale programs (İmrohoroglu et al., 1995; Conesa and Krueger, 1999; De Nardi et al., 1999; Benabou, 2000; Fuster et al., 2003; Krueger and Kubler, 2006; Hong and Ríos-Rull, 2007; Kotlikoff et al., 2007; Attanasio et al., 2007; Krusell et al., 2010; Huggett

and Parra, 2010; Scheuer and Slemrod, 2021). In particular, we emphasize how the interaction of end-of-life motives and social security incentives results in persistent labor supply distortions and welfare losses. Third, by examining a policy that redistributes resources from the married to the singles, we clarify the incentives built into the existing system and advance the understanding of the institutional drivers of observed gaps in the earnings and labor supply of diverse subsets of the population (Krueger and Perri, 2006; Heathcote et al., 2010b,a; Golosov et al., 2016; Stantcheva, 2017; Daruich and Fernández, 2020; Miller and Bairoliya, 2021).

2 Motivation

To motivate our quantitative analysis, we document three instructive facts. First, we illustrate how households engage in activities consistent with end-of-life motives. Irrespective of marital status and the presence of children, most households exhibit both intent to bequeath wealth and readiness to hold significant liquidity up to the end through will preparation.[1] Second, we document that many individuals hold life insurance, consistent with the presence of end-of-life motives across various familial structures. The coverage of life insurance, however, varies across these marital status and parental groups. Finally, we show that the structure and timing of Social Security entitlements imply large penalties for many households that, despite contributing to the system, are shut out from its end-of-life provisions.

1. End-of-life motives and family characteristics: facts about intent and preparedness.

We use data from the Health and Retirement Study (HRS) to establish facts related to intent and preparation to bequeath wealth. Our HRS sample covers years 1996-2020 (Waves 3-15) and consists of 85,764 men over the age of 50 born between 1926 and 1970 for whom marital and parental statuses are known.[2]

Intent to leave assets by wealth quantile, marital status, and parental status. To gauge and compare the intent to leave behind assets, we use an HRS question asking workers about the likelihood they will leave bequests of various sizes.[3] Workers are first asked, "What is the

[1] Marital status is defined based upon the current relationship. This means that some individuals that we classify as single may have been previously married. Among singles without children, roughly 75 percent have never been married; among singles with children over 80 percent of men are divorced (65.7 percent) or widowed (15.2 percent). Children are defined as having any living child with whom the individual is still in contact. The full distribution of detailed marital statuses is shown in Figure A1

[2] Questions about will preparation are available beginning in Wave 6 (2002). Details about sample selection are in Appendix A1

[3] In Wave 3, the order of questions is different. Workers are asked the probability of leaving any bequest before being asked about a bequest of $10,000 or more.

probability that you will leave a bequest of $10,000 or more?" Follow-up questions are determined based on the response. If the individual reports a probability strictly above 0, he or she is asked "What is the probability that you will leave a bequest of $100,000 or more?"[4] If an individual, on the other hand, reports a 0 probability of leaving a bequest of at least $10,000, that individual is asked for the probability of leaving any bequest.

We use these questions to construct an indicator of whether or not an individual has an intention to leave a bequest. The indicator takes a value of 1 if the worker, between ages 50 and 70, reports at least a 95 percent chance of leaving a bequest. This design addresses three issues. First, we limit the sample to workers between ages 50 and 70; this helps us focus on bequest intentions rather than higher probabilities due to reaching a very old age with more wealth than expected.[5] Second, we set a conservative 95 percent threshold so that we get as close as possible to measuring the deliberate intention to leave a bequest rather than a future chance based only on how things turn out. Finally, we focus on the likelihood of leaving any bequest rather than leaving an exact monetary amount as we want to measure the general intention to leave some wealth behind. Therefore, a decline in the probability measure can be interpreted as a weaker intent to leave a bequest rather than a lower probability of ending up alive in old age with unplanned excess wealth above a specific amount.

We compute this likelihood measure for various marital and parental status types after conditioning on a given quintile of wealth. Our findings, in Figure 1, highlight that wealth is an important indicator of the intent to leave bequests. The influence of wealth is larger than marital status or the presence of children and, once we condition on the quintile of the wealth distribution, there is little or no difference across marital status and children in the share of workers who intend to leave a bequest.

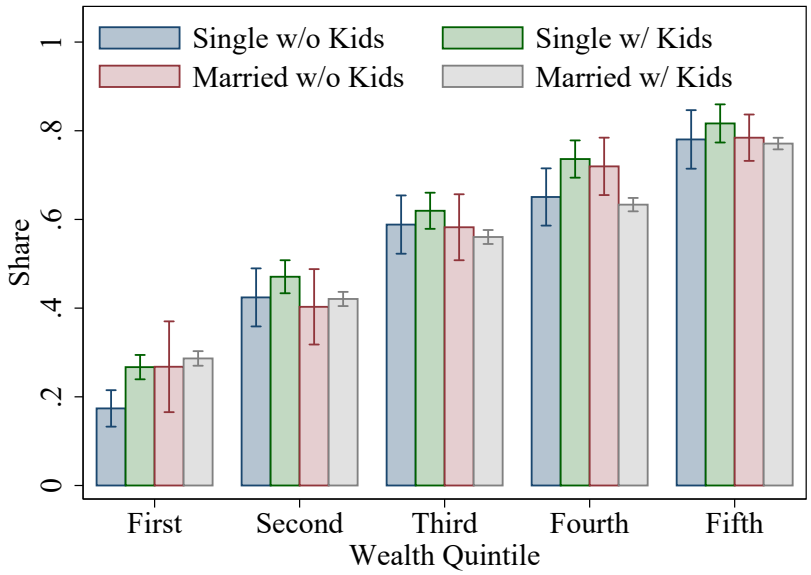
Similarities in preparing a valid will and variation in beneficiaries across marital and parental status. The intent to leave a bequest indicates the workers may have a subjective desire. It does not, however, demonstrate that these individuals engage in concrete preparations to leave behind wealth. To establish that such preparation does occur, we document patterns related to a key activity of estate planning: making a legally binding will. Figure 2 plots the share of men, by marital-parental status and wealth quintile, who make such provisions as well as who the designate as beneficiaries of this will.

Figure 2a plots the share of individuals who have a legally binding will to dispose of assets upon death. The share increases with wealth but, after controlling for wealth, there are no significant differences by marital status and children. In the fifth quintile of the wealth distribution,

[4]Beginning in Wave 6, those who report a positive probability of leaving a bequest of $100,000 or more are asked the probability they will leave a bequest of $500,000 or more. We will not use this question in our analysis.

[5]This restriction limits our sample to 60,803 observations between the ages of 50 and 70.

Figure 1: Bequest Intentions by Family Structure and Wealth



Source: Health and Retirement Study. authors' calculations

Notes: Figure shows the share of workers age 50-70 of a given marital-parental status and wealth quintile who report over a 95 percent probability of leaving a bequest of any amount. 95 percent confidence intervals shown. Results of the Tukey-Kramer pairwise mean comparison test are shown in Table A2.

around 70 percent of individuals, irrespective of marital and parental status, have a written will to determine how their assets will be distributed. In the median wealth quintile, the share is about 55 percent while in the lowest wealth quintile, it is around 25 percent. Tests of pairwise differences (Tukey-Kramer tests) cannot reject the null hypothesis that, in each wealth group, similar shares of individuals make a valid will whether they are married or have children (see Table A2).

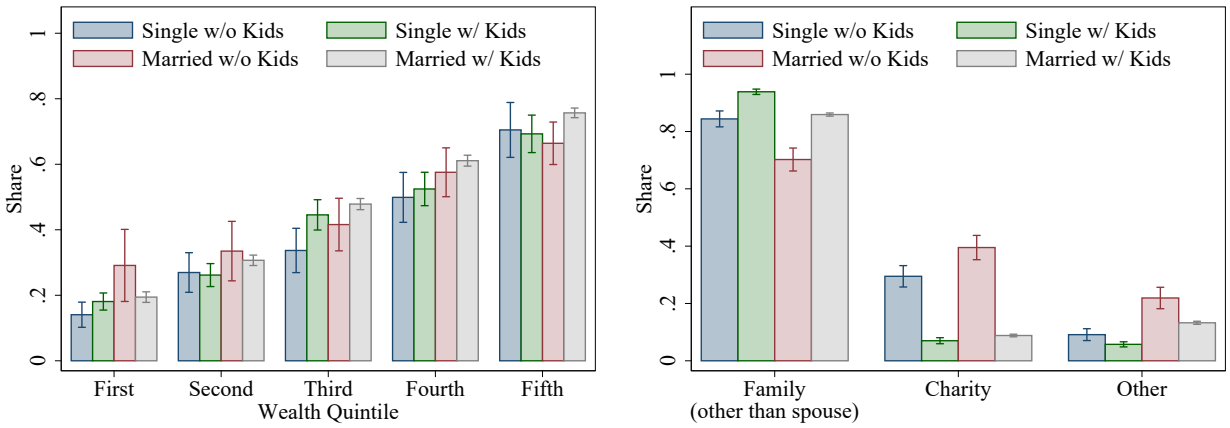
Given this intent and preparation to leave resources behind, to whom do individuals plan to leave them? First, many individuals make provisions in their wills for immediate family, such as children, grandchildren, parents, and siblings.[6] The tendency to make provisions for the family is slightly more pronounced among those with children. Around 90 percent of men with children make provisions for family, as opposed to 60-80 percent of those without children. Second, leaving money to charity is a recurrent consideration, especially for those without children. Between 30 and 40 percent of individuals without children make provisions for charity in their wills, in contrast with less than 10 percent of those with children.

[6] An individual's spouse is the default recipient of wealth if the individual dies without a will. A valid will is only needed to deviate from the legal requirements of probate.

Figure 2: Will Coverage and Beneficiaries by Family Structure and Wealth

(a) Valid Will

(b) Beneficiaries



Source: Health and Retirement Study. authors' calculations

Notes: Panel (a) shows the share of individuals age 50+ of a given marital-parental status and wealth quintile who say they have a written and witnessed will or trust. 95 percent confidence intervals shown. Results of the Tukey-Kramer pairwise mean comparison test are shown in Table A2. Panel (b) shows the share of individuals age 50+ of a given marital-parental status and who have a valid will who report having made provisions for a given group in said will. 95 percent confidence intervals shown.

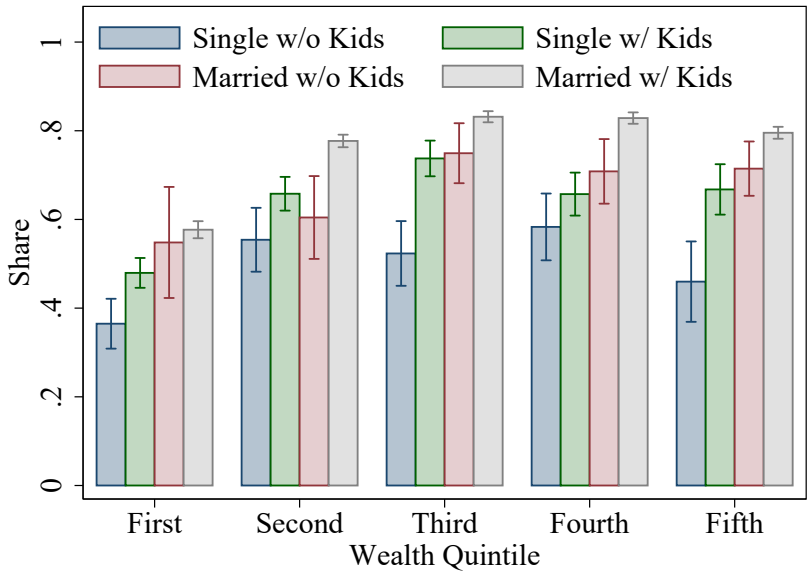
2. Life insurance markets provide a way to satisfy these end-of-life motives: differences across familial groups

Private life insurance markets allow individuals to purchase contracts that pay out to a designated beneficiary in the event of the death of the policyholder. Once again using data from the HRS, we document the prevalence of these contracts–which can transfer resources similarly to a bequest–across individuals with heterogeneous family structures.[7] The share of individuals who hold a life insurance policy by marital and parental status as well as wealth quintile is shown in Figure 3.

Consistent with the analysis above in which we show that intention and preparation to leave wealth extend beyond those in marital relationships and those with children, a significant share of men hold at least one life insurance policy. The average coverage within the population is nearly 70 percent. This share ranges from roughly 40 percent among the poorest singles with no children to 80 percent for wealthy married individuals with children, and, unlike results related to bequest intentions and will preparations, pairwise mean comparison tests are often able to reject the null hypothesis that these shares are the same across groups. Notably, among those in the third wealth quintile, the life insurance coverage among married individuals with children is 27 percentage

[7]Questions about life insurance start in Wave 6. Therefore, the sample for these calculations covers years 2002-2020.

Figure 3: Life Insurance Coverage by Family Structure and Wealth



Source: Health and Retirement Study. authors' calculations

Notes: Figure shows the share of individuals age 50+ of a given marital-parental status and wealth quintile who hold a life insurance policy. 95 percent confidence intervals shown. Results of the Tukey-Kramer pairwise mean comparison test are shown in Table A2.

points higher than coverage for singles without children (see Table A2 for this and other results).

The variation in coverage rates across groups may be due to differences in preferences, market prices, or other constraints. We revisit this issue in the context of our model.

3. The uneven costs and benefits of Social Security: despite similar contributions, some households get much more value out of the S.S. system than others.

The previous facts illustrate that the intent and readiness to leave behind resources are widespread, and these intention do not apply exclusively to people with children or spouses. The United States Social Security system, however, treats married individuals differently by providing benefits that extend beyond the death of the beneficiary.

We focus on two types of benefits paid by the US Social Security System: (i) *Old-Age Benefits* which provide annuity income to retirees from the time of claiming until death; and (ii) *Survivors Benefits* which continue to pay a deceased beneficiary's benefits to their spouse after their death.

Old-Age Benefits are an increasing and concave function of the average indexed monthly earnings. Up to a maximum taxable amount, higher income in working life translates to higher benefits in retirement. The progressive nature of the system implies that high-income individuals receive lower replacement rates on their earnings. Spouses of primary earners can claim spousal benefits

on the earnings record of the primary earner.[8] These benefits are up to 50 percent of the benefits of the primary earner and are contingent on the primary earner having claimed S.S. benefits. Individuals first become eligible for benefits at the age of 62 and become eligible for full benefits at the normal retirement age (NRA). Claiming Social Security benefits before the NRA entails lower pension payments for a longer period of time. Delaying pension claims until beyond the NRA (up until age 70) entitles workers to larger payments, albeit for a shorter period of time. These penalties/credits for early/delayed claiming also apply to spousal benefits. Spouses who claim prior to the NRA incur a penalty while spouses who delay their claims up to age 70 receive a credit.

Widows and widowers of eligible workers[9] can receive *Survivors Benefits* of 77.5 percent of the beneficiary's basic SS benefits as early as age 60 or full benefits at the normal retirement age.[10][11] They will continue to receive these benefits until their own death. Moreover, spouses who were living with the deceased are eligible for a small, one-time lump-sum payment in the month when the death occurs.[12]

Figure 4 shows how benefits received and contribution made vary by income and marital status. Within each marital group, both the contributions and benefits increase with income. As expected, higher-income individuals pay more into the system and receive more in benefits than lower-income ones. Interesting differences emerge when we compare households with different marital statuses after conditioning on income. Even when people make similar contributions, benefits vary with marital status because spousal and survivors benefits are only available to married couples. By design, the S.S. system gives significantly higher benefits to married couples than to singles with similar incomes and contributions.

Differences in the effective entitlements partly depend on the timing of the benefits. While spousal benefits are received during the life of the beneficiary, survivors benefits are paid out to a spouse after the death of the beneficiary. These payments, represented by the white portion of the benefits bar for married couples in the graph, account for roughly 30 percent of the gap between the benefits of singles and married. Despite similar contributions to the S.S. system, these entitlements are available only to individuals in a legal marital relationship and there are no provisions to direct them to other potential beneficiaries even if an individual (single or married) may wish to do so.

[8] Spouses elect whether to claim benefits on their own earnings record or that of the primary earner.

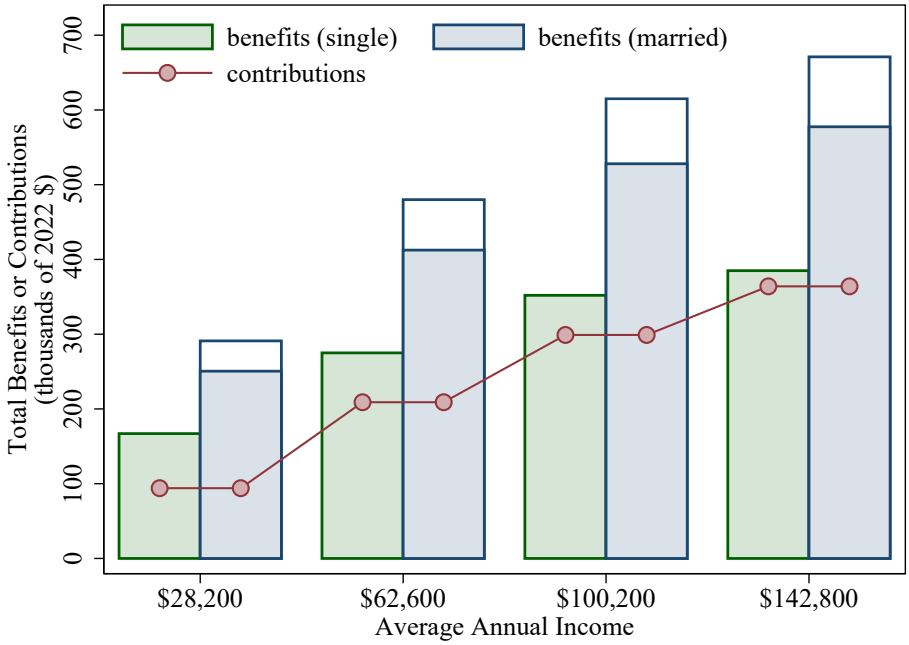
[9] Whether a worker's record is eligible to pay out survivors benefits depends on the work history of the deceased beneficiary: the number of credits needed for survivor benefit eligibility varies by age with younger workers requiring fewer credits. No worker is required to have more than 40 credits–roughly 10 years of work.

[10] While widows and widowers are the largest group claiming SS Survivors Benefits, children and dependent parents of the deceased beneficiary are also eligible for benefits. Divorced spouses are eligible as long as their marriage to be beneficiary lasted at least ten years, and they do not remarry before the age of 60.

[11] There are ways to "game" the system here. For example, a surviving spouse may claim these scaled-down benefits at age 60 and switch to benefits on their own record later.

[12] This payment is $255. If there is no living spouse, this benefit is given to children who are eligible for benefits of the deceased worker's record.

Figure 4: Social Security Contributions and Benefits by Marital Status



Source: Steuerle and Smith (2023) Tables 1-4, 9-12

Notes: Figure shows variation in total SS contributions and benefits by average annual income and marital status. Green bars represent the old-age benefits of a single individual of a given income while the blue bars represent the old-age, spousal, and survivors benefits of a single earner, married household of a given income. Within the blue bars for married households, the shaded area is old-age + spousal benefits (50% of old-age benefits) while the white portion is survivors benefits. Red line shows the total contributions paid for each income level. Benefits and taxes are for a cohort who turned 65 in 1995.

3 An Analytical Example

The discussion above suggests that single individuals exhibit end-of-life motives similar to married ones but participate less in the private life insurance market and are treated differently by the Social Security system. Is this evidence that the system should be changed to accommodate those motives? One reason for caution is that tweaking the system might not be necessary if individuals can easily acquire private life insurance (LI) contracts and household already optimally choose LI holdings that meet their end-of-life needs. In that case, private provision of LI would be sufficient to de-annuatize S.S. entitlements.

To emphasize the interaction between end-of-life motives and savings behavior, we consider a two-period economy populated by a (unit) continuum of agents who differ in their survival probability π. Survival probabilities are distributed over the interval (π_0, π_1), where $\pi_0 < \pi_1$. An agent's type and the distribution of agents on the interval (π_0, π_1) are public knowledge and everyone observes the realization of the random survival process. Agents receive an endowment e

in period 1 and derive utility from consumption in each of the two periods of life (c_1, c_2), as well as from leaving behind unconsumed wealth (b) after death, an end-of-life motive. Two contracts are available to transfer resources inter-temporally: (i) agents can buy a risk-free bond with gross return A, and (ii) agents can acquire life insurance, which returns $\frac{A}{1-\pi}$ upon death. We model life insurance contracts because they are a common way for households to de-annuitize their wealth. The majority of the adult U.S. population holds some form of LI contract; this is in stark contrast to the scant popularity of private annuity contracts. The return on a LI contract depends on the survival probability of the buyer, and insurance pools are type-specific. Agents can buy but not sell bonds and LI contracts (zero lower bounds).

The government administers a mandatory annuity (Social Security) that pays a fixed transfer T to all agents and is financed through a flat tax τ levied on period 1 endowments. Social Security receipts are transferred to period 2 with a gross return A, the same as bonds. The life cycle utility of an agent is

$$u(c_1) + \pi\beta U(c_2) + (1-\pi)\beta v(b)$$

where $u(c_1)$ and $U(c_2)$ are period-specific utilities from consumption, and $v(b)$ denotes the value of end-of-life liquidity. The utility functions are concave, continuous, twice differentiable, and satisfy the Inada conditions. Agents take taxes, transfers, and returns on bonds and LI as given. The problem of an agent of type π is to choose a consumption sequence $\{c_1, c_2\}$ and bequest b to maximize lifetime utility. That is,

$$\max_{\{c_1,c_2,b,l,s\}} u(c_1)+\pi\beta U(c_2) + (1-\pi)\beta v(b)$$

$$s.t.$$

$$c_1 + \frac{1-\pi}{A} LI + s \leq e(1-\tau)$$

$$c_2 \leq As + T$$

$$b \leq As + LI$$

$$c_1, c_2, b, s, LI \geq 0$$

where LI is the amount of life insurance and s is the savings held in the risk-free bond. If Consumers face a fair price, the solution to this problem (see Hosseini, 2008) implies that all households would buy some life insurance as long as the S.S. tax τ is sufficiently high: we maintain this assumption, given the empirical observation that most U.S. households hold some life insurance. The interpretation is straightforward: given high social security benefits and the lower marginal utility of consumption, agents choose higher life insurance to increase the bequest and reduce the marginal utility of the end-of-life motive.

Consumption and bequests. The solution is a function of the survival type π of the buyer. If a type π agent holds LI (i.e. $LI(\pi) > 0$), the optimality condition is:

$$\frac{u'(c_1(\pi))}{A\beta} = U'(c_2(\pi)) = v'(b(\pi))$$

Optimality implies that agents use bonds to transfer resources to the second period so that $s(\pi) > 0$. Bonds are fungible and can finance b even though their conditional return is lower in the case of early death. Agents would consider the corner solution $s = 0$ only if government annuities were overly generous (carry a higher return than bonds), a case which is not empirically relevant.

Lower survival probabilities should be associated with higher LI coverage since agents benefit from the higher conditional returns of LI relative to bonds. Given actuarially fair prices, agents with lower π would hold relatively more LI if supplied by the market. However, this implication of the model does not align with the empirical observation in Figure 3, which shows that both the frequency and value of LI policies are higher in demographic groups with higher π. In fact, not only do low π individuals exhibit lower LI coverage but, conditional on holding insurance, have much smaller policies. This observation holds regardless of wealth: in Table A2 we test pairwise differences in coverage and values conditional on wealth for demographic groups with different survival rates and find that higher π is associated with more LI. This finding, in conjunction with the analytical results above, suggests that the market might not supply low π individuals with fairly priced LI contracts. In our quantitative analysis, we verify this conjecture by explicitly estimating the price of different LI policies (by survival risk and level of coverage).

If some LI policies do not fairly reflect the observable π, life-cycle optimality implies a corner solution:

$$\frac{u'(c_1(\pi))}{A\beta} = U'(c_2(\pi)) > v'(b(\pi)$$

The inequality above would, for example, hold if private markets offer contracts priced in a way that leads low π agents to use bonds to smooth marginal utilities. That is, the demand for LI by agents with low π might be constrained by prices that reflect risks above and beyond what is captured in their mortality rates. This phenomenon, discussed in Chambers et al. (2011), suggests that non-actuarially fair LI prices constrain the portfolio choice of low π agents. Our quantitative analysis indicates that a combination of over-annuitization through Social Security, and LI under-coverage through incomplete private provision results in large welfare losses for many individuals, especially the non-married and the less educated.

4 Quantitative Framework

Because most households exhibit behavior consistent with end-of-life motives, the design of Social Security and the potential incompleteness of private insurance markets may induce significant distortions in life-cycle choices. To explore the magnitude of these impacts and their welfare costs, we develop and estimate a dynamic model of labor supply, life insurance, and retirement decisions. Our analysis emphasizes life-cycle incentives under the current policy regime with a focus on the differential treatment of single and married individuals. To this purpose we carefully model retirement benefits from Social Security so that the model matches key life-cycle outcomes.

The life-cycle spans ages $t = 25, 26, ..., 99$. Individuals differ in their permanent education type (e) and marital status (q). Marital circumstances are summarized by a pair $q = (m, \iota)$ where m indicates if the agent is single or married, and ι denotes the age gap between spouses. Non-permanent (evolving) heterogeneity includes stochastic labor productivity (η_t), employment status (λ_t), health status (μ_t), assets (a_t), social security wealth (a_t^{ss}) and SS benefits' application status (b_{t-1}^{ss}).

Agents choose labor supply (h_t), consumption (c_t) and savings (a_{t+1}), We explicitly model the SS benefit application (b_t^{ss}) of couples and singles and allow them to purchase one-period LI contracts (i_t^v) with face value $v_t^{e,m}$.[13] Individuals make decisions in each time period t and adjust their behavior in response to changes in wages, employment, health, and survival. Given a vector of states $\left(e, q, \eta_t, \lambda_t, \mu_t, a_t, a_t^{ss}, b_{t-1}^{ss}\right)$, they maximize the present discounted value of lifetime utility.

4.1 Preferences

Agents derive utility in period t from consumption c_t and leisure l_t. The within-period utility is non-separable in these arguments. We account for the decline in expenditures at retirement through a combination of (1) unexpected health shocks causing unplanned retirement, and (2) consumption-leisure complementarities in utility (French and Jones, 2011; French, 2005; Casanova, 2010, see, for example,). The period utility is

$$U(c_t, l_t) = \frac{1}{1-\rho}\left(\left(\frac{c_t}{\zeta_t}\right)^{\nu} l_t^{1-\nu}\right)^{1-\rho}.$$

Relative risk aversion is dictated by the parameter ρ, while ν is the weight on consumption and ζ_t is a consumption-equivalence scale. The utility of married households is multiplied by two to

[13] Social Security application is a one-time decision and cannot be reversed.

account for spousal utility from consumption and leisure. The leisure enjoyed in period t is

$$l_t = \bar{l} - h_t - \phi_P(t)\mathbb{I}\{h_t > 0\} - \phi_H(\mu_t, t), \tag{1}$$

where $\bar{l}$ is the endowment of leisure in each period, h_t is hours worked, the function ϕ_H represents leisure lost to bad health and ϕ_P is the cost of employment participation (positive if $h_t > 0$). We set the time cost of poor health using estimates in Jones and Li (2023) and assume the following functional form for the time costs of working:[14]

$$\phi_t = \frac{\exp(\phi_0 + \phi_1 t + \phi_2 t^2)}{1 + \exp(\phi_0 + \phi_1 t + \phi_2 t^2)} \tag{2}$$

At the end of life, an individual values wealth A_t^q (see De Nardi, 2004):

$$\Omega(A_t^q) = \frac{\theta}{1-\rho}\left(A_t^q + \kappa\right)^{(1-\rho)\nu}. \tag{3}$$

End-of-life wealth A_t^q includes all assets a_t, the face value of any purchased life insurance, $i_t^v \times v_t^{e,m}$, plus Social Security survivors benefits, if eligible. Eligibility for survivors benefits depends on marital status and the age gap between spouses, q (we discuss details in Section 4.5.1). The coefficient θ measures the intensity of the end-of-life motive, while κ determines the curvature of the function. Higher θ increases the marginal utility of an extra unit of end-of-life wealth and higher κ makes them more similar to a luxury good.

4.2 Health, Medical Expenditures and Mortality

Each period, individuals are subject to an exogenous health process (μ_t) that affects their medical expenditures (ρ_{mt}), survival probability (π_t^s) as well as their time endowment. The transition rule across health states depends on current health, education, and age. Medical expenditures for determined by current health state, age, education and marital status.[15] Individuals are also subject to mortality shocks. The probability of survival depends on age, education, marital status and current health.

[14]The best health state in Jones and Li (2023) corresponds to our first two health states; our worst state maps into their fair/poor group. Health costs vary by education. The age-education time cost of poor health for our worst health group, at age 25, is roughly 15% of the time endowment for non-college graduates and 40% for college graduates.

[15]We follow Borella et al. (2019) in modeling and estimating these medical expenditures which climb quickly toward the end of life. In robustness section 7.1 we allow for persistence and uncertainty in these expenditures beyond that is captured by health status. We find that our benchmark welfare results remain robust to this alternate specification.

4.3 Employment and Wages

Unemployment shocks are an important driver of claiming and retirement behavior among older Americans (Bairoliya and McKiernan, 2021). Individuals in the model experience unemployment shocks with probability π^{λ}. Unemployment implies lower productivity and wage-scarring effects.

$$\pi^{\lambda}_{t+1} = prob(\lambda_{t+1} = 1)$$

Hourly wages follow a deterministic education-specific age profile $\omega(e,t)$ and depend on two stochastic components: employment status (λ_t) and an auto-regressive component η_t.[16]

$$\begin{aligned} w_t &= \xi(\lambda_t)\exp(\omega(e,t)+\eta_t) \\ \eta_t &= \rho^w \eta_{t-1} + \epsilon^w_t \\ \epsilon^w_t &\sim N(0,\sigma^2_{\epsilon^w}) \end{aligned} \tag{4}$$

Upon realization of a shock $\lambda_t = 1$, an individual can immediately re-enter the labor market with a wage penalty, ξ. This captures the short average duration of unemployment spells.

4.4 Life Insurance

Each period, households can purchase a one-period life insurance policy for a premium $\rho\,(v^{e,m}_t)$. The life insurance policy pays out the face value $v^{e,m}_t$ in the event of death in the next period. We allow for individuals to choose from two levels of insurance (high and low): $v^{e,m}_t \in \{\underline{v}^{e,m}_t, \overline{v}^{e,m}_t\}$. These levels of life insurance coverage are permitted to vary by age, t, education, e, and marital status, m. Note that we do not model the life insurance markets or take a stance on the determination of prices by the insurers: actuarially fair or not. Rather, we infer these prices indirectly by the observed coverage and value of insurance held by individuals in the data. For this reason, these premiums are not to be interpreted as strictly those a life insurance provider would offer an individual but rather the shadow price of life insurance coverage which may reflect things such as access differences across these groups. We detail our estimation strategy for prices in Section 5.

4.5 Social Security

The Social Security system in the U.S. provides retirement incentives at the time when these benefits become available. Benefits are computed in several steps. First, the earnings of the 35

[16] This specification delivers realistic wage scarring effects after unemployment spells.

highest earning years are averaged into an index – Average Indexed Monthly Earnings (AIME). The AIME increases by working an additional year if earnings in that year are higher than the lowest earnings embedded in it and are also capped at a threshold. We let a_t^{ss} be the Social Security wealth (an annualized measure of AIME). Then, the Social Security wealth evolution is approximated in the model by the following rule:

$$a_{t+1}^{ss} = \max\{[a_t^{ss} + \max\{0, (w_t h_t - a_t^{ss})/35\}], a^{\max}\} \tag{5}$$

In Equation (5), $a^{\max}$ is the threshold at which the Social Security wealth is capped and $w_t h_t$ denotes annual earnings in period t. In (5) we assume that the high earnings year replaces an average earnings year. Modeling the actual system would require keeping track of the entire earnings history which is computationally infeasible. In the second step, we use a piece-wise linear function to convert the AIMEinto the Primary Insurance Amount (PIA), which determines the Social Security benefits:

$$\begin{aligned} pia(a_t^{ss}) = 0.90 \times \min\{a_t^{ss}, b_0\} + 0.32 \times \min\{\max\{a_t^{ss} - b_0, 0\}, b_1 - b_0\} \\ +0.15 \times \max\{a_t^{ss} - b_1, 0\} \end{aligned} \tag{6}$$

The Social Security system in the model provides several work disincentives at older ages. For instance, the Social Security wealth a_t^{ss} is revised upwards only if current earnings are greater than average past earnings (as shown in 5). Therefore, staying in the labor market for longer by working fewer hours may not increase the benefits for the individuals in the model.[17] Additional work disincentives are due to the penalty/reward system associated with the timing of the SS application and earnings test, as we discuss below.

Adjustments and SS Benefit. Social Security benefits, ssb_t, depend on the PIA defined above as well as on two possible adjustments: (1) a penalty/credit for claiming early/late (Γ_t); and (2) a claw-back of benefits for those who continue working while claiming benefits (Υ_t).

$$ssb_t = pia(a_t^{ss}) * \Gamma_t - \Upsilon_t \tag{7}$$

[17] In practice, the highest 35 years of covered earnings are used to compute AIME. If the individual has not yet worked for 35 years, some zeros are included in the average, and any positive earnings, including part-time work, will increase the AIME.

Adjustment (1): Early/Late Claiming Penalty. SS benefits can be claimed without penalty at the normal retirement age (t_{NRA}).[18] However, individuals can claim benefits with some penalty starting from the Early Retirement Age (t_{ERA}) of 62. For every year before the NRA that these benefits are claimed, the amount received is permanently reduced by the early claiming penalty. Individuals can delay their benefit claim beyond NRA. In that case, future benefits are permanently increased by the delayed claiming credit.

It has been argued in the literature (Heiland and Yin, 2014; Gruber and Wise, 2005) that while the benefit reductions due to early claim are actuarially fair, the delayed claim benefit increase does not fully compensate the beneficiary for the loss of benefits in previous periods. This structure of the Social Security system provides strong incentives to not delay benefit claims. The actual incentives may depend on a variety of other factors such as an individual's subjective mortality expectations, heterogeneous discount factors, and more. In the model, penalties show up as a percentage decrease, γ_t^{ss}, for each year prior to the normal retirement age that a worker claims; credits show up as a percentage increase for each year after the normal retirement age that a worker delays claiming.

$$\Gamma_t = \begin{cases} 1 - (t_{NRA} - t^{ss}) * \gamma_t^{ss} & \text{if } t^{ss} < t_{NRA} \\ 0 & \text{if } t^{ss} = t_{NRA} \\ 1 + (t^{ss} - t_{NRA}) * \gamma_t^{ss} & \text{if } t^{ss} > t_{NRA} \end{cases} \tag{8}$$

Adjustment (2): Earnings Test. The SS system has an earnings test that taxes the labor income of beneficiaries above a certain threshold y_t^{ss} at a rate τ_t^{ss}, until the age of 70. For each additional dollar earned above the threshold, SS benefits are reduced by τ_t^{ss}, until all benefits are taxed away:

$$\Upsilon_t = \min\{pia\left(a_t^{ss}\right), \max\{0, w_t h_t - y_t^{et}\}\tau_t^{et}\}$$

Υ_t denotes benefits lost through the earnings test. Taxed benefits are credited back through permanent increases in future benefits, which is implemented in the model through increases in the

[18]The NRA differs slightly across birth cohorts. For the sample used in this analysis, the average NRA is 65. Later cohorts have NRA of 66 or 67.

Social Security wealth as shown below:[19]

$$ssb_{t+1} = pia(a^{ss}_{t+1}) * \left[1 + \left(\frac{\Upsilon_t}{ssb_t}\right)\gamma^{ss}_t\right]$$
$$a^{ss*}_{t+1} = pia^{-1}(ssb_{t+1}) \tag{9}$$

where γ^{ss}_t is the same reduction/increment factor that is used for determining penalty/credit for early/late benefit applications as discussed earlier. The work incentives introduced by the earnings test crucially depend on γ^{ss}_t. In combination with the application age requirements, the earnings test may reinforce the incentives to retire upon reaching the claiming age.

4.5.1 Marriage Related Benefits

Spousal Benefits

Married households receive additional income through Social Security spousal benefits. Spouses of household heads are entitled to up to 50 percent of the head's benefits depending upon the age benefits are claimed. We assume that all spouses claim together, and thus the size of the spousal benefits received is a function of the head's age at SS claiming, t^{ss}, and of the age gap between spouses, ι. The total SS benefits received by a household are $\delta^q_t ssb_t$, where δ^q_t is determined as:

$$\delta^q_t = \begin{cases} 1.0 & \text{if } m = \text{single or } m = \text{married}, t^{ss} - \iota < t_{ERA} \\ 1.5 * [1 - (t_{NRA} - (t^{ss} - \iota)) * \gamma^{ss}_t)] & \text{if } m = \text{married}, \ t_{ERA} \leq t^{ss} - \iota < t_{NRA} \\ 1.5 & \text{if } m = \text{married}, \ t^{ss} - \iota \geq t_{NRA} \end{cases} \tag{10}$$

Singles and married individuals whose spouse is not yet eligible for benefits ($t^{ss} - \iota < t_{ERA}$) receive no additional spousal benefits. Married individuals for whom the spouse's age is above the normal retirement age, receive the additional 50 percent of benefits. Married individuals whose wives are between 62 and 65 at the time of claiming receive benefits penalized by the early retirement penalty.Spousal benefits do not accrue delayed retirement credits and are maximized at the spouse's normal retirement age.

[19]This is a simplification. The benefits are typically adjusted upon reaching the NRA. The earnings test was removed for workers over the NRA starting in the year 2000. Since SS rules have been changing over time, the restrictions pertaining to the sample used in this analysis are based on the SSA.

Survivors Benefits

Upon death, married individuals may also pass part of their SS entitlements on to their spouses. These survivors benefits enter into the end-of-life wealth of individuals, A_t^q, which takes the form:

$$A_t^q = \begin{cases} a_t + \sum_{j=t-\iota}^{T} \frac{1}{1+r} \pi_{j+1}^s ssb_t & \text{if } m = \text{married}, t - \iota \geq 62 \\ a_t & \text{otherwise} \end{cases} \tag{11}$$

In addition to any leftover assets, a_t, end-of-life wealth is a function of Social Security wealth if the individual is married and the spouse is over the age of 62. Survivors benefits are calculated as the present value of the stream of benefits a spouse would receive from the time of the death of the household head until the end of her own life. Therefore, the present value is a function of the household head's age t and the spousal age gap, ι (see Bairoliya and McKiernan, 2021).

4.6 Budget Constraint

The household budget constraint can be summarized as:

$$c_t + i_t^v \times \rho(v_t^{e,m}) + \rho_{mt}^{e,m}(\mu) + a_{t+1} = a_t + W\left(y_t, y_{st}, ra_t, \tau\right) + \delta_t^q ssb_t + tr_t \tag{12}$$

Where i_t^v is an indicator representing the decision to purchase a one-period life insurance contract. The price of this contract, $\rho(v_t^{e,m})$, is dependent on the face value of life insurance chosen, $v_t^{e,m}$. Labor income, y_t, is a function of the hourly wage and work hours chosen by the individual. We let spousal income for married households be a function of the head's age, health status, and labor income, as follows:

$$y_{st} = f\left(t, \mu_t, w_t h_t\right) \tag{13}$$

There is a simple no-borrowing constraint on assets:

$$a_{t+1} \geq 0 \ \forall t \tag{14}$$

4.7 Government

The government taxes individuals with a proportional payroll tax, τ_t^{ss}, and labor income taxes, τ. The payroll tax τ_t^{ss} includes both the Social Security duties and the Medicare tax. The Social Security payroll duty is 6.2 percent on income up until the maximum taxable amount, a^{max}, while the Medicare tax is 1.45 percent on total labor income.

We adopt a smooth functional form for the labor income tax that allows for negative tax rates to account for Earned Income Tax Credit (EITC).[20] We let the function vary with education and estimate the following function from PSID data:

$$\tau = 1 - \lambda y^{-\xi}.$$

The government guarantees a minimum consumption level (Hubbard et al., 1995) $c_t \geq \bar{c}$. Government transfers, denoted as tr_t, bridge the gap between the minimum level of consumption and an individual's liquid resources, that is:

$$tr_t = \min\{0, \underline{c} - (a_t + W_t + \delta_t^q ssb_t)\}, \tag{15}$$

where W_t is the total disposable household income defined in equation (12). This approximates federal safety-net programs such as the Supplemental Nutritional Assistance Program (SNAP), the Supplemental Security Income (SSI), and the Temporary Assistance for Needy Families (TANF).

4.8 Recursive Formulation

The period t individual state vector is $z_t = \left(e, q, \eta_t, \lambda_t, \mu_t, a_t, a_t^{ss}, b_{t-1}^{ss}\right)$. Individuals solve a finite-horizon Markov problem where they choose a sequence of consumption $\{c(z_t)\}_{t=1}^T$, hours $\{h(z_t)\}_{t=1}^T$ and SS benefit application rules $\{b^{ss}(z_t)\}_{t=1}^T$ to maximize their expected discounted utility subject to the exogenous processes for health and employment transitions, survival and wage risk, a set of budget, borrowing, and time constraints, a government transfer rule, and policies for taxes and Social Security.

4.8.1 Stages of the life cycle

The life cycle of an individual between ages 25 and 99 consists of three phases. The first is the *employment* phase between ages 20 and 61 when individuals make consumption, savings, and employment decisions.[21] The second stage is the *retirement choice* phase between ages 62 and 69 when individuals can make Social Security application decisions (b_t^{ss}). Finally, there is a *retirement* phase when individuals make only consumption and savings decisions.

Employment phase. The problem of a household head in the initial phase is:

[20] Supplemental appendix S7 provides details of our estimated tax function.

[21] We do not allow individuals to claim disability benefits and only estimate the model for individuals who claim Social Security through the non-disability route.

$$V(a_t, a_t^{ss}, \eta_t, \lambda_t, \mu_t) = \max_{c_t, h_t, a_{t+1}, i_t^v} \Bigg\{ U(c_t, l_t) + \beta \pi_{t+1}^s \Big[EV(a_{t+1}, a_{t+1}^{ss}, \eta_{t+1}, \lambda_{t+1}, \mu_{t+1}) \Big] + \beta(1 - \pi_{t+1}^s) \Omega(A_{t+1}^q) \Bigg\} \quad s.t.$$

$$c_t + i_t^v \times \rho(v_t^{e,m}) + \rho_{mt}^{e,m}(\mu) + a_{t+1} = a_t + W\left(y_t, y_{st}, ra_t, \tau\right) + tr_t$$
$$1), (5\text{-}9), (14), \text{ and } c_t \geq \bar{c}.$$

where $y_t + y_{st} + ra_t$ is the total pre-tax income and $W(., \tau)$ is the level of post-tax income, given tax rate τ. The expectation is taken with respect to wages, employment, and health uncertainty.

Claiming phase. Starting at age 62, individuals can make a benefit claim. The claim is a one-time decision and benefits are based on the age at which the individuals claim for the first time. If an individual enters a period as a non-claimer, they must choose whether or not to claim benefits during this period, as shown below:

$$V(a_t, a_t^{ss}, \eta_t, \lambda_t, \mu_t, b_{t-1}^{ss} = 0) = \max \left\{ V^{b_t^{ss}=0}, V^{b_t^{ss}=1} \right\}$$

where the value of postponing the claim, $V^{b_t^{ss}=0}$, is

$$V^{b_t^{ss}=0}(a_t, a_t^{ss}, \eta_t, \lambda_t, \mu_t, b_{t-1}^{ss} = 0) = \max_{c_t, h_t, a_{t+1}, i_t^v, b_t^{ss}} \Bigg\{ U(c_t, l_t) + \beta \pi_{t+1}^s \Big[EV(a_{t+1}, a_{t+1}^{ss}, \eta_{t+1}, \lambda_{t+1}, \mu_{t+1}, b_t^{ss} = 0) \Big] + \beta(1 - \pi_{t+1}^s) \Omega(A_{t+1}^q) \Bigg\} \quad s.t.$$

$$c_t + i_t^v \times \rho(v_t^{e,m}) + \rho_{mt}^{e,m}(\mu) + a_{t+1} = a_t + W(y_t, y_{st}, ra_t, \tau) + tr_t,$$
$$1), (5\text{-}9), (14), \text{ and } c_t \geq \bar{c}.$$

The value of filing the claim in the current period is:

$$V^{b_t^{ss}=1}(a_t, a_t^{ss}, \eta_t, \lambda_t, \mu_t, b_{t-1}^{ss}=0) = \max_{c_t, h_t, a_{t+1}, i_t^v, b_t^{ss}} \Bigg\{ U(c_t, l_t) + \beta \pi_{t+1}^s \Big[EV(a_{t+1}, a_{t+1}^{ss}, \eta_{t+1}, \lambda_{t+1}, \mu_{t+1}, b_t^{ss}=1) \Big] + \beta(1-\pi_{t+1}^s)\Omega(A_{t+1}^q) \Bigg\} \quad s.t.$$

$$c_t + i_t^v \times \rho_{it}^{e,m}(v) + \rho_{mt}^{e,m}(\mu) + a_{t+1} = a_t + W(y_t, y_{st}, ra_t, \tau) + tr_t + \delta_t^q ssb_t,$$
1), (5-9), (14), and $c_t \geq \bar{c}$.

Retirement phase. At age 70, if an individual has still not claimed, they automatically start receiving their benefits and (if applicable) spousal benefits. Their value function is:

$$V(a_t, a_t^{ss}, \mu_t) = \max_{c_t, a_{t+1}, i_t^v} \Bigg\{ U(c_t, l_t) + \beta \pi_{t+1}^s EV(a_{t+1}, a_{t+1}^{ss}, \mu_{t+1}) + \beta(1-\pi_{t+1}^s)\Omega(A_{t+1}^q) \Bigg\} \quad s.t.$$

$$c_t + i_t^v \times \rho(v_t^{e,m}) + \rho_{mt}^{e,m}(\mu) + a_{t+1} = a_t + W(y_{st}, ra_t, \tau) + \delta_t^q ssb_t + tr_t,$$
(1), (6), (14) and $c_t \geq \bar{c}$.

5 Estimation

We estimate the model on a sample of male household heads born between 1931 and 1935. Estimation proceeds in two-steps, along the lines of Gourinchas and Parker (2002), following the details outlined in Bairoliya and McKiernan (2021). In the first step, we combine several data sets—including the Panel Study of Income Dynamics (PSID), the Health and Retirement Study (HRS), and the Household Component of the Medical Expenditure Panel Study (MEPS)—to estimate processes that can be identified without imposing the restrictions of the dynamic life-cycle model. These include health transitions, out-of-pocket medical expenditures, survival probabilities, family structure and spousal income, wages, unemployment probabilities, the tax function, and the exogenous rate of return on assets.

In the second step, we use initial conditions drawn from data for the cohort born between 1931-1935, our structural model, and the parameters from the first step to estimate the preference

parameter vector $\{\beta, \rho, \nu, \theta, \kappa, \phi_H(t, \mu_t), \phi_P(t), \rho(v_t)\}$ by education and marital status. In this step, we employ the Method of Simulated Moments (MSM) and use the model to simulate the life-cycle of a representative cohort of individuals and find the parameter values that allow simulated moments to best match (as measured by a GMM criterion function) the data profiles for this cohort of male household heads. The data that inform the structural estimation process are composed of the following 45 moments for each education and marital status group (refer to Appendix S9 for the construction of data moments):

1. Labor market participation of male household heads within 5 year age bins between 25 and 99 (ages 25-29, 30-34, 35-39, 40-44, 45-49, 50-54, 55-59, 60-64, 65-69), resulting in 9 moment conditions for each education and marital group.

2. Hours worked, conditional on participation, of male household heads within 5 year age bins between 25 and 99 like above, resulting in 9 moment conditions for each education and marital group.

3. Median assets of male household heads within 5 year age bins between 25 and 99 (ages 25-29, 30-34, 35-39, 40-44, 45-49, 50-54, 55-59, 60-64, 65-69, 70-74, 75-79, 80-84, 85-99), resulting in 13 moment conditions for each education and marital group.

4. Life insurance coverage rates for male household heads within 10 year age bins between ages 25 and 99 (ages 25-35, 35-45, 45-55, 55-65, 65-75, 75-85, 85-99), resulting in 7 moment conditions for each education and marital group.

5. Median face value of life insurance policies held by male household heads within roughly 10 year age bins between ages 25 and 99 like above, resulting in 7 moment conditions for each education and marital group.

5.1 Estimation: First Step

Health, mortality and medical expenditures: Health can take three values in the model, $\mu_t = \{\text{excellent}, \text{good}, \text{poor}\}$. We identify these health states from the self-reported health status variable in the MEPS.[22] We estimate health transitions across these states through an ordered probit of

[22]The Medical Expenditure Panel Survey asks respondents to report their health on a 1-to-5 scale: 1 is "Excellent", 2 is "Very Good", 3 is "Good", 4 is "Fair", and 5 is "Poor". We convert the 5-point scale to a 3-point scale, grouping individuals with "Very Good" and "Good" score into the good health category and those with "Fair" and "Poor" scores into the poor health category. We could instead use a frailty index (see Hosseini et al., 2022). As discussed in Miller and Bairoliya (2021), self-rated health is predictive of mortality even after controlling for other health conditions and behaviors, which suggests that people may have private info about their overall state above and beyond what is recorded in the frailty index.

current health on the previous year's health status, education, and a quadratic function of age. Survival probabilities are also obtained from the MEPS. We estimate the age-, education-, marital status-, and health-specific profiles by running an ordered probit model of a death indicator on health status, a quadratic in age, education, and marital status. Finally, we estimate annual out-of-pocket medical expenditures from the Health and Retirement Study data using all years from 1992-2018, following the same approach as detailed in Borella et al. (2019). Supplemental Appendices S1 and S2 provide futher details.

Family structure: Family structure determines two parameters for married men: the consumption equivalence scale, ζ_t, and the gap between spouses, ι. In addition, married men receive spousal income (see Bairoliya and McKiernan, 2021). We construct the consumption equivalence scale by education and marital status using family statistics from the PSID. Single households have an equivalence scale of 1. The equivalence scale of married households depends on the presence of a spouse and the average number of children living in the household for each age-education type. Given family size, values of ζ_t are set using the OECD equivalence scale, which gives a weight of 1 to the household head, 0.5 to the spouse, and 0.3 to each child. For married couples, the age gaps between the male household head and their spouse are based on the distribution of age gaps for the cohort at hand. We use four age gap states $(0, 1, 4, 8)$ to describe this distribution and assign the mass at each point from PSID data. About 8.7 percent of married couples have no age gap, 26.2 percent have an age gap of one year, 46.1 percent have an age gap of four years, and 19 percent have an age gap between spouses of eight years. Spousal income y_{st} is estimated from the PSID as a function of the age, education, health, and labor income of the household head. Refer to Supplemental Appendix S3 for further details.

Labor productivity, employment shocks and wage scarring: Wages are comprised of an age and education profile and a persistent shock. The age and education function, as well as parameters of the AR(1) shock process, are estimated on a PSID sample (Bairoliya and McKiernan, 2021). The probability that a worker is separated from the labor market is independent of education and marital status. We set the employment shock, λ, at 0.1 to match the separation rate in the JOLTS. The wage penalty associated with a bad employment shock, ξ, is modeled as a percentage of income. The penalty is estimated from the PSID following the literature on wage scarring and set to $\xi = 0.86$ (see, for example, Jacobson et al., 1993; Huff Stevens, 1997; Huckfeldt, 2016). To estimate the displacement penalty, the log of hourly wages is regressed on dummies representing years since displacement as well as a vector of control variables including a quadratic in age and a quadratic in experience. The penalty is set to the percentage drop in annual wages that displaced workers experience on average. Supplemental Appendices S4 - S6 provide further details of our

estimation.

Social security. Modeling the rich detail of the U.S. Social Security System, described in Section 4.5, requires a set of parameters. Table 1 shows these parameters based on the 1998 rules from the U.S. Social Security Administration. A first subset of parameters, b_0, b_1, and a^{max}, determine the value of Social Security wealth and benefits. The maximum wealth at which benefits are capped is a^{max} and is set at \$68,400. The parameters b_0 and b_1 define the bend points of the Social Security benefits formula, $g(\cdot)$. These points are set to \$5,724 and \$34,500. There is no variation in these parameters based on the claiming age. A second subset of parameters is based on the earnings test. Before the normal retirement age, earnings above \$9,120 are taxed at a rate of 50 percent. After the NRA, earnings above \$14,500 are taxed at 33 percent. The normal retirement age varies with birth cohort and is age 65 for our benchmark birth cohort (born in 1931-1935). The last parameter in Table 1 defines the penalty for early claiming (or the benefit for delayed claiming). Benefits decrease by 6.7 percent for each year prior to the NRA the worker claims. After the normal retirement age, benefits increase by 5.5 percent for each year of delays in claiming benefits.

Table 1: Social Security Benefit Formula

Parameter	Value*	
	before the NRA	after the NRA
a^{max}	68,400	68,400
b_0	5,724	5,724
b_1	34,500	34,500
Earnings Test		
y^{et}	9,120	14,500
τ^{et}	0.50	0.33
γ_t^{ss}	0.067	0.055

*1998 rules from the SSA and those pertaining to the 1931-1935 birth cohort.

5.2 Estimation: Second Step

We allow the discount rate, bequest motives, time endowment, time costs, and life insurance prices to vary across different marital and education groups. However, we keep the consumption weight and risk aversion parameter the same across all groups. This approach helps capture the differences in wealth accumulation, labor supply, and life insurance holdings among individuals based on their marital and education status in a parsimonious manner. Estimating end-of-life motives for each group is central to our analysis. By allowing the discount rate to vary, we can

better capture differences in early-life wealth growth, while adjusting bequest motives enables us to capture varying rates of depletion of capital stock later in life across these groups.

Table 2 shows our structural parameter estimates for $(\beta^{e,m}, \theta_{beq}^{e,m}, \kappa_{beq}^{e,m}, \bar{l}^{e,m})$. The estimated discount factor is higher for married men than singles and for college graduates than those without a college degree (for singles). Together, the bequest parameters $\theta_{beq}^{e,m}$ and $\kappa_{beq}^{e,m}$ govern the strength of terminal bequest motive. While our estimates indicate variation in $\theta_{beq}^{e,m}$ and $\kappa_{beq}^{e,m}$ across groups, it is important to note that both the marginal propensity to bequeath out of wealth and the wealth thresholds for individuals to leave positive bequests are non-linear function of the parameters $(\nu, \rho, \beta^{e,m}, \bar{l}^{e,m}, \theta_{beq}^{e,m}, \kappa_{beq}^{e,m})$.

We jointly estimate preference parameters and the cost of LI holdings, $\rho\,(v_t^{e,m})$. To this purpose, we match the coverage of life insurance and the face value of policies held. We emphasize that the estimates of LI shadow prices (presented in Figure 5 for each education-marital status group) reflect both the pecuniary costs of the contracts and potential constraints to access such contracts. Two key observations can be made about these estimates. First, the prices per dollar of LI are increasing in both age and the face value of the policy. This is consistent with the notion that prices reflect the expected payout of the policy. As the expected payout increases with falling survival probabilities and the value of the policy, one would expect higher prices. Second, comparisons to the actuarially fair price, which incorporates variation in life expectancies across groups–shown as the dashed line in Figure 5–indicate that married individuals face prices that are approximately fair. This is in sharp contrast with the market for non-married individuals, where the high face value contracts are more expensive than implied by survival rates, particularly for non-college individuals. Deviations from the actuarially fair prices may be due to asymmetric information or other frictions; regardless of their origin, these price deviations result in limits to LI coverage and in imperfect end-of-life insurance.

5.3 Model fit

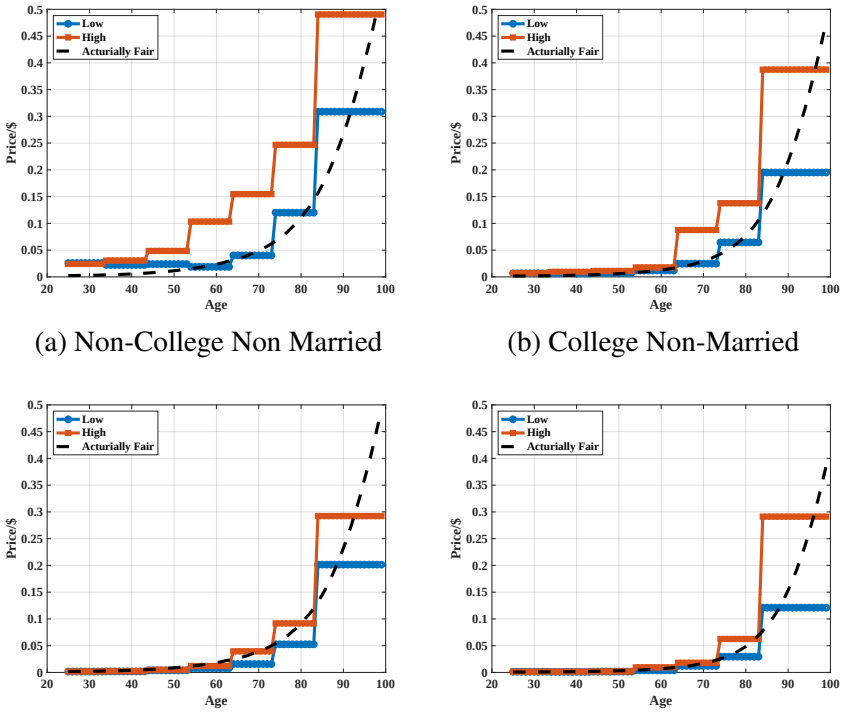
Figures 6 and 7 show the observed and simulated life insurance coverage rates as well as the median face value of life insurance policies (conditional on having a positive value). The model has sufficient flexibility to directly match these moments.

Further, Appendix Figure A6 plots observed and simulated participation rates by age, education, and marital status and documents labor market fit. Model values are within the 95% confidence intervals and exhibit realistic patterns of decline over the life-cycle. This is true also around the retirement phase when incentives change significantly and models do not always fit empirical observations. While we only target median wealth in our estimation, Figure 8 shows that the model also roughly matches the age evolution of the 25^{th} and 75^{th} percentile of the wealth distribution.

Table 2: Preference Parameters

Parameter	Description	Singles		Married	
		Non-College	College	Non-College	College
ρ	relative risk aversion	3.340	3.340	3.340	3.340
ν	consumption weight	0.578	0.578	0.578	0.578
β	discount factor	0.941	0.967	1.014	1.003
$\bar{l}$	time endowment	5429	4856	5749	6173
θ_{beq}	bequest intensity	1.295	1.118	1.231	1.437
κ_{beq} (in 000s)	bequest curvature	1.102	2.084	1.566	2.102

Figure 5: Estimated Life-Insurance Price per Dollar of Coverage
by Education Marital Status



(a) Non-College Non Married

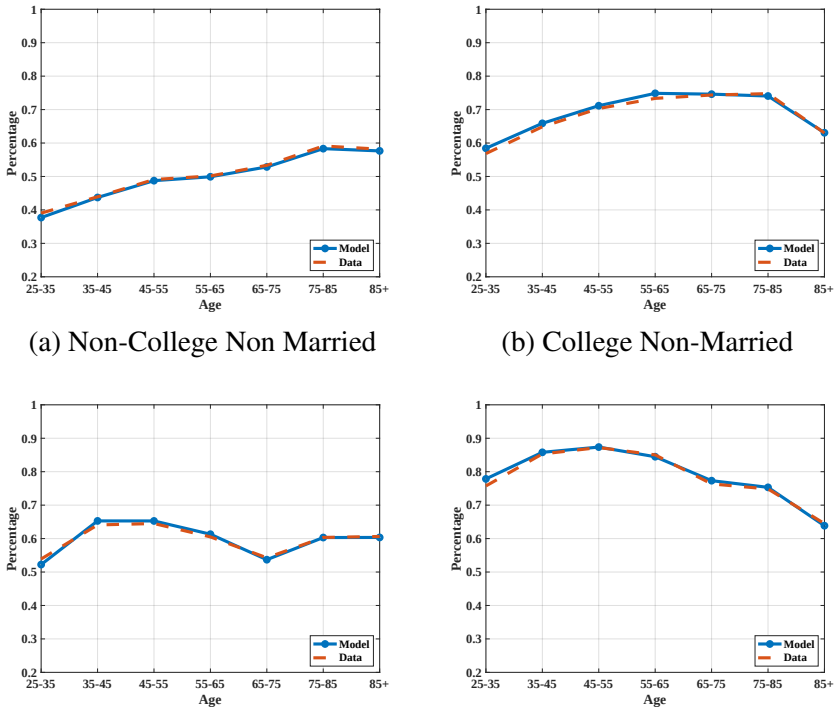
(b) College Non-Married

(c) Non-College Married

(d) College Married

The model closely matches variation in levels and growth, including the key observation of no dis-saving in old age. Finally, even though not in our estimation, Appendix Figure A9 compares observed and simulated claiming rates at four junctures: at the ERA, between ERA and NRA, at the

Figure 6: Life-Insurance Coverage and Face Value by Education Marital Status



(a) Non-College Non Married

(b) College Non-Married

(c) Non-College Married

(d) College Married

NRA, and after the NRA. The canonical model that we use is known to marginally overstate delay and understate early claims; however, magnitudes are accurate and provide a good approximation of empirical observations.

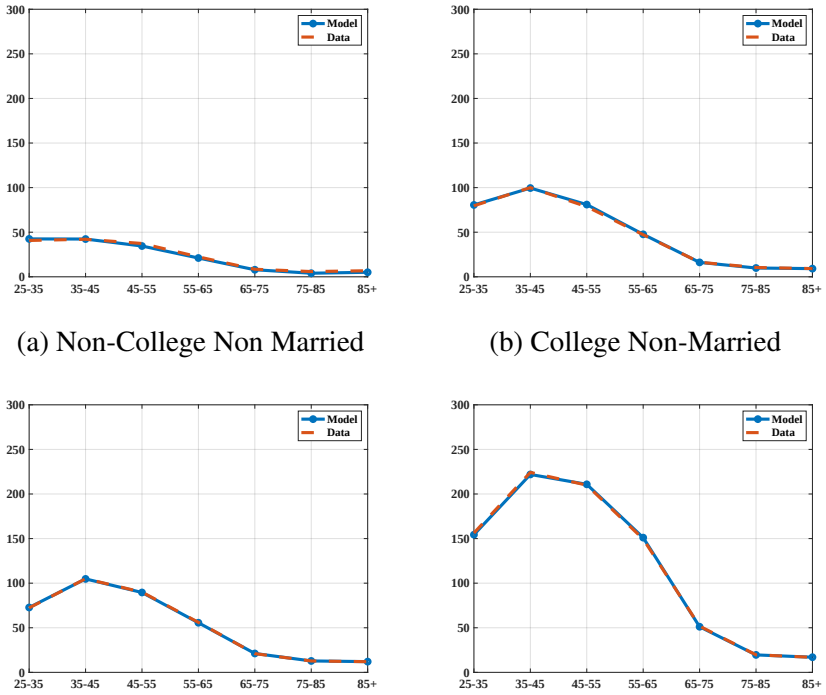
6 Universal Transfer Reform

We examine a reform that introduces a universal end-of-life payment disbursed regardless of marital status. Specifically, we study a system that replaces part of the annuity benefit with an end-of-life lump sum transfer. In this experiment, we hold government outlays on SS lifetime payouts at the benchmark level. Denoting the total expected present value of the end-of-life lump sum transfers as PV^D, we impose the following constraint:

$$\sum_{i=1}^{N}\sum_{t=62}^{100}\left[(1-d_{it})\alpha^{ss}ssb^{*}_{it}\right] + PV^D = \sum_{i=1}^{N}\sum_{t=62}^{100}\left[(1-d_{it})ssb^{*}_{it}\right] + SVB \tag{16}$$

$$\text{Where: } PV^D = \sum_{i=1}^{N} PV^D_i; \quad SVB = \sum_{i=1}^{N} \mathbb{I}_{\{m_i=1\}} SVB_i$$

Figure 7: Life-Insurance Face Value by Education Marital Status



(a) Non-College Non Married

(b) College Non-Married

(c) Non-College Married

(d) College Married

In the expression above, the d_{it} are death indicators based on simulated health transitions and survival probabilities, ssb^*_{it} is the benchmark value of the annuity benefits paid to individual i at age t, SVB_i is the survivor benefits paid to the surviving spouse of a married individual upon death (if they have already claimed their benefits), α^{ss} is the scaling factor that scales down the value of the annuity benefits in the experiment, and PV_i^D is the lump sum paid out to the original SS program recipient at the end of life (regardless of their claiming status). The constraint guarantees that Social Security disbursements are the same as in the baseline.

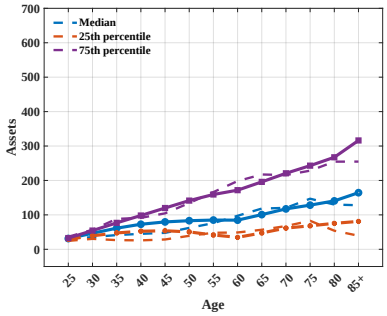
Note that the lump-sum amount varies in the cross-section of individuals because it is set to a multiple of their primary insurance amount (PIA). The multiple is such that the expected payout, in the form of annuity and lump sum, is the same as the expected entitlement upon initial receipt of benefits in the benchmark economy. Surviving spouses receive a marginally smaller annuity in the experiments, which is proportional to the annuity received by the main beneficiary before death.

6.1 The Trade-Off Between Annuities and EOL Transfers

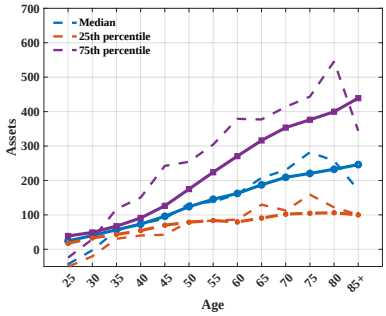
First, we study the life-cycle and welfare impacts of such a reform a fixed value of the lump-sum end-of-life payment. In this experiment, all individuals are guaranteed, upon death, the equiv-

Figure 8: Model Fit: Wealth Distribution

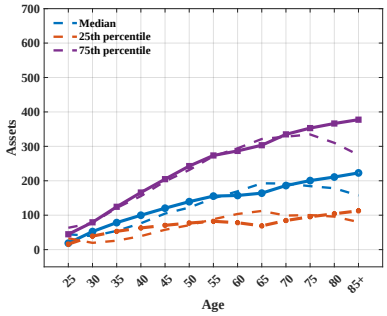
(a) Non-College, Non Married



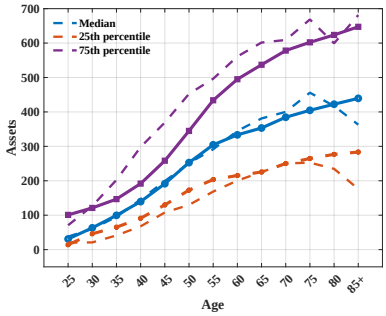
(b) College, Non Married



(c) Non-College, Married



(d) College, Married



alent of five and a half years' worth of their PIA entitlement; the transfer is financed through a smaller per-year benefit, and the percentage drop in the annuity is such that the total expected benefits remain unchanged. In practice, the introduction of the transfer requires a marginal cut of roughly 4% in the annuity payments. The cut is similar to what individuals would give up if they brought forward retirement by around 7 months relative to the normal retirement age. The utility value of the end-of-life transfer is derived from the function described in Equation (3). We do not model intergenerational linkages and the transfer does not add to the consumption of surviving beneficiaries. For this reason, the estimates of welfare changes may be viewed as a lower bound.

6.1.1 Cumulative Life-Cycle Impacts

Table 3 shows the effects of the reform on labor supply, asset holdings, average consumption, and welfare (in consumption equivalent values, CEV). We report results by education and marital status. Despite the small departure from the baseline SS system, the impacts are sizable.

Singles: Welfare grows significantly among non-married individuals, who now have guaranteed access to part of their entitlement even if they die before claiming SS annuity income. The average CEV change is 3.34% for the relatively poor non-college singles. The reform induces them to de-

Table 3: Universal Transfer Reform

	Singles		Married	
	Non-College	College	Non-College	College
Participation	0.06	0.59	1.04	0.54
Hours	0.19	0.81	0.82	0.25
Assets	-23.33	-19.61	3.68	3.23
Consumption	2.27	2.04	0.24	-0.08
LI coverage	-12.74	-12.90	-14.91	-3.71
Welfare	3.34	1.88	-0.54	-0.32

Notes: Experiment entails a lump sum payment to every individual. The payment is 5.5-year's worth of their PIA at the time of death. The transfer is introduced in conjunction with a 4% reduction in the annual flow of SS benefits. This guarantees that the lifetime SS outlay of each individual remains at its benchmark level.

lay retirement as these individuals no longer need to claim their annuities in order to avoid losing everything in case of premature death. For instance, benefit claiming at age 62 (earliest age of pension eligibility) goes down by 60 percentage points for non-college singles (refer to appendix figure A14). College-educated singles experience a relatively smaller CEV gain of 1.88% on average with most individuals exhibiting gains slightly over 1.5 percent (see Figure 9). Both college and non-college singles experience an increase in labor force participation and hours worked at older ages. This effect is somewhat stronger for the college group. The strong labor supply response of both groups, especially close to years around SS benefit eligibility, is partly due to the novel ability, under the reform, to circumvent the earnings-test claw-back that applies after claiming SS benefits (refer to appendix figure A10). Delaying SS claims allows these earners to supply more hours at older ages without being subject to the implicit earnings tax of the claw-back of the original system.

Regardless of education, singles reduce their savings and life insurance holdings significantly under the reform (refer to appendix Figures A11 and A13). The drop in asset holdings can be as large as a quarter of the baseline values at old ages, reflecting a diminished need for end-of-life hoarding of resources. This finding is related to the ongoing debate on the origins of the high saving rates among the elderly and suggests that a non-trivial share of their wealth might be due to end-of-life considerations that extend beyond their own death. A policy that liquidates some of their entitlements upon death appears effective in reducing the asset holdings of older non-married individuals.Remarkably life insurance holdings go down substantially (between to 5 and 10 percentage points) over the life-cycle.

Married individuals: Welfare changes are negative but small among married individuals: -0.54% CEV for non-college households and -0.32% CEV among the college-educated.

Non-college workers tend to be less wealthy and, by delaying SS claims, are able to participate in the labor market at older ages and with more hours. This group has the strongest labor supply response in the experiment : this suggests that the design of the current SS system generates large distortions in their behavior.

The reform induces married individuals to decrease savings during their working lives but increase them later in life (Figure A11). Under the baseline Social Security model, surviving spouses receive marital benefits only if the surviving spouse is of retirement age; therefore, any end-of-life transfer motive would introduce an incentive to save during middle life and insure against death during their working years. The reform, however, makes everyone eligible for a transfer that is not conditional on being of retirement age: this significantly reduces the need to stash away resources earlier in life.

Conditional on surviving to older ages, married individuals receive marginally lower survivor benefits under the new regime, which may lead to the accumulation of additional resources for precautionary motives after retirement. We do find evidence for this offsetting force, but mostly among the less educated workers (Figure A11). The changes in saving behavior are smaller in magnitude among married individuals, who tend to be richer. This is most apparent for college-educated couples, whose dynamic choices are not overly dependent on SS benefits and therefore exhibit smaller responses to the reform.

Figure 9: Welfare: Universal Transfer Reform

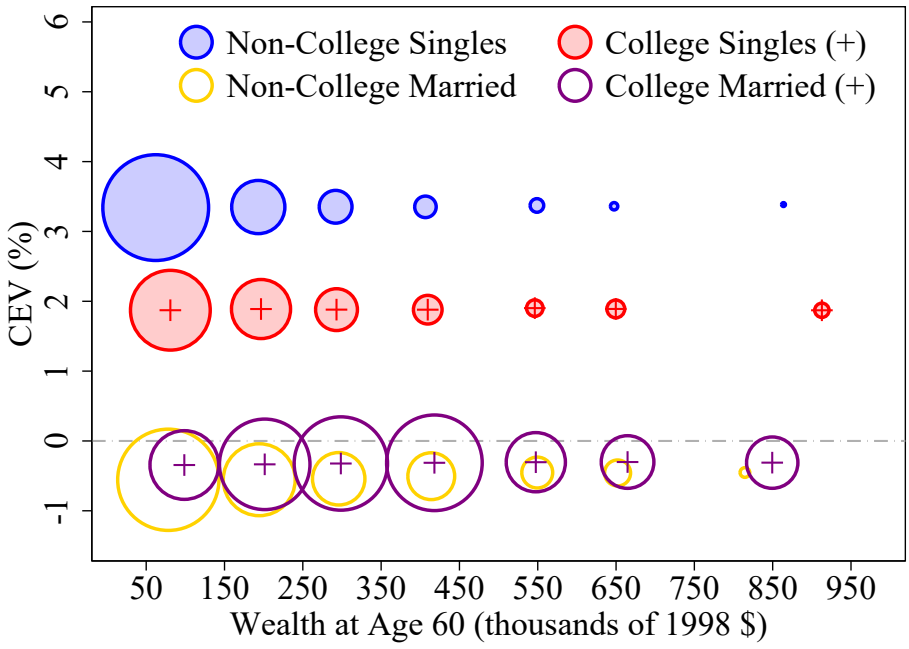


Figure 9 provides an overview of the distribution of welfare changes from the Universal Transfer Reform across different worker groups. The size of the bubbles conveys information about the relative size of workers of different marital status and wealth. The key takeaway from this plot is that the vast majority of workers has higher welfare under the reform: for some groups, like the non-college single workers, the gains are extremely large. As expected, the married workers are the one who earn less, or marginally lose, from the reform since they have to give up the survivor benefits of the status quo SS policy. However, even among married individuals, the majority exhibits positive welfare changes after the reform. Among those who exhibit a marginally lower welfare (college married) the changes are fairly small and indicate that the flexibility of the reform does not fully make up for the loss of the survivor benefits in the status quo SS system. Nonetheless, their losses are not large and confirm that the reform induces positive welfare changes on average.

6.1.2 Valuing the Reform at Different Junctures of the Life-Cycle

To shed light on the mechanics of the reform and convey information on its value at different stages of the life cycle, we compute measures of the benefits' worth (in dollar terms) at different ages. Figure A15 plots the flow value of accruable benefits in the benchmark (legacy) economy and under the universal transfer reform at each age between 25 and 99 for workers of different education and marital states. Accruable means that the benefit would be payable upon death at a specific age.

The first striking difference is that the value of benefits starts growing from the beginning of the working life in the reform experiment. Unlike the legacy system, in which all benefits are conditional on claiming (or reaching age 62 in the case of married individuals for access to survivors benefits), the reform guarantees an end-of-life transfer to every worker regardless of their age at death. Moreover, the transfer grows with a worker's PIA, which explains the positive slope of the benefits flow measure between ages 25 and 60. In contrast, the flow value of the benchmark SS policy is stuck at zero until claiming occurs after 60.

The claiming patterns across groups are themselves very different. Each panel of A14 shows a jump in benefit flows at the average age at which the annuity is claimed. what is this referring too? It is apparent that in all groups the claiming ages change in response to the reform, sometimes by several years. However, while single individuals delay their claims, married workers move their claims earlier. These strong responses capture the relationship between claiming and liquidity preferences. While singles have reduced incentive to start claiming in order to preserve entitlements and save for transfers, married individuals have higher incentives to do so due to the removal of the more generous survivor's benefits.

The flatness of benefit flows after retirement reflects the constant annuity income and the fact that the PIA no longer changes after retirement, so the final transfer is also fixed. The gaps be-

tween benchmark and experiment accruable benefits at each age are large and can reach almost ten thousand dollars after retirement. However, note that while the flow benefits under the reform always remain higher than the benchmark for singles, it is not the case for married individuals. There are three distinct regions in this case as evident from figures A15c and A15d. Up until age 62, the reform clearly dominates as flow value of benefits under the reform remains higher. However, starting age 62, the legacy system favors the married individuals more given the higher level of spousal benefits. Finally, conditional on surviving to older ages, the reform dominates again as the value of survivors benefits (which is based on the remaining life-span of the surviving spouse), under the legacy system, is much lower at older ages. This is in stark contrast with the reform where the size of end-of-life transfers remains fixed after retirement.

The large welfare gains under the reform, especially for singles, follow from the fact that benefits can be accrued at young ages and before reaching retirement age, as shown in Figure A15. For many people, this means that entitlements are not fully lost due to premature death. This difference explains why the reform is so valuable even though the total government outlays are unchanged.

To further understand the value of the reform, we provide an illustrative example of a hypothetical agent in the model facing both the benchmark and the universal transfer reform economies. Figure A16 shows the present value at age 25 (discounted at a rate of 2%) of total SS benefits received conditional on death at a specific age between 25 and 99. In this Figure, we plot the PDV of total benefits by marital status and set the claiming age to 65 for both. The calculation is done for a fixed AIME of $52,000$ dollars to keep things simple. Under the benchmark SS policy, singles are assumed to claim at 62 (the modal age in the benchmark simulation). In the universal transfer experiment, singles are assumed to adjust their claiming to age 65 (the new modal claiming age in the experiment). The married are assumed to claim at 65 in both cases.

For singles, the reform results in a much higher PDV at every age. Under the benchmark SS policy, the PDV of benefits is stuck at zero until claiming. While for singles the PDV in the legacy system never catches up to the experiment's PDV, for married individuals this occurs after retirement since the reform does imply a marginal redistribution of resources from the married to the unmarried.

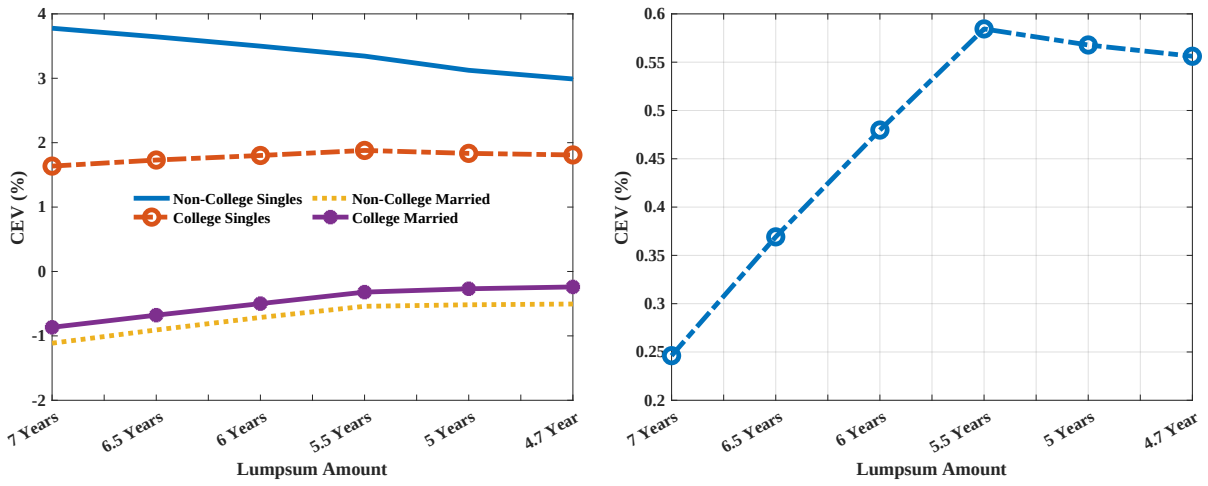
In Figure A18 we refine the present value analysis by considering three scenarios for the delay in the first age of benefits claiming under the universal reform (from 62 to 65, from 62 to 70, and from 65 to 70). This analysis is meant to convey information about the present value of the reform for households that change their claiming patterns differently. The main lesson is that the present value of the reform, relative to the benchmark, becomes large with the delay in claims. That is, households that delay claiming more tend to extract more value from the policy reform. This observation is instructive to make sense of the key forces triggered by the simple policy we

consider: bundling a transfer that is paid at any age and an annuity that starts after retirement, the system assuages concerns about loss of entitlements while preserving the approximate value of old-age income. In fact, the plots show that conditional on living very long lives (say, age 80 and beyond) the presence of the transfer becomes very valuable again since people don't have to worry about depleting their assets to satisfy the warm glow constraint. In this sense, the transfer appears to be most valuable for the very young and the very old.

6.2 Establishing the Value of Different EOL Transfers

Different subsets of workers will associate different values to alternative combinations of annuity and EOL transfer. In particular, the value of EOL transfers varies with their size and the necessary adjustment in annuities. To examine who likes larger EOL transfers, Figure 10 plots changes in welfare (CEV) for different amounts of the end-of-life lump-sum transfer. Experiments for each EOL transfer size are done in an expenditure neutral way where the size of the annuity received is decreased in order to finance the lump-sum payment. The penalty ranges from nothing at all (to provide an EOL payment of 4.7 years of benefits for all workers, financed in whole by the baseline survivors benefits of the married workers) up to 16 percent of the annuity value in order to finance an EOL transfer of 7 years of benefits for all workers.

Figure 10: Welfare: Variation Across EOL Transfers



(a) (b)

Notes: First panel reports average CEV of each education-marital status group across different EOL transfer sizes. Each EOL size is associated with a decrease in annuity payments ranging from 0 (for a lump-sum of 4.7 years of benefits) to 16 percent (for a lump-sum of 7 years of benefits). Second panel reports average CEV of the simulated population across these different EOL transfer sizes.

Across all values for the end-of-life transfer, we observe a consistent pattern. While married

individuals experience small welfare losses due to the experiment, singles–particularly those without college education–experience large gains. The gains for the singles increase with the size of lump-sum payment while the CEV for married workers decreases for higher levels of EOL transfers. This finding highlights the role that redistribution from married to single individuals plays in the reform.

As we find significant heterogeneity in the value that different workers associate to combinations of EOL transfers and annuities, one natural question is what amount of EOL lump sum transfer would result in the highest average welfare change.

We explore this question by solving the model for different bundles of EOL transfer and annuity, and computing the average change in CEV across all workers. Second panel of figure 10 shows that CEV is locally maximized at a lump-sum value of roughly 5.5 years of benefits–the value of end-of-life transfer considered in our original universal transfer reform.

7 Robustness And Sensitivity Exercises

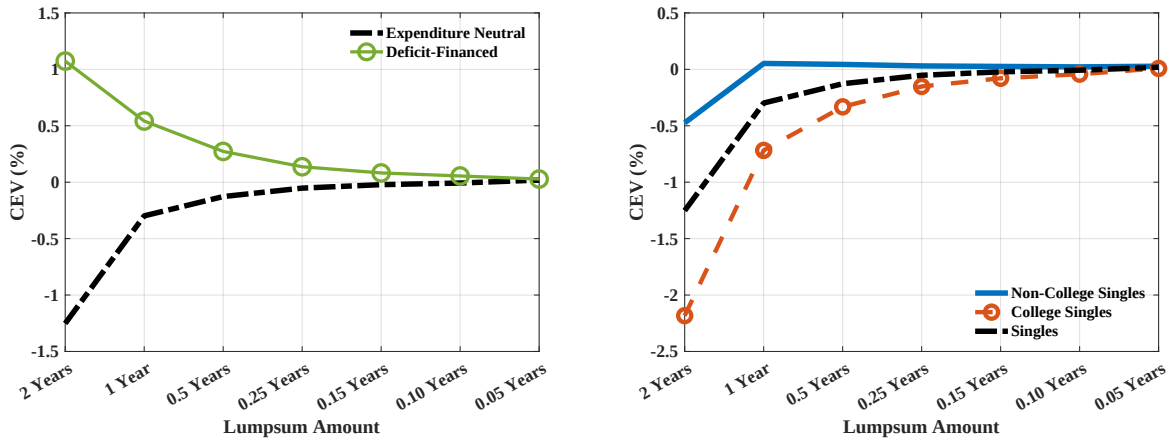
7.1 Robustness: Medical Expenditure Uncertainty

Note that in our benchmark model, bequest motives for different education-marital groups are identified from the rate of wealth decumulation at older ages. Given that the risk of extreme medical expenditure events at older ages could be partly driving why the elderly run down their wealth slowly (De Nardi et al., 2010, 2021), we estimate a model with an alternate specification for these expenditures. Following the literature (French and Jones, 2011), we estimate our structural model with the following process for medical expenditures:

$$\begin{aligned}
\log \rho_{mt}^{e,m} &= m(\mu_t, t, e, m) + \sigma(\mu_t, t, e, m) \times \varphi_t \\
\varphi_t &= \zeta_t^m + \epsilon_t^m, \quad \epsilon_t^m \sim N(0, \sigma_\epsilon^m) \\
\zeta_t^m &= \rho^m \zeta_{t-1}^m + \nu_t^m, \quad \nu_t^m \sim N(0, \sigma_\nu^m)
\end{aligned}$$

We find that our welfare results are not impacted qualitatively (somewhat stronger quantitatively) under this alternate modeling assumption.

Figure 12: Welfare Across Experiments: Non-Universal Transfer Reform



(a) (b)

Notes: Panel (a) shows the average CEV of singles across two sets of policy experiments 1) individuals are paid lump sum without corresponding decrease in annuity (deficit financed) and 2) individuals are paid lump sum with a corresponding decrease in annuity (expenditure neutral). Panel (b) presents CEV of all-singles, college-singles and non-college singles for a policy which pays out a lump sum worth some years of annual social security benefits at the end of life. The annuity benefits are reduced in each experiment to make lump sum payments revenue neutral and ranges up to 42 percent for a EOL transfer of two years of benefits.

7.2 Sensitivity: Singles Only Reform

Figure 12 shows results for this experiment featuring only single workers. In panel (b) we show how CEV varies for EOL transfers valued between 2 and 0.05 years. Unlike the positive relationship between the EOL transfer and welfare we highlighted for singles in the universal transfer reform of section 6, considering only singles makes the CEV decreasing in higher levels of the lump-sum — a result driven by the drop in the annuity needed to finance the larger lump-sum payments. Panel (a) demonstrates this point; when the end-of-life is debt-financed rather than financed through a drop in the annuity, we find a monotonic positive relationship between CEV and the size of the transfer established previously. However, when the reforms are required to be expenditure-neutral, the relationship between EOL transfers and CEV seems non-monotonic — welfare increases at lower levels of EOL transfers and decreases at higher levels. While singles value the liquidity provided by the EOL transfer, they also highly value the size of the annuity they receive. Moreover, the cost in terms of annuity reduction, is much larger for even a small EOL transfer.

8 Conclusion

Uncertainty about lifespan duration introduces a preference for end-of-life liquidity. This motive has implications for consumption and labor supply over the life cycle. To assess the quantitative importance of these motives, we study the way individuals respond to changes in social security entitlements. In particular, we examine the effects of a policy reform whereby individuals are allowed to cash in a small part of their entitlements upon death to satisfy a desire for end-of-life liquidity. This policy change induces large responses that we quantify in a dynamic model of consumption and labor supply.

Using a variety of data sources, we estimate the baseline model under the current SS system and show that it can rationalize household choices along all relevant margins, including the timing of benefit claims and late-life asset holdings. Our analysis suggests that the illiquid nature of the current Social Security system has a strong impact on the life cycle choices of program participants and distorts their life-cycle decisions.

To illustrate the magnitude of these distortions, we consider a policy that pays out part of the entitlements at the end of life, wherever that may occur while preserving the expected value of benefits. Bundling a guaranteed end-of-life entitlement and an annuity into the SS benefit significantly changes labor supply and saving patterns over the entire life cycle. Policy impacts are heterogeneous in the population and tend to be larger for single individuals with less education. The welfare changes induced by the policy are large and positive across marital and education groups and come close to being Pareto-improving ex-post.

By focusing on a moderate departure from the current system, we show that small tweaks that guarantee partial access to entitlements in case of premature death can have strong impacts and bring about significant welfare gains. Our work highlights how late-life liquidity needs that extend up to the time of death can shape behaviors much earlier in the life cycle and affect the value of existing policies. More generally, our findings suggest that end-of-life preferences for liquidity may exert a large influence on individual choices.

References

Abel, A. (1986). Capital accumulation with adverse selection and uncertain lifetimes. *Econometrica*, 54:1079–97.

Attanasio, O., Kitao, S., and Violante, G. L. (2007). Global demographic trends and social security reform. *Journal of Monetary Economics*, 54(1):144–198.

Auerbach, A. J. and Kotlikoff, L. J. (1985). Life insurance of the elderly: adequacy and determinants. Technical report, National Bureau of Economic Research.

Auerbach, A. J. and Kotlikoff, L. J. (1991). The adequacy of life insurance purchases. *Journal of financial Intermediation*, 1(3):215–241.

Bairoliya, N. and McKiernan, K. (2021). Revisiting retirement and social security claiming decisions. *Available at SSRN 3896031*.

Benabou, R. (2000). Unequal societies: Income distribution and the social contract. *American Economic Review*, 91(1):96–129.

Bernheim, B. D., Forni, L., Gokhale, J., and Kotlikoff, L. J. (2003). The mismatch between life insurance holdings and financial vulnerabilities: evidence from the health and retirement study. *American Economic Review*, 93(1):354–365.

Borella, M., De Nardi, M., and Yang, F. (2019). Are marriage-related taxes and social security benefits holding back female labor supply? *NBER Working Paper No. w26097*.

Casanova, M. (2010). Happy together: A structural model of couples' joint retirement choices. *Unpublished Manuscript, Department of Economics, UCLA, Los Angeles*.

Chambers, M., Schlagenhauf, D., and Young, E. R. (2011). Why aren't more families buying life insurance? *Center for Retirement Research at Boston College Working Paper*, (2011-7).

Conesa, J. C. and Krueger, D. (1999). Social security reform with heterogeneous agents. *Review of Economic dynamics*, 2(4):757–795.

Daruich, D. and Fernández, R. (2020). Universal basic income: A dynamic assessment. Technical report, National Bureau of Economic Research.

Davidoff, T., Brown, J. R., and Diamond, P. A. (2005). Annuities and individual welfare. *American Economic Review*, 95(5):1573–1590.

De Nardi, M. (2004). Wealth inequality and intergenerational links. *The Review of Economic Studies*, 71(3):743–768.

De Nardi, M., French, E., and Jones, J. B. (2010). Why do the elderly save? the role of medical expenses. *Journal of political economy*, 118(1):39–75.

De Nardi, M., French, E., and Jones, J. B. (2016). Savings after retirement: A survey. *Annual Review of Economics*, 8:177–204.

De Nardi, M., French, E., Jones, J. B., and McGee, R. (2021). Why do couples and singles save during retirement?

De Nardi, M., French, E., Jones, J. B., and McGee, R. (2021). Why do couples and singles save during retirement?

De Nardi, M., Imrohoroğlu, S., and Sargent, T. J. (1999). Projected us demographics and social security. *Review of Economic dynamics*, 2(3):575–615.

De Nardi, M., Jones, J., and French, E. (2010). Why do the elderly save? *Journal of Political Economy*, 118(1):39–75.

Dynan, K. E., Skinner, J., and Zeldes, S. P. (2002). The importance of bequests and life-cycle saving in capital accumulation: a new answer. *The American Economic Review*, 92(2):274–278.

French, E. (2005). The effects of health, wealth, and wages on labor supply and retirement behavior. *The Review of Economic Studies*, 72(2):395–427.

French, E., De Nardi, M., Jones, J. B., Baker, O., and Doctor, P. (2006). Right before the end: Asset decumulation at the end of life. *Economic Perspectives*, 30(3):2–13.

French, E. and Jones, J. B. (2011). The effects of health insurance and self-insurance on retirement behavior. *Econometrica*, 79(3):693–732.

Fuster, L., İmrohoroğlu, A., and İmrohoroğlu, S. (2003). A welfare analysis of social security in a dynastic framework. *International Economic Review*, 44(4):1247–1274.

Golosov, M., Troshkin, M., and Tsyvinski, A. (2016). Redistribution and social insurance. *American Economic Review*, 106(2):359–386.

Gourinchas, P.-O. and Parker, J. A. (2002). Consumption over the life cycle. *Econometrica*, 70(1):47–89.

Gruber, J. and Wise, D. (2005). Social security programs and retirement around the world: Fiscal implications, introduction and summary. Technical report, National Bureau of Economic Research.

Heathcote, J., Perri, F., and Violante, G. L. (2010a). Unequal we stand: An empirical analysis of economic inequality in the united states, 1967–2006. *Review of Economic dynamics*, 13(1):15–51.

Heathcote, J., Storesletten, K., and Violante, G. L. (2010b). The macroeconomic implications of rising wage inequality in the united states. *Journal of political economy*, 118(4):681–722.

Heiland, F. and Yin, N. (2014). Have we finally achieved actuarial fairness of social security retirement benefits and will it last?

Hong, J. H. and Ríos-Rull, J.-V. (2007). Social security, life insurance and annuities for families. *Journal of Monetary Economics*, 54(1):118–140.

Hong, J. H. and Rios-Rull, J.-V. (2012). Life insurance and household consumption. *American Economic Review*, 102(7):3701–3730.

Hosseini, R. (2008). A competitive model of annuity and life insurance with nonexclusive contracts. Technical report. mimeo.

Hosseini, R., Kopecky, K. A., and Zhao, K. (2022). The evolution of health over the life cycle. *Review of Economic Dynamics*, 45:237–263.

Hubbard, R. G. and Judd, K. L. (1987). Social security and individual welfare: Precautionary saving, borrowing constraints, and the payroll tax. *The American Economic Review*, pages 630–646.

Hubbard, R. G., Skinner, J., and Zeldes, S. P. (1995). Precautionary saving and social insurance. Technical report, National Bureau of Economic Research.

Huckfeldt, C. (2016). Understanding the scarring effect of recessions. *Working Paper*.

Huff Stevens, A. (1997). Persistent effects of job displacement: The importance of multiple job losses. *Journal of Labor Economics*, 15:165–188.

Huggett, M. and Parra, J. C. (2010). How well does the us social insurance system provide social insurance? *Journal of Political Economy*, 118(1):76–112.

İmrohoroglu, A., Imrohoroglu, S., and Joines, D. H. (1995). A life cycle analysis of social security. *Economic theory*, 6(1):83–114.

Jacobson, L. S., LaLonde, R. J., and Sullivan, D. G. (1993). Earnings losses of displaced workers. *The American Economic Review*, 83:685–709.

Jones, J. B., De Nardi, M., French, E., McGee, R., and Rodgers, R. (2020). Medical spending, bequests, and asset dynamics around the time of death.

Jones, J. B. and Li, Y. (2023). Social security reform with heterogeneous mortality. *Review of Economic Dynamics*, 48:320–344.

Kopczuk, W. and Lupton, J. P. (2007). To leave or not to leave: The distribution of bequest motives. *The Review of Economic Studies*, 74:207–235.

Kotlikoff, L. J., Smetters, K., and Walliser, J. (2007). Mitigating america's demographic dilemma by pre-funding social security. *Journal of monetary Economics*, 54(2):247–266.

Krueger, D. and Kubler, F. (2006). Pareto-improving social security reform when financial markets are incomplete!? *American Economic Review*, 96(3):737–755.

Krueger, D. and Perri, F. (2006). Does income inequality lead to consumption inequality? evidence and theory. *The Review of Economic Studies*, 73(1):163–193.

Krusell, P., Mukoyama, T., and Şahin, A. (2010). Labour-Market Matching with Precautionary Savings and Aggregate Fluctuations. *The Review of Economic Studies*, 77(4):1477–1507.

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

Lockwood, L. M. (2018). Incidental bequests and the choice to self-insure late-life risks. *American Economic Review*, 108(9):2513–2550.

Miller, R. and Bairoliya, N. (2021). Health, longevity, and welfare inequality of older americans. *The Review of Economics and Statistics*, pages 1–45.

Scheuer, F. and Slemrod, J. (2021). Taxing our wealth. *Journal of Economic Perspectives*, 35(1):207–30.

Stantcheva, S. (2017). Optimal taxation and human capital policies over the life cycle. *Journal of Political Economy*, 125(6):1931–1990.

Steuerle, C. E. and Smith, K. E. (2023). Social security & medicare lifetime benefits and taxes: 2022. Technical report, Urban Institute.

Yaari, M. E. (1965). Uncertain lifetime, life insurance, and the theory of the consumer. *The Review of Economic Studies*, 32(2):137–150.

Online Appendix

A Motivation

A1 Sample Selection

As discussed in the main text, the empirical facts use data from the Health and Retirement Study (HRS). The Health and Retirement Study is a longitudinal study of Americans over the age of 50. We use Waves 3-15 of 2020 version 1 in this work; these waves cover the years 1992 through 2020. Our final sample focuses on a cohort of men born between 1926 and 1970 who are over the age of 50 and for whom marital status and parental status are known. We also will be using two sub-samples when computing the statistics in Section 2: the bequest intention sub-sample and the life insurance and will sub-sample. The bequest intention sub-sample add the additional restriction of limiting the age to be between 50 and 70; the life insurance and will sub-sample limits the Waves used to 6-15, covering years 2002-2020. Table A1 details this sample selection.

Table A1: Stylized Facts Sample Selection

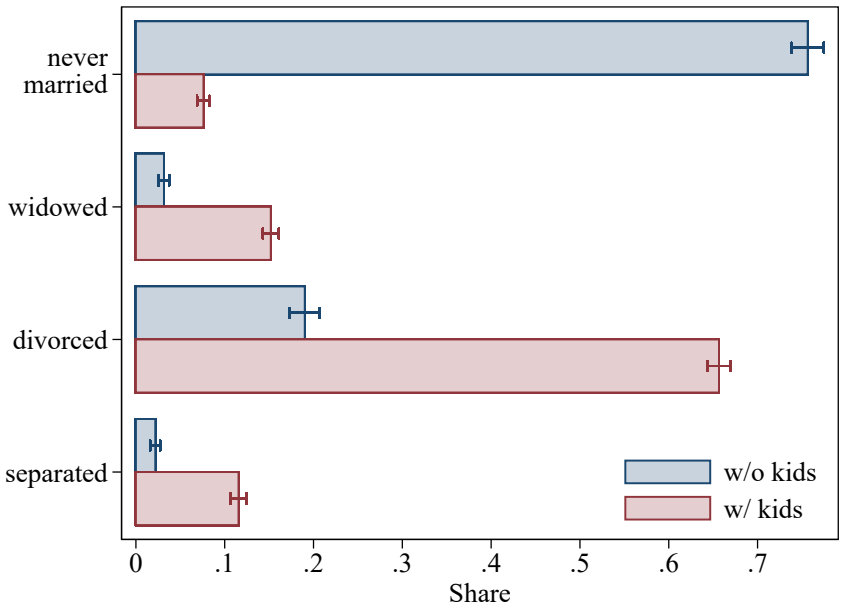
Main Sample	Observations
HRS 2020 version 1	636,180
Waves 3-15	551,356
Born 1926-1970	430,404
Men	193,206
Age 50+	87,400
Marital, parental status known	85,764
Bequest Intention Sub-Sample	85,764
Age 50-70	60,803
Life Insurance and Will Sub-Sample	85,764
Waves 6-15	68,927

We separate this sample into four groups based upon marital and parental status: single with no children, single with children, married with no children, married with children. Married individuals are those who are currently legally married (spouse present or absent) and those who are partnered; singles individuals are separated, widowed, divorced, or never married; an individual is classified as having children is he has at least one living child with whom he is currently in touch. Using current marital status means that it is possible that individuals who we classify as single has been previously married. Figure A1 shows how singles are distributed across more detailed marital status groups based upon whether they have children. We see that nearly 80 percent of singles without children have never been married while over 80 percent of singles with children are either

divorced or widowed. Additionally, using a measure of children based upon living children who are currently in touch may mean we classify an individual as not having children if a child is deceased or not in communication. We prefer this measure as it allows us to include step-children who would not be captured in a measure of biological children. Even so, our results are robust to using a measure of children that counts whether an individual has ever has a biological child born to him.

Within the main sample, 6.4% are single with no children, 15.9% are single with children, 3.1% are married with no children, and 74.6% are married with children.

Figure A1: Distributions of Detailed Marital Status Among Singles



Notes: Figure shows the distribution of detailed marital statuses among those who are classified as single, condition on either having children or not having children.

A2 Additional Results and Figures

In the main text, we present Figures 1 and 2 showing how the share of individuals who have bequest intentions and make choices to prepare. We also perform a Tukey pairwise mean comparison test to study whether there is a statistically significant difference between the means of each of the marital-parental status groups. These results are shown in Table A2. The table shows the point estimate for the difference in the mean of the outcome variable (bequest intention, presence of a valid will) for a given pair of marital-parental status groups in a given wealth quintile. A significant result for this point estimate means that we are able to reject the null hypothesis that the means are the same across groups; an insignificant results means that we can reject the null hypothesis that these means are the same.

Table A2: Results of Tukey-Kramer Pairwise Mean Comparison Test

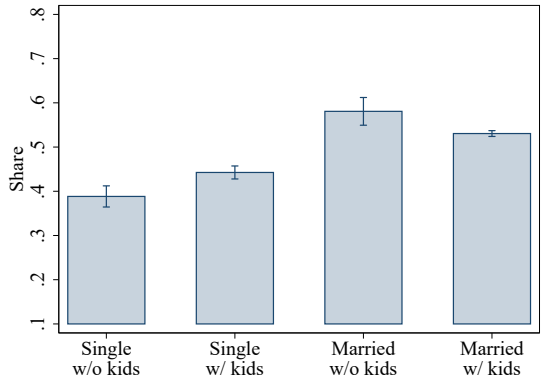
	Wealth Quintile				
	Q1	Q2	Q3	Q4	Q5
Bequest Intentions					
single w/ kids – single w/o kids	0.074^{***}	0.026	0.063	0.103^{***}	0.029
	(0.022)	(0.030)	(0.033)	(0.033)	(0.033)
married w/o kids – single w/o kids	0.029	-0.054	-0.023	0.041	-0.044
	(0.040)	(0.043)	(0.042)	(0.039)	(0.035)
married w/ kids – single w/o kids	0.101^{***}	-0.031	0.010	0.014	-0.012
	(0.020)	(0.028)	(0.028)	(0.028)	(0.028)
married w/o kids – single w/ kids	-0.045	-0.080	-0.085^{*}	-0.061	-0.073^{*}
	(0.037)	(0.037)	(0.036)	(0.034)	(0.029)
married w/ kids – single w/ kids	0.027	-0.057^{***}	-0.052^{**}	-0.089^{***}	-0.041
	(0.012)	(0.016)	(0.018)	(0.019)	(0.020)
married w/ kids – married w/o kids	0.072	0.023	0.033	-0.028	0.032
	(0.036)	(0.035)	(0.032)	(0.029)	(0.023)
Valid Will					
single w/ kids – single w/o kids	0.023	-0.025	0.066	0.030	0.009
	(0.019)	(0.030)	(0.036)	(0.037)	(0.040)
married w/o kids – single w/o kids	0.089^{*}	0.023	0.035	0.028	-0.033
	(0.036)	(0.043)	(0.045)	(0.043)	(0.041)
married w/ kids – single w/o kids	0.018	0.004	0.093^{**}	0.138^{***}	0.080^{*}
	(0.018)	(0.027)	(0.031)	(0.032)	(0.033)
married w/o kids – single w/ kids	0.066	0.049	-0.031	-0.002	-0.041
	(0.033)	(0.037)	(0.038)	(0.037)	(0.033)
married w/ kids – single w/ kids	-0.004	0.029	0.027	0.108^{***}	0.072^{**}
	(0.011)	(0.016)	(0.020)	(0.022)	(0.023)
married w/ kids – married w/o kids	-0.070	-0.020	0.058	0.110^{***}	0.113^{***}
	(0.032)	(0.035)	(0.034)	(0.031)	(0.025)
Life Insurance					
single w/ kids – single w/o kids	0.117^{***}	0.110^{***}	0.162^{***}	0.092^{**}	0.192^{***}
	(0.026)	(0.030)	(0.029)	(0.030)	(0.039)
married w/o kids – single w/o kids	0.190^{***}	0.065	0.178^{***}	0.132^{***}	0.226^{***}
	(0.049)	(0.043)	(0.036)	(0.035)	(0.039)
married w/ kids – single w/o kids	0.187^{***}	0.189^{***}	0.267^{***}	0.262^{***}	0.331^{***}
	(0.024)	(0.027)	(0.025)	(0.026)	(0.032)
married w/o kids – single w/ kids	0.073	-0.045	0.015	0.040	0.034
	(0.045)	(0.037)	(0.031)	(0.030)	(0.032)
married w/ kids – single w/ kids	0.069^{***}	0.079^{***}	0.105^{***}	0.170^{***}	0.139^{***}
	(0.015)	(0.016)	(0.016)	(0.018)	(0.023)
married w/ kids – married w/o kids	-0.004	0.125^{***}	0.090^{***}	0.130^{***}	0.105^{***}
	(0.044)	(0.034)	(0.027)	(0.025)	(0.024)

Notes: Table shows the difference in the means of the outcome variable (bequest intention, presence of a valid will) for each pair of marital-parental status groups and each wealth quintile. Null hypothesis of the Tukey-Kramer test is that the means are equivalent, $H_0 = \mu_1 - \mu_2 = 0$, for any two groups 1 and 2. Standard errors in parentheses. *p< 0.1, **p< 0.05, ***p< 0.01

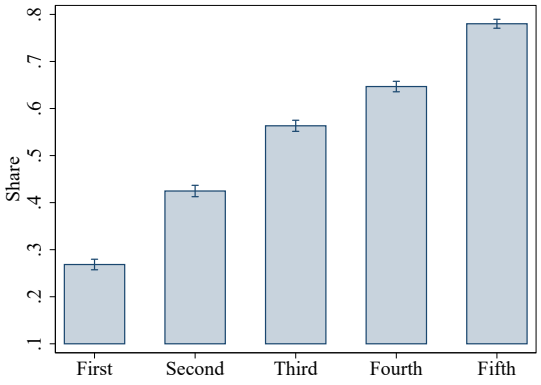
Additionally, in the main text, we report the bequest intentions and writing of a legal will by marital-parental status conditional on wealth quintile. Here (Figures A2 - A4) we show the results if we do not condition on wealth. These figures highlight the importance of wealth in interpreting the importance of marital-parental status in these choices. In each case, Panel (a) which shows the variation by marital-parental group seems to show the importance of different family structures. However, Panel (b) of all figures show an even stronger relationship by wealth. As shown in the main text, once we control for wealth quintile, there is little variation by family structure.

Figure A2: Bequest Intentions by Marital Status, Wealth

(a) Marital-Parental Status



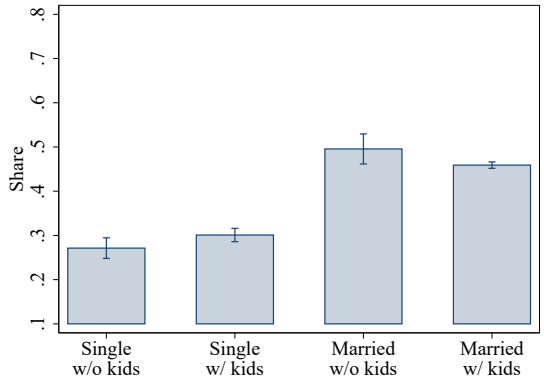
(b) Wealth Quintile



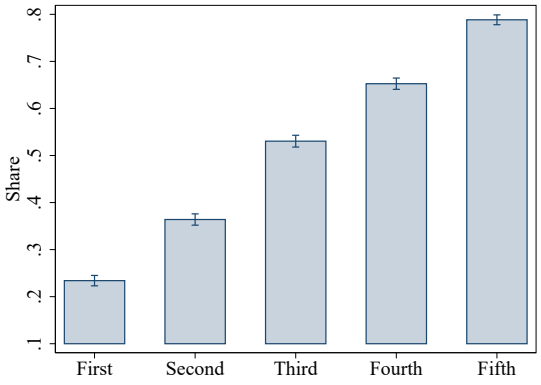
Notes: Figure shows the share of workers age 50-70 of a given marital-parental status [Panel (a)] or wealth quintile [Panel (b)] who report over a 95 percent probability of leaving a bequest of any amount. 95 percent confidence intervals shown.

Figure A3: Valid Will by Marital Status, Wealth

(a) Marital-Parental Status



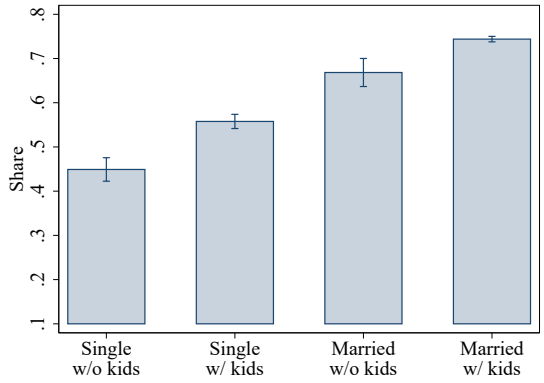
(b) Wealth Quintile



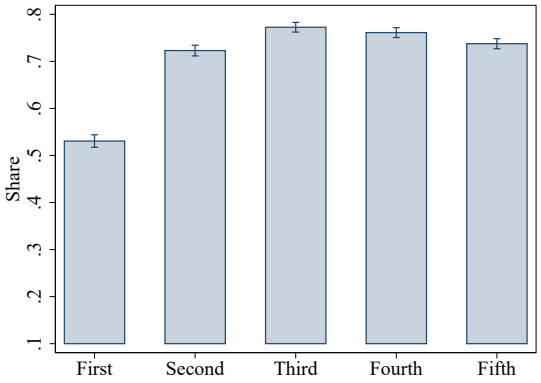
Notes: Figure shows the share of workers age 50+ of a given marital-parental status [Panel (a)] or wealth quintile [Panel (b)] who report they have a written and witnessed will or trust. 95 percent confidence intervals shown.

Figure A4: Life Insurance Coverage by Marital Status, Wealth

(a) Marital-Parental Status



(b) Wealth Quintile



Notes: Figure shows the share of workers age 50+ of a given marital-parental status [Panel (a)] or wealth quintile [Panel (b)] who report currently holding at least one life insurance policy (not including burial insurance). 95 percent confidence intervals shown.

B Robustness: Medical Expenditure Uncertainty

Out-of-pocket (OOP) medical expenditures are measured from years 1992-2018 of the Health and Retirement Study (HRS) and include spending on hospital and nursing costs; doctor, dentist, and outpatient surgery costs; average monthly prescription drug costs; and health health care and special facilities costs plus any expenses incurred in the final year of life. Expenses covered by insurance (public or private) are not included in this measure. As HRS is collected every two years, medical expenditures variables are divided by 2 to obtain an average annual cost.[23]

To estimate these OOP medical expenditures, we follow the two-stage procedure in De Nardi et al. (2010); French and Jones (2011); Borella et al. (2019). We first run the following fixed-effect regression on a sample of men aged 70 and above:

$$\ln(m_{it}) = f_i + X_{it}^{'}\beta_{it} + u_{it} \tag{A1}$$

where m_{it} is the measure of these out-of-pocket medical expenditure, f_i is the individual-specific fixed effect, and X_{it} in a set of explanatory variables including a fourth-order polynomial in age, a dummy variable indicating an individual is in poor health and this indicator interacted with a quadratic polynomial in age, and a dummy variable taking a value of 1 interacted with a quadratic polynomial in age. Interactions were only included when they were statistically significant. Notably, we investigated including education variables and its interaction with marital status and age but these were not statistically different from 0 in this first stage. In the second stage, the residuals and fixed effects from this equation are regressed on dummies for marital status, education level (and their interaction, and birth-year cohort. Table A3 shows these first and second stage regression results.

In the final step, the squared residuals are regressed on a quadratic polynomial in age, a third-

[23] In Wave 2A, the reference period for medical spending questions was the previous 12 months. Therefore, in this wave, we do not divide by two.

Table A3: Estimated Coefficients from Logarithm of Out-of-Pocket Medical Expenditures

First Stage	Coefficient	(std. err.)
Age	-2.438	(7.328)
$Age^2/10^2$	3.793	(13.56)
$Age^3/10^3$	-2.491	(11.11)
$Age^4/10^4$	0.585	(3.405)
bad health	5.681**	(2.836)
bad health*Age	-0.152**	(0.0714)
bad health*$Age^2/10^2$	0.104**	(0.0448)
married*$Age^3/10^3$	0.0288***	(0.00478)
married*$Age^4/10^4$	-0.0218***	(0.00579)
constant	62.15	(148.1)
Second Stage		
married	-0.113***	(0.0324)
college	0.197***	(0.0503)
married*college	0.0254	(0.0543)
born 1941-1945	-0.0452	(0.0423)
born 1936-1940	0.00842	(0.0347)
born 1926-1930	0.0153	(0.0351)
born 1921-1925	-0.204***	(0.0359)
born 1916-1920	-0.321***	(0.0405)
born 1900-1915	-0.571***	(0.0406)
constant	0.125***	(0.0372)
N	36680	
R^2 first stage	0.051	
R^2 second stage	0.027	

Source: Health and Retirement Study, authors' calculation

Notes: First Stage is fixed effects, Second Stage is OLS

Standard errors in parentheses, clustered at the individual level

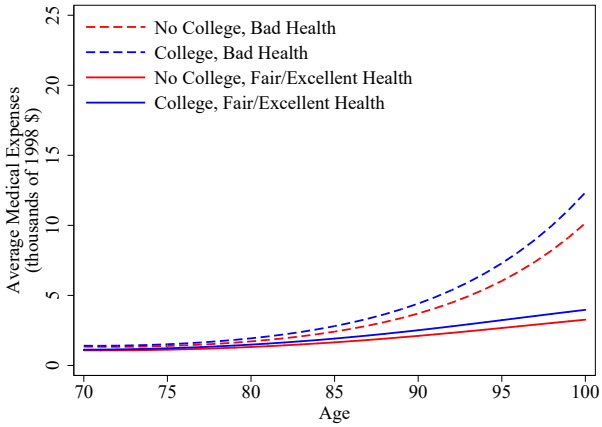
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

order polynomial in age interacted with bad health, a third order polynomial in age interacted with education, a dummy for being married, an interaction between marital status and education, and cohort dummies to compute the variance. The expected medical expenditures by age, education, marital status, and health are computed using the deterministic profiles from the first two stages plus half of the variance from this final step. These expected expenditures are shown in Figure A5.

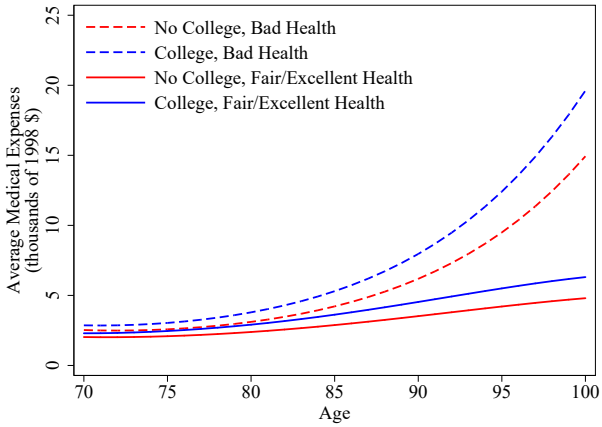
C Additional Figures and Tables

Figure A5: Expected Out-of-Pocket Medical Expenditures

(a) Single



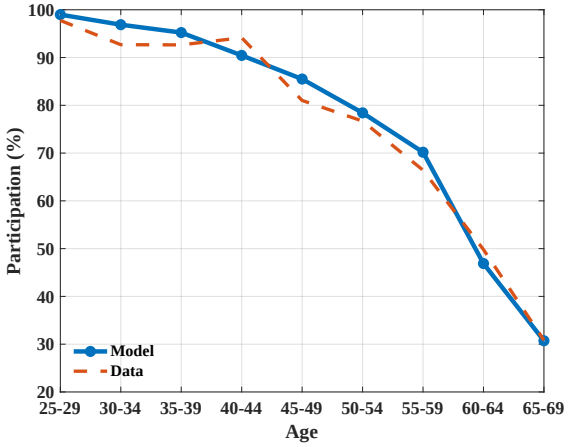
(b) Married



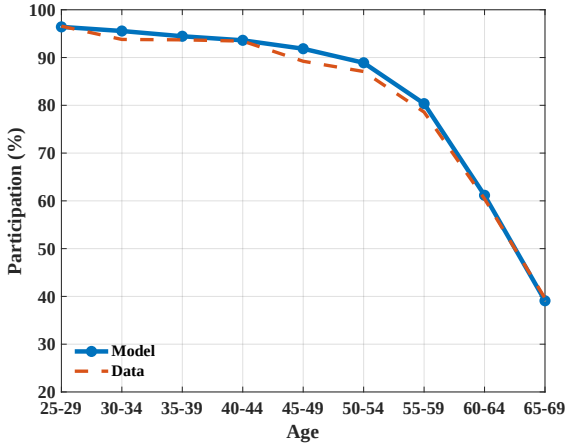
Source: Health and Retirement Study, authors' calculations

Figure A6: Model Fit: Labor force participation rates

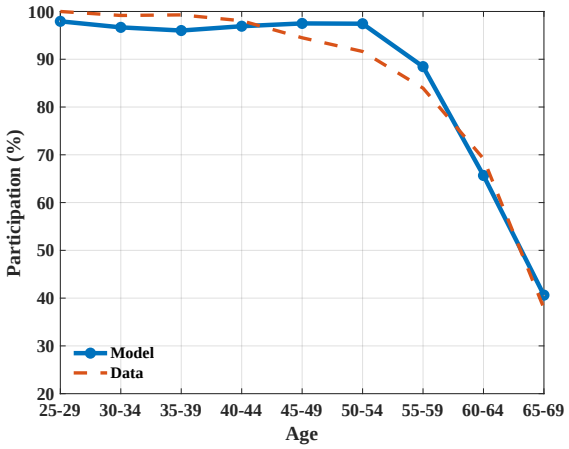
(a) Non-College, Non Married



(b) College, Non Married



(c) Non-College, Married



(d) College, Married

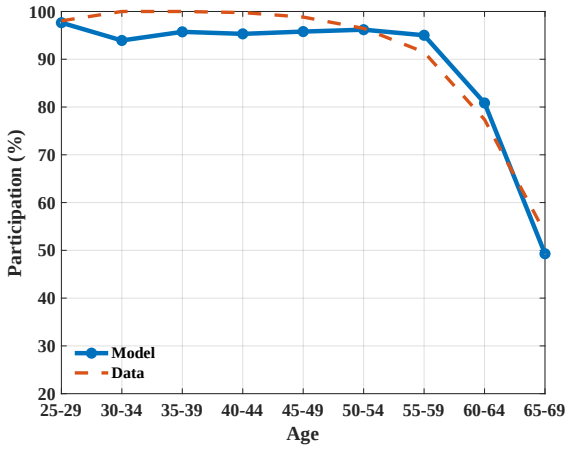
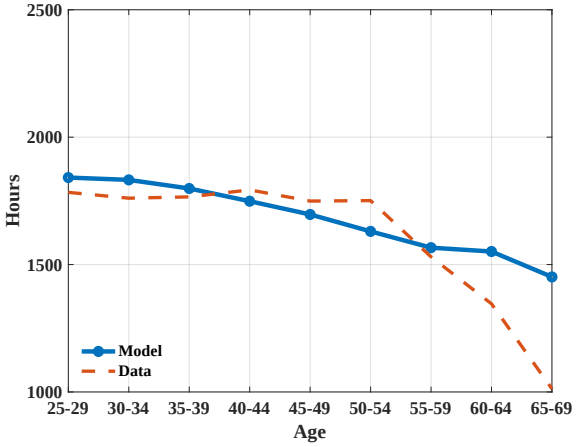
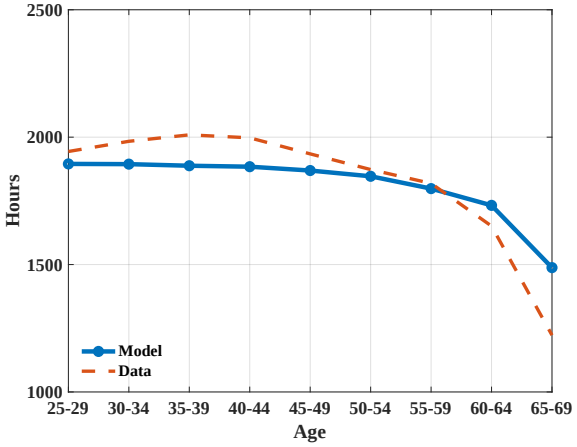


Figure A7: Model Fit: Hours Worked

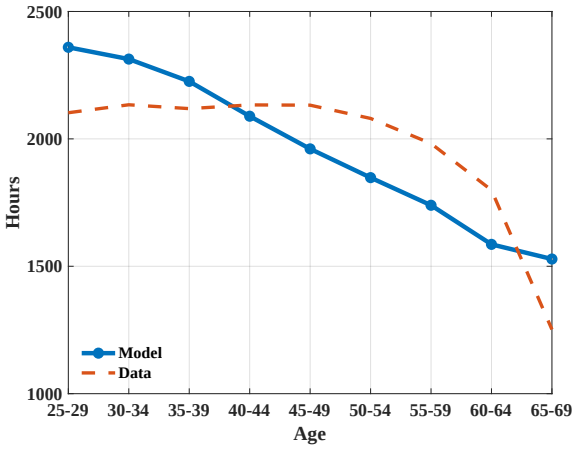
(a) Non-College, Non Married



(b) College, Non Married



(c) Non-College, Married



(d) College, Married

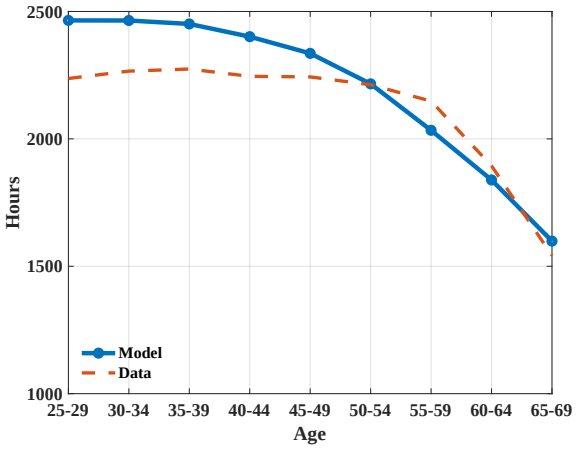
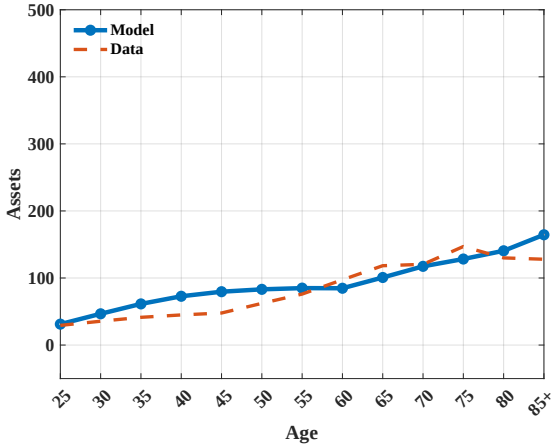
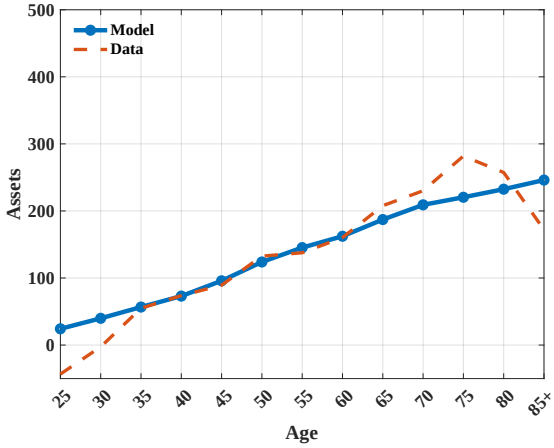


Figure A8: Model Fit: Wealth

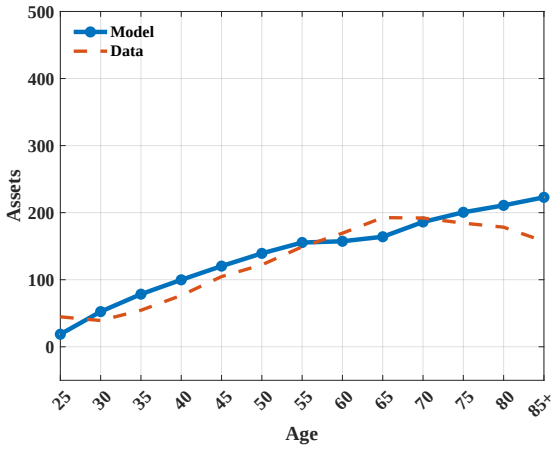
(a) Non-College, Non Married



(b) College, Non Married



(c) Non-College, Married



(d) College, Married

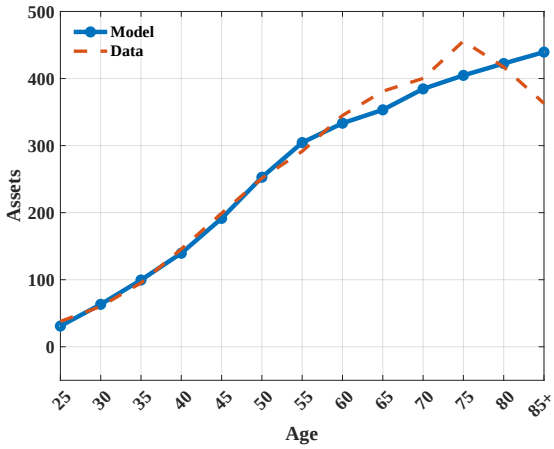


Figure A9: Model Fit: SS Claiming Behavior

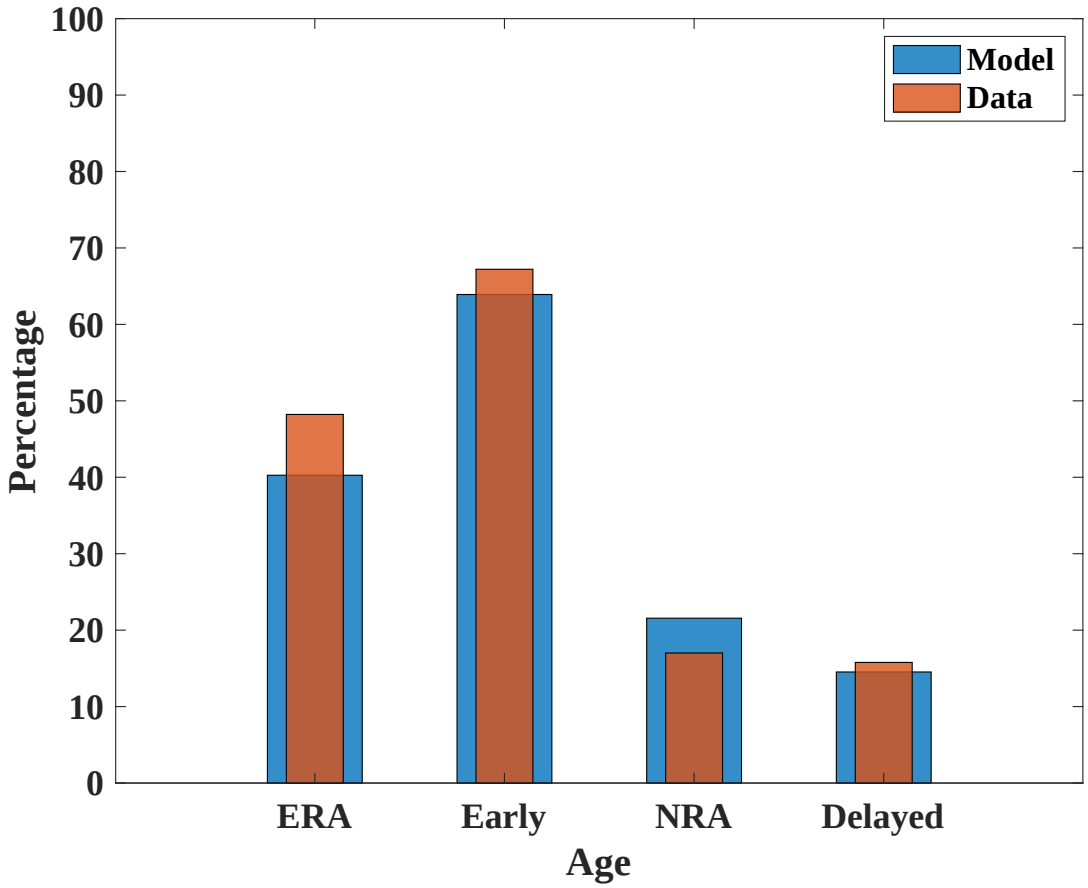
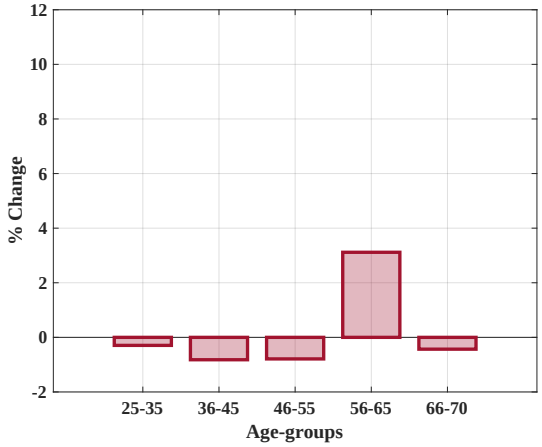
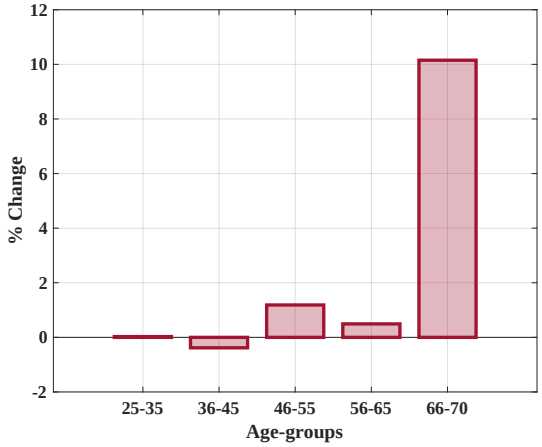


Figure A10: Life cycle Hours: Universal Transfer Reform

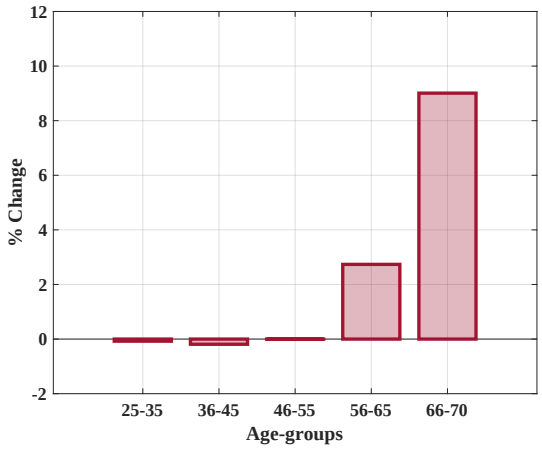
(a) Non-College, Non Married



(b) College, Non Married



(c) Non-College, Married



(d) College, Married

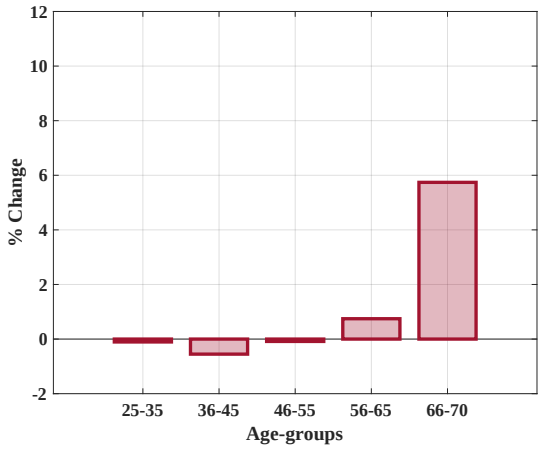
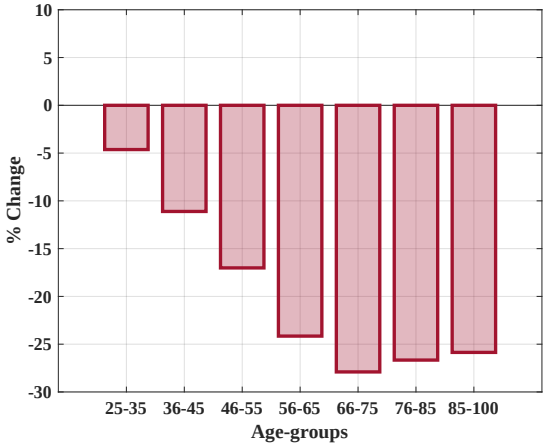
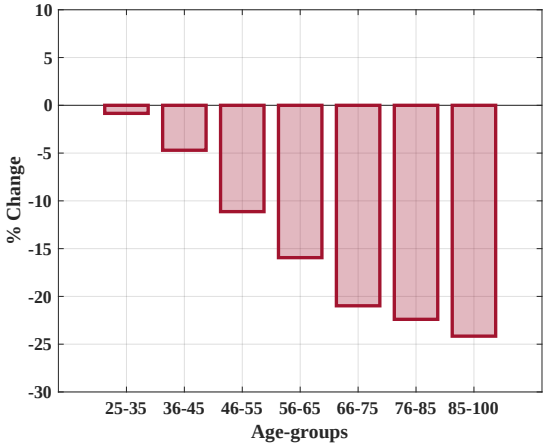


Figure A11: Life cycle Assets: Universal Transfer Reform

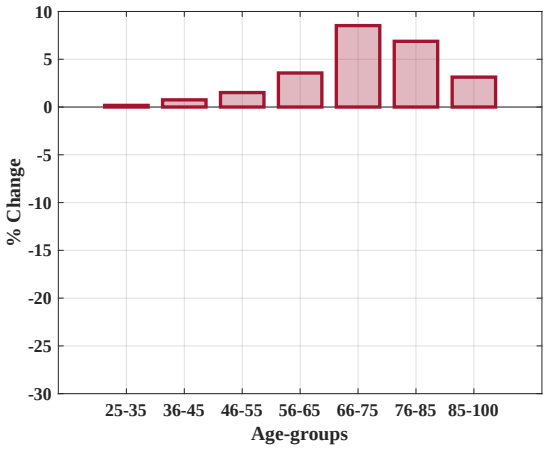
(a) Non-College, Non Married



(b) College, Non Married



(c) Non-College, Married



(d) College, Married

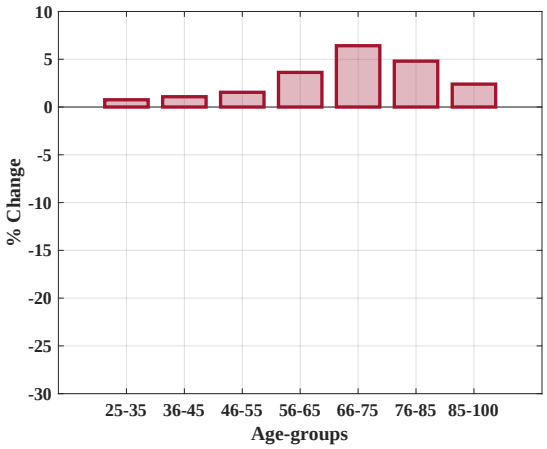
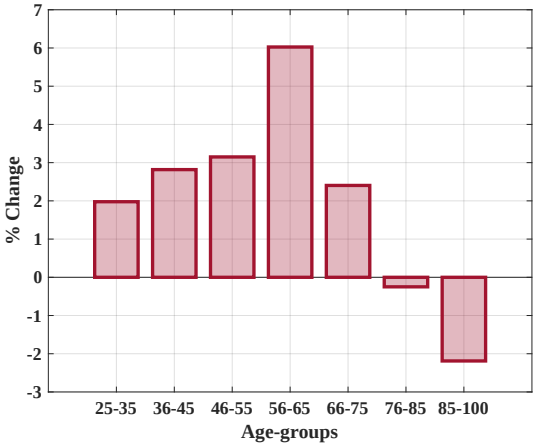
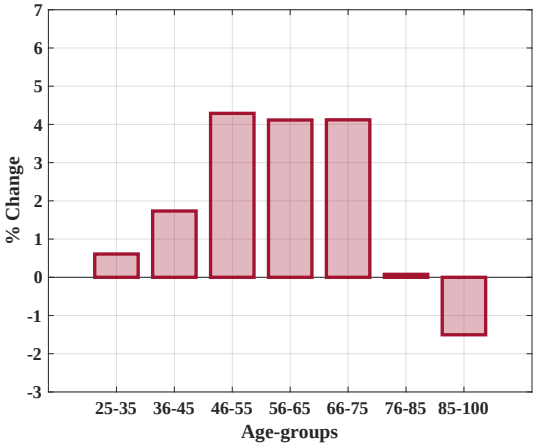


Figure A12: Life cycle Consumption: Universal Transfer Reform

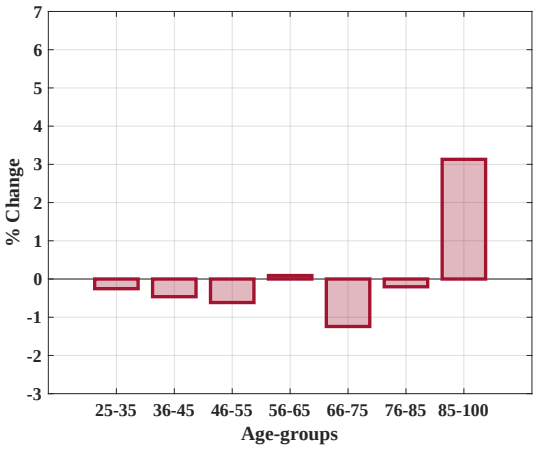
(a) Non-College, Non Married



(b) College, Non Married



(c) Non-College, Married



(d) College, Married

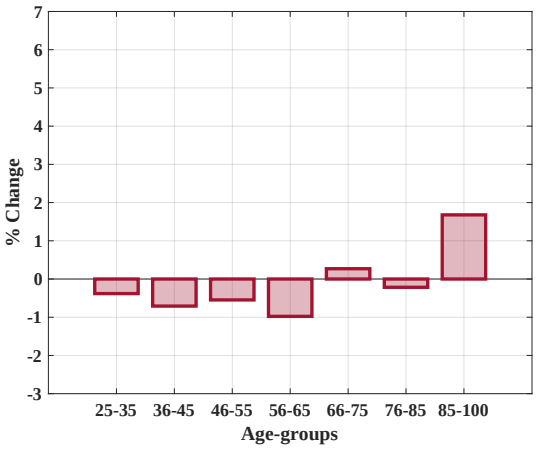
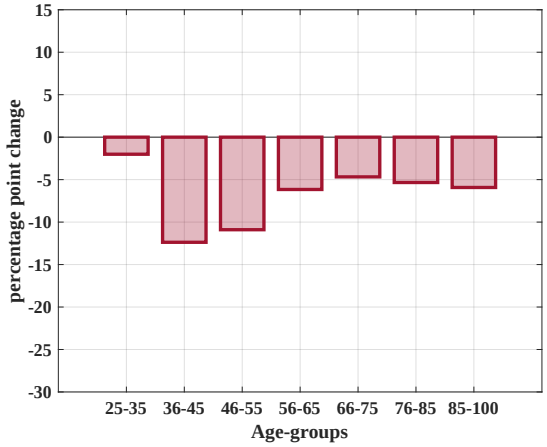
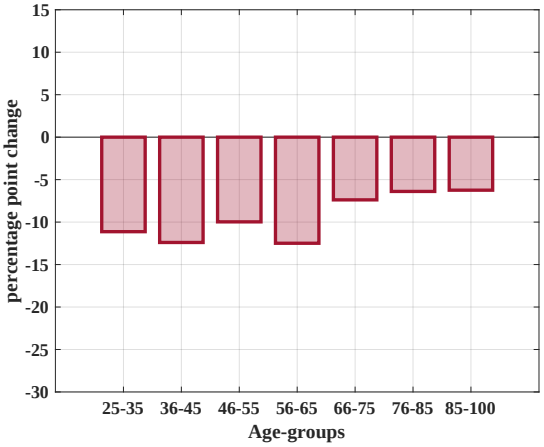


Figure A13: Life cycle LI Holdings: Universal Transfer Reform

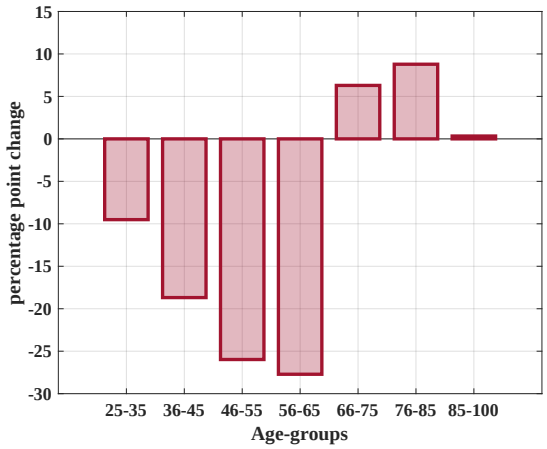
(a) Non-College, Non Married



(b) College, Non Married



(c) Non-College, Married



(d) College, Married

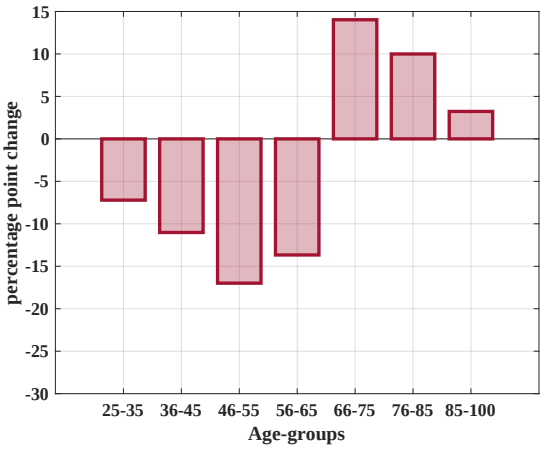
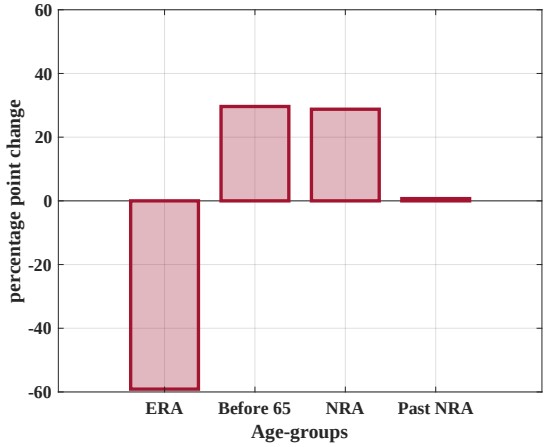
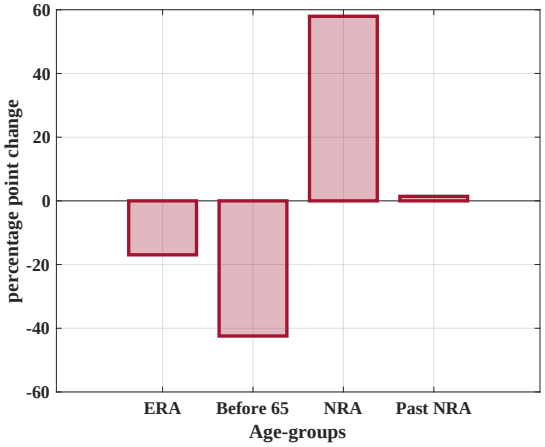


Figure A14: Social Security Claiming Behavior: Universal Transfer Reform

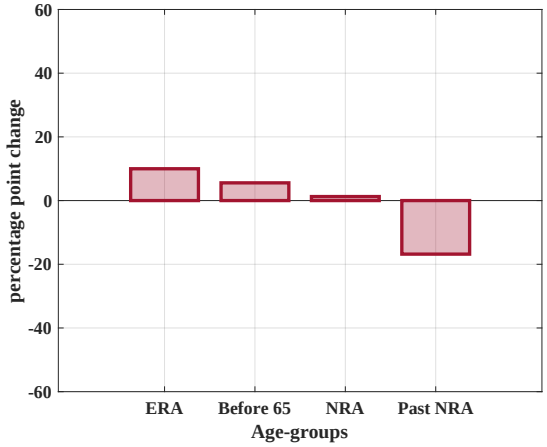
(a) Non-College, Non Married



(b) College, Non Married



(c) Non-College, Married



(d) College, Married

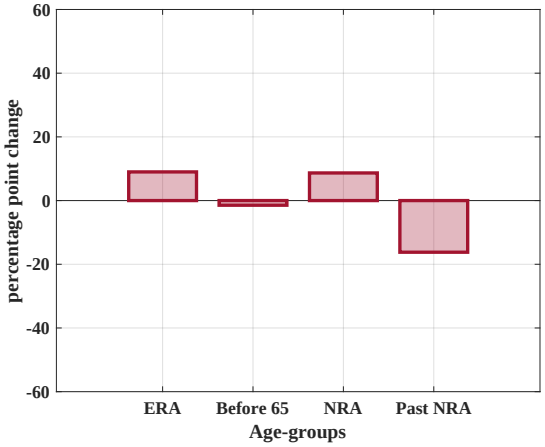
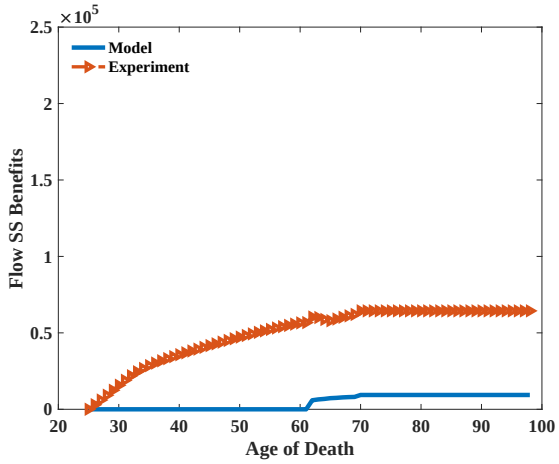
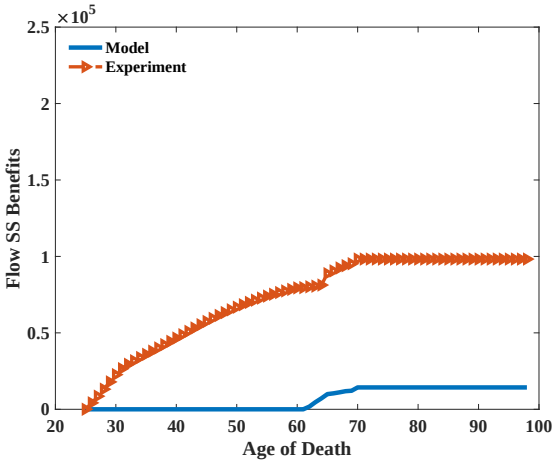


Figure A15: Flow SS Benefits by Age: Universal Transfer Reform

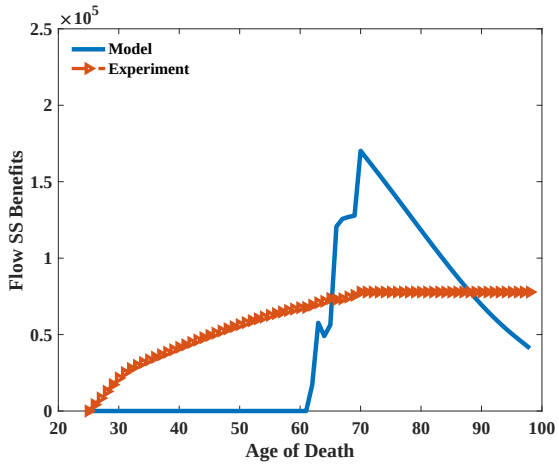
(a) Non-College, Non Married



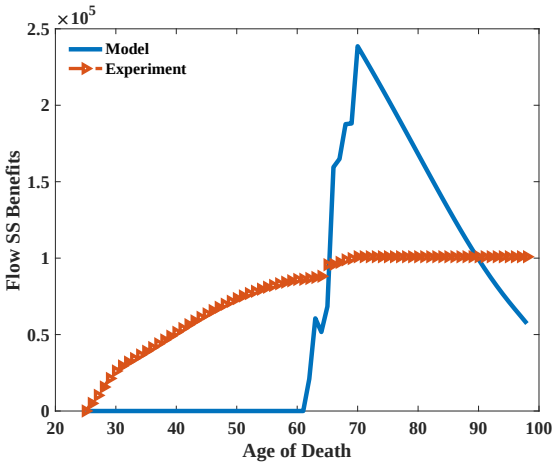
(b) College, Non Married



(c) Non-College, Married

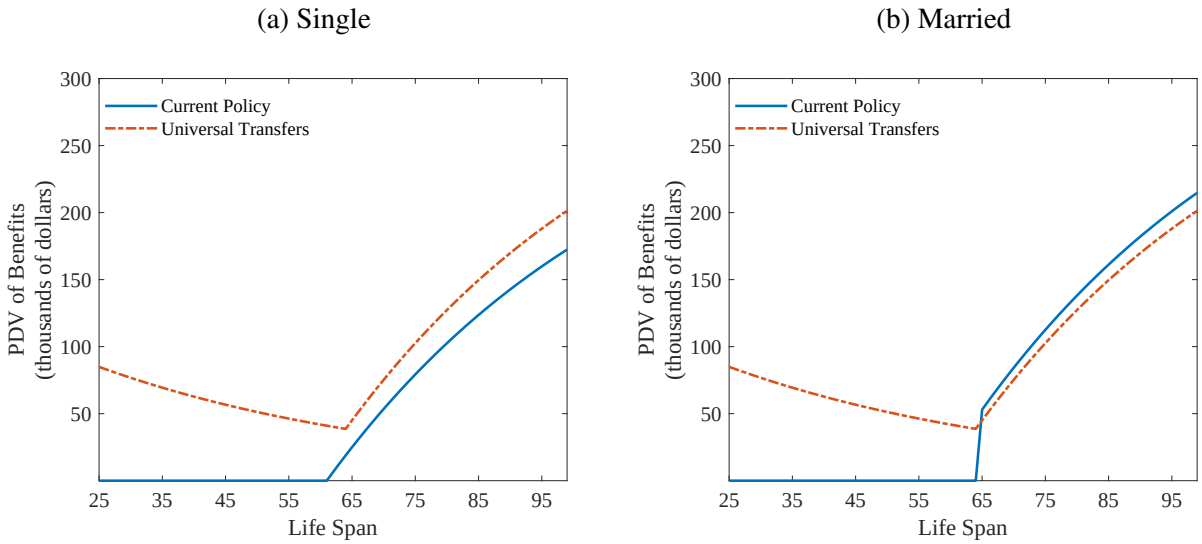


(d) College, Married



Notes: The figure plots the size of annual Social Security benefits by age, conditional on dying at the end of that period. In the benchmark, the flow is zero until upon claiming. In the counterfactual experiment, individuals receive a lump sum upon dying which is 5.5 times the PIA at the dying age.

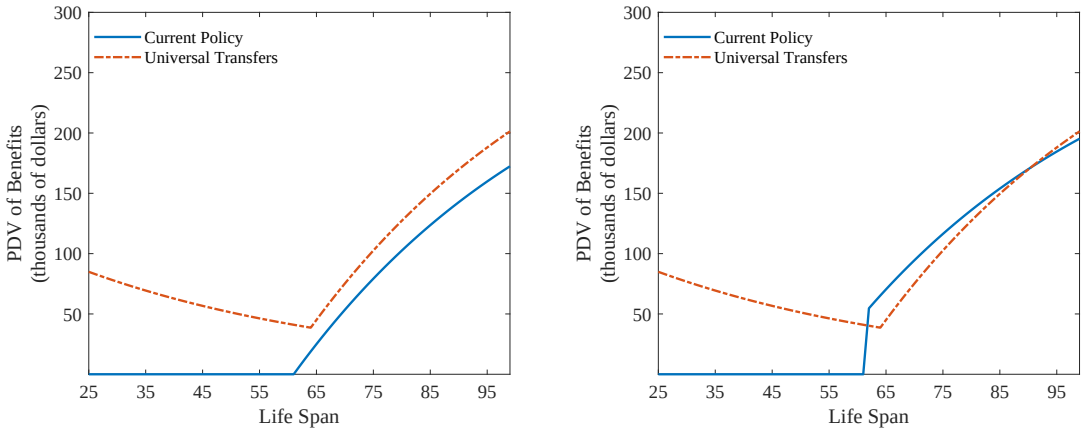
Figure A16: Present Value of SS Benefits Received: Current Policy and Universal Transfer Reform



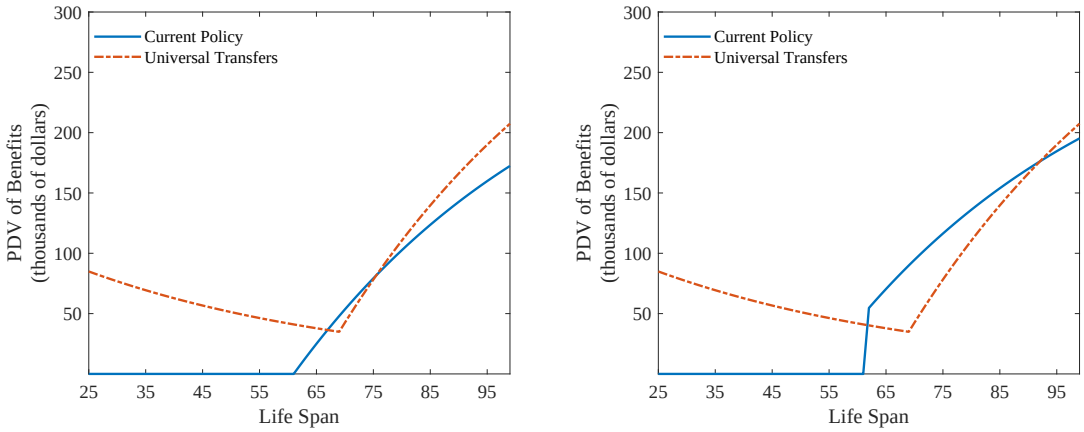
Notes: The figure plots the present value at age 25 of SS benefits received (old-age + survivors benefits) by lifespan. This calculation is done for a fixed AIME of $52,000 annually and discounted at a rate of 2%. Singles workers are assumed to claim benefits at age 62 in the current policy and delay claims to age 65 in the Universal Survivors Benefits experiment; married workers are assumed to claim benefits at age 65 and not change claiming decisions due to the reform. In the current policy, married couples receive Survivors Benefits equal to 6 years of PIA if they die after claiming. In the Universal Survivors Benefits experiment, all workers are eligible to leave Survivors Benefits equal to 5 years of PIA for any premature death and annual old-age benefits are decreased by 0.5%.

Figure A18: Present Value of SS Benefits: Alternative Social Security Claiming Ages

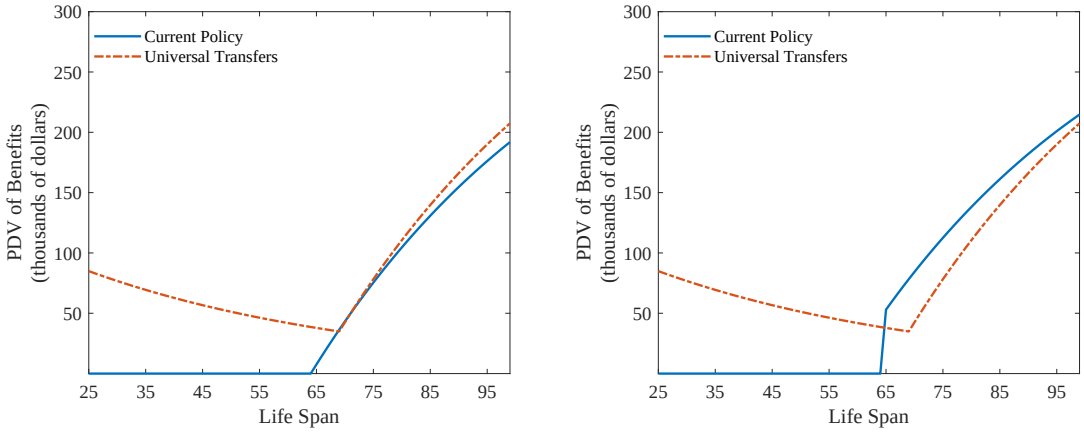
(a) Delay SS Claiming from Age 62 to Age 65



(b) Delay SS Claiming from Age 62 to Age 70



(c) Delay SS Claiming from Age 65 to Age 70



Notes: The figure plots the present value at age 25 of SS benefits received (old-age + survivors benefits) by lifespan. This calculation is done for a fixed AIME of $52,000 annually and discounted at a rate of 2%. In the current policy, married couples receive Survivors Benefits equal to 6 years of PIA if they die after claiming. In the Universal Survivors Benefits experiment, all workers are eligible to leave Survivors Benefits equal to 5 years of PIA for any premature death and annual old-age benefits are decreased by 0.5%.

D Supplemental Appendix

S1 Health Transitions and Survival Probabilities

We estimate health transitions using Household Component of Medical Expenditure Panel Survey on a population between the ages of 20 and 90. The coefficient estimates are provided in Table S1. We estimate the survival probabilities from MEPS by running an ordered probit regression of death indicator on self rated health interacted with a third order age polynomial as well as interactions of college and married indicators on the same population. Figure S1 provides the estimated survival probabilities.

Table S1: Ordered Probit Estimates of Health Transitions

	Coefficient Estimates
Health Lag=2	0.924***
Health Lag=3	2.118***
Age	0.0200***
$Age^2/10^2$	-0.00939
$Age^3/10^3$	-0.000500
College=1	-0.329***
cut1	0.635***
cut2	2.613***
Observations	188,226
Pseudo R^2	0.1622

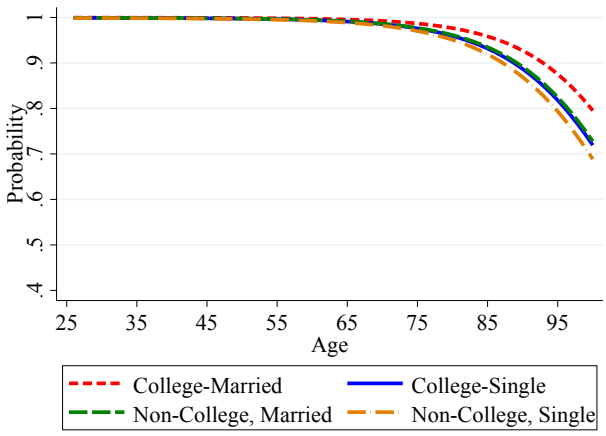
$^{*}\ p < 0.05$, $^{**}\ p < 0.01$, $^{***}\ p < 0.001$

S2 Medical Expenditures

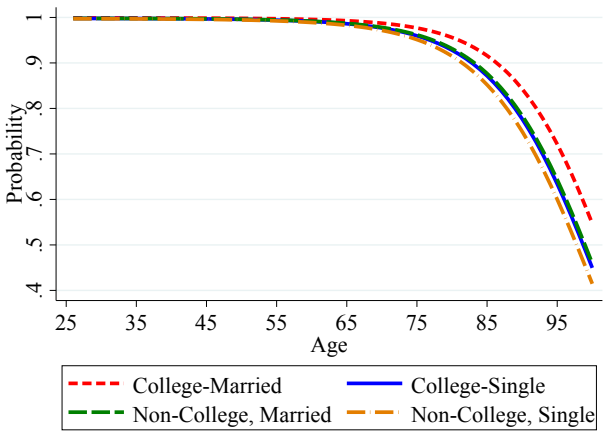
We estimate annual out-of-pocket medical expenditures from the Health and Retirement Study data using all years from 1992-2018, following the same approach as detailed in Borella et al. (2019). We use their approach to estimate medical expenditures for each education and marital groups. We allow these expenditures to vary health. However, using the approach in Borella et al. (2019), we allow expenditures to vary by excellent/good and poor health only. In other words, we combine the first two groups into one. There are two key details in the estimation of these profiles from the data. First, it is important to include medical expenses in the last year of life. This comes from the HRS exit interviews. Second, expenses covered by public or private insurance are not included in the measure, because they are not directly incurred by the individual. Figure S2 reports the estimated expenditures for male household heads, born between 1931-1935, by age and health, for different education and marital status.

Figure S1: Survival Probabilities by Health, Education and Marital Status

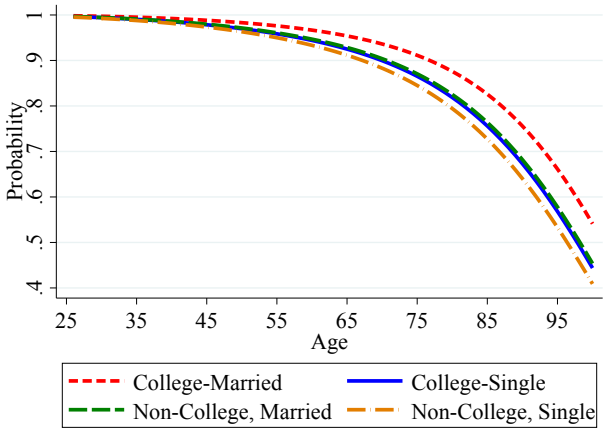
(a) Excellent Health



(b) Good Health



(c) Fair Health



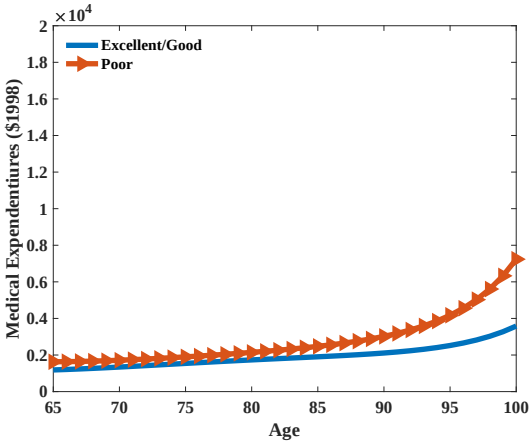
S3 Family Structure

Family structure in the model differs by education type. For the cohort of men born between 1931 and 1935, the share of the population in a married or co-habitating relationship is roughly stable over the life cycle. For this reason, we keep marriage a fixed state in the model, determined at age 25, based on the initial condition draws from the data and the married share is set to about 60 percent for non-college educated individuals and roughly 76 percent for college educated men.

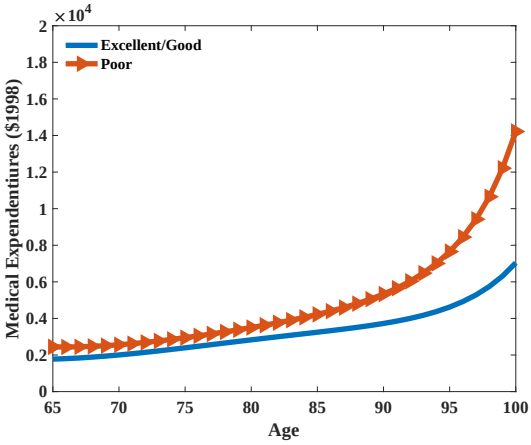
Figure S3b shows the measure of children used in the model to estimate the household consumption equivalence scale. The figure shows how the number of children living in the household varies across the life cycle, and by education of the household head. This measure takes into account all children ages 17 and under who are in the household at any point in time. The number of children is hump-shaped over the life cycle; the number of children peaks between the ages of 30 and 40 and then declines to 0 by age 60. Second, the profiles differ across education. Those without a college degree peak at around 1.2 children while those with a college degree peak at slightly above 1.5 children.

Figure S2: Annual Medical Expenditures by Health, Education and Marriage
Men born in 1931-1935

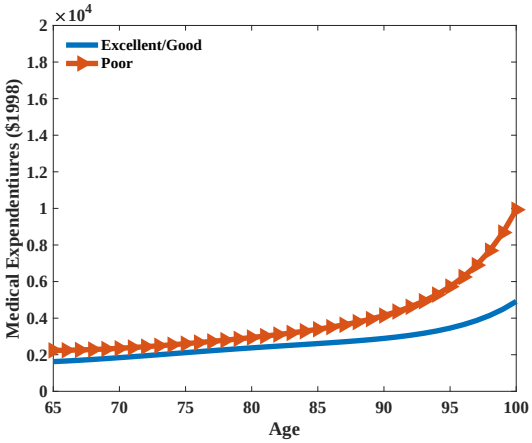
(a) Non-College, Non-Married



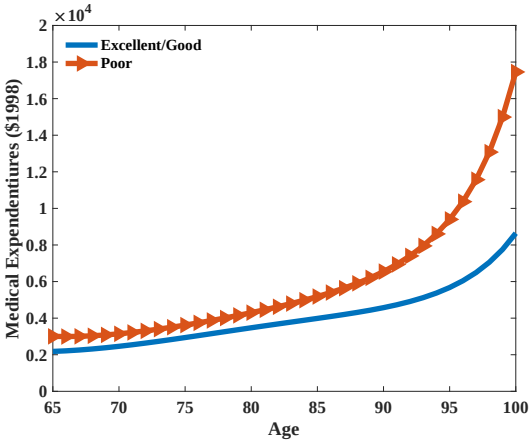
(b) College, Non-Married



(c) Non-College, Married



(d) College, Married

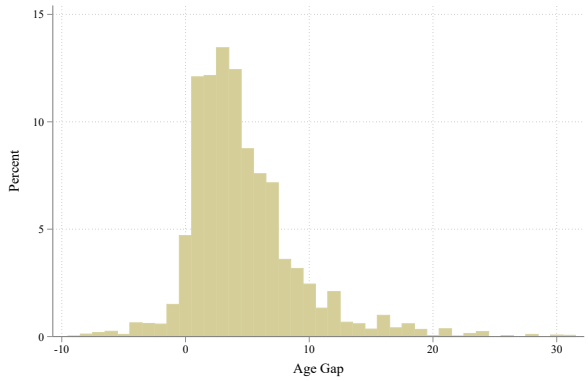


Married couples have access to spousal benefits through the Social Security system; these benefits depend not only on the age of the worker but also on the age of the spouse. Therefore, the gap between the ages of the spouses is very important. Figure S3a shows the distribution of this age gap for married couples born between 1926 and 1940. The average gap for this group is roughly 4 years.

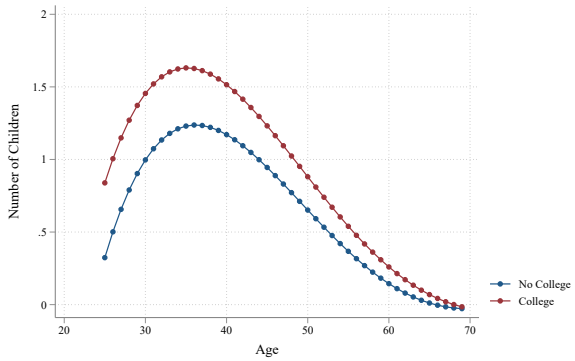
The distribution shows 95 percent of the married couples have a positive age gap – meaning the male head is often older than their spouse. However, this age gap varies largely across the distribution. To represent this distribution in a computationally feasible way, we include four age gaps (0 years, 1 year, 4 years, and 8 years) in the model. This allows us to capture that while most couples have age gaps between 0 and 4 years, there are also many couples with large age gaps. Estimated shares show that 8.7 percent of married couples have no age gap, 26.2 percent have an age gap of one year, 46.1 percent have an age gap of four years, and 19 percent have an age gap between spouses of eight years.

Figure S3: Family Structure

(a) Distribution of Spousal Age Gap



(b) Number of Children



S4 Spousal Income

We include spousal income in the budget constraint of the married worker since we model only male household heads. Because of the high fraction of married household heads, the estimation of spousal income is important to understand the budget constraints faced by individuals. We first estimate how spousal income varies based upon characteristics of the head of household:

$$y^s_{it} = X'_{it}\beta + \varepsilon_{it} \tag{A2}$$

where X'_{it} is a vector of control variables including a fourth order polynomial in the age of the household head, and indicator for whether the head of household is in poor health (both in levels and interacted with age), an indicator for whether the head of household attended college (both in levels and interacted with age), and the labor income of the head of household (in thousands). This regression is run for a sample of married individuals born between 1926 and 1940. We then use the estimated coefficients to impute spousal income in the model.

$$\hat{y}^s_{it} = X'_{it}\hat{\beta} \tag{A3}$$

By estimating how spousal income varies based on characteristics of the household head, we capture impact of assortative matching and differing probabilities of marriage across education levels and health status. The estimated coefficients are included in Table S2.

S5 Wages

Data on wages is taken from PSID. In PSID, the hourly wage is calculated as annual earnings divided the annual hours worked. Observations less than 50 percent of the minimum wage or over $500 (1998 dollars) are dropped to control for minimum wage and high wage outliers. The sample is increased to include all workers born between 1901 and 2000 in order to have sufficient observations through the life cycle. Because wages are only observed for those who participate in the labor market, we use a Heckman selection model to estimate hourly wage profiles.

In the first step, a probit regression estimates the probability of labor market participation and

Table S2: Coefficients of Spousal Income Regression

	coefficient	standard error
age	-1.581	(1.369)
age*age/(10^2)	8.024	(5.730)
age*age*age/(10^4)	-12.265	(7.859)
age*age*age*age/(10^6)	5.900	(3.540)
college	17.071**	(7.951)
college*age/(10^2)	-19.268**	(9.199)
poor health	4.842	(23.379)
poor health*age/(10^2)	-2.369	(28.361)
labor income (thousands)/(10^2)	9.842*	(5.211)
N	59,491	
R^2	0.165	

Source: Panel Study of Income Dynamics, authors' calculations
Notes: Dependent variable is labor income received by spouse (thousands). Standard errors are in parentheses.
*p<0.10, **p<0.0.5, ***p<0.01

computes the inverse mills ratio. This ratio is then used as an additional regressor in the regression of wages. In our approach, labor market participation is explained by number of children fully interacted with a dummy for a college education, marital status, and cohort dummies. A fourth-order age polynomial as well as this polynomial interacted with the college education dummy are included as independent variables in the selection equation. Table S3 shows that in this estimation, the Mills ratio is statistically significant, indicating the presence of selection in OLS estimates.

In the second stage, wage profiles are estimated by regressing the natural log of the hourly wage on a fourth-order polynomial in age, the interaction of this polynomial with the college indicator, and dummies for cohorts. We experimented with including marital status in this regression, but estimates were statistically insignificant once we controlled for education. Therefore, we focus on estimates for age profiles which vary only by age. Table S3 includes the coefficients of this regression; figure S4a shows the age and education specific profiles.

These estimates are used to compute predicted wages for individuals in the sample. Residuals between hourly wages in the data and these predicted wages are used to estimate the AR(1) process governing the income risk. Specifically, using these residuals, the parameters of the income process are estimated using Method of Simulated Moments with the identity matrix as the weighting matrix. this process delivers estimates of $\rho = 0.987$ and $\sigma^2 = 0.018$ for the persistence and standard deviation, respectively.

S6 Employment Shocks

There are two main parameters which govern the employment shocks in the model: the probability of receiving the shock, λ, and the wage penalty associated with receiving the shock, ξ. Figure S4 shows the estimation of these parameters.

Table S3: Coefficients of Wage Estimation Equation

	coefficient	standard error
age	0.302***	(0.027)
age*age/(10^2)	-0.9383***	(0.096)
age*age*age/(10^4)	1.309***	(0.145)
age*age*age*age/(10^6)	-0.681***	(0.080)
college*age	-0.038***	(0.005)
college*age*age/(10^2)	0.289***	(0.031)
college*age*age*age/(10^4)	-0.552***	(0.066)
college*age*age*age*age/(10^6)	0.326***	(0.045)
cohort 1911-1920	-0.016	(0.027)
cohort 1921-1930	0.042*	(0.026)
cohort 1931-1940	-0.013	(0.026)
cohort 1941-1950	-0.026	(0.026)
cohort 1951-1960	-0.035	(0.026)
cohort 1961-1970	-0.094***	(0.026)
cohort 1971-1980	-0.114***	(0.027)
cohort 1981-1990	-0.184***	(0.028)
cohort 1991-2000	-0.234***	(0.046)
constant	-0.925***	(0.280)
Inverse Mills Ratio	-0.297***	(0.046)
N	97,896	

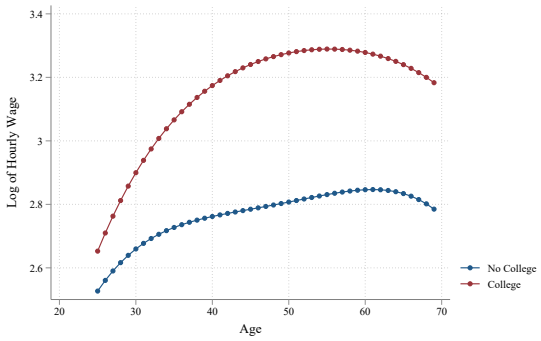
Source: Panel Study of Income Dynamics, authors' calculations

Notes: Estimated coefficients of Heckman two-step model. Dependent variable of wage equation: natural log of the hourly wage. Dependent variable of selection equation: labor force participation. Standard errors are in parentheses.

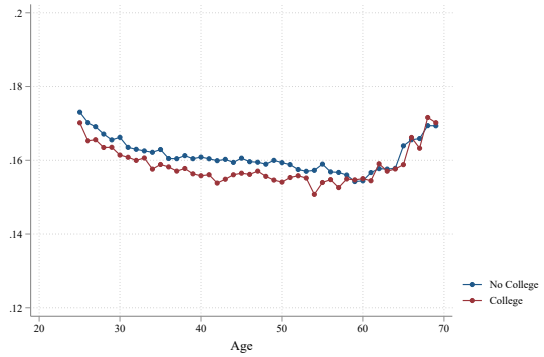
*p<0.10, **p<0.0.5, ***p<0.01

Figure S4: Productivity, Unemployment Share and Wage Impact of Unemployment

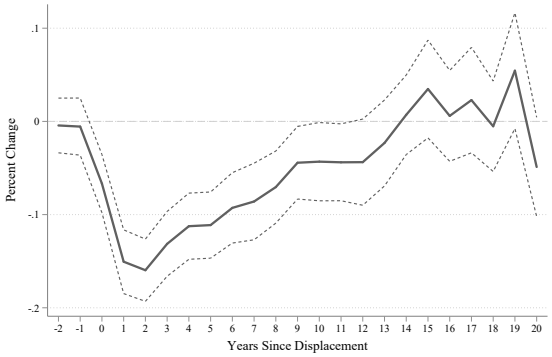
(a) Productivity Profiles



(b) Unemployment Probabilities



(c) Wage Scarring



We set the probability of receiving the unemployment shock based upon the annual separation rate. We use a combination of data from the Job Openings and Labor Turnover Survey (JOLTS) and the Current Population Survey (CPS) to construct education and age-specific separation rates. Specifically, we use JOLTS to construct industry level separation rates as the ratio of total separations to total employment within the industry. Education and age specific rates are calculated as the weighted average of these industry separation rates. We weight them by using the employment share of industries for each age-education pair. This calculation is shown in figure S4b.

These separation rates are not only similar across education groups but also stable over the life cycle. For this reason, we choose to model the employment shock as independent of age, education, and marital status. While the youngest and oldest workers experience slightly higher separation rates, this increase in only 1-2 percentage points. The annual separation rates is roughly 15-16 percent. Similar literature often uses an annual separation rate of 10 percent. In order to remain consistent with previous literature, and since our estimate is only slightly higher, we use a value of $\lambda = 0.1$ in this work.

To construct the wage impact of receiving the unemployment shock, $\xi(\lambda)$, we follow the labor literature which estimates the impact of a displacement on the re-employment earnings of individuals. To measure this, we run the regression in Equation A4.

Table S4: Parameters of the Tax Function

	λ	ξ
Single	1.333	0.042
Married	1.370	0.043

$$y_{it} = x'_{it}\beta + \sum_{k\geq -2}^{20} \delta^k D_{it}^k + \alpha_i + \gamma_t + \varepsilon_{it} \tag{A4}$$

where y_{it} represents the log wages of individual i in time t, x_{it} is a vector of control variables including education and a quadratic term in experience, D_{it}^k is a series of dummies which identify displaced workers the k-th years relative to displacement and α_i, γ_t represent individual and time fixed effects. coefficients, δ^k, results of this regression are shown in figure S4c. Similar to other literature, these results show that wages drop roughly 14-20 percent for those who experienced a displacement relative to those who did not experience a displacement. This impact of displacement is persistent with wages only recovering (relative to the non-displaced) between 5 to 10 years after displacement. We take a value of 14 percent drop in wages due to an unemployment spell, or $\xi(\lambda) = 0.86$. We, additionally estimated these wage scarring by education and found the impact of displacement to be similar across education.

S7 Taxes

PSID includes information on taxes paid up until 1991 and cover tax years up through 1990. In order to have individuals throughout the life cycle, we extend the sample to those workers between the ages of 1916 and 1945. In order to estimate the parameters of the taxation function, we regress the natural log of total family income net of income taxes on a constant and the natural log of family pre-tax income for each education level and marital status. Total taxable income of the family is measured as the sum of labor and Social Security income of the household head and the spouse and other family members (if present). Federal tax liability is constructed based upon the taxable income of the family as well as exemptions and the tax table used.

In order to maximize the sample size for measurement at each education level and marital status, we focus on estimating these parameters independently from age. The estimated parameters are shown in Table S4.

S8 Initial Distribution

The initial distribution across states is calculated by taking random draws of the empirical distribution for men between the ages of 25 and 30. Summary statistics for the initial conditions are shown in Table S5.

Table S5: Summary Statistics for Initial Conditions

	overall	no college, single	college, single	no college, married	college, married
share	1.00	0.18	0.15	0.23	0.43
assets (thousands of 1998 $)					
mean	48.966	24.724	63.105	37.326	60.276
std. dev	91.224	50.408	118.925	64.495	104.31
hourly wage (1998 $)					
mean	15.28	12.65	15.12	13.78	17.06
std. dev	10.24	5.95	13.67	10.55	9.80

S9 Target Profiles

The target life-cycle profiles of participation, hours, and assets used are estimated from the Panel Study of Income Dynamics (PSID) covering years 1968 - 2016 (annually 1968-1995; biannually 1996-2016) using the method from French (2005).

In each year t, PSID respondents are asked to report their labor market variables for year $t-1$. Participation is a measured as a dummy variable taking a value of 1 if the respondent worked at least 300 hours (25 hours per month) in the previous year. Hours is measured as a continuous, self-reported variable of total hours worked in the previous year.

We use wealth observed in the 1984, 1989, and 1994, 1999, 2001, 2003, 2005, and 2007 PSID wealth surveys as well as the 2010-2016 PSID family files. Household assets are measured as the value of farm and business, vehicles, checking and savings accounts, real estate, stocks, mutual funds, IRAs, Keoghs, and investment trusts net any debts. We drop observations with either negative household wealth or household wealth above $1 million (1998 dollars).

To measure mean life-cycle profiles, we run the following fixed effect regression:

$$Y_{it} = f_i + \sum_{c\in\{0,1\}}\sum_{m\in\{0,1\}}\left(\sum_{k=1}^{J}\beta_c^m \mathbb{I}\left[age_{it}=k\right]\times\mathbb{I}\left[married_i=m\right]\times\mathbb{I}\left[college_i=c\right]\right) + \sum_{f=1}^{F}\beta_f\mathbb{I}\left[family\ \ size_{it}=f\right]+\beta_u u_t+\varepsilon_{it}$$

where Y_{it} represents a data observation of the outcome variation (labor force participation, hours, assets) of individual i in time period t; f_i is an individual fixed effect, $\mathbb{I}\left[age=k\right]\times\mathbb{I}\left[married_i=m\right]\times\mathbb{I}\left[college_i=c\right]$ are education and marital status specific age-group dummies; $\mathbb{I}\left[family\ \ size_{it}=f\right]$ are family size dummies; u_t is the unemployment rate in period t. The estimated parameters are $\{\beta_{nc}^{nm},\beta_c^{nm},\beta_{nc}^m,\beta_c^m,\beta_f,\beta_u\}$.

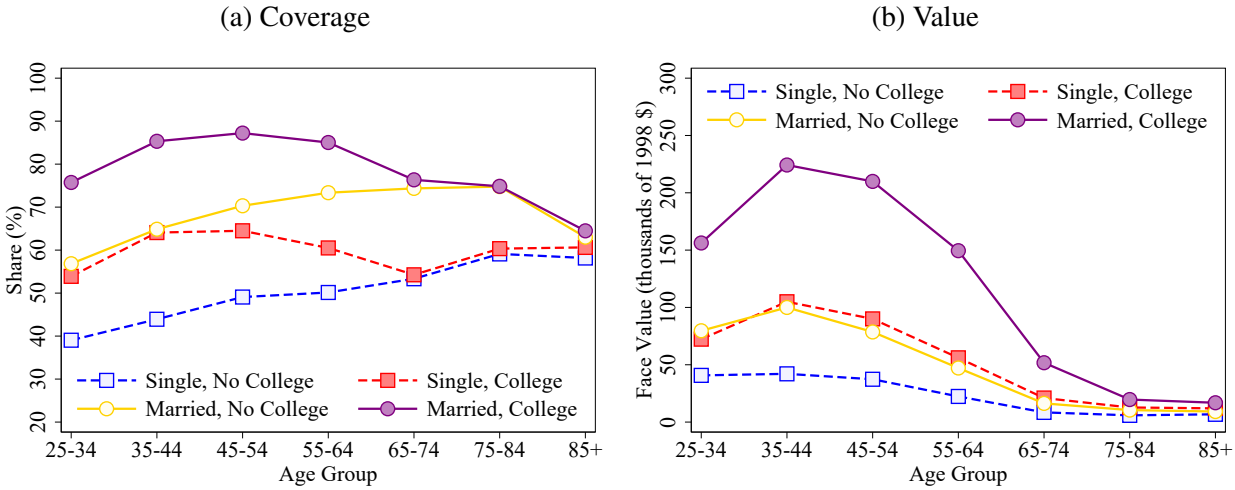
The education-marital status specific age effects from this estimate regression are shown in

Figures A6, A7, and A8 of the main text. In these profiles, we make three adjustments to control for year, family size, and cohort effects: (1) the unemployment rate is set to 6.5%, (2) family size is fixed at 1 for single individuals and 4 for married individuals, and (3) for each education-marital status group, we use the individual specific effect of the same group who are born in 1930 and between the ages of 45 and 49.

For the wealth moments, we report distributional (25th percentile, median, and 75th) moments in addition to the mean (Figure 8). These are computed by running a quartile regression on the set of education and marital status specific age dummies.

The life insurance moments used in the estimation (shown in Figures **??** and **??**) are computed from years 1989 through 2019 of the Survey of Consumer Finances (SCF). Coverage is constructed as a dummy variable which takes a value of 1 if the respondent reports holding at least one term or whole life insurance policy. The face value is a continuous variable taking a value of the total insurance value of all policies. For policies that build cash value, we subtract the cash value from the total face value. We then compute the share of individuals of a given age group-education-marital status group who hold policy as well as the median face value of policies with these groups.

Figure S5: Target Life Insurance Moments by Age, Education, and Marital Status



Source: SCF 1989-2019, authors' calculations

Notes: figure presents the the share of individuals who hold a life insurance policy and the median face value of all life insurance policies held by individuals of a given age, education, and marital status. Value of life insurance held is calculated as the face value net any cash value that has been accumulated (for whole LI policies). All values are in thousands of 1998 $.

These share align well with statistics in Hong and Rios-Rull (2012).[24] It is important to note that we are unable to separate privately purchased life insurance policies from those that are employer provided. Therefore, the coverage of life insurance as well as the face values may include plans that are included as employer benefits rather than optimally purchased. [25]

[24]Hong and Rios-Rull (2012) uses data from the Stanford Research Institute for the year 1990.

[25]Hong and Rios-Rull (2012) are able to separate group plans from purchased plans in the SRI. They find that over 70% of life insurance value is voluntary.