End-of-Life Liquidity

Neha Bairoliya[*]

Neha.Bairoliya@marshall.usc.edu

Giovanni Gallipoli[†]

giovanni.gallipoli@ubc.ca

Kathleen McKiernan[‡]

kathleen.mckiernan@vanderbilt.edu

July 19, 2023

Abstract

Uncertainty about one's own 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. We illustrate its quantitative importance using a model of consumption, labor supply, and retirement decisions and show that the illiquid and uncertain nature of U.S. social security entitlements implies large welfare losses. Reducing the value of retirement annuities in exchange for a guaranteed amount upon death would have significant effects on life-cycle choices, especially among unmarried individuals with low education.

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

JEL Classification Numbers: J14, J26, E21, H55

---

[*]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. In such circumstances, a consumer will value liquidity at the time of death and derive utility from it. Such end-of-life preference for liquidity is not unlike a warm-glow motive; however, it need not be interpreted that way as it arises more generally from the objectives and constraints of expected utility maximization over an uncertain lifespan (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; DeNardi et al., 2016; Jones et al., 2020; DeNardi et al., 2021).

Given this premise, it is perhaps surprising that many social security systems, including the U.S. one, impose limits on the liquidity of entitlements so that they cannot be easily cashed out upon, or before, death. Even in the case of surviving spouses, it is not possible to claim entitlements in the form of lump sum transfers but they must be paid out as a flow extending up to the survivor's demise. This design guarantees that resources are disbursed sequentially when earnings fade for most people, delivering insurance against the survival risk of program participants.

Such restrictions prevent the diminution of old-age income due to early redemption but are not without consequences and may curtail the value of SS entitlements. Perhaps more importantly, the illiquidity of expected entitlements upon death may force individuals to engage in behaviors that circumvent the constraints. Thus, non-trivial distortions in labor supply and consumption may ensue.

For an illustration, consider a person with an end-of-life motive who has little or no other resources. Absent alternative ways to access their entitlements, this person might draw early Social Security payments to set aside funds while they can. This response is an imperfect way to liquidate entitlements and minimize the risk of losing everything upon premature death. This behavior would be associated with significantly different life-cycle consumption, labor market participation, and hours worked. In this example, introducing a partial pay-out upon death would mitigate the distortions due to the potential loss of entitlements at the end of life.

The notion that liquidity considerations are important for labor supply and retirement is not new (see Gustman and Steinmeier, 2005b; French, 2005; French and Jones, 2011; Gallipoli and Turner, 2011, and references therein). Unlike previous studies, we emphasize the large costs of restricting the ability to liquidate entitlements at the end of life and document the distortions they imply in life-cycle settings. To examine these motives and the mechanisms through which they operate, we develop a model of retirement choices (see also Gustman and Steinmeier, 2005a; French, 2005; French and Jones, 2011; Gustman and Steinmeier, 2015; Bairoliya, 2019; Jones and Li, 2020; Pashchenko and Porapakkarm, 2018; Bairoliya and McKiernan, 2021). We focus on the

U.S. context and consider a reform that would allow social security participants to liquidate part of their entitlement upon death. The reform does not facilitate access to funds at younger ages and, therefore, does not introduce front-loading in the timing of payouts. The conservative policy design mitigates concerns about unintended incentives towards dis-investments through early redemption while expanding access to end-of-life benefits to all individuals, as opposed to the current system that only provides widows and widowers with continuing annuities after the death of a beneficiary.

We deliberately avoid modeling a more radical reform for three reasons: first, our focus is on the role of preferences for end-of-life liquidity; second, allowing early withdrawals might encounter stronger opposition and be subject to criticisms of perverse incentives; lastly, we want to show that even a small departure from the *status quo* would induce strong life cycle responses.

An attribute of the reform we examine is that end-of-life entitlements are independent of marital status. This enhances their value for individuals who would otherwise lose them completely upon death and might, therefore, have an incentive to claim early and sock away funds while alive. Replacing part of the annuity with a small end-of-life entitlement enables all individuals to preserve a guaranteed share of their locked-in funds. This includes people whose motives are other than benefiting their spouses, increasing the utility they derive from the SS system irrespective of marital choices. In this sense, the reform can be viewed as an incremental step to equalize the value of the program across participants. Under the current United States Social Security system, FICA taxes are levied on all workers and used to finance survivor benefits. However, survivor benefits are available only to married individuals so that tax revenues from 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 liquidity preference, the reform guarantees a more even treatment of beneficiaries.

In the main counterfactual experiment, we consider a policy that allows all individuals to claim the equivalent of five and one-half years worth of social security annuity benefits upon death; the entitlement comes at the cost of lower per-year benefits and is calculated to keep aggregate government outlays constant. Given the details of the existing SS program, the end-of-life transfer is financed through a 4 percent decrease in yearly payments. This drop is equivalent to the benefit cut an individual would get, under the current system, if anticipating retirement by roughly seven months (SS provisions establish that, in the case of early retirement, a benefit is reduced by $5/9$ of one percent for each month before the normal retirement age).

Making the lump sum transfer available to all workers increases the flexibility of the system for both single and married individuals. For singles, the policy provides access to a guaranteed minimum claim in exchange for a small 4 percent decrease in the annuity; this claim is valid upon death even if it occurs before retirement. By the same token, a key advantage for married individuals is that the end-of-life benefit can be claimed at any time, even before the spouse is of

3

retirement age. The value of the annuity of the initial claimant is also decreased by 4 percent and, conditional on current mortality patterns, the reform only marginally reduces the expected value of the survivor benefits. This is because survivor annuities are received for an average of 6 years after spousal death while the lump sum in the baseline experiment amounts to 5.5 times yearly benefits, which almost fully compensates married individuals for their expected annuity reduction.

One lesson from this experiment is that the behavioral and welfare impacts of an incremental policy reform can be large if it affects the liquidity provisions of the SS program. The reform would bring about non-trivial welfare gains in expected terms (ex-ante). In addition, it would come close to being Pareto improving (ex-post): nearly 70 percent of households would experience welfare gains post-policy with the exception of some married individuals, whose welfare would be very marginally lower than under the current system.

As one might expect, gains from this reform would be larger for unmarried workers who acquire a benefit that does not dissolve upon death while costing only a small drop in their annuity. Gains also tend to be larger among single individuals with lower incomes and education as they gain flexibility and access to this lump-sum at younger ages. The welfare changes in the broader group of all married SS recipients, irrespective of education, are also on average positive. This is more surprising since the reform involves the universal extension of benefits that were originally directed only to married survivors. The main reason for the large returns from the universal end-of-life liquidation is that the benefit does not depend on having claimed SS annuity income and, therefore, it guarantees partial access to entitlements even in the event of early death.

Despite its limited nature, the reform triggers considerable labor supply responses along both the extensive and intensive margins. Single individuals delay their SS benefit claims, which boosts the value of their guaranteed end-of-life benefits as well as their expected annuity. Under the reform, singles do not have to worry about the loss of entitlements upon death: this means that they can delay claiming and, therefore, can increase their hours with no concern about the earnings test claw-back. Married individuals also delay their claims due to the increased flexibility introduced by the the reform. Since they can rely on an end-of-life liquidation no matter how short their life, these individuals increase both participation and hours.

We find that the post-reform behavior of single individuals is consistent with weaker incentives for old-age savings. Since they no longer need to set aside resources for late-life liquidity, their asset holdings after retirement are reduced between 20 and 40 percent (with larger proportional drops among the less educated, indicating a tighter constraint under the current system). In contrast, the reform induces only small changes in the behavior of college-educated couples, who are the least liquidity-constrained group and do not materially respond to small tweaks to Social Security. The saving responses of married individuals without a college education are larger but asymmetric over their life cycle. During their working lives, less educated couples exhibit lower

savings relative to the baseline; after retirement, they accumulate more wealth than in the baseline. This is because, under current SS arrangements, married workers can only leave survivor benefits after claiming, which pushes them to accumulate assets to satisfy end-of-life liquidity motives in case of premature death. The reform, instead, makes these workers eligible for a lump sum benefit upon death regardless of age, which means they do not need to save as much in their early life. However, conditional on surviving into older age, surviving spouses receive slightly lower annuities in the experiment and have stronger incentives to hold assets in late life for self-insurance, which explains the marginally higher asset holdings among less educated married couples.

The results of the main counterfactual experiment demonstrate the impact of a specific experiment in which individuals receive 5.5 times benefits at death in exchange for a 4 percent reduction in annuity benefits. However, we also demonstrate that the welfare impacts remain similar across various combinations of end-of-life transfers (and annuity payment reductions which finance these transfers). Single workers, especially those without a college degree, experience welfare gains for all experiments; these gains are increasing in the size of the lump-sum. Married workers continue be indifferent or experience gains which are slightly negative–highlighting the impact of redistribution between married and single individual in driving the welfare results. When we consider an additional reform focused only on singles, providing even a small end-of-life transfers comes at the cost of large decreases in annuity payments. While single workers value this end of life transfer, the decrease in the annuity is very costly and leads to welfare losses.

Given the importance of redistribution between married and single individuals, we turn attention back to the universal transfer reform. By using married individuals to help finance these end-of-life transfers, a reform which trades a portion of the old-age annuity for a lump-sum payment at death achieves average welfare gains. These welfare gains vary by the size of the lump-sum considered; average gains are largest for this reform which combines a end-of-life transfer of 5.5 years of benefits with a 4 percent decline in old-age annuity payments.

Our work makes three contributions. First and foremost, we present evidence that end-of-life liquidity motives may lead different types of households to change their labor supply and consumption behaviors in peculiar ways; in fact, we show that liquidity needs that extend up to the time of death can shape behaviors much earlier in the life cycle (Dynan et al., 2002; French et al., 2006; Kopczuk and Lupton, 2007; DeNardi et al., 2016; Jones et al., 2020; DeNardi et al., 2021).

Second, we provide new insights into the welfare implications of reforming large programs like the Social Security system (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; Huggett and Parra, 2010; Scheuer and Slemrod, 2021) with a special focus on liquidity considerations. In particular, we document how the interaction of end-of-life liquidity motives and social security incentives can result in large welfare losses.

Third and last, by examining a policy that redistributes resources from the married to the singles, our work emphasizes some of the incentives built into existing systems and advances our understanding of little-known drivers of observed gaps in the earnings and labor supply of different 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 Background: United States Social Security

The United States Social Security system provides an annuity payment to older workers as well as survivor benefits that allow married workers to pass on benefits to their spouses in the event of death.

2.1 Old-Age Benefits

The Social Security system provides a flow of retirement income that starts at the time of claiming and continues until the death of the beneficiary. Benefits are a progressive function of the average indexed monthly earnings. Up to a maximum taxable amount, higher income during working life translates to higher benefits during retirement. The progressivity of the formula means that high-income individuals receive lower replacement rates on their earnings than lower-income workers. Spouses of primary earners can claim *spousal benefits* on the earnings record of the primary earner.[1] These benefits may be up to 50 percent of the benefits of the primary earner and are contingent on the primary earner having claimed Social Security 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. The structure of the U.S. Social Security system, therefore, creates a trade-off between the number of years pension payments are received and their size.

2.2 Survivors Benefits

A key objective of the United States Social Security system is providing survivor benefits to family members in the case of the death of a beneficiary. These benefits are available to spouses–

---

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

and, potentially, past spouses–, children, and parents of an eligible deceased beneficiary. However, the size of benefits received varies with their relationship to the covered individual and limits the ability to make late-life transfers.

Social Security benefits are based upon the number of work credits accumulated[2], whether a worker's record is eligible to pay out survivor benefits depends upon 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.

Spouses are the largest group receiving SS survivor benefits with roughly 4 million spouses receiving benefits through this channel.[3] Widows and widowers of eligible workers can receive reduced benefits of 77.5 percent of the beneficiaries basic SS benefits[4] as early as age 60 or full benefits at the normal retirement age.[5] Moreover, spouses who were living with the deceased are eligible for a one-time lump-sum payment in the month when the death occurs.[6] Benefits are also available to divorced spouses provided that the marriage lasted at least ten years.[7]

In addition to survivor benefits for spouses, benefits may be paid out to children and dependent parents of a beneficiary. Children under age 18 may receive benefits equivalent to 75 percent of the beneficiary's primary insurance amount. These benefits expire once the child is no longer a minor. Additionally, if the deceased beneficiary provided at least 50 percent of the parents' support, surviving parents over the age of 62 may receive survivor's benefits. These benefits are 82.5 percent for a single surviving parent or 75 percent each for two surviving parents.

There is a maximum family benefit amount[8] that can be paid out on a beneficiary's earnings record.[9] This maximum amount ranges from 150 percent to 180 percent of the primary insurance amount. If the total survivor's benefits are greater than this limit, payouts to all recipients are reduced proportionally.

---

[2]Workers can accumulate up to 4 credits annually based upon their wages and self-employment income. For example, in 2023, workers can earn 1 credit for each $1,640 earned or the maximum of 4 credits per year if they earn at least $6,560.

[3]Details:https://www.ssa.gov/benefits/survivors/onyourown.html#h4;https://www.ssa.gov/pubs/EN-05-10084.pdf

[4]Basic Social Security benefits = Primary Insurance Amount

[5]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.

[6]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.

[7]Divorced spouses lose eligibility for these spousal benefits if they remarry before the age of 60.

[8]Details: https://www.ssa.gov/OACT/COLA/familymax.html

[9]This amount excludes any benefits paid out to divorced spouses.

3 Model

We develop and estimate a dynamic model of labor supply and retirement. To capture the nature of life cycle incentives under different policy regimes, we carefully model retirement benefits from Social Security to match key outcomes of the current U.S. system.

We study the choice of labor supply (h_t), consumption (c_t), savings (a_{t+1}) and the SS benefit application (b_t^{ss}) of couples and singles. Individuals make decisions in each time period t and adjust their behavior in response to changes in wages, employment, health, and survival.

The life-cycle spans ages $t = 25, 26, ..., 99$. Individuals are heterogeneous with respect to both permanent and evolving states. They are permanently different in their 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 if married.

Evolving states include stochastic labor productivity (η_t), employment status (λ_t), health status (μ_t), assets (a_t), social security wealth (a_t^{ss}) and application status (b_{t-1}^{ss}). Given this vector of states $\left(e, q, \eta_t, \lambda_t, \mu_t, a_t, a_t^{ss}, b_{t-1}^{ss}\right)$, individuals choose consumption, labor supply, and Social Security benefit application decisions (if eligible) to maximize the present discounted value of lifetime utility.[10] Below we describe different elements of the model in detail.

3.1 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[11] and is defined as

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

The parameter ρ dictates relative risk aversion; ν is the weight on consumption; ζ_t is a consumption equivalence scale. The utility of married households is multiplied by two to account for spousal utility from consumption and leisure. The amount of 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

---

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

[11]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,).

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:[12]

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

At the end of life, an individual values wealth A_t^q, as in 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 amounts to any remaining assets, a_t, plus Social Security survivor benefits, if eligible. Eligibility for survivor benefits depends on marital status and the age gap between spouses, q.[13] The coefficient θ measures the intensity of the end-of-life motive, and κ changes the curvature of the function. Higher θ increases the marginal utility of an extra unit of end-of-life resources and higher κ makes them more similar to a luxury good.

3.2 Health and Mortality

Each period, individuals are subject to an exogenous process that affects their survival probability as well as their time endowment. The transition rule across health states depends on current health, education, and age. The probability to transition between two health states i and j is:

$$\pi_{t+1}^{\mu_{ij}} = prob(\mu_{t+1} = j | \mu_t = i, e, t+1)$$

Individuals are also subject to mortality shocks. The probability of survival depends on age and current health, as shown below:

$$\pi_{t+1}^s = prob(s_{t+1} = 1 | \mu_t, m, t+1)$$

3.3 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

---

[12]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.

[13]Details on survivor benefits are discussed in Section 3.4.1

effects.

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

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.[14]

$$w_t = \xi(\lambda_t) \exp(\omega(e, t) + \eta_t) \tag{4}$$
$$\eta_t = \rho^w \eta_{t-1} + \epsilon_t^w$$
$$\epsilon_t^w \sim N(0, \sigma_{\epsilon^w}^2)$$

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.

3.4 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 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:

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

---

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

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.[15] 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).

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

Adjustment (1): Early/Late Claiming Penalty. SS benefits can be claimed without penalty at the normal retirement age (t_{NRA}).[16] 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's 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.

---

[15]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.

[16]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.

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

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:

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

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 Social Security wealth as shown below:[17]

$$ssb_{t+1} = pia(a_{t+1}^{ss}) * \left[1 + \left(\frac{\Upsilon_t}{ssb_t} \right) \gamma_t^{ss} \right]$$

$$a_{t+1}^{ss*} = pia^{-1}(ssb_{t+1}) \qquad (9)$$

where γ_t^{ss} 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_t^{ss}. In combination with the application age requirements, the earnings test may reinforce the incentives to retire upon reaching the claiming age.

3.4.1 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 $\delta_t^q ssb_t$, where δ_t^q is determined as:

---

[17]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.

$$\delta_t^q = \begin{cases} 1.0 & \text{if } m = \text{single or } m = \text{married}, t^{ss} - \iota < t_{ERA} \\ 1.5 * [1 - (t_{NRA} - (t^{ss} - \iota)) * \gamma_t^{ss})] & \text{if } m = \text{married}, \ t_{ERA} \leq t^{ss} - \iota < t_{NRA} \\ 1.5 & \text{if } m = \text{married}, \ t^{ss} - \iota \geq t_{NRA} \end{cases}$$

(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.

Survivor 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:

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

(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).

3.5 Budget Constraint

The household budget constraint can be summarized as:

$$c_t + a_{t+1} = a_t + W(y_t, y_{st}, \bar{r}a_t, \tau) + \delta_t^q ssb_t + tr_t$$

(12)

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:

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

3.6 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). 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).

3.7 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.

3.7.1 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.[18] 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:

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

$$
a_{t+1} = a_t + W(y_t, y_{st}, ra_t, \tau) + tr_t - c_t,
$$
$$
1), (5\text{-}9), (14), \text{ and } c_t \geq \bar{c}.
$$

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:

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

---

[18]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.

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

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

$$a_{t+1} = a_t + W(y_t, y_{st}, ra_t, \tau) + tr_t - c_t,$$
$$1), (5\text{-}9), (14), \text{ and } c_t \geq \bar{c}.$$

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

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

$$a_{t+1} = a_t + W(y_t, y_{st}, ra_t, \tau) + tr_t + \delta_t^q ssb_t - c_t,$$
$$1), (5\text{-}9), (14), \text{ and } c_t \geq \bar{c}.$$

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:

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

$$a_{t+1} = a_t + W(y_{st}, ra_t, \tau) + \delta_t^q ssb_t + tr_t - c_t,$$
$$(1), (6), (14) \text{ and } c_t \geq \bar{c}.$$

4 Estimation

We estimate the model on a sample of male household heads born between 1931 and 1935. Estimation proceeds in two-steps (Gourinchas and Parker, 2002). 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 programming model. We call this vector of estimates Φ: it includes health transitions, 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 relevant cohort, our structural model, and the parameters from the first step to estimate the preference parameter vector $\Theta = \{\beta, \rho, \nu, \theta, \kappa, \phi_H(t, \mu_t), \phi_P(t)\}$ by education and marital status. In this step, we employ the Method of Simulated Moments (MSM).

4.1 Estimation First Step

Health and Mortality. 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, .[19] We estimate health transitions across these states through an ordered probit of 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. Since the MEPS does not sample the institutionalized population, we adjust these profiles to match life expectancy at age 65 for both education groups in our benchmark birth cohort (born between 1931 and 1935).[20]

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

---

[19]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.

[20]Data on LE: https://www.ssa.gov/policy/docs/workingpapers/wp108.html

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.

Labor productivity. 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).

Employment shocks and wage scarring. 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.

Social security. Modeling the rich detail of the U.S. Social Security System, described in Section 3.4, 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.

4.2   Estimation Second Step

Having recovered the vector of exogenous data generating processes Φ and the vector of preference parameters Θ, we numerically solve for the decision rules $c(z_t, \Phi, \Theta)$, $h(z_t, \Phi, \Theta)$, and $b^{ss}(z_t, \Phi, \Theta)$ using backward induction. Then, we use estimates of Φ and of initial conditions z_0 to simulate the life-cycle profiles for a large set of individuals. Lastly, we employ an MSM criterion function to find the $\hat{\Theta}$ that minimizes the distance between simulated profiles and data observations. We match the following moments to estimate of Θ by education (college or no college) and marital status groups (single or married):

1. Labor market participation of male household heads between ages 25 and 69, resulting in 180 moment conditions.

2. Log of hours worked, conditional on participation, of male household heads between age 25 and 69, resulting in 180 moment conditions.

3. Mean assets of male household heads between ages 25 and 69, resulting in 180 moment conditions.

These moments add up to a total of 540 conditions.[21] The MSM estimate of $\hat{\Theta}_{MSM}$ solves:

$$\hat{\Theta}_{MSM} = \text{argmin } \tilde{g}(\Theta, \Phi) W_T \tilde{g}(\Theta, \Phi)$$

---

[21]Life-cycle profiles are estimated by fitting a fourth-order polynomial in age and controls for education and marital status (in levels plus interaction with each other and age) for the 1931-1935 cohort. We estimate such polynomials for participation, hours, and wealth profiles. See Bairoliya and McKiernan (2021) for a discussion of wealth measures.

where

$$\tilde{g}(\Theta, \Phi) = \begin{bmatrix} \frac{1}{N} \sum_{i=1}^{N} \{p_{it} - \tilde{p}_t^{e,m}(z_{it}, \Theta, \Phi)\} \\ \frac{1}{N} \sum_{i=1}^{N} \{\log h_{it|p_{it}>0} - \log \tilde{h}_{t|p_t>0}^{e,m}(z_{it}, \Theta, \Phi)\} \\ \frac{1}{N} \sum_{i=1}^{N} \{a_{it} - \tilde{a}_t^{e,m}(z_{it}, \Theta, \Phi)\} \end{bmatrix}$$

$$t = \{1, ..., T\}; \ \ e \in \{\text{non-college, college}\}; \ \ m \in \{\text{single,married}\}.$$

W_T allows for an optimal weighting matrix given by the inverse of a consistent estimate of the covariance matrix of data moments. However, efficient choice of weighting matrix may introduce finite sample biases. Therefore we adopt the following non-optimal weighting matrix in the structural estimation:

$$W_T = \left[diag\left(var\left(\frac{1}{\sqrt{N}} \sum_{i=1}^{N} m_{it} \right) \right) \right]^{-1}$$

where m_{it} is a vector of data moments.

4.3 Model fit

Figure B.1 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. Figure B.2 describes asset holdings over the life cycle, by education and marital status. The model closely matches variation in levels and growth, including the key observation of no dis-saving in old age. Finally, Figure B.3 compares observed and simulated claiming rates at four junctures: at the ERA, between ERA and NRA, at the 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.

5 Counterfactual Experiments

In the numerical experiments, we hold government outlays on SS lifetime payouts at the benchmark level. Denoting the total present expected value of the end-of-life lump sum transfers as

PV^D, we impose the following constraint in each experiment:

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

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

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.

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.

5.1 Universal transfer reform

We begin by examining 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, all individuals are guaranteed, upon death, the equivalent 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 (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. Table 2 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.

Table 2: Universal Transfer Reform

	Singles		Married	
	Non-College	College	Non-College	College
Participation	-0.64	1.99	0.59	0.47
Hours	-1.85	1.85	1.44	0.55
Assets	-41.17	-23.76	1.67	-0.34
Consumption	1.50	-1.24	-0.21	-0.33
Welfare	4.79	1.35	-0.45	0.21

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.

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 positive and a little below five percentage points for the relatively poor non-college singles. The reform induces them to delay retirement as these individuals no longer need to claim their annuities in order to avoid losing everything in case of premature death. Labor force participation and hours worked at older ages changes differently between college and non-college singles. Non-college singles reduce participation and hours worked since the income effect from the reform dominates in this group. College-educated singles experience a smaller CEV gain of 1.35% on average with most individuals exhibiting gains slightly over 1 percent (see Figure 1). They also experience a large increase in hours worked. Their stronger labor supply response is partly due to the novel ability, under the reform, to circumvent the earnings-test claw-back that applies after claiming SS benefits (for these workers the substitution effect is strong enough to offset the income effect). Delaying SS claims allows these higher 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, non-married individuals reduce their savings significantly under the reform. This is especially true after retirement (Figure B.5). The drop in asset holdings can exceed $1/3$ 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.

Married individuals. Welfare changes are not as large among married individuals. The change is slightly negative (-0.45% CEV) for non-college households, while it is small and positive

(0.21% CEV) among the college-educated. Since married individuals tend to live longer than the unmarried, the ability to plan their retirement in a less constrained way carries positive returns in terms of welfare.

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 a relatively stronger labor supply response in the experiment (as opposed to the college workers): 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 B.5). 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 begin 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 B.5). 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 1 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 (non-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.

Figure 1: Welfare: Universal Transfer Reform
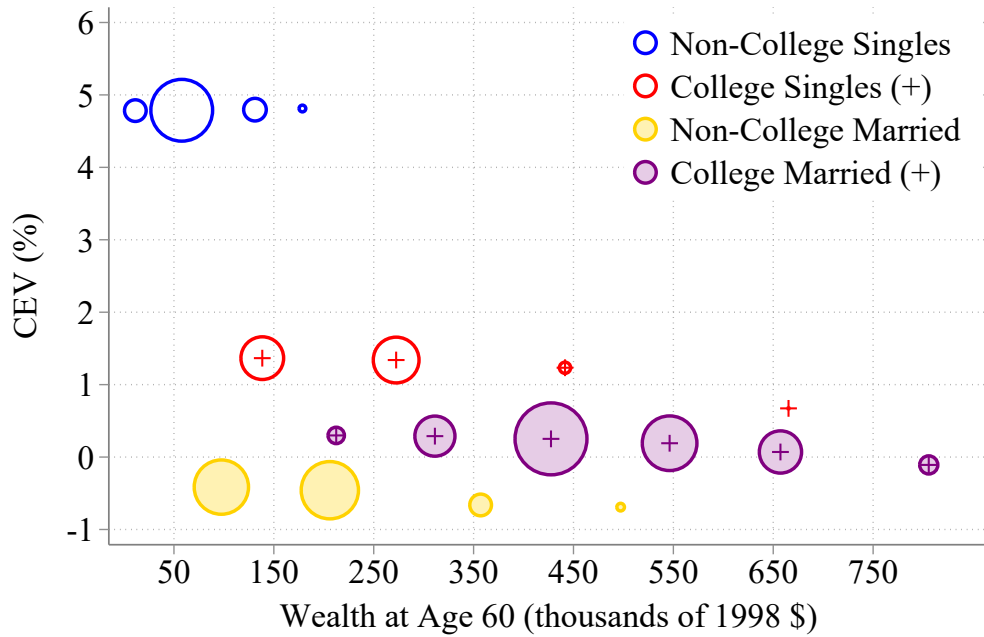
Notes: Figure B.18 shows this CEV calculation by wealth decile for each of these four education-marital status groups.

5.2 The value of the reform over 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 B.8 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, 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 B.7 shows a jump in benefit flows at the average age at which the annuity is claimed. 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.

The flatness of the 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 between benchmark and experiment accruable benefits at each age are large and can reach almost ten thousand dollars after retirement.

The large welfare gains under the reform follow from the fact that benefits can be accrued at young ages and before reaching retirement age, as shown in Figure B.8. 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. Figure B.9 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 B.11 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.

5.3 The Value of Larger 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 2 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 4.8 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 2: Welfare: Variation Across EOL Transfers
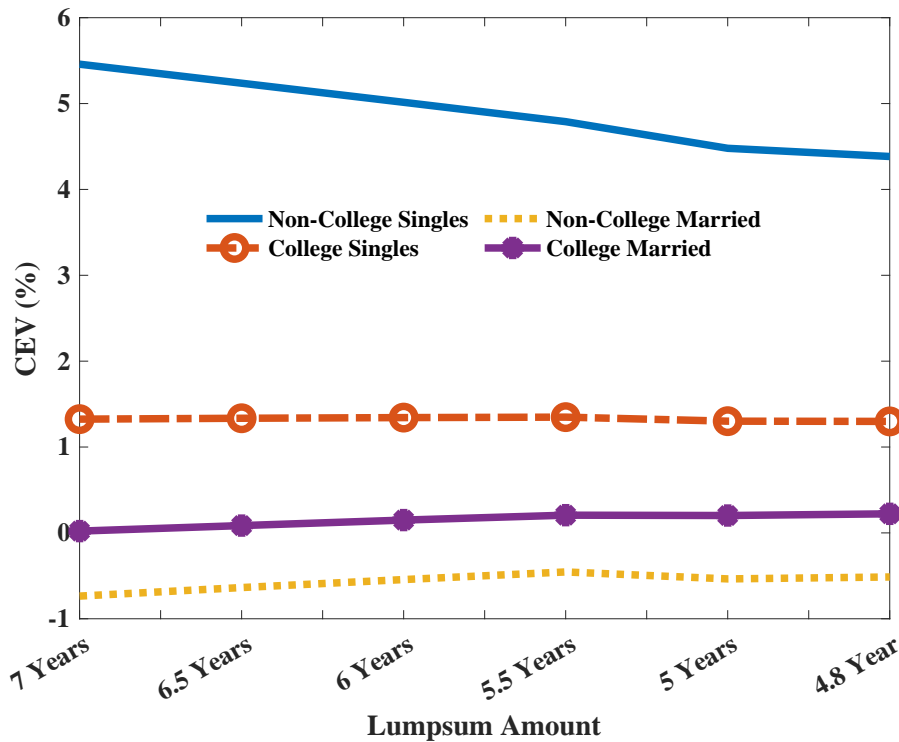
Notes: Figure reports average CEV of each education-marital status group across different EOL transfer sizes. Each EOL size is associated with a decrease in annuity payments ranging from 0 (for a lump-sum of 4.8 years of benefits) to 16 percent (for a lump-sum of 7 years of benefits)

Across all values for the end-of-life transfer, we observe a consistent pattern. While married individuals are nearly indifferent or 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.

We highlight this redistribution by performing an experiment that focuses only on single individuals. This experiment is more expensive as the decrease in the annuity required to keep the the reform expenditure-neutral must be larger when the baseline survivors benefits of married individuals are not used to finance the reform. For example, to finance an end-of-life transfer of one year of benefits to single workers, their annuity payments must decrease by 18 percentage points.

Figure 3: Welfare Across Experiments: Non-Universal Transfer Reform

(a)   (b)
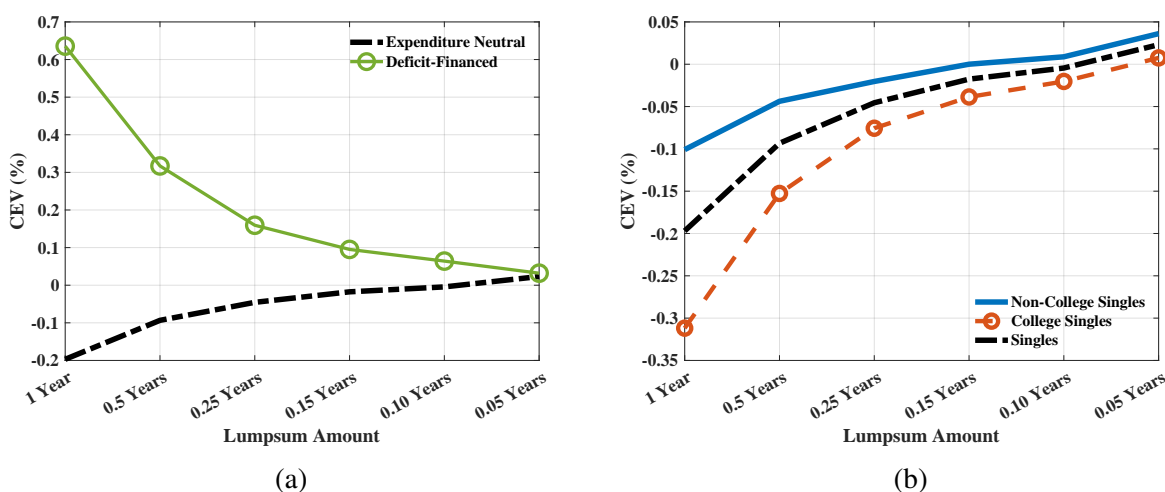
Notes: Panel (a) shows the average CEV of singles across two sets of policy experiments 1) individuals are paid lump sum without corresponding decrease in annuity (deficit financed) and 2) individuals are paid lump sum with a corresponding decrease in annuity (expenditure neutral). Panel (b) presents 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 18 percent for a EOL transfer of 1 year of benefits.

Figure 3 shows results for this experiment featuring only single workers. In panel (b) we show how CEV varies for EOL transfers valued between 1 0.05 years. Unlike the positive relationship between the EOL transfer and welfare we highlighted for singles in the universal transfer reform of section 5.1, considering only singles makes the CEV decreasing in the size 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 the positive relationship between CEV and the size of the transfer established previously. However, when the reforms are required to be expenditure-neutral, the negative relationship between EOL transfers and CEV appears. While singles value the liquidity provided by the EOL transfer, they also highly value the size of the annuity they receive. However, since benefits are concentrated among poorer households, they are unable to afford even the small

cost associated to the reform.

5.4 Back to the Universal Transfer Reform: What EOL Transfer has the Highest Average CEV?

The previous section documents 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. Figure 5 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.

Figure 5: Average CEV by EOL Transfer
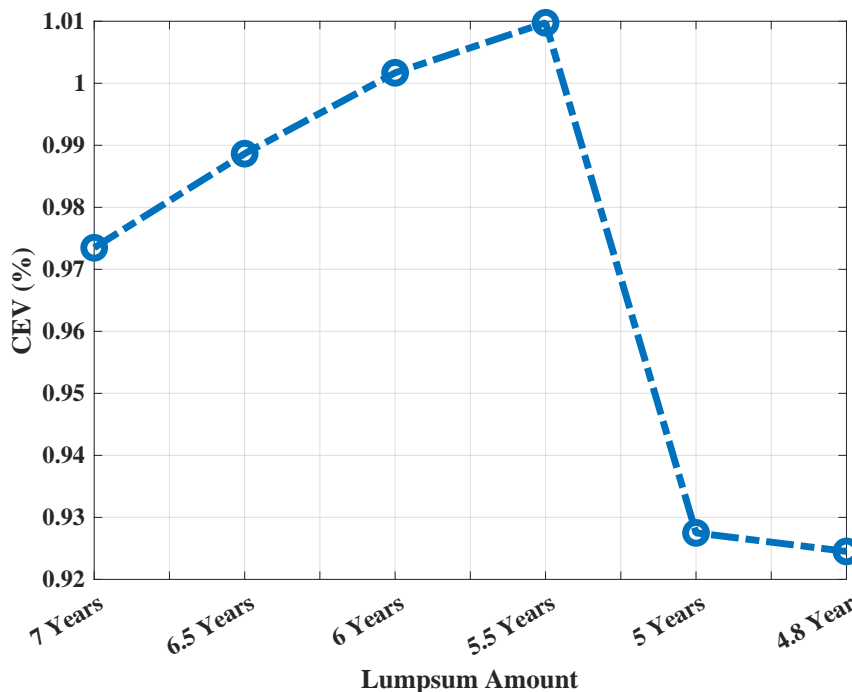
Notes: Figure reports average CEV of the simulated population across different EOL transfer sizes. Each EOL size is associated with a decrease in annuity payments ranging from 0 (for a lump-sum of 4.8 years of benefits) to 16 percent (for a lump-sum of 7 years of benefits)

6 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

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

Bairoliya, N. (2019). Pension plan heterogeneity and retirement behavior. *European Economic Review*, 116:28–59.

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.

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

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.

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

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

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

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

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., DeNardi, M., Jones, J. B., Baker, O., and Doctor, P. (2006). Right before the ecnd: 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.

Gallipoli, G. and Turner, L. (2011). Social security, endogenous retirement, and intrahousehold cooperation. Technical report, UBC.

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.

Gustman, A. and Steinmeier, T. L. (2005a). The social security early retirement age in a structural model of retirement and wealth. *Journal of Public Economics*, 89(2-3):441–63.

Gustman, A. and Steinmeier, T. L. (2015). Effects of social security policies on benefit claiming, retirement and saving. *Journal of Public Economics*, 129(C):51–62.

Gustman, A. L. and Steinmeier, T. L. (2005b). The social security early entitlement age in a structural model of retirement and wealth. *Journal of Public Economics*, 89(2):441–463.

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.

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., 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., DeNardi, 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. (2020). Social security reform with heterogeneous mortality. *FRB Richmond Working Paper No. 20-09*.

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.

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

Pashchenko, S. and Porapakkarm, P. (2018). Accounting for social security claiming behavior. *Working Paper*.

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.

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

Appendix

A   CEV Calculation

Let $\{c_t^*, l_t^*\}_{t=1}^T$ denote benchmark optimal choices of consumption and leisure respectively and $\{c_t^p, l_t^p\}_{t=1}^T$ denote optimal choices in the policy world.

Then preferences with end-of-life motives can be mapped to preferences without them using a simple scaling factor for each individual, in both benchmark and policy worlds, in the following way:

$$\sum_{t=1}^T E\beta^t \left[s_t \frac{\left(c_t^{p\nu}(1-l_t^p)^{1-\nu} \right)^{1-\rho}}{1-\rho} + (1-s_t)\Omega(a_{t+1}^p) \right] = \sum_{t=1}^T E\beta^t \left[s_t \frac{\left(c_t^{p\nu}(1-l_t^p)^{1-\nu} \right)^{1-\rho}}{1-\rho} \right] \times \kappa_1$$
(17)

$$\sum_{t=1}^T E\beta^t \left[s_t \frac{\left(c_t^{*\nu}(1-l_t^*)^{1-\nu} \right)^{1-\rho}}{1-\rho} + (1-s_t)\Omega(a_{t+1}^*) \right] = \sum_{t=1}^T E\beta^t \left[s_t \frac{\left(c_t^{*\nu}(1-l_t^*)^{1-\nu} \right)^{1-\rho}}{1-\rho} \right] \times \kappa_2$$
(18)

Given that we can write lifetime utility including an additive end-of-life liquidity motive as the product of lifetime utility without the end-of-life flow utility and a multiplicative constant, CEV (τ) in our framework is computed as follows:

$$\sum_{t=1}^T E\beta^t \left[s_t \frac{\left(c_t^{p\nu}(1-l_t^p)^{1-\nu} \right)^{1-\rho}}{1-\rho} \right] \times \kappa_1 = \sum_{t=1}^T E\beta^t \left[s_t \frac{\left([(1+\tau)\,c_t^*]^\nu \,(1-l_t^*)^{1-\nu} \right)^{1-\rho}}{1-\rho} \right] \times \kappa_2$$

$$\sum_{t=1}^T E\beta^t \left[s_t \frac{\left(c_t^{p\nu}(1-l_t^p)^{1-\nu} \right)^{1-\rho}}{1-\rho} \right] \times \kappa_1 = (1+\tau)^{\nu(1-\rho)} \sum_{t=1}^T E\beta^t \left[s_t \frac{\left(c_t^{*\nu}(1-l_t^*)^{1-\nu} \right)^{1-\rho}}{1-\rho} \right] \times \kappa_2$$

$$\sum_{t=1}^T E\beta^t \left[s_t \frac{\left(c_t^{p\nu}(1-l_t^p)^{1-\nu} \right)^{1-\rho}}{1-\rho} + (1-s_t)\Omega(a_{t+1}^p) \right] = (1+\tau)^{\nu(1-\rho)} \sum_{t=1}^T E\beta^t \left[s_t \frac{\left(c_t^{*\nu}(1-l_t^*)^{1-\nu} \right)^{1-\rho}}{1-\rho} + (1-s_t)\Omega(a_{t+1}^*) \right]$$

$$V^P = (1+\tau)^{\nu(1-\rho)} V^*$$

$$\implies \tau = \left(\frac{V^p}{V^*} \right)^{1/\nu(1-\rho)} - 1$$

Where V^* is the optimal lifetime utility of an agent in the benchmark world and V^p denotes the lifetime utility of an agent in the policy experiment.

B   Additional Figures and Tables

Figure B.1: Model Fit: Labor force participation rates

(a) Non-College, Non Married

(b) College, Non Married

(c) Non-College, Married

(d) College, Married
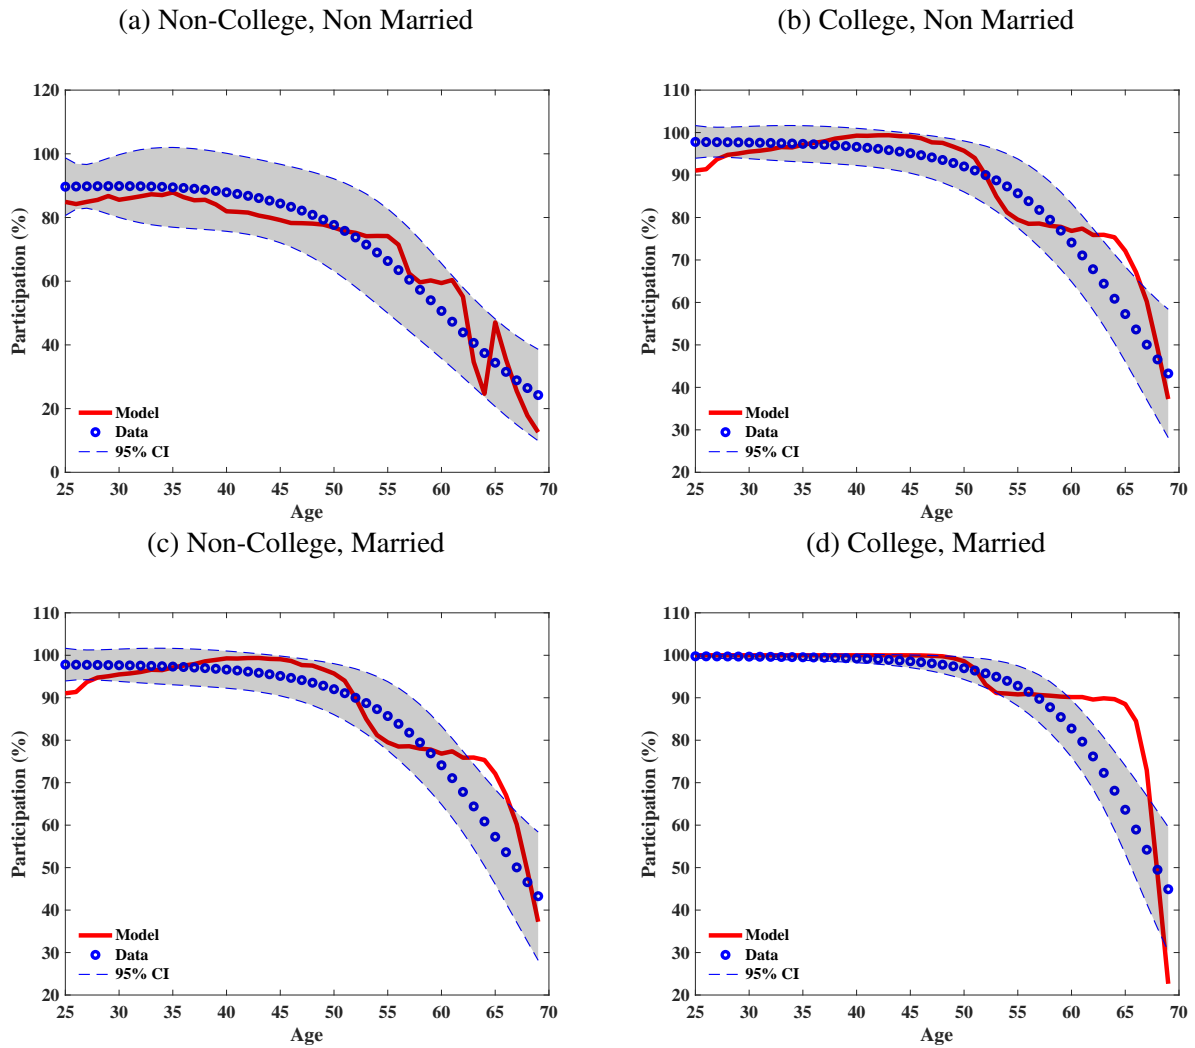
Figure B.2: Model Fit: Wealth

(a) Non-College, Non Married

(b) College, Non Married

(c) Non-College, Married

(d) College, Married
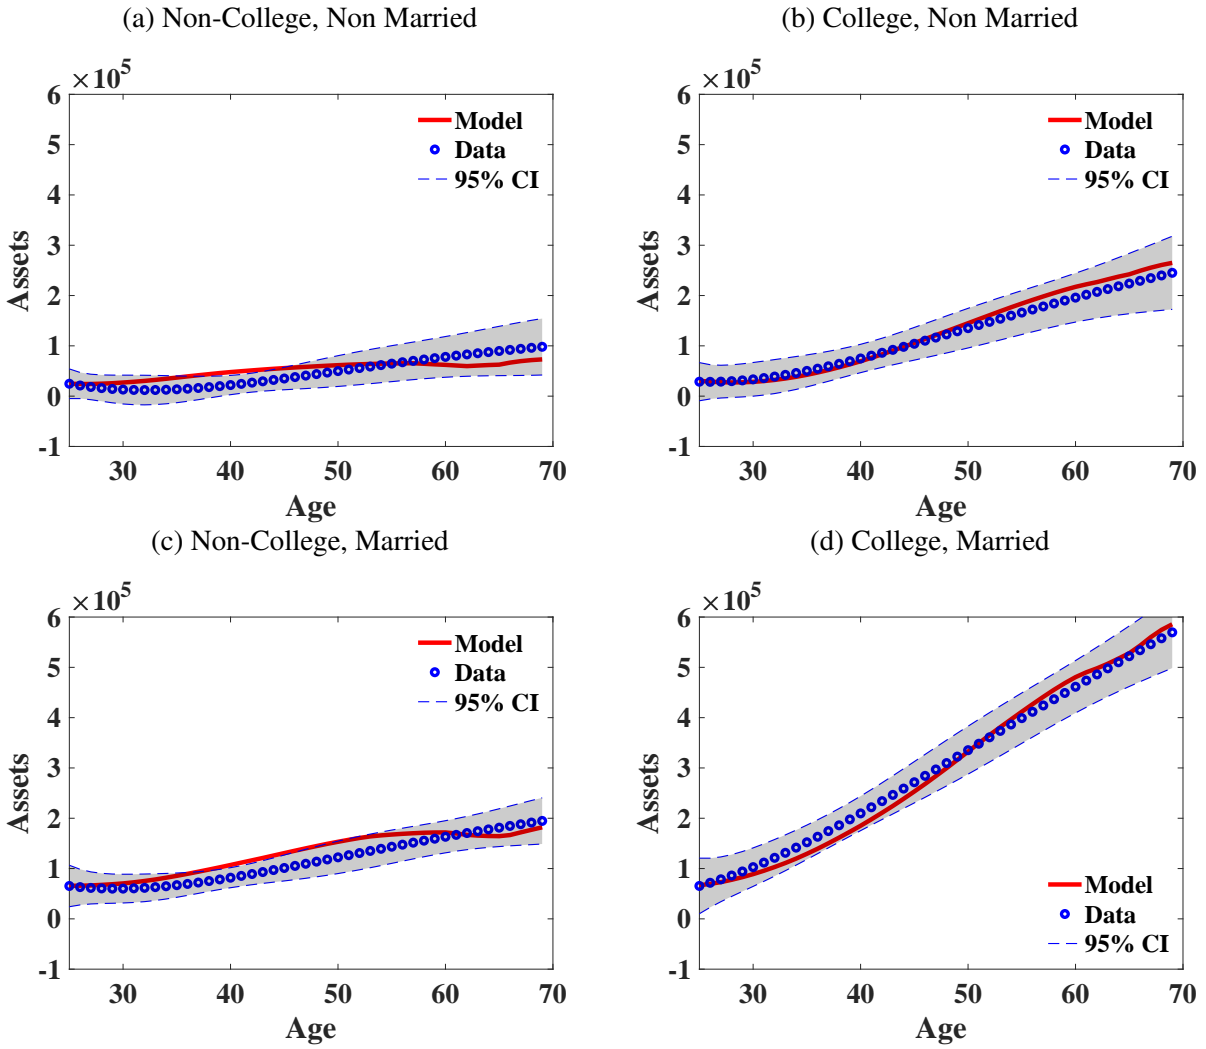
Figure B.3: Model Fit: SS Claiming Behavior
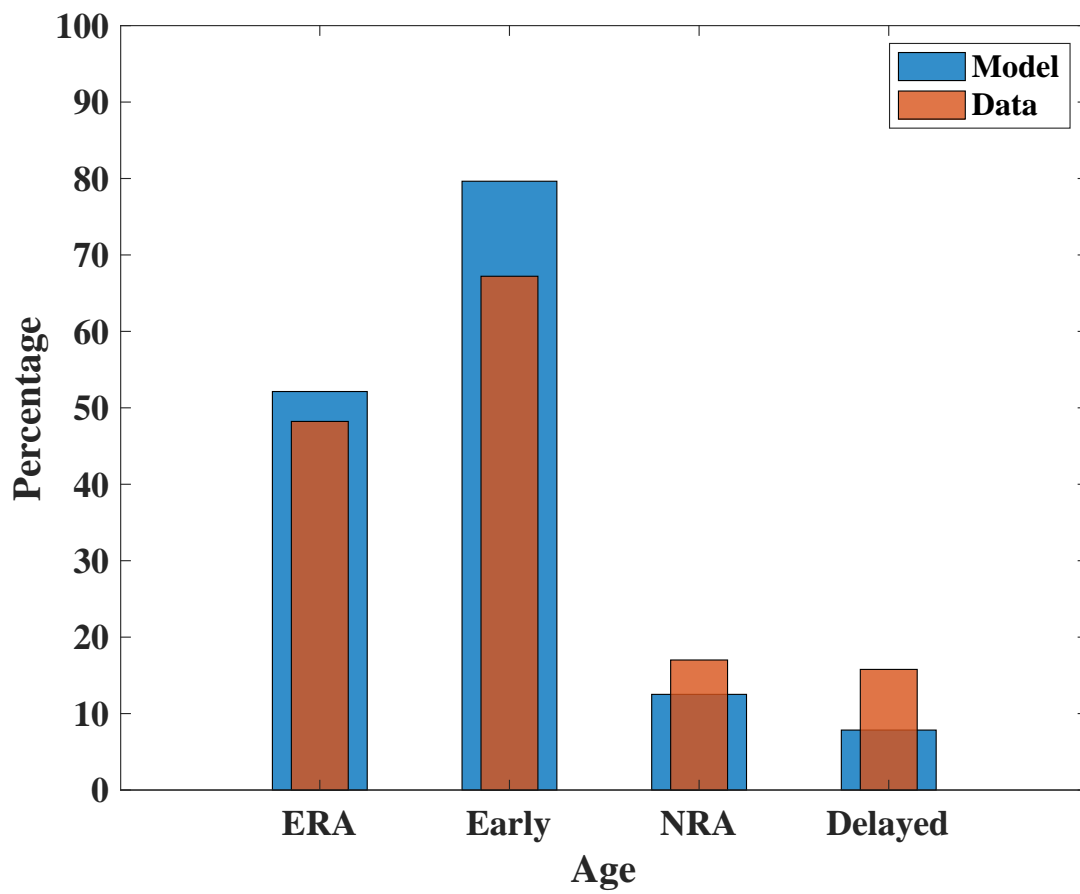
Figure B.4: Life cycle Hours: Universal Transfer Reform
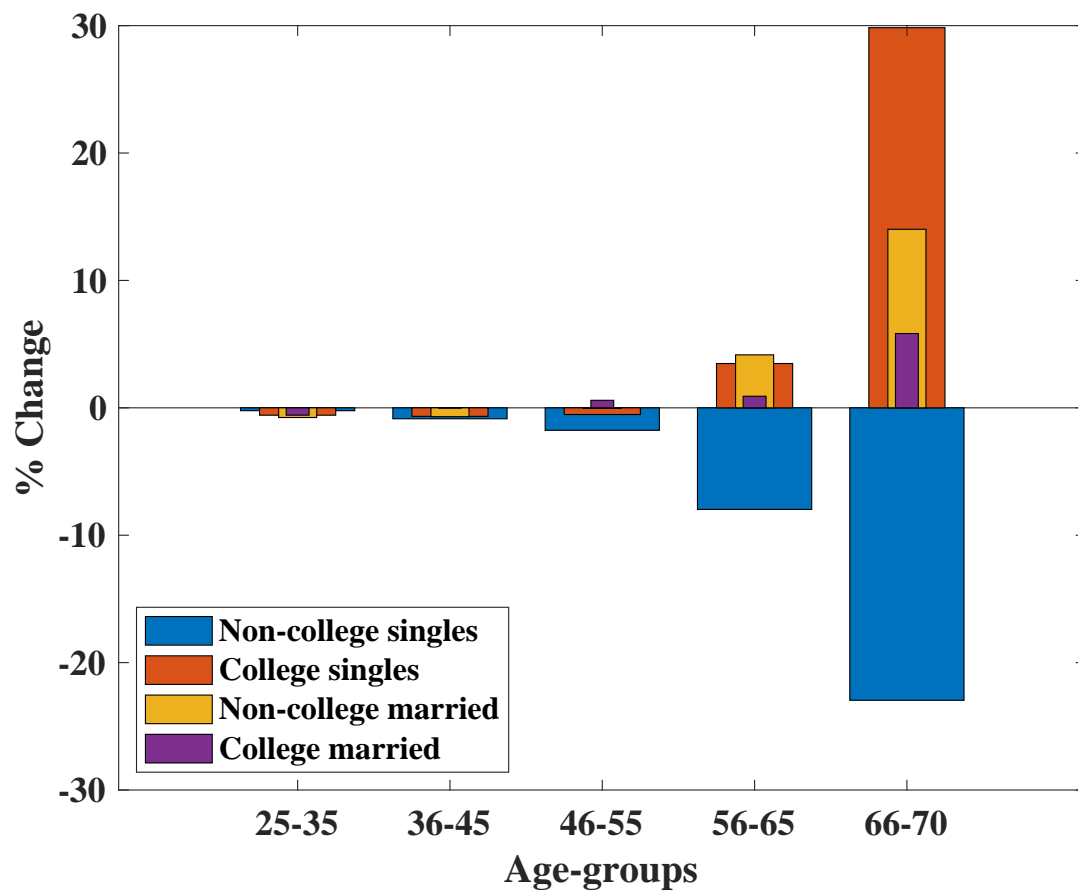
Figure B.5: Life cycle Assets: Universal Transfer Reform
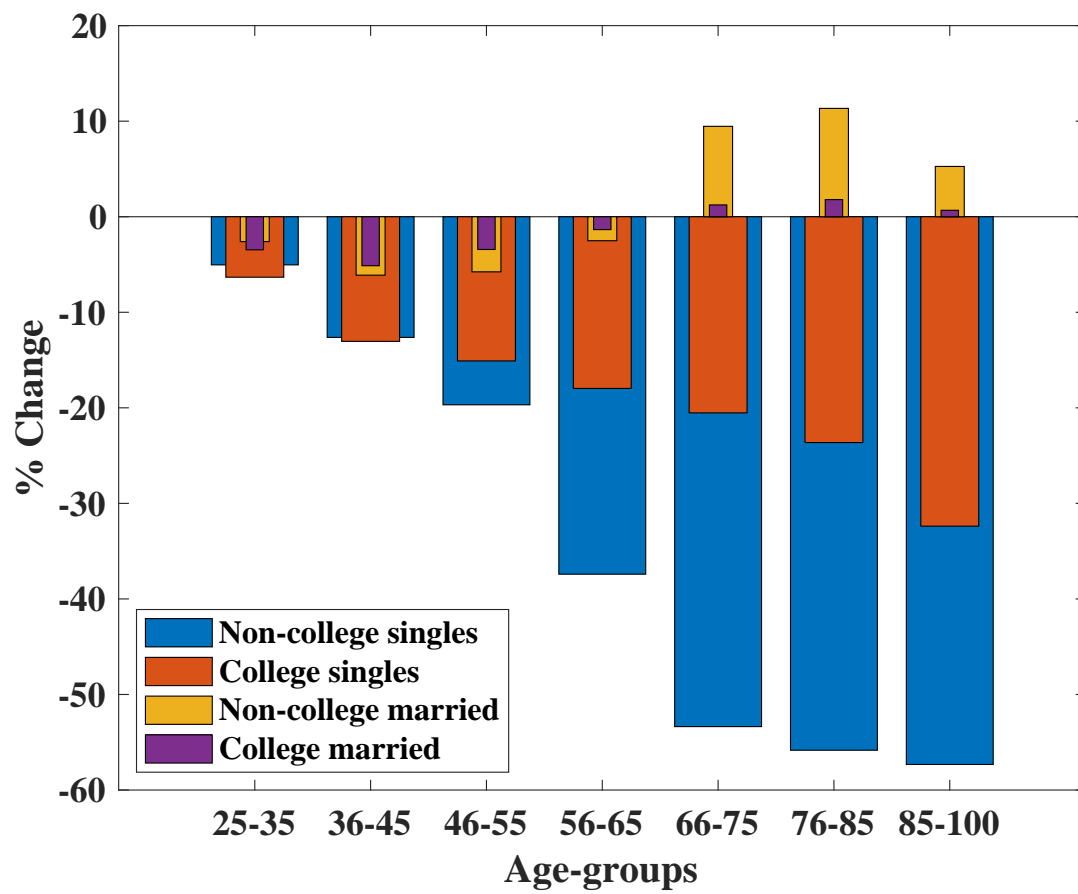
Figure B.6: Life cycle Consumption: Universal Transfer Reform
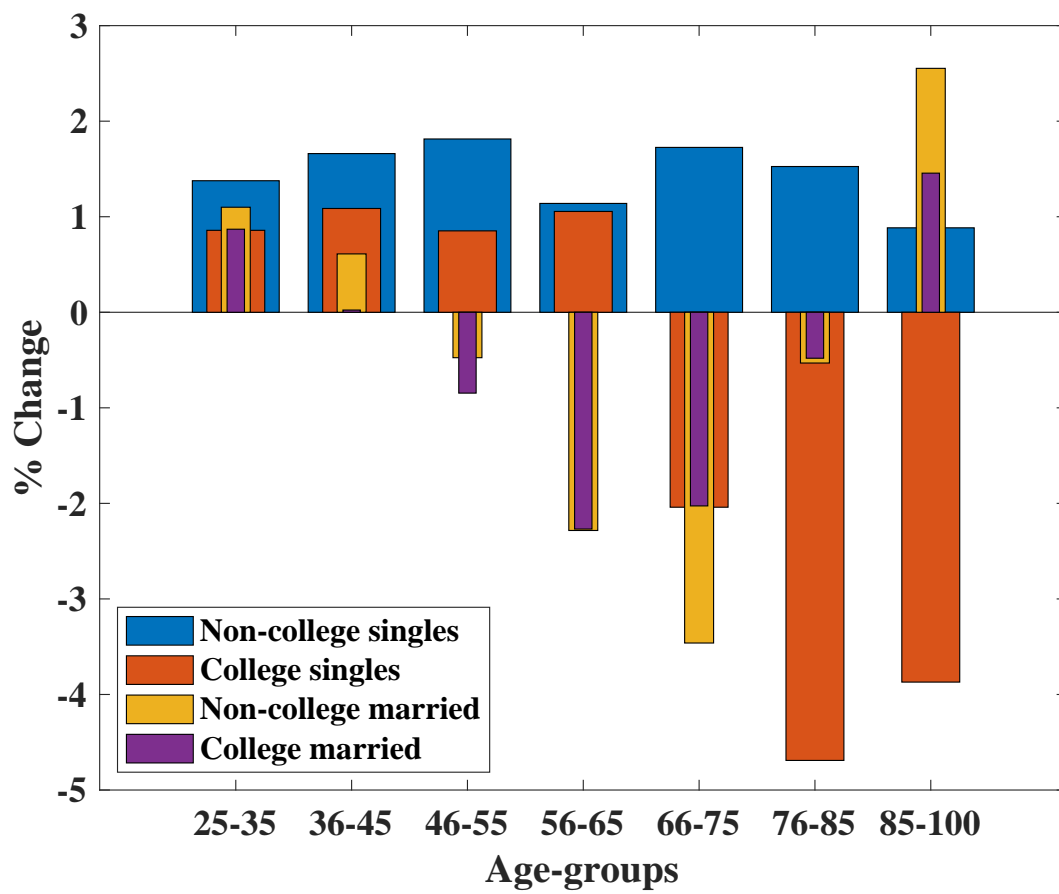
Figure B.7: Social Security Claiming Behavior: Universal Transfer Reform

(a) Non-College, Non Married
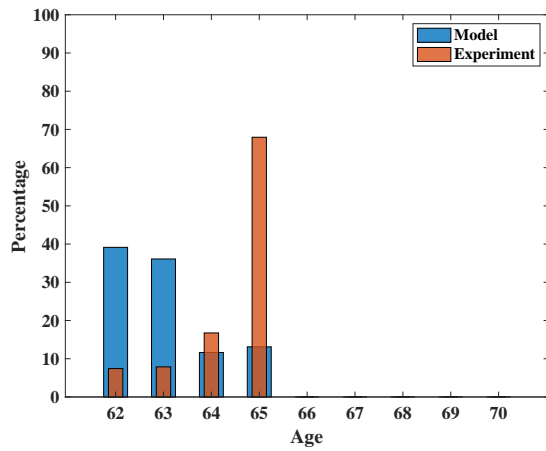
(b) College, Non Married
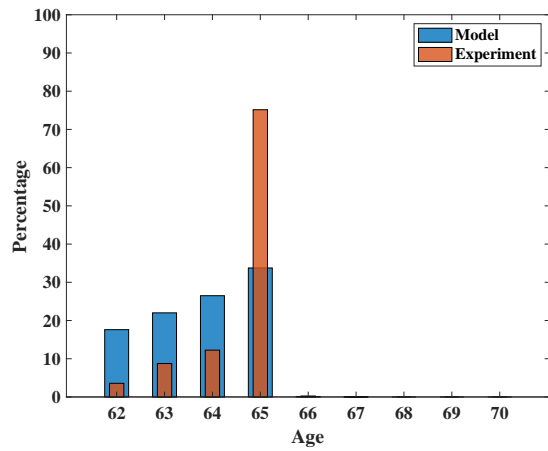
(c) Non-College, Married
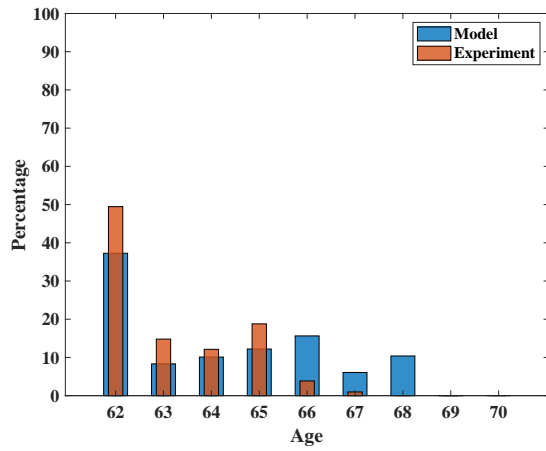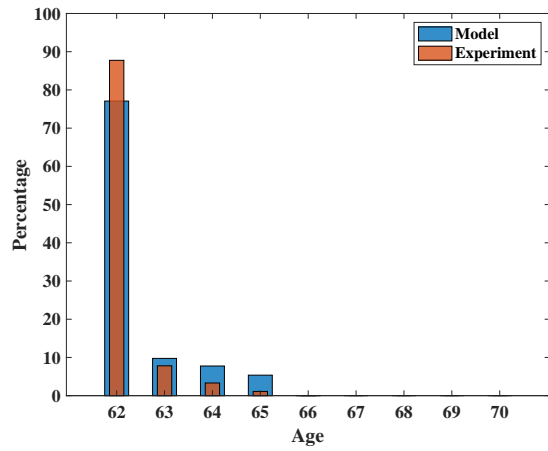
(d) College, Married

Figure B.8: Flow SS Benefits by Age: Universal Transfer Reform

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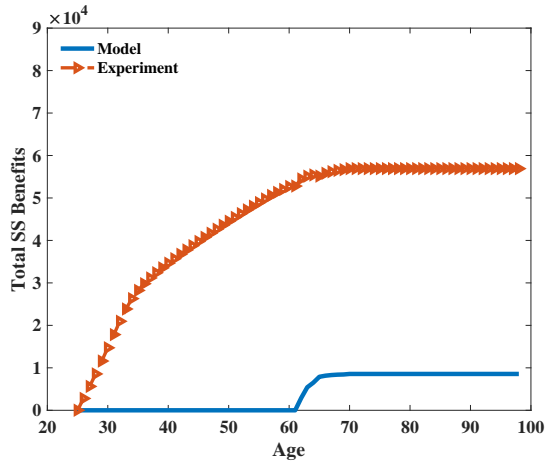
(b) College, Non Married
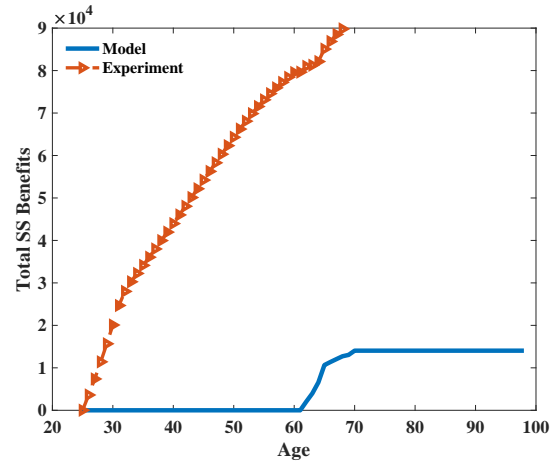
(c) Non-College, Married
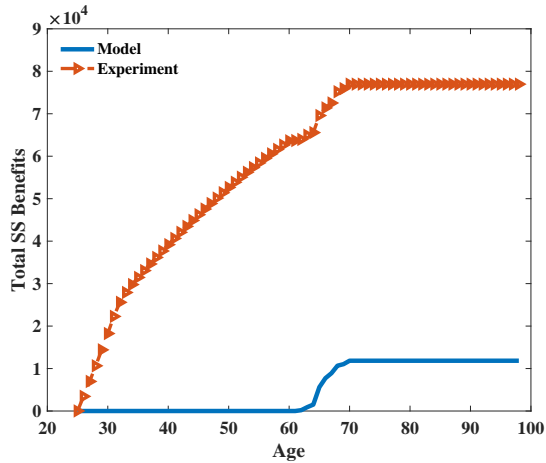
(d) College, Married
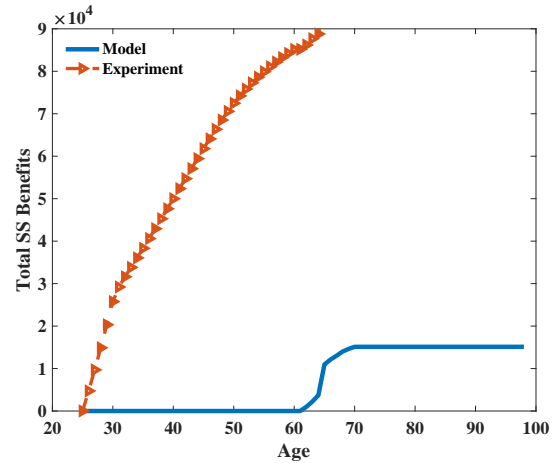
Notes: The figure plots the size of annual Social Security benefits by age, conditional on dying at the end of that period. In the benchmark, the flow is zero until upon claiming. In the counterfactual experiment, individuals receive a lump sum upon dying which is 5.5 times the PIA at the dying age.

Figure B.9: Present Value of SS Benefits Received: Current Policy and Universal Transfer Reform
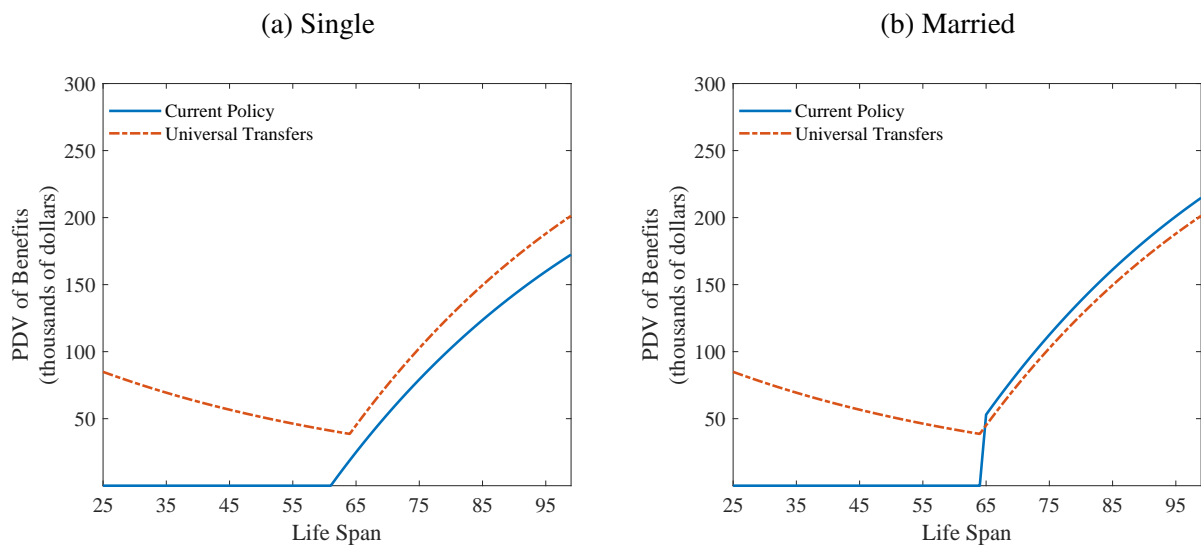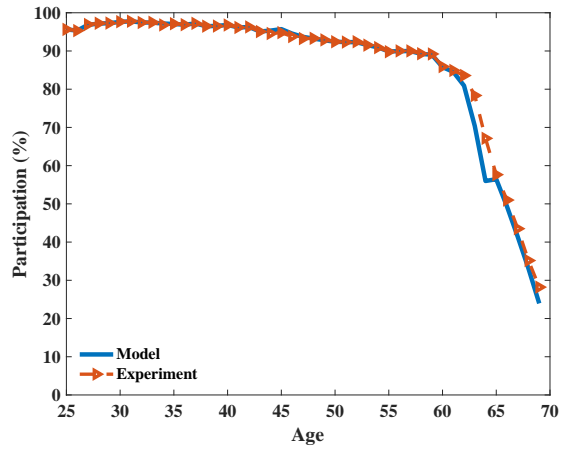
(a) Single

(b) Married

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 B.11: Present Value of SS Benefits: Alternative Social Security Claiming Ages

(a) Delay SS Claiming from Age 62 to Age 65
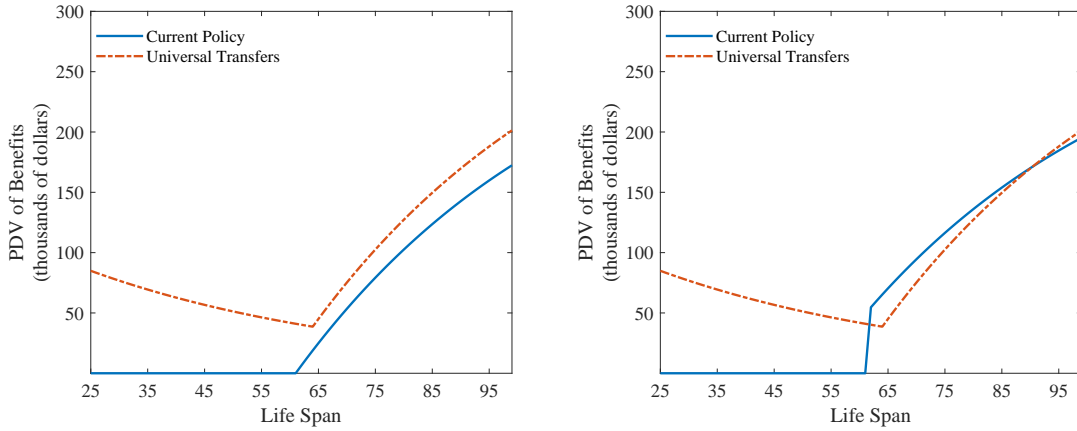
(b) Delay SS Claiming from Age 62 to Age 70
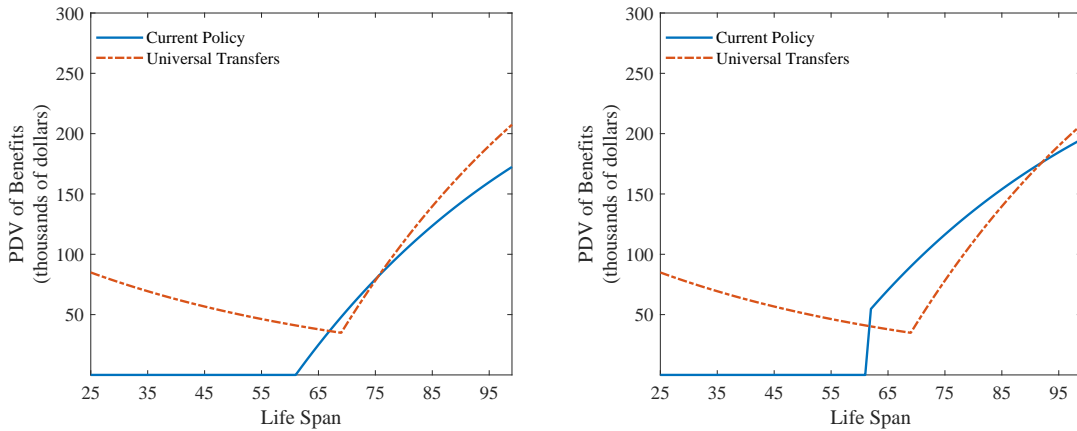
(c) Delay SS Claiming from Age 65 to Age 70
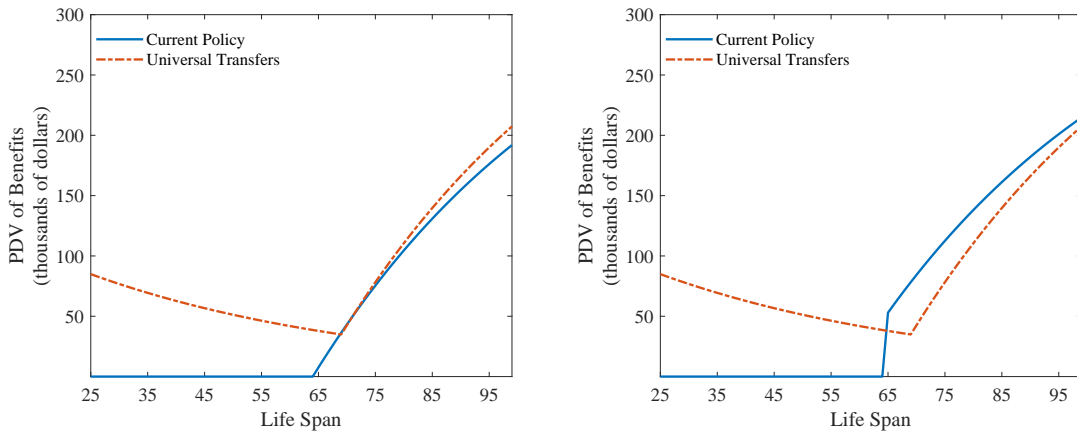
Notes: The figure plots the present value at age 25 of SS benefits received (old-age + survivors benefits) by lifespan. This calculation is done for a fixed AIME of $52,000 annually and discounted at a rate of 2%. In the current policy, married couples receive Survivors Benefits equal to 6 years of PIA 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%.

Figure B.13: Life Cycle Labor Supply: Non Universal Transfer Reform

(a) Participation: Non-College, Singles
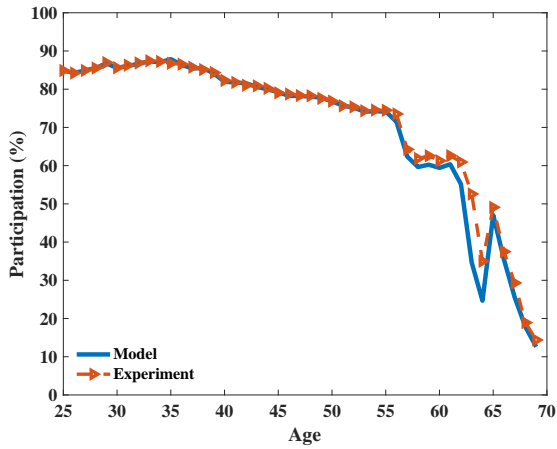
(b) Participation: College, Singles

(c) Hours: Non-College, Singles
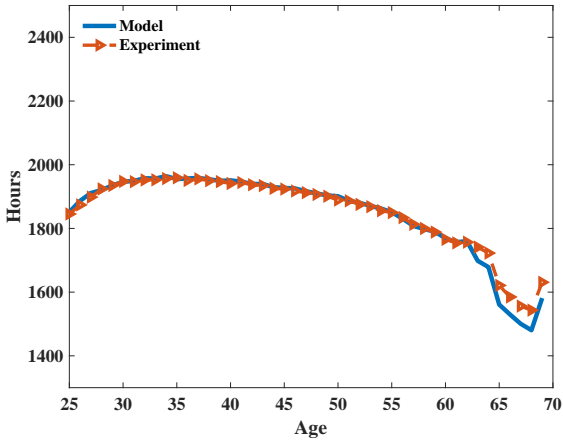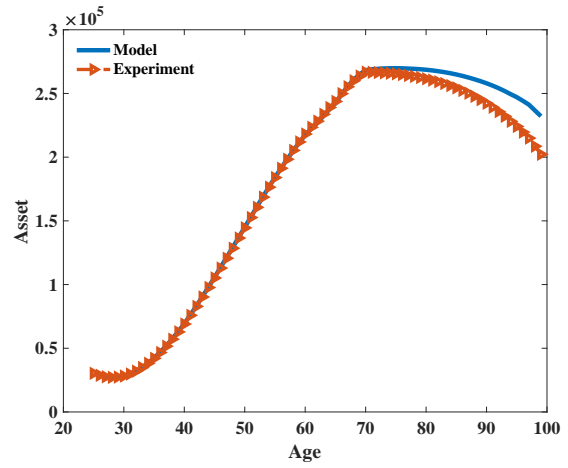
(d) Hours: College, Singles
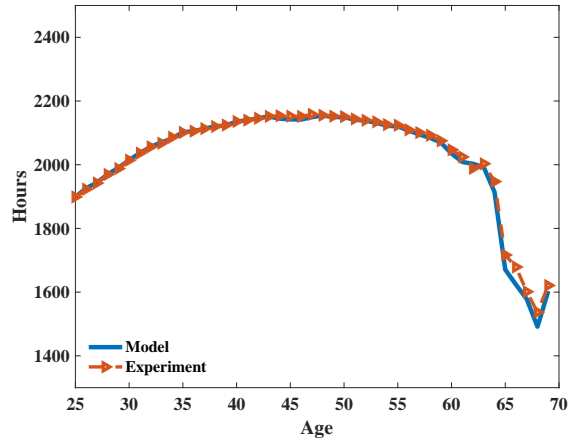
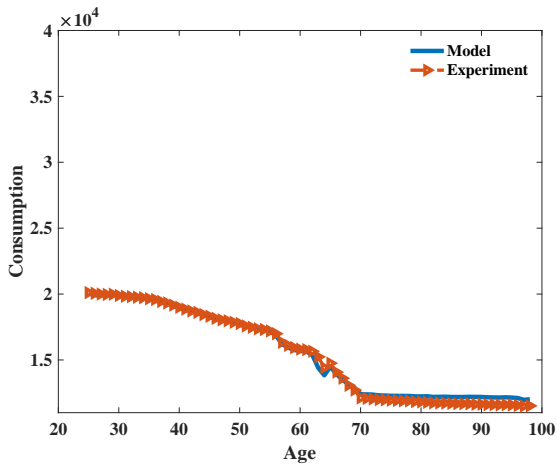Figure B.14: Life cycle Assets and Consumption: Non Universal Transfer Reform
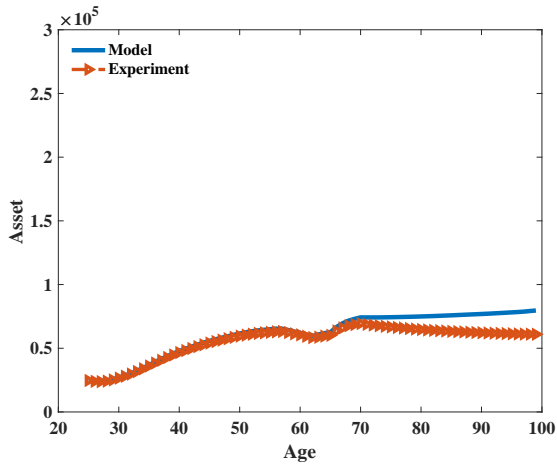
(a) Asset: Non-College, Singles

(b) Asset: College, Singles

(c) Consumption: Non-College, Singles
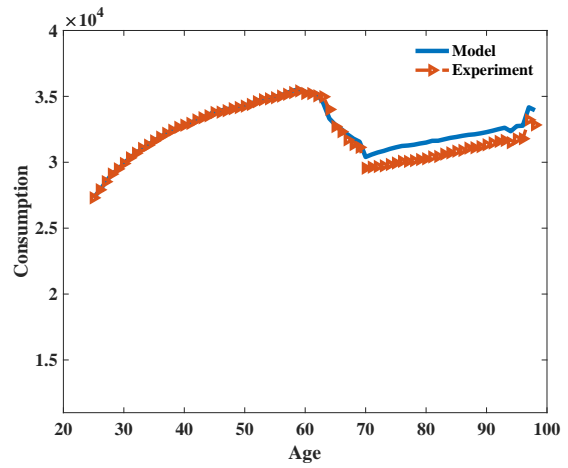
(d) Consumption: College, Singles

Figure B.15: Percent Change in Assets and Consumption: Non Universal Transfer Reform
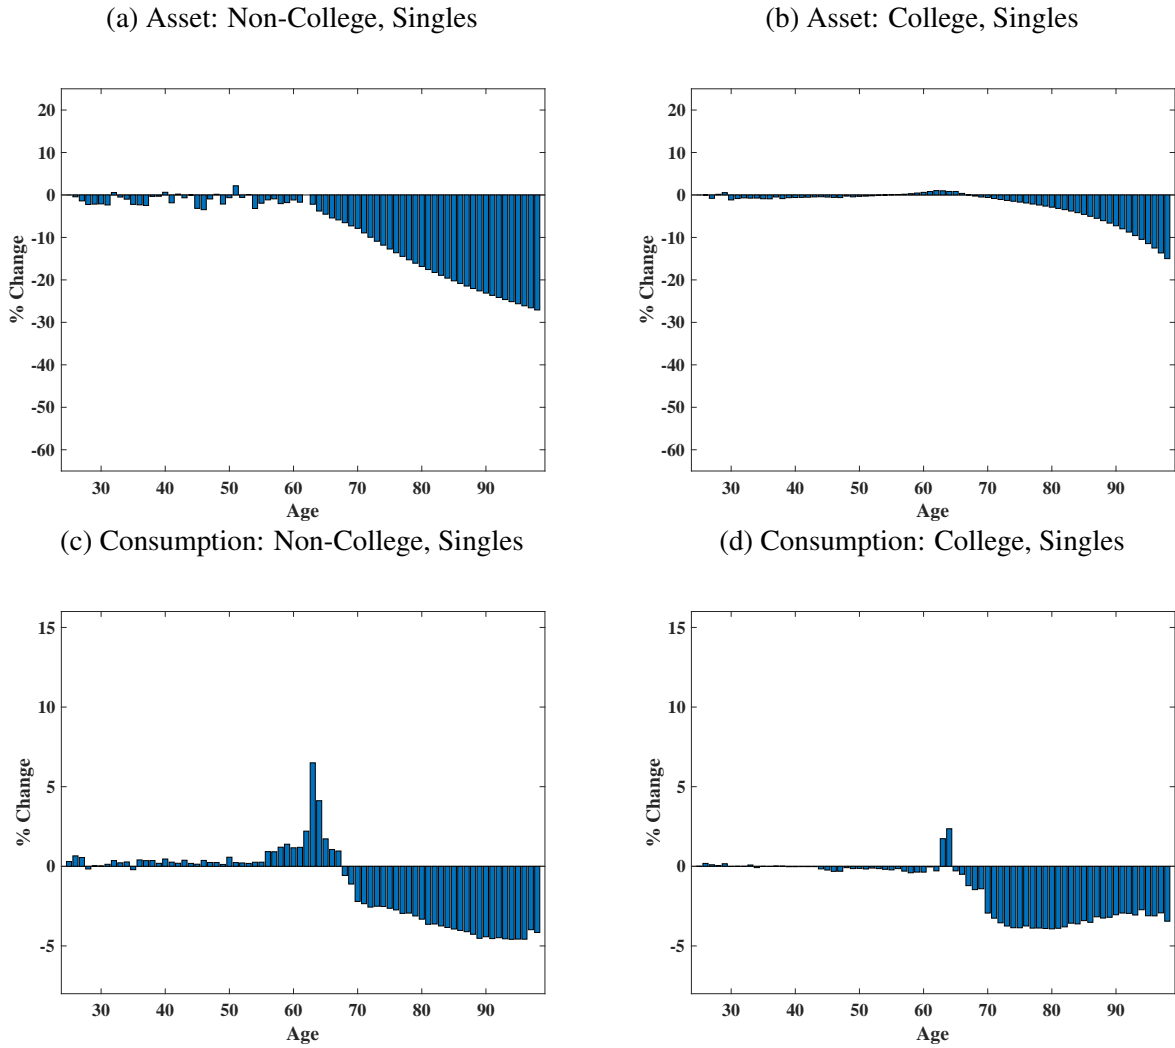
(a) Asset: Non-College, Singles

(b) Asset: College, Singles

(c) Consumption: Non-College, Singles

(d) Consumption: College, Singles

Figure B.16: Social Security Claiming Behavior: Non Universal Transfer Reform

(a) Non-College, Singles
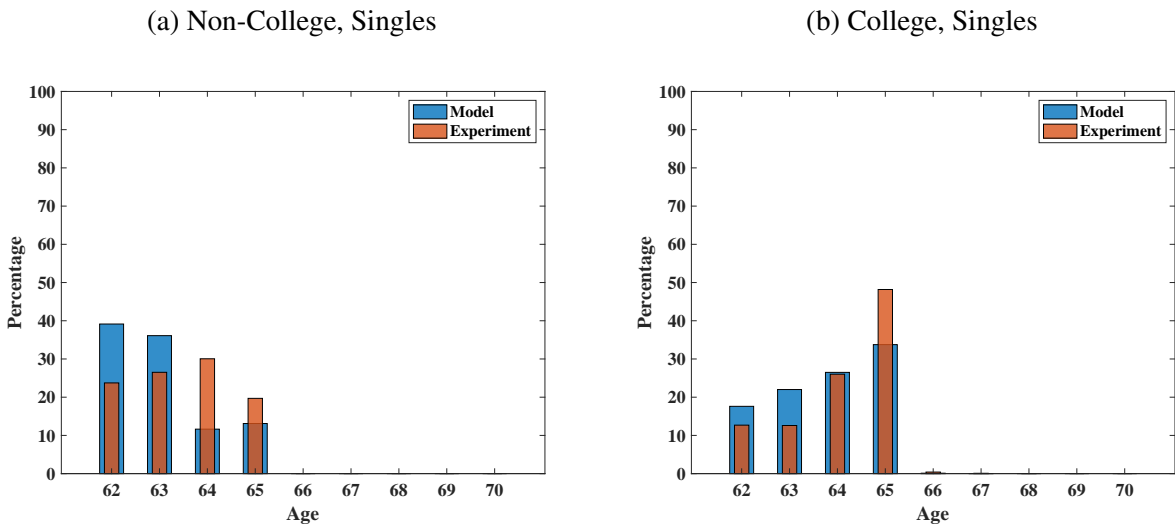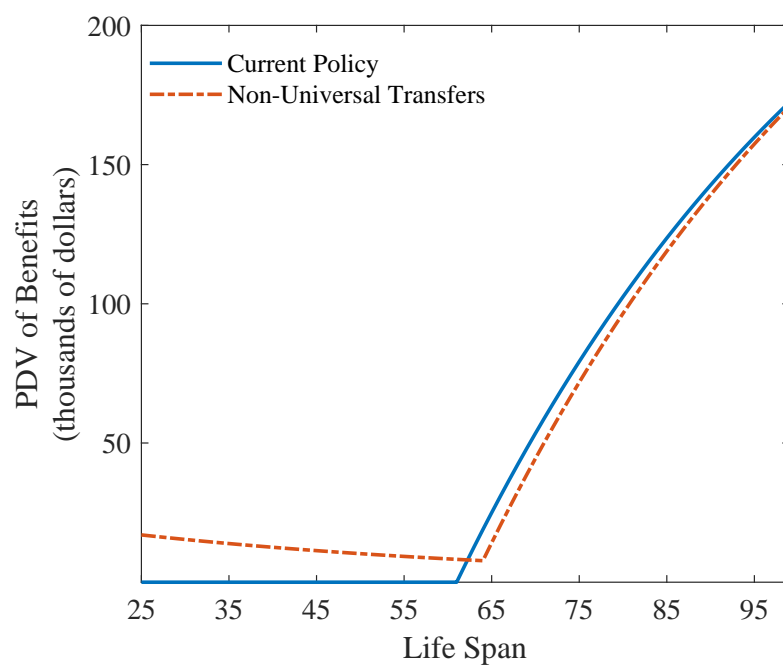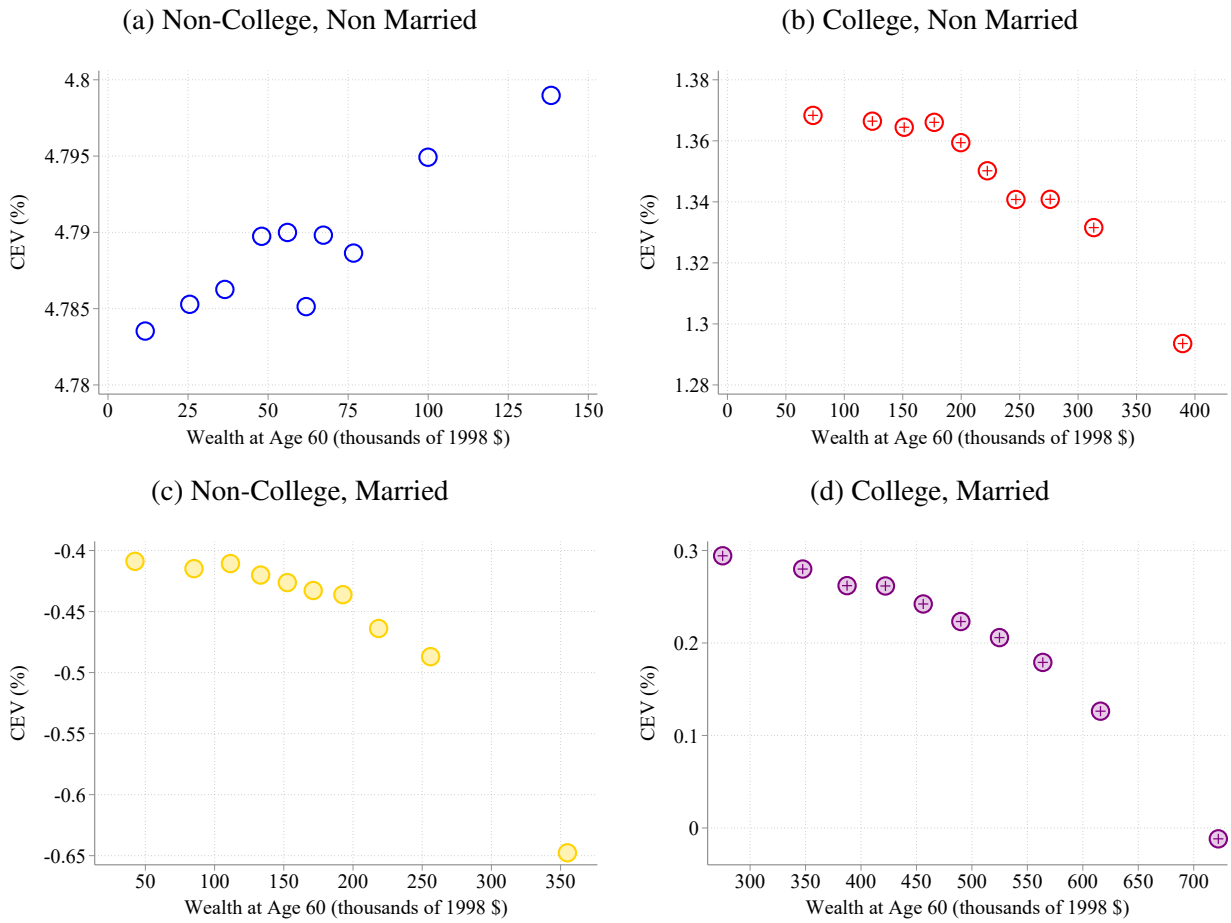
(b) College, Singles

Figure B.17: Present Value of SS Benefits Received: Current Policy and Non Universal Transfer Reform

Notes: Panel (a) presents CEV of all-singles, college-singles and non-college singles for a policy which pays out a lump

Figure B.18: Average Welfare by Wealth Decile

Notes: Each figure plots the average CEV by wealth decile for each education-marital status group