

De Nardi, Mariacristina; French, Eric; Jones, John Bailey; McGee, Rory

Working Paper

Why do couples and singles save during retirement?

IFS Working Paper, No. W21/12

Provided in Cooperation with:

The Institute for Fiscal Studies (IFS), London

Suggested Citation: De Nardi, Mariacristina; French, Eric; Jones, John Bailey; McGee, Rory (2021) : Why do couples and singles save during retirement?, IFS Working Paper, No. W21/12, Institute for Fiscal Studies (IFS), London,
<https://doi.org/10.1920/wp.ifs.2021.1221>

This Version is available at:

<https://hdl.handle.net/10419/242911>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Mariacristina De Nardi
Eric French
John Bailey Jones
Rory McGee

21/12

Working paper

Why do couples and singles save during retirement?

Why Do Couples and Singles Save During Retirement?

Mariacristina De Nardi, Eric French, John Bailey Jones
and Rory McGee*

May 13, 2021

Abstract

While the savings of retired singles tend to fall with age, those of retired couples tend to rise. We estimate a rich model of retired singles and couples with bequest motives and uncertain longevity and medical expenses. Our estimates imply that while medical expenses are an important driver of the savings of middle-income singles, bequest motives matter for couples and high-income singles, and generate transfers to non-spousal heirs whenever a household member dies. The interaction of medical expenses and bequest motives is a crucial determinant of savings for all retirees. Hence, to understand savings, it is important to model household structure, medical expenses, and bequest motives.

*De Nardi: University of Minnesota, Federal Reserve Bank of Minneapolis, CEPR, and NBER. French: University of Cambridge, CEPR, and IFS. Jones: Federal Reserve Bank of Richmond. McGee: UWO and IFS. We thank Wilbert van der Klaauw, Monica Costa Dias, Tom Crossley, Ami Ko, Melanie Lührmann, Tatyana Koreshkova, Victor Rios-Rull and many seminar participants for helpful comments and Karolos Arapakis, Justin Kirschner, Jeremy McCauley, Ali Merali and Rachel Rodgers for research assistance. We gratefully acknowledge support from NSF grant SES-2044748, the Social Security Administration through the Michigan Retirement Research Center (MRRRC grant UM17-15), the Economic and Social Research Council (Centre for Microeconomic Analysis of Public Policy at the Institute for Fiscal Studies (RES-544-28-50001), and Grant Inequality and the insurance value of transfers across the life cycle (ES/P001831/1)). The views expressed in this paper are those of the authors and not necessarily those of the National Bureau of Economic Research, the Federal Reserve Bank of Minneapolis, the Federal Reserve Bank of Richmond, the IFS, or the CEPR.

1 Introduction

The savings of couples and singles behave very differently over the retirement period. While the savings of singles tend to fall with age, those of retired couples tend to increase or stay constant until one of the spouses dies. Why is this the case? Do couples face different medical spending and longevity risks? How important are risks to the surviving spouse? Do couples have different attitudes toward bequests? Answering these questions helps understand the savings of all retirees. Understanding why households save, in turn, is essential to predict how households would react to policy reforms and other changes.

Previous work has mainly focused on singles. These studies find that medical expenses are an important reason why single retirees hold wealth, but that the desire to leave bequests might also be important. Much less is known about the reasons why retired couples save.

Couples differ from singles in several notable ways. Being in a couple allows its members to pool their longevity and medical expense risks, potentially reduces medical spending due to spousal care giving, but also exposes each member to their spouse's risks, including the income loss and high medical expenses that often accompany a spouse's death. A couple also cares about leaving resources to the surviving spouse, and potentially about leaving bequests to other heirs. Disentangling these motivations is important because couples account for about 50% of people age 70 and older, and tend to be richer.

We start our analysis by documenting couples' and singles' savings, medical expenditures, and longevity. We then construct a rich model of retiree saving that incorporates heterogeneity in life expectancy and medical expenses, and includes the large jump in the medical expenses that is typically observed before a death. We model means-tested social insurance and bequest motives, both upon death of the last survivor and upon the death of one of the people in a couple, in which case net worth is optimally split between the surviving spouse and other heirs. We estimate our model using the Method of Simulated Moments (MSM) and Assets and Health Dynamics of the Oldest Old (AHEAD) data. Finally, we use our estimated model to evaluate the extent to which retirees' savings are driven by medical expenses, bequest motives, and their interaction.

Our paper provides multiple contributions. We establish new facts on the savings of singles and couples during retirement: While the savings of retired singles stay roughly constant or fall during retirement, those of couples increase as long as both spouses are alive. Net worth drops sharply (by \$160,000 on average) when the first spouse dies. A large share of this drop is explained by high medical expenses before

death, but not the majority, which is accounted for by transfers to non-spousal heirs. By the time the second spouse dies, much of the couples' wealth has vanished.

We use our estimated model to quantify the effects of various saving motives during retirement for couples and singles. Our results for singles reinforce earlier findings: most singles mainly save to self-insure against future medical expenses, but at the top of the permanent income (PI) distribution, bequests are also important, especially when interacting with medical expense risk.

Our results for couples reveal several new findings. First, bequest motives have large effects on the savings of older couples in the top two permanent income terciles. That is, at age 84, couples in the highest PI tercile would hold 26% less wealth (that is, \$325,000 instead of \$437,000) if they did not have bequest motives. Couples in the middle PI tercile would hold about 14% less (\$172,000 instead of almost \$200,000). In contrast, couples at the bottom, just like singles at the bottom and median PI level, would barely change their wealth. Thus, more couples than singles save to leave bequests, and bequest motives affect their savings by a much larger extent. This is for at least two reasons. First, couples are wealthier and we estimate bequests to be a luxury good. Second, couples not only receive direct utility from leaving bequests at the time of one spouse's death, but also indirect utility from the expectation that the last surviving spouse will leave bequests at his or her death.

Second, in the presence of bequest motives, medical expenses have a smaller effect on the saving of couples at all permanent income levels. For instance, at age 84, the median wealth of couples in the top PI tercile would fall by less than 10% (from \$437,000 to \$417,000) if they did not face any medical expenses. This suggests that medical expenses are only one of the reasons why they continue to save at very advanced ages.

Third, while the savings of couples are driven more by bequests than by medical expenses, the interaction of medical expenses (including for the surviving spouse) and bequest motives has large effects. At age 84, the median wealth of couples in the top PI tercile would fall by more than 65% (from \$437,000 to \$150,000) if they did not face any medical expenses *and* had no bequest motives.

Because wealth saved for bequests also help insure against medical spending, the incremental effects of bequest motives or medical spending, even when summed, can be much smaller than their joint effect. This is due to the way bequest and precautionary saving motives interact. When medical expenses are uncertain, households that reduce their consumption to finance future medical expenses may die with unspent wealth. By giving these unspent funds value, bequest motives reduce the opportunity cost of self-insuring through saving. Moreover, medical expenses introduce uncertainty over whether a household will actually leave a bequest. This might

generate precautionary saving aimed at increasing the likelihood that a bequest is made. A third way in which these motives interact is that medical expenses have a negative wealth effect, which tends to decrease both savings and bequests.

Fourth, couples save to provide resources to the surviving spouse after the death of the first spouse. This is important because there are considerable differences in the life expectancy of each spouse. At age 70 married women and men expect to live 15.8 years and 11.5 years, respectively. In addition, while the average life expectancy of the couple is 13.7 years, the oldest survivor expects to live for 17.9 years. During their remaining years alone, survivors face higher medical spending (including from an elevated risk of entering a nursing home) and receive lower income. As a result, couples wish to leave resources to the surviving spouse upon a partner's death. This has a large effect on savings at all PI levels. If couples instead placed no weight on the survivor, at age 84 their median wealth in the middle PI tercile would fall by more than 17% (from \$199,000 to \$165,000). Studies omitting this risk will underestimate this precautionary saving motive.

Fifth, when one member of a high-PI couple dies, even though most of the estate is usually left to surviving spouse, large bequests to non-spousal heirs are often distributed. Changes to saving motives thus affect both the saving of couples and singles and the division of wealth between the surviving spouse and non-spousal heirs. For example, when bequest motives are removed, couples save less, but when one spouse dies all the of estate is left to the survivor. In fact, we find that removing bequest motives often increases the amount of wealth held by widows and widowers.

Taken together, our results show that bequest motives play a larger role for the savings of couples than those of singles and for the savings of high-PI households than those of low-PI households, and that the interaction of precautionary savings and bequests are most important for couples and high-PI households. In contrast, providing for the surviving spouse is important for all couples.

We reach similar conclusions when we look at all households, couples and singles, together. For median wealth, and at lower percentiles, the effects of removing medical spending exceed the effects of eliminating bequest motives. At higher asset percentiles, and for the mean, bequest motives have larger effects. The interaction of bequest and precautionary saving motives is large across the entire asset distribution.

The rest of the paper is organized as follows. Section 2, summarizes the findings of key related papers and motivates our main modeling choices. Section 3 briefly introduces our data and displays the key facts that we seek to understand. Section 4, introduces our model and section 5 details our estimation procedure. Section 6 displays important features of the data affecting households' saving decisions and evaluates our model fit. To build additional trust in the predictions of our esti-

mated model Section 7 evaluates the performance of our model in terms of important additional aspects of the data that we do match at the estimation stage. Section 8 evaluates the impact of medical expenses and bequest motives for singles and couples and Section 9 concludes.

2 Related literature and modeling choices

The goal of this paper is to document and explain the differences in the saving behavior of couples and singles and their contribution to aggregate retirement savings. Although much of the structural work on retiree saving has focused on singles,¹ a few papers, including Braun, Kopecky and Koreshkova [9] and Nakajima and Telyukova [43], include couples.

Our paper builds on these contributions by modeling what happens when a spouse dies. As Poterba, Venti, and Wise [49], we find that net worth falls when households lose a spouse and that this fall is too large to be attributable to medical and other end-of-life expenses. A novel aspect of our framework is that we allow households to leave bequests to non-spousal heirs when only one of the spouses dies. This feature proves important in matching the fall in wealth observed at the time of the first spouse's death.

Bequests to heirs also occur when the final member of a household dies. The desire to leave such bequests has received considerable attention as a potential explanation why households retain high levels of wealth at very old ages, for instance by Ameriks, Caplin, Laufer, and Van Nieuwerburg [4], Ameriks, Briggs, Caplin, Shapiro, and Tonetti [3], Dynan, Skinner, and Zeldes [23], and Lockwood [41]. In addition, De Nardi, French, and Jones [19] show this kind of bequest motives help fit the asset and Medicaid recipiency profiles of singles.

Besides modeling household demographic transitions and bequest motives towards the surviving spouse and other heirs, we take into account health and medical expense risk, social insurance for couples and singles, and heterogeneity in life expectancy by PI and marital status. These features are important because households face potentially large out-of-pocket medical and nursing home expenses (Braun, Kopecky, and Koreshkova [9], French and Jones [27, 28], Palumbo [47], Feenberg and Skinner [25], and Marshall, McGarry, and Skinner [42]), which generate precautionary savings (as in Kopecky and Koreshkova [38] and Laitner Silverman, and Stolyarov [39]). But these risks are partially insured by means-tested programs such as Medicaid and Supplemental Security Income, which provide strong saving dis-

¹De Nardi, French and Jones [20] provide a detailed literature review.

incentives (Hubbard, Skinner and Zeldes [32] De Nardi, French, and Jones [18]). Finally, previous work has shown that high-income individuals live longer than low-income individuals (Attanasio and Emmerson [6], Rios Rull and Pijoan Mas) and that married people live longer than single people.

3 Savings, medical spending, and Medicaid

3.1 The AHEAD dataset

We use the AHEAD dataset, which begins with households of non-institutionalized individuals age 70 or older in 1993/94 and subsequently surveys their survivors every two years. Appendix A describes the details of our sample selection and data work. Three important decisions are that (i) to abstract from labor supply decisions, we only consider retired households; (ii) to be consistent with the transitions in the model we drop household who either get married or divorced during the sample period; and (iii) due to the well-known under-reporting of assets and medical spending in the initial wave, we drop the 1993/94 data. Our 1996 data include 4,634 households, of whom 1,388 are couples and 3,246 are singles. This represents 24,274 household-year observations for which at least one household member was alive.

We distinguish three types of household: single man, single woman, and couple. Single people include those who are divorced, never married, or widowed when first observed, and those who become widowed over the sample period.

3.2 Permanent Income

At retirement time, permanent income (PI) captures substantial household ex ante heterogeneity: Households with different PI ranks receive different flows of retirement income and face different processes for health, mortality and medical expenses. A contribution of this paper is distinguishing heterogeneity from risks.

To estimate a household's PI, we first sum all of its annuitized income sources (Social Security benefits, defined benefit pension benefits, veterans benefits and annuities).² Because there is a roughly monotonic relationship between lifetime earnings and our annuity income measure, this measure is also a good indicator of its income during the working period. We then construct a PI measure comparable across households of different ages and sizes. To do so, we regress annuity income on a household fixed effect, dummies for household structure, a polynomial in age,

²Since we model means-tested social insurance from SSI and Medicaid explicitly in our model (through a consumption floor), we do not include SSI transfers.

and interactions between these variables. The rank order of each household's estimated fixed effect provides our measure of its PI. This is a time-invariant measure that follows the household even after one of its members dies. See Appendix B for details.

3.3 Savings, medical expenses, and bequest

Our measure of wealth (or net worth, or assets) is the sum of all assets less mortgages and other debts.³ Figure 1 displays median net worth, conditional on birth cohort and PI tercile, for different configurations of couples and singles. We break the data into 4 cohorts containing people who in 1996 were ages 72-77, 78-83, 84-89, and 90-102. For clarity we display data for two cohorts in this picture (Appendix I displays all of our data, as well as model fit). To construct these profiles, we calculate medians for each cohort for those alive in each calendar year (an unbalanced panel). The line for each cohort starts at the cohort's average age in 1996.

Panel (a) displays net worth for current singles, that is all people who are single at a given age and in a given cohort. It highlights that the savings of elderly singles depend on their income (which is predetermined at retirement). Individuals with the lowest lifetime incomes reach retirement with little wealth and run them down as they age. After that point, they rely on annuitized income and government insurance to fund their consumption and medical expenses. Elderly singles in the middle and top of the income distribution also run down their wealth as they age, but very slowly, and carry some wealth into very old ages.

Panel (b) shows the net worth profiles that would arise if we were to limit the sample to only people who were initially single, without including those who become single upon being widowed. Because the newly single tend to have more wealth, the trajectories in Panel (a) slope down less than those in Panel (b); the two sets of trajectories are nonetheless quite similar.

Panel (c) of Figure 1 plots the net worth of current couples, including couples as long as they remain intact but dropping them when they become single. Panel (c) reveals several interesting results. First, relative to singles in the same PI tercile, couples in the lowest PI tercile reach retirement with more wealth. Second, the asset profiles of the lower- and middle-income surviving couples display no decumulation. Third, couples in the top PI tercile increase their net worth until almost age 90.

³The AHEAD has information on the value of housing and real estate, autos, liquid assets (which include money market accounts, savings accounts, Treasury bills, etc.), individual retirement accounts, Keoghs, stocks, the value of a farm or business, mutual funds, bonds, and other assets.




Figure 1: Median wealth, AHEAD data. Panel (a): current singles. Panel (b): initial singles. Panel (c): current couples. Panel (d): initial couples. Each line shows medians for a cohort-PI cell over the period 1996-2014, plotted against average cohort ages. Thicker lines denote higher PI groups.

Panel (d) of Figure 1 reports median wealth for the population of those who are initially in a couple but might have subsequently lost a spouse, in which case we report the wealth of the surviving spouse. Relative to intact couples, initial couples in the higher PI terciles decumulate their wealth more quickly. Some of this decumulation may reflect the higher decumulation rates of singles, but some is also due to the loss of wealth that occurs when one of the spouses dies.

To formalize and quantify this observation, we study the difference in net worth between two sets of couples: those in which one spouse dies and those that are similar in multiple dimensions 6 years prior to death, except for the fact that they do not experience a spousal death during a certain time period.⁴ This matching analysis accounts for observable initial heterogeneity across couples and thus identifies the impact of death and the declining health that precedes it.

We use these data to estimate the following event study specification:

$$a_{i,t} = f_i + \sum_{j=-4}^4 (g_j + d_j D_i) \times 1\{t - T_i = j\} + e_{i,t} \quad (1)$$

where $a_{i,t}$ denotes the wealth of household i in calendar year t , and D_i indicates whether i belongs to the group that does not lose a spouse ($D_i = 0$) or to the group that does ($D_i = 1$). We normalize the date of death, T_i , to occur at year 0 and follow the households for three waves before and two waves after the death. The coefficients of interest are the parameters $\{d_j\}$, which show the extent to which wealth rise or fall for households who experience the death of a spouse, relative to households in the matched sample who experience only a placebo death at date T_i .⁵

Panel (a) of Figure 2 reports our estimated death-related asset decline (the values of d_j). On average, the wealth of households experiencing a death declines by an additional \$160,000 compared to those not experiencing it. While the decline begins up to 6 years in advance of death, a substantial share of the decline occurs in the final two years. Part of this decline reflects a gradual worsening of health and a related increase in medical expenses. The rest reflects bequests made at the time of death of the first spouse and changes in consumption. The gap between the treatment and control groups continues to widen during the post-death years. This may reflect additional death-related effects, for instance the loss of a spouse's income, or other differences between intact couples and singles.

⁴We match on PI, initial wealth, and age. See Appendix C for details and Jones, De Nardi, French, McGee and Rodgers [35] for additional outcomes.

⁵The coefficients g_j and d_j are defined relative to their value 6 years before death (the omitted indicator category). In interpreting the summation, recall that AHEAD interviews occur every other year.

Some of these asset declines are explained by high medical expenses prior to death. The AHEAD contains high-quality data on what the household spends out-of-pocket on private insurance premia, drug costs, hospital stays, nursing home care, doctor visits, dental visits, and outpatient care, including those incurred during the last year of life. To construct our measure of out-of-pocket medical spending, we sum all of these components, along with funeral, other end-of-life expenses, and imputed Medicare insurance contributions.

Panel (b) of Figure 2 shows that when a member of the household dies, there is a sharp increase in medical expenses: they are \$8,000 larger during each of the two years preceding the death of a spouse, compared with similar couples experiencing no death. Over the 6 years preceding death, the total difference in medical expenses between these two households surpasses \$27,000. Thus high end-of-life medical expenses can explain $\frac{27,000}{160,000} = 17\%$ of the average fall in wealth. The effect for singles is even larger, totalling almost \$40,000.




Figure 2: Average changes around the death of a spouse, initial couples, AHEAD/MCBS data. Panel (a): Mean asset decline around the death of a spouse. Panel (b): Mean increase in the sum of out-of-pocket medical and death related expenses. Solid lines: estimates, shaded region: 95% confidence interval. Death dates are centered at year 0.

If not medical and other end-of-life expenses, what explains the drop in wealth at the time of death of a spouse? When one spouse dies, a family member of the deceased, usually the spouse, is asked how the estate is split between the surviving spouse and other heirs, such as children. Comparing the wealth of the newly single with non-spousal transfers shows that most of the wealth is left to the surviving

spouse, but that, on average \$79,000 are transferred to other heirs around the time of death of the first spouse.

Transfers to other heirs at the time of death of a spouse are large and support the view that bequest motives are important. Although we do not require our model to match the decline in wealth around the death of a spouse or transfers to other heirs, our estimated structural model matches these aspects of the data very well, and thus helps validate its predictions.

3.4 Medicaid recipiency




Figure 3: Medicaid recipiency by age, income, and cohort.

Because previous work has shown that means-tested social insurance programs affects saving, it is important to match the extent to which retirees use such programs. Figure 3 plots Medicaid recipiency rates conditional on age, PI tercile, and cohort for all households in our sample (the AHEAD has reliable measures of Medicaid recipiency, but not spending). It shows that the lowest-income retirees are most likely to end up on Medicaid and that, and consistently with the asset-tested nature of these programs (high-asset households are not eligible) and the asset profiles shown in Figure 1, Medicaid recipiency increases with age.

4 The model

A retired household maximizes utility by choosing savings, current and future consumption, and, upon death of his or her spouse, the split of wealth between

himself or herself and other heirs. These choices also determine final bequests upon the death of the last household member.

Households make these choices at household head age t , $t = t_r, t_r + 1, \dots, T + 1$, where t_r is the initial period that we consider and T is the maximum potential lifespan. For tractability, we assume that wives are always two years younger than their husbands, so that one age is sufficient to characterize a household. Consistent with the AHEAD data frequency, our time period is two years long.

The household begins the period in one of the following possible marital statuses: a couple, a newly widowed man or woman, or a single man or single woman. Single men, single women, and couples use the cash on hand they have at the beginning of a period to consume and save. After that, mortality, health and medical expense shocks occur, income net of taxes is received, and government transfers take place. At this point, the amount of cash-on-hand that is carried into the next period is known, and the household enters next period.

People who are newly widowed have an additional decision, which is how to divide their estate between themselves and bequests to others. Once that decision is made, they become single men or women, with wealth equal to the remainder of their estates. When the final member of a household dies, all remaining net worth goes to its heirs. Consistent with reality, our timing implies that medical costs associated with death are collected before any bequests can be made and that Medicaid pays bills that are incurred even when a patient dies.

4.1 Preferences

The per-period utility functions for singles and couples are given by

$$u^S(c) = \frac{c^{1-\nu}}{1-\nu}, \quad u^C(c) = 2 \frac{(c/\eta)^{1-\nu}}{1-\nu}, \quad 1 < \eta \leq 2, \quad \nu \geq 0, \quad (2)$$

respectively, where c is total consumption and the parameter η determines the extent to which couples enjoy economies of scale in the transformation of consumption goods to consumption services. The household weighs future utility with the factor β .

The household derives utility $\theta_j(b)$ from leaving bequest b , where $\theta_0(b)$ is the utility from bequests when there are no surviving members in the household, while $\theta_1(b)$ is the utility from bequests when there is a surviving spouse. It takes the form

$$\theta_j(b) = \phi_j \frac{(b + \kappa_j)^{1-\nu}}{1-\nu}, \quad (3)$$

where κ_j determines the curvature of the bequest function, and ϕ_j determines its intensity. De Nardi [16] used this functional form to help explain savings over the

life cycle. It supports several interpretations of the “bequest motive:” dynastic or “warm glow” altruism (as in Becker and Tomes [7], Abel and Warshawsky [1], or Andreoni [5]); strategic motives (as in Bernheim, Schleifer and Summers [8] or Brown [10]); or some form of utility from wealth itself, as in (Carroll [12] and Hurd [33]).

4.2 Sources of uncertainty and budget constraints

4.2.1 Income

Because there is much less income uncertainty during retirement than during the working period, we simplify the model by assuming that the household’s non-asset income at time t , $y_t(\cdot)$, is a deterministic function of the household’s permanent income, I , age, and family structure f_t (single man, single woman, or couple).

$$y_t(\cdot) = y(I, t, f_t). \quad (4)$$

We do not include received bequests as a source of income, because very few households aged 70 or older receive them.

4.2.2 Health and survival uncertainty:

Health and survival are individual, rather than household-level, variables and we use gender, $g \in \{h, w\}$, to differentiate between men and women. A person’s health status, hs^g , indicates whether he or she is in a nursing home, in bad health, or in good health. The transition probabilities for a person’s future health status depend on that person’s current health status, permanent income, age, gender and marital status

$$\pi_t(\cdot) = \Pr(hs_{t+1}|hs_t, I, t, g, f_t). \quad (5)$$

Survival depends on the same variables. Let $s_t(I, g, hs_t, f_t)$ denote the probability that an individual alive at age t survives to age $t + 1$.

4.2.3 Medical expense uncertainty

We use m_{t+1} to denote the sum of medical spending that is either paid out-of-pocket by the household or covered by Medicaid between periods t and $t + 1$. While we treat this total as exogenous, the division of expenses between the household and Medicaid depends on the household’s financial resources, the total expenses that must be covered, and the level of the consumption floor. In other words, m_{t+1} gives

the household's maximum possible medical spending obligation, but because of social insurance the amount faced by poorer households may be much smaller.

We allow m_{t+1} to depend on the health status of each family member at both the beginning and end of the period, permanent income, age, household's family structure (also differentiating single men and single women) at the beginning and end of each period, and an idiosyncratic component, ψ_{t+1} .

$$\begin{aligned}\ln m_{t+1} = & m(hs_t^h, hs_t^w, hs_{t+1}^h, hs_{t+1}^w, I, t+1, f_t, f_{t+1}) \\ & + \sigma(hs_t^h, hs_t^w, hs_{t+1}^h, hs_{t+1}^w, I, t+1, f_t, f_{t+1}) \times \psi_{t+1}.\end{aligned}\quad (6)$$

We normalize the variance of ψ_{t+1} to 1.

Allowing medical expenses to depend on the household's composition and health status at both the beginning of a period (which was realized at the very end of the previous period) and the period's end (which will be carried over into the subsequent period) allows us to capture the jump in medical spending that occurs when a family member dies – that is f_t changes – and to incorporate the impact of two subsequent health realizations on medical spending. The latter better helps us account for the cost of prolonged periods of bad health or nursing home stays. Our timing implies that the value of m_{t+1} is not known to the household when it decides how much to consume between periods t and $t+1$.

Following Feenbreg and Skinner [25] and French and Jones [27], we assume that ψ_{t+1} can be decomposed as

$$\psi_{t+1} = \zeta_{t+1} + \xi_{t+1}, \quad \xi_{t+1} \sim N(0, \sigma_\xi^2), \quad (7)$$

$$\zeta_{t+1} = \rho_m \zeta_t + \epsilon_{t+1}, \quad \epsilon_{t+1} \sim N(0, \sigma_\epsilon^2), \quad (8)$$

where ξ_{t+1} and ϵ_{t+1} are serially and mutually independent. We discretize ξ and ζ , using the methods described in Tauchen [52].

4.2.4 Budget constraints

Let a_t denote net worth at the beginning of period t and r denote their constant pre-tax rate of return. Total post-tax income is given by $\Upsilon(r a_t + y_t(\cdot), \tau_{f_t})$, with the vector τ_{f_t} summarizing the tax code, which depends on family structure. Define the resources available *before* government transfers as

$$\tilde{x}_t = a_t + \Upsilon(r a_t + y_t(\cdot), \tau_{f_t}) - m_t, \quad (9)$$

To capture Medicaid and SSI, we assume that government transfers bridge the gap between a minimum consumption floor and the household's financial resources,

$$tr_t(\tilde{x}_t, f_t) = \max \left\{ 0, c_{min}(f_t) - \tilde{x}_t \right\}, \quad (10)$$

where we allow the guaranteed consumption level c_{min} to vary with family structure.

To save on state variables we follow Deaton [17] and sum the household's financial resources *after* government transfers into cash-on-hand:

$$x_t = a_t + \Upsilon(r a_t + y_t(\cdot), \tau_{f_t}) - m_t + tr_t(\tilde{x}_t, f_t). \quad (11)$$

Households divide their cash-on-hand between consumption and savings:

$$a_{t+1} = x_t - c_t, \quad (12)$$

$$c_t \in [c_{min}(f_t), x_t], \quad \forall t. \quad (13)$$

Equation (13) ensures that consumption is at least as high as the consumption floor and that savings are non-negative.

Next period's resources before transfers and next period's cash-on-hand can then be expressed in terms of this period's cash-on-hand and are, respectively,

$$\tilde{x}_{t+1} = (x_t - c_t) + \Upsilon(r(x_t - c_t) + y_{t+1}(\cdot), \tau_{f_{t+1}}) - m_{t+1}, \quad (14)$$

$$x_{t+1} = x_t - c_t + \Upsilon(r(x_t - c_t) + y_{t+1}(\cdot), \tau_{f_{t+1}}) - m_{t+1} + tr_{t+1}(\tilde{x}_{t+1}, f_{t+1}). \quad (15)$$

4.3 Recursive formulation

Let $f_t = S$ indicate a single-person household. The value function for a single person of age t and gender g is

$$\begin{aligned} V_t^g(x_t, hs_t, I, \zeta_t) = \max_{c_t} & \left\{ u^S(c_t) + \beta s_t(I, g, hs_t, S) \right. \\ & \times E_t \left(V_{t+1}^g(x_{t+1}, hs_{t+1}, I, \zeta_{t+1}) \right) \\ & \left. + \beta[1 - s_t(I, g, hs_t, S)]E_t\theta_0(x_{t+1}) \right\}, \end{aligned} \quad (16)$$

subject to Equations (4)-(8), and (13)-(15).

A newly-single person – one who was part of a couple in the previous period and single now – distributes bequests towards other heirs before making savings and consumption decisions as a single person:

$$V_t^{ng}(x_t, hs_t, I, \zeta_t) = \max_{b_t} \left\{ \theta_1(b_t) + V_t^g(x_t - b_t, hs_t, I, \zeta_t) \right\}, \quad (17)$$

subject to equation

$$b_t \in [0, x_t - c_{min}(f_t)], \quad (18)$$

which prohibits the surviving spouse from using bequests to become eligible for government transfers.

The value function for couples ($f_t = C$) can be written as

$$\begin{aligned}
V_t^C(x_t, hs_t^h, hs_t^w, I, \zeta_t) = \max_{c_t} \Bigg\{ & u^C(c_t) \\
& + \beta s_t(I, w, hs_t^w, C) s_t(I, h, hs_t^h, C) E_t \left(V_{t+1}^C(x_{t+1}, hs_{t+1}^h, hs_{t+1}^w, I, \zeta_{t+1}) \right) \\
& + \beta s_t(I, w, hs_t^w, C) [1 - s_t(I, h, hs_t^h, C)] E_t \left(V_t^{nw}(x_{t+1}^w, hs_{t+1}^w, I, \zeta_{t+1}) \right) \\
& + \beta [1 - s_t(I, w, hs_t^w, C)] s_t(I, h, hs_t^h, C) E_t \left(V_{t+1}^{nh}(x_{t+1}^h, hs_{t+1}^h, I, \zeta_{t+1}) \right) \\
& + \beta [1 - s_t(I, w, hs_t^w, C)] [1 - s_t(I, h, hs_t^h, C)] E_t \theta_0(x_{t+1}) \Bigg\}, \tag{19}
\end{aligned}$$

subject to Equations (4)-(8) and (13)-(14). The dating of the continuation value for new widows, $V_t^{nw}(\cdot)$, reflects that wives are one model period (two years) younger than their husbands. We solve our model numerically: see Appendix E for more details.

5 Estimation

We adopt a two-step strategy to estimate the model. In the first step, we estimate or calibrate those parameters that, given our assumptions, can be cleanly identified outside our model. These include health transitions, out-of-pocket medical expenses, and mortality rates from raw demographic data. In addition, we fix the discount factor β at an annual value of 0.97, and we set the consumption floor for couples to be a multiple of the consumption floor for singles according to statutory rules for Medicaid and Supplemental Social Insurance.

In the second step, we estimate the rest of the model's parameters, which include risk aversion, the consumption equivalence scale, bequest parameters, and the consumption floor for singles,

$$\Delta = (\nu, \eta, \phi_0, \phi_1, \kappa_0, \kappa_1, c_{min}(f_t = S)),$$

with the method of simulated moments (MSM), taking as given the parameters that were estimated in the first step. In particular, we find the parameter values that allow the simulated life-cycle decision profiles to "best match" (as measured by a GMM criterion function) those from the data. Appendix E provides more details on the mechanics of our MSM procedure.

Because our goal is to explain why retirees save so much and do so at rates that differ by income, we match moments of the distribution of net worth by cohort, age, and PI tercile. Because we wish to study differences in savings patterns of couples and singles, we match asset profiles for the singles and couples separately. Finally, because Medicaid is an important program insuring the medical expenses and consumption of the poor we also match Medicaid recipiency. More specifically, the moment conditions that comprise our estimator are given by

1. The 25th percentile, median, and 75th percentile of asset holdings by PI tercile-cohort-year for current singles (including bequests).
2. The 25th percentile, median, and 75th percentile of asset holdings by PI tercile-cohort-year for those who are currently couples with both members currently alive.⁶
3. Medicaid recipiency by PI tercile-cohort-year for all households currently alive.

Appendix F contains a detailed description of these moment conditions, the weighting matrix in our GMM criterion function, and the asymptotic distribution of our parameter estimates.

When there is a death in a couple, we drop the household from the current couples' moments and add the surviving spouse to the appropriate moment condition for current singles. This is consistent with our assumption that singles differ only in their state variables, regardless of marital history. Such an assumption also holds well in the data, meaning that, conditional on our state variables the savings of recent singles are very similar to those of the people we were single in the past. (See Figure 1.) This is perhaps not surprising, given that the vast majority of single people in our sample were married at some point in the past.

Our estimation targets do not include the falls in wealth that follow the first spouse's death because we use them to validate our estimated model by applying the same estimation procedure to the AHEAD data and our model-simulated data.

When estimating our model, we face two well-known problems. First, in a cross-section, older households were born in earlier years than younger households and, due to secular income growth, have lower lifetime incomes. Because of this, the asset levels of households in older cohorts will likely be lower as well. As a result, comparing older households born in earlier years to younger households born in later years leads to understated asset growth. We address this problem by starting

⁶We also include bequests for the small number of couples where both spouses die between adjacent waves.

our simulations with initial conditions that come from the data (including wealth, income, age, and so on). This allow us to capture much of the heterogeneity across cohorts. Second, lower-income households and singles tend to die at younger ages than higher-income households and couples. The average survivor in a cohort thus has higher lifetime income, and thus more wealth, than the average deceased member of the same cohort. This “mortality bias” is more severe at older ages, when a greater share of the cohort members are dead. As a result, not accounting for mortality bias leads to overstated asset growth. We address this problem by modeling mortality as a function of key observables that include gender, permanent income and marital status. Our simulated profiles thus incorporate the same mortality bias as the data.

6 Estimation results

In section 3 we report some of the key facts about savings and Medicaid recipiency that we require our estimated model to match. In this section we report some of the most relevant features of our first-step model estimates, and then discuss our second-step estimates.

6.1 First-step estimation results

As they are key elements affecting saving behavior, the most important features of our first-step estimates pertain to life expectancy and nursing home risk, income and its drop when one of the spouses dies, and medical spending. Appendix G describes our first-step estimation procedures in detail.

6.1.1 Life expectancy and nursing home risk

We estimate health transitions and mortality rates simultaneously by fitting the transitions observed in the AHEAD to a multinomial logit model. We allow the transition probabilities to depend on age, sex, marital status, current health status, PI, and interactions of these variables.

Table 1 shows life expectancies at age 70 for singles and couples, respectively, that we obtain when we use our estimated transition probabilities to simulate demographic histories, beginning at age 70, for different gender-PI-health-family structure combinations. It shows that rich people, women, married people, and healthy people live much longer than their poor, male, single, and sick counterparts. For instance, a single man at the 10th permanent income percentile and in a nursing home expects to live only 3.0 more years, while a single woman at the 90th percentile and in good

Income Percentile	Nursing Home	Men		Women		Good Health	All
		Bad Health	Good Health	Nursing Home	Bad Health		
Singles							
10	3.0	6.9	8.7	4.1	11.3	13.2	10.2
50	3.0	7.8	10.3	4.1	12.3	14.9	11.5
90	2.9	8.1	10.9	3.8	12.5	15.4	12.0
Couples							
10	2.7	7.8	9.8	4.0	12.1	14.1	11.3
50	2.8	9.4	12.2	4.0	13.7	16.3	13.4
90	2.7	10.4	13.5	3.9	14.6	17.3	14.5
Single Men							9.0
Married Men							11.5
Single Women							13.9
Married Women							15.8
Oldest Survivor*							17.9
Probability that Oldest Survivor is Woman							63.7%

Notes: * Oldest survivor among households that were a couple at age 70.

Table 1: Life expectancy in years, conditional on reaching age 70

health expects to live 15.4 more years. The far right column of the top two panels shows average life expectancy conditional on permanent income, averaging over both genders and health states. Singles at the 10th percentile of the permanent income distribution live on average 10.2 years, while singles at the 90th percentile live on average 12.0 years.

People in couples live about 2 years longer than singles: single women live on average 13.9 years versus 15.8 for married women but, conditional on PI and health, the differences in longevity are much smaller. Thus, married people live longer than singles, but a significant part of the difference is explained by the fact that married people tend to have higher PI and to be in better health. The bottom part of Table 1 shows the expected years of remaining life for the oldest survivor in a household when both the man and the woman are 70. On average the last survivor lives 17.9 more

years. The woman is the oldest survivor 63.7% of the time.

Income Percentile	Men		Women		
	Bad Health	Good Health	Bad Health	Good Health	All
Singles					
10	23.6	25.3	35.8	37.9	32.8
50	22.8	24.8	35.5	38.2	32.5
90	20.3	22.8	32.2	35.8	30.1
Couples					
10	17.3	19.2	34.4	37.0	28.7
50	16.6	18.8	34.1	37.3	28.7
90	14.6	16.8	31.4	34.5	26.3
Single Men					26.4
Married Men					19.5
Single Women					37.2
Married Women					36.3

Table 2: Probability of ever entering a nursing home, conditional on being alive at age 70

Table 2 shows that single men and women face on average a 26% and 37% chance of being in a nursing home for an extended stay (at least 60 days in a year), respectively, while married men and women face on average a 20% and 36% chance of being in a nursing home for an extended stay. Married people are much less likely to transition into a nursing home at any age, but married people, especially women, often become single as their partner dies. Furthermore, married people tend to live longer than singles, so they have more years of life to potentially enter a nursing home. Compared to gender or marital status, permanent income and age 70 health have a smaller effect on ever being in a nursing home. This is because those with high permanent income, or in good health, are less likely to be in a nursing home at each age but they tend to live longer.

6.1.2 Income

As noted above, we model non-asset income as a function of PI, age, and family structure. Figure 4 presents predicted income profiles for those at the 20th and 80th percentiles of the PI distribution. The average income of couples ranges from about \$14,000 at the 20th percentile to over \$30,000 at the 80th. As a point of comparison, median wealth holdings for these two groups at age 74 are \$70,000 and \$330,000, respectively.

To evaluate how household income drops when one of the spouses dies, Figure 4 displays three scenarios for each PI level, all commencing from the income of a couple. Under the first scenario, the household remains a couple until age 100. Under the second one, the man dies at age 80. Under the third one, the woman dies at age 80. Our estimates imply that couples in which the husband dies at age 80 suffer a 40% decline in income, while couples in which the wife dies at 80 suffer a 30% decline in income.




Figure 4: Income, conditional on permanent income and family structure. All households begin as couples, then either stay in a couple or switch to being a single man or single woman at age 80

The income losses that occur at the death of a spouse reflect the fact that even though Social Security and private defined benefit pensions have survivors' benefits, these benefits replace only a fraction of the deceased spouse's income. By way of example, a wife can base her Social Security benefits either on her own Social Security contributions, or on those of her husband (assuming the husband is retired). Should the husband die, his benefits would cease. If the wife's benefits were based on her husband's contributions, after his death her individual benefit would double, from

50% to 100% of the husband’s benefit. In spite of this, the spouse would receive only $(100\%/[100\% + 50\%]) = 67\%$ of the household’s original Social Security benefits.⁷ If the wife’s benefits were based on her own contributions, her individual benefit would increase only if her husband’s contributions were higher than hers. The relative decrease in household benefits would then depend on each spouse’s contributions. Our fixed effects regression estimates, which capture the average drop in income at the time of a spouse’s death, average over all these cases.

6.1.3 Medical spending

Because our model explicitly accounts for Medicaid payments, which depend on a household’s savings, the measure of medical expenses we need is the sum of out-of-pocket spending by the household and Medicaid payments.⁸ Although the AHEAD contains detailed data on out-of-pocket medical spending, it does not contain Medicaid payments. A key empirical contribution of this paper is to construct the sum of these two components by combining out-of-pocket medical spending information in the AHEAD with Medicaid payment data from the Medicare Current Beneficiary Survey (MCBS).

The MCBS contains extremely high quality administrative and survey information on both Medicaid payments and out-of-pocket medical spending (De Nardi, French, Jones, and McCauley [21]). Its main limitation for our purposes is that the data are collected at the individual, rather than the household level: although the MCBS contains marital status, it lacks information on the medical spending or health of the spouse. To exploit the strengths of each dataset, we use the conditional mean matching procedure described in Appendix D to impute a Medicaid payment for each individual in each period in the AHEAD. The procedure preserves both the mean and the distribution of combined medical spending, conditional on Medicaid recipiency, age, income, out-of-pocket spending and other health and medical utilization variables. The regression of Medicaid payments on these variables has an R^2 statistic of 0.67, suggesting that our predictions are accurate.

We model medical spending as a function of a polynomial in age, marital status, health of each spouse (at the beginning and the end of the period), interactions of these variables, and persistent and transitory spending shocks (see Equation (6)).

⁷These calculations also operate in reverse, with “husband” and “wife” exchanged. To perform the calculations, we make several assumptions, including that both spouses begin receiving benefits at the normal retirement age. See <https://socialsecurity.gov/planners/retire/yourspouse.html> for more details of calculation of spousal benefits.

⁸A large share of the elderly’s medical spending is covered by Medicare co-pays, which are not means-tested. Our measure is net of those co-pays.

To estimate these profiles, we use fixed effects rather than ordinary least squares for two reasons. First, differential mortality causes the composition of our sample to vary with age, whereas we are interested in how medical expenses vary for the same individuals as they grow older. Second, cohort effects are ubiquitous and using fixed effects we construct the profiles for our model using the mean fixed effect for the cohort with an average age of 74 in 1996.

We estimate the medical spending persistence parameter ρ_m and the variance of the transitory and persistent medical spending shocks using a standard error components model. Our estimates imply that approximately 40% of the cross-sectional variation in log medical spending is explained by observables. Of the remaining cross-sectional variation, 40% comes from the persistent shock and 60% from the transitory shock. Our estimated value of ρ_m is 0.85. See Appendix G.4 for details.




Figure 5: Mean medical spending. Panel (a): current singles. Panel (b): current couples.

Figure 5 shows model-predicted mean medical spending for the AHEAD cohort aged 72-77 in 1996, where we simulate the households through age 100.⁹ Three general trends are apparent. First, medical expenses rise rapidly with age, in part because older individuals are more likely to reside nursing homes or die.¹⁰ Second, spending rises only modestly with PI, if at all. Although our underlying coefficients show that medical expenses rise with PI when health is held constant, lower-income households are often in worse health. Finally, the spending of singles (left panel) is roughly half that of couples (right panel).

⁹Because of the structure of our medical spending model, which requires two periods of health realizations, the simulation results start at age 76.

¹⁰Our simulations include end-of-life spending.

Figure 6 shows out-of-pocket medical spending. We find out-of-pocket expenses by simulating our estimated structural model and calculating Medicaid payments. Subtracting these payments from total medical spending (shown in Figure 5) yields out-of-pocket spending. Because Medicaid covers a larger share of medical expenses in poorer households, out-of-pocket spending has a strong income gradient. Likewise, as medical expenses rise with age, the share covered by Medicaid rises as well. This leads out-of-pocket medical spending to rise more slowly with age than total spending, especially among the low-income. For instance, among singles in the bottom PI tercile, out-of-pocket medical spending stays between \$3,700 and \$5,400; among those at the top, spending rises by a factor of 3, from \$5,600 to \$18,000.




Figure 6: Mean out-of-pocket medical spending. Panel (a): current singles. Panel (b): current couples.

To find the implications of our estimated spending process for lifetime medical spending, we update the calculations in Jones, De Nardi, French, McGee and Kirschner [34] to use our latest estimates. In particular, we simulate a large number of demographic and medical spending histories and calculate discounted sums at each age. Figure 7 shows that, as of age 70, single households will on average incur \$63,000 of medical expenses over the remainders of their lives, and that nearly 10% of them will incur medical expenses in excess of \$126,000. The corresponding values for initial couples are \$129,000 for the mean and \$234,000 for the 90th percentile. Some of this total may be picked up by Medicaid, but it is nonetheless significant. The spending statistics for singles rise with age until age 90. This is not merely an artifact of mortality bias. Although older people have less time to live, Figure 5 shows that medical expenses rise with age. In other words, at any point in time, most survivors have yet to incur the bulk of their lifetime expenses. As De Nardi,

French, and Jones [18] and others have pointed out, this “backloading” of medical expenses gives it the potential to be an important driver of saving. In contrast, as couples become singles their medical expenses fall.




Figure 7: Mean and 90th percentile of remaining lifetime medical spending for surviving households, initial singles and initial couples

6.2 Second-step estimation results

Table 3 reports our estimated preference parameters. Our estimate of ν , the coefficient of relative risk aversion, is 3.7, a value similar to that estimated for retired singles in De Nardi, French, and Jones [18], and to those typically used in the life cycle literature.

Our estimate of consumption equivalence scale η is almost identical to the “modified” OECD scale. It also lies within the confidence interval estimated by Hong and Rios-Rull [31], who estimate it using data on life insurance holdings and a structural life-cycle model of consumption and saving decisions. After reviewing a variety of estimates, Fernandez-Villaverde and Krueger [26] argue in favor of similar economies of scale. In combination with our estimated value of ν , our estimated η implies that the consumption of a couple must be 1.63 times that of a single to equate the marginal utility of consumption across the two groups.¹¹ Intuitively, η is identified

¹¹To see this note that the marginal utilities of singles and couples are $\frac{\partial u^S(c^S)}{\partial c} = (c^S)^{-\nu}$ and

ν : coefficient of RRA	3.698 (0.09073)
η : consumption equivalence scale	1.514 (0.2036)
ϕ_0 : bequest intensity, single (in 000,000s)	133.3 (25.5)
κ_0 : bequest curvature, single (in 000s)	9,175 (1,020)
ϕ_1 : bequest intensity, surviving spouse	7,581 (1,411)
κ_1 : bequest curvature, surviving spouse (in 000s)	244.7 (25.77)
$c_{min}(f = 1)$: annual consumption floor, singles	4,108 (113.8)

Notes: Standard errors in parentheses. We set $c_{min}(f = 2) = 1.5 \cdot c_{min}(f = 1)$.

Table 3: Estimated second-step parameters

from the extent to which couples save, relative to singles, to self-insure future risks.

Our estimates for θ_0 and κ_0 govern the strength of terminal bequest motives. Panel (a) of Figure 8 plots their implications in terms of the share of resources that would be left as bequests by single households who know that they will die next period and by married couples who know that both members of their couple will die next period for sure. In that case, single households with less than \$29,600 in resources will leave no bequests,¹² while those having \$100,000 will leave over 75% of their resources in bequests. In contrast, couples require almost \$50,000 in resources before they choose to leave bequests to other heirs. These differences in behavior for singles and couples do not come from bequest parameters, which we constrain to be the same when there are no survivors. Rather, they are due to the difference in household size between couples and singles and to the presence of equivalence scales. Together, these parameters imply that bequests are luxury goods for singles (i.e., the threshold level of wealth above which households leave bequests is higher) and that they are even more of a luxury good, but are weaker for couples. Panel

¹² $\frac{\partial u^C(c^C)}{\partial c} = \left(\frac{2}{\eta^{1-\nu}}\right)(c^C)^{-\nu}$. Equating these marginal utilities yields $\frac{c^C}{c^S} = \left(\frac{2}{\eta^{1-\nu}}\right)^{1/\nu}$.

¹²This is an annual value. In our two-year framework, the threshold will (approximately) double. See Appendix H.




Figure 8: Estimated Bequest Motives. Panel (a): expenditure share allocated to bequests for couples (solid) and singles (dashed) facing certain death in the next period. Panel (b): comparing our results for singles (dashed red), with those in De Nardi, French and Jones[18] (DFJ), Lockwood [41], Ameriks, Briggs, Caplin, Shapiro and Tonetti [3] (ABCST), and Lee and Tan [40] (L&T).

(b) of Figure 8 compares our estimated bequest motives for singles with those from several previous papers and displays that our estimated bequest motives for singles are of similar strength and slightly more of a luxury good than those estimated by Lockwood [41] and Lee and Tan [40].

Our estimated consumption floor implies that the consumption of a single household is bounded below at \$4,108 per year. The floor for couples is set to 150% of this value. Our estimated floor is best interpreted as an “effective” consumption floor that accounts for transactions costs, stigma, and other determinants of Medicaid usage.

6.3 Model fit

We require our model to match the observed heterogeneity in savings and Medicaid recipiency by age, cohort, income, and wealth levels. More specifically, we target the 25th percentile, the median, and the 75th percentile of wealth, conditional on PI tercile, by cohort and age, for both couples and singles. We also target average Medicaid recipiency by cohort, age, and PI tercile.

These targets help separately identify household precautionary savings, bequest motives, and the degree to which bequest motives are luxury goods.

Figure 9 plots median net worth by age, PI, and birth cohort for current couples

(Panel (a)) and current singles (Panel (b)), in the data (solid line) and from our model (dashed line). For clarity, we only show targeted moments for two of our birth cohorts in each panel of these set of graphs, but the model fit is similarly good across cohorts (See Appendix I).

Figure 9 highlights several important features of saving behavior during retirement. First, higher-PI households dissave more slowly than lower-PI households throughout retirement and, second, that this force is stronger for couples, as long both members of the couple stay alive. More specifically, singles in the lowest PI tercile have almost no wealth, while couples in the same PI tercile hold onto their retirement savings until age 90. Singles in the middle PI tercile start decumulating their wealth from the time that we start tracking them (age 75), while couples of the same age and PI tercile display flat asset trajectories over much of the period that we observe them. Finally, households in the highest PI tercile, and especially current couples, have highest savings during their retirement years. Our estimated model matches all of these aspects of the data.




Figure 9: Median wealth. Panel (a): current singles. Panel (b): current couples. Solid lines: cohorts aged 72-77 and 84-89 in 1996. Dashed lines: model simulations.

Figure 10 compares the Medicaid recipiency profiles generated by our model (dashed line) to those in the data (solid line) and shows that our model also matches the patterns of Medicaid usage across age and household PI. Matching these features of Medicaid recipiency means that our model is able to capture household reliance on social insurance, in addition to matching the risk of the catastrophic medical spending that leads higher-income households to utilize Medicaid. Matching Medicaid recipiency rates is important for identifying the consumption floor, which in




Figure 10: Medicaid Recipiency. Solid lines: cohorts aged 72-77 and 84-89 in 1996. Dashed lines: model simulations.

turn is important to evaluate the role of precautionary savings because a low consumption floor implies high consumption risk, and thus a high role of precautionary savings. The role of precautionary savings is in turn important for understanding bequest motives, since if precautionary savings are high, the role for bequest motives is weaker.

Appendix I shows all the moments we match and our model’s fit. Matching multiple asset percentiles across multiple PI terciles requires the model to reproduce saving behavior across a wide variety of income flows, household structures and initial asset holdings. High-PI households and couples save at higher rates than their low-PI and single counterparts, and this is especially visible at the 25th and 75th asset percentiles as shown in Figures 18 and 19. Among couples at the top of the PI distribution, wealth rises with age at the 75th percentile of wealth, but are constant at the 25th percentile. Among singles at the bottom of the PI distribution, wealth is constant at the 75th percentile, while at the 25th percentile it falls rapidly, or is 0 throughout. This variation disentangles the saving motives operating within our model.

7 Model validation

To build additional trust in the predictions of our estimated model, we turn to comparing its implications with some aspects of the data that are important given our question but are not targeted by our estimation procedure. These aspects of the data are medical spending and net worth dynamics around the time of death of one

of the spouses in a couple.




Figure 11: Average change in out-of-pocket medical expenses around death of a spouse, initial couples, relative to similar households with no death. Estimates (solid blue line) and their 95% confidence interval (shaded region) from the AHEAD data and differences-in-differences estimates (dashed red line) from model-simulated data. Death dates are centered at year 0.

Our first validation exercise focuses on the rise in out-of-pocket medical spending around the death of a spouse and is shown in Figure 11. In this graph, we use the same matched-differences-in-differences estimation as described in Figure 2 for both actual and model-generated data. Unlike Figure 2, this figure excludes Medicaid payments and focuses on the out-of-pocket component because it is this part that captures the extent to which asset declines at death are driven by medical expenses. Figure 11 shows that our model generates a trajectory of out-of-pocket medical expenses around the time of death that matches almost exactly the one observed in the data. This is notable for two reasons. First, because we did not require the model to match it by construction. Second, because although the sum of out-of-pocket and Medicaid payments is exogenous in our model, a household's out-of-pocket spending component is partly determined by their wealth, which determines Medicaid eligibility and thus the share of expenses paid by Medicaid. Consequently, the fact that our model closely replicates the dynamics of out-of-pocket medical spending through the dynamics of saving and means-tested transfers, demonstrates that we endogenously generate the correct exposure of wealth to medical expense risk.

Our second validation exercise returns to the asset dynamics of couples when one of their spouses dies. Here, we display results not only for our full sample of matched initial couples (Panel (a)) but also for those in the top half of the 1996




Figure 12: Average change in wealth around death of a spouse for initial couples (Panel (a)) and initial couples in the top half of the 1996 wealth distribution (Panel (b)). Differences-in-differences estimates (solid blue line) and their 95% confidence interval (shaded region) from the AHEAD data and differences-in-differences estimates (dashed red line) from model-simulated data. Death dates are centered at year 0.

initial wealth distribution (Panel (b)). We add this panel because matching well the asset dynamics of wealthier households around death implies that our luxury good specification of bequest motives matches households' saving behavior of wealthier households. Figure 12 shows that the net worth dynamics around a spousal death generated by our model lie well within the 95% confidence interval of their AHEAD data counterpart, even when we condition by household wealth. For the full sample of initial couples both the data and the simulations predict a decline of \$160,000 while for the sample of wealthier households we replicate the decline of \$255,000 around death.

Our model thus captures well both medical spending and net worth dynamics around the time of death of a spouse across the wealth distribution, which implies that it also matches the extent to the decline in wealth at death which is due to changes in out-of-pocket medical expenses as opposed to bequests to other heirs than the spouse. These features of the data, in turn, have important implications about the relative importance of precautionary savings and bequest motives, forces that help identify risk aversion and bequest motive parameters (given the consumption floor, which is pinned down by Medicaid recipiency).

8 What drives savings?

To quantify the savings determinants of retired households, we perform several decomposition exercises, calculating the changes in total retiree wealth that occur as we switch off various saving motives. In doing so, we assume that household wealth accumulation prior to retirement is unchanged. Because of this, our decompositions provide a lower bound on the savings changes, compared to an environment in which people can re-optimize their savings at younger ages. Our approach allows us to focus on the drivers of post-retirement savings and to better understand policy reforms that are not announced far in advance.

More specifically, to perform our decomposition, we fix our model parameters at their estimated values and change one feature of the model at a time, recompute the optimal saving decisions, and compare the resulting asset accumulation profiles to those generated by our baseline model. To make the differences clearer, we display the simulated median asset profiles for our youngest cohort, which has an average age of 74 in 1996. When couples lose a spouse and become singles, we drop them from the analysis. We do the same when singles die. Therefore, the profiles embed the same mortality bias that is in the data.




Figure 13: Median wealth by permanent income tercile: baseline model (dashed lines) and model with no medical expenses (solid lines), initial singles (Panel (a)) and current couples (Panel (b)).

We start by asking whether medical expenses are important drivers of saving behavior. To answer this question, we zero out both expected and realized medical expenses and report the corresponding asset trajectories.

Panel (a) of Figure 13 shows that the poorest singles, in the bottom PI tercile, save almost nothing to cover medical expenses: with little wealth, they choose to rely on Medicaid to pay for their medical needs. It is singles in the middle PI tercile who are most sensitive to medical spending: at age 84, removing medical expenses would cause their wealth to fall by nearly 40%, from \$87,000 to \$54,000. These individuals have enough wealth for medical expenses to pose a concern, but not enough for bequest motives to dominate. At the top PI tercile, eliminating medical spending has a smaller effect on savings. At age 84, their wealth would fall by 7%, from \$238,000 to \$221,000. The wealth they hold to leave bequests allows them to self-insure against medical expenses (as noted by Dynan, Skinner and Zeldes [23]). Interestingly, after age 90 the savings of richer singles are higher without medical expenses. This is because eliminating medical spending leaves households with more resources to consume or save.

Panel (b) shows that the effect of medical spending on couples' wealth is relatively small at every PI tercile, suggesting that medical expenses are only one of the reasons why they continue to hold wealth at very advanced ages. The largest effect occurs for couples in the top PI tercile, for whom net worth at age 84 declines from \$437,000 to \$417,000, by 5%. Comparing Panel (a) and Panel (b) shows that medical expenses have a larger effect on the savings of singles than of couples.




Figure 14: Median wealth by permanent income tercile: baseline model (dashed lines) and model with no bequest motives (solid lines). Singles (Panel (a)) and couples (Panel (b)).

We next consider the role of our estimated bequest motives by shutting down the utility from leaving any bequests (that is, we set $\phi_0 = \phi_1 = 0$) while holding all other

preference parameters constant. Figure 14 reports the resulting effects on savings. For singles, eliminating bequest motives has little effect on the median savings of the bottom two PI terciles, although they do reduce median savings at the top tercile from \$238,000 to \$197,000 (17%) at age 84.

For couples, bequest motives are more important and have a large effect on savings for the top two PI terciles. Panel (b) of Figure 14 shows while savings barely change for couples in the bottom third of the PI distribution, eliminating bequest motives reduces savings at age 84 from \$199,000 to \$172,000 (14%) in the middle tercile and from \$437,000 to \$325,000 (26%) in the top tercile.

Thus, more couples than singles save to leave bequests, and bequest motives affect their savings by a much larger extent. This is for at least two reasons. First, couples are wealthier, and bequests are a luxury good. Second, couples not only receive direct utility from leaving bequests at the time of one spouse’s death, but also indirect utility from the expectation that the last surviving spouse will leave bequests at her or his death.




Figure 15: Median wealth by permanent income tercile: baseline model (dashed lines) and model with no bequest motives and no medical expenses (solid lines). Singles (Panel (a)) and couples (Panel (b)).

While medical expenses and bequest motives taken each in isolation have relatively modest effects on the saving behavior of most households, their interaction could have a much larger impact. To quantify the importance of these interactions, we eliminate both medical expenses and bequest motives.

The resulting savings profiles are displayed in Figure 15. Because singles in the bottom and middle PI tercile are too poor for bequest motives to actively influ-

ence their saving behavior (as seen in Figure 14), the effects on their savings from eliminating both bequest motives and medical expenses are very similar to those in the case in which we only eliminate medical expenses (Figure 13). In contrast, for those in the top PI tercile, for whom the bequest motive is active, savings falls from \$238,000 to \$82,000 (by 66%). This drop is almost three times larger than the sum of the drops in the previous two experiments (24%).

This is due to the way bequest and precautionary saving motives interact. When medical expenses are uncertain, households that reduce their consumption to finance future medical expenses may end up dying with unspent wealth. In the absence of bequest motives, these accidental bequests have no value. Bequest motives reduce the opportunity cost of self-insuring through saving.¹³ At the same time, medical expenses introduce uncertainty over the size of the household's bequest. This might generate precautionary saving aimed at increasing the size of the bequest.

Panel (b) of Figure 15 displays the corresponding results for couples. The savings of couples in the bottom PI tercile fall from \$62,000 to \$41,000 (by 34%) at age 84. For couples in the middle PI tercile, savings falls from \$199,000 to \$99,000 (by 50%) at age 84. Because bequest motives are active for this group, unlike for the corresponding singles, the interaction of bequest and precautionary saving motives is much larger for middle PI couples than singles. This interaction is also very strong for couples in the top PI tercile. Rather than accumulating wealth in retirement, couples at the top tercile would run down their wealth holdings from \$437,000 to \$150,000 (66%) by age 84.

Finally, couples save to provide resources to the surviving spouse after the death of the first spouse. This is important because while at age 70 married men live on average 11.5 years and women live 15.8 years (or on average 13.7 years), the oldest survivor lives on average 17.9 years. During these remaining years the survivor faces lower income and high medical spending, including medical spending from an elevated risk of being in a nursing home when single. Studies omitting this risk will underestimate this precautionary savings motive.

To quantify the effect of this force, Figure 16 compares our baseline savings with those implied by our model when we shut down the weight on the utility of the surviving spouse. It shows that the desire to leave resources to the surviving spouse upon a partner's death has a large effect on savings at all PI levels. For instance, among couples in the lowest PI tercile, median wealth at age 84 would fall from \$62,000 in the benchmark to \$46,000 when there is no concern for the welfare of the surviving spouse, a 26% drop. Median-asset couples in the second PI tercile at age

¹³A similar argument applies to saving against longevity risk and for the decision to self-insure instead of purchasing formal Long Term Care insurance (See Lockwood [41] for an example).




Figure 16: Median wealth for couples by permanent income tercile: baseline model (dashed lines) and model with no weight on the surviving spouse (solid lines).

84 hold \$199,000 in the benchmark economy and \$165,000 in this counterfactual, a 17% drop. At the highest PI tercile, median wealth falls from \$437,000 to \$368,000 at the same age, a 16% drop.

Taken together, Figures 13-16 show that medical spending is a more important determinant of savings for singles than for couples. In contrast, bequest motives affect the savings of couples and high-PI households more than those of singles and low-PI households. The figures also show that the interaction of precautionary savings and bequests are most important for couples and high-PI households. Moreover, providing for the surviving spouse is important for all couples.

For clarity, Figures 13-16 only display profiles for households whose composition does not change. Thus, couples do not experience a decline in net worth upon death of the first spouse and singles do not include the newly widowed, who tend to be richer. To assess the saving motives across all retired households simultaneously, Table 8 aggregates the preceding simulations.¹⁴

The first column of Table 8 shows summary statistics for our baseline specification. The distribution of wealth is right skewed, with the mean (\$372,300) more than twice as large as the median (\$150,100). The second column shows the effects of eliminating medical expenses. Consistent with Figure 13, medical expenses have little effect on the saving of the wealthiest: the 75th percentile falls by 2.7% and the

¹⁴Formally, we take the same joint distribution of household initial conditions as in the data and simulate a panel accounting for the realisations of individual mortality, health and medical expenditure shocks. We use all retired households alive at each age in the simulated panel, except for their wealth when first observed which we take as given – and calculates summary statistics for wealth.

	Percentage Change from Baseline				
	Baseline Wealth	No Medical Expenses	No Bequest Motives	No Bequests or Medical Expenses	No Weight on Surviving Spouse
25th Percentile	50.0	-61.4%	8.3%	-61.9%	-83.3%
Median	150.1	-20.9%	6.9%	-45.6%	-36.33%
75th Percentile	391.9	-2.7%	-8.4%	-42.8%	-25.3%
Mean	372.3	-3.0%	-17.4%	-44.2%	-27.4%

Notes: Wealth in baseline simulations are expressed in thousands of 2014 dollars.

Table 4: Aggregate Implications of Retiree Saving Motives

mean, which is dominated by the richest households, falls only 3.0%. The effects are much larger in the lower half of the wealth distribution: the median falls by 20.9%. When medical expenses are eliminated, middle-wealth singles reduce their wealth and middle-wealth couples transfer a larger fraction of their estates from surviving spouses to non-spousal heirs.

The effects reverse in the third column, where we eliminate bequest motives. Bequests are luxury goods and their elimination causes mean wealth to fall by 17.4%. In contrast, eliminating bequest motives leads median wealth to *increase* by 6.9%. Couples save less when bequest motives are removed, but when one spouse dies all of the estate is left to the surviving spouse. For middle-wealth households, the latter effect dominates. Taken together, the second and third column reinforce our earlier findings that medical expenses are more important for poorer households and singles, while bequest motives are more important for the rich.

The fourth column also reinforces Figures 13-16, in that the effect of eliminating medical expenses and bequest motives together is large and exceeds the sum of the individual changes. This is a key finding because the previous literature has focused on the role of precautionary saving motives due to medical spending without a bequest motive. Given that our results present robust evidence of bequest motives, and that their interaction with medical expenses has a very large impact on savings, our results show that it is essential to model both bequest motives and medical spending.

The final column of Table 8 shows that setting to zero the weight that couples

place on the surviving spouse also has large effects. This is not surprising because this experiment leads couples to bequeath all their wealth to non-spousal heirs upon the death of the first spouse, but highlights that it is important to model couples and spousal deaths to fully capture life cycle saving motives.

9 Conclusions

This paper documents new facts about the retirement savings and risks of couples and singles. It then develops and estimates a rich structural model that contains these risks and explains their savings.

In terms of new facts, we compare the saving behavior of couples and singles. Households with high PI, and especially couples, tend to accumulate wealth, whereas those with low PI, and especially singles, tend to decumulate it. Furthermore, upon the death of a spouse, there are large drops in wealth and bequests are often distributed.

To understand the relative roles of risks versus bequest motives as drivers of these savings decisions, it is crucial to properly measure the risks faced by couples and singles. While medical expenses for couples and singles are not very different on a per capita basis until someone dies, medical spending jumps at the time of death of a spouse. At the same time, household income falls, exposing households to a substantial loss in financial resources at death of a spouse.

We embed these estimated risks within a rich structural model of savings that incorporates heterogeneity in financial resources, life expectancy, and medical expense risks, including the large jump in medical expenses that is typically observed before a death. Our model also accounts for means-tested social insurance and bequest motives. It allows households to leave bequests upon death of the last survivor. In addition, and very importantly, households can choose to distribute bequests when the first member of the couple dies (net worth is split between the surviving spouse and other heirs).

We estimate our model using MSM, targeting the savings of both couples and singles and Medicaid recipiency rates by cohort, permanent income, and wealth levels. Our model fits all of these important features of the data well. In addition, although we do not require our model to match the dynamics of out-of-pocket medical spending and wealth around the time of death of a spouse, our model also matches those well. This validates key predictions of our model.

Finally, we use our estimated model to measure the extent to which the savings of retirees are driven by medical expenses, bequest motives, and their interaction. We uncover several new findings. Medical expenses are more important for singles,

whereas bequest motives are more important for couples. This is due to two main reasons. First, couples not only receive direct utility from leaving bequests at the time of one spouse's death, but also indirect utility from the expectation that the last surviving spouse will leave bequests upon their own death. Second, bequests are a luxury good, and couples are wealthier. We also find that introducing a bequest motive reduces the effects of medical expenses on savings. This is because savings intended for medical spending that go unspent due to early death are still valued when bequest motives are present. This finding shows that we need to model both medical spending and bequests. Lastly, although medical spending is important, the desire to transfer resources to one's heirs other than their spouse is key to understanding the wealth drops at the death of a first spouse.

Hence, to understand the drivers of savings, it is essential to model the bequest motives and medical spending of both couples and singles. Why people save has important policy implications. Our results suggest that couples and high-PI singles can easily self-insure against medical spending risk because they save to leave bequests. Low-PI singles are well insured through Medicaid and do not need to save for medical expenses. In contrast, middle-PI singles set aside significant amounts of wealth for precautionary purposes because they are not rich enough to want to leave bequests, but are too rich to qualify for Medicaid. It is thus the middle-PI singles who would respond most to changes in public health insurance.

Medicaid also includes rules determining its relative generosity for couples and singles. Evaluating whether these rules are well designed requires carefully modeling and quantifying the saving motives of all retirees.

References

- [1] Andrew B Abel and Mark Warshawsky. Specification of the joy of giving: Insights from altruism. *The Review of Economics and Statistics*, 70(1):145, 1988.
- [2] Joseph G. Altonji and Lewis M. Segal. Small sample bias in gmm estimation of covariance structures. *Journal of Business and Economic Statistics*, 14(3):353–366, 1996.
- [3] John Ameriks, Joseph Briggs, Andrew Caplin, Matthew D. Shapiro, and Chris Tonetti. Long-term care utility and late in life saving. *Journal of Political Economy*, 128(6):2375–2451, 2020.
- [4] John Ameriks, Andrew Caplin, Steven Laufer, and Stijn Van Nieuwerburgh. The joy of giving or assisted living? using strategic surveys to separate bequest and precautionary motives. *Journal of Finance*, 66(2):519–561, 2011.
- [5] James Andreoni. Giving with impure altruism: Applications to charity and ricardian equivalence. *Journal of Political Economy*, 97:1447–1458, 1989.
- [6] Orazio P. Attanasio and Carl Emmerson. Mortality, health status, and wealth. *Journal of the European Economic Association*, 1(4):821–850, 2003.
- [7] Gary S. Becker and Nigel Tomes. Human capital and the rise and fall of families. *Journal of Labor Economics*, 4(3):1–39, 1986.
- [8] B. Douglas Bernheim, Andrei Shleifer, and Lawrence H. Summers. The strategic bequest motive. *Journal of Political Economy*, 93(5):1045–1076, 1985.
- [9] R. Anton Braun, Karen A. Kopecky, and Tatyana Koreshkova. Old, sick, alone and poor: A welfare analysis of old-age social insurance programs. *The Review of Economic Studies*, 84(2):580–612, 2017.
- [10] Meta Brown. End-of-life transfers and the decision to care for a parent. mimeo, 2003.
- [11] Moshe Buchinsky. Recent advances in quantile regression models: A practical guideline for empirical research. *Journal of Human Resources*, 33:88–126, 1998.
- [12] Christopher D. Carroll. Why do the rich save so much? NBER Working Paper 6549, 1998.

- [13] Gary Chamberlain. Comment: Sequential moment restrictions in panel data. *Journal of Business & Economic Statistics*, 10(1):20–26, 1992.
- [14] Martin David, Roderick J. A. Little, Michael E. Samuhel, and Robert K. Triest. Alternative methods for cps income imputation. *Journal of the American Statistical Association*, 81(393):29–41, 1986.
- [15] James B. Davies and Antony F. Shorrocks. The distribution of wealth. In Anthony B. Atkinson and François Bourguignon, editors, *Handbook of Income Distribution*, pages 605–675. Handbooks in Economics, vol. 16. Amsterdam; New York and Oxford: Elsevier Science, North-Holland, 2000.
- [16] Mariacristina De Nardi. Wealth inequality and intergenerational links. *Review of Economic Studies*, 71(3):743–768, 2004.
- [17] Angus Deaton. Saving and liquidity constraints. *Econometrica*, 59(5):1221–1248, 1991.
- [18] Mariacristina De Nardi, Eric French, and John Bailey Jones. Why do the elderly save? the role of medical expenses. *Journal of Political Economy*, 118(1):39–75, 2010.
- [19] Mariacristina De Nardi, Eric French, and John Bailey Jones. Medicaid insurance in old age. *American Economic Review*, 106(11):3480–3520, 2016.
- [20] Mariacristina De Nardi, Eric French, and John Bailey Jones. Savings after retirement: A survey. *Annual Review of Economics*, 8(1):177–204, 2016.
- [21] Mariacristina De Nardi, Eric French, John Bailey Jones, and Jeremy McCauley. Medical spending of the u.s. elderly. *Fiscal Studies*, 37(3-4):717–747, 2016.
- [22] Darrell Duffie and Kenneth J. Singleton. Simulated moments estimation of markov models of asset prices. *Econometrica*, 61(4):929–952, 1993.
- [23] Karen E. Dynan, Jonathan Skinner, and Stephen P. Zeldes. The importance of bequests and life-cycle saving in capital accumulation: A new answer. *American Economic Review*, 92(2):274–278, 2002.
- [24] Itzik Fadlon and Torben Heien Nielsen. Family health behaviors. *American Economic Review*, 109(9):3162–91, 2019.
- [25] Daniel Feenbergs and Jonathan Skinner. The risk and duration of catastrophic health care expenditures. *Review of Economics and Statistics*, 76:633–647, 1994.

- [26] Jesús Fernández-Villaverde and Dirk Krueger. Consumption over the Life Cycle: Facts from Consumer Expenditure Survey Data. *The Review of Economics and Statistics*, 89(3):552–565, 2007.
- [27] Eric French and John Bailey Jones. On the distribution and dynamics of health care costs. *Journal of Applied Econometrics*, 19(4):705–721, 2004.
- [28] Eric French and John Bailey Jones. The effects of health insurance and self-insurance on retirement behavior. *Econometrica*, 79(3):693–732, 2011.
- [29] Miguel Gouveia and Robert P Strauss. Effective federal individual income tax functions: An exploratory empirical analysis. *National Tax Journal*, pages 317–339, 1994.
- [30] Nezih Guner, Remzi Kaygusuz, and Gustavo Ventura. Income taxation of us households: Facts and parametric estimates. *Review of Economic Dynamics*, 17(4):559–581, 2014.
- [31] Jay H. Hong and José-Víctor Ríos-Rull. Life Insurance and Household Consumption. *The American Economic Review*, 102(7):3701–3730, 2012.
- [32] R. Glenn Hubbard, Jonathan Skinner, and Stephen P. Zeldes. Precautionary saving and social insurance. *Journal of Political Economy*, 103(2):360–399, 1995.
- [33] Michael D. Hurd. Mortality risk and bequests. *Econometrica*, 57(4):779–813, 1989.
- [34] John Bailey Jones, Mariacristina De Nardi, Eric French, Rory McGee, and Justin Kirschner. The lifetime medical spending of retirees. *Economic Quarterly-Federal Reserve Bank of Richmond*, 104(3):103, 2018.
- [35] John Bailey Jones, Mariacristina De Nardi, Eric French, Rory McGee, and Rachel Rodgers. Medical spending, bequests, and asset dynamics around the time of death. *Economic Quarterly*, Fourth Quarter:135–157, 2020.
- [36] Thomas F. Juster, James P. Smith, and Frank Stafford. The measurement and structure of household wealth. Mimeo, 1998.
- [37] Ruud H. Koning. Kernel: A gauss library for kernel estimation. <http://www.xs4all.nl/>

- [38] Karen Kopecky and Tatyana Koreshkova. The impact of medical and nursing home expenses and social insurance policies on savings and inequality. *American Economic Journal: Macroeconomics*, 6:29–72, 2014.
- [39] John Laitner, Dan Silverman, and Dmitriy Stolyarov. The role of annuitized wealth in post-retirement behavior. *American Economic Journal: Macroeconomics*, 10(3):71–117, 2018.
- [40] Siha Lee and Kegon T. K. Tan. Bequest Motives and the Social Security Notch. October 2019.
- [41] Lee M. Lockwood. Incidental bequests and the choice to self-insure late-life risks. *American Economic Review*, 108(9):2513–50, 2018.
- [42] Samuel Marshall, Kathleen M. McGarry, and Jonathan S. Skinner. The risk of out-of-pocket health care expenditure at the end of life. Working Paper 16170, National Bureau of Economic Research, 2010.
- [43] Makoto Nakajima and Irina A Telyukova. Home equity in retirement. *International Economic Review*, 61(2):573–616, 2020.
- [44] Whitney K. Newey. Generalized method of moments specification testing. *Journal of Econometrics*, 29(3):229–256, 1985.
- [45] Whitney K. Newey and Daniel L. McFadden. Large sample estimation and hypothesis testing. In Robert Engle and Daniel L. McFadden, editors, *Handbook of Econometrics, Volume 4*. Elsevier, Amsterdam, 1994.
- [46] Ariel Pakes and David Pollard. Simulation and the asymptotics of optimization estimators. *Econometrica*, 57(5):1027–1057, 1989.
- [47] Michael G. Palumbo. Uncertain medical expenses and precautionary saving near the end of the life cycle. *Review of Economic Studies*, 66:395–421, 1999.
- [48] Jorn-Steffen Pischke. Measurement error and earnings dynamics: Some estimates from the PSID validation study. *Journal of Business & Economics Statistics*, 13(3):305–314, 1995.
- [49] James M. Poterba, Steven F. Venti, and David A. Wise. The composition and drawdown of wealth in retirement. *Journal of Economic Perspectives*, 25:95–117, 2011.

- [50] James Powell. Estimation of semiparametric models. In Robert Engle and Daniel L McFadden, editors, *Handbook of Econometrics, Volume 4*. Elsevier, Amsterdam, 1994.
- [51] Susann Rohwedder, Steven J. Haider, and Michael Hurd. Increases in wealth among the elderly in the early 1990s: How much is due to survey design? *Review of Income and Wealth*, 52(4):509–524, 2006.
- [52] George Tauchen. Finite state markov-chain approximations to univariate and vector autoregressions. *Economics Letters*, 20(2):177–181, 1986.

Appendix A: Sample selection and data handling

Our main dataset is the AHEAD data, which is a cohort within the Health and Retirement Study (HRS).

To keep the dynamic programming problem manageable, we assume a fixed difference in age between spouses, and we take the average age difference from our data. In our sample, husbands are on average 3 years older than their wives; in the model we assume the difference is one period, or 2 years. To keep the data consistent with this assumption, we drop all households where the wife is more than 4 years older or 10 years younger than her husband.

We do not use 1993/94 assets, nor medical expenses, due to underreporting issues (Rohwedder, Haider, and Hurd [51]) and thus begin with 6,047 households. Because we only allow for household composition changes through death, we drop households where an individual enters a household or an individual leaves the household for reasons other than death. Fortunately, attrition for reasons other than death is a minor concern in our data. We drop 401 households who get married, divorced, were same sex couples, or who report making other transitions not consistent with the model, 753 households who report earning at least \$3,000 in any period, 171 households with a large difference in the age of husband and wife, and 87 households with no information on the spouse in a household. As a result, we are left with 4,634 households, of whom 1,388 are couples and 3,246 are singles. This represents 24,274 household-year observations where at least one household member was alive.

An advantage of the AHEAD relative to other datasets is that it provides panel data on health status and nursing home stays. We assign individuals a health status of “good” if self-reported health is excellent, very good or good, and a health status of “bad” if self-reported health is fair or poor. We assign individuals to the nursing home state if they were in a nursing home at least 120 days since the last interview (on average 60 days per year) or if they spent at least 60 days in a nursing home before the next scheduled interview and died before that scheduled interview.

The AHEAD has information on the value of housing and real estate, autos, liquid assets (which include money market accounts, savings accounts, T-bills, etc.), IRAs, Keoghs, stocks, the value of a farm or business, mutual funds, bonds, “other” assets and investment trusts less mortgages and other debts. One problem with asset data is that the wealthy tend to underreport their wealth in virtually all household surveys (Davies and Shorrocks [15]). This could lead us to underestimate asset levels at all ages. However, Juster, Smith, and Stafford [36] show that the wealth distribution of the AHEAD matches up well with aggregate values for all but the richest 1 % of households. Given that we match the quantiles (conditional on permanent income)

rather than means, underreporting at the very top of the wealth distribution should not affect our results.

Social Security and pension benefits are included in our measure of annual income, hence differences in Social Security and pension wealth appear in our model as differences in the permanent income measure we use to predict annual income (and other state variables).

Appendix B: Computing permanent income

We assume that log income evolves as following

$$\ln y_{it} = \kappa(t, f_{it}) + h(I_i) + \omega_{it}, \quad (20)$$

where $\kappa(t, f_{it})$ is a flexible function of age t and family structure f_{it} (i.e., couple, single man or single woman) and ω_{it} represents measurement error. The variable I_i is the household's percentile rank in the permanent income (PI) distribution. Since it is a summary measure of lifetime income at retirement, it should not change during retirement and is thus a fixed effect over our sample period. However, income could change as households age and potentially lose a family member.

To estimate Equation (20) we first estimate the fixed effects model

$$\ln y_{it} = \kappa(t, f_{it}) + \alpha_i + \omega_{it}, \quad (21)$$

which allows us to obtain a consistent estimate of the function $\kappa(t, f_{it})$. Next, note that as the number of time periods over which individual i is observed (denoted T_i) becomes large,

$$\text{plim}_{T_i \rightarrow \infty} \frac{1}{T_i} \sum_{t=1}^{T_i} [\ln y_{it} - \kappa(t, f_{it}) - \omega_{it}] = \frac{1}{T_i} \sum_{t=1}^{T_i} [\ln y_{it} - \kappa(t, f_{it})] = \alpha_i = h(I_i). \quad (22)$$

We thus calculate the PI ranking I_i for every household in our sample by taking the percentile ranking of $\frac{1}{T_i} \sum_{t=1}^{T_i} [\ln y_{it} - \hat{\kappa}(t, f_{it})]$, where $\hat{\kappa}(t, f_{it})$ is the estimated value of $\kappa(t, f_{it})$ from Equation (21). Put differently, we take the mean residual per person from the fixed effects regression (where the residual includes the estimated fixed effect), then take the percentile rank of the mean residual per person to construct I_i .

However, we also need to estimate the function $h(I_i)$, which converts the estimated index I_i back to a predicted level of income that can be used in the dynamic programming model. To do this we estimate the function

$$[\ln y_{it} - \kappa(t, f_{it})] = h(I_i) + \omega_{it} \quad (23)$$

where the function $h(I_i)$ is a flexible functional form. In practice we model $\kappa(t, f_{it})$ as a third order polynomial in age, dummies for family structure, and family structure interacted with an age trend. When estimating Equation (23), we replace the function $\kappa(t, f_{it})$ with its estimated value. We model $h(I_i)$ as a fifth order polynomial in our measure of permanent income percentile.

Given that we have, for every member of our sample, t, f_{it} , and estimates of I_i and the functions $\kappa(., .), h(.)$, we can calculate the predicted value $\ln \hat{y}_{it} = \hat{\kappa}(t, f_{it}) + \hat{h}(\hat{I}_i)$. It is $\ln \hat{y}_{it}$ that we use when simulating the model for each household. A regression of $\ln y_{it}$ on $\ln \hat{y}_{it}$ yields a R^2 statistic of .74, suggesting that our predictions are accurate.

Appendix C: Matching methodology

We match each household that suffers a death to a household that did *not* experience a death at the same date or within the window shown in our graphs, but did experience a death within 6-10 years. (Fadlon and Nielson [24] impose a similar restriction in their analysis of health shocks.) More specifically, for each household we take a matched control household from the set of households who have the same initial (1996) household composition, have a PI percentile within 15 percentage points and where their 1996 wealth differs by no more than 2.5% (or \$5,000 in levels). We then use a random number generator to select a single control household from this set giving each household an equal probability of being selected. Our event study specification then uses a sample of 476 households who experience the death of a spouse (2,383 household-year observations) matched to 476 households who do not experience a death in the sample window (2,624 household-year observations). Jones, De Nardi, French, McGee and Rodgers [35] present more results on this.

Appendix D: Imputing Medicaid payments and out-of-pocket medical expenses

Our goal is to estimate the data generating process for the sum of Medicaid payments and out-of-pocket expenses: this is the variable $\ln(m)$ in Equation (6) of the main text. The AHEAD data contain information on out-of-pocket medical spending, but not on Medicaid payments. Fortunately, the Medicare Current Beneficiary Survey (MCBS) has extremely high quality information on both Medicaid payments and out-of-pocket medical spending. One drawback of the MCBS, however, is that although it has information on marital status and household income, it does not include information on the medical spending or health of the spouse.

We use a two-step imputation procedure to exploit the strengths of both datasets.

First step of imputation procedure

We use the MCBS to infer Medicaid payments for recipients, conditional on observable variables that exist in both the MCBS and the AHEAD datasets. Let oop_{it} denote out-of-pocket medical expenses for AHEAD person i at time t , and let Mcd_{it} denote the dollar value of Medicaid payments.

To impute Mcd_{it} , we follow David, Little, Samuhel and Triest [14] and French and Jones [28] and use the following predictive mean matching regression approach. First, using every member of the MCBS sample with a positive Medicaid indicator (i.e., a Medicaid recipient), we regress the variable of interest Mcd on the vector of observable variables z , yielding $Mcd = z\beta + \varepsilon$. Second, for each individual j in the MCBS we calculate the predicted value $\widehat{Mcd}_{jt} = z_{jt}\hat{\beta}$, and for each member of the sample we calculate the residual $\hat{\varepsilon}_{jt} = Mcd_{jt} - \widehat{Mcd}_{jt}$. Third, we sort the predicted value \widehat{Mcd}_{jt} into deciles and keep track of all values of $\hat{\varepsilon}_{jt}$ within each decile.

The variables z_{jt} include nursing home status, number of nights spent in a nursing home, an age polynomial, total household income, marital status, self-reported health, race, visiting a medical practitioner (doctor, hospital or dentist), out-of-pocket medical spending, education and death of an individual. Because the measure of medical spending in the AHEAD is medical spending over two years, we take two-year averages of the MCBS data to be consistent with the structure of the AHEAD. The regression of Mcd_{jt} on z_{jt} yields a R^2 statistic of .67, suggesting that our predictions are accurate.

Second step of imputation procedure:

Next, for every individual i in the AHEAD sample with a positive Medicaid indicator, we impute $\widehat{Mcd}_{it} = z_{it}\hat{\beta}$ using the value of $\hat{\beta}$ estimated using the MCBS. Then we impute ε_{ic} for each member of the AHEAD sample by finding a random individual j in the MCBS with a value of \widehat{Mcd}_{jt} in the same decile as \widehat{Mcd}_{it} in the AHEAD, and set $\varepsilon_{it} = \hat{\varepsilon}_{jt}$. The imputed value of Mcd_{it} is $\widetilde{Mcd}_{it} = \widehat{Mcd}_{it} + \varepsilon_{it}$, and the imputed value of the sum m_{it} is $\widetilde{m}_{it} = oop_{it} + \widetilde{Mcd}_{it}$.

As David, Little, Samuhel and Triest [14] point out, our imputation approach is equivalent to hot-decking when the “ z ” variables are discretized and include a full set of interactions. The advantages of our approach over hot-decking are two-fold. First, many of the “ z ” variables are continuous, and it seems unwise to discretize them. Second, we use a large number of observable variables “ z ” because we find that adding extra variables greatly improving goodness of fit when imputing Medicaid payments. Even a small number of variables generates a large number of hot-decking cells because hot-decking uses a full set of interactions. Thus, in this context, hot

decking is too data intensive.

We predict Medicaid payments for 3,756 household-wave observations with Medicaid recipiency in the AHEAD and report the results of this imputation exercise in Table 5. Households where the last surviving member died between the previous and current waves of the sample have the largest imputed Medicaid payments. In contrast, couples have the smallest Medicaid payments per individual, but Medicaid payments at the household level are larger than for either single men or women. The imputed Medicaid payments for each spouse in a couple are approximately equal and the results for couples are not driven by the expenditures of only the husband or wife. In Table 5, new widows and widowers (those whose spouse has died between the current and preceding waves of the AHEAD sample) are included in the rows for single men and single women – on average the Medicaid payments of dead spouses are less than 20% of the single households’ Medicaid payments.

The distribution of imputed Medicaid expenditures in the AHEAD is close to that in the MCBS. The mean imputed Medicaid expenditure, Mcd_{it} , for Medicaid recipients in the AHEAD is \$14,050. This is lower than the corresponding value of \$16,000 in the MCBS. This difference is due to the distribution of observable variables, z , in both the MCBS and AHEAD. In particular, the number of nights in a nursing home is 20% lower in the AHEAD than in the MCBS, which is in part due to the fact that AHEAD respondents were not in a nursing home when they answered the survey for the first time.

Family Structure	Number of Medicaid Eligible Households	Mean	Standard Deviation
Dead	1,040	21,800	31,500
Couples	287	15,300	25,700
Single Men	351	11,600	21,500
Single Women	2,078	12,300	21,800

Table 5: Imputed Medicaid payments for Medicaid beneficiaries in AHEAD

Appendix E: Computing optimal decision rules and MSM estimates

We compute the value functions by backward induction. We start with singles. At time- T , we find the value function and decision rules by maximizing Equation (16), subject to the relevant constraints, using $V_{T+1}^g = \theta_0(x_t - c_t)$, $g = h, w$. This yields the value function V_T^g . We then find the value function and decision rules at time $T - 1$ by solving Equation (16) with V_T^g . Continuing this backward induction yields decision rules for periods $T - 2, T - 3, \dots, t_r$.

We find the value function for new widows and widowers, V_t^{ng} , and the split of the estate between spousal and non-spousal heirs, by solving Equation (17), subject to the relevant constraints. These calculations utilize the value function for singles, V_t^g , described immediately above.

We find the time- T value function for couples, V_T^c by maximizing Equation (19), subject to the relevant constraints and the value function for the singles, and setting the time- $T + 1$ continuation value to $\theta_0(x_t - c_t)$. This yields the value function V_T^c and the decision rules for time T . We then find the decision rules at time $T - 1$ by solving Equation (19) using V_T^c and $V_T^{ng}, g = h, w$. Continuing this backward induction yields decision rules for time $T - 2, T - 3, \dots, 1$.

More specifically, to solve our value functions, we discretize the persistent component and the transitory components of the health shock, ξ and ζ , using the methods described in Tauchen [52]. We assume a finite number of permanent income categories and we break the space of cash-on-hand into a finite number of grid points. We use linear interpolation within the grid and linear extrapolation outside of the grid to evaluate the value function at cash-on-hand values that we do not directly compute. The end result is a set of decision rules for each possible combination of cash-on-hand, PI, health status, and persistent health shock (ζ).

The mechanics of our MSM approach are as follows. We compute life-cycle histories for a large number of artificial households. Each of these households is endowed with an initial value of the state vector $(t, f_t, x_t, I, hs_t^h, hs_t^w)$ that is drawn from the data for 1996. In addition, each household is assigned the entire health and mortality history that is recorded for this same household in the AHEAD data. This way we generate attrition in our simulations that exactly follows the attrition in the data (including by initial wealth and mortality). The simulated medical expenditure shocks ζ and ξ are Monte Carlo draws from discretized versions of our estimated shock processes. Our decision rules, combined with the initial conditions and simulated shocks, allows us to simulate a household's wealth, medical expenses, health, and mortality.

We construct life-cycle profiles from the artificial histories in the same way that we compute them from the real data. We use these profiles to construct moment conditions and evaluate the match using our GMM criterion. As done when constructing the figures from the AHEAD data, we drop cells with fewer than 10 observations from the moment conditions. We search over the parameter space for the values that minimize the criterion. Appendix F details our moment conditions, the weighting matrix in our GMM criterion function, and the asymptotic distribution of our parameter estimates.

Appendix F: Moment conditions and the asymptotic distribution of parameter estimates

We estimate the parameters of our model in two steps. In the first step, we estimate the vector χ , the set of parameters that can be estimated without explicitly using our model. In the second step, we use the method of simulated moments (MSM) to estimate the remaining parameters, which are contained in the $M \times 1$ vector $\Delta = (\nu, \eta, \phi_0, \phi_1, \kappa_0, \kappa_1, c_{min}(f_t = S))$.

Our estimate, $\hat{\Delta}$, of the “true” preference vector Δ_0 is the value of Δ that minimizes the (weighted) distance between the estimated life cycle profiles for wealth and Medicaid recipiency in the data and those generated by the model.

For each calendar year $t \in \{t_0, \dots, t_T\} = \{1996, 1998, \dots, 2014\}$,¹⁵ we match the 25th, 50th and 75th percentile of wealth for 3 permanent income terciles in 4 birth year cohorts, both for singles and couples. Because the 1996 (period- t_0) distribution of simulated wealth is bootstrapped from the 1996 data distribution, we match wealth for the period 1998 to 2014, hence 9 waves. In addition, we require each cohort-income-age cell have at least 15 observations to be included in the GMM criterion. As a result, we end up with 414 asset targets.

We construct the moment conditions for wealth by building on French and Jones [28] (useful references include Buchinsky [11] and Powell [50]). Suppose that household i of family type f belongs to birth cohort c and permanent income tercile p . Let $a_{fcqt}^q(\Delta, \chi)$ denote the model-predicted q th asset quantile for households in household i ’s group at time t .¹⁶ Assuming that observed wealth have a continuous conditional density, a_{fcqt} will satisfy

$$\Pr \left(a_{it} \leq a_{fcpt}^q(\Delta_0, \chi_0) \mid f, c, p, t, \text{household } i \text{ observed at } t \right) = \pi_q,$$

¹⁵Because we do not allow for macro shocks, in any given cohort, t is used only to identify the individual’s age.

¹⁶The elements of χ include the permanent income boundaries.

where π_q is the probability value associated with quantile q – when q denotes a median, π_q would equal 1/2. The preceding equation can be rewritten as a moment condition. In particular, applying the indicator function produces

$$E(1\{a_{it} \leq a_{fcpt}^q(\Delta_0, \chi_0)\} - \pi_q | f, c, p, t, \text{household } i \text{ observed at } t) = 0.$$

Letting \mathcal{I}_p denote the values contained in the p th permanent income tercile, we can convert this conditional moment equation into an unconditional one (e.g., Chamberlain [13]):

$$\begin{aligned} E\left(1\{a_{it} \leq a_{fcpt}^q(\Delta_0, \chi_0)\} - \pi_q\right) \times 1\{f_{it} = f\} \times 1\{c_i = c\} \times 1\{I_i \in \mathcal{I}_p\} \\ \times 1\{\text{household } i \text{ observed at } t\} | t = 0, \end{aligned} \quad (24)$$

for $q \in \{.25, .5, .75\}$, $f \in \{\text{single, couple}\}$, $c \in \{1, 2, 3, 4\}$, $p \in \{1, 2, 3\}$, and $t \in \{t_1, t_2, \dots, t_T\}$.

We also match Medicaid recipiency rates. We divide individuals into 4 cohorts, match data from 9 waves, and stratify the data by permanent income, but combine couples and singles. Let $\bar{u}_{cpt}(\Delta, \chi)$ denote the model-predicted recipiency rate for households in cohort c and permanent income tercile p at time t . Let u_{it} be the $\{0, 1\}$ indicator that equals 1 when household i receives Medicaid. The associated moment condition is

$$\begin{aligned} E\left([u_{it} - \bar{u}_{cpt}(\Delta_0, \chi_0)] \times 1\{c_i = c\} \times 1\{I_i \in \mathcal{I}_p\}\right. \\ \left.\times 1\{\text{individual } i \text{ observed at } t\} | t\right) = 0 \end{aligned} \quad (25)$$

for $c \in \{1, 2, 3, 4\}$, $p \in \{1, 2, 3\}$, and $t \in \{t_1, t_2, \dots, t_T\}$.

To summarize, the moment conditions used to estimate our model consist of: the moments for asset quantiles described by Equation (24) and the moments for the Medicaid recipiency rates described by Equation (25). We have a total of $J = 505$ moment conditions.

Suppose we have a dataset of I independent individuals that are each observed at up to T separate calendar years. Let $\varphi(\Delta; \chi_0)$ denote the J -element vector of moment conditions described immediately above, and let $\hat{\varphi}_I(\cdot)$ denote its sample analog. Letting $\widehat{\mathbf{W}}_I$ denote a $J \times J$ weighting matrix, the MSM estimator $\hat{\Delta}$ is given by

$$\operatorname{argmin}_{\Delta} \frac{I}{1 + \tau} \hat{\varphi}_I(\Delta; \chi_0)' \widehat{\mathbf{W}}_I \hat{\varphi}_I(\Delta; \chi_0),$$

where τ is the ratio of the number of observations to the number of simulated observations.

In practice, we estimate χ_0 using the approach described in the main text. Computational concerns, however, compel us to treat χ_0 as known in the analysis that follows.

Under the regularity conditions stated in Pakes and Pollard [46] and Duffie and Singleton [22], the MSM estimator $\hat{\Delta}$ is both consistent and asymptotically normally distributed:

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

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

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

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

$$\mathbf{D} = \frac{\partial \varphi(\Delta; \chi_0)}{\partial \Delta'} \Big|_{\Delta=\Delta_0} \quad (26)$$

is the $J \times M$ gradient matrix of the population moment vector; and $\mathbf{W} = \text{plim}_{I \rightarrow \infty} \{\widehat{\mathbf{W}}_I\}$. Moreover, Newey [44] shows that if the model is properly specified,

$$\frac{I}{1 + \tau} \hat{\varphi}_I(\hat{\Delta}; \chi_0)' \mathbf{R}^{-1} \hat{\varphi}_I(\hat{\Delta}; \chi_0) \rightsquigarrow \chi^2_{J-M},$$

where \mathbf{R}^{-1} is the generalized inverse of

$$\begin{aligned} \mathbf{R} &= \mathbf{P}\mathbf{S}\mathbf{P}, \\ \mathbf{P} &= \mathbf{I} - \mathbf{D}(\mathbf{D}'\mathbf{W}\mathbf{D})^{-1}\mathbf{D}'\mathbf{W}. \end{aligned}$$

The asymptotically efficient weighting matrix arises when $\widehat{\mathbf{W}}_I$ converges to \mathbf{S}^{-1} , the inverse of the variance-covariance matrix of the data. When $\mathbf{W} = \mathbf{S}^{-1}$, \mathbf{V} simplifies to $(1 + \tau)(\mathbf{D}'\mathbf{S}^{-1}\mathbf{D})^{-1}$, and \mathbf{R} is replaced with \mathbf{S} .

Even though the optimal weighting matrix is asymptotically efficient, it can be biased in small samples. (See, for example, Altonji and Segal [2].) Hence, we use a “diagonal” weighting matrix, as suggested by Pischke [48]. This diagonal weighting scheme uses the inverse of the matrix that is the same as \mathbf{S} along the diagonal and has zeros off the diagonal of the matrix.

We estimate \mathbf{D} , \mathbf{S} , and \mathbf{W} with their sample analogs. For example, our estimate of \mathbf{S} is the $J \times J$ estimated variance-covariance matrix of the sample data. When

estimating this matrix, we use sample statistics, so that $a_{pqt}(\Delta, \chi)$ is replaced with the sample median for group pqt .

One complication in estimating the gradient matrix \mathbf{D} is that the functions inside the moment condition $\varphi(\Delta; \chi)$ are non-differentiable at certain data points; see equation (24). This means that we cannot consistently estimate \mathbf{D} as the numerical derivative of $\hat{\varphi}_I(\cdot)$. Our asymptotic results therefore do not follow from the standard GMM approach, but rather the approach for non-smooth functions described in Pakes and Pollard [46], Newey and McFadden [45] (section 7), and Powell [50].

To find \mathbf{D} , it is helpful to rewrite equation (24), for the case of medians, as:

$$\Pr \left(f_{it} = f \& c_i = c \& I_i \in \mathcal{I}_p \& \text{individual } i \text{ observed at } t \right) \times \\ \left[\int_{-\infty}^{a_{fcpt}^{\text{median}}(\Delta_0, \chi_0)} f(a_{it} \mid f, c, I_i \in \mathcal{I}_p, t) da_{it} - \frac{1}{2} \right] = 0. \quad (27)$$

It follows that the rows of \mathbf{D} for the median asset moments are given by

$$\Pr \left(f_{it} = f \& c_i = c \& I_i \in \mathcal{I}_q \& \text{individual } i \text{ observed at } t \right) \times \\ f(a_{fcpt}^{\text{median}} \mid f, c, I_i \in \mathcal{I}_p, t) \times \frac{\partial a_{fcpt}^{\text{median}}(\Delta_0; \chi_0)}{\partial \Delta'}. \quad (28)$$

In practice, we find $f(a_{fcpt}^{\text{median}} \mid f, c, p, t)$, the conditional p.d.f. of wealth evaluated at the median a_{fcpt}^{median} , with a kernel density estimator written by Koning [37]. The derivatives for the 25th and 75th percentiles are found in analogous fashion.

Appendix G: First stage estimates

In the first-step estimation procedure, we estimate or calibrate those parameters that, given our assumptions, can be cleanly identified outside our model. In particular, we estimate health transitions and mortality rates from raw demographic data, estimate income profiles from the AHEAD, and calibrate the returns to saving and income taxes using external sources. We use our measure of total medical expenses (which combines data from the AHEAD and the MCBS) to estimate the process for medical expenses as a function of exogenous household-level state variables which are observed in the AHEAD.

We outline the estimation and calibration of these objects next.

G.1 Health and mortality

Each person's health status, hs^g , has four possible values: dead; in a nursing home; in bad health; or in good health. We allow transition probabilities for an individual's health depend on his or her current health, age, household composition f , permanent income I , and gender g .¹⁷ It follows that the elements of the health transition matrix are given by

$$\pi_{i,j,k}(t, f_{it}, I_i, g) = \Pr(hs_{i,t+2}^g = k \mid hs_{i,t}^g = j; t, f_{it}, I_i, g), \quad (29)$$

with the transitions covering a two-year interval, as the AHEAD interviews every other year.¹⁸ We estimate health/mortality transition probabilities by fitting the transitions observed in the AHEAD to a multinomial logit model.¹⁹

We enumerate each possible health status as follows

$$hs = \begin{cases} 0 & \text{Dead} \\ 1 & \text{Nursing Home} \\ 2 & \text{Bad Health} \\ 3 & \text{Good Health.} \end{cases} \quad (30)$$

Our multinomial logit assumption gives the following expression for health and mortality transitions

$$\pi_{i,j,k}(t, f_{it}, I_i, g) = \frac{e^{x_{it}\beta_k}}{\sum_{h=0}^3 e^{x_{it}\beta_h}}, \quad (31)$$

with β_k denoting the coefficient vector for future outcome k . The coefficient for death, β_0 , is normalized to zero. We jointly estimate health transitions and survival

¹⁷We do not allow health transitions to depend on medical spending. The empirical evidence on whether medical spending improves health, especially at older ages, is surprisingly mixed; see, for example, the discussion in De Nardi, French and Jones [19]. Likely culprits include reverse causality – sick people have higher expenditures – and a lack of insurance variation – almost every retiree gets Medicare.

¹⁸As discussed in De Nardi, French and Jones [19], one can fit annual models of health and medical spending to the AHEAD data. The process becomes significantly more involved, however, especially when accounting for the dynamics of two-person households.

¹⁹We do not control for cohort effects. Instead, our estimates are a combination of period (cross-sectional) and cohort probabilities. While our AHEAD sample covers 18 years, it is still too short to track a single cohort over its entire post-retirement lifespan. This may lead us to underestimate the lifespans expected by younger cohorts as they age. Nevertheless, lifespans have increased only modestly over the sample period. Accounting for cohort effects would have at most a modest effect on our estimates.

probabilities at the individual level using a maximum likelihood estimator. We allow the transition probabilities to depend on x_{it} which includes age, sex, current health status, marital status, and permanent income. In particular we include a third order age polynomial, indicators for gender and marital status (interacted with an age quadratic), single woman (interacted with a first order age polynomial), contemporaneous indicators for health (interacted with age), and a quadratic in permanent income (interacted with a first order polynomial in age and marital status).

Using the estimated transition probabilities, we simulate demographic histories, beginning at age 70, for different gender-PI-health combinations. We report moments of these histories in Section 6.1.1.

G.2 Income

Our procedure for estimating the income profiles described in section 6.1.2 is part of the procedure for inferring permanent income described in Appendix B. We set the annual rate of return on saving to 4%.

G.3 The tax system

We parameterize the tax system using the functional form developed by Gouveia and Strauss [29]. Average tax rates for total pre-tax income, y , are given by

$$t(y, \tau_{f_t}) = b_{f_t} [1 - (s_{f_t} y^{p_{f_t}} + 1)^{\frac{-1}{p_{f_t}}}], \quad (32)$$

where b_{f_t} , s_{f_t} and p_{f_t} are parameters of the tax system that can differ between couples and singles, and $\tau_{f_t} = (b_{f_t}, s_{f_t}, p_{f_t})$ is their union. We use the estimates of τ for married and singles households with no children provided by Guner et al. [30]. After-tax income $\Upsilon(y, \tau_{f_t})$ is then

$$\Upsilon(y, f_t) = (1 - t(y, \tau_{f_t})) \times y. \quad (33)$$

To reflect differences in the generosity of Medicaid, and other means tested social insurance programs available to the elderly, between couples and singles we impose that the consumption floor for couples is 150% of the value we estimate for singles.

G.4 Medical spending

Our medical spending measure is the sum of expenditures paid out-of-pocket plus those paid by Medicaid (see Appendix D for its construction). The AHEAD's

measure of medical spending is backward-looking: in each wave, the household is asked about the expenses it incurred since the previous interview.

Let $m_{i,t}$ denote the expenses incurred by household i between ages $t-1$ and t . We observe household's health at the beginning and the end of this interval, that is, at the time of the interview conducted at age $t-1$ and at the time of the interview conducted at age t . We assume that medical expenses depend upon a household's PI, its family structure at both $t-1$ and t , the health of its members at both dates, and an idiosyncratic component $\psi_{i,t}$:

$$\ln m_{i,t} = m(hs_{i,t-1}^h, hs_{i,t-1}^w, hs_{i,t}^h, hs_{i,t}^w, f_{i,t-1}, f_{i,t}, I_i, t) + v_{i,t}, \quad (34)$$

$$v_{i,t} = \sigma(hs_{i,t-1}^h, hs_{i,t-1}^w, hs_{i,t}^h, hs_{i,t}^w, f_{i,t-1}, f_{i,t}, I_i, t) \times \psi_{i,t}. \quad (35)$$

We normalize the variance of ψ to 1, hence $\sigma^2(\cdot)$ gives the conditional variance of v .

Including current and lagged family structure indicators allows us to account for the jump in medical spending that occurs in the period a family member dies. Likewise, including health indicators for both periods allows us to distinguish persistent health episodes from transitory ones.

To show how we estimate $m(\cdot)$ and $\sigma(\cdot)$, write Equation (6) as

$$\ln m_{it} = x_{1i}\beta_1 + x_{2it}\beta_2 + \vartheta_i + \varsigma_{it}, \quad (36)$$

where x_{1i} denotes a vector of time-invariant variables, x_{2it} denotes a vector of time-varying variables, ϑ_i is an unobserved person-specific term, and ς_{it} captures any remaining variation. We assume that $E(\varsigma_{it} | \vartheta_i) = 0$.

We estimate Equation (36) in three steps. First, we regress log medical spending on the time-varying factors in Equation (34), namely age, household structure, and health, and interaction terms (such as gender and PI interacted with the time varying variables) using a fixed effects estimator. In particular we regress log medical spending on a fourth order age polynomial, indicators for single man (interacted with an age quadratic), single woman (interacted with an age polynomial), the contemporaneous and lagged values of indicators for {man in bad health, married man in a nursing home, single man in a nursing home, woman in bad health, married woman in a nursing home, single woman in a nursing home}, whether the man died (interacted with age and permanent income), whether the woman died (interacted with age, and permanent income).

Because fixed effects regression cannot identify the effects of time-invariant factors, which are combined into the estimated fixed effects, in the second step we collect the residuals from the first regression, inclusive of the estimated fixed effects, and regress them on the time-invariant factors, namely a quadratic in permanent

income and a set of cohort dummies. The level of $m(\cdot)$ is set to be consistent with the outcomes of the cohort aged 72-76 in 1996.

A key feature of our spending model is that the conditional variance and the conditional mean of medical spending depends on demographic and socioeconomic factors, through the function $\sigma(\cdot)$ shown in Equation (35). In the third step of our estimation procedure, we use the coefficients for $m(\cdot)$ in hand, found in the previous two steps, to back out the residual v from Equation (6). To find $\widehat{\sigma^2(\cdot)}$, we square the residuals and regress v^2 on the demographic and socioeconomic variables in Equation (35).

We assume that $\psi_{i,t}$ can be decomposed as

$$\psi_{i,t} = \zeta_{i,t} + \xi_{i,t}, \quad \xi_{i,t} \sim N(0, \sigma_\xi^2), \quad (37)$$

$$\zeta_{i,t} = \rho_m \zeta_{i,t-2} + \epsilon_{i,t}, \quad \epsilon_{i,t} \sim N(0, \sigma_\epsilon^2), \quad (38)$$

where $\xi_{i,t}$ and $\epsilon_{i,t}$ are serially and mutually independent. With the variance of $\psi_{i,t}$ normalized to 1, σ_ξ^2 can be interpreted as the fraction of idiosyncratic variance due to transitory shocks.

We estimate the parameters of Equations (7) and (8) using a standard error components method. Although the estimation procedure makes no assumptions on the distribution of the error terms $\psi_{i,t}$, we assume normality in the simulations. [27] shows that if the data are carefully constructed, normality captures well the far right tail of the medical spending distribution.²⁰

Approximately 40% of the cross-sectional variation in log medical spending is explained by observables, which are quite persistent. Of the remaining cross-sectional variation, 40% comes from the persistent shock ζ and 60% from the transitory shock ξ . In keeping with the results in Feenber and Skinner [25], French and Jones [27] and De Nardi, French and Jones [18], we estimate substantial persistence in the persistent component, with $\rho_m = 0.85$.

²⁰To help us match the distribution of medical spending, we bottom code medical spending at 10% of average medical spending. French and Jones [27] also bottom code the data to match the far right tail of medical spending. Because we include Medicare B payments in our medical spending measure, which most elderly households pay, for the vast majority of households these bottom coding decisions are not important.

Appendix H: The marginal propensity to consume

Consider an individual with no income and no medical expenses. In the last period of life, this individual solves

$$\max_c \frac{1}{1-\nu} c^{1-\nu} + \beta \phi_0 \frac{1}{1-\nu} ((1+r)(a-c) + \kappa_0)^{1-\nu},$$

where c denotes consumption and a denotes beginning-of-period-wealth. The first-order condition for an interior solution is

$$c^{-\nu} = \beta \phi_0 (1+r) ((1+r)(a-c) + \kappa_0)^{-\nu},$$

Define

$$\varphi = [\beta \phi_0 (1+r)]^{1/\nu}. \quad (39)$$

We can then solve for bequests, $b = (1+r)(a-c)$:

$$\varphi \left(a - \frac{b}{1+r} \right) = b + \kappa_0 \Rightarrow b = \frac{1+r}{1+r+\varphi} (\varphi a - \kappa_0). \quad (40)$$

It immediately follows that for intentional bequests, the marginal propensity to bequeath out of wealth, $\frac{\partial}{\partial a} \left(\frac{b}{1+r} \right)$, is

$$MPB = \frac{\varphi}{1+r+\varphi}. \quad (41)$$

With non-negative bequests, however, it follows from Equation (40) that the propensity to bequeath is identically zero unless

$$a > \underline{a} = \frac{\kappa_0}{\varphi}. \quad (42)$$

It is widely accepted that MPB (or $MPC = 1 - MPB$) and \underline{a} provide a better characterization of the bequest motive than ϕ_0 and κ_0 (see, e.g., De Nardi et al., 2010 or Lockwood, 2018). When one expects to die at a future date rather than next period, the effective threshold where the bequest motive becomes operative will be higher than \underline{a} .

There is a one-to-one mapping between the vectors (MPB, \underline{a}) and (ϕ_0, κ_0) , and thus if one of these objects can be identified, the other one can be as well. Using equations (42), we know $\kappa_0 = \varphi \underline{a}$. It then follows from equations (39) and (41) that:

$$\begin{aligned} \varphi &= (1+r) \frac{MPB}{1-MPB} \\ \Rightarrow \phi_0 &= \frac{\varphi^\nu}{\beta(1+r)} = \left[(1+r) \frac{MPB}{1-MPB} \right]^\nu \frac{1}{\beta(1+r)}. \end{aligned}$$

We also need to map parameters between models with one-year and two-year period lengths. An important intermediate step in this process is to consider the sequence of 1-year problems where the individual lives two periods and then dies. We will use hats to designate parameters for this one-year model, and unmarked parameters for the two-year framework. To simplify the algebra, suppose that $\hat{\beta}(1 + \hat{r}) = 1$, so that \hat{c} is the same in the two periods of life. The present value budget constraint is

$$a = \hat{c} \left(\frac{2 + \hat{r}}{1 + \hat{r}} \right) + \frac{\hat{b}}{(1 + \hat{r})^2},$$

which can be rewritten as

$$\hat{c} = \left(\frac{1 + \hat{r}}{2 + \hat{r}} \right) \left(a - \frac{\hat{b}}{(1 + \hat{r})^2} \right). \quad (43)$$

If the marginal propensity to bequeath out of starting wealth is $\widehat{MPB} = \frac{\partial}{\partial a} \left(\frac{\hat{b}}{(1 + \hat{r})^2} \right)$, the marginal propensity to consume on a **annual** basis, found by differentiating equation (43), is:

$$\widehat{MPC} = \frac{\partial \hat{c}}{\partial a} = \frac{1 + \hat{r}}{2 + \hat{r}} \left[1 - \widehat{MPB} \right]. \quad (44)$$

Next, consider the asset threshold \hat{a} . Recall that in the final period of life, wealth below \hat{a} are consumed. Let \tilde{a} denote the asset threshold **two** periods prior to death. By definition, when wealth two periods prior to death equal \tilde{a} , $\hat{b} = 0$. Equation (43) then implies that the asset threshold **one** period prior to death is

$$\hat{a} = (1 + \hat{r}) \left(\tilde{a} - \frac{1 + \hat{r}}{2 + \hat{r}} \tilde{a} \right) = \frac{1 + \hat{r}}{2 + \hat{r}} \tilde{a}. \quad (45)$$

To complete the mapping, suppose that $1 + r = (1 + \hat{r})^2$, $MPB = \widehat{MPB}$, and $a = \tilde{a}$. Equations (44) and (45) then imply that

$$\widehat{MPC} = \frac{1 + \hat{r}}{2 + \hat{r}} MPC, \quad (46)$$

$$\hat{a} = \frac{1 + \hat{r}}{2 + \hat{r}} a. \quad (47)$$

In other words, when we operate at a one-year frequency, the marginal propensity to consume and the asset threshold are (more or less) half their two-year counterparts.

Appendix I: Additional model fits

Figure 17 compares the Median asset profiles generated by the model (dashed line) to those in the data (solid line) for all of the birth cohorts and permanent income terciles that we match in estimation. We reproduce profiles for the cohorts reported in the main text, but also show profiles for the cohorts that were omitted. In all cohorts, the model matches differences in the level and trajectory of wealth by age and household permanent income.




Figure 17: Median wealth. Top panel: current couples. Bottom panel current singles: Solid lines: cohorts aged 72-77, 84-89 in 1996 in the left-hand panel and ages 78-83 and 90+ in 1996 in the right-hand side panel. Dashed lines: model simulations.

Figure 18 plots the corresponding data for the 75th percentile of the asset distribution. Comparing the 75th wealth percentiles to the median ones towards the beginning of the retirement period, when we start tracking our sample, shows that




Figure 18: 75th wealth percentile. Top panel: current couples. Bottom panel current singles: Solid lines: cohorts aged 72-77, 84-89 in 1996 in the left-hand panel and ages 78-83 and 90+ in 1996 in the right-hand side panel. Dashed lines: model simulations.

the 75th asset percentiles are about twice the medians for the two highest PI groups and much larger than twice the median for the lowest PI tercile. In addition, the 75th asset percentiles show even more evidence of wealth accumulation during retirement than do median wealth, both for couples and singles. Finally, the 75th percentiles suggest that current couples tend to save more during their retirement than singles, even given similar initial wealth. As with the medians, the model does a good job of matching the 75th percentile of asset holdings across PI groups and ages, both for singles and couples. For example, it reproduces the flatter profiles of lower PI groups, while still generating eventual wealth holdings of more than \$1,000,000 for the top terciles.

Figure 19 plots the corresponding data for the 25th percentile of the asset distribution. Comparing the 25th wealth percentiles to the medians towards the beginning




Figure 19: 25th wealth percentile. Top panel: current couples. Bottom panel current singles: Solid lines: cohorts aged 72-77, 84-89 in 1996 in the left-hand panel and ages 78-83 and 90+ in 1996 in the right-hand side panel. Dashed lines: model simulations.

of the retirement period, when we start tracking our sample, shows that the 25th asset percentiles are about half the median ones for the two highest PI groups, with the exception of the lowest PI tercile for singles, where it is almost always 0. In addition, the 25th asset percentiles show even more evidence that couples and singles' wealth evolves differently as the wealth holdings of singles decline. As with the 75th percentile, it seems that current couples tend to save more during their retirement than singles, even given similar initial wealth. Our model, in addition to matching the median and the 75th percentile of asset holdings by age, PI, and cohort, also does a good job of generating the differential deaccumulation we see for the 25th percentile.

Finally, Figure 20 compares the Medicaid recipiency profiles generated by the model (dashed line) to those in the data (solid line) for the same birth cohorts and permanent income terciles and shows that our model matches important patterns of




Figure 20: Medicaid Recipiency. Solid lines: cohorts aged 72-77, 84-89 in 1996 in panel (a) and ages 78-83 and 90+ in 1996 in panel (b). Dashed lines: model simulations.

Medicaid usage. We reproduce the cohorts reported in the main text and add the cohorts not shown there. The model matches difference in the level of usage by age and household permanent income for all cohorts. It does especially well in capturing the run-up in Medicaid use by wealthier households at very old ages.

Because individuals who are single, poor or sick die at younger ages, in an unbalanced panel the sample composition changes over time, introducing mortality bias. Figure 21 compares the mortality bias found in the AHEAD data (left panel) with that predicted by the model (right panel). The solid lines show median wealth for all households (both singles and couples) observed at a given point in time, even if all members died in a subsequent wave, that is, the unbalanced panel. The dashed lines show median wealth for the subsample of households with at least one member still alive in the final wave, that is, the balanced panel. They show that the asset profiles for households who survived to the final wave (the balanced panel) have much more of a downward slope. The difference between the two sets of profiles confirms that people who died during our sample period tended to have lower wealth than the survivors.

Our model matches the extent of the mortality bias found in the data, without being required to do so by construction. As in the data, restricting the profiles to long-term survivors reveals much more asset decumulation.




Figure 21: Mortality Bias. Left panel: data. Right panel: model simulations. Solid lines: all household in the data at a given point in time. Dashed lines: households with at least one member surviving to 2014.