

Couples' and Singles' Savings After Retirement

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

February 10, 2019

Abstract

Not only do retired couples hold more assets than singles, but high-income couples grow their savings during most of their retirement period. Why are they saving so much? Do they expect high medical expenses, do they deeply care about the welfare of their surviving spouse upon death, or do they want to leave bequests to others? What happens to savings when one of the members of the couple dies and why? We build a model of retired couples and singles facing uncertain longevity and medical expenses in which couples and singles can have different bequest motives. Both might care about heirs, but couples might also care about their surviving spouse. We use the AHEAD data and the method of simulated moments to estimate our model and disentangle the importance of saving motives of couples and singles.

*Mariacristina De Nardi: UCL, Federal Reserve Bank of Chicago, IFS, CEPR, and NBER, e-mail: denardim@nber.org. Eric French: UCL, CEPR, and IFS, e-mail: eric.french.econ@gmail.com. John Bailey Jones: Federal Reserve Bank of Richmond, jbjones.albany@gmail.com. McGee: UCL and IFS, rory.mcgee.13@ucl.ac.uk. We thank Wilbert van der Klaauw, Monica Costa Dias, Tom Crossley, for helpful comments and Karolos Arapakis, Jeremy McCauley, and Ali Merali for research assistance, and we gratefully acknowledge support from the ERC, grant 614328 “Savings and Risks”, the Social Security Administration through the Michigan Retirement Research Center (MRRC grant UM17-15), and 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 Federal Reserve Bank of Chicago, the Federal Reserve Bank of Richmond, the IFS, the CEPR, or the NSF.

We also evaluate the effects of a Medicaid expansion on the savings and bequests of couples and singles.

1 Introduction

There is considerable agreement on the reasons why single retirees hold assets. Recent work has shown that much of their saving behavior can be explained by the possibility of expensive medical conditions at older ages (De Nardi, French, and Jones [22, 23]), or by bequest motives that act as luxury goods (De Nardi [18]). A number of recent studies seek to disentangle these two motives by including other features of the data, such as survey responses (Ameriks et al. [4] and [3]), long-term care insurance purchases (Lockwood [40]), or Medicaid usage (De Nardi et al. [24]).

Less attention has been given to the saving of retired couples. Couples tend to be richer than singles and face a more complex set of risks. To what extent do members of a couple care about their own consumption and medical expenses, as opposed to leaving bequests to their surviving spouse and other heirs? What happens to savings when one of the members of the couple dies, and why?

In the Assets and Health Dynamics of the Oldest Old (AHEAD) dataset, about 50% of individuals aged 70 or older are in a couple, while about 50% are single. Because retired couples are not only numerous, but also richer, they hold a large share of the elderly's wealth. Better understanding why they save matters, including in the assessment of policy reforms.

Couples and singles differ in important ways. Being in a couple allows its members to pool their longevity and medical expense risks, 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 (Braun et al. [11]). A couple might also care about leaving bequests to the surviving spouse, in addition to other heirs.

We start our analysis of couples by establishing facts about their retirement savings and the medical expenses and longevity risks that they face. We then develop a model of couples and singles during retirement, in a framework that incorporates observed heterogeneity in life expectancy and medical expenses, and that explicitly models means-tested social insurance. We allow for bequest motives to other heirs for both couples and singles, and for altruism towards the surviving spouse for couples. We estimate our model using the Method of Simulated Moments (MSM). Finally, we use the model to evaluate the extent to which the savings of retirees are driven by bequest motives towards the other spouse, bequest motives towards other heirs, and medical expenses. We also evaluate the effects of a Medicaid expansion on

the savings and bequests.

In the first part of our analysis we document that while the savings of retired singles tend to stay roughly constant or fall over time, those of couples increase as long as both spouses are alive. We also find that assets drop sharply at the time of death of each spouse and that, by the time the second spouse dies, a large fraction of the wealth of the original couple has vanished. Depending on the specification, assets decline \$30,000-\$60,000 at the time of an individual's death. A large share of this drop, but not all of it, is explained by the high medical expenses at the time of death. For example, out of pocket medical spending plus death expenses are approximately \$20,000 during the year of death (whereas medical spending is \$6,000 per year for similarly aged people who do not die).

Our second-step estimates allow us to evaluate to what extent the risk sharing and economies of scale of a couple help insure against longevity and medical-expense risk and to what extent couples get a better or worse deal from publicly provided health insurance. They will also allow us to estimate bequest motives towards the surviving spouse, in contrast to bequests to children or others. Finally, it will enable us to study in what ways the responses of a couple will differ from the responses of a single person when, for example, public health insurance becomes more limited, or its quality worsens, or some of its means-testing criteria become tighter.

2 Related Literature

Much of the previous structural literature only studies singles. In a previous paper (De Nardi, French, and Jones [23]) we show that post-retirement medical expenses and government-provided insurance are important to explaining the saving patterns of U.S. single retirees at all income levels, including high permanent-income individuals who keep large amounts of assets until very late in life. These savings patterns are due to two important features of out-of-pocket medical expenses. First, out-of-pocket medical and nursing-home expenses can be large. Second, average medical expenditures rise very rapidly with age and permanent income. Medical expenses that rise with age provide the elderly with a strong incentive to save, and medical expenses that rise with permanent income encourage the rich to be more frugal. In other work, we showed that heterogeneous life expectancy is important to matching the savings patterns of retired elderly singles (De Nardi, French

and Jones [22]).

Poterba et al. ([48]) show that relatively little dissaving occurs amongst retirees whose family composition does not change, but that assets fall significantly when households lose a spouse. We find similar results. Furthermore, as in French et al. ([30]), we document significant drops in assets when the last member of the household passes away.

Previous literature had shown that high income individuals live longer than low income individuals (see Attanasio and Emmerson [6] and Deaton and Paxson [20]). This means that high income households must save a larger share of their lifetime wealth if they are to smooth consumption over their retirement. Differential mortality rates thus provide a potential explanation for why high income households have higher savings rates than low income households. We extend the analysis along this dimension by explicitly modeling the interaction of life expectancy for individuals in couples and the differential life expectancy for couples and singles.

Even in presence of social insurance such as Medicare and Medicaid, households face potentially large out-of-pocket medical and nursing home expenses (see French and Jones [32, 31], Palumbo [46], Feenberg and Skinner [29], and Marshall et al. [41]). The risk of incurring such expenses might generate precautionary savings, over and above those accumulated against the risk of living a very long life ([39]).

Hubbard, Skinner and Zeldes [36] argue that means-tested social insurance programs such as Supplemental Security Income and Medicaid provide strong incentives for low income individuals not to save. De Nardi et al. [25] finds that these effects extend to singles in higher permanent income quintiles as well. In this paper, to allow for these important effects, we model means-tested social insurance explicitly for both singles and couples.

Bequest motives could be another reason why households, and especially those with high permanent income, retain high levels of assets at very old ages (Dynan, Skinner and Zeldes [27] and Americks et al. [4]). De Nardi [18] and Castañeda et al. [16] argue that bequest motives are necessary to explain why the observed distribution of wealth is more skewed and concentrated than the distribution of income (Quadrini et al. [21]). De Nardi et al [24] shows that bequest motives help fit both assets and Medicaid reciprocity profiles for singles. We allow for a richer structure of bequest motives, in that couples might want to leave resources to the surviving spouse, children and other heirs, while singles might want to leave bequests to children and other heirs.

Previous quantitative papers on savings have used simpler models that omit one or more of these features. Hurd [37] estimates a structural model of bequest behavior in which the time of death is the only source of uncertainty. Palumbo [46] focuses on the effect of medical expenses and uncertain lifetimes, but omits bequests. Dynan, Skinner and Zeldes [27, 28] consider the interaction of mortality risk, medical expense risk and bequests, but use a stylized two-period model. Moreover, none of these papers model household survival dynamics, assuming instead that households (while “alive”) always have the same composition. In contrast, we explicitly model household survival dynamics: when the first household member dies, assets are optimally split among the surviving spouse and other heirs. Although Hurd [38] extends his earlier model to include household survival dynamics, he omits medical expense risk, and, in contrast to his earlier work, he does not estimate his model.

Our work also complements the one on retirement behaviour by couples by Blau and Gilleskie [10], Casanova [15], and Gallipoli and Turner [34].

Including couples and simultaneously considering bequest motives, social insurance, uncertain medical expenses, and uncertain life expectancy is important for at least two reasons. First, Dynan, Skinner and Zeldes [28] argue that explaining why the rich have high savings rates requires a model with precautionary motives, bequest motives and social insurance. Second, simultaneously considering multiple savings motives allows us to identify their relative strengths. This is essential for policy analysis. For example, the effects of estate taxes depend critically on whether rich elderly households save mainly for precautionary reasons, or mainly to leave bequests (see for example Gale and Perozek [33]).

3 Data and key facts

We use data from the Asset and Health Dynamics Among the Oldest Old (AHEAD) dataset. The AHEAD is a sample of non-institutionalized individuals, aged 70 or older in 1993. These individuals were interviewed in late 1993/early 1994, and again in 1996, 1998, 2000, 2002, 2004, 2006, 2008, and 2010.¹ Appendix A describes the details of our sample selection and data work.

¹We do not use 1994 assets, nor medical expenses, due to underreporting (Rohwedder et al. [50]).

We only consider retired households to abstract from the retirement decisions and focus on the determinants of savings and consumption. 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 break the data into 5 cohorts. The first cohort consists of individuals that were ages 72-76 in 1996; the second cohort contains ages 77-81; the third ages 82-86; the fourth ages 87-91; and the final cohort, for sample size reasons, contains ages 92-102. Even with the longer age interval, the final cohort contains relatively few observations. In the interest of clarity, we exclude this cohort from our graphs, but we use all cells with at least 10 observations when estimating the model.

We use data for 8 different years; 1996, 1998, 2000, 2002, 2004, 2006, 2008, and 2010. We calculate summary statistics (e.g., medians), cohort-by-cohort, for surviving individuals in each calendar year—we use an unbalanced panel. We then construct life-cycle profiles by ordering the summary statistics by cohort and age at each year of observation. Moving from the left-hand-side to the right-hand-side of our graphs, we thus show data for four cohorts, with each cohort's data starting out at the cohort's average age in 1996.

Since we want to understand the role of income, we further stratify the data by post-retirement permanent income (PI). Hence, for each cohort our graphs usually display several horizontal lines showing, for example, median assets in each PI group in each calendar year. These lines also identify the moment conditions we use when estimating the model.

Our PI measure can be thought of as the level of income if there were two people in the household at age 70. We measure post-retirement PI using non-asset non-social insurance annuitized income and the methods described in appendix B . The method maps the relationship between current income and PI, adjusted for age and household structure. The income measure includes the value of Social Security benefits, defined benefit pension benefits, veterans benefits and annuities. Since we model means-tested social insurance from SSI and Medicaid explicitly through our consumption floor, we do not include SSI transfers. Because there is a roughly monotonic relationship between lifetime earnings and the income variables that we use, our measure of post-retirement PI is also a good measure of lifetime permanent income.

3.1 Savings and social insurance

The following figures display median assets, conditional on birth cohort and income quintile, for different configurations of couples and singles.

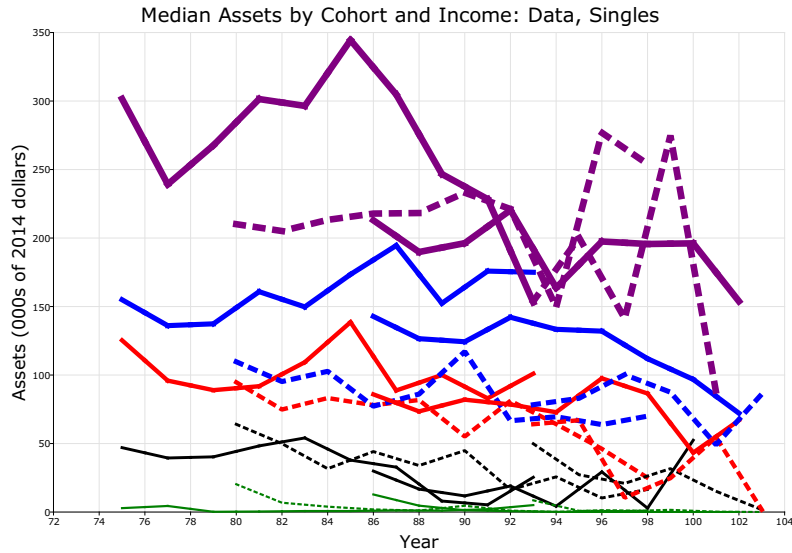


Figure 1: Assets by age, income, and cohort for currently singles

Figure 1 displays the net worth (or assets, from now on, since the retirees have very little debt) of the unbalanced panel of people who are currently single, that is, either enter the sample because single at the start of the period, or join it later as they become single (the picture would look very similar if we were to limit the sample to only people who were initially single only, see Appendix D). The figure displays the important observation that the savings of elderly singles depends on their income (which is predetermined at retirement). Elderly singles with the lowest lifetime incomes reach retirement with little to no asset and accumulate no additional assets as they age. Rather, they rely on annuitized income and government insurance to fund their consumption and their medical expenses. In the middle of the income distribution, these elderly singles start our retirement with some assets, but then they run them down. At the top of the income distribution, instead, assets of the survivors stay roughly constant as they age and survive. In our previous work (De Nardi, French, and Jones [23] and De Nardi, French, and Jones [24]) we have shown that the age profiles of medical expenses in

old age, which increase with age and income, and the availability of means-tested social insurance (such as Medicaid and SSA) are key drivers of these patterns and that bequest motives appear to explain little of the lack or slow decumulation of the higher income singles.




Figure 2: Assets by age, income, and cohort for intact couples

Figure 2 displays the assets of “intact couples,” that is couples that start out and continue as such. Thus, if one or both components die, we drop them from the figure. This figure displays some important facts. First, conditional on the same income quintile, even couples in the lowest income quintiles reach retirement with some assets. Second, the asset profiles of the lower to middle income surviving couples display no decumulation, even at the lowest levels of income. Third, the highest income intact couples actually have increasing savings in age for much of their retirement, with drops in their assets only starting to appear after age 85.

Figure 3 reports median assets for the population of those who are initially in a couple. Some of these couples lose one spouse during the period in which we observe them, in which case we report the assets of the surviving spouse. Relative to intact couples, initial couples in the higher PI quintiles decumulate their assets more quickly. Some of this decumulation is due to the loss of assets that occurs when one of the spouses dies.




Figure 3: Median assets for households initial couples. Each line represents median assets for a cohort-income cell, traced over the time period 1996-2010. Thicker lines refer to higher permanent income groups. Solid lines: cohorts ages 72-76 and 82-86 in 1996. Dashed lines: ages 82-86 and 92-96 in 1996.

To formalize and quantify this observation, we turn to documenting the savings profiles of couples when one of their members dies. To do so, we perform regression analysis.²

Figure 4 reports the predicted assets of a couple starting out, respectively, in the top PI income quintile, and in the bottom one, and display the assets under three scenarios. Under the first one, the couple remains intact until age 90. Under the second one, the men dies at age 80, while under the third one, the woman dies at age 80. For both PI levels, assets stay either increase or

²We regress assets on a household fixed effect, dummies for household status (couple, single man, single woman, everyone dead), a polynomial in age, and interactions between these and permanent income. Including the household specific fixed effect means that we can track wealth of the same households, when one member of the household dies. The fixed effect includes all time invariant factors, including permanent income, so we cannot include permanent income directly. To include permanent income directly, we then regress the estimated fixed effect on the permanent income of that household.




Figure 4: Assets, conditional on permanent income and family structure, for a couple (top lines for each PI) and in the case in which there is a switch to becoming a single men or single woman (other two lines at each PI) at age 80. The survivors to the first death then die at age 90.

stay roughly constant if both partners are alive, but display a significant drop at the death of one of the spouses. On average, wealth declines approximately \$60,000 at the time of the husband's death and approximately \$30,000 at the time of the wife's death.

The simpler exercise of merely tabulating the wealth decline when one member of the household dies suggests smaller estimates of the wealth decline at the time of death of a spouse. This procedure suggests a decline of wealth of \$30,000 at the death of a spouse. A large share of this drop, but not all of it, is explained by the high medical expenses at the time of death. For example, out of pocket medical spending plus death expenses are approximately \$20,000 during the year of death (whereas medical spending is \$6,000 per year for similarly aged people who do not die).

To measure the decline in wealth at the time of death of the final member of the family, we exploit the exit surveys, which include information on the

heirs' reports of the value of the estate.³ In addition, we also use data from post-exit interviews, which are follow up surveys of heirs, to better measure wealth held in the estate. The results suggest even larger declines in wealth when the final member of the household dies. For those at the top of the income distribution, the death of a single person results in a wealth decline of approximately \$110,000, whereas the decline is closer to \$70,000 for those at the bottom of the income distribution.

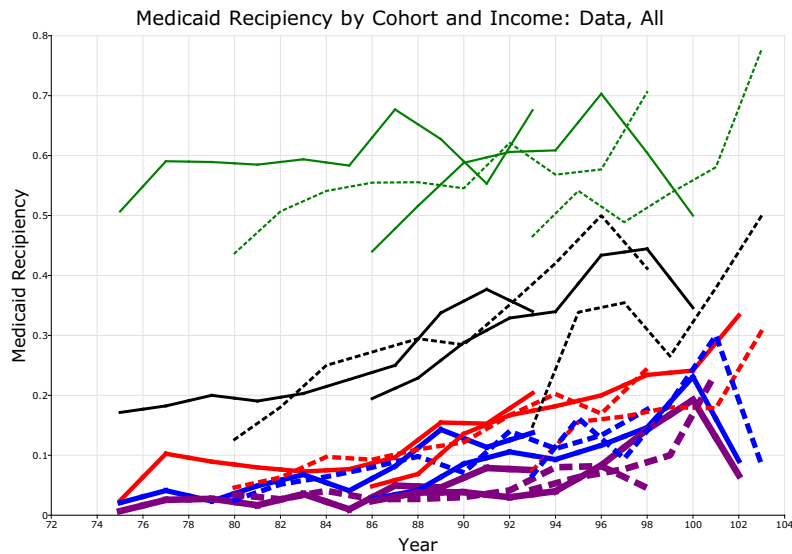


Figure 5: Medicaid reciprocity by age, income, and cohort

Because previous work has shown that means-tested programs are very important determinants of savings (see our related literature section) and because our goal is understanding the savings of the retirees, it is important to take into account the extent to which these retirees benefit from mean-tested government insurance programs that help pay for their medical expenses, such as Medicaid. Figure 5 plots Medicaid reciprocity rates conditional on age, income, and cohort for all of the retirees in our sample. Several patterns are worth pointing out. First, the lines are flipped compared to the asset graphs. It is the lowest income retirees that are more likely to end up on Medicaid. Second, and consistently with the asset-tested nature of these

³Some couples lose both members between successive AHEAD waves, but this is a relatively infrequent occurrence.

programs (high assets people are not eligible) and the asset profiles that we have described over the retirement period, Medicaid reciprocity of increases fast earlier in their retirement for retirees with intermediate values of income, while it only tends to increase at very advanced ages for higher income people.

3.2 Health and mortality

Income Percentile	Nursing Home	<u>Men</u>		<u>Women</u>		All	
		Bad Health	Good Health	Nursing Home	Bad Health		Good Health
Singles							
10	3.03	6.92	8.68	4.07	11.29	13.18	10.17
50	3.02	7.78	10.29	4.05	12.29	14.86	11.51
90	2.91	8.11	10.94	3.80	12.51	15.37	11.99
Couples							
10	2.73	7.83	9.82	3.95	12.10	14.05	11.30
50	2.77	9.39	12.18	3.99	13.74	16.27	13.37
90	2.74	10.39	13.50	3.88	14.59	17.28	14.45
Single Men							8.99
Married Men							11.52
Single Women							13.88
Married Women							15.77
Oldest Survivor							17.90
Probability that Oldest Survivor is Woman							63.7%

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

Table 1 shows life expectancies for singles and couples at age 70, respectively.⁴ Rich people, women, married people, and healthy people live much longer than their poor, male, single, and sick counterparts.

⁴We estimate health transitions and mortality rates simultaneously by fitting the tran-

More specifically, a single man at the 10th permanent income percentile in a nursing home expects to live only 3.0 more years, while a single woman at the 90th percentile in good 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 over all health states. It shows that singles at the 10th percentile of the permanent income distribution live on average 10.2 years, whereas singles at the 90th percentile live on average 12.0 years. It also highlights that a 70 year old man at the 10th permanent income percentile in a nursing home married to a 70 year old woman with the same health expects to live only 2.7 more years, while a 70 year old married woman at the 90th percentile in good health (married to a 70 year old man also in good health) expects to live 17.3 more years.

Conditional only on gender, people in couples live about 2 years longer than singles: single 70 year old 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 much 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 number of years of remaining life of the oldest survivor in a household when both the man and the woman are 70. On average the last survivor lives an additional 17.9 years. In that case, woman is the oldest survivor 63.7% of the time.

Table 2 shows that single men and women face on average a 26.4% and 37.2% chance of being in a nursing home for an extended stay, respectively, while married men and women face on average a 19.5% and 36.3% 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 often become single as their partner dies and as they age. Furthermore, married people tend to live longer than singles, so they have more years of life to potentially enter a nursing home. Permanent income has only a small effect on ever being in a nursing home. Those with a high permanent income are less likely to be in a nursing home at each age, but they tend to live longer.

sitions observed in the HRS to a multinomial logit model. We allow the transition probabilities to depend on age, sex, current health status, marital status, permanent income, as well as polynomials and interactions of these variables. Using the estimated transition probabilities, we simulate demographic histories, beginning at age 70, for different gender-PI-health combinations. All tables use the appropriate distribution of people over state variables to compute the number which is object of interest.

Income Percentile	<u>Men</u>		<u>Women</u>		All
	Bad Health	Good Health	Bad Health	Good Health	
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.

3.3 Income

The goal of this subsection is to illustrate how income changes for couples when one of the two spouses dies. We model income as a function of permanent income.⁵ Figure 6 presents predicted income profiles for those at the 20th and 80th percentiles of the permanent income distribution. For each permanent income level, we display three scenarios, all commencing from the income of a couple. Under the first scenario, the household remains a couple

⁵More specifically, income is a third order polynomial in age, dummies for family structure, family structure interacted with an age trend, and a fifth order polynomial in permanent income. The estimates use a fixed effects estimation procedures, where the fixed effect is a transformation of initial permanent income. Hence the regression results can be interpreted as the effect of changing age or family structure for the same household: see appendix B for details of the procedure.




Figure 6: Income, conditional on permanent income and family structure. Figure assumes all households begin as couples, then either stay in a couple, or switch to being to a single men or single woman at age 80.

until age 100. Under the second one, the man dies at age 80. Under the third one, the woman dies at age 80.

Couples’ average annual income ranges from about \$14,000 to over \$30,000 per year in the 20th and 80th percentiles of the PI distribution, respectively. As a point of comparison, median wealth holdings for these two groups are \$70,000 and \$330,000 at age 74, respectively.

Our estimates suggest 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. These income losses at the death of a spouse reflect the fact that although both Social Security and defined benefit pensions have spousal benefits, these benefits replace only a fraction of the deceased spouse’s income.

More specifically, people can receive benefits either based on their own history of Social security contributions (in which case they are a “retired worker”), or based on their spouse’s or former spouse’s history (in which case they receive the “spouse” or “widow” benefit). A married person who never worked can receive 50% of their spouse’s benefit if their spouse is

alive and is a “retired worker”. The same person can receive up to 100% of their spouse’s benefit if their “retired worker” spouse has died. Thus the household benefit can receive $100\%+50\%=150\%$ of the former worker spouse’s benefit when alive and 100% of the former worker spouse’s benefit when either spouse has died. Thus, after the death of a household member the household would maintain $(100\%/150\%)=67\%$ of the original Social Security benefit and would experience a 33% drop in benefits.

In contrast, a person who earned the same amount as their spouse will not receive a spousal or widow’s benefit. In this case, both spouses in the couple will receive 100% of their own “retired worker” benefit, which is based off of their own earnings history. After the death of a spouse, the household benefit will be $(100\%/(100\%+100\%))=50\%$ of its level when both were alive.

To perform our calculations, we make several assumptions, including that both spouses begin receiving benefits at the normal retirement age. In practice, there are many modifications to this rule, including those to account for the age at which the beneficiary and spouse begin drawing benefits.⁶ Our regression estimates capture the average drop in income at the time of death of a spouse, averaging over those who retire at different ages and have different claiming histories.

3.4 Medical spending

The goal of this subsection is to illustrate how medical expenses change for couples when one of the two spouses dies. In our model, we explicitly model Medicaid, a means-tested program that helps insure medical expenses of the needy. Thus, the appropriate measure of medical expense risk that we have to measure in the data are out-of-pocket medical costs (net of insurance co-pays that are not based on one’s resources), plus Medicaid payments.

The AHEAD data contains 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 Medicaid payments plus out of pocket medical spending. It uses a mixture of both administrative and survey data and is described in greater detail in De Nardi et al [25]. One drawback of the MCBS, however, is that although it has information on marital status and household income, it does not have information

⁶See <https://socialsecurity.gov/planners/retire/yourspouse.html> for more details of calculation of spousal benefits.

on the medical spending or health of the spouse.

To exploit the strength of both datasets, we use the conditional mean matching procedure described in Appendix C to obtain our measured out-of-pocket medical expenditures plus Medicaid payments for the AHEAD data from the MCBS data. The procedure preserves both the mean as well distribution of medical spending, conditional on age, income, out of pocket spending and other variables.

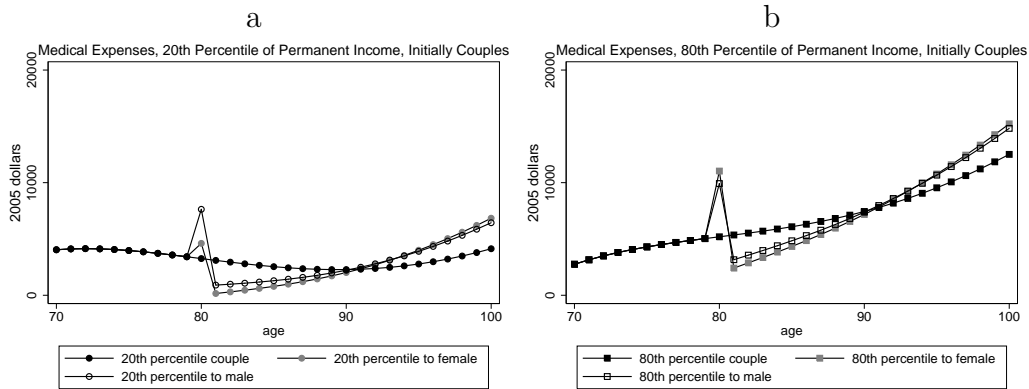


Figure 7: Mean out-of-pocket medical expenditures by age, conditional on permanent income and family structure. Figure assumes all households begin as couples, then potentially change to a single men or single woman at age 80.

Figure 7 displays our predicted medical expense profiles for those at the 20th (left panel) and 80th (right panel) percentiles of the permanent income distribution for those in bad health.⁷

For each permanent income level, we present three scenarios, all of which start out with 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. The jump in medical spending shown at age 80 represents the elevated medical spending during the year of

⁷We model the logarithm of medical spending as a fourth order polynomial in age, dummies for family structure, the health of each family member (including whether that family member just died), family structure interacted with an age trend, and a fifth order polynomial in permanent income. As with the income estimate, we use fixed effects estimation procedures, so the regression results can be interpreted as the effect of changing age or family structure for the same household.

death of a family member and amounts to almost \$5,000 on average. Both average medical spending and the jump at time of death are larger for those with higher permanent income.

The figure shows that, before age 80, average annual medical spending hovers around \$4,000 per year for the couples in both income groups. For those at the 20th percentile of the PI distribution, medical expenses change little with age. However, for couples in the 80th percentile, medical expenses rise to well over \$10,000 per year by age 95.

4 The model

Consider a retired household with family structure f_t (either a single person or a couple), seeking to maximize its expected lifetime utility at household head age t , $t = t_r, t_r + 1, \dots, T + 1$, where t_r is the retirement age, while T is the maximum potential lifespan.

We use w to denote women and h to denote men. Each person's health status, hs^g , $g \in \{h, w\}$, can vary over time. The person can be in a nursing home, in bad health, or in good health.

For tractability, we assume a fixed age gap between the husband and the wife in a couple, so that one age is sufficient to characterize the household. To be consistent with the data frequency, and to reduce computation time, our time period is two years long. We assume that wives are two years younger than their husbands.

4.1 Preferences

Households maximize their utility by choosing savings, bequests and current and future consumption. The annual discount factor is given by β . Each period, the household's utility depends on its total consumption, c , and the health status of each member. The within-period utility function for a single is given by

$$u(c, hs) = (1 + \delta(hs)) \frac{c^{1-\nu}}{1-\nu}, \quad (1)$$

with $\nu \geq 0$. When $\delta(\cdot) = 0$, health status does not affect utility.

The flow utility function for couples is:⁸

$$u^c(c, h_s^h, h_s^w) = [1 + \delta(h_s^h) + 1 + \delta(h_s^w)] \frac{(c/\eta)^{1-\nu}}{1-\nu}, \quad (2)$$

where $1 < \eta \leq 2$ determines the extent to which couples enjoy economies of scale in the transformation of consumption goods to consumption services.

When a household member dies, the estate can be left to the surviving spouse (if there is one) or to other heirs, including the household's children. Estates are subject to estate taxes, but the exemption level during the time period in our sample is above the actual assets of the vast majority of the households in our sample. For this reason, we abstract from explicitly modelling estate taxation. When the final member of the household dies bequests are net of their medical expenses.

We indicate with b the part of the estate that does *not* go to the surviving spouse, and assume that the deceased member of the household derives utility $\theta_j(b)$ from leaving that part of the estate to heirs other than the spouse. The subscript j indicates whether there is a surviving spouse or not, and whether one or two people have just died. In particular, $\theta_0(b)$ gives the utility from bequests for a single person with no surviving spouse, $\theta_1(b)$ gives the utility from bequests when there is a surviving spouse, and $\theta_2(b)$ gives the utility from bequests when both spouses die at the same time. The bequest function takes the form

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

where k_j determines the curvature of the bequest function, and ϕ_j determines its intensity. It was first used to help explain savings over the life cycle by De Nardi [18] and can support several interpretations of the “bequest motive:” dynastic or “warm glow” altruism (as in Becker and Tomes [8], Abel and Warshawsky [1], or Andreoni [5]); strategic motives (as in Bernheim, Schleifer and Summers [9] or Brown [12]); or some form of utility from wealth itself, as in (Carroll [14] and Hurd [37]).

⁸Mazzocco [42] shows that under full commitment, the behavior of a couple can be characterized by a unique utility function if the husband and wife share identical discount factors, identical beliefs and Harmonic Absolute Risk Aversion utility functions with identical curvature parameters.

4.2 Timing

The timing is the following: the household begins the period with a given marital status, which is just widowed man or woman, single man, single woman, or couple.

The people who are just widowed first optimize over bequests left upon the death of the second spouse. After distributing spousal bequests, they become single men or women with a set of given state variables.

Single men, single woman, or couples use the cash on hand they have at the beginning of the period to consume and save. After that, mortality shocks, health shocks, and medical shocks happen, income net of taxes is received, and government transfers take place. As a result, the amount of cash on hand that is carried over into the next period is also realized.

Thus, in case of death, households where the final member(s) has just died leave all their resources to their heirs. When one of the members of the couple dies, the survivor, before taking any other decisions, splits cash on hand at the beginning of the period into bequests left after the death of the first spouse and cash on hand that he or she actually gets to use to consume and save. Our timing assumes, consistently with reality, the medical costs associated with death are collected before any bequests can be made and that Medicaid pays bill that are incurred even when the patient dies.

4.3 Technology and sources of uncertainty

We assume that non-asset income at time t , $y_t(\cdot)$, is a deterministic function of the household's permanent income, I , age, family structure, and gender if single.

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

We do not include received bequests as a source of income, because very few households aged 65 or older receive them.

There are several sources of uncertainty:

1) Health status uncertainty. The transition probabilities for a person's future health status depend on that person's permanent income, age, gender, current health status and current marital status:

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

2) Survival uncertainty. Let $s(I, t, g, h_{s_t}, f_t)$ denote the probability that an individual of gender g is alive at age $t + 1$, conditional on being alive at

age t , having time- t health status hs , enjoying household permanent income I , and having family structure f_t .

3) Medical expense uncertainty at the household level. We define m_{t+1} as the flow of all out-of-pocket medical expenses, including insurance premia, and medical expenses covered by the consumption floor incurred between periods t and $t + 1$. We assume that medical expenses depend upon the health status of each family member (wife w and husband h), household permanent income, family structure, gender if single, age, and an idiosyncratic component, ψ_{t+1} .

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

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

Following Feenberg and Skinner [29] and French and Jones [31], 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 quadrature methods described in Tauchen [51].

4.4 Recursive formulation

Let a_t denote assets at period t and let r denote the constant pre-tax rate of return earned on those assets. Total post-tax income is given by $\Upsilon(r a_t + y_t(\cdot), \tau)$, with the vector τ summarizing the tax structure.

Define the available resources *before* government transfers as

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

Following Hubbard et al. [35, 36], we assume that the transfers bridge the gap between a minimum consumption floor and the household's financial resources. Consistently with the main Medicaid and SSI rules, we can then express government transfers as

$$\begin{aligned} tr_t(\tilde{x}_t, f_t, a_t, y_t) &= \max \left\{ 0, c_{min}(f_t) - (\tilde{x}_t - a_d(f_t, a_t) - y_d(f_t, y_t, a_t)) \right\}, \\ a_d(f_t, a_t) &= \min\{a_t, a_d(f_t)\}, \\ y_d(f_t, y_t, a_t) &= \min\{y_t + ra_t, y_d(f_t)\}, \end{aligned} \quad (10)$$

where $a_d(f_t)$ and $y_d(f_t)$ are the statutory disregard limits for assets and income, respectively. We allow the guaranteed consumption level c_{min} , the asset disregard a_d , and the income disregard y_d to all vary with family structure.

To save on state variables we follow Deaton [19] and express the household's financial resources in terms of cash-on-hand:

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

Households choose consumption and the remaining wealth carried forward to period $t + 1$ is

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

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

$$\tilde{x}_{t+1} = (x_t - c_t) + \Upsilon(r(x_t - c_t) + y_{t+1}(\cdot), \tau) - m_{t+1}, \quad (13)$$

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

The value function for a single individual of gender g can thus be written in terms of cash on hand

$$\begin{aligned} V_t^g(x_t, h_{s_t}, I, \zeta_t) &= \max_{c_t} \left\{ u(c_t, h_{s_t}) + \beta s(I, t, g, h_{s_t}, f_t) \right. \\ &\quad \times E_t \left(V_{t+1}^g(x_{t+1}, h_{s_{t+1}}, I, \zeta_{t+1}) \right) \\ &\quad \left. + \beta(1 - s(I, t, g, h_{s_t}, f_t)) E_t \theta_0(x_{t+1}) \right\}, \end{aligned} \quad (15)$$

subject to equations (4)-(8), (13), (14), and the following constraint that ensures that consumption is at least as high as its floor and that savings are non-negative

$$c_{min}(f_t) \leq c_t \leq x_t, \quad \forall t. \quad (16)$$

The newly widowed (single) individual of gender g has to first distribute bequests after the death of the first spouse before he or she can make savings and consumption decisions as a single person

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

subject to equation

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

which prohibits the surviving spouse from using bequests to become eligible for government transfers. The term ω measures the weight a couple gives to the surviving spouse, as opposed to the intact couple or other heirs, when it takes expectations over future outcomes.

The value function for couples can be written as

$$\begin{aligned} V_t^c(x_t, h s_t^h, h s_t^w, I, \zeta_t) = \max_{c_t} \left\{ & u^c(c_t, h s_t^h, h s_t^w) \right. \\ & + \beta s(I, t, w, h s_t, 1) s(I, t, h, h s_t, 1) E_t \left(V_{t+1}^c(x_{t+1}, h s_{t+1}^h, h s_{t+1}^w, I, \zeta_{t+1}) \right) \\ & + \beta s(I, t, w, h s_t, 1) (1 - s(I, t, h, h s_t, 1)) E_t \left(V_t^{nw}(x_{t+1}^w, h s_{t+1}, I, \zeta_{t+1}) \right) \\ & + \beta (1 - s(I, t, w, h s_t, 1)) s(I, t, h, h s_t, 1) E_t \left(V_{t+1}^{nh}(x_{t+1}^h, h s_{t+1}, I, \zeta_{t+1}) \right) \\ & \left. + \beta (1 - s(I, t, w, h s_t, 1)) (1 - s(I, t, h, h s_t, 1)) E_t \theta_2(x_{t+1}) \right\}, \quad (19) \end{aligned}$$

subject to equations (4)-(8), (13), (14), and (16). The dating of the continuation value for widows, $V_t^w(\cdot)$, reflects that wives are one model period (two years) younger than their husbands.

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. In particular, we estimate health transitions, out-of-pocket medical expenses, and mortality rates from raw demographic data. We calibrate the household economies of scale parameter. We set the consumption floor for couples to be a multiple of the consumption floor for singles, that we estimate.

In the second step, we estimate the rest of the model’s parameters (discount factor, risk aversion, health preference shifter, the consumption floor for singles, and bequest parameters)

$$\Delta = (\beta, \nu, \delta, \omega, c_{min}(f_t = s), \phi_0, \phi_1, \phi_2, k_0, k_1, k_2)$$

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 simulated life-cycle decision profiles to “best match” (as measured by a GMM criterion function) the profiles from the data.

Because our underlying motivations are to explain why elderly individuals retain so many assets and to explain why individuals with high income save at a higher rate, we match median assets by cohort, age, and permanent income. Because we wish to study differences in savings patterns of couples and singles, we match profiles for the singles and couples separately. Because Medicaid is an important program insuring medical expenses and consumption of the poorest, we also match Medicaid reciprocity.

In particular, the moment conditions that comprise our estimator are given by

1. Median asset holdings by PI-cohort-year for the currently singles who are still alive when observed.
2. Median asset holdings by PI-cohort-year for those who are currently couples with both members currently alive.

When there is a death in a couple, the surviving spouse is included in the current singles’ profile of the appropriate age, cohort, and permanent income cell; in keeping with our assumption that all singles differ only in their state variables.⁹

⁹This assumption appears to hold remarkably well in the data, in that conditional on our state variables, the savings and Medicaid recipients of recently singles are very similar to those of the people we were already single in the past. This is perhaps not surprising given that the vast majority of single people in our sample were married at some point.

3. Medicaid reciprocity by PI-cohort-year for everyone who is currently alive.

The cells for net worth are computed as follows.¹⁰ Household i has family structure f_t ; which indexes households that are currently couples and those that either are currently singles. We sort type- f households in cohort c by their permanent income levels, separating them into $Q = 5$ quintiles. Suppose that household i 's permanent income level falls in the q th permanent income interval of households in its cohort and family structure.

For Medicaid reciprocity, we construct cells for all individuals, whether they are married or not, but conditioning by PI and cohort and we match Medicaid reciprocity by cell.

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 a value of the state vector $(t, f_t, x_t, I, hs_t^h, hs_t^w, \zeta_t)$ drawn from the data distribution for 1996, and each is assigned the entire health and mortality history realized by the household in the AHEAD data with the same initial conditions. This way we generate attrition in our simulations that mimics precisely the attrition relationships in the data (including the relationship between initial wealth and mortality).

We discretize the asset grid and, using value function iteration, we solve the model numerically. This yields a set of decision rules, which, in combination with the simulated endowments and shocks, allows us to simulate each individual's assets, medical expenditures, health, and mortality. We compute assets from the artificial histories in the same way as we compute them from the real data. We use these profiles to construct moment conditions, and evaluate the match using our GMM criterion. We search over the parameter space for the values that minimize the criterion. Appendix F contains a detailed description of our moment conditions, the weighting matrix in our GMM criterion function, and the asymptotic distribution of our parameter estimates.

When estimating the life-cycle profiles, and subsequently fitting the model to those profiles, we face two well-known problems. First, in a cross-section, older households were born in an earlier year than younger households and

¹⁰Simulated agents are endowed with asset levels drawn from the 1996 data distribution, and thus we only match asset data 1998-2010. As was done when constructing the figures from the HRS data, we drop simulated cells with fewer than 10 observations from the moment conditions.

thus have different lifetime incomes. Because lifetime incomes of households in older cohorts will likely be lower than the lifetime incomes of younger cohorts, the asset levels of households in older cohorts will likely be lower also. Therefore, comparing older households born in earlier years to younger households in later years leads to understate asset growth. Second, households with lower income and wealth tend to die at younger ages than richer households. Therefore, the average survivor in a cohort has higher lifetime income than the average deceased member of the cohort. As a result, “mortality bias” leads the econometrician to overstate the average lifetime income of members of a cohort. This bias is more severe at older ages, when a greater share of the cohort members are dead. Therefore, “mortality bias” leads to overstate asset growth.

We use panel data to overcome these first two problems. Because we are tracking the same households over time, we are obviously tracking members of the same cohort over time. Similarly, we do separate sets of simulations for each cohort, so that the (initial) wealth and income endowments behind the simulated profiles are consistent with the endowments behind the empirical profiles. As for the second problem, we explicitly simulate demographic transitions so that the simulated profiles incorporate mortality effects in the same way as the data, both for couples and singles.

6 Estimation results

Section 3 includes results from our first-step estimation, that we use as inputs for our structural model, and the outputs that we require our model to match. In this section, we report our second-step parameter estimates for our structural model, the model fit, and we discuss the model’s identification.

6.1 Second-step results

7 Conclusions

Over one-third of total wealth in the United States is held by households over age 65. This wealth is an important determinant of their consumption and welfare. As the U.S. population continues to age, the elderly’s savings will only grow in importance.

β : discount factor, biennial	0.7436 (0.0494)
ν : RRA coefficient	2.936 (0.2315)
$\delta(hs^g = 1)$: preference shifter, nursing home	0.3949 (0.1082)
$\delta(hs^g = 2)$: preference shifter, bad health	0.5272 (0.1110)
ω : weight on the surviving spouse	1.025 (0.2246)
ϕ_0 : bequest intensity, single	12.10 (2.227)
κ_0 : bequest curvature, single (in 000s)	14,020 (1,489)
ϕ_1 : bequest intensity, surviving spouse	39.34 (11.48)
κ_1 : bequest curvature, surviving spouse (in 000s)	183,000 (45,750)
ϕ_2 : bequest intensity, both spouses	0.4190 (0.7187)
κ_2 : bequest curvature, both spouses (in 000s)	7,115 (20,540)
$c_{min}(f = 2)$: consumption floor, couples	6,089 (NA)
$c_{min}(f = 1)$: consumption floor, singles	3,672 (345.3)

Notes: Standard errors in parentheses. NA refers to the fixed ratio of couples' and singles' consumption floors

Table 3: Estimated parameters from our second-step estimation

Retired U.S. households, especially those who are part of a couple and have high income, decumulate their assets at a slow rate and often die with large amounts of assets, raising the questions of what drives their savings behavior and how their savings would respond to policy reforms.

We develop a model of optimal lifetime decision making and estimate key properties of the model. We find that singles live less long than people who are part of a couple, but are more likely to end up in a nursing home in any given year. For that reason, singles also have higher medical spending, per person, than people who are part of a couple. We also find that assets drop sharply with the death of a spouse. By the time the second spouse dies, a large fraction of the wealth of the original couple has vanished, with the wealth falls at the time of death of each spouse explaining most of the decline. A large share of these drops in assets is explained by the high medical expenses at the time of death. This suggests that a large fraction of all assets held in retirement are used to insure oneself against the risk of high medical and death expenses.

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, and Matthew D. Shapiro. Long-term care utility and late in life saving. Mimeo, 2017.
- [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] Orazio P. Attanasio and Hilary Williamson Hoynes. Differential mortality and wealth accumulation. *Journal of Human Resources*, 35(1):1–29, 2000.
- [8] Gary S. Becker and Nigel Tomes. Human capital and the rise and fall of families. *Journal of Labor Economics*, 4(3):1–39, 1986.
- [9] B. Douglas Bernheim, Andrei Shleifer, and Lawrence H. Summers. The strategic bequest motive. *Journal of Political Economy*, 93(5):1045–1076, 1985.
- [10] David M. Blau and Donna Gilleskie. The role of retiree health insurance in the employment behavior of older men. *International Economic Review*, 49(2):475–514, 2008.
- [11] 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.
- [12] Meta Brown. End-of-life transfers and the decision to care for a parent. mimeo, 2003.
- [13] Moshe Buchinsky. Recent advances in quantile regression models: A practical guideline for empirical research. *Journal of Human Resources*, 33:88–126, 1998.
- [14] Christopher D. Carroll. Why do the rich save so much? NBER Working Paper 6549, 1998.
- [15] Maria Casanova. Happy together: A structural model of couples’ retirement choices”. Mimeo, 2011.
- [16] Ana Castañeda, Javier Díaz-Giménez, and José-Víctor Ríos-Rull. Accounting for u.s. earnings and wealth inequality. *Journal of Political Economy*, 111(4):818–857, 2003.

- [17] 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.
- [18] Mariacristina De Nardi. Wealth inequality and intergenerational links. *Review of Economic Studies*, 71(3):743–768, 2004.
- [19] Angus Deaton. Saving and liquidity constraints. *Econometrica*, 59(5):1221–1248, 1991.
- [20] Angus Deaton and Christina Paxson. Mortality, education, income, and inequality among american cohorts. In David A. Wise, editor, *Themes in the Economics of Aging*, pages 129–165. NBER Conference Report Series. Chicago and London: University of Chicago Press, 2001.
- [21] Javier Díaz-Giménez, Vincenzo Quadrini, and José-Víctor Ríos-Rull. Dimensions of inequality: Facts on the U.S. distributions of earnings, income and wealth. *Federal Reserve Bank of Minneapolis Quarterly Review*, 21(2):3–21, 1997.
- [22] Mariacristina De Nardi, Eric French, and John B. Jones. Life expectancy and old age savings. *American Economic Review: Papers and Proceedings*, 99(2):110–115, 2009.
- [23] Mariacristina De Nardi, Eric French, and John B. Jones. Why do the elderly save? the role of medical expenses. *Journal of Political Economy*, 118(1):39–75, 2010.
- [24] Mariacristina De Nardi, Eric French, and John B. Jones. Medicaid insurance in old age. *American Economic Review*, 106(11):3480–3520, 2016.
- [25] Mariacristina De Nardi, Eric French, John B. Jones, and Jeremy McCauley. Medical spending of the u.s. elderly. *Fiscal Studies*, 37(3-4):717–747, 2016.
- [26] Darrell Duffie and Kenneth J. Singleton. Simulated moments estimation of markov models of asset prices. *Econometrica*, 61(4):929–952, 1993.

- [27] 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.
- [28] Karen E. Dynan, Jonathan Skinner, and Stephen P. Zeldes. Do the rich save more? *Journal of Political Economy*, 112(2):397–444, 2004.
- [29] Daniel Feenberg and Jonathan Skinner. The risk and duration of catastrophic health care expenditures. *Review of Economics and Statistics*, 76:633–647, 1994.
- [30] Eric French, Mariacristina De Nardi, John B. Jones, Olyesa Baker, and Phil Doctor. Right before the end: Asset decumulation at the end of life. *Economic Perspectives*, 30(3):2–13, 2006.
- [31] Eric French and John Bailey Jones. On the distribution and dynamics of health care costs. *Journal of Applied Econometrics*, 19(4):705–721, 2004.
- [32] Eric French and John Bailey Jones. The effects of health insurance and self-insurance on retirement behavior. Michigan Retirement Research Center Working Paper 2007-170, 2007.
- [33] William G. Gale and Maria G. Perozek. Do estate taxes reduce saving. In John Shoven, editor, *Rethinking Estate and Gift Taxation*. Brookings Institution Press, 2001.
- [34] Giovanni Gallipoli, Jutong Pan, and Sarah Turner. Household responses to individual shocks: Disability and labour supply”. Mimeo, 2013.
- [35] R. Glenn Hubbard, Jonathan Skinner, and Stephen P. Zeldes. Expanding the life-cycle model: Precautionary saving and public policy. *American Economic Review*, 84:174–179, 1994.
- [36] R. Glenn Hubbard, Jonathan Skinner, and Stephen P. Zeldes. Precautionary saving and social insurance. *Journal of Political Economy*, 103(2):360–399, 1995.
- [37] Michael D. Hurd. Mortality risk and bequests. *Econometrica*, 57(4):779–813, 1989.

- [38] Michael D. Hurd. Mortality risk and consumption by couples. Working Paper 7048, National Bureau of Economic Research, 1999.
- [39] John Laitner and Dmitriy Stoloyarov. Health status and post-retirement wealth decumulation. Mimeo, 2013.
- [40] Lee Lockwood. Incidental bequests: Bequest motives and the choice to self-insure late-life risks. Mimeo, 2018.
- [41] 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.
- [42] Maurizio Mazzocco. Household intertemporal behavior: A collective characterization and empirical tests. Mimeo, 2003.
- [43] Whitney K. Newey. Generalized method of moments specification testing. *Journal of Econometrics*, 29(3):229–256, 1985.
- [44] 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.
- [45] Ariel Pakes and David Pollard. Simulation and the asymptotics of optimization estimators. *Econometrica*, 57(5):1027–1057, 1989.
- [46] Michael G. Palumbo. Uncertain medical expenses and precautionary saving near the end of the life cycle. *Review of Economic Studies*, 66:395–421, 1999.
- [47] 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.
- [48] 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.
- [49] James Powell. Estimation of semiparametric models. In Robert Engle and Daniel L. McFadden, editors, *Handbook of Econometrics, Volume 4*. Elsevier, Amsterdam, 1994.

- [50] 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.
- [51] 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

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. 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 begin with 6047 households. After dropping 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, 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.

A key advantage of the AHEAD relative to other datasets is that it provides panel data on health status, including nursing home stays. We assign individuals a health status of “good” if self-reported health is excellent, very good or good, and are assigned 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 (or 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.

We do not include pension and Social Security wealth for four reasons. First, we wish to maintain comparability with other studies (Hurd [37], Attanasio and Hoynes [7] for example). Second, because it is illegal to borrow against Social security wealth pension and difficult to borrow against most forms of pension wealth, Social Security and pension wealth are much more illiquid than other assets. Third, their tax treatment is different from other assets. Finally, differences in Social Security and pension are captured in our model by differences in the permanent income measure we use to predict annual income.

One important problem with our asset data is that the wealthy tend

to underreport their wealth in virtually all household surveys (Davies and Shorrocks [17]). This will lead us to understate asset levels at all ages. However, Juster et al. (1999) 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 medians (conditional on permanent income), underreporting at the very top of the wealth distribution should not seriously affect our results.

Appendix B: Inferring Permanent Income

We assume that log income follows the process

$$\ln y_{it} = \kappa_1(t, f_{it}) + h(I_i) + \omega_{it} \quad (20)$$

where $\kappa_1(t, f_{it})$ is a flexible functional form of age t and family structure f_{it} (i.e., couple, single men, or single woman) and ω_{it} represents measurement error. The variable I_i is the household's percentile rank in the permanent income 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. Our procedure to estimate equation (20) is to first estimate the fixed effects model

$$\ln y_{it} = \kappa_1(t, f_{it}) + \alpha_i + \omega_{it} \quad (21)$$

which allows us to obtain a consistent estimate of the function $\kappa_1(t, f_{it})$. Next, note that as the number of time periods over which we can measure income and other variables for individual i (denoted T_i) becomes large,

$$plim_{T_i \rightarrow \infty} \frac{1}{T_i} \sum_{t=1}^{T_i} [\ln y_{it} - \kappa_1(t, f_{it}) - \omega_{it}] = \frac{1}{T_i} \sum_{t=1}^{T_i} [\ln y_{it} - \kappa_1(t, f_{it})] = h(I_i). \quad (22)$$

Thus we calculate the percentile ranking of permanent income I_i for every household in our sample by taking the percentile ranking of $\frac{1}{T_i} \sum_{t=1}^{T_i} [\ln y_{it} - \hat{\kappa}_1(t, f_{it})]$, where $\hat{\kappa}_1(t, f_{it})$ is the estimated value of $\kappa_1(t, f_{it})$ 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 . This gives us a measure of the percentile ranking of permanent income I_i . However, we also need to estimate the function $h(I_i)$, which gives a mapping

from 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

$$[\ln y_{it} - \hat{\kappa}_1(t, f_{it})] = h(I_i) + \omega_{it} \quad (23)$$

where the function $h(I_i)$ is a flexible functional form.

In practice we model $\kappa_1(t, f_{it})$ as a third order polynomial in age, dummies for family structure, and family structure interacted with an age trend. 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 $\kappa_1(\cdot, \cdot), h(\cdot)$, we can calculate the predicted value of $\ln \hat{y}_{it} = \hat{\kappa}_1(t, f_{it}) + \hat{h}(\hat{I}_i)$. 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: Imputing Medicaid plus Out of Pocket Medical Expenses

Our goal is to measure the data generating process of the sum of Medicaid payments plus out of pocket expenses: this is the variable $\ln(m_{it})$ in equation (6) of the main text. If the household is drawing Medicaid benefits, then the household will spend less than $\ln(m_{it})$ on out of pocket medical spending (Medicaid picking up the remainder).

The AHEAD data contains 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 Medicaid payments plus 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 have information on the medical spending or health of the spouse. Here we explain how to exploit the best of both datasets.

We use a three step estimation procedure. First, we use the MCBS to infer Medicaid payments for recipients, conditional on the observable variables that exist in both datasets. Second, we impute Medicaid payments in the AHEAD data using conditional mean matching procedure (which is a procedure very similar to hotdecking). Third, we estimate the data generating process for out of pocket plus Medicaid medical spending in the HRS.

First Step Estimation Procedure:

We use the MCBS to infer Medicaid payments for recipients, conditional on the observable variables that exist in both the MCBS and the HRS

datasets. Define oop_{it} as the out of pocket medical expenses and m_{it} is the sum of out of pocket plus Medicaid payments that we wish to plug in the model, and so Med_{it} is the dollar value of Medicaid payments.

To impute Med_{it} in (which is missing in the HRS), we follow David et al. (1986) and French and Jones (2011) and use the following predictive mean matching regression approach. First, for every member of the the MCBS sample with a positive Medicaid indicator (i.e., a Medicaid recipient), we regress the variable of interest Med on the vector of observable variables z , yielding $Med = z\beta + \varepsilon$. Second, for each individual j in the MCBS we calculate the predicted value $\hat{Med}_{jt} = z_{jt}\hat{\beta}$, and for each member of the sample we calculate the residual $\hat{\varepsilon}_{jt} = Med_{jt} - \hat{Med}_{jt}$. Third, we sort the predicted value \hat{Med}_{jt} into deciles and keep track of all values of $\hat{\varepsilon}_{jt}$ within each decile.

In practice we include in z_{jt} nursing home status, the 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 HRS is medical spending over two years, we take two year averages of the MCBS data to be consistent with the structure of the HRS. A regression of Med_{jt} on z_{jt} yields a R^2 statistic of .67, suggesting that our predictions are accurate.

Second Step Estimation Procedure:

Next, for every individual i in the HRS sample with a positive Medicaid indicator, we impute $\hat{Med}_{it} = z_{it}\hat{\beta}$ using the values of $\hat{\beta}$ estimated using the MCBS. Then we impute ε_{it} for each member of the HRS sample by finding a random individual j in the MCBS with a value of \hat{Med}_{jt} in the same decile as \hat{Med}_{it} in the HRS, and set $\varepsilon_{it} = \hat{\varepsilon}_{jt}$. The imputed value of Med_{it} is $\hat{Med}_{it} + \varepsilon_{it}$.

As David et al. (1986) 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 ”. We find that adding extra variables are very important for improving goodness of fit when imputing medicaid payments. Even a small number of variables generate a large number of hot-decking cells, as hot-decking uses a full set of interactions. Thus, in this context, hot decking is too data

intensive.

We predict Medicaid expenditures for 3,756 Medicaid eligible households (i.e., Medicaid beneficiaries) in the HRS and report the results of this imputation exercise in Table 4. 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 expenditures at the household level are larger than for either single men or women. The imputed Medicaid payments for both spouses 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 4 new widows and widowers who’s spouse has died between waves of the HRS 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 household’s Medicaid expenditures.

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 4: Imputed Medicaid payments for Medicaid beneficiaries in HRS.

The distribution of imputed Medicaid expenditures in the HRS is close to the distribution observed in the MCBS. Mean imputed Medicaid expenditures, Med_{it} , for Medicaid recipients in the HRS is \$14,050. This is lower than the corresponding value of \$16,XXXK in the MCBS. This small difference is due to the distribution of observable variables, z , in both the MCBS and HRS. In particular, the number of nights in a nursing home is 20X% lower in the HRS than in the MCBS, due in part to the fact that the initial HRS sample is for the non-institutionalized population who were not in a nursing home.

Third Step Estimation Procedure:

Our value of m_{it} in the HRS is thus $oop_{it} + Med_{it}$, where Med_{it} was

calculated in the second step. Define $X_{it} = (hs_t^h, hs_t^w, I, g, t, f_t, f_{t-1})$. Recall that equation (6) of the main text is given by $\ln(m_{it}) = m(X_{it}) + \sigma(X_{it}) \times \psi_t$.

Using HRS, we use fixed effects estimation procedures to estimate

$$\ln(m_{it}) = \alpha_i + m(X_{it}) + \epsilon_{it}. \quad (24)$$

which gives $\hat{m}(X_{it})$. Next we construct the estimated residuals $(\widehat{\epsilon_{it} + \alpha_i}) = \ln(m_{it}) - \hat{m}(X_{it})$ and estimate the regression (using OLS and no fixed effects):

$$(\widehat{\epsilon_{it} + \alpha_i})^2 = h(X_{it}) + \zeta_{it}. \quad (25)$$

This allows us to recover the variance of ϵ_{it} conditional on X_{it} .

In order to infer the variance of medical expenditures conditional on X_{it} , note that from equations (6) and (??)

$$\sigma(X_{it}) \times \psi_{it} = \ln(m_{it}) - m(X_{it}) = (\epsilon_{it} + \alpha_i). \quad (26)$$

so we can obtain the conditional variance by noting that $E[\psi_{it}^2] = 1$ and so

$$[\sigma(X_{it})]^2 = E[(\epsilon_{it} + \alpha_i)^2 | X_{it}] \quad (27)$$

where we estimate $E[(\epsilon_{it} + \alpha_i)^2 | X_{it}] = h(X_{it})$ in equation (25).

Appendix D: Other savings graphs




Figure 8: Median assets for households initially single. Each line represents median assets for a cohort-income cell, traced over the time period 1996-2010. Thicker lines refer to higher permanent income groups. Solid lines: cohorts ages 72-76 and 82-86 in 1996. Dashed lines: ages 82-86 and 92-96 in 1996.

Figure 8 displays the asset profiles for households that were single when first observed in the AHEAD, classified according. Median assets are increasing in permanent income, with the 74-year-olds in the highest PI income of the singles holding about \$240,000 in median assets, while those at the lowest PI quintiles holding essentially no assets. Over time, those with the highest PI tend to hold onto significant wealth well into their nineties, those with the lower PIs never save much, while those in the middle PIs display quite some asset decumulation as they age.

Appendix E: Outline of the computation of the value functions and optimal decision rules

We compute the value functions by backward induction. We start from the singles, find their time T value function and decision rules by maximizing equation (17) subject to the relevant constraints, and $V_{T+1}^g = \theta_0(x_t - c_t)$, $g = h, w$. This yields the value function V_T^g and the decision rules for time T . We then find the decision rules at time $T - 1$ by solving equation (17) with V_T^g . Continuing this backward induction yields decision rules for time $T - 2, T - 3, \dots, 1$.

We find the decisions for couples by maximizing equation (19), subject to the relevant constraints and the value function for the singles, and setting V_{T+1}^c to the appropriate bequest motive value. 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 $V_T^c, V_T^g, g = h, w$. Continuing this backward induction yields decision rules for time $T - 2, T - 3, \dots, 1$.

We discretize the persistent component and the transitory components of the health shock, and interest rate into Markov Chain following Tauchen and Hussey (1991). We assume a finite number of permanent income categories. We take cash-on-hand to lay into a finite number of grid points.

Given each level of permanent income of the household, we solve for decision rules for each possible combination of cash-on-hand, income, health status, and persistent component of the health shock. We use linear interpolation within the grid and linear extrapolation outside of the grid to evaluate the value function at points that we do not directly compute.

For the singles, for simplicity of notation, here for the most part we drop any reference to the missing spousal variables everywhere, and just use similar notation to the one for couples (except for the number of arguments in the function). In the code, we have different names, say for example for the survival of single people and couples.

The value function for the singles, $g = h, w$ and $f_t = 0$ is given by:

$$V_t^g(x_t, hs_t, I, \zeta_t) = \max_{c_t, x_{t+1}} \left\{ u(c_t, hs_t) + \beta \left(1 - s(hs_t, I, g, t, 0) \right) \theta_0(x_t - c_t) + \right. \\ \left. \beta s(hs_t, I, g, t, 0) \left[\sum_{k=1}^{d_m} \sum_{l=1}^{d_\zeta} \sum_{n=1}^{d_\xi} Pr(hs_{t+1} = hs_k | hs_t, I, g, t, 0) Pr(\zeta_{t+1} = \zeta_l | \zeta_t) Pr(\xi_{t+1} = \xi_n) \right. \right. \\ \left. \left. V_{t+1}^g \left(x_{t+1}(k, l, n), hs_{t+1}(k), I, \zeta_{t+1}(l) \right) \right] \right\}.$$

Subject to:

$$z_{t+1} = x_t - c_t + y \left(r(x_t - c_t) + y_{t+1}(I, 0), \tau \right)$$

$$x_t > c_{\min}(g), \quad c_t \leq x_t, \quad \forall t,$$

$$\ln(m_{t+1}(k, l, n)) = m^g(hs_{t+1}(k), I, t + 1) + \sigma^g(hs_{t+1}(k), I, t + 1) \psi_{t+1}(l, n)$$

$$\psi_{t+1}(l, n) = \zeta_{t+1}(l) + \xi_{t+1}(n),$$

$$s_{t+1}(k, l, n) = z_{t+1} - m_{t+1}(k, l, n)$$

$$tr_{t+1}(k, l, n) = \max \left\{ 0, c_{\min}(f_t) - \max(0, s_{t+1} - a_d(i)) \right\}$$

$$x_{t+1}(k, l, n) = s_{t+1}(k, l, n) + tr_{t+1}(k, l, n)$$

$\forall x_t$ COH level, determine maximum consumption (and hence savings)

$\forall c_t \in (c_{\min}, x_t)$, compute $u^g(c, m)$ and $\theta_0(x_t - c_t)$.

$\forall (g, t, I)$ For each gender, age, PI, compute savings = z_{t+1}

$\forall hs_t, \zeta_t$ For each health state and pers. medex shock TODAY

$\forall hs_{t+1}(k), \zeta_{t+1}(l), \xi(n)$ tomorrow's shocks

$x_{t+1}(k, l, n)$ compute tomorrow's COH for each state

Interpolate and extrapolate $V_{t+1}^g(x_{t+1}(k, l, n), hs_{t+1}(k), I, \zeta_{t+1}(l))$

Compute

$$\beta s(hs_t, I, t) \sum_{k=1}^{d_m} \sum_{l=1}^{d_\zeta} \sum_{n=1}^{d_\xi} \Pi(k, l, n) V_{t+1}^g(x_{t+1}(k, l, n), hs_{t+1}(k), I, \zeta_{t+1}(l))$$

$$\beta (1 - s(hs_t, I, t)) \phi_0(x_t - c_t) + u^g(c_t, hs_t)$$

$W(c_t, x_t, I, hs_t, \zeta_t)$ = sum of the two lines just above

$\max_{c_t} W(c_t, x_t, I, hs_t, \zeta_t)$

$c_t^g(x_t, I, hs_t, \zeta_t)$ = maximizer

$V_t^g(x_t, I, hs_t, \zeta_t)$ = maximum

The value function for the couples is given by:

$$\begin{aligned}
V_t^c(x_t, hs_t^w, hs_t^h, I, \zeta_t) &= \max_{c_t, x_{t+1}, b_t^w, b_t^h} \left\{ u^c(c_t, hs_t^h, hs_t^w) + \right. \\
&+ \beta \left(1 - s(hs^w, I, w, t, 1) \right) \left(1 - s(hs^h, I, h, t, 1) \right) \theta_2(x_t - c_t) + \\
&\beta s(hs^w, I, w, t, 1) (1 - s(hs^h, I, h, t, 1)) \left(\theta_1(b_t^h) + \omega \sum_{k=1}^{d_m} \sum_{l=1}^{d_\zeta} \sum_{n=1}^{d_\xi} Pr(\zeta_{t+1} = \zeta_l | \zeta_t) \right. \\
&Pr(\xi_{t+1} = \xi_n) Pr(hs_{t+1}^w = hs_k | hs_t^w, I, w, t, 1) V_{t+1}^w(x_{t+1}^w(k, l, n), hs_{t+1}^w(k), I, \zeta_{t+1}(l)) \left. \right) + \\
&\beta (1 - s(hs^w, I, w, t, 1)) s(hs^h, I, h, t, 1) \left(\theta_1(b_t^w) + \omega \sum_{k=1}^{d_m} \sum_{l=1}^{d_\zeta} \sum_{n=1}^{d_\xi} Pr(\zeta_{t+1} = \zeta_l | \zeta_t) \right. \\
&Pr(\xi_{t+1} = \xi_n) Pr(hs_{t+1}^h = hs_{kh} | hs_t^h, I, h, t, 1) V_{t+1}^h(x_{t+1}^h(k, l, n), hs_{t+1}^h(k), I, \zeta_{t+1}(l)) \left. \right) + \\
&\beta s(hs^w, I, w, t, 1) s(hs^h, I, h, t, 1) \\
&\left(\sum_{kh=1}^{d_m} \sum_{kw=1}^{d_m} \sum_{l=1}^{d_\zeta} \sum_{n=1}^{d_\xi} V_{t+1}^c(x_{t+1}(k_h, k_w, l, n), hs_{t+1}^w, hs_{t+1}^h, I, \zeta_{t+1}(l)) \right. \\
&Pr(\zeta_{t+1} = \zeta_l | \zeta_t) Pr(\xi_{t+1} = \xi_n) Pr(hs_{t+1}^h = hs_{kh} | hs_t^h, I, h, t, 1) Pr(hs_{t+1}^w = hs_{kw} | hs_t^w, I, w, t, 1) \left. \right) \left. \right\} \\
\end{aligned} \tag{28}$$

subject to

$$\begin{aligned}
x_t &> c_{\min}(c), \quad c_t \leq x_t, \quad \forall t \\
0 &< b_t \leq x_t - c_t, \quad \forall t, \quad i = h, w, \\
y_{t+1}^c &= y(c, I, t + 1) \\
y_{t+1}^i &= y(i, I, t + 1) \\
z_{t+1}^c &= \{x_t - c_t + y(r(x_t - c_t) + y_{t+1}^c, \tau)\} \\
z_{t+1}^i &= \{x_t - c_t + y(r(x_t - c_t) + y_{t+1}^i, \tau)\}
\end{aligned}$$

$$\begin{aligned}
\ln m_{t+1}^c(k_h, k_w, l, n) &= m^c(hs_{t+1}^h(k_h), hs_{t+1}^w(k_w), t + 1, I) \\
&+ \sigma^c(hs_{t+1}^h(k_h), hs_{t+1}^w(k_w), I, t + 1) \psi_{t+1}(l, n)
\end{aligned}$$

$$\ln m_{t+1}^g(k, l, n) = m^g(g, hs_{t+1}^i(k), t+1, I) + \sigma^g(g, hs_{t+1}^g(k), I, t+1)\psi_{t+1}(l, n)$$

$$\psi_{t+1}(l, n) = \zeta_{t+1}(l) + \xi_{t+1}(n),$$

$$x_{t+1}^c(k_h, k_w, l, n) = \max\{z_{t+1}^c - m_{t+1}^c(k_h, k_w, l, n), c_{\min}(c)\}$$

$$x_{t+1}^g(j, k, l, n) = \max\{z_{t+1}^i - m_{t+1}^i(k, l, n) - b_t, c_{\min}(i)\}$$

Computation

$\forall(x_t)$

$\forall c_t \in (c_{\min}, x_t)$

$\theta_2(x_t - c_t)$

$\forall(b_t^i \in (0, x_t))$

$\theta_1(b_t)$

$\forall(m_t^h, m_t^w)$

$u(c_t, m_t^h, m_t^w)$

$\forall t, \zeta_t$

sumC: expected value next period in case both survive

sumH: expected value next period in case husband dies

sumW: expected value next period in case wife dies

bequestUCM: utility from bequests when both die

$$W(c_t, x_t, I, m_t^h, m_t^w, \zeta_t) = u(c_t, m_t^h, m_t^w) + \beta s_{m,I,t}^w s_{m,I,t}^h \text{sumC}$$

$$+ \beta s_{m,I,t}^w (1 - s_{m,I,t}^h) \text{sumH}$$

$$+ \beta (1 - s_{m,I,t}^w) s_{m,I,t}^h \text{sumW} + \beta (1 - s_{m,I,t}^w) (1 - s_{m,I,t}^h) +$$

$$\beta (1 - s_{m,I,t}^w) (1 - s_{m,I,t}^h) \text{bequestUCM}$$

$$\max_{c_t} W(c_t, x_t, I, m_t^h, m_t^w, \zeta_t)$$

$$c_t^c(x_t, I, m_t^h, m_t^w, \zeta_t) = \text{maximizer}$$

$$V_t^c(x_t, I, m_t^h, m_t^w, \zeta_t) = \text{maximum}$$

In the calculation of next period's expected value in the case in which one spouse dies (sumH and sumW), we compute the optimal decision rule for bequests, conditional on each consumption choice. In the end when we get the decision rule for consumption, we automatically get the optimal decision rule of bequest corresponding to each state $b_t^c(x_t, I, m_t^h, m_t^w, \zeta_t)$

Appendix F: Moment Conditions and the Asymptotic Distribution of Parameter Estimates

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 assets found in the data and the simulated profiles generated by the model. For each calendar year $t \in \{1, \dots, T\}$, we match median assets for 5 permanent income quintiles in 5 birth year cohorts, both for singles and couples, leading to a total of $50T$ moment conditions. Sorting households into quintiles also requires us to estimate the 1/5th, 2/5th, 3/5th and 4/5th permanent income quantiles for each birth year cohort and initial family type. This produces a total of 40 nuisance parameters, which we collect into the vector γ . Each of these parameters has its own moment condition.

The way in which we construct these moment conditions builds on the approach described in French and Jones [32]. Useful references include Buchinsky [13] and Powell [49]. Consider first the permanent income quantiles. Let $q \in \{1, 2, \dots, Q - 1\}$ index the quantiles. Assuming that the permanent income distribution is continuous, the π_q -th quantile of permanent income for initial family type f of cohort c , $g_{\pi_q}(c, f)$, is defined as

$$\Pr(I_i \leq g_{\pi_q}(c, f) | c, f) = \pi_q. \quad (29)$$

In other words, the fraction of households with less than g_{π_j} in permanent income is π_j . Using the indicator function, the definition of π_j -th conditional quantile can be rewritten as

$$E([1\{I_i \leq g_{\pi_q}(c, f)\} - \pi_j] \times 1\{c_i = c\} \times 1\{f_i = f\}) = 0, \quad (30)$$

for $c \in \{1, 2, \dots, C\}$, $f \in \{\text{single}, \text{couple}\}$, $q \in \{1, 2, \dots, Q - 1\}$.

The more important set of moment conditions involves the permanent income-conditional age-asset profiles. Suppose that household i 's permanent income level falls in the q th permanent income interval of households in its cohort, i.e.,

$$g_{\pi_{q-1}}(c, f) \leq I_i \leq g_{\pi_q}(c, f). \quad (31)$$

We assume that $\pi_0 = 0$ and $\pi_Q = 1$, so that $g_{\pi_0}(c, f) \equiv -\infty$ and $g_{\pi_Q}(c, f) \equiv \infty$. Let $a_{cft}(\Delta, \chi)$ be the model-predicted median observed asset level for group cft . Recall that the median is just the 1/2 quantile. Assuming that observed assets have a continuous conditional density, we arrive at the following moment condition:

$$E(1\{\tilde{a}_{it} \leq a_{cft}(\Delta_0, \chi_0)\} - 1/2 | c, f, q, t, \text{household observed at } t) = 0. \quad (32)$$

Equation (32) is merely equation (??) in the main text, adjusted to allow for “missing” as well as deceased households, as in French and Jones [31]. Using indicator function notation, we can convert this conditional moment equation into an unconditional one:

$$E\left([1\{\tilde{a}_{it} \leq a_{cft}(\Delta_0, \chi_0)\} - 1/2] \times 1\{c_i = c\} \times 1\{f_i = f\} \times 1\{g_{\pi_{q-1}}(c, f) \leq I_i \leq g_{\pi_q}(c, f)\} \times 1\{\text{household observed at } t\} \mid t\right) = 0, \quad (33)$$

for $c \in \{1, 2, \dots, C\}$, $f \in \{\text{single, couple}\}$, $q \in \{1, 2, \dots, Q\}$, $t \in \{t_1, t_2, \dots, t_T\}$.

Suppose we have a data set of I independent households that are each observed at T separate calendar years. Let $\varphi(\Delta, \gamma; \chi_0)$ denote the $(50T + 40)$ -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 $(50T + 40) \times (50T + 40)$ weighting matrix, the MSM estimator $(\hat{\Delta}', \hat{\gamma}')'$ is given by

$$\arg \min_{\{\Delta, \gamma\}} \frac{I}{1 + \tau} \hat{\varphi}_I(\Delta, \gamma; \chi_0)' \widehat{\mathbf{W}}_I \hat{\varphi}_I(\Delta, \gamma; \chi_0),$$

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

In practice, we estimate χ_0 as well, 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 regularity conditions stated in Pakes and Pollard [45] and Duffie and Singleton [26], the MSM estimator $\hat{\theta}$ is both consistent and asymptotically normally distributed:

$$\sqrt{I} \left(\begin{pmatrix} \hat{\Delta} \\ \hat{\gamma} \end{pmatrix} - \begin{pmatrix} \Delta_0 \\ \gamma_0 \end{pmatrix} \right) \rightsquigarrow N(0, \mathbf{V}),$$

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;

$$\mathbf{D} = \begin{bmatrix} \frac{\partial \varphi(\Delta, \gamma; \chi_0)}{\partial \Delta'} \Big|_{\Delta=\Delta_0} & \frac{\partial \varphi(\Delta, \gamma; \chi_0)}{\partial \gamma'} \Big|_{\gamma=\gamma_0} \end{bmatrix} \quad (34)$$

is the $(50T + 40) \times (9 + 40)$ gradient matrix of the population moment vector; and $\mathbf{W} = \text{plim}_{T \rightarrow \infty} \{\widehat{\mathbf{W}}_I\}$. Moreover, Newey [43] shows that if the model is properly specified,

$$\frac{I}{1 + \tau} \hat{\varphi}_I(\hat{\Delta}, \hat{\gamma}; \chi_0)' \mathbf{R}^{-1} \hat{\varphi}_I(\hat{\Delta}, \hat{\gamma}; \chi_0) \rightsquigarrow \chi_{50T-9}^2,$$

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

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

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} . But even though the optimal weighting matrix is asymptotically efficient, it can be severely biased in small samples. (See, for example, Altonji and Segal [2].) We thus use a “diagonal” weighting matrix, as suggested by Pischke [47]. The 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 $(50T + 40) \times (50T + 40)$ estimated variance-covariance matrix of the sample data. When estimating preferences, we use sample statistics, so that $a_{cftqt}(\Delta, \chi)$ is replaced with the sample median for group $cftqt$. When computing the chi-square statistic and the standard errors, we use model predictions, so that the sample medians for group $cftqt$ is replaced with its simulated counterpart, $a_{cftqt}(\hat{\Delta}, \hat{\chi})$.

One complication in estimating the gradient matrix \mathbf{D} is that the functions inside the moment condition $\varphi(\Delta, \gamma; \chi)$ are non-differentiable at certain data points; see equations (30) and (33). 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 [45], Newey and McFadden [44] (section 7) and Powell [49].

In finding \mathbf{D} , it proves useful to partition $\varphi(\cdot)$ into the $50T$ -element vector $\varphi_{\Delta}(\cdot)$, corresponding to the moment conditions described by equation (33) and the $40 - element$ vector $\varphi_{\gamma}(\cdot)$, corresponding to the moment conditions described by equation (30). Using this notation, we can rewrite equation

(34) as

$$\mathbf{D} = \begin{bmatrix} \frac{\partial \varphi_{\Delta}(\Delta_0, \gamma_0; \chi_0)}{\partial \Delta'} & \frac{\partial \varphi_{\Delta}(\Delta_0, \gamma_0; \chi_0)}{\partial \gamma'} \\ \frac{\partial \varphi_{\gamma}(\Delta_0, \gamma_0; \chi_0)}{\partial \Delta'} & \frac{\partial \varphi_{\gamma}(\Delta_0, \gamma_0; \chi_0)}{\partial \gamma'} \end{bmatrix}, \quad (35)$$

and proceed element-by-element.

It immediately follows from equation (30) that

$$\frac{\partial \varphi_{\gamma}(\Delta_0, \gamma_0; \chi_0)}{\partial \Delta'} = \mathbf{0}. \quad (36)$$

To find $\frac{\partial \varphi_{\gamma}(\Delta_0, \gamma_0; \chi_0)}{\partial \gamma'}$, we rewrite equation (30) as

$$\Pr(c_i = c \ \& \ f_i = f) \times [F(g_{\pi_q}(c, f)|c, f) - \pi_j] = 0, \quad (37)$$

where $F(g_{\pi_q}(c, f)|c, f)$ is the c.d.f. of permanent income for family type- f members of cohort c evaluated at the π_j -th quantile. Differentiating this equation shows that $\frac{\partial \varphi_{\gamma}(\Delta_0, \gamma_0; \chi_0)}{\partial \gamma'}$ is a diagonal matrix whose diagonal elements are given by

$$\Pr(c_i = c \ \& \ f_i = f) \times f(g_{\pi_q}(c, f)|c, f). \quad (38)$$

In practice we find $f(g_{\pi_q}(c, f)|c, f)$, the conditional p.d.f. of permanent income evaluated at the π_j -th quantile, with a kernel density estimator.

To find $\frac{\partial \varphi_{\Delta}(\Delta_0, \gamma_0; \chi_0)}{\partial \Delta'}$ and $\frac{\partial \varphi_{\Delta}(\Delta_0, \gamma_0; \chi_0)}{\partial \gamma'}$, it is helpful to rewrite equation (33) as

$$\Pr(c_i = c \ \& \ f_i = f \ \& \ \text{household observed at } t) \times \int_{g_{\pi_{q-1}}(c, f)}^{g_{\pi_q}(c, f)} \left[\int_{-\infty}^{a_{cfqt}(\Delta_0, \chi_0)} f(\tilde{a}_{it}|c, f, I_i, t) d\tilde{a}_{it} - \frac{1}{2} \right] f(I_i|c, f) dI_i = 0, \quad (39)$$

It follows that the rows of $\frac{\partial \varphi_{\Delta}(\Delta_0, \gamma_0; \chi_0)}{\partial \Delta'}$ are given by

$$\Pr(c_i = c \ \& \ f_i = f \ \& \ g_{\pi_{q-1}}(c, f) \leq I_i \leq g_{\pi_q}(c, f) \ \& \ \text{household observed at } t) \times f(a_{cfqt}|c, f, g_{\pi_{q-1}}(c, f) \leq I_i \leq g_{\pi_q}(c, f), t) \times \frac{\partial a_{cfqt}(\Delta_0, \gamma_0; \chi_0)}{\partial \Delta'}. \quad (40)$$

Proceeding similarly, it can be shown that each row of $\frac{\partial \varphi_{\Delta}(\Delta_0, \gamma_0; \chi_0)}{\partial \gamma'}$ has the following two non-zero elements:

$$\Pr(c_i = c \ \& \ f_i = f \ \& \ \text{household observed at } t) \times \begin{pmatrix} -f(g_{\pi_{q-1}}(c, f)|c, f) [F(a_{cfqt}|c, f, g_{\pi_{q-1}}(c, f), t) - 1/2] \\ f(g_{\pi_q}(c, f)|c, f) [F(a_{cfqt}|c, f, g_{\pi_q}(c, f), t) - 1/2] \end{pmatrix}', \quad (41)$$

with

$$f(g_{\pi_0}(c, f)|c, f) [F(a_{cft}|c, f, g_{\pi_0}(c, f), t) - 1/2] \equiv 0, \quad (42)$$

$$f(g_{\pi_Q}(c, f)|c, f) [F(a_{cft}|c, f, g_{\pi_Q}(c, f), t) - 1/2] \equiv 0. \quad (43)$$

In practice, we find $F(a_{cft}|c, f, I, t)$, the conditional c.d.f. of assets evaluated at the median a_{cft} by finding kernel estimates of the mean regression of $1\{a_{it} \leq a_{cft}\}$ on I (holding c, f and, t fixed).