Intrahousehold Inequality, Savings, and Long-Term Care in Old Age*

Minsu Chang† Ami Ko‡

December 7, 2025

Abstract

The distribution of resources between spouses and the extent to which household savings reflect each spouse's motives are fundamental to analyzing married retirees' welfare. This paper provides one of the first quantitative assessments of cross-sectional heterogeneity in gender inequality among retired couples. We develop and estimate a non-unitary model of couples in retirement. We find that women's bargaining positions have strengthened on average, pushing their mean consumption shares closer to parity, but this improvement has been accompanied by widening dispersion across households. Within couples, wives on average have greater incentives to save due to their stronger precautionary motives. Consequently, policies that encourage faster dissaving of retirement wealth disadvantage wives relative to husbands, revealing distributional consequences that would be obscured under a unitary framework.

*An early version of the paper was circulated under the title "Marital Transitions, Housing, and Long-Term Care in Old Age". We thank Bertrand Achou, Daniel Barczyk, Mariacristina De Nardi, Eric French, John Bailey Jones, Soojin Kim, Karen Kopecky, Siha Lee, Hamish Low, Rory McGee, Makoto Nakajima, Weilong Zhang, as well as seminar and conference participants at Western Ontario, McMaster, Johns Hopkins, UWisconsin-Madison, University of Cambridge, Richmond Fed, Barcelona GSE Summer Forum, Econometric Society Meeting, Dynamic Structural Econometrics, ASU Empirical Micro Conference, KBER-Macro at Seoul National University, Montreal Series on the Economics of Aging, SITE, Sogang, and KAIST for helpful comments. Financial support from Korea Bureau of Economic Reserach in the Institute of Economic Research of Seoul National University is gratefully acknowledged.

†Seoul National University, Department of Economics. minsuchang@snu.ac.kr.

‡Georgetown University, Department of Economics. ami.ko@georgetown.edu.

1 Introduction

Individual welfare in retirement depends critically on access to adequate savings and long-term care. For married couples, these choices are made jointly, so the balance of power within the household plays a central role in shaping each spouse's well-being. How resources are allocated between spouses, and the extent to which household savings reflect each spouse's motives, are fundamental to analyzing retirees' welfare. Old-age support policies can also generate differential welfare effects depending on each spouse's preferences and risks. For instance, means-tested programs that create spend-down incentives may disproportionately disadvantage the spouse with a stronger desire to save. Yet when outcomes are evaluated only at the household level, these distributional consequences are obscured.

Addressing these questions about intrahousehold inequality and individual welfare requires moving beyond a unitary framework, which treats the household as a single decision maker and masks gender differences in preferences and power. This paper develops and estimates a collective framework for studying couples' decisions in retirement. Instead of relying on a single household utility function, we specify each spouse's preferences separately and assume that decisions are made by maximizing a weighted sum of these preferences. The key input to this process is each spouse's Pareto weight, which reflects relative bargaining power and is shaped by features of the environment. This framework allows us to examine how within-household gender inequality in retirement varies across households, insights that unitary models cannot deliver. We identify variation in Pareto weights using income shares as a distribution factor, a variable that influences the bargaining process but not preferences or the budget set. We also exploit shifts in household structure, such as a spouse's death or nursing home entry, to identify the level of the Pareto weights. The paper's first main innovation is thus to empirically implement a collective model for analyzing married households' savings and long-term care decisions in retirement.

A second innovation deepens our intrahousehold analysis by treating housing as a distinctive asset in retirees' savings decisions. Although a large share of retirement wealth is held in housing, much of the existing literature treats all assets as liquid. Accounting for the unique features of housing is especially important when analyzing heterogeneous savings motives between spouses. Because housing assets are harder to dissave than liquid assets, the value of homeownership may depend on each spouse's longevity expectations, which vary by gender. Housing also plays a central role in old-age policies that exempt home equity from means tests, such as Medicaid. The welfare implications of these policies may differ across spouses, depending on who benefits relatively more from the exemption. By modeling

housing explicitly, we capture how asset composition interacts with intrahousehold dynamics and policy design in shaping individual welfare.

A third innovation lies in studying not only couples' but also singles' savings decisions in retirement, which is especially helpful in identifying gender differences in preferences. In our dynamic model of retirees facing mortality risk, married individuals transition into singlehood after the death of a spouse. It is therefore necessary to incorporate singles' behavior when solving the intertemporal problem for married households. In addition, we use singles' behavior to help identify gender differences in preferences. While couples' decisions reflect both individual preferences and relative bargaining power, singles' choices reveal preferences directly, free from intrahousehold dynamics. This allows us to more clearly disentangle gender-based heterogeneity in preferences.

Using data from the Health and Retirement Study (HRS), we begin by examining whether married retirees' consumption and savings behaviors are consistent with the unitary model and whether the patterns instead support a collective framework. Under the unitary model, conditional on total household resources, household choices should not vary with distribution factors such as individual income shares. In contrast to this prediction, we find that the wife's income share is negatively correlated with household consumption and positively correlated with household savings, conditional on other observables including total household resources. These patterns are inconsistent with the unitary model but are consistent with a collective framework in which the wife's bargaining power rises with her relative income share. This interpretation aligns with the idea that wives, on average, prefer higher savings because they expect to live longer and are more likely to outlive their husbands, limiting their access to spousal care in old age.

We next document a series of descriptive patterns that guide our structural modeling choices. Using the HRS, we show that couples and singles differ sharply in how they hold and decumulate housing assets. Homeownership remains high among couples well into advanced ages, while singles rapidly dissave housing. We document lack of spousal care and higher nursing home risk for singles as a potential reason for their lower homeownersihp. Institutional rules amplify these differences by marital status. Medicaid exempts the home for married nursing home residents but can recover costs from single beneficiaries' housing. Using state variations, we confirm that this asymmetry discourages homeownership among singles relative to couples, consistent with Greenhalgh-Stanley (2012).

We also show gender differences among single retirees that help identify preference heterogeneity underlying household behavior. Single women in the low- and mid-PI groups save

more than observably similar men, whereas single women in the high-PI group save significantly less than their male counterparts. This pattern suggests that savings motives differ by gender in ways that vary with income.

Building on the descriptive findings, we develop a dynamic model of retirees' savings, housing, and long-term care decisions under health and mortality risk. The household head is age 65 in the initial period and may be married or single. For married households, we use a collective framework in which husbands and wives have separate preferences, and the couple's choices maximize a weighted sum of their values. Preferences are defined over non-housing consumption and housing services with gender-specific weights, and incorporate (dis)utility from providing spousal care as well as bequest motives toward heirs and a surviving spouse. Retirees either keep their home or sell and become renters, with renting modeled as an absorbing state. Means-tested social insurance programs, representing Medicaid, are incorporated and treat housing differently by marital status: single nursing home residents must effectively use their housing wealth to finance care, whereas married households benefit from a homestead exemption when a spouse remains in the home.

For couples, Pareto weights depend on household permanent income (PI) and the wife's income share, allowing intrahousehold bargaining power to vary across households. The dependence on income follows the standard collective literature (e.g., Chiappori (1988), Chiappori (1992)), and the use of relative income as a distribution factor is consistent with a large body of work on working-age households, including Browning, Bourguignon, Chiappori, and Lechene (1994), Lise and Yamada (2019), and Lise and Seitz (2011). Because each spouse has their own preferences, the model delivers separate value functions for husbands and wives, which we use to assess policies from an individual, rather than solely household, perspective.

We estimate the model using indirect inference, selecting auxiliary models that capture the data features most informative for identifying the structural parameters. To identify the Pareto weights, we exploit variation in the wife's income share. Because income shares affect choices only through bargaining power, differences in choices across households with different shares identify how the weight responds to the distribution factor. To recover how bargaining power varies with PI, we exploit how the effect of income share on consumption changes across the PI distribution. Finally, we use behavior changes around nursing home entry as particularly useful data patterns about the level of the Pareto weights.

Beyond the Pareto weights, caregiving preferences are primarily shaped by observed spousal care rates; gender differences in housing choices among singles help discipline gender-specific housing preferences; the wealth trajectories of single retirees inform bequest motives toward

children; and the late-life savings behavior of married couples provides important information about altruism toward a surviving spouse.

The estimation results reveal substantial cross-sectional heterogeneity in wives' Pareto weights, with their bargaining position increasing with their income share and decreasing with PI. Regarding gender differences in saving motives, we find that in the low- and middle-PI groups, women have stronger saving incentives driven by precautionary motives due to longer life expectancy and greater long-term care risk, whereas in the high-PI group, men save more because of their stronger bequest motives toward children.

Using the estimated model, we first quantify intrahousehold consumption inequality in retirement. Comparing couples who entered retirement in 1998–2000 with those who did so in 2012–2014 reveals two offsetting compositional changes: the younger cohort contains a larger share of high-PI households where women's bargaining power is weaker, but also more households with higher wives' income share which positively affects their bargaining power. On net, the second force dominates: the distribution of wives' Pareto weights shifts to the right, from a mean of 0.45 to 0.48. This shift translates into a modest rise in wives' consumption shares during marriage (from 0.482 to 0.495). At the same time, we find that dispersion in Pareto weights and consumption shares increased in the more recent cohort.

Next, we examine intrahousehold heterogeneity in savings motives. To understand how much couples' savings reflect the different incentives of spouses, we conduct counterfactuals in which either the wife or the husband receives the entire Pareto weight. Under either scenario, households' precautionary savings motives increase. This is because the spouse with full power never provides spousal care due to substantial caregiving disutility, which raises the household's exposure to formal long-term care risks. Consequently, in both scenarios, savings rise relative to the baseline, in which the mean Pareto weight is around 0.45 and both spouses provide substantial care. Which scenario leads to a larger increase in savings varies across PI groups. Low- and mid-PI households save more when the wife holds all power, consistent with her stronger precautionary motives, while high-PI households save more when the husband holds all power, reflecting his stronger bequest motives toward children.

Finally, we use the model to evaluate how Medicaid's treatment of housing assets in its means test affects retirees' welfare. We find substantial heterogeneity, with particularly strong effects for initial singles, for women, and for low-PI retirees. Importantly, our collective approach reveals meaningful within-household differences, with wives typically experiencing larger welfare changes due to their stronger precautionary savings motives. These gender disparities would be invisible under a unitary approach, underscoring the value of a collective model for evaluating policy reforms.

This paper contributes to a large empirical literature on savings in retirement.[1] Prior work has made substantial progress in identifying why retirees save, emphasizing precautionary motives arising from medical expense risk and bequest motives.[2] However, nearly all of this literature focuses on single retirees. Only a small set of studies examine couples' savings (e.g., Braun, Kopecky, and Koreshkova (2017), Nakajima and Telyukova (2020), De Nardi, French, Jones, and McGee (2025)), and these typically rely on unitary frameworks that are based on aggregate household preferences. Our contribution is to incorporate intrahousehold heterogeneity into the analysis of retirees' savings. By estimating a collective model that separates each spouse's savings motives and bargaining power, we show how savings motives differ within married households and how these differences are aggregated into household-level decisions through their respective Pareto weights.

Our analysis also connects to the literature on collective models, which began with the foundational work of Chiappori (1988), Chiappori (1992) and Apps and Rees (1988). In contrast to unitary frameworks in which the household is characterized by a unique utility function, collective approaches specify individual preferences and aggregate them using relative bargaining weights that place allocations on the Pareto frontier. Cross-sectional variation in Pareto weights is captured by their dependence on the economic environment, such as wages and non-labor income. The presence of a distribution factor (Browning and Chiappori (1998), Chiappori, Fortin, and Lacroix (2002), Blundell, Chiappori, and Meghir (2005)), such as income shares or sex ratios, facilitates identification because it affects only the Pareto weights and not preferences or the budget set. A large empirical literature exploits such cross-sectional variation to estimate collective models of the working-age population.[3] Closely related to our study is Lise and Seitz (2011), who use a collective model to estimate intrahousehold consumption inequality among working-age households in the U.K. We contribute to this literature by being the first, to the best of our knowledge, to estimate a collective model of retirees and to measure intrahousehold inequality in old age.

In addition, our paper contributes to the growing literature on home equity in retirement. Venti and Wise (2004) found that retirees typically do not liquidate home equity to support general nonhousing consumption unless they experience the death of a spouse or enter into a nursing home. Using an estimated life-cycle savings model, Nakajima and Telyukova (2020)

[1] See De Nardi, French, and Jones (2016a) for an excellent survey.

[2] Some examples on precautionary savings motives include works by Hubbard, Skinner, and Zeldes (1995), Palumbo (1999), De Nardi, French, and Jones (2010) and Kopecky and Koreshkova (2014b), and those on bequest motives include Hurd (1989), De Nardi (2004), De Nardi, French, and Jones (2010), Lockwood (2018) and Ameriks, Caplin, Laufer, and Van Nieuwerburgh (2011).

[3] For example, see Browning, Bourguignon, Chiappori, and Lechene (1994), Lundberg, Pollak, and Wales (1997), and Blundell, Chiappori, Magnac, and Meghir (2007).

found that homeowners dissave more slowly than renters because they have a preference for staying in their own home as long as possible and cannot easily borrow against it. A recent work by McGee (2019) used UK data to estimate a retirement savings model that incorporates house price shocks. Additionally, Achou (2021) studied, in the context of single retirees only, the value of Medicaid's homestead exemption. We contribute to this literature by studying how asset composition interacts with intrahousehold dynamics and policy design to shape individual welfare within couples.

Finally, our work is related to the literature on old-age caregiving. Papers by Barczyk and Kredler (2018), Ko (2022) and Mommaerts (2016) use an intergenerational life-cycle savings model to study long-term care arrangements between elderly parents and adult children. A recent work by Barczyk, Kredler, and Fahle (2022) studies how housing assets can be used by parents as a commitment device to leave larger bequests and to elicit caregiving behaviors from their children. While our model accounts for care by adult children in a reduced-form manner, our contribution lies in endogenizing spousal caregiving, which is the primary source of family insurance for married retirees. We find that intrahousehold bargaining is central to determining spousal care and, consequently, the household's precautionary savings motives.

The rest of this paper proceeds as follows. Section 2 presents descriptive evidence. Section 3 presents the model. Section 4 presents our data and estimation results. Section 5 presents the main results. Section 6 concludes.

2 Data and descriptive analysis

2.1 Data

The main dataset for this paper comes from the Health and Retirement Study (HRS) which has surveyed a representative sample of Americans over the age of 50 every two years since 1992. We use biennial interview waves from 1998 to 2014. We only consider individuals who were responsive and aged 65 and older in 1998; retired by 2000; and did not miss any interviews while alive.[4] The main empirical sample consists of 6,475 respondents and 35,865 respondent-wave observations. Appendix A provides details about the data and Table A.1 provides summary statistics of the sample.

[4]Respondents who move into a nursing home remain in the sample. We exclude individuals that leave the sample for reasons other than death to avoid conflating economic behavior with selective survey attrition.

2.2 Construction of permanent income (PI)

We closely follow De Nardi, French, Jones, and McGee (2025) in specifying household PI. We assume that log household income evolves according to

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

where i indexes household, $\kappa(t, f_{it})$ is a flexible function of household head's age t and family structure f_{it} (i.e., couple, single man, or single woman), $h(I_i)$ is specified as a flexible polynomial in I_i, and ω_{it} captures measurement error. From the estimated fixed effects, we obtain $\hat{I}_i$, which represents the rank order of each household's estimated fixed effect and $\hat{h}(\hat{I}_i)$.[5] As in De Nardi, French, Jones, and McGee (2025), this household's PI measure is invariant to changes in the family structure.

For our subsequent analysis, we also require a measure of the wife's share of household income among couples. To do so, we regress the household-level PI measure, $\hat{h}(\hat{I}_i)$, on the average incomes of the husband and the wife over the sample period.[6] Using the estimated coefficients from this regression, we define the wife's income share for a household as the fraction of predicted household $\hat{h}(\hat{I}_i)$ attributable to the wife's average income. The mean value of this wife's income share across households is 0.34.

When we consider a finite number of PI values, we partition households into three groups based on terciles of $\hat{h}(\hat{I}_i)$. We refer to these groups as low PI, mid PI, and high PI, corresponding to the first through third terciles. We likewise form three groups for the wife's income share based on terciles of its sample distribution.[7]

2.3 Evidence on collective decision-making in retirement

We first examine whether married retirees' consumption and savings behaviors are consistent with a unitary model, a widely used framework in the literature. If unitary models provide a good approximation of couples' decision-making during retirement, then distribution factors, i.e., variables that influence the intrahousehold bargaining process but do not directly affect preferences or the budget constraint, should have no effect on household behavior.

[5] Additional details are provided in Online Appendix B.

[6] In the HRS, all incomes are reported at the individual level except for capital and other income which are reported at the household level. To define each spouse's individual income, we divide these two measures by two.

[7] The 33rd percentile cutoff is 0.24, and the 67th percentile cutoff is 0.35.

We use the wife's income share, defined in Section 2.2, as the distribution factor. This choice follows a large body of work that employs relative earnings as a distribution factor for working-age couples (e.g., Browning, Bourguignon, Chiappori, and Lechene (1994), Lise and Yamada (2019), and Lise and Seitz (2011)). While numerous studies have tested and rejected the income-pooling implication of unitary models for the working population, evidence for retired couples remains limited.[8]

Table 1 reports the association between the wife's income share and household consumption, conditional on observable characteristics including assets, household current and permanent income, age, health, number of children, education, insurance coverage, and birth cohort fixed effects.[9] Consumption is measured using the HRS Consumption and Activities Mail Survey (CAMS) subsample, combining nondurable, durable, housing, and transportation spending. Column (1) uses the consumption rate (consumption to income ratio), and column (2) uses log consumption. A 10–percentage point increase in the wife's income share is associated with a 3.4–percentage point decline in the consumption rate and a 2.9% decrease in consumption. Table C.4 in the Online Appendix use savings rate and savings in dollars as the dependent variable. Consistent with Table 1, they show that on average, the wife's income share is positively correlated with household savings.

As discussed later in Section 4, we revisit this relationship using auxiliary regressions from the indirect inference analysis, which include a smaller set of controls limited to variables incorporated into our structural model and disaggregate the impact of the wife's income share by PI group. The results again provide evidence that household behavior responds to individual income shares.

Together, these results provide descriptive evidence consistent with a collective model of household behavior in retirement, in which the wife's bargaining power rises with her relative income share. This motivates the structural collective framework developed in the next section, which allows us to quantify intrahousehold inequality and assess how policy changes affect individual welfare within a married household.

2.4 Marital status and savings in retirement

To analyze intrahousehold dynamics among couples, it is important to also understand singles' behavior. Because couples may transition into singlehood following a spouse's death, their savings and long-term care decisions necessarily reflect expectations about how assets

[8] See Chiappori and Mazzocco (2017) for an excellent review of the literature.
[9] See Table C.3 in the Online Appendix for the full list of controls.

Table 1: Wife's income share and household consumption

	(1) Consumption rate	(2) Log of consumption
Wife's income share	-0.34**	-0.29***
	(0.16)	(0.11)
Observations	1701.00	1701.00
Adjusted R^2	0.40	0.42
Mean of dependent variable	1.39	10.87

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the household level (in parentheses). Sample consists of married respondents from HRS 1998-2014. Consumption questions were asked only to a randomly selected subsample. See Table C.3 in the Online Appendix for coefficients on other covariates.

will affect their value as single retirees. At the same time, singles' savings choices, which are free from intrahousehold bargaining, reveal gender-specific preferences that underlie household decisions. The recent study by De Nardi, French, Jones, and McGee (2025) provides detailed evidence on how total savings, combining liquid and housing assets, differ between couples and singles. We build on this literature by separately examining housing assets by marital status and by analyzing singles' dissaving behavior by gender.

Figure 1 presents homeownership rates by current marital status and PI terciles.[10] Singles are those not currently married, most of whom are widows or widowers.[11] Because individuals transition from married to single through spousal death and exit the sample upon death, the data are not balanced. The third column reports the share of singles among survivors at each age. Among married households, homeownership declines only modestly with age, with about 75% still owning homes at age 90. In contrast, singles exhibit rapid housing dissaving, and by age 90 only about half remain homeowners. Figure C.3 in the Online Appendix shows similar patterns using the median housing asset share. This stark divergence indicates that singles' faster dissaving is concentrated in housing: as Figure C.4 shows, non-housing asset trajectories are similar for couples and singles.

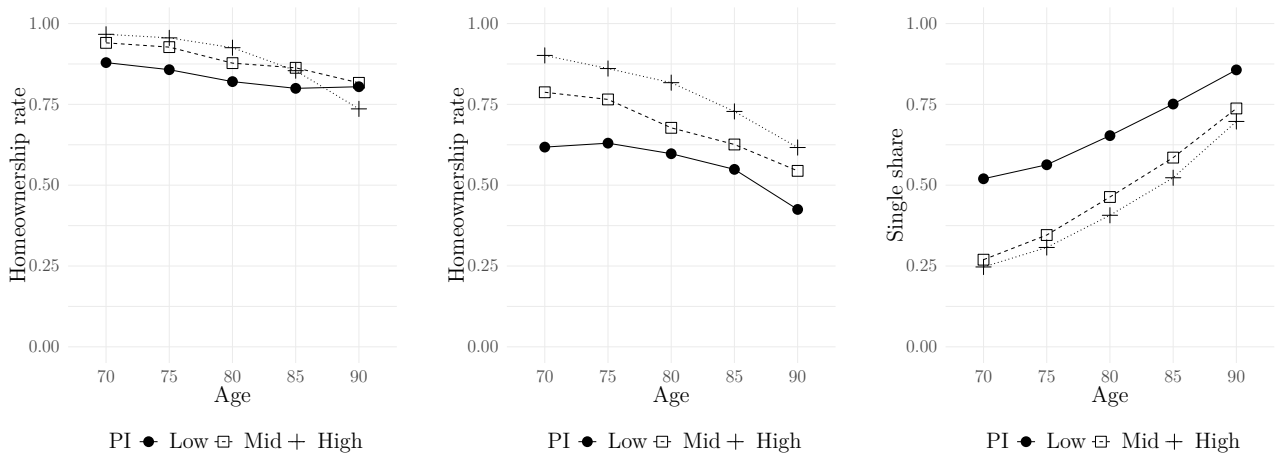
Exploiting the panel nature of the data, we conduct an event study analysis to examine changes in the homeownership rate around spousal death. Details are given in the Online Appendix C.1. The event study shows a substantial decline in homeownership following spousal death: relative to four years before the event, homeownership is about 12.5 percentage points lower two years after and 16 percentage points lower four years after.[12]

[10] Figure C.2 in the Online Appendix shows results by birth cohort.

[11] Among single retirees, about 80% are widowed, 6% never married, and 14% are divorced or have an unknown marital ending.

[12] Table C.5 in the Appendix uses the same sample as in the event study and reports living arrangements

Figure 1: Homeownership rates by marital status



Couples Singles Share of singles

Notes: PI represents permanent income, grouped into terciles. Sample Figure C.2 in the Online Appendix reports the results conditional on birth cohorts.

The descriptive patterns presented so far suggest that couples are more likely to hold housing assets, which leads to slower dissaving of retirement wealth. One possible reason is differential long-term care arrangements by marital status. Table 2 shows that for married disabled individuals with a healthy spouse, 75% receive spousal care, with similarly high caregiving rates for both husbands and wives, with limited reliance on nursing homes (26%) or children (25%, mostly complementing spousal care). When both spouses are disabled, accounting for only 22% of married cases, care provided by children becomes most common (49%). For singles, care from children dominates (63%), and nursing home use is higher (38%) relative to couples. Lack of spousal care and increased risk of moving into a nursing home might contribute to singles' lower value for homeownership.[13]

Another potential mechanism is Medicaid's differential treatment of housing assets in its means test by marital status. In the U.S., formal long-term care costs are costly, with median nursing home costs exceeding $90,000 annually in 2017. Medicaid is a means-tested public health insurance program that covers long-term care costs for impoverished retirees.[14] As the

four years after spousal death. Among those that liquidate their house and become non-homeowners, 67% move into rented units, 24.4% move into a nursing home, and 7.6% move into their children's home.

[13]This, in turn, might increase singles' incentive to sell their house even before the nursing home shock is realized. Once an individual becomes disabled and requires a prolonged stay in a nursing home, there is limited opportunity to enjoy the proceeds from selling the home, since the scope for discretionary consumption in a nursing home is very small. Consequently, the increased nursing home risk can induce singles to sell their home in advance.

[14]Medicare, which is a universal public health insurance program for retirees does not cover prolonged stays at nursing homes.

Table 2: Long-term care arrangements by marital status

	Married		Single
	Spouse healthy	Spouse disabled	
	(1)	(2)	(3)
Caregiving by spouse	0.75	0.39	—
:*by wife*	0.76	0.41	—
:*by husband*	0.74	0.36	—
Caregiving by children	0.25	0.49	0.63
Nursing home care	0.26	0.24	0.38
Observations	1366	397	4553

Notes: The sample consists of retirees that have long-term care needs. We classify individuals as having long-term care needs if they report at least two limitations in activities of daily living (ADLs) or exhibit cognitive impairment.

primary payer for long-term care in the U.S., Medicaid excludes the home from eligibility tests as long as it remains a "home." However, when a recipient has no living spouse in the home or an unmarried recipient moves into a nursing home without intent to return, Medicaid can place liens on the property and recover long-term care costs from the estate upon death (Department of Health and Human Services, 2005). Greenhalgh-Stanley (2012) uses state-level policy variations and finds that such liens reduce singles' homeownership relative to couples. Using a longer sample period with more policy changes than Greenhalgh-Stanley (2012), we estimate the differential effect of Medicaid's estate recovery on singles versus couples (details are in Appendix C.2).[15] We find that these liens indeed lower singles' homeownership by 8.9 percentage points relative to couples, consistent with earlier findings.

We conclude this descriptive section by documenting how savings differ by gender among single retirees. The sample for Table 3 consists of single retirees found in the CAMS. As in Table 1, we include various observables including assets, current and permanent income, age, health, number of children, education, insurance coverage, and birth cohort fixed effects.[16] Column (1) indicates that single women consume a smaller share of their resources than single men, a pattern that is most pronounced in the lowest PI group and diminishes at higher PI levels as shown in Column (2). Turning to savings behavior, Column (3) shows only a small difference by gender. However, allowing the relationship to vary with PI in Column (4) reveals that single women in the low and mid PI groups save more than observably similar men (although these estimates are less precise), whereas single women in the high PI group save significantly less. Columns (5) and (6), which use changes in total assets, display

[15] Our sample covers 1998–2014, during which 25 states adopted TEFRA liens, compared to 19 states by 2004 in Greenhalgh-Stanley (2012).

[16] Table C.6 of the Online Appendix reports the full list.

Table 3: Gender and singles' savings

	(1) Consumption rate	(2) Consumption rate	(3) Savings rate	(4) Savings rate	(5) Changes in assets	(6) Changes in assets
Female	-0.14***	-0.26**	0.17	0.29	1566.94	5481.49
	(0.04)	(0.11)	(0.12)	(0.18)	(5307.78)	(3667.26)
PI (mid) × Female		0.20*		0.19		8190.54
		(0.12)		(0.25)		(7719.08)
PI (high) × Female		0.18		-0.70**		-24923.02
		(0.12)		(0.33)		(17509.37)
Observations	3530.00	3530.00	13692.00	13692.00	13692.00	13692.00
Adjusted R^2	0.26	0.26	0.14	0.14	0.22	0.22
Mean dep. var	1.08	1.08	-0.40	-0.40	-16073.37	-16073.37

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the household level (in parentheses). Sample consists of single respondents from HRS 1998-2014. Consumption questions were asked only to a randomly selected subsample. See Table C.6 in the Online Appendix for coefficients on other covariates.

a similar structure. Although the asset-based estimates are noisier, they again suggest relatively higher asset accumulation for single women in the low and mid PI groups and substantially lower accumulation for single women in the high PI group, relative to their male counterparts.

3 Model

The model presented in this section describes retirees' housing, long-term care arrangement, and consumption-savings decisions in the face of health and mortality shocks. Time, $t = 1, 2, ..., T$, is discrete and corresponds to the age of the household head. The first period represents age 65, and the final period corresponds to age 99, after which individuals are assumed to die with certainty. For married couples, the husband is the household head, and the wife is two years younger. As the HRS interviews are conducted biennially, each period spans two years. Individuals' discount factor is denoted by β. For married households, we use a collective household model to describe couples' decision making process.[17] The model incorporates welfare programs including Medicaid as a lower bound on consumption. Table D.1 in the Online Appendix summarizes model variables.

[17] An alternative non-unitary framework is the limited commitment approach, in which large shocks to either spouse's outside option may cause a constraint to bind, trigger a change in Pareto weights, and, if no feasible reallocation satisfies both constraints, lead to divorce. However, income shocks in retirement are small given the dominance of stable sources such as Social Security and pensions, and there were no major changes to divorce laws or Social Security rules during our sample period. As pointed out in Ligon, Thomas, and Worrall (2002) and Lise and Yamada (2019), limited commitment models make different predictions from collective models only when large income shocks occur. Consistent with small income shocks in retirement, divorce is exceedingly rare in the HRS data at only about 0.5%.

3.1 Timing

At the start of each period, mortality shocks are realized. In the event of survival, health risks are first realized. If there is a sick member in the household, long-term care arrangements are determined which could be either spousal care (available only for married households), nursing home care, or informal care provided by adult children. Homeowners then decide whether to sell their home, and renters choose housing services. The government makes transfers to guarantee a minimum consumption floor. Finally, the household chooses consumption.

3.2 Preferences

Gender is denoted by $g \in \{W, M\}$, where W refers to women and M to men. For married households, we specify each member's preferences separately rather than using an aggregate household utility function. This methodological individualism is essential for analyzing intrahousehold inequality and aligns with neoclassical microeconomic analysis, which is based on individual preferences rather than those of a group (Chiappori, 1992). Flow utility in period t for a married individual of gender g is given by

$$U_t^g = u^g(c_t^g, h_t^g) - \psi^g P_t^g \tag{1}$$

The function u^g represents gender-specific utility from non-housing and housing consumption. As is common in the literature (e.g., Lise and Seitz (2011), Blundell, Chiappori, and Meghir (2005)), we allow for gender-specific preferences.[18] c_t^g denotes private non-housing consumption, and h_t^g represents housing services. Although housing is typically modeled as a public good (i.e., $h_t^g = h_t^{g'}$), we allow them to differ when one spouse resides in a nursing home, as such residency significantly constrains access to housing services. The indicator variable P_t^g equals 1 if individual g provides spousal care and 0 otherwise. If both spouses are healthy, then no spousal care is needed, so $P_t^g = P_t^{g'} = 0$. If one spouse g' is disabled and the other g is healthy, the healthy spouse may provide care. If spousal care is provided ($P_t^g = 1$), then the care provider obtains disutility of ψ^g. While ψ^g is expected to be positive, reflecting the disutility from providing care to a disabled spouse, we allow it to take either sign.[19] If both spouses are disabled, they are assumed to be too ill to provide care, so

[18] When household members share identical preferences, collective models collapse to a unitary model. See Vermeulen (2002) for discussions about when unitary models emerge as a special case of a collective model.

[19] An alternative specification allows the care recipient to derive utility from receiving spousal care. Point identification of this recipient parameter in addition to ψ^g is hard given the data variations we have. We therefore focus on the caregiver's preference ψ^g, which is allowed to be negative or positive.

$P_t^g = P_t^{g'} = 0$.

For singles, the flow utility does not include the caregiving term, and its functional form is given by

$$u^g(c_t, h_t) = \frac{{c_t}^{1-\gamma} - 1}{1-\gamma} + \sigma^g \frac{{h_t}^{1-\gamma} - 1}{1-\gamma}, \tag{2}$$

where c_t and h_t represent consumption of non-housing and housing goods, respectively. As discussed later, we allow for economies of scale in consumption by marital status. Following De Nardi, French, Jones, and McGee (2025), we assume that an individual's preferences over a given bundle of consumption do not depend on marital status.[20] This assumption also facilitates identification, as discussed in Section 4.2.1.

The functional form in (2) follows Campbell and Cocco (2007), who impose a common curvature on non-housing and housing consumption and introduce a relative weight on housing. Because identifying both the housing curvature and the housing preference weight is difficult with our data, we likewise use the same curvature for housing; further discussion is provided in Section 4.1. We introduce preference heterogeneity across genders through the housing preference weight σ^g, rather than through risk aversion. This choice is also motivated by identification considerations and is explained in detail in Section 4.1.

When there is no surviving spouse, the bequests go to children (or other heirs). As we do not model these agents' preferences, we adopt a warm-glow approach and specify bequest preferences as

$$v^{S,g}(b_t) = \delta^g \frac{(a^g + b_t)^{1-\gamma} - 1}{1-\gamma}, \tag{3}$$

where b_t denotes the total remaining household assets, including both liquid and housing assets. This is a commonly used functional form to model bequest motives when preferences of the inheritors are not specified (e.g., De Nardi (2004), De Nardi, French, and Jones (2010), and Lockwood (2018)). The parameter δ^g represents the magnitude of bequest motives, and a^g represents the threshold of consumption level below which no bequests are left under conditions of perfect certainty. We allow for gender differences in bequest preferences, which we estimate structurally.

When a deceased individual is survived by a spouse, data from the HRS exit interviews show that the vast majority of bequests are left to the surviving spouse, with an average of 85.6% of the estate. We therefore assume that the surviving spouse retains the remaining

[20] Voena (2015) also assumes consumption utility remains the same after getting divorced.

household assets. As the surviving spouse's preferences are specified, we model bequest motives of a married individual as pure altruism,

$$v^{C,g}(x_t) = \alpha V_t^{S,g'}(x_t), \tag{4}$$

where x_t is the household state vector, described in Section 3.10, and $V_t^{S,g'}(x_t)$ is the surviving spouse's continuation value as a single retiree, which depends on remaining household assets, among other elements. The altruism parameter α reflects the extent to which a married individual values the well-being of their surviving spouse at the time of their own death. While it may vary by gender, we estimate an average level for identification considerations, which we detail in Section 4.2.1. The use of pure altruism in modeling bequest motives is commonly used when the inheritor's preferences are specified (e.g., Kaplan (2012) and Ko (2022)).

An alternative way to specify bequest preferences upon the death of a married individual is to allow for bequests not only to the surviving spouse but also to children (or other heirs). For instance, the data show that while small, 14.4% of the remaining assets go to non-spousal heirs at the time of the first spouse's death. While De Nardi, French, Jones, and McGee (2025) explicitly consider determination of these "side bequests" using a unitary framework, identification becomes challenging in our collective framework as we have to define preferences over these side bequests for both the deceased and the surviving spouse.[21] In the Online Appendix D.2, we outline how side bequests could be incorporated into our model and explain why the available data do not allow credible identification without strong assumptions. As such, we abstract from side bequests and focus instead on bequests to the surviving spouse, which dominate empirically and offer an opportunity to study within-household altruism at death, an aspect possible only in a collective framework.

3.3 Consumption

We assume non-housing consumption depends on nursing home (NH) residency:

$$c_t = \begin{cases} \hat{c}_t & \text{if not in NH} \\ c_{nh} & \text{if in NH} \end{cases} \tag{5}$$

[21] In contrast, De Nardi, French, Jones, and McGee (2025) do not face these problems as they have a unitary framework and endow households with a single aggregate utility.

When the individual is not in a nursing home, consumption is chosen freely and denoted by $\hat{c}_t$. In contrast, when the individual resides in a nursing home, consumption is fixed at a basic level c_{nh}. This reflects the institutional reality that nursing home services are delivered as a bundled package covering not only medical and housing services, but also meals, basic amenities, and routine assistance with daily activities, which substitute for the non-housing goods and services individuals would otherwise choose themselves. As a result, individuals have limited opportunities to enjoy additional consumption beyond what the facility provides. Papers such as Kopecky and Koreshkova (2014a) similarly model residents in publicly financed nursing homes as receiving a fixed level of consumption.[22]

The household consumption expenditure is given as

$$\chi_t = \begin{cases} [(\hat{c}_t^W)^\rho + (\hat{c}_t^M)^\rho]^{\frac{1}{\rho}} & \text{for couples} \\ \hat{c}_t & \text{for singles} \end{cases} \tag{6}$$

$\rho \geq 1$ means there are economies of scale for couples' consumption. When none of the spouses is in a nursing home, then each spouse gets their own private consumption, denoted by $\hat{c}_t^g$ for $g \in \{W, M\}$. If only one spouse is in a nursing home, then that spouse's consumption expenditure is zero as they cannot consume beyond basic c_{nh}, and the other spouse gets the entire household consumption expenditure. If both spouses are in a nursing home, then the household will optimally choose $\chi_t = 0$.

3.4 Housing

Each family enters period t with initial housing assets denoted by $\tilde{h}_{t-1} \geq 0$. Strictly positive values of $\tilde{h}_{t-1}$ imply the household is a homeowner, and $\tilde{h}_{t-1} = 0$ means the household is a renter. Switching from a renter to a homeowner is rare in retirement, so we assume renting is an absorbing state.[23] Homeowners decide whether they sell home ($D_t = 1$) or keep it ($D_t = 0$). We do not allow homeowners to switch houses unless they become a renter, i.e., there is no downsizing. This assumption is grounded in the empirical study by Venti and Wise (2004), which finds that, on average, there is no reduction in housing equity among retirees who continue to own homes. Significant reductions in housing equity are observed

[22] In theory, the consumption value from nursing home care might depend on whether services are publicly or privately financed. However, papers such as Lockwood (2018) find that public care aversion is modest in the HRS. For example, the consumption value (including housing) provided by publicly financed care facilities is estimated at \$19,100 per year, which is very close to the \$20,000 benchmark used for privately financed facilities.

[23] In the HRS data, less than 10% of retirees purchase a new home.

only when individuals sell and discontinue homeownership. The no-downsizing assumption has become standard in structural models that study housing decisions in retirement in the U.S.[24]

After homeowners' liquidation decision, housing assets are updated as the following:[25]

$$\tilde{h}_t = \begin{cases} \tilde{h}_{t-1} & \text{if } \tilde{h}_{t-1} = 0 \text{ or } (\tilde{h}_{t-1} > 0 \text{ and } D_t = 0) \\ 0 & \text{if } \tilde{h}_{t-1} > 0 \text{ and } D_t = 1 \end{cases} \tag{7}$$

Households with $\tilde{h}_t = 0$ decide on renting services, denoted by R_t.

Individual consumption of housing services depends on homeownership, marital status, and nursing home residency:

$$h_t = \begin{cases} \omega_{couple}\tilde{h}_t & \text{if not in NH, homeowner, and couple} \\ \omega_{single}\tilde{h}_t & \text{if not in NH, homeowner, and single} \\ R_t & \text{if not in NH and renter} \\ h_{nh} & \text{if in NH} \end{cases} \tag{8}$$

If an individual is a homeowner, then the individual's housing consumption is equal to the home's value multiplied by the homeownership premium which varies by marital status. This premium is lower for singles ($\omega_{couple} > \omega_{single}$) due to their higher likelihood of moving into their child's home, representing a greater opportunity cost of homeownership.[26] The lower premium also accounts for potentially higher maintenance costs for singles (de Ruijter, Treas, and Cohen, 2005). For renters, they derive utility from the rented housing service R_t. For individuals residing in nursing homes, housing consumption is set to a fixed value h_{nh}, reflecting the standardized living arrangements provided by these facilities.[27]

[24] Nakajima and Telyukova (2020) assume that for homeowners, the housing decision is whether to move out of the house and rent or to stay in it. Achou (2021) similarly assumes that retirees do not switch homes unless they move to rental housing. Barczyk, Kredler, and Fahle (2022) adopt a related assumption, restricting home purchases beyond a certain age while allowing only sales.

[25] Housing assets are risk-free in our model. Admittedly, this is a simplifying assumption. However, our primary goal is to understand different homeownership incentives between couples and singles. As house price shocks are aggregate shocks that should have similar effects on housing decisions of couples and singles, we abstract from them.

[26] In the HRS data, 18% of singles who sell their house upon spousal death move into their children's home.

[27] As previously mentioned, this approach is consistent with Kopecky and Koreshkova (2014a), who model residents in publicly financed nursing homes as receiving a fixed level of consumption that includes housing. Lockwood (2018) finds only modest differences in the consumption value (including housing) between publicly and privately financed facility care.

Housing expenditure in each period is

$$e(\tilde{h}_t, R_t) = \begin{cases} \delta \tilde{h}_t & \text{if } \tilde{h}_t > 0 \\ (r+\delta) R_t & \text{if } \tilde{h}_t = 0 \end{cases} \tag{9}$$

where δ is the depreciation rate, and r is the real interest rate. Liquidating housing assets worth of $\tilde{h}$ incurs transaction costs $\tau \tilde{h}$.

3.5 Household income

As described in Section 2.2, household income evolves according to

$$\ln y_{it} = \kappa(t, f_{it}) + h(I_i) + \omega_{it}, \tag{10}$$

where $\kappa(t, f_{it})$ is a flexible function of age t and family structure f_{it} (i.e., couple, single man, or single woman), $h(I_i)$ is specified as a flexible polynomial in I_i, and ω_{it} captures measurement error. Using the estimated household PI, we also obtain the wife's income share as the fraction of predicted household PI attributable to the wife's average income over the sample period.[28] To reduce computational burden, we consider three values of PI and three values of the wife's income share based on terciles of their respective sample distribution.

3.6 Health and mortality risk

We consider three health statuses: $s_t \in \{\text{healthy, require long-term care, dead}\}$. Health transition probabilities follow a Markov chain. As in De Nardi, French, Jones, and McGee (2025), they depend on the individual's current health, age, gender, permanent income, and marital status,

$$Pr(s_{t+1} | s_t, age_t, sex, I, f_{it}). \tag{11}$$

The health transition process is exogenous and does not depend on the receipt of care. This is based on previous studies which find that the evolution of elderly health is largely unaffected by the receipt of care (Finkelstein and McKnight, 2008; Card, Dobkin, and Maestas, 2008).

[28] Details are provided in Online Appendix B.

3.7 Long-term care arrangements

When individuals become disabled, they require either informal or formal long-term care. To maintain model tractability, we assume that disabled individuals receive only one type of care at a time, as in Ko (2022) and Mommaerts (2016). For single individuals, the major source of long-term care is their adult children, as shown in Table 2.[29] Previous studies have shown that individuals generally prefer family-provided care over formal long-term care services (Barczyk and Kredler, 2018; Mommaerts, 2016; Ko, 2022). Based on these findings, we assume that disabled singles will rely on care provided by their children if it is "available" ($ic_{child} = 1$) and will resort to formal care only when it is not ($ic_{child} = 0$).[30] We abstract from the possibility that children help their parents pay for formal care services, as such financial transfers are negligible in the data.[31]

For married households, spousal care becomes another significant source of informal care. If one spouse becomes disabled and the other spouse remains healthy, the household collectively decides whether the healthy spouse will provide spousal care, in which case no other type of care is needed. Empirically, 75% of such households rely on spousal care, as shown in Table 2. If the household decides that the healthy spouse will not provide care, then the disabled spouse will use formal care, rather than relying on care from children. This assumption is grounded in the finding that when the healthy spouse does not provide care to their disabled partner, the use of formal care services is more substantial than the help received from children.[32] Because the alternative to spousal care is formal care, spousal care helps avoid costly long-term care expenditures, thereby benefiting the household. The care provider experiences (dis)utility, denoted ψ^g, which reflects forgone leisure and other non-monetary costs associated with caregiving.

[29] Papers such as Barczyk and Kredler (2018), Ko (2022), and Mommaerts (2016) consider children as active agents in long-term care arrangement decisions. While endogenizing children's care decisions would be ideal, we adopt a more reduced-form approach, as our focus is on the intrahousehold dynamics of a couple. Extending the model to incorporate intergenerational dynamics would introduce unnecessary complexity given our focus.

[30] The HRS survey asks, "Suppose in the future, you needed help with basic personal care activities like eating or dressing. Will your daughter/son be willing and able to help you over a long period of time?" If the answer is positive for any of the respondent's children, we assume informal care from children is available ($ic_{child} = 1$). Otherwise, we assume it is not ($ic_{child} = 0$), and the individual uses nursing home care whenever they become sick. We have verified that the reported beliefs align reasonably with actual informal care receipt from children.

[31] Among disabled parents who use formal long-term care, the mean annual financial transfer received from their children is just \$300 (Ko, 2022).

[32] When the healthy spouse does not provide care, the extensive usage rate is 24% for nursing homes and 19% for care from children. Furthermore, conditional on usage, disabled individuals receive about 4 hours of daily care from their children, while those in nursing homes spend an average of 355 nights over a two-year period, suggesting that nursing homes become the primary source of care in the absence of spousal care.

When both spouses become disabled, which are relatively rare due to differences in the onset of long-term care needs by gender, care from children empirically becomes the major source of assistance, as indicated in Table 2. As with singles, we assume that the disabled couple will rely on care provided by their children if it is available. Otherwise, they resort to formal care rather than providing care themselves, as we assume they are too sick to care for each other.[33]

We focus on nursing home care as the sole form of formal long-term care, based on the following grounds. First, nursing homes are the largest financial long-term care risk faced by the elderly: Kemper, Komisar, and Alecxih (2005/2006) find that a person turning 65 in 2005 can expect to spend about $47,000 on formal care before death, most of which is on facility care at about $39,000. Second, the lion's share of Medicaid long-term care spending (about 70%) is attributed to nursing homes. Third, Medicaid's asset test rules regarding housing change dramatically by marital status for nursing home residency, but not for paid home care use.

3.8 Medicaid and government transfers

Government guarantees a minimum consumption floor through means-tested welfare programs such as Medicaid and Supplemental Security Income. To simplify notations, we define the household's cash-at-hand before government transfers as

$$\tilde{a}_t = a_t + y_t + \underbrace{I[D_t = 1](1-\tau)\tilde{h}_{t-1} - e(\tilde{h}_t, R_t)}_{\text{net proceeds from housing decisions}} - \underbrace{m_t}_{\text{cost of NH}} \tag{12}$$

$D_t = 1$ means the household sells home. $e(\cdot,\cdot)$ is the housing expenditure. m_t represents the cost of nursing home care and is strictly positive only for nursing home users.[34]

The amount of government transfers that singles receive is

$$tr_t = \begin{cases} \max\{0, \bar{a}_{nh=0} - \tilde{a}_t\} & \text{if single and not in NH} \\ \max\{0, \bar{a}_{nh=1} - (\tilde{a}_t + (1-\tau)\tilde{h}_{t-1})\} & \text{if single and in NH} \end{cases} \tag{13}$$

[33] Table 2 shows that among disabled individuals married to a partner who is also disabled, over one-third report receiving spousal care. However, on the intensive margin, the number of care hours is only about one third of the care hours provided by a healthy spouse.

[34] We abstract from private long-term care insurance which pays for formal long-term care costs before Medicaid. In the data, the private insurance ownership rate is less than 10%. In contrast to a common misconception, Medicare does not pay for prolonged nursing home stays.

For non-institutionalized singles, the government offers the homestead exemption by excluding housing assets from the eligibility test. Only the liquid assets at hand $\tilde{a}_t$ are considered, and the government makes transfers to ensure the consumption floor for non-institutionalized singles denoted by $\bar{a}_{nh=0}$. For singles in a nursing home, their post-sales housing assets are counted. Eligible single homeowners must first liquidate their home to pay for nursing home care, and then the government makes transfers to ensure the consumption floor for institutionalized singles $\bar{a}_{nh=1}$.[35]

The amount of government transfers that couples receive is

$$tr_t = \begin{cases} \max\{0, 1.5\bar{a}_{nh=0} - \tilde{a}_t\} & \text{if couple and none in NH} \\ \max\{0, \bar{a}_{nh=0} + \bar{a}_{nh=1} - \tilde{a}_t\} & \text{if couple and one in NH} \\ \max\{0, \bar{a}_{nh=1} - \tilde{a}_t\} & \text{if couple and both in NH} \end{cases} \tag{14}$$

The consumption floor for non-institutionalized couples is assumed to be 1.5 times the consumption floor for non-institutionalized singles, $\bar{a}_{nh=0}$, as in De Nardi, French, Jones, and McGee (2025).[36] The key difference from equation (13) is that housing assets are disregarded by Medicaid when a married individual ends up in a nursing home. This is consistent with the Medicaid rules which provide the homestead exemption to married households in a nursing home, as explained in Section 2.

3.9 Asset accumulation law

Non-housing assets evolve according to

$$a_{t+1} = (1 + r)(\tilde{a}_t + tr_t - \chi_t), \tag{15}$$

where χ_t is the household consumption expenditure described in equation (6). We assume there is no borrowing. Furthermore, we abstract from reverse mortgages which enable households to borrow against their home equity. The choice is based on the empirical fact that

[35] As nursing home residents already receive basic food and housing from the facility, we have $\bar{a}_{nh=0} > \bar{a}_{nh=1}$. Furthermore, to keep the computation tractable, we do not model different states and assume the Medicaid estate recovery program is effective. As of 2014, estate recovery programs were implemented in all 50 states, plus the District of Columbia, with around one half of them incorporating TEFRA liens. If we were to relax the effectiveness of the estate recovery program, the program's impact on retirees' housing decisions would be reduced.

[36] Later, we calibrate $\bar{a}_{nh=1}$ to be zero as in Lockwood (2018). Hence, the consumption floor for institutionalized couples would be zero regardless of the constant in front of $\bar{a}_{nh=1}$.

almost no retirees use such loans: according to Nakajima and Telyukova (2017), only 1.9% of eligible homeowners were reported to use reverse mortgages in 2013.

3.10 Recursive formulation

This section provides a recursive formulation for a couple's problem. We derive that of singles in the Online Appendix D.3. A married household's state vector in t is

$$x_t = (a_t, \tilde{h}_{t-1}, s_t^W, s_t^M; I, ic_{child}), \tag{16}$$

where a_t is the non-housing wealth, $\tilde{h}_{t-1}$ is the housing wealth at the beginning of the period, and s_t^g is the health status of each spouse, $g \in \{W, M\}$. Time-invariant state variables are permanent income I and the availability of informal care from children ic_{child}.

The household's choice vector is

$$q_t = (\hat{c}_t^W, \hat{c}_t^M, D_t, R_t, P_t^W, P_t^M), \tag{17}$$

where $\hat{c}_t^g \geq 0$ is private non-housing consumption for member $g \in \{W, M\}$, $D_t \in \{0, 1\}$ represents the house selling choice relevant only for homeowners ($\tilde{h}_t > 0$), $R_t \geq 0$ is the rent choice relevant only for non-homeowners, and $P_t^g \in \{0, 1\}$ indicates spousal care provision of member $g \in \{W, M\}$.

Denote the survival probability of member g in period t by p_t^g, which also depends on current health, PI, and marital status, but is suppressed for notational simplicity. The couple's value function is given as

$$\begin{aligned} V_t^C(x_t) \quad = \max_{q_t} \quad & \lambda(I,z)U_t^W + (1-\lambda(I,z))U_t^M \\ & + \beta p_t^W p_t^M E(V_{t+1}^C(x_{t+1})|x_t, q_t) \\ & + \beta p_t^W (1-p_t^M)(\lambda(I,z) + \alpha(1-\lambda(I,z)))E(V_{t+1}^{S,W}(x_{t+1})|x_t, q_t) \\ & + \beta (1-p_t^W) p_t^M (\alpha\lambda(I,z) + (1-\lambda(I,z)))E(V_{t+1}^{S,M}(x_{t+1})|x_t, q_t) \\ & + \beta (1-p_t^W)(1-p_t^M) E(\lambda(I,z)v^{S,W}(b_{t+1}) + (1-\lambda(I,z))v^{S,M}(b_{t+1})|x_t, q_t) \end{aligned} \tag{18}$$

subject to budget constraints. The Pareto weight $\lambda \in (0, 1)$ reflects the relative bargaining power of the wife. Collective models assume that these weights vary with the environment, capturing cross-sectional differences in intrahousehold bargaining power (Chiappori, 1988,

1992).[37] For the working population, these weights are modeled as functions of wages and household non-labor income. Since we focus on retired households, we exclude wages but include household permanent income, I. In addition, we allow the Pareto weights to depend on the wife's permanent income share, z. This variable serves as a distribution factor, which refers to variables that affect the Pareto weights but not preferences nor the budget set (Browning, Bourguignon, Chiappori, and Lechene, 1994). As shown in papers like Chiappori, Fortin, and Lacroix (2002) and Blundell, Chiappori, and Meghir (2005), the presence of distribution factors is crucial for identification, as we elaborate in Section 4.2.1. The use of relative retirement income as a distribution factor is not only supported by our descriptive results in Section 2.3, but also aligns with a large body of work that employs relative earnings as a distribution factor in studies of the working-age population, including Browning, Bourguignon, Chiappori, and Lechene (1994), Lise and Yamada (2019), and Lise and Seitz (2011).

When both spouses are alive in the next period, then the continuation value is that of a couple, where the expectation operator is taken with respect to health statuses of the next period. When the husband dies and the wife survives, then the husband cares about the wife's continuation value. If the altruism parameter satisfies $\alpha = 1$, then the couple's continuation value is simply the surviving wife's value as a single, $V^{S,W}$. On the other extreme, if the altruism parameter is zero, then the couple places only the wife's Pareto weight, λ, on her value as a single retiree in case of the husband's death. Similarly, when the wife dies and the husband survives, the couple values the husband's value function as a widower, $V^{S,M}$, using a weight that combines the husband's Pareto weight, $1-\lambda$, and the wife's weight, λ, multiplied by her altruism toward her surviving husband, α. If both spouses die, then remaining household assets go to their children or other heirs, where $b_{t+1} = a_{t+1} + (1-\delta)\tilde{h}_t$, in which case each member derives warm-glow utility according to $v^{S,g}$.

The model is solved backwards. For $t = T$, the terminal value for the couple is given as

$$V_T^C(x_T) = \lambda(I,z)v^{S,W}(b_T) + (1-\lambda(I,z))v^{S,M}(b_T). \tag{19}$$

For $t < T$, with singles' value functions in hands (details in Online Appendix D.3), we can solve the couple's problem and obtain V_t^C.

For our analysis of intrahousehold inequality, it is useful to define each married individual's

[37] Cross-sectional variation in Pareto weights is essential for a non-trivial collective model. If the weight is fixed, the model assumes that intrahousehold bargaining power does not depend on prices or the environment, making it indistinguishable from a unitary framework. See Vermeulen (2002) for a discussion of when the unitary model becomes a special case of the collective model.

value function. For $g \in \{W, M\}$, given a sequence of choices, let $V_t^{C,g}$ denote the individual value from being married. Their terminal values are given as

$$V_T^{C,W}(x_T) = v^{S,W}(b_T), \tag{20}$$
$$V_T^{C,M}(x_T) = v^{S,M}(b_T). \tag{21}$$

For $t < T$, given a sequence of choices, $V_t^{C,W}$ and $V_t^{C,M}$ are given as

$$\begin{aligned} V_t^{C,W}(x_t) &= U_t^W + \beta p_t^W p_t^M E(V_{t+1}^{C,W}(x_{t+1})|x_t, q_t) + \beta p_t^W (1-p_t^M) E(V_{t+1}^{S,W}(x_{t+1})|x_t, q_t) \\ &+ \beta(1-p_t^W) p_t^M \alpha E(V_{t+1}^{S,M}(x_{t+1})|x_t, q_t) + \beta(1-p_t^W)(1-p_t^M) E(v^{S,W}(b_{t+1})|x_t, q_t) \\ V_t^{C,M}(x_t) &= U_t^M + \beta p_t^W p_t^M E(V_{t+1}^{C,M}(x_{t+1})|x_t, q_t) + \beta p_t^W (1-p_t^M) \alpha E(V_{t+1}^{S,W}(x_{t+1})|x_t, q_t) \\ &+ \beta(1-p_t^W) p_t^M E(V_{t+1}^{S,M}(x_{t+1})|x_t, q_t) + \beta(1-p_t^W)(1-p_t^M) E(v^{S,M}(b_{t+1})|x_t, q_t) \end{aligned} \tag{22}$$

Note that at optimal allocations, we have $V_t^C(x_t) = \lambda(I,z) V_t^{C,W}(x_t) + (1-\lambda(I,z)) V_t^{C,M}(x_t)$.

4 Estimation

To estimate our model, we employ a two-step estimation procedure. In the first step, we fix or estimate parameters outside the model. In the second step, we use indirect inference to estimate the remaining structural parameters within the model.

4.1 First-stage estimation

Health and mortality risks. The model assumes that health transition probabilities follow an exogenously given Markov process, where health status in the next period is determined by current health, age, gender, permanent income, and marital status. We estimate the transition probabilities using maximum likelihood estimation with a flexible logit specification. Table 4 summarizes simulated life expectancy and long-term care risks at age 65. Life expectancy rises with PI, whereas years with disability decline with PI. Conditional on initial marital status and PI, women live longer than men and spend more years in need of long-term care. Conditional on gender and PI, initially married retirees live longer than their single counterparts, while the difference in long-term care years is negligible.

Risk aversion. We assume the same curvature on non-housing and housing consumption, as shown in equation (2). Without it, we would have to identify both housing curvature and housing preference weight, which is difficult from the data on savings and asset composition:

Table 4: Simulated life expectancy and years with disability at age 65

Initial marital status	Gender	PI=Low	PI=Mid	PI=High
A. Life expectancy (years)				
Married	Man	12.7	16.9	17.8
Married	Woman	17.9	22.0	22.8
Single	Man	11.0	14.7	15.8
Single	Woman	15.3	19.3	20.1
B. Years with disability				
Married	Man	2.4	1.7	1.4
Married	Woman	3.8	2.5	2.2
Single	Man	2.5	1.8	1.3
Single	Woman	4.1	2.6	2.2

Notes: The table reports simulated long-term care and life expectancy from age 65 onward, conditional on marital status at age 65, gender, and PI.

both higher curvature and higher weight would lead to a higher housing asset share. We therefore apply the same curvature to non-housing and housing consumption, as in Campbell and Cocco (2007). Furthermore, we do not allow gender differences in risk aversion, i.e., the same γ is used for both men and women. This choice follows from the identification structure. Risk aversion and bequest motives are difficult to identify from wealth data alone, as both increase savings.[38] We therefore fix one and estimate the other. Because the literature provides stronger consensus on risk aversion among retirees than on bequest motives, we fix risk aversion and estimate bequest motives.[39] However, there is little empirical evidence on gender-specific risk aversion in retirement. We therefore assume $\gamma = 3$ across genders.

Pareto weight inputs. We consider three values of PI (I), low, mid, and high, based on the terciles of the estimated distribution of households' fixed effects, as described in Section 2.2. We parameterize the Pareto weight as

$$\lambda(I, z) = \frac{\exp(\beta_0 + \beta_I \ln(\hat{y}(I)) + \beta_z z)}{1 + \exp(\beta_0 + \beta_I \ln(\hat{y}(I)) + \beta_z z)}, \tag{23}$$

[38] As emphasized in De Nardi, French, and Jones (2016c), distinguishing precautionary saving from bequest-driven saving based only on wealth data is inherently difficult. Papers such as Lockwood (2018) bring in additional data (e.g., long-term care insurance demand) precisely to separate the two forces.

[39] For example, simulation papers such as Hubbard, Skinner, and Zeldes (1995), Engen, Gale, and Uccello (1999), Scholz, Seshadri, and Khitatrakun (2006), and Brown and Finkelstein (2008) all use a value of 3 for the coefficient of relative risk aversion. Structural estimates are typically in a similar range; De Nardi, French, and Jones (2016b) estimate 2.83, De Nardi, French, and Jones (2010) 3.84, De Nardi, French, Jones, and McGee (2025) 3.7, and Lockwood (2018) 4.5. In contrast, estimates of bequest motive strength vary more widely, which we discuss in Section X.

which is similar to the specification in Lise and Seitz (2011), who model the wife's share of household full income for working-age couples.[40] For each PI group, we set $\hat{y}(I)$ as the median income among married households as predicted by equation (10). We also consider three values of the wife's income share (z), calculated as the median within each tercile of its distribution. The resulting representative values are 0.17, 0.30, and 0.44. The parameters $(\beta_0, \beta_I, \beta_z)$ are estimated internally which we describe in the next section.

Economies of scale. The OECD modified equivalence scale assigns a value of 1 to the household head and 0.5 to the spouse. Based on this, we set the parameter on economies of scale in consumption for couples at 1.5. We set the homeownership premium for married couples so that the same magnitude of economies of scale applies to consumption and housing services. Given this condition, we set married couples' homeownership premium at 2.725 and that of singles at 2.162 so that the average homeownership premium is 2.5 as in Nakajima and Telyukova (2020).

Other values. We set the per-capita consumption floor for nursing home residents to zero (Lockwood, 2018). We set the consumption value and housing value of nursing home services to be 40% each of the consumption floor for non-nursing home residents, $\bar{a}_{nh=0}$, which we estimate in the second stage. We set the transaction cost of selling a house at 7% of the house value, following Gruber and Martin (2003). We set the depreciation rate for housing assets at 1.7% per year following Nakajima and Telyukova (2020). Following Brown and Finkelstein (2008), we use 3% time preference rate per year and 3% annual real interest rate. For formal care prices, we use the average rates in 2008, which were \$230 per day for nursing home care (MetLife, 2008).

4.2 Second-stage estimation via indirect inference

We estimate the remaining model parameters using indirect inference, targeting auxiliary models selected based on the identification arguments discussed below. Although the parameters are jointly identified from information contained in all targets, certain data patterns are especially informative for specific parameters, which is what we highlight below.

[40] In their static collective framework, they specify the sharing rule of the full income as $\frac{\exp(x'\psi)}{1+\exp(x'\psi)}$, where the vector x contains the wife's potential share of household earnings, the log of full household income, and the spousal age gap.

4.2.1 Identification

Identification of Pareto weights In collective models, Pareto weights are generally not identified from preferences unless certain conditions are met. We begin by explaining how the presence of a distribution factor identifies variation in the Pareto weights with respect to their arguments, as shown in Blundell, Chiappori, and Meghir (2005). Because we specify the Pareto weight as $\lambda(I, z)$, where I denotes PI and z represents the wife's income share, this amounts to identifying the derivatives, $\partial\lambda/\partial I$ and $\partial\lambda/\partial z$.

Identifying $\partial\lambda/\partial z$ is straightforward as income shares satisfy the definition of a distribution factor, meaning they influence household choices solely through their impact on bargaining power. Therefore, variation in household decisions across different income shares directly identifies $\partial\lambda/\partial z$. Specifically, we use how the household consumption rate changes in z, which we further elaborate below.

Identifying $\partial\lambda/\partial I$ is more challenging, because I enters both the budget and the Pareto weight. Our strategy exploits how the effect of the distribution factor varies with permanent income, together with the parametric structure of $\lambda(I, z)$. Suppose we observe the cross-partial derivative, $\partial^2\lambda/\partial I\partial z$. Then, by integrating this object with respect to z, we can recover $\partial\lambda/\partial I$. The problem is that we do not directly observe the cross-partial derivative of λ itself. However, as household choices depend on λ, their response to z and the way this response varies with I reflects the cross-partial derivative of λ.

To capture this relationship, we target the coefficients from an auxiliary regression of the household consumption rate on the interaction term between I and z.[41] The key difference from the regression in Section 2 is that here we focus on how the effect of the distribution factor varies with PI, as the interaction term identifies the cross-partial effect of I and z.[42] The empirical results in Table 5 show that household consumption declines with the wife's income share, consistent with Section 2, but this negative effect *weakens* as PI increases. Mechanisms such as means-tested social insurance would predict the opposite pattern: at low PI, the consumption floor would bind and mute the influence of distributional factors, while at high PI, the household becomes less constrained and more responsive to changes in bargaining power. Observing the reverse pattern instead suggests a countervailing force. Under our parametric assumption of $\lambda(I, z) = \frac{e(I,z)}{1+e(I,z)}$ where $e(I, z) = \exp(\beta_0 + \beta_I I + \beta_z z)$, the result can be rationalized if $\beta_I < 0$. The inequality would imply the wife's bargaining

[41] We use the consumption rate to remove the mechanical scaling effect of higher income. Empirical results using log consumption are qualitatively similar and are reported in Table E.1 of the Online Appendix.

[42] Furthermore, the regression includes only controls incorporated into the structural model so that it can be replicated using simulated data.

Table 5: Auxiliary OLS I (empirical)

	(1)	
	Consumption rate	
Wife's income share	-0.81**	(0.34)
PI (mid) × Wife's income share	0.62	(0.39)
PI (high) × Wife's income share	0.72*	(0.38)
Observations	1915.00	
Adjusted R^2	0.42	
Mean of dependent variable	1.24	

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the household level (in parentheses). This table reports the auxiliary regression estimated on the empirical data. For the auxiliary specification to be replicated on simulated data, the regression includes only observables incorporated in the structural model, including ages, liquid and housing assets, health, availability of informal long-term care from children, and household income.

power declines with household income, as found in Lise and Seitz (2011). Accordingly, the regression coefficients on the interaction terms provide key information for identifying β_I.[43]

We also identify the level of the Pareto weights, which amounts to identifying β_0. Earlier works on collective models focus on identifying only the derivatives of bargaining power, not its level. Some studies have shown identification of the location of the Pareto weights by using combined data on singles and couples and assuming certain preferences do not change with marital status (e.g., Vermeulen (2006), Browning, Chiappori, and Lewbel (2013)). Similar to these studies, we assume the same preference over general and housing consumption between married and single individuals, as shown in equation (2). In addition, we exploit variation in household decisions arising from the nursing home entry of one spouse. This event constrains the institutionalized spouse's consumption and fixes their current-period utility. As a result, the household's decisions in that period primarily reflect the preferences of the community spouse. Crucially, the *change* in behavior due to the nursing home shock provides information about the *pre-shock* Pareto weights: households in which the community spouse held a high weight prior to entry would show minimal behavioral change, whereas households with lower pre-shock weights would exhibit more pronounced adjustments.

Based on this argument, we target the coefficients from an auxiliary regression that relates changes in the household consumption rate to nursing home entry. Specifically, we regress

[43]Within our model, another possible explanation for the pattern in Table 5 is that women have weaker bequest motives toward children when they die as single retirees. In that case, in high-PI households, for whom bequests are more relevant, consumption might increase when wives' bargaining power increases. However, as described below, gender-specific bequest motives are primarily identified from singles' savings profiles which are free from intrahousehold dynamics. Therefore, the variation in Table 5 primarily informs β_I, the parameter governing how bargaining power varies with PI.

Table 6: Auxiliary OLS II (empirical)

	(1) Consumption rate	
Spouse enters NH x Spouse is husband	-0.28	(0.17)
Observations	400.00	
Adjusted R^2	0.48	
Mean of dependent variable	1.32	

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the household level (in parentheses). This table reports the auxiliary regression estimated on the empirical data. For the auxiliary specification to be replicated on simulated data, the regression includes only observables incorporated in the structural model, including ages, liquid and housing assets, health, availability of informal long-term care from children, and household income.

household consumption rate on whether the spouse enters a nursing home interacted with whether the spouse is a husband. The empirical coefficient on the interaction term is reported in Table 6. The household consumption rate reduces further by 0.28 when the husband enters and the wife gains more control, relative to the case when it is the wife that enters the nursing home. The result suggests that, prior to nursing home entry, wives on average hold a lower Pareto weight than husbands.

Identification of remaining preference parameters To help identify the caregiving preference parameters (ψ^W, ψ^M), we exploit spousal care rates conditional on gender. While spousal care rates reflect both caregiving preferences and bargaining power within the household, we discipline the impact of the latter using the auxiliary models described above.

The parameters (σ^W, σ^M) reflect the relative preference for housing consumption. Singles' homeownership rates conditional on gender and PI are particularly informative in identifying these parameters. .

The parameters (δ^g, a^g) govern the bequest motives in the absence of a spouse. To help identify these parameters, we use the median total assets of singles conditional on gender, PI terciles, age groups, and birth cohort. We have eight age groups, spanning ages 65 to 96, with each group covering a 4-year age bin. We consider four birth cohorts based on ages in 1998.

While we allow for gender-specific bequest motives among singles, we assume a common spousal altruism parameter α for married individuals. This is due to identification considerations. Gender-specific bequest motives toward children can be identified from the wealth trajectories of single men versus single women, given that we use the same risk aversion for both genders. However, such clean variation does not exist for couples, as we only observe

household-level savings that reflect a combination of both spouses' preferences. We therefore estimate an average altruism toward the spouse. Couples' wealth profiles are particularly useful in identifying α, and we use the median total assets of couples conditional on PI terciles, age groups, and birth cohort.

To help identify the consumption floor, we use the Medicaid recipiency rates of couples, single men, and single women.

4.2.2 Estimation procedure

We implement indirect inference in a Bayesian way (Fernandez-Villaverde, Rubio-Ramirez, and Schorfheide, 2016) using a multiple-block Metropolis–Hastings algorithm. Details are provided in Section E of the Online Appendix. The algorithm splits the parameters into different blocks, as different subsets of parameters are primarily informed by different sets of auxiliary models, as described in Section 4.2.1. In simulations, we account for cohort heterogeneity by matching wealth moments conditional on birth cohort. To address survival bias, we incorporate mortality heterogeneity into the model and assign simulated households the same mortality histories observed in the data.

4.2.3 Estimation results

Table 7 reports the structural parameter estimates obtained via indirect inference. Caregiving disutility is positive for both genders, indicating that providing care to a disabled spouse is costly. While the estimates appear small, this reflects the high curvature of the CRRA utility function with coefficient three. To interpret the magnitude, we compute for each gender, how much an individual with that gender's median non-housing consumption among married households would be willing to forgo to avoid providing care. For husbands, this value is \$36,192 per year, while for wives it is smaller at \$33,306, despite the higher estimate of ψ^W. This is because wives have lower median consumption and therefore higher marginal utility, so a given utility loss corresponds to a smaller dollar equivalent. These values fall within the range of monetary caregiving costs reported in the literature: Feinberg, Reinhard, Houser, and Choula (2011) estimate the annual cost of informal care at \$10,678; Ernst and Hay (1994) calculate the net cost of informal care for an Alzheimer's patient to be \$33,038 per year; and Skira (2015) estimates middle-aged women's cost of providing long-term care to their disabled mother at \$82,363 per year.

The estimated utility weights on housing consumption are similar across genders. In contrast, the parameters governing bequest motives toward children differ substantially by

Table 7: Estimates from indirect inference

Parameter	Description	Estimate
ψ^W	Caregiving disutility, wives (10^{-8})	0.498
		[0.498, 0.554]
ψ^M	Caregiving disutility, husbands (10^{-8})	0.179
		[0.172, 0.190]
σ^W	Weight on housing consumption, women	0.207
		[0.082, 0.207]
σ^M	Weight on housing consumption, men	0.213
		[0.149, 0.218]
δ^W	Bequest intensity, single women (10^6)	0.496
		[0.470, 0.637]
a^W	Bequest curvature, single women (10^6)	12.481
		[12.341, 12.493]
δ^M	Bequest intensity, single men (10^6)	0.253
		[0.146, 0.262]
a^M	Bequest curvature, single men (10^6)	5.660
		[5.130, 5.702]
α	Altruism toward surviving spouse	0.121
		[0.083, 0.385]
β_0	Pareto weight – level	3.556
		[3.503, 3.557]
β_I	Pareto weight – PI	-0.606
		[-0.625, -0.603]
β_z	Pareto weight – income share	9.061
		[9.009, 9.333]
c_{nh}	Monthly consumption (10^2)	10.137
		[10.097, 10.273]

Notes: We report posterior medians as our point estimates, with 90% credible intervals shown in brackets.

gender. To interpret the magnitude, we follow the practice in the literature and compute the asset threshold below which individuals do not leave any bequests when death in the next period is certain; this threshold is $\underline{a}^W$ = \$78,842 for women and $\underline{a}^M$ = \$44,767 for men. These values are higher than what is typically found in the literature, suggesting our model predicts weaker bequest motives.[44] The higher bequest motive toward children that we find for men is consistent with the wealth profiles in Figure 2, where single men in the high–PI group decumulate their wealth at a noticeably slower pace than single women. It is also consistent with the regression results in Table 3, which show that among high–PI single individuals, women save less than men. Together, these findings suggest that for low- and

[44]For instance, De Nardi, French, and Jones (2010) find values of \$31,500-\$43,000, and De Nardi, French, Jones, and McGee (2025) find around \$30,000.

Table 8: Wife's Pareto weight estimates

	Low z	Middle z	High z
Low PI	0.2590	0.5199	0.7886
Middle PI	0.2017	0.4391	0.7294
High PI	0.1488	0.3513	0.6510

mid-PI groups, women's stronger precautionary saving motives lead them to save more than men, whereas for high–PI individuals, women's weaker bequest motives lead them to save less.

To the best of our knowledge, the literature provides no direct estimates of altruism toward surviving spouse. Most structural bequest models with couples either (i) treat the household as a single decision maker with a joint bequest motive, or (ii) focus on transfers to children and abstract from a distinct altruistic component toward the surviving spouse. In contrast, our model delivers an estimated spousal altruism parameter around 0.12. This is smaller than the altruism toward children typically found in the literature: for instance, Lee and Seshadri (2019) find 0.32, Gayle, Golan, and Soytas (2022) obtain 0.80, and Caucutt and Lochner (2020) find 0.86.

The negative value of β_I implies that wives' bargaining power decreases with PI. This is similar to Lise and Seitz (2011), who estimate a sharing rule within a collective model and find that women's consumption share decreases with household income. It is also consistent with Blundell, Chiappori, Magnac, and Meghir (2007), who show that as total resources increase, a larger share of the additional income accrues to the husband, thereby reducing the wife's relative share. The estimate of β_z is positive, suggesting that women's bargaining power increases with their income share. While the literature typically fixes the level of Pareto weights by assuming equal weights for certain households (e.g., those with mean income and equal shares), we estimate this value by exploiting changes in consumption around nursing home entry. Table 8 reports the wife's bargaining power conditional on PI and her income share. The mean λ is 0.45 in our estimation sample, which is higher than what is found for the working-age population. For instance, Voena (2015) finds approximately 0.30 for the wife's Pareto weight by leveraging employment responses to changes in U.S. divorce laws. The table shows that the Pareto weight varies more steeply across z than across PI. There is substantial cross-sectional heterogeneity in women's bargaining power, which we further discuss in Section 5.1.

Our estimated monthly consumption floor is \$1,014. This is close to Mommaerts (2016) who find an annual floor around \$12,000, but higher than De Nardi, French, Jones, and McGee (2025) who find an annual floor around \$4,100.

Table 9: Fit of of auxiliary regression coefficients

Dependent variable	Regressor	Data	Model
Auxiliary OLS I			
Consumption rate	Wife's income share	-0.81	-0.77
Consumption rate	PI (mid) × Wife's income share	0.62	0.19
Consumption rate	PI (high) × Wife's income share	0.72	0.76
Auxiliary OLS II			
Consumption rate	Spouse enters NH × Spouse is husband	-0.28	-0.27

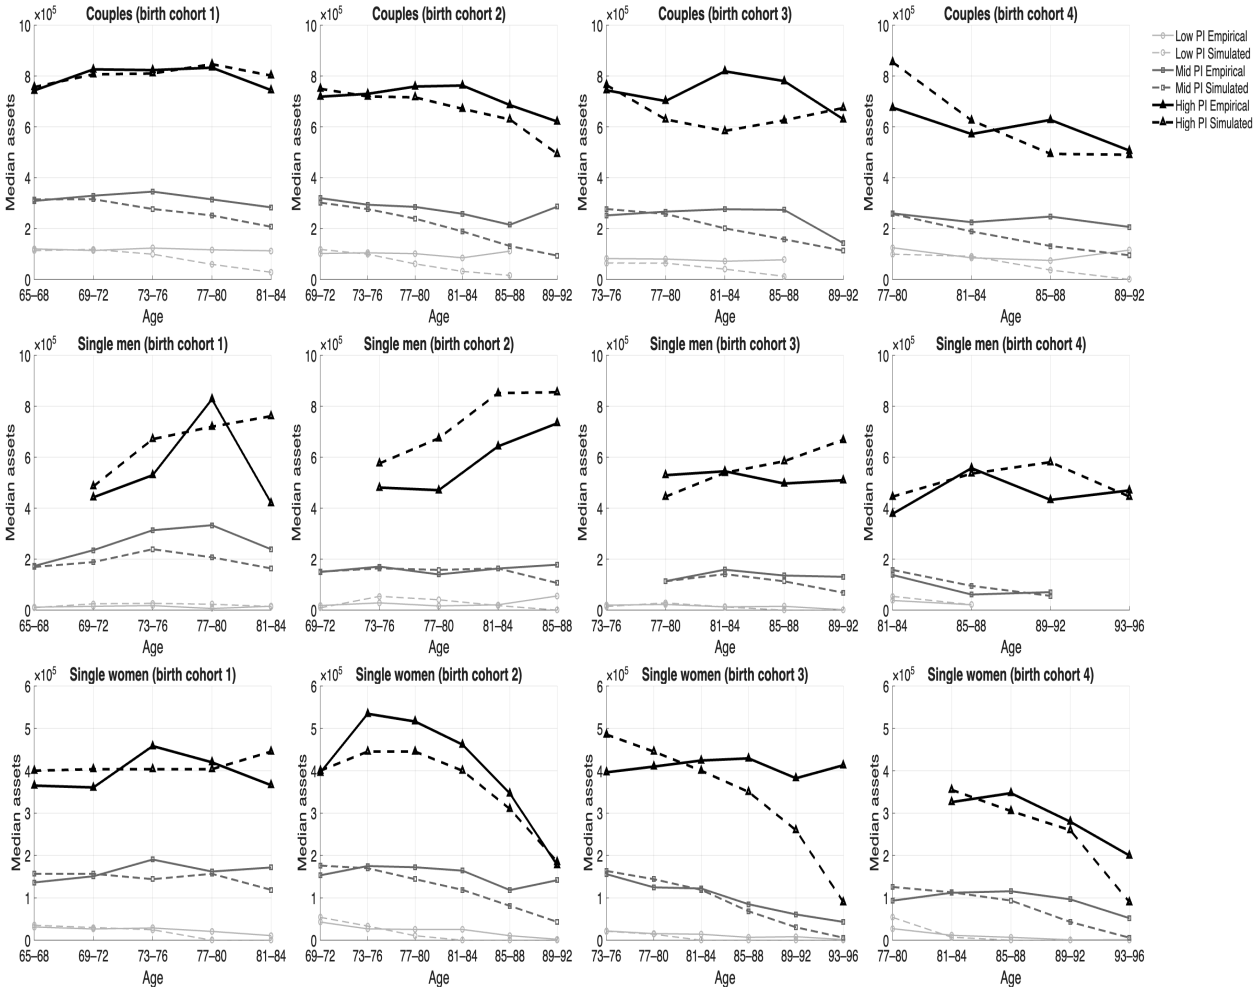
Table 10: Fit of auxiliary moments (rates)

Moment	Data	Model
Spousal caregiving rates		
By wife	0.76	0.66
By husband	0.74	0.87
Medicaid recipiency rate		
Couples	0.07	0.03
Single males	0.14	0.08
Single females	0.18	0.18
Homeownership rate of single males		
Low PI	0.41	0.26
Mid PI	0.61	0.49
High PI	0.74	0.68
Homeownership rate of single females		
Low PI	0.47	0.20
Mid PI	0.64	0.48
High PI	0.76	0.70

4.3 Model fit

Targeted objects Tables 9 and 10, together with Figure 2, summarize the fit of all auxiliary models used in the indirect inference estimation. The estimated model replicates most auxiliary regression coefficients that are informative for identifying the Pareto weights, although it overstates the effect of the wife's income share on lowering consumption for mid-PI households. The model matches Medicaid recipiency rates well for single women, but understates these rates for single men and couples. Figure 2 shows that the model provides a reasonable approximation to wealth evolution conditional on birth cohort, current marital status, and PI.

Figure 2: Fit of auxiliary moments (wealth)



Notes: The figure reports empirical (solid lines) and simulated (dashed lines) wealth conditional on birth cohort, current marital status, PI, and age. Darker colors indicate higher PI. Birth cohorts are defined based on household ages in 1998.

Untargeted objects For external validation, we also assess the model's fit for several untargeted statistics. For instance, regressing the consumption rate on the wife's income share (without conditioning on PI-group heterogeneity) on the simulated data yields a coefficient of –0.39, which is close to the empirical coefficient of –0.34 shown in Table 1 of Section 2. We also compare the model to empirical long-term care arrangements reported in Table 2 of Section 2. Some categories are intentionally omitted because they are not feasible within the model. For example, when a married individual becomes sick and the spouse is healthy, the model allows only spousal care or nursing home care, not care provided by children; and the model does not allow a disabled spouse to provide care.[45] Even with the restriction that

[45]These assumptions are empirically grounded in the data patterns, as justified in Section 3.

Table 11: Fit of untargeted long-term care arrangements

	Married		Single
	Spouse healthy	Spouse disabled	
	(1)	(2)	(3)
Caregiving by spouse	0.75	—	—
	[0.75]		
Caregiving by children	—	0.55	0.63
		[0.49]	[0.63]
Nursing home care	0.25	0.45	0.37
	[0.26]	[0.24]	[0.38]

Notes: Entries without brackets report model-simulated rates. Entries in square brackets report empirical rates from Table 2 .

only one type of care can be used at a time, the estimated model performs reasonably well in matching care arrangements by marital status.

5 Main results

5.1 Intrahousehold inequality in retirement

This section uses the estimated model to quantify intrahousehold inequality in retirement. We construct initial conditions from the 1998 and 2000 HRS waves using households in which the head is aged 65–69, and then simulate their choices forward over the life cycle. We also construct initial conditions from the 2012 and 2014 waves. Simulations based on the two cohorts would reveal how intrahousehold inequality changed between couples who entered retirement in 1998–2000 and those who did so 14 years later in 2012–2014.

For each couple, we compute the wife's share of household consumption over the marriage using non-housing consumption, c_t^g, as in equation (1). We exclude housing consumption (h_t^g) because it is mostly a public good.[46] We use the model-generated consumption paths for the husband and wife during periods in which both spouses are alive as a couple. We then compute the present-discounted value of each spouse's non-housing consumption over these periods. The wife's share is defined as the ratio of her present-discounted consumption to the household total.

[46]The exception is when one spouse enters a nursing home. However, even in that case intratemporal allocation decisions do not vary with bargaining power as the nursing home spouse receives a fixed level of housing services. Including housing would mechanically dampen differences and obscure consumption inequality within the household.

Table 12: Composition of (PI, z) bins by cohort

Old cohort

	Low z	Middle z	High z	Total (PI)
Low PI	0.1069	0.0833	0.0706	0.2608
Middle PI	0.1445	0.0995	0.1196	0.3636
High PI	0.1095	0.1022	0.1640	0.3757
Total (z)	0.3609	0.2849	0.3542	1.0000

Young cohort

	Low z	Middle z	High z	Total (PI)
Low PI	0.0826	0.0622	0.1021	0.2470
Middle PI	0.1047	0.0451	0.1192	0.2690
High PI	0.1459	0.0892	0.2489	0.4840
Total (z)	0.3332	0.1966	0.4702	1.0000

Notes: Each panel reports shares in (PI, z) bins. The last column shows totals by PI and the last row shows totals by z.

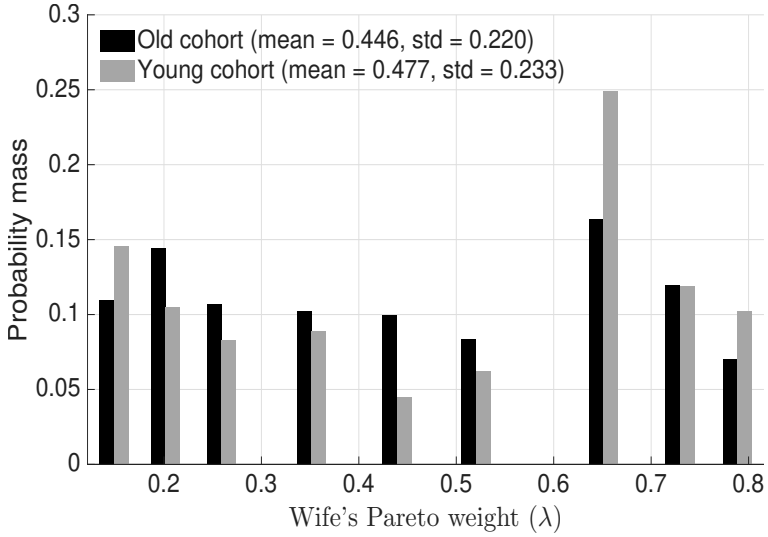
Table 12 provides useful background by reporting the composition of PI and wife's income share (z) groups for the two cohorts.[47] The share classified as low PI decreased only slightly from 26% to 25%, while the middle-PI share declined by almost 10 percentage points and the high-PI share increased by a similar amount. Because PI is negatively associated with women's bargaining power, this shift in the PI distribution can exacerbate consumption inequality within married households. For the wife's income share, the low-share group decreased modestly from 36% to 33%, the middle-share group fell by about 9 percentage points, and the high-share group rose by 12 percentage points. Because a higher income share raises a wife's Pareto weight, the increase in the high-share group can, in contrast, reduce within-household consumption inequality. Table 8 shows that variation in women's Pareto weights is larger across z groups than across PI groups, suggesting that women's bargaining position is likely to have improved.

Figure 3 plots the distribution of wives' Pareto weights for the two cohorts. The younger cohort exhibits a rightward shift, with the mean rising from 0.446 to 0.477. This pattern is consistent with the changes in the composition of z groups documented earlier. However, the younger cohort also shows slightly greater dispersion, with the standard deviation increasing from 0.220 to 0.233. This aligns with Table 12, which reports a higher share of High-PI/Low-z households from the younger cohort, the group in which the wife's bargaining weight is lowest.

Figure 4 illustrates how differences in Pareto weights translate into realized consumption allocations. For the older cohort, the mean is 0.482, implying that women who entered retirement in 1998–2000 received about 48.2% of household consumption over the remainder

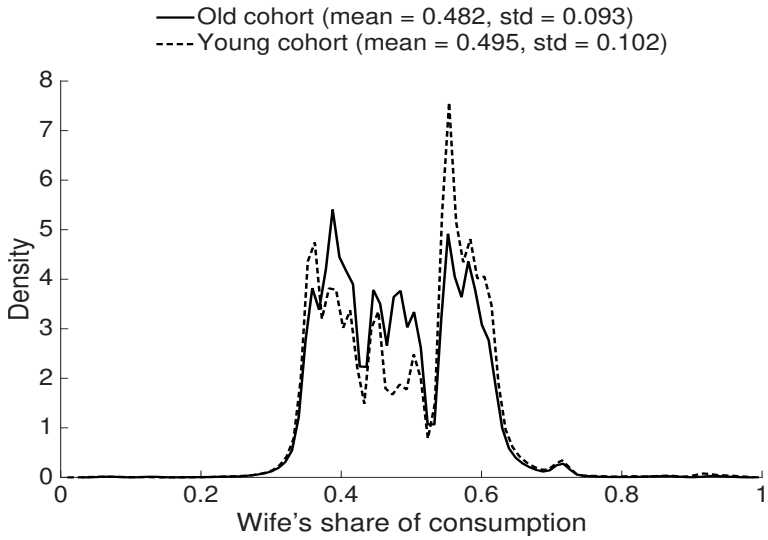
[47] As described in Section 2.2, we define the tercile cutoffs using the HRS estimation sample, which includes the older cohort shown in the left panel of Table 12. We then apply these same cutoffs to the younger cohort, shown in the right panel, who were too young to be included in the estimation sample, to ensure consistent classification into three permanent income and three wife's income share groups.

Figure 3: Distribution of wife's Pareto weight



Notes: The solid line plots the distribution of wives' Pareto weights for couples in which the husband was aged 65–69 in the 1998–2000 HRS waves ("Old cohort"). The dashed line shows the corresponding distribution for couples with husbands aged 65–69 in the 2012–2014 waves ("Young cohort").

Figure 4: Distribution of wife's share of consumption



Notes: The solid line plots the distribution of the wife's share of consumption over the marriage for couples in which the husband was aged 65–69 in the 1998–2000 HRS waves ("Old cohort"). The dashed line shows the corresponding distribution for couples with husbands aged 65–69 in the 2012–2014 waves ("Young cohort").

of the marriage. For the younger cohort, the distribution shifts slightly to the right, with the mean rising to 0.495.[48] Although the shift is modest, it indicates that wives in the more recent cohort receive a larger share of household consumption. Similarly with the Pareto weights, the younger cohort exhibits a greater dispersion, with the standard deviation increasing modestly form 0.093 to 0.102.

To summarize, wives in the more recent retirement cohort have, on average, a better bargaining position, though inequality across households has widened.

5.2 Intrahousehold heterogeneity in savings motives

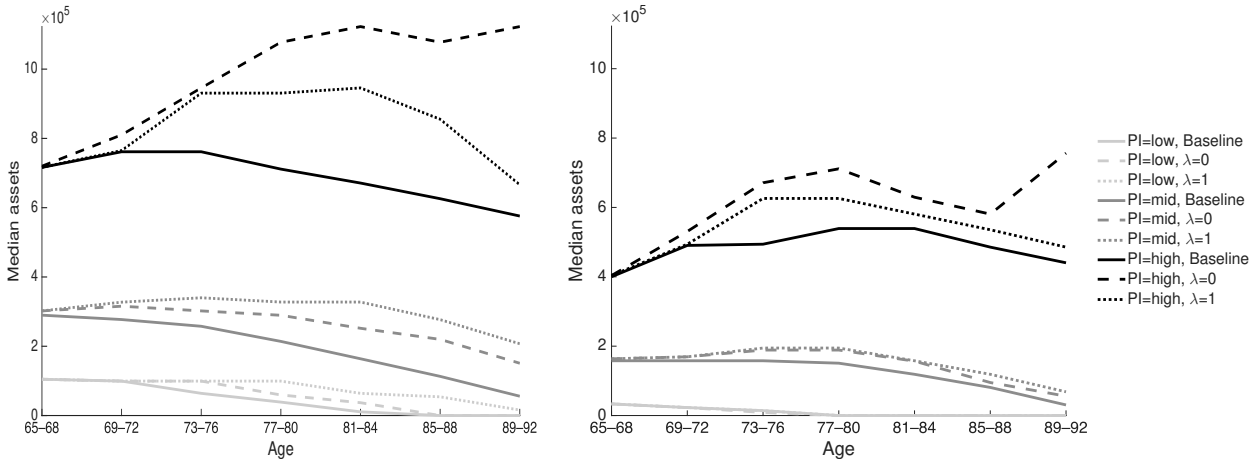
We now examine how much couples' observed savings reflect the wife's motives versus the husband's. While the existing literature has studied different savings motives (precautionary versus bequest motives) across household types and across the PI distribution, it has not been able to quantify intrahousehold heterogeneity in savings motives between spouses. This is infeasible in a unitary model, which treats the couple as having a single set of preferences. Our estimation results indicate meaningful heterogeneity in savings motives by gender. Women have stronger precautionary motives because they face longer longevity and higher long-term care risks, whereas bequest motives toward children, relevant primarily for high-PI retirees, are weaker for women than for men. How these differences translate into couples' observed savings depends on the Pareto weights, which have an estimated mean around 0.45. To isolate each spouse's savings motives, we conduct counterfactuals in which either the wife or the husband receives the entire Pareto weight.[49]

Figure 5 reports how median asset profiles change under these counterfactuals, conditional on current marital status, PI, and age. The left panel shows the wealth profiles of currently married individuals. When either spouse receives all bargaining power, couples' savings increase across all PI groups relative to the baseline. Changes in spousal care provision and the resulting increase in long-term care expenditure risk are central to this pattern. In both counterfactuals, the spouse whose utility is weighted never provides spousal care because the utility cost of caregiving is substantial. Although the unweighted spouse increases their informal care, the couple's overall reliance on formal long-term care services rises. As shown in Table 11, spousal care is an important source of insurance for both genders within a married household. Therefore, when one spouse essentially shuts down care provision, the

[48] For comparison, Lise and Seitz (2011) estimate that in 2000, working-age women in childless U.K. couples had a consumption share of 0.40. To our knowledge, there are no comparable estimates for retirees.

[49] We use initial conditions from the 1998–2000 HRS waves ("Old cohort" in Table 12), who were also part of the model estimation sample.

Figure 5: Counterfactual wealth trajectories



Couples Singles

Notes: The left panel reports results for currently married retirees, and the right panel reports results for currently single retirees. Each panel plots median assets by PI and age, with darker colors indicating higher PI. Solid lines show the baseline; dashed lines show the counterfactual in which the husband's Pareto weight is one ($\lambda = 0$); and dotted lines show the counterfactual in which the wife's Pareto weight is one ($\lambda = 1$).

household's exposure to long-term care expenditures increases. This strengthens precautionary saving motives, especially for high-PI households whose probability of qualifying for means-tested Medicaid is low. Consequently, the increase in savings is the greatest in the high-PI group under both counterfactuals.

Which counterfactual leads to a larger increase in savings varies across PI groups. For the low- and mid-PI groups, savings increase more when the wife holds all bargaining power (the dotted lines lie above the dashed lines). In these groups, bequests are less relevant, so precautionary motives dominate, which are stronger for women than men. As a result, when the wife's preferences are fully weighted, low- and mid-PI households save more than when the husband's preferences are fully weighted. In contrast, for high-PI households, savings rise more when the husband has all bargaining power (the dashed line lies above the dotted line). This pattern is driven by men's stronger bequest motives toward children, an effect that becomes more pronounced as the household ages.

The right panel in Figure 5 shows the corresponding results for individuals who are currently single. As couples accumulate more assets under both counterfactuals, surviving partners enter singlehood with higher wealth, and this effect grows with age as more households experience the transition from couplehood to singlehood. As a result, average assets

among singles rise systematically with age in these simulations.

5.3 Welfare effects of Medicaid housing rules

This section considers alternative treatments of housing assets by Medicaid's means test. Our measure of welfare is the compensating wealth variation (CWV), the initial wealth transfer that makes the economic agent indifferent to the counterfactual regime. Online Appendix F.1 provides details on how we compute the CWV. Importantly, the collective model uncovers spouse-specific welfare effects that a unitary model would mask, revealing gender-differentiated policy consequences hidden by relying solely on household-level CWV.

For the counterfactuals, we vary the recipient of the homestead exemption used in Medicaid's means test, which is granted only to couples in the baseline. To measure the fiscal impact, we compute the change in the present-discounted value of government expenditures over households' life cycles. Using the 1998 and 2000 HRS waves, we construct initial conditions for households entering retirement, both initially single and initially married, with a household head aged 65–69.

First, we study the welfare impact when Medicaid no longer provides the homestead exemption to couples in a nursing home, with welfare impacts reported in the first column of Table 13.[50] Households entering retirement as married are, on average, worse off by $658. By gender, wives lose about $800, while husbands' loss is $569. By allowing couples to remain homeowners longer, the homestead exemption slows dissaving and is particularly valuable for women, whose longer life expectancy strengthens their precautionary savings motives. Table F.3 in the Online Appendix reports welfare impacts for initially married couples conditional on PI, z, and gender. The losses are concentrated among low- and middle-PI households for whom means-tested Medicaid is most relevant. The largest losses occur for households with low PI and high z, where the wife's bargaining power is highest. In these households, women are less likely to provide spousal care, so when husbands become disabled they are more likely to enter nursing homes. Without the homestead exemption, these households must spend down their housing assets before qualifying for Medicaid, which substantially worsens their welfare.

Initially singles are not affected by the policy because they did not qualify for the exemption in the baseline. The household-level average welfare loss is $379. When we instead treat single individuals and each spouse within a couple as separate persons, the resulting

[50] Table F.1 in the Online Appendix reports the corresponding changes in Medicaid recipiency conditional on current marital status and gender.

individual-level average welfare loss is higher at $500. This difference is in part driven by wives' larger losses, which receive less weight in the household aggregation due to their lower Pareto weights on average. The savings in Medicaid spending amount to about $630 per household.

Next, we extend the homestead exemption to singles, such that all retirees' housing assets are disregarded by Medicaid's means test for nursing home use. Results are reported in Column (2) of Table 13. The more lenient version of Medicaid generates a positive yet modest impact for those who enter retirement as married. For initially singles, the welfare gains are substantial: $3,777 for women and $1,785 for men. The larger gains for women reflect the fact that the exemption lets them retain homeownership longer, which is particularly valuable given their longer life expectancy and hence stronger incentives to save. Table F.2 in the Online Appendix reports the welfare impact for initially singles by PI and gender, showing again that the effects are concentrated among low- and middle-PI households. The average welfare gain is about $1,410 per household, and the additional Medicaid spending is $1,240 per household.

Finally, in the third column of Table 13, we reform Medicaid so that the homestead exemption is granted to singles in a nursing home but not to couples. The impact on initially singles is the same as in the previous counterfactual, as they are newly granted the homestead exemption under both counterfactuals. For initially couples, the average welfare loss is nearly as large as when the homestead exemption is removed altogether, suggesting that retirees who enter retirement as married place a higher value on retaining the exemption early in retirement, while they are still married. As before, the welfare loss is larger for initially married women than for men. The welfare impact is positive on average ($1,032 at the household level and $406 per person), reflecting initially singles' greater valuation of the homestead exemption relative to couples. Singles lack access to spousal care, leading to higher formal long-term care risks. This makes the homestead exemption, which preserves their housing assets while qualifying for public long-term care insurance, more valuable than it is for couples.

Taken together, these counterfactuals show that Medicaid's housing rules generate heterogeneous welfare consequences, with particularly strong effects for initial singles, for women, and for low-PI retirees. Importantly, our collective approach reveals meaningful within-household differences, with wives typically experiencing larger welfare changes due to their stronger precautionary savings motives. These gender disparities would be invisible under a unitary approach, underscoring the value of a collective model for evaluating policy reforms.

Table 13: Welfare effects of counterfactual Medicaid

Homestead exemption:	(1) Nobody	(2) Everybody	(3) Singles
Initially couples	-658	38	-618
: *Wives*	-800	59	-798
: *Husbands*	-569	-3	-502
Initially single women	0	3777	3777
Initially single men	0	1785	1785
All (household-weighted)	-379	1410	1032
All (individual-weighted)	-500	901	406
Δ Medicaid	-631	1240	562

Notes: For each counterfactual in the columns, the table reports mean CWV for couples entering retirement as married, broken down by gender (wives and husbands), and for individuals who are initially single, also distinguished by gender. Δ Medicaid is the dollar change in per-household Medicaid spending relative to the baseline. The row All (household-weighted) reports the mean CWV evaluated at the household level. The row All (individual-weighted) reports the mean CWV evaluated at the individual level, weighting each married and single retiree equally.

6 Conclusion

While many retirees spend a substantial share of retirement as married, little is known about intrahousehold bargaining dynamics in old age. This paper provides one of the first quantitative assessments of cross-sectional heterogeneity in gender inequality among retired couples. By estimating a collective model of couples in retirement, we find that women's bargaining positions are higher than values typically estimated for working-age households and have strengthened on average, pushing their consumption shares closer to parity. At the same time, dispersion in bargaining power across households has widened. Beyond cross-sectional inequality, we find meaningful intrahousehold heterogeneity in savings motives: wives on average have stronger precautionary incentives to save due to longer life expectancy and higher long-term care risk. Consequently, policies that encourage faster dissaving of retirement wealth disadvantage wives relative to husbands, revealing distributional consequences that would be obscured under a unitary framework.

References

Achou, B., 2021, Housing in Medicaid: Should it Really Change?, *American Economic Journal: Economic Policy* forthcoming.

Ameriks, John, Andrew Caplin, Steven Laufer, and Stijn Van Nieuwerburgh, 2011, The Joy of Giving or Assisted Living? Using Strategic Surveys to Separate Public Care Aversion from Bequest Motives, *The Journal of Finance* 66, 519–561.

Apps, Patricia, and Ray Rees, 1988, Taxation and the Household, *Journal of Public Economics* 35, 355–369.

Barczyk, D., and M. Kredler, 2018, Evaluating Long-Term-Care Policy Options, Taking the Family Seriously, *Review of Economic Studies* 85, 766–809.

Barczyk, D., M. Kredler, and S. Fahle, 2022, Save, Spend or Give? A Model of Housing, Family Insurance, and Savings in Old Age, *Review of Economic Studies* forthcoming.

Blundell, Richard, Pierre-Andre Chiappori, Thierry Magnac, and Costas Meghir, 2007, Collective Labour Supply: Heterogeneity and Non-Participation, *Review of Economic Studies* 74, 417–445.

Blundell, Richard, Pierre-André Chiappori, and Costas Meghir, 2005, Collective Labor Supply with Children, *Journal of Political Economy* 113, 1277–1306.

Braun, R. A., K. A. Kopecky, and T. Koreshkova, 2017, Old, Sick, Alone, and Poor: A Welfare Analysis of Old-Age Social Insurance Programmes, *Review of Economic Studies* 84, 580–612.

Brown, J. R., and A. Finkelstein, 2008, The Interaction of Public and Private Insurance: Medicaid and the Long-Term Care Insurance Market, *American Economic Review* 98(3), 1083–1102.

Browning, Martin, François Bourguignon, Pierre-André Chiappori, and Valérie Lechene, 1994, Income and Outcomes: A Structural Model of Intrahousehold Allocation, *Journal of Political Economy* 102, 1067–1096.

Browning, Martin, and Pierre-André Chiappori, 1998, Efficient Intra-household Allocations: A General Characterization and Empirical Tests, *Econometrica* 66, 1241–1278.

Browning, Martin, Pierre-André Chiappori, and Arthur Lewbel, 2013, Estimating Consumption Economies of Scale, Adult Equivalence Scales, and Household Bargaining Power, *The Review of Economic Studies* 80, 1267–1303.

Campbell, John Y., and Joao F. Cocco, 2007, How do house prices affect consumption? Evidence from micro data, *Journal of Monetary Economics* 54, 591–621.

Card, D., C. Dobkin, and N. Maestas, 2008, The Impact of Nearly Universal Insurance Coverage on Health Care Utilization: Evidence from Medicare, *American Economics Review* 98, 2242–2258.

Caucutt, Elizabeth, and Lance Lochner, 2020, Early and late human capital investments, borrowing constraints, and the family, *Journal of Political Economy* 128, 1065–1147.

Chiappori, Pierre-André, 1988, Rational Household Labor Supply, *Econometrica* 56, 63–90.

Chiappori, Pierre-André, 1992, Collective Labor Supply and Welfare, *Journal of Political Economy* 100, 437–467.

Chiappori, Pierre-André, Bernard Fortin, and Guy Lacroix, 2002, Marriage market, divorce legislation, and household labor supply, *Journal of Political Economy* 110, 37–72.

Chiappori, Pierre-Andre, and Maurizio Mazzocco, 2017, Static and Intertemporal Household Decisions, *Journal of Economic Literature* 55, 985–1045.

De Nardi, M, 2004, Wealth Inequality and Intergenerational Links, *Review of Economic Studies* 71, 743–768.

De Nardi, Mariacristina, Eric French, and Bailey Jones, John, 2016a, Savings after Retirement: A Survey, *Annual Review of Economics* 8, 177–204.

De Nardi, M., E. French, and J.B. Jones, 2010, Why Do the Elderly Save? The Role of Medical Expenses, *Journal of Political Economy* 118, 39–75.

De Nardi, Mariacristina, Eric French, and John Bailey Jones, 2016b, Medicaid Insurance in Old Age, *American Economic Review* 106, 3480–3520.

De Nardi, Mariacristina, Eric French, and John Bailey Jones, 2016c, Savings After Retirement: A Survey, *Annual Review of Economics* 8, 177–204 First published online August 8, 2016.

De Nardi, Mariacristina, Eric French, John Bailey Jones, and Rory McGee, 2025, Why Do

Couples and Singles Save during Retirement? Household Heterogeneity and Its Aggregate Implications, *Journal of Political Economy* 133, 750–792 Electronically published January 28, 2025.

de Ruijter, Esther, Judith K. Treas, and Philip N. Cohen, 2005, Outsourcing the Gender Factory: Living Arragements and Service Expenditures on Female and Male Tasks, *Social Forces*.

Department of Health and Human Services, 2005, Medicaid Treatment of the Home: Determining Eligibility and Repayment for Long-Term Care, *Medicaid Eligibility for Long-Term Care Benefits* Policy Brief 2.

Engen, Eric M., William G. Gale, and Cori R. Uccello, 1999, The Adequacy of Retirement Saving, *Brookings Papers on Economic Activity* 1999, 65–165.

Ernst, Richard, and Joel Hay, 1994, The US Economic and Social Costs of Alzheimer's Disease Revisited, *American Journal of Public Health* 84, 1261–1264.

Feinberg, Lynn, Susan Reinhard, Ari Houser, and Rita Choula, 2011, Valuing the Invaluable: 2011 Update The Growing Contributions and Costs of Family Caregiving, *Insight on the Issues* 51, AARP Public Policy Institute.

Fernandez-Villaverde, J., J. Rubio-Ramirez, and F. Schorfheide, 2016, *Handbook of Macroeconomics*vol. 2 chap. Solution and Estimation Methods for DSGE Models, pp. 527–724.

Finkelstein, A., and R. McKnight, 2008, What did Medicare do? The initial impact of Medicare on mortality and out of pocket medical spending, *Journal of Public Economics* 92, 1644–1668.

Gayle, George-Levi, Limor Golan, and Mehmet Soytas, 2022, What is the source of the intergenerational correlation in earnings?, *Journal of Monetary Economics* 129, 24–45.

Greenhalgh-Stanley, N, 2012, Medicaid and the Housing and Asset Decisions of the Elderly: Evidence from Estate Recovery Programs, *Journal of Urban Economics* 72, 210–224.

Gruber, J., and R. Martin, 2003, Precautionary Savings and the Wealth Distribution with Illiquid Durables, International Finance Discussion paper.

Hubbard, R. G, J. Skinner, and S.P. Zeldes, 1995, Precautionary Saving and Social Insurance, *Journal of Political Economy* 103, 360–399.

Hurd, M. D., 1989, Mortality Risk and Bequests, *Econometrica* 57, 779–813.

Kaplan, G., 2012, Moving Back Home: Insurance against Labor Market Risk, *Journal of Political Economy* 120, 446–512.

Kemper, P., H. Komisar, and L. Alecxih, 2005/2006, Long-Term Care Over an Uncertain Future: What Can Current Retirees Expect?, *Inquiry* 42, 335–350.

Ko, A., 2022, An Equilibrium Analysis of the Long-Term Care Insurance Market, *Review of Economic Studies* 89, 1993–2025.

Kopecky, K.A., and T. Koreshkova, 2014a, The Impact of Medical and Nursing Home Ex-

penses on Savings, *American Economic Journal: Macroeconomics* 6, 29–72.

Kopecky, Karen A., and Tatyana Koreshkova, 2014b, The Impact of Medical and Nursing Home Expenses on Savings, *American Economic Journal: Macroeconomics* 6, 29–72.

Lee, Sang Yoon, and Ananth Seshadri, 2019, On the intergenerational transmission of economic status, *Journal of Political Economy* 127, 855–921.

Ligon, Ethan, Jonathan P. Thomas, and Tim Worrall, 2002, Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies, *Review of Economic Studies* 69, 209–244.

Lise, Jeremy, and Shannon Seitz, 2011, Consumption Inequality and Intra-household Allocations, *Review of Economic Studies* 78, 328–355.

Lise, Jeremy, and Ken Yamada, 2019, Household Sharing and Commitment: Evidence from Panel Data on Individual Expenditures and Time Use, *The Review of Economic Studies* 86, 2184–2219.

Lockwood, L., 2018, Incidental Bequests and the Choice to Self-Insure Late-Life Risks, *American Economic Review* 108, 2513–2550.

Lundberg, Shelly J., Robert A. Pollak, and Terence J. Wales, 1997, Do Husbands and Wives Pool Resources? Evidence from the United Kingdom Child Benefit, *Journal of Human Resources* 32, 463–480.

McGee, R., 2019, Old Age Savings and House Price Shocks, *Working paper*.

MetLife, 2008, National Survey of Nursing Home and Assisted Living Costs, MetLife Mature Market Institute.

Mommaerts, C., 2016, Long-Term Care Insurance and the Family, Working paper.

Nakajima, M., and I. A. Telyukova, 2017, Reverse Mortgage Loans: A Quantitative Analysis, *Journal of Finance* 72, 911–950.

Nakajima, M., and I. A. Telyukova, 2020, Home Equity in Retirement, *International Economic Review* 61, 573–616.

Palumbo, M.G., 1999, Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle, *Review of Economic Studies* 66, 395–421.

Sauer, Robert, and Christopher Taber, 2021, Understanding women's wage growth using indirect inference with importance sampling, *Journal of Applied Econometrics* 36, 453–473.

Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun, 2006, Are Americans Saving "Optimally" for Retirement?, *Journal of Political Economy* 114, 607–643.

Skira, M., 2015, Dynamic Wage and Employment Effects of Elder Parent Care, *International Economic Review* 56, 63–93.

Venti, S. F., and D. A. Wise, 2004, Aging and Housing Equity: Another Look, *Perspectives on the Economics of Aging* pp. 127–175 University of Chicago Press.

Vermeulen, Frederic, 2002, Collective household models: Principles and main results, *Journal of Economic Surveys* 16, 533–564.

Vermeulen, Frederic, 2006, A Collective Model for Female Labour Supply with Nonparticipation and Taxation, *Journal of Population Economics* 19, 99–118.

Voena, Alessandra, 2015, Yours, Mine, and Ours: Do Divorce Laws Affect the Intertemporal Behaviorof Married Couples?, *American Economic Review*.

Warshawsky, Mark J., and Ross A. Marchand, 2017, Improving the System of Financing Long-Term Services and Supports for Older Americans, Working Paper.

Online Appendix (Not For Publication)

Title

Minsu Chang and Ami Ko

December 7, 2025

A Data and sample construction

This section provides details about the HRS sample, which is constructed from the biennial interview waves covering 1998-2014. We only consider individuals who were responsive and aged 65 and older in 1998; retired by 2000; and did not miss any interviews while alive. The main empirical sample consists of 6,475 respondents and 35,865 respondent-wave observations.

We measure housing assets as net value of primary residence.[51] We define homeowners as having strictly positive housing assets. We define renters as non-homeowners who are not residing in a nursing home. We measure non-housing assets as the sum of vehicles, businesses, IRA and Keogh accounts, stocks, mutual funds, investment trusts, checking, savings, money market accounts, CDs, bonds and T-bills.

The HRS asks respondents about their use of formal long-term care services including nursing homes and paid home care. It also asks about the receipt of informal long-term care, defined as unpaid help associated with managing daily tasks. We observe the identity of informal caregivers, most of whom are either spouses or adult children.[52]

We classify individuals as having long-term care needs if they report at least two limitations in activities of daily living (ADLs) or exhibit cognitive impairment. The availability of informal care provided by children is a dummy variable which is equal to one if a respondent says the number of children they believe will provide care when necessary exceeds zero.

Table A.1 reports the summary statistics of the sample used for descriptive analysis as well as structural estimation.

[51] Most retirees in our data have zero outstanding mortgages: the mean ratio of mortgage to housing value is 5%.

[52] Due to inconsistencies in questions in the 1998 wave, we only use data starting in 2000 to measure spousal caregiving.

Table A.1: Summary statistics of the HRS sample

	(1) Single			(2) Married		
	Mean	Median	Std.	Mean	Median	Std.
Age	80.53	81.00	6.68	77.41	77.00	5.90
Female	0.79	1.00	0.41	0.50	0.00	0.50
Housing assets	101606	44700	202082	187710	136000	256124
Non-housing assets	203415	23430	762834	502049	163155	1161399
Homeowner	0.59	1.00	0.49	0.88	1.00	0.33
Renter	0.31	0.00	0.46	0.11	0.00	0.31
Have LTC needs	0.24	0.00	0.43	0.13	0.00	0.33
Household income	31848	20444	63015	64367	45552	81832
Wife's income share	.	.	.	0.31	0.30	0.13
Believe children will provide LTC	0.74	1.00	0.44	0.76	1.00	0.43
Observations	18895			16970		

Notes: The HRS sample (1998–2014) used which consists of retired individuals aged 65 and older in 1998. Observations are at the individual–wave level.

B Permanent income construction

B.1 Baseline approach

We closely follow De Nardi, French, Jones, and McGee (2025) in specifying household income. We assume that log household income evolves according to:

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

where $\kappa(t, f_{it})$ is a flexible function of age t and family structure f_{it} (i.e., couple, single man, or single woman), and ω_{it} captures measurement error. The term I_i denotes the individual's percentile rank in the PI distribution, which is fixed over time. We begin by estimating a fixed effects model:

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

We then compute I_i as the percentile rank of the individual's sample mean of $\ln y_{it} - \hat{\kappa}(t, f_{it})$, using $\hat{\kappa}(t, f_{it})$ from the fixed effects estimates. Next, we estimate $h(I_i)$ by fitting:

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

where $h(I_i)$ is specified as a fifth-order polynomial in the estimated PI percentile. The function $\kappa(t, f_{it})$ is modeled as a third-order polynomial in age, with family structure dummies and interactions with age trends. The predicted income is then given by:

$$\ln \hat{y}_{it} = \hat{\kappa}(t, f_{it}) + \hat{h}(\hat{I}_i).^{53}$$

As in De Nardi, French, Jones, and McGee (2025), the index I_i does not vary with marital transition.

For our specification of Pareto weights, we also need to define the wife's share of the household income for couples. To relate individual income components to the household-level permanent income measure constructed above, we estimate the following household-level regression:

$$\hat{h}(\hat{I}_i) = \beta_0 + \beta_1 \text{income}_{\text{husband},i} + \beta_2 \text{income}_{\text{wife},i} + \epsilon_i,$$

where $\text{income}_{j,i}$ denotes the average income of household member $j \in \{\text{husband}, \text{wife}\}$ over

[53] Based on this approach, a regression of $\ln y_{it}$ on $\ln \hat{y}_{it}$ gives an R^2 of 0.728. Also, we find no evidence of positive serial correlation of the residuals from this regression (across any two waves, serial correlation is -0.034).

the sample period. Based on the estimated coefficients, we define the wife's income share as:

$$\frac{\hat{\beta}_2 \cdot \text{income}_{\text{wife},i}}{\hat{\beta}_1 \cdot \text{income}_{\text{husband},i} + \hat{\beta}_2 \cdot \text{income}_{\text{wife},i}}.$$

The mean of the resulting income share across households is 0.34.

B.2 Alternative approach

Alternatively, we could construct permanent income at the individual level by estimating $\ln \hat{y}_{it}$ separately for four family categories: married male, married female, single male, and single female. We then compute household permanent income as the sum of individual permanent incomes.

Specifically, we estimate the following model using individual-level income data y_{it}, incorporating both spouses for each married household:

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

where α_i denotes an individual fixed effect.

We then obtain predicted values as $\ln \hat{y}_{it} = \hat{\kappa}(t, f_{it}) + \hat{h}(\hat{I}_i)$, as in the baseline approach. Based on these estimates, we define the wife's share of household permanent income as:

$$\frac{\exp(\hat{h}(\hat{I}^{\text{wife}}))}{\exp(\hat{h}(\hat{I}^{\text{husband}})) + \exp(\hat{h}(\hat{I}^{\text{wife}}))}.$$

The correlation between the wife's income share from the baseline approach and the alternative method is 0.879. However, the alternative approach requires tracking an additional state variable in the model. For the sake of parsimony, we proceed with the household-level approach presented in the previous subsection.

C Descriptive analysis appendix

C.1 Homeownership around spousal death

We conduct an event study analysis to examine changes in the homeownership rate around spousal death. For this exercise, the sample consists of married individuals who were homeowners four years prior to their spouse's death and who remain alive four years afterward. From this sample, we estimate

$$Y_{it} = \sum_{k=-4}^{4} \beta_k I_{itk} + X_{it} + \tau_t + \varepsilon_{it} \tag{24}$$

where the dependent variable is one if individual i is still a homeowner and zero otherwise. I_{itk} represents an indicator for being k years since the spousal death. The vector X_{it} includes gender, birth cohort, an indicator for having a child, long-term care needs, income, and non-housing assets. τ_t represents year fixed effects. Figure C.1 reports the estimates of β_k. The results show that there is a statistically significant reduction in homeownership rates following spousal death.[54]

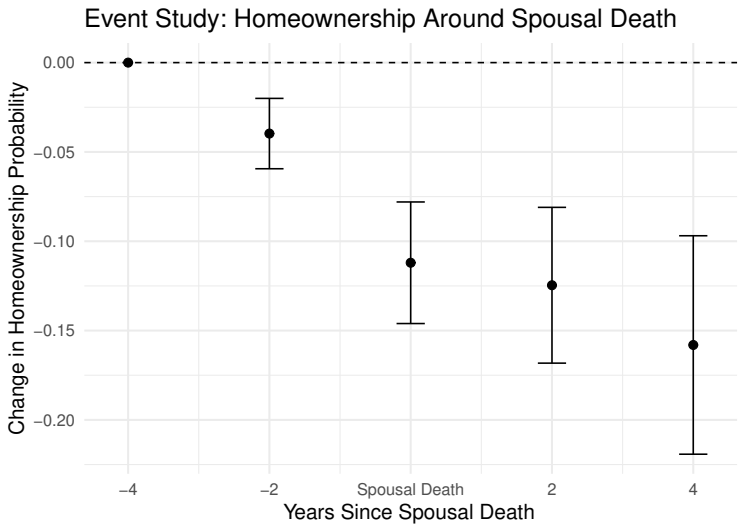
C.2 Medicaid's treatment of the home

When determining Medicaid eligibility, the home is an excluded resource regardless of its value (Department of Health and Human Services, 2005). However, if a recipient no longer has a living spouse in the home or an unmarried recipient moves into a nursing home without the intention to return, the house loses its protected status as it is no longer a "home".[55] In principle, there are two ways Medicaid can make recipients pay for nursing home expenses from their housing wealth when the house is no longer a home. The first is to deny coverage by counting the housing value against Medicaid eligibility. However, as individuals can easily obtain an exemption by stating their intent to return from a nursing home, in practice this rarely happens (Department of Health and Human Services, 2005). The second is to place liens on the home and recover Medicaid-paid long-term care costs from the recipient's estates upon death. The most popular component of the program is known as a TEFRA (Tax Equity and Fiscal Responsibility Act) lien, which is placed on the homes of permanently

[54] Table C.5 in the Appendix uses the same sample as in the event study and reports living arrangements four years after spousal death. Among those that liquidate their house and become non-homeowners, 67% move into rented units, 24.4% move into a nursing home, and 7.6% move into their children's home.

[55] Medicaid also counts the value of the home if it is transfered within a specified period (3-5 years) prior to applying for Medicaid long-term care.

Figure C.1: Homeownership around spousal Death



Notes: The sample consists of individuals who were homeowners four years prior to their spouse's death and who remain alive four years afterward. The figure reports the estimates of β_k in equation (24) along with their 95% confidence interval. Standard errors are clustered at the household level.

institutionalized beneficiaries to ensure that states can recover their costs from the estates (Greenhalgh-Stanley, 2012). Consistent with Medicaid's treatment of housing assets for couples, TEFRA liens cannot be placed when there is a surviving spouse living in the home (Greenhalgh-Stanley, 2012).

We estimate

$$Y_{ist} = \beta D_{ist}^{TEFRA \times Single} + \delta D_{st}^{TEFRA} + \eta D_{ist}^{Single} + \rho X_{ist} + \gamma_s + \phi_t + u_{ist} \tag{25}$$

The outcome variable is the homeownership status of individual i living in state s in year t. D_{st}^{TEFRA} is an indicator variable for whether a TEFRA lien is in place. D_{ist}^{Single} is an indicator variable for whether the individual is single. $D_{ist}^{TEFRA \times Single}$ is an indicator variable for whether the individual is single *and* a TEFRA lien is in place. The coefficient β therefore represents how TEFRA liens affect homeownership of singles relative to couples. We control for state and year fixed effects and include individual demographics contained in X_{ist}. We estimate equation (25) using restricted HRS data that contain geographic information.

Table C.1 reports the results. We find that TEFRA liens discourage homeownership among singles relative to couples. Specifically, a state's implementation of TEFRA liens reduces singles' homeownership rate by 8.9 percentage points. The results are consistent

Table C.1: TEFRA liens and homeownership

	Own home	
TEFRA x Single	-0.089	(0.036)
TEFRA	0.036	(0.054)
Single	-0.175	(0.025)
Mean of dependent variable	0.68	
Observations	1947	

Notes: Standard errors are in parentheses and are clustered at the household-state level. Restricted HRS data from years 1998-2014 used. Linear probability model is used. Not all controls shown including age, gender, number of children, income, non-housing assets, and year, state, and birth cohort fixed effects. The sample consists of respondents living in states that adopted TEFRA liens before and during the sample period.

with Greenhalgh-Stanley (2012) who also find that Medicaid estate recovery programs reduce singles' homeownership relative to couples.[56]

To further examine the relationship between Medicaid estate recovery programs and retirees' housing decisions, we use variations in estates recovered. We define the estate recovery rate as the ratio of estates recovered to total Medicaid nursing home expenses in a given state-year combination.[57] This is not our preferred empirical strategy as the estates recovered are endogenous to recipients' saving behavior. Nevertheless, we examine the correlation between the amount of estates recovered and homeownership. Table C.2 reports the results. While the estimates are less accurate, we find that a higher estate recovery rate is associated with a lower homeownership rate among singles, consistent with Table C.1.

Table C.2: Estate recovery rate and homeownership

	Own home	
ERP rate x Single	-0.065	(0.076)
ERP rate	-0.072	(0.083)
Single	-0.125	(0.159)
Age	-0.003	(0.004)
Female	-0.002	(0.057)
Have children	0.052	(0.105)
Income (in $100K)	0.058	(0.011)
Non-housing assets (in $100K)	0.003	(0.001)
Mean of dependent variable	0.66	
Observations	573	

Notes: Standard errors are in parentheses and are clustered at the household-state level. HRS 2002-2010 used. Linear probability model is used. Year, state, and birth cohort fixed effects are included. The sample consists of respondents living in states that have TEFRA liens in place.

Table C.3: Wife's income share and household consumption

	(1)		(2)	
	Consumption rate		Log of consumption	
Wife's income share	-0.34**	(0.16)	-0.29***	(0.11)
Husband age	-0.00	(0.01)	-0.01	(0.00)
Wife age	-0.02***	(0.01)	-0.02***	(0.00)
Housing assets from prev. wave (in 100K)	0.18***	(0.02)	0.14***	(0.01)
Liquid assets from prev. wave (in 100K)	0.02***	(0.01)	0.01***	(0.00)
Household income (in 100K)	0.18**	(0.08)	0.15***	(0.05)
PI (low)	0.00	(.)	0.00	(.)
PI (mid)	-0.79***	(0.10)	0.26***	(0.04)
PI (high)	-1.75***	(0.11)	0.30***	(0.05)
Husband has LTC needs	-0.10	(0.06)	-0.05	(0.05)
Wife has LTC needs	-0.00	(0.09)	-0.02	(0.05)
Number of children	-0.02**	(0.01)	-0.01	(0.01)
Husband's education years	0.02**	(0.01)	0.01***	(0.00)
Wife's education years	0.02	(0.01)	0.02**	(0.01)
Out-of-pocket medical spending	0.00**	(0.00)	0.00**	(0.00)
Husband has Medicare	-0.08	(0.24)	0.06	(0.08)
Wife has Medicare	-0.17	(0.25)	0.10	(0.06)
Husband has private LTC insurance	0.11	(0.09)	0.09	(0.06)
Wife has private LTC insurance	-0.07	(0.08)	-0.03	(0.05)
Health: Excellent	0.00	(.)	0.00	(.)
Health: Very good	-0.01	(0.06)	-0.01	(0.04)
Health: Good	-0.01	(0.07)	-0.01	(0.04)
Health: Fair	-0.04	(0.07)	-0.03	(0.04)
Health: Poor	-0.05	(0.10)	-0.03	(0.06)
Constant	3.83***	(0.60)	11.65***	(0.25)
Observations	1701.00		1701.00	
Adjusted R^2	0.40		0.42	
Mean of dependent variable	1.39		10.87	

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the household level (in parentheses). The table reports coefficient estimates not included in Table 1 of the main text. Not all controls are shown, including birth cohort fixed effects.

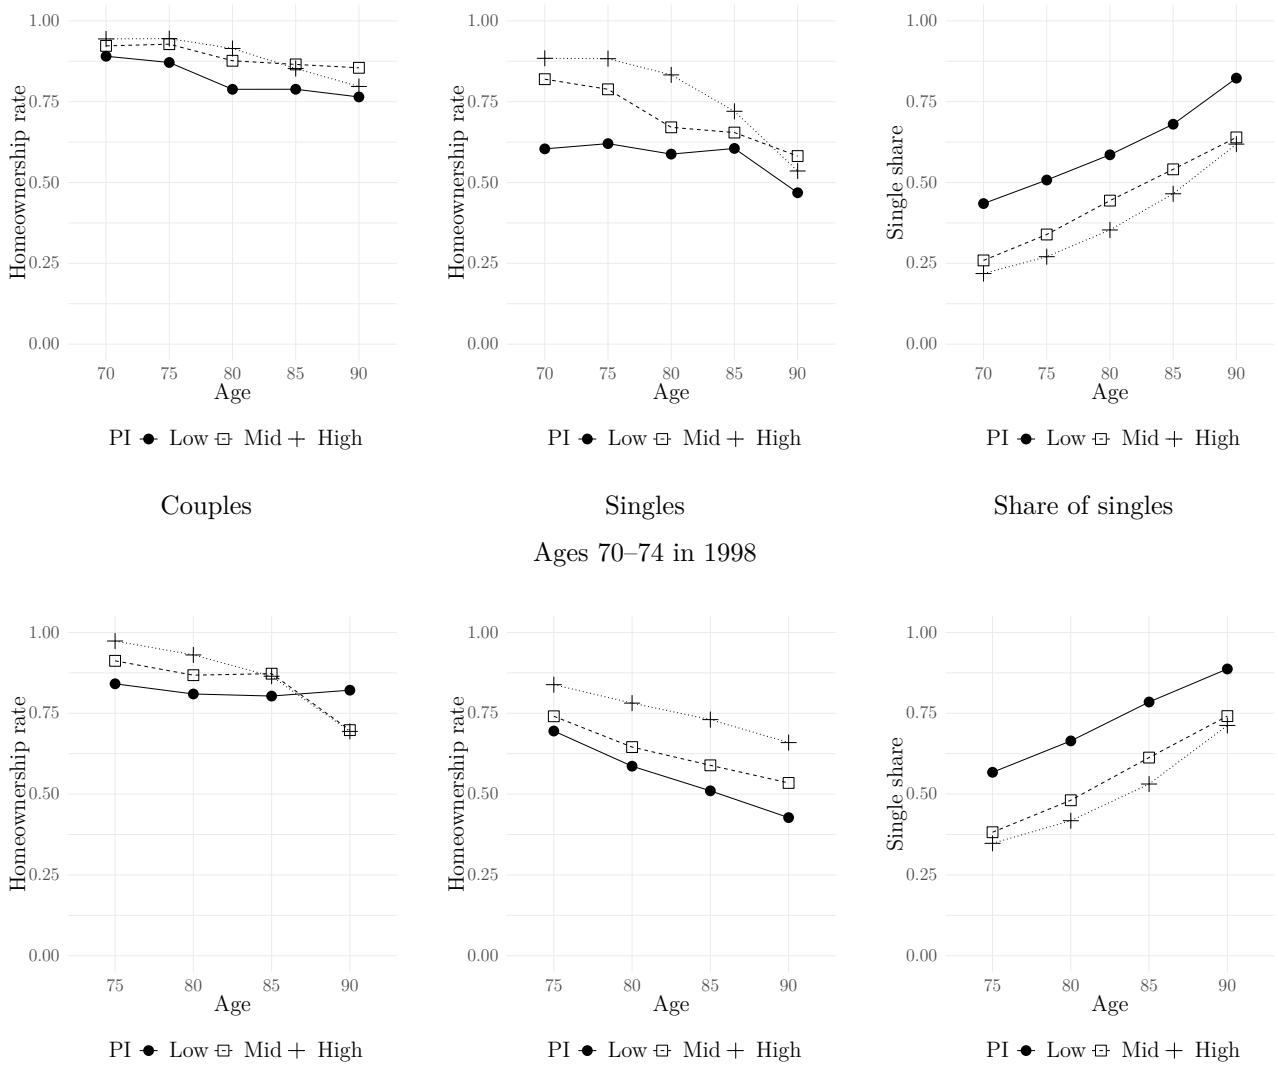
Table C.4: Wife's income share and household savings

	(1)		(2)	
	Savings rate		Change in assets	
Wife's income share	1.23**	(0.59)	79837.54***	(28990.34)
Husband age	-0.03*	(0.02)	-1922.44**	(856.61)
Wife age	-0.01	(0.02)	-232.00	(835.86)
Housing assets from prev. wave (in 100K)	-0.27***	(0.07)	-5897.52*	(3473.77)
Liquid assets from prev. wave (in 100K)	-0.56***	(0.04)	-33511.31***	(2151.49)
Household income (in 100K)	3.65***	(0.43)	201368.05***	(24643.85)
PI (low)	0.00	(.)	0.00	(.)
PI (mid)	0.41**	(0.19)	9027.10	(7429.20)
PI (high)	1.30***	(0.28)	44314.10***	(14125.98)
Husband has LTC needs	-0.14	(0.20)	-1740.91	(9892.96)
Wife has LTC needs	-0.52**	(0.24)	-20541.81*	(10725.23)
Number of children	-0.01	(0.03)	-1123.98	(1491.07)
Husband's education years	0.03	(0.02)	2425.58**	(1158.67)
Wife's education years	-0.07**	(0.04)	-2021.05	(1494.12)
Out-of-pocket medical spending	-0.00*	(0.00)	-0.45**	(0.19)
Husband has Medicare	-0.90*	(0.49)	-39867.14	(31201.72)
Wife has Medicare	0.98*	(0.55)	63235.67**	(30527.79)
Husband has private LTC insurance	0.72*	(0.39)	41086.54**	(17702.85)
Wife has private LTC insurance	-0.49	(0.32)	-29104.53*	(15692.82)
Health: Excellent	0.00	(.)	0.00	(.)
Health: Very good	-0.08	(0.22)	-12021.24	(13026.30)
Health: Good	-0.21	(0.21)	-20428.02*	(12233.97)
Health: Fair	-0.41*	(0.23)	-28743.98**	(12831.03)
Health: Poor	-0.45*	(0.26)	-39060.87***	(13511.32)
Constant	2.82**	(1.27)	124364.23*	(68125.82)
Observations	11008.00		11008.00	
Adjusted R^2	0.12		0.18	
Mean of dependent variable	-0.27		-13782.93	

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the household level (in parentheses). Not all controls are shown, including birth cohort fixed effects.

C.3 Additional tables and figures for descriptive analysis

Figure C.2: Homeownership rates

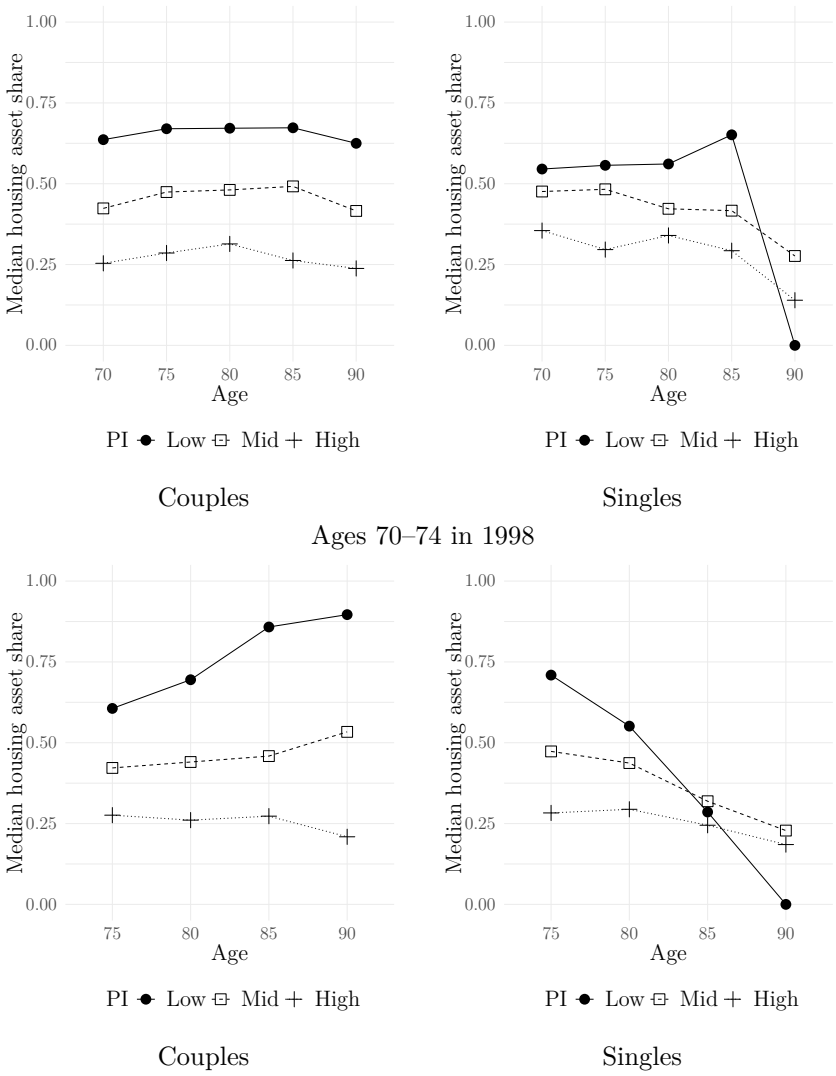


Notes: Compared to Figure 1, this figure conditions on individuals' age in 1998, focusing on those aged 70–74 and 75–79 at baseline. Although our sample includes both younger and older cohorts, we do not report them here due to their smaller sample sizes.

[56] According to Greenhalgh-Stanley (2012), the homeownership rate for singles in TEFRA lien states decreases by 3 percentage points compared to couples. This finding is based on analyzing the AHEAD cohort, which consists of elderly individuals aged seventy and older in 1993, and tracking them until 2004 or their death.

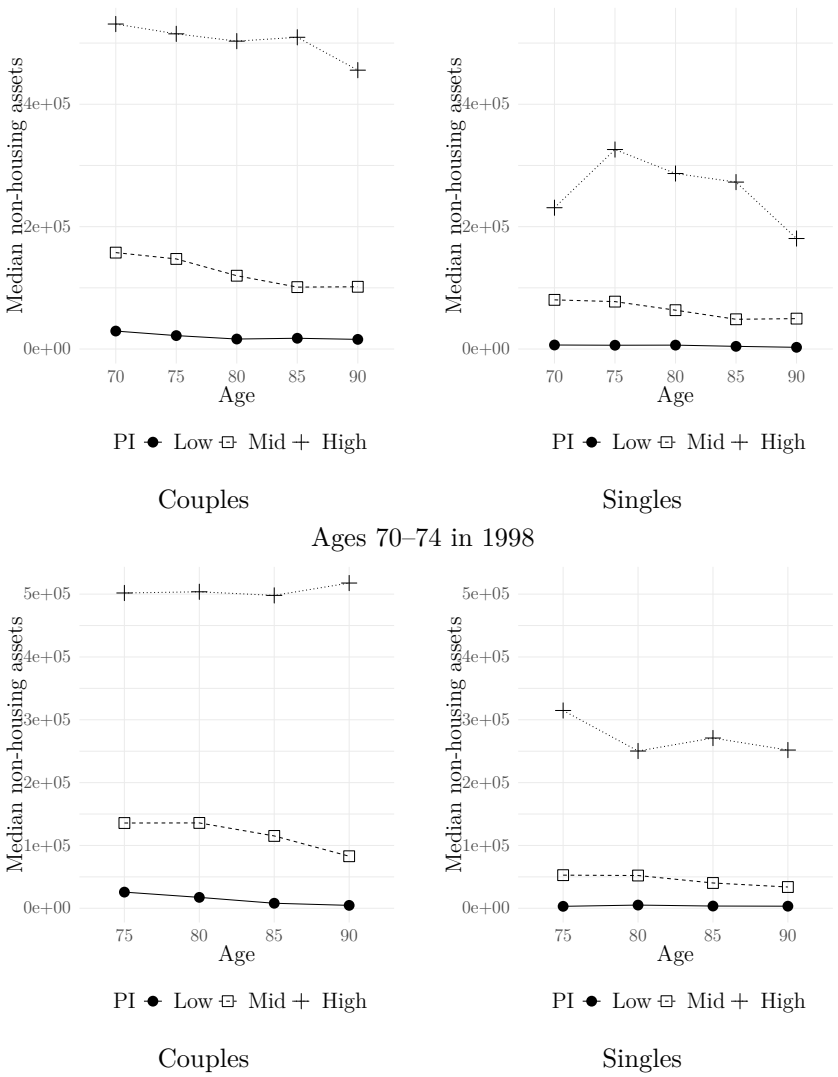
[57] We use Warshawsky and Marchand (2017) for information about the estates recovered. We obtain Medicaid nursing home expenses from Medicaid.gov. The estate recovery rate is computed for 2002-2010.

Figure C.3: Median housing asset share



Notes: The figure shows the median housing asset share by birth cohort, current marital status, permanent income terciles, and age group. Housing asset share is defined as the ratio of housing assets to total assets.

Figure C.4: Median non-housing assets



Notes: The figure shows the median non-housing assets by birth cohort, current marital status, permanent income terciles, and age group.

Table C.5: Living arrangements after spousal death

	Non homeowners	Still homeowners
Move into nursing home	0.248	0.080
Move into kids' home	0.083	0.017
Move into rented units	0.669	–
Observations	109	349

Notes: The table reports living arrangements four years after spousal death for the sample used in the event study analysis reported in Figure C.1.

Table C.6: Gender and singles' savings

	(1) Consumption rate	(2) Consumption rate	(3) Savings rate	(4) Savings rate	(5) Changes in assets	(6) Changes in assets
Female	-0.14***	-0.26**	0.17	0.29	1566.94	5481.49
	(0.04)	(0.11)	(0.12)	(0.18)	(5307.78)	(3667.26)
Age	-0.02***	-0.02***	-0.01	-0.01	-730.86*	-773.93*
	(0.00)	(0.00)	(0.01)	(0.01)	(417.19)	(416.27)
Housing assets from prev. wave (in 100K)	0.15***	0.15***	-0.40***	-0.40***	-10325.05***	-10283.45***
	(0.02)	(0.02)	(0.06)	(0.06)	(2405.14)	(2392.35)
Liquid assets from prev. wave (in 100K)	0.02***	0.02***	-0.72***	-0.73***	-40679.29***	-40749.71***
	(0.00)	(0.00)	(0.04)	(0.04)	(1753.09)	(1743.65)
Household income (in 100K)	0.13**	0.14***	3.79***	3.73***	222332.95***	219945.57***
	(0.05)	(0.05)	(0.33)	(0.33)	(20940.06)	(20962.48)
PI (low)	0.00	0.00	0.00	0.00	0.00	0.00
	(.)	(.)	(.)	(.)	(.)	(.)
PI (mid)	-0.62***	-0.79***	0.34***	0.21	6093.82	105.13
	(0.05)	(0.11)	(0.11)	(0.24)	(4071.39)	(7779.27)
PI (high)	-1.26***	-1.41***	1.83***	2.42***	59012.68***	80055.99***
	(0.07)	(0.12)	(0.18)	(0.36)	(8111.61)	(17201.46)
Have LTC needs	-0.08*	-0.07	-0.44***	-0.44***	-17652.66***	-17694.00***
	(0.05)	(0.05)	(0.11)	(0.11)	(4487.41)	(4479.06)
Number of children	0.01	0.01	-0.05***	-0.05***	-2524.39***	-2561.71***
	(0.01)	(0.01)	(0.02)	(0.02)	(789.25)	(790.23)
Education years	0.03***	0.03***	-0.00	-0.01	289.77	194.83
	(0.01)	(0.01)	(0.01)	(0.01)	(637.68)	(636.36)
Out-of-pocket medical spending	0.00**	0.00**	-0.00	-0.00	-0.01	-0.01
	(0.00)	(0.00)	(0.00)	(0.00)	(0.11)	(0.11)
Have Medicare	0.03	0.02	-0.30	-0.30	1484.56	1555.38
	(0.09)	(0.09)	(0.53)	(0.53)	(13543.05)	(13545.58)
Have private LTC insurance	0.00	-0.00	0.21	0.24	12771.73	13949.82*
	(0.03)	(0.03)	(0.16)	(0.16)	(8322.04)	(8335.73)
Health: Excellent	0.00	0.00	0.00	0.00	0.00	0.00
	(.)	(.)	(.)	(.)	(.)	(.)
Health: Very good	-0.00	-0.01	-0.12	-0.12	-9554.95	-9201.35
	(0.05)	(0.05)	(0.26)	(0.26)	(11280.09)	(11250.54)
Health: Good	-0.00	-0.00	-0.25	-0.25	-22657.27**	-22720.76**
	(0.05)	(0.05)	(0.25)	(0.25)	(10893.33)	(10880.59)
Health: Fair	-0.01	-0.01	-0.22	-0.22	-19513.75*	-19458.32*
	(0.05)	(0.05)	(0.26)	(0.26)	(10998.95)	(10981.93)
Health: Poor	0.08	0.07	-0.17	-0.17	-16736.13	-16915.86
	(0.08)	(0.08)	(0.28)	(0.28)	(11431.89)	(11435.04)
PI (low) × Female		0.00		0.00		0.00
		(.)		(.)		(.)
PI (mid) × Female		0.20*		0.19		8190.54
		(0.12)		(0.25)		(7719.08)
PI (high) × Female		0.18		-0.70**		-24923.02
		(0.12)		(0.33)		(17509.37)
Constant	2.78***	2.87***	0.87	0.90	63967.12*	65556.60*
	(0.28)	(0.29)	(0.84)	(0.84)	(34549.75)	(34443.33)
Observations	3530.00	3530.00	13692.00	13692.00	13692.00	13692.00
Adjusted R^2	0.26	0.26	0.14	0.14	0.22	0.22
Mean dep. var	1.08	1.08	-0.40	-0.40	-16073.37	-16073.37

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the household level (in parentheses). The table reports coefficient estimates not included in Table 3 of the main text. Sample consists of single respondents from HRS 1998-2014. Consumption questions were asked only to a randomly selected subsample. Not all controls are shown, including birth cohort fixed effects.

D Model appendix

D.1 Model Notation and Definitions

Table D.1 lists the key variables, functions, and parameters used in the model.

Table D.1: Model notation

Symbol	Definition
Choice variables	
$D \in \{0,1\}$	House selling choice for homeowners: keep (0) or sell (1)
$R \geq 0$	Rented housing service
$P^g \in \{0,1\}$	Spousal care from the member $g \in \{W, M\}$: no care (0) or care (1)
$\chi \geq 0$	Household consumption expenditure
State variables	
t	Household head's age
$a \geq 0$	Non-housing assets
$\tilde{h} \geq 0$	Housing assets. $\tilde{h} > 0$ implies homeowner, $\tilde{h} = 0$ renter.
s	Health status: healthy, require long-term care, or dead
I	Permanent income
ic_{child}	Availability of informal care from children: available (1) or not available (0)
Functions	
u^g	gender g-specific utility over general consumption and housing services
$v^{S,g}$	Bequest utility of single retiree of gender g
$v^{C,g}$	Bequest utility of married retiree of gender g
Utility parameters	
ψ^g	caregiver g's utility from care provision by member g
σ^g	Housing consumption utility scale of gender g
γ	Consumption and housing CRRA coefficient
δ^g, a^g	Bequest utility parameters that enter $v^{S,g}$
α	Altruism towards surviving spouse that enters $v^{C,g}$
Others	
$\beta_0, \beta_I, \beta_z$	Pareto weight parameters
c_{nh}, h_{nh}	Basic consumption and housing value from nursing home care
ρ	Economies of scale for married households' consumption
$\omega_{couple}, \omega_{single}$	Homeownership premium for couples and singles
δ	Depreciation rate for housing assets
r	Real interest rate
τ	Home transaction cost
m	Formal long-term care cost
$\bar{a}_{nh=0}$	Per-capita consumption floor for non-nursing home residents
$\bar{a}_{nh=1}$	Per-capita consumption floor for nursing home residents

D.2 Incorporation of side bequests and its challenges

In this appendix, we outline how side bequests could be incorporated into our model and explain why the available data do not allow credible identification without strong assumptions.

If a deceased individual is survived by a spouse, we assume that a fraction of the bequest is allocated to children, following De Nardi, French, Jones, and McGee (2025). The crucial difference is that in their unitary framework, the surviving spouse determines the size of the side bequest. Specifically, the surviving spouse chooses how much to allocate to children to maximize the sum of their utility from the side bequest and their continuation value as a single retiree. In contrast, our collective framework requires us to explicitly specify how both the deceased and the surviving spouses value side bequests. The amount of side bequests will be set while both spouses are alive, as the outcome of a collective decision-making process. Under our full commitment framework, we assume that these predetermined side bequests are enforceable, effectively treating them as if specified in a will.

We model the bequest preferences of a deceased individual who is survived by a spouse as

$$v^{C,g}(x_t) = v^{\text{dec},g}(\tau \cdot b_t) + \alpha^g V^{S,g'}((1-\tau) \cdot b_t), \tag{26}$$

where τ is the predetermined share of the bequest allocated to children, and b_t is the remaining household assets at the first spouse's death. The function $v^{\text{dec},g}$ represents the deceased individual's utility from leaving bequests to children when there is a surviving spouse. While these preferences may coincide with $v^{S,g}$, the utility from leaving bequests to children as a single individual, this is not obvious. As in the main text, we allow the deceased to care for the surviving spouse, captured via the pure altruism parameter, α^g. The surviving spouse's value depends on elements other than remaining assets after side bequests, which we suppress for notational simplicity.[58]

In our collective model, we also need to specify how the surviving spouse values side bequests. That is, if the husband dies first and a predetermined fraction τ of remaining wealth goes to children, the surviving wife may derive some utility from this transfer. From her perspective, however, this transfer is not a real bequest as she is still alive, so it more closely resembles an inter-vivos transfer. Therefore, we do not model her utility as $v^{S,g}(\cdot)$,

[58] While it is plausible that $v^{\text{dec},g}$ depends only on total assets, the surviving spouse's continuation value $V^{S,g'}$ depends on the composition of assets (liquid vs. housing). A full incorporation of side bequests would therefore require modeling how side bequests are funded. We abstract from this additional layer because our point is to demonstrate that preferences over side bequests are already difficult to identify even before accounting for asset composition.

which reflects bequest utility upon her own death. Nor is it obvious that her valuation of side bequests is identical to that of her deceased husband, $v^{\text{dec},g}$. We therefore denote the surviving spouse's preferences for side bequests in this context as $v^{\text{surv},g}$.

It is not feasible to identify the altruism parameter toward the spouse, α^g, alongside both $v^{\text{dec},g}$ and $v^{\text{surv},g}$ given the data we have. Even under the simplifying assumption that preferences for leaving bequests to children at one's own death are the same regardless of whether there is a surviving spouse, i.e., $v^{\text{dec},g} = v^{S,g}$, identification of α^g and $v^{\text{surv},g}$ remains infeasible. That is, we would not be able to tell how much of the side bequests is coming from the deceased's altruism for the surviving spouse vs. the surviving spouse's preference for side bequests (or inter-vivos transfers from her perspective).

If we impose a further assumption that the surviving spouse's preferences for side bequests are the same as her preferences for terminal bequests at her own death, $v^{\text{surv},g} = v^{S,g}$, then identification becomes feasible. Specifically, we can use the evolution of singles' savings to obtain information about $v^{S,g}$ and the evolution of couples' savings to identify altruism toward the spouse.[59] However, this is a strong assumption that is unlikely to hold in reality.[60] We therefore choose to abstract from side bequests. Instead, we focus on bequests left to the surviving spouse, which dominate empirically and allow us to study altruism toward the surviving spouse, a component only identifiable in a non-unitary framework.

D.3 Recursive formulation for singles' problem

We provide a recursive formulation for a single household's problem. In each period, a single household's state vector is given as

$$x_t = (a_t, \tilde{h}_{t-1}, s_t^g; I, ic_{child}) \tag{27}$$

where a_t is the non-housing wealth, $\tilde{h}_{t-1}$ is the housing wealth at the beginning of the period, and s_t^g is the health status of single gender $g \in \{W, M\}$. Time-invariant state variables are permanent income I and the availability of informal care from children ic_{child}.

The household's choice vector is $q_t = (\hat{c}_t, D_t, R_t)$ where $\hat{c}_t \geq 0$ is private non-housing consumption, $D_t \in \{0, 1\}$ represents the house selling choice, and $R_t \geq 0$ is the rent choice.

[59] Of course, couples' savings decisions would also be influenced by Pareto weights and other preference parameters. We can use the data variations explained in Section 4.2.1 to identify these parameters.

[60] For example, De Nardi, French, Jones, and McGee (2025) find that in their unitary framework, individuals' preferences for terminal bequests are stronger than those for side bequests.

Denote the survival probability of gender g in period t by p_t^g which also depends on current health, PI, and marital status as stated in equation (11). A recursive formulation for a single household's problem for gender g is

$$V_t^{S,g}(x_t) = \max_{q_t} \quad u^g(c_t, h_t) + \beta p_t^g E[V_{t+1}^{S,g}(x_{t+1})|x_t, q_t] + \beta(1 - p_t^g)v^{S,g}(b_{t+1}) \tag{28}$$

subject to budget constraints. $V_t^{S,g}$ denotes the value function of a single retiree of gender g in t. If they survive, they continue the next period again as a single retiree, where the expectation operator is taken with respect to health statuses of the next period. If they die, their remaining assets are left as bequests to children or other heirs ($b_{t+1} = a_{t+1} + (1-\delta)\tilde{h}_t$), in which case they derive bequest utility of $v^{S,g}(b_{t+1})$. The terminal value function is given as $V_T^{S,g}(x_T) = v^{S,g}(b_T)$.

Table E.1: Auxiliary OLS I - Robustness

	(1)	
	Log of consumption	
Wife's income share	-0.74***	(0.24)
PI (mid) × Wife's income share	0.55*	(0.29)
PI (high) × Wife's income share	0.65**	(0.31)
Observations	1915.00	
Adjusted R^2	0.51	
Mean of dependent variable	10.88	

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the household level (in parentheses). The same regression specification is used as in Table 5, except that the dependent variable in this table is log of consumption.

E Estimation appendix

E.1 Auxiliary model

Compared to Table 5 in the main text, Table E.1 in this Online Appendix uses log consumption as the dependent variable. The results are robust in that the household consumption declines with the wife's income share, and the negative effect weakens as PI increases.

E.2 Estimation procedure

We implement indirect inference by following the limited information Bayesian approach as in Fernandez-Villaverde, Rubio-Ramirez, and Schorfheide (2016). Let Θ denote the parameters of interest and $\hat{\mathbf{m}}$ denote the vector of M parameters of auxiliary models from the empirical data. The objective function of $\hat{\mathbf{m}}$ conditional on Θ is

$$f(\hat{\mathbf{m}}|\Theta) = -\left(\hat{\mathbf{m}} - \mathbf{m}(\Theta)\right)' G \left(\hat{\mathbf{m}} - \mathbf{m}(\Theta)\right),$$

where $\mathbf{m}(\Theta)$ is the parameters of auxiliary models based on the simulated data given Θ and G is a weighting matrix. Bayes' theorem tells us that the posterior density $f(\Theta|\hat{\mathbf{m}})$ is proportional to the product of $f(\hat{\mathbf{m}}|\Theta)$ and the prior $p(\Theta)$:

$$f(\Theta|\hat{\mathbf{m}}) \propto f(\hat{\mathbf{m}}|\Theta)p(\Theta).$$

We use uniform prior $p(\theta)$ which serves as putting bounds on parameter values.[61]

Table E.2: Uniform priors

Parameter	[min, max]	Parameter	[min, max]
ψ^h	[-1.0e-8, 1.0e-8]	σ^h	[0.0, 1.0]
ψ^w	[-1.0e-8, 1.0e-8]	σ^w	[0.0, 1.0]
δ^h	[0.1e6, 12.0e6]	a^h	[0.1e6, 100.0e6]
δ^w	[0.1e6, 12.0e6]	a^w	[0.1e6, 100.0e6]
α	[0.0, 1.0]	β_I	[-1.0, 0.0]
c_{nh}	[200.0, 2,000.0]	β_z	[0.1, 20.0]
		β_0	[-5.0, 10.0]

Notes: [min, max] reports the domain of the uniform distribution.

We use the multiple-block Metropolis Hastings algorithm to obtain a sequence of random samples from the posterior distribution. We partition the parameter vector into two blocks: one block (block 1) contains the parameters that are central to governing the regression coefficients and the rate moments, and the other block (block 2) contains the parameters central to governing wealth moments. We denote the objects from auxiliary models to match as $y \equiv [\hat{\mathbf{m}}_1, \hat{\mathbf{m}}_2]$. $\hat{\mathbf{m}}_1$ contains the targets reported in Tables 9 and 10, and $\hat{\mathbf{m}}_2$ includes the wealth profiles reported in Figure 2. We break the parameter space into two blocks as follows: $\Theta = (\Theta^1, \Theta^2)$ where $\Theta^1 = (\psi^W, \psi^M, \beta_I, \beta_z, \beta_0, c_{nh})$ and Θ^2 is for the remaining model parameters. Starting from an initial value $\Theta_0 = (\Theta_0^1, \Theta_0^2)$, the algorithm works as follows. For iteration $j = 1, \ldots, M$ and for block $k = 1, 2$:

- Propose a value $\tilde{\Theta}^k$ for the kth block, conditional on Θ_{j-1}^k for the kth block and the current value of the other block (Θ^{-k}). Θ^{-k} stands for the remaining block except for the kth block.[62]
- Compute the acceptance probability $\alpha^k = \min\left\{1, \frac{f(\tilde{\Theta}^k|\Theta^{-k},y)}{f(\Theta_{j-1}^k|\Theta^{-k},y)}\right\}$. Update the kth block as

$$\Theta_j^k = \begin{cases} \tilde{\Theta}^k & \text{w.p. } \alpha^k \\ \Theta_{j-1}^k & \text{w.p. } (1-\alpha^k) \end{cases}$$

For each iteration, we first update the block 1's parameters conditional on the previous iteration's value for the blocks 1 and 2. Then we sequentially update the remaining model

[61] In cases where prior knowledge or beliefs regarding parameters are available, informative priors can be employed to integrate them into Bayesian estimation. The selection of suitable informative prior distributions remains a subject of debate, representing an avenue for future research.

[62] In our application, $\Theta^{-1} = \Theta^2$ and $\Theta^{-2} = \Theta^1$.

parameter block conditional on the updated block 1.[63]

We apply the multiple-block RWMH algorithm to simulate draws from the posterior $f(\Theta|\hat{\mathbf{m}})$ with uniform priors. The posterior distribution is characterized by a sequence of 4000 draws.

E.3 Model fit

[63]For block 1, we begin with a uniform weighting matrix G, but increase the weights on the regression coefficients relative to the other rate moments, reflecting our priority to match these coefficients more closely as they carry the key information in identifying Pareto weights. Similar weighting strategies have been used in the literature that emphasize putting additional weights on objects that are key in identifying economically important parameters (Sauer and Taber (2021)). For block 2, we use a diagonal weighting matrix based on the inverse of the estimated variance–covariance matrix of the dollar-valued moments.

F Counterfactual appendix

F.1 Welfare measure

Our measure of welfare is the compensating wealth variation (CWV), the initial wealth transfer that makes the economic agent indifferent to the counterfactual regime. Singles' CWV is denoted by $\Delta^{S,g}$, defined as

$$V_t^{S,g}(a_t + \Delta^{S,g}) = \tilde{V}_t^{S,g}(a_t) \tag{29}$$

for $t = 1$ and gender $g \in \{W, M\}$, where $\tilde{V}^{S,g}$ is the value under the counterfactual.

Couples' household-level CWV is denoted by Δ^C, defined as

$$V_t^C(a_t + \Delta^C) = \tilde{V}_t^C(a_t) \tag{30}$$

for $t = 1$, where $\tilde{V}^C$ is the household's value under the counterfactual.

To analyze intrahousehold effects, we compute analogous measures for each spouse $g \in \{W, M\}$, with $\Delta^{C,g}$ defined by

$$V_t^{C,g}(a_t + \Delta^{C,g}) = \tilde{V}_t^{C,g}(a_t). \tag{31}$$

Because the household's value is an aggregation of the spouses' values, these CWVs are linked by

$$V_t^C(a_t + \Delta^C) = \lambda V_t^{C,W}(a_t + \Delta^{C,W}) + (1-\lambda) V_t^{C,M}(a_t + \Delta^{C,M}). \tag{32}$$

Linearizing around baseline wealth yields an intuitive relationship:

$$\Delta^C = \frac{\lambda \Delta^W \frac{\partial V^{C,W}}{\partial a} + (1-\lambda)\Delta^M \frac{\partial V^{C,M}}{\partial a}}{\lambda \frac{\partial V^{C,W}}{\partial a} + (1-\lambda)\frac{\partial V^{C,M}}{\partial a}}. \tag{33}$$

In words, the couple's compensating wealth equivalent is a weighted average of the wife's and husband's, where the weights reflect each spouse's marginal value of wealth. Using (31) and (32), one can easily show that $\min(\Delta^W, \Delta^M) \leq \Delta^C \leq \max(\Delta^W, \Delta^M)$.

F.2 Additional tables for counterfactual analysis

Table F.1: Percent changes (%) in Medicaid recipiency rate

Homestead exemption:	(1) Nobody	(2) Everybody	(3) Singles
Currently married	-4.1	-0.3	-4.0
Currently single women	1.6	0.6	2.3
Currently single men	0.9	3.4	4.1

Notes: The table reports percent changes from the baseline Medicaid recipiency rate, conditional on current marital status and gender.

Table F.2: Welfare effects of homestead exemption for initially singles

	Mean welfare effect
Low PI	W: 6029
	M: 2108
Middle PI	W: 1208
	M: 2138
High PI	W: 396
	M: 0

Notes: Mean compensating wealth equivalents for initially single women (W) and men (M) when singles are granted the homestead exemption.

Table F.3: Intrahousehold welfare effects of counterfactual Medicaid

Counterfactual: Homestead exemption given to nobody			
	Low z	Middle z	High z
Low PI	C: -23	C: -730	C: -8082
	W: -41	W: -614	W: -7650
	M: -20	M: -772	M: -9173
Middle PI	C: -505	C: -115	C: 0
	W: -1799	W: 0	W: 0
	M: 637	M: -162	M: 0
High PI	C: 0	C: 0	C: 0
	W: 0	W: 0	W: 0
	M: 0	M: 0	M: 0
Counterfactual: Homestead exemption given to everybody			
	Low z	Middle z	High z
Low PI	C: 143	C: 35	C: 65
	W: 191	W: 0	W: 118
	M: 132	M: 104	M: -236
Middle PI	C: 10	C: 0	C: 126
	W: 10	W: 0	W: 253
	M: 34	M: 0	M: -111
High PI	C: 0	C: 0	C: 0
	W: 0	W: 0	W: 0
	M: 0	M: 0	M: 0
Counterfactual: Homestead exemption given to singles only			
	Low z	Middle z	High z
Low PI	C: 31	C: -482	C: -8082
	W: 33	W: -372	W: -7626
	M: 0	M: -590	M: -8933
Middle PI	C: -494	C: 0	C: 15
	W: -1989	W: 0	W: 17
	M: 774	M: 0	M: 0
High PI	C: 0	C: 0	C: 0
	W: 0	W: 0	W: 0
	M: 0	M: 0	M: 0

Notes: For each (PI, z) bin, mean compensating wealth equivalents reported for initially couples (C), broken down by wives (W) and husbands (M).