Debt as Insurance: Precautionary Borrowing and the Co-Holding Puzzle in Retirement*

Jared Buck†

November 2026

Job Market Paper

Abstract

Households increasingly retire with mortgage debt while holding liquid assets that earn far less than the mortgage rate. This paper explains why such co-holding is rational and shows what it implies for policy analysis that nets debt against assets. I build a life-cycle model of retirement in which households hold housing, a mortgage, and a liquid asset, and face health risk, medical expense risk, a Medicaid means test, and costly refinancing, and I estimate it on the Health and Retirement Study by the simulated method of moments. The model reproduces co-holding among mortgagors at ages 65 and 75. A Shapley decomposition attributes about half of mortgagors' co-holding to medical expense risk and close to a third to the cost of re-accessing home equity. In experiments that eliminate the Medicaid homestead exemption and crash house prices by 30%, welfare effects vary widely within net-worth quintiles, and a netted representation of co-holders misstates individual welfare effects by a quarter to a half of their size while leaving aggregates roughly unchanged. Net worth is a reasonable guide to a policy's total cost but a poor guide to its incidence.

Keywords: co-holding; precautionary saving; precautionary debt; mortgages; retirement; Medicaid.

JEL: D14, D15, E21, G51, I13.

*I am grateful to my supervisor, Tatyana Koreshkova, and to my committee for guidance and support, and to participants at the 2026 Canadian Economics Association Annual Meeting for helpful comments. All errors are my own.

†Department of Economics, Concordia University. Email: jared.buck@mail.concordia.ca.

1 Introduction

The share of US households retiring with debt has risen sharply since 1998, and the increase has been driven mostly by mortgages. At the same time, nearly half of these households carry liquid assets in excess of three months of income. Since the return on liquid assets is low relative to mortgage interest rates, these households pay a premium to hold cash beside their debt rather than pay the debt down. The coexistence of low-return liquid assets and more expensive debt is a long-standing puzzle (Tobin 1957), and the rational explanations developed for it were built for revolving credit: when access to a credit line is uncertain, drawing on it now insures against losing it later (Fulford 2015; Druedahl and Jørgensen 2018). Mortgage debt removes that motive. The proceeds arrive once at origination and the lender cannot recall the loan, so there is no access to protect. What takes its place is irreversibility: once principal is repaid, it can only be recovered through a new origination that is costly or, for many retirees, unavailable. In companion work (Buck 2026), I show in a three-period model that this structure makes co-holding rational: households that borrow through closed-end contracts against resources arriving later keep part of the proceeds liquid when the expenses ahead are uncertain. The mechanism is particularly relevant in retirement. Most people begin retirement as home-owners and become more likely to sell with age, so the home is an illiquid asset against future needs, and because retirees usually do not have labour income, they are less likely to qualify for new loans. This paper builds a quantitative life-cycle model that explains mortgage co-holding throughout retirement, estimates it on the Health and Retirement Study, and uses it to ask what representing households by net worth alone would miss.

In the model, households carry liquid assets, mortgages, housing wealth, and earn income deterministically given a type and marital status; the risks are health and medical expenses. These resources are then used to make consumption, housing, saving, and borrowing decisions simultaneously. The model is estimated by the simulated method of moments on the Health and Retirement Study (HRS). It reproduces the portfolio facts of retired homeowners: ownership, mortgage holding and balances, the income gradient of co-holding, and the share of mortgagors who hold a liquid buffer, about half in the model as in the data. Co-holding as a share of all households is slightly below the data at 75, because the model has too few mortgagors at that age rather than too few buffers among them. Its main misfit is liquid wealth late in life, which rises in the model after 85 because the proceeds of a home sale have nowhere to go but liquid assets, while in the HRS it stays flat.

Co-holding in the model is self-insurance through precautionary saving and precautionary debt, and medical-expense risk is the main source. Among households with a mortgage

between 65 and 95, 45 per cent co-hold in the estimated model and removing expense risk leads to a 29 percentage point drop. Costly access to home equity matters less on its own, producing co-holding among 11 per cent of mortgagors (compared to 15 per cent for expense risk alone), but it shapes how households insure: given the risks, removing refinancing costs lowers co-holding among mortgagors by 19 percentage points, because households that can re-borrow cheaply pay the mortgage down and borrow when the need arises instead of holding cash beside it. In the Shapley decomposition, expense risk accounts for about half of co-holding among mortgagors and refinancing costs for close to a third. Health risk and the risk of losing access to credit contribute little on their own but amplify the others.

In this paper, I consider two classes of experiments, Medicaid reform and a housing price crash, to see how household portfolios respond, and to infer whether net-worth is a sufficient statistic to study welfare. The experiments reveal that while net-worth is a reasonable guide to the total cost of a policy or shock, it is a poor guide to understanding how much each household gains or loses. Within net-worth quintiles, welfare effects range from nothing to tens of thousands of dollars measured in terms of liquid-wealth equivalence. A 30% housing price crash results in welfare costs of \$0.57 to \$1.01 per dollar of wealth lost, depending on portfolio composition. Additionally, representing co-holders by their net positions misstates welfare losses among these groups by between one quarter to one half while leaving aggregate results relatively similar.

This paper contributes to the literature on the co-holding puzzle. The observation that households borrow at high rates while saving at low ones dates to Tobin (1957) and was revived by the credit card debt puzzle (Gross and Souleles 2002). Rational explanations rest on two mechanisms. The first is a demand for liquidity: when some consumption can only be paid for with liquid assets, households hold cash even while revolving debt (Telyukova and Wright 2008; Telyukova 2013). The second is uncertain access to credit: when a line of credit can be cut, drawing it down now insures against losing it later (Ludvigson 1999; Fulford 2015; Druedahl and Jørgensen 2018). Both were built for revolving credit. Closed-end debt pays out once at origination, so there is no line to protect; what remains is the irreversibility of repaid principal, and this paper shows that this friction alone sustains co-holding in a quantitative model of retirement. The move from credit cards to mortgages also matters for the stakes: in the SCF, mortgage co-holders pay several times the annual carry cost that credit-card co-holders do (Buck 2026).

A parallel literature attributes co-holding to behavioural frictions: mental accounting and earmarked savings (Thaler 1985; Shefrin and Thaler 1988; Sussman and O'brien 2016), self-control problems within and between household members (Bertaut, Haliassos, and Reiter 2009; Meier and Sprenger 2010; Vihriälä 2025; Medina and Pagel 2025), and beliefs about the

cost of debt (Stango and Zinman 2009). Measurement work shows the puzzle's size depends heavily on definitions and data frequency (Gathergood and Weber 2014; Greene and Stavins 2023; Islam and Picault 2026; Gathergood and Olafsson 2026). The model here abstracts from these frictions, not because they are absent, but to establish how much of retirees' co-holding a fully rational benchmark explains before appealing to preferences or cognition; the decomposition in section 5 is an answer to that question.

Beyond the co-holding literature (Telyukova 2013; Fulford 2015; Druedahl and Jørgensen 2018), this paper connects to the work on hand-to-mouth (HtM) (Kaplan, Violante, and Weidner 2014; Kaplan and Violante 2014; Aguiar, Bils, and Boar 2025). Specifically, the retirement model features a house that is treated as an illiquid asset. In this way, the paper can bridge the gap between HtM and the co-holding model due to Druedahl and Jørgensen (2018). Co-holders and the HtM live on opposite sides of the household problem: co-holders pay a premium to access liquidity whereas the HtM prefer to live constrained rather than pay interest on debt. In the Medicaid experiments, this paper shows that co-holders and the HtM generally move in different directions. This also relates to the wider literature on portfolio choices featuring debt and illiquid assets (Campbell and Cocco 2003; Cocco 2005; Yao and Zhang 2005; Chetty, Sándor, and Szeidl 2017; Berger et al. 2018). Existing research focuses on the price of converting illiquid housing wealth. In this paper, the cost of converting housing wealth, together with the risk it may become unavailable, is what makes retirees hold cash beside their mortgages.

Another margin that affects portfolio composition is insurance. Hubbard, Skinner, and Zeldes (1995) and De Nardi, French, and Jones (2010) both show that means-testing discourages saving. This paper shows that Medicaid doesn't just discourage liquid savings, but it also determines what form savings take. Medicaid is means tested across income net of expenses and asset holdings with a homestead exemption. The model has all three tests, so I can evaluate each margin, but because income is exogenous, I focus on assets and the homestead exemption. Intuitively, tightening the homestead exemption (decreasing the cutoff) should encourage homeowners to choose higher leverage to qualify for coverage. However, because the asset test is quite strict (2 thousand dollars for singles and 3 thousand dollars for couples) and the only place to store loan benefits is in liquid assets, tightening the homestead exemption under 250 thousand dollars quickly disqualifies most from eligibility. Households still choose higher leverage in the counterfactuals however, this works through self-insurance channels rather than by meeting eligibility criteria. Softening the asset test does the opposite: eligibility rises and co-holding falls, but the relation is not monotone. At intermediate limits, households that could save keep their liquid wealth at or below the limit to stay eligible, so the test itself discourages saving. Once the limit is high enough not to

bind (60 thousand dollars for singles and 90 thousand for couples), that penalty disappears, and median liquid wealth rises from about 10 to 15 thousand dollars while co-holding edges back up. A sufficiently generous test stops distorting saving.

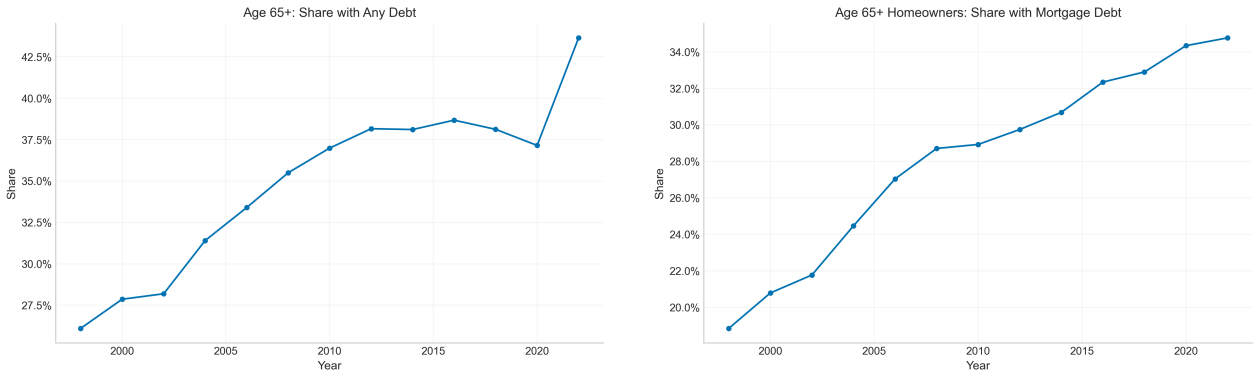
Home equity is also one of the most important savings devices for retirees and the greatest share of household savings tends to be the house (Nakajima and Telyukova 2017, 2020). This means that retirees may be particularly sensitive to housing shocks. The literature on the 2008 crash establishes that housing shocks transmit through collateral and liquidity rather than wealth, and that distressed borrowers benefit more from cash flow relief rather than restored equity however, the work has focused on working aged households (Foote, Gerardi, and Willen 2008; Mian, Rao, and Sufi 2013; Agarwal et al. 2017; Ganong and Noel 2020). To see how a housing shock affects retirees, I impose a one-time, permanent 30% fall in house prices at age 67 and compare responses under different beliefs about the crash. The responses are small. Co-holding falls by one to two percentage points and the mortgagor rate is almost unchanged. Median leverage among mortgagors rises from 0.40 to 0.44–0.45, well below the 0.57 a mechanical revaluation of the pre-crash loan would imply, so mortgagors pay down after the crash. Compared to an unexpected shock, anticipation moves two margins in opposite directions as fewer households hold a mortgage, but for those who do, they borrow more against the pre-crash value, raising leverage at the 90th percentile from 0.82 to 0.88 and the share of underwater households from 1.1% to 1.7%.

The rest of the paper is structured as follows. Section 2 describes the HRS data and documents some key observations, section 3 sketches the model, section 4 outlines the calibration methodology, 5 discusses the calibration results, section 6 runs experiments on Medicaid and a housing price crash, and section 7 concludes.

2 Data

Before presenting the quantitative model, this section documents several descriptive patterns in the data that motivate the analysis. Standard models of retirement often emphasize wealth decumulation and the repayment of debt (Lusardi, Mitchell, and Oggero 2020; De Nardi, French, and Jones 2010; Amromin, Huang, and Sialm 2007; Browning and Crossley 2001). By contrast, the data suggest that retirement behaviour is also shaped by the interaction of liquid assets, debt obligations, housing, and expenditure risk. The goal of this section is not to establish causality, but rather to highlight a set of empirical facts that motivate the analysis, and some which the model might explain. The descriptive analysis

uses HRS data from 1998 to 2022 and restricts the sample to households aged 65–95.[1]



(a) Share with Any Debt

(b) Share of Homeowners with a Mortgage

Figure 1: Debt Trends over Time

Note: These panels plot average share of HRS households with debt from 1998-2022. Panel (a) plots the share of any reported debt and panel (b) plots the share with mortgage debt.

Contrary to the conventional view that households enter retirement debt-free, many older households continue to carry debt well into retirement. Figure 1a shows that the share of households aged 65 and older with any debt rose steadily over the sample period. Figure 1b shows a similar pattern for mortgage debt among homeowners, suggesting that much of the rise in debt is tied to housing-related borrowing. This is consistent with existing evidence that mortgage debt has become increasingly common later in life (Lusardi, Mitchell, and Oggero 2020; Chen, Liu, and Munnell 2023; Bédard and Michaud 2021; Friedberg, Sun, and Webb 2023).

At the intensive margin, mortgage balances have also risen. Figure 2 reports the mean and selected percentiles of real mortgage debt among mortgagors, expressed in 2022 dollars. Across much of the distribution, mortgage balances increased from 1998 through roughly 2010, followed by a flattening period or modest decline in more recent years. Part of this increase appears to reflect cohort replacement rather than only within-cohort aging. Figure 3 shows that younger birth cohorts tend to carry more debt at a given age than earlier cohorts, suggesting that debt in retirement has become more prevalent across successive generations.

These trends are not unique to older households, and the literature points to several forces that may have contributed to the rise in debt, especially mortgage debt. Better access to credit, higher house prices, and the expansion of refinancing opportunities have

1. Earlier waves are excluded to maintain consistency in the construction and interpretation of household balance sheet measures across survey years. The sample is also restricted to ages 65–95 because the number of households above age 95 declines quickly, and the quantitative model is intended to capture a comparable 30-year retirement horizon.

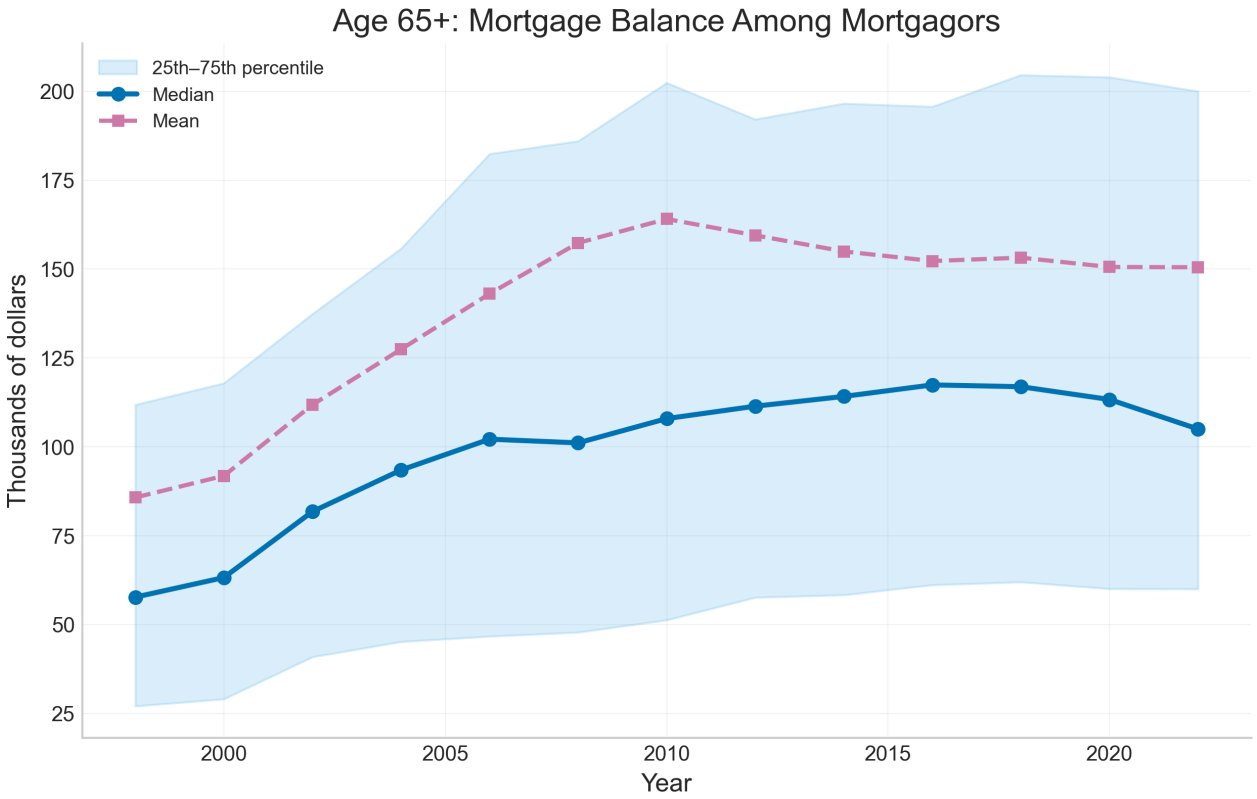


Figure 2: Real Mortgage Debt among Mortgagors

Note: This plot shows the mean, median, and 1st to 3rd quartile spread of HRS mortgage balances from 1998-2022.

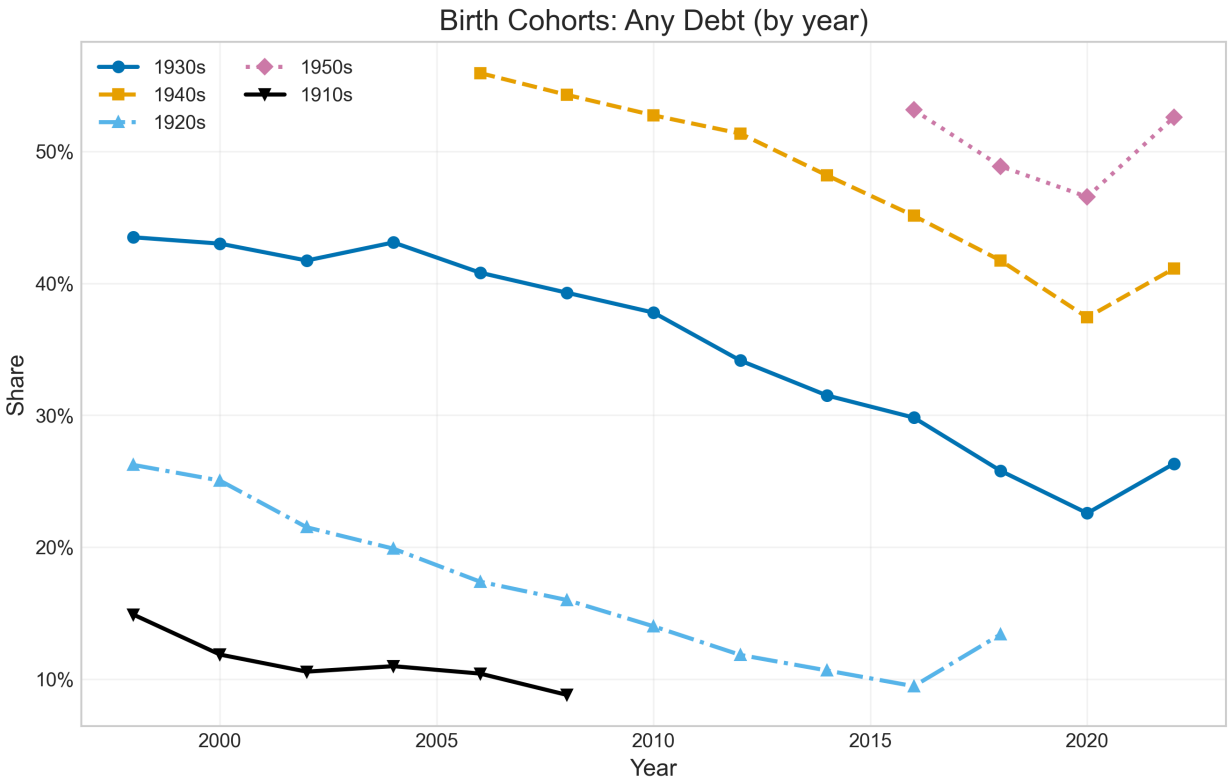


Figure 3: Debt Holdings by Birth Cohort

Note: This figure plots average share of HRS households with debt by cohort from 1998-2022.

all likely played a role (Lusardi, Mitchell, and Oggero 2020). In addition, the low interest rate environment following the Great Recession reduced the cost of mortgage borrowing for many households. At the same time, rising debt in retirement occurred alongside declining homeownership. Figure 4 compares homeownership rates and debt prevalence over time. This adds an important layer to the descriptive picture. On one hand, higher house prices may require larger loans for households that purchase homes. On the other hand, higher house values may relax borrowing constraints for existing owners by increasing home equity. Yet these channels alone do not explain why homeownership itself declines over the same period. One possible interpretation is that rising house prices changed both affordability and borrowing incentives, contributing to a shift in the composition of owners and renters at older ages.

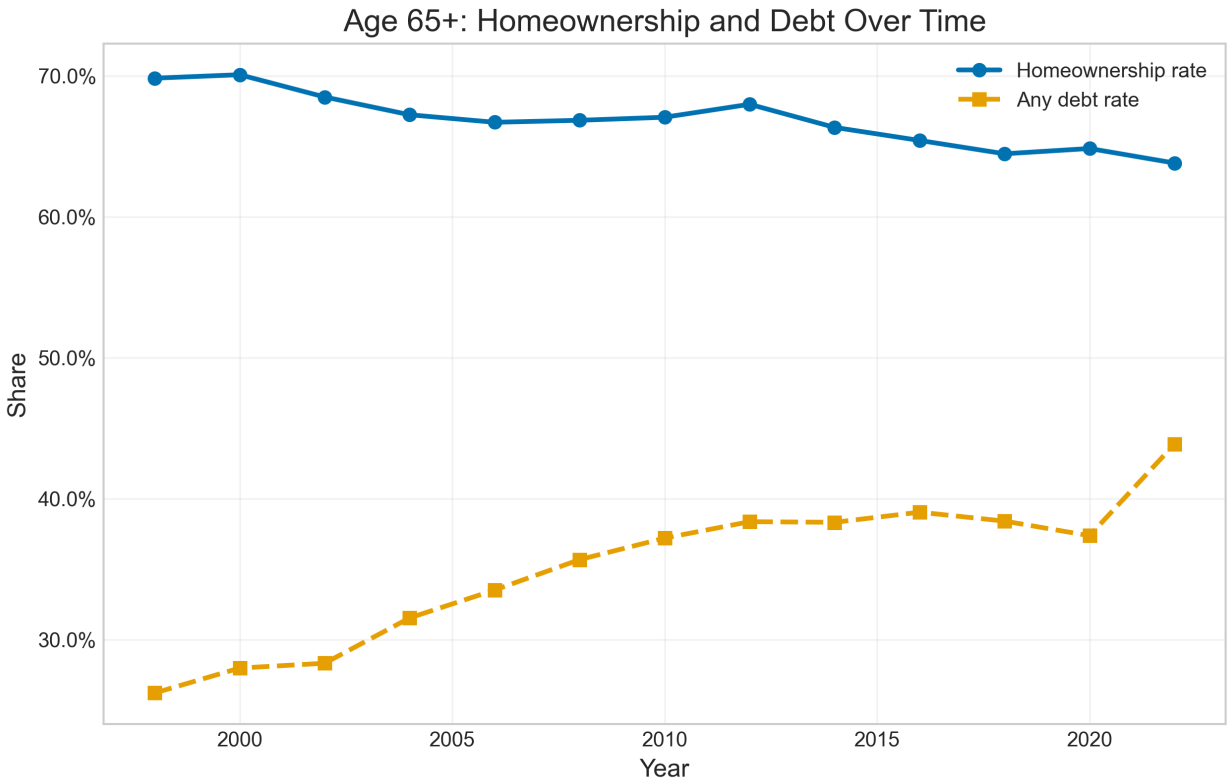


Figure 4: Homeownership and Debt over Time

Note: This figure plots the homeownership and debt share among all households in the HRS from 1998-2022.

Nevertheless, older households are more and more likely to carry mortgage debt at more sizeable balances. At a glance, this is not necessarily a problem. If indeed it is the case that market conditions exogenous to the household have driven the increase, there may be no cause for concern. For example, if older households have cheaper debt secured by appreciating

assets, then from a risk perspective, they may be no worse off. Figure 5 would support this view showing that leverage has not increased substantially over time. However, figure

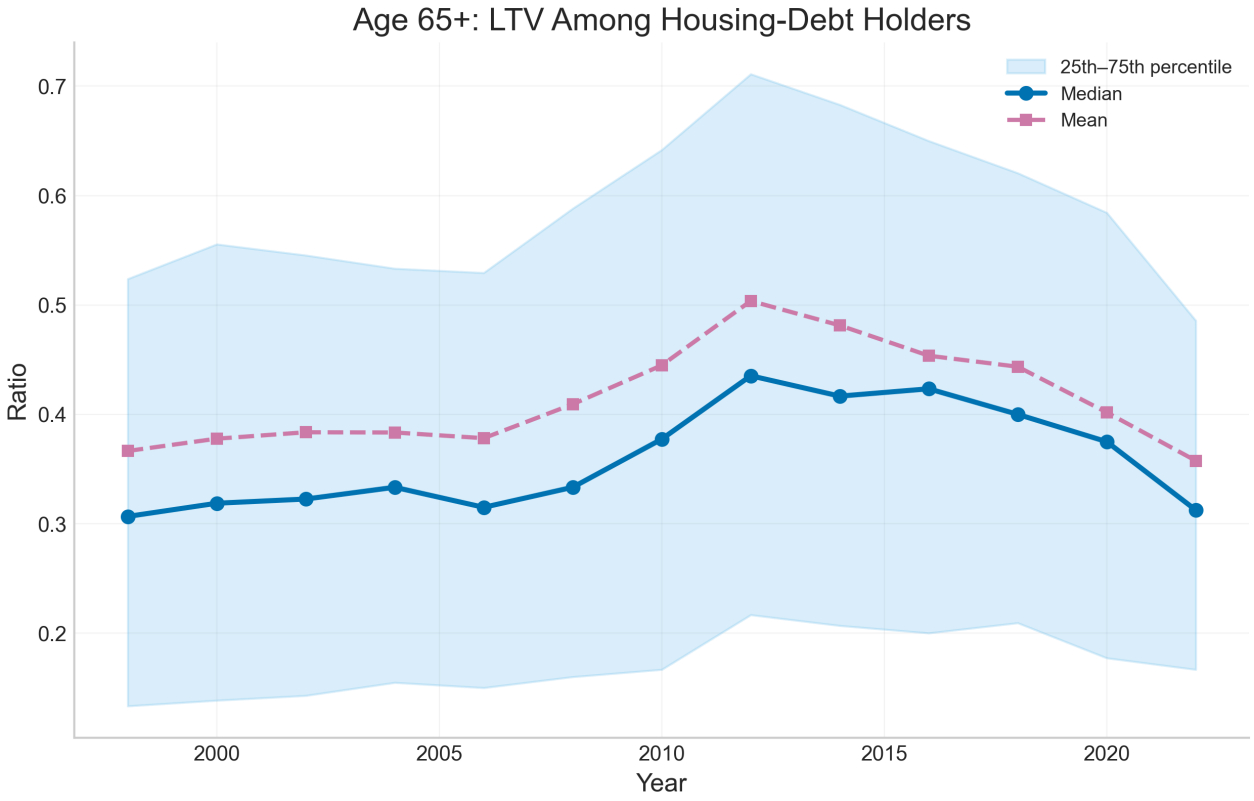


Figure 5: Loan to Value (LTV) over Time

Note: This figure plots mean, median, and 1st to 3rd quartile spreads of loan-to-value among home-owners with housing debt in the HRS from 1998-2022.

6 suggests many mortgagors hold only modest liquid buffers relative to their obligations. This reveals that the key issue is not necessarily rising risk or leverage, but the difficulty of managing mortgage debt when wealth is tied up in housing and expenditures are uncertain.

In the SCF, co-holding at the extensive margin peaks between ages 50–60 and persists beyond age 80 (Buck 2026). Figure 7 reveals what the age pattern looks like for households aged 65+ in the HRS across three different asset thresholds. The estimates here come from the HRS not the SCF, so the estimates should be slightly different. However, another reason my HRS estimates suggest higher levels of co-holding than in the SCF is because the liquid asset measures differ. In addition to transaction accounts, here, CDs, bonds, and stocks are added to the liquid set, so the same households here would have higher balances by inclusion. Nevertheless, regardless of whether a 1 month, 3 month, or 6 month income threshold is chosen, the overall pattern persists, and co-holding is an important characterization of behaviour late in life.

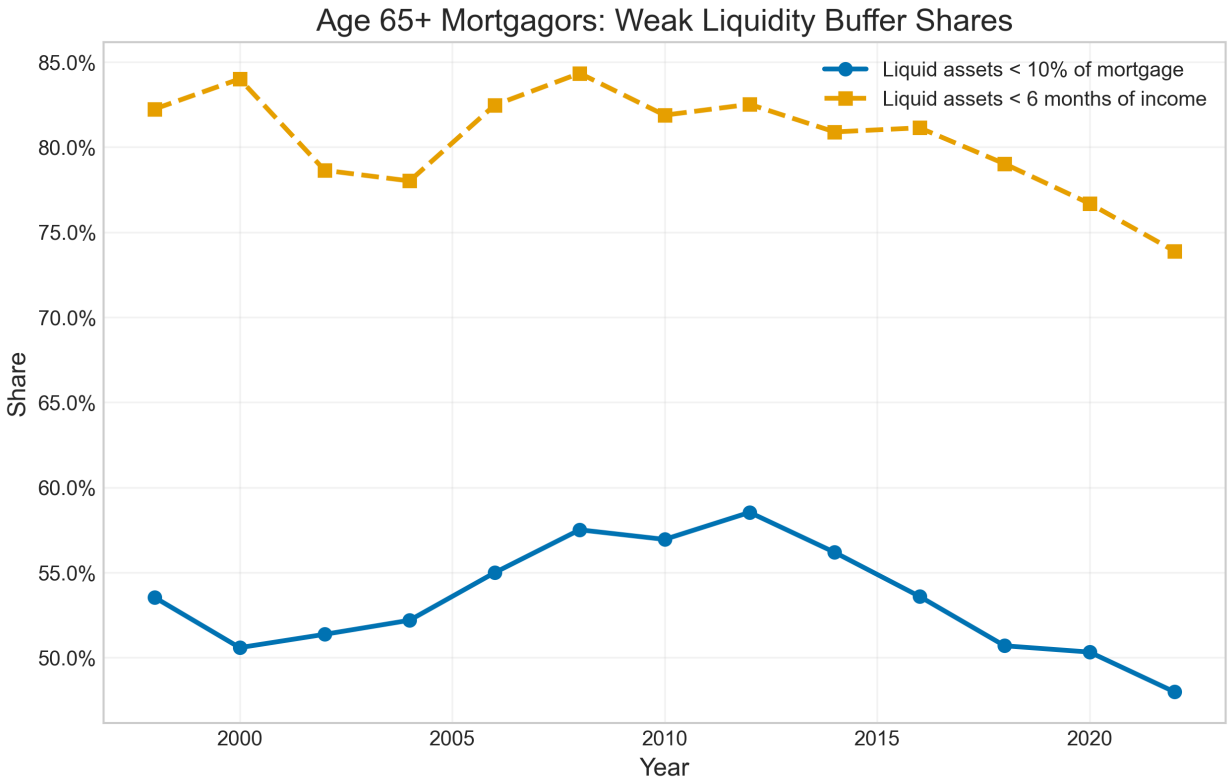
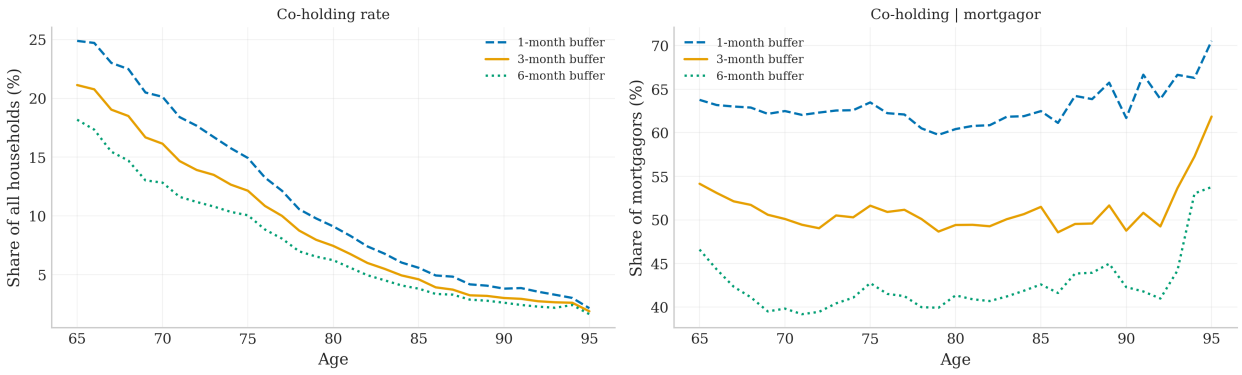


Figure 6: Weak Liquidity Buffer Shares

Note: This figure plots liquidity buffers in the HRS from 1998-2022 by share below the specified thresholds.



(a) Co-holding Share

(b) Co-holding Share Among Mortgagors

Figure 7: Co-holding by age

Note: This figure plots the mortgage co-holding share in the HRS from 1998-2022. Panel (a) shows the co-holding share across the entire sample and panel (b) shows the co-holding share among mortgagors. In the plots, co-holders are defined across three different liquid asset thresholds: liquid assets in excess of 1-month's income, in excess of 3-months of income, and in excess of 6-months of income.

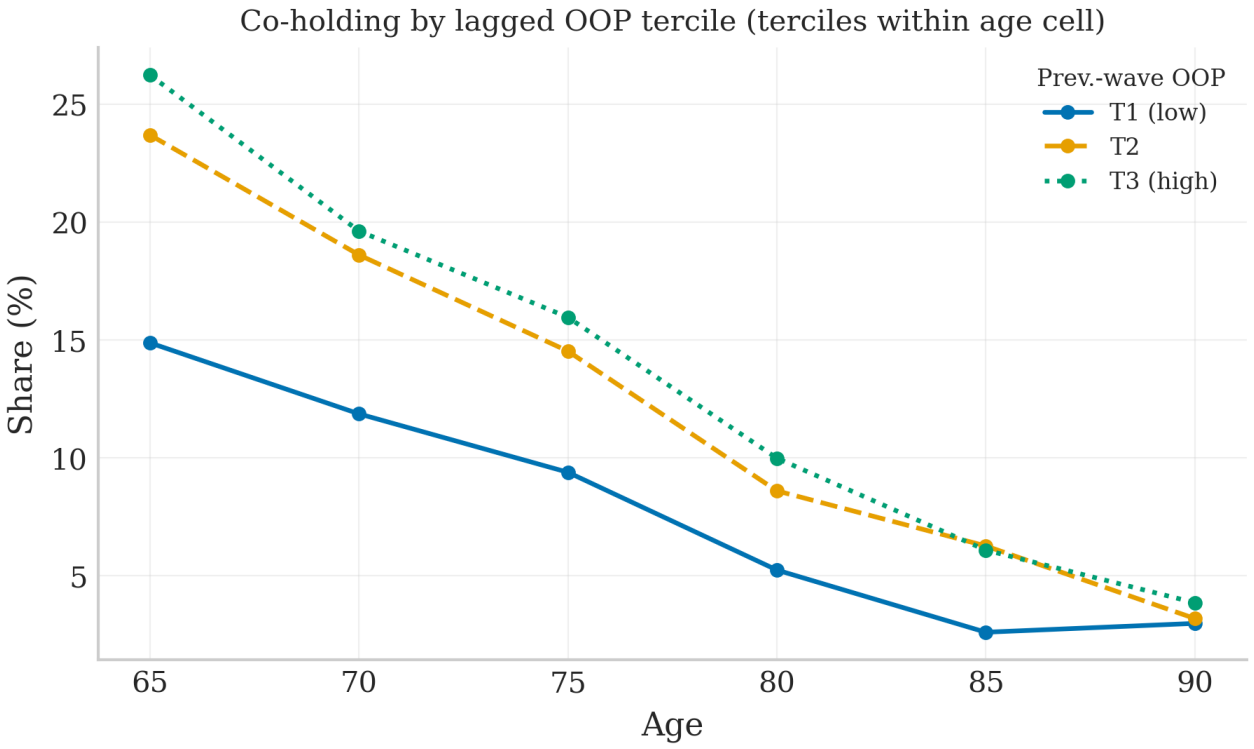
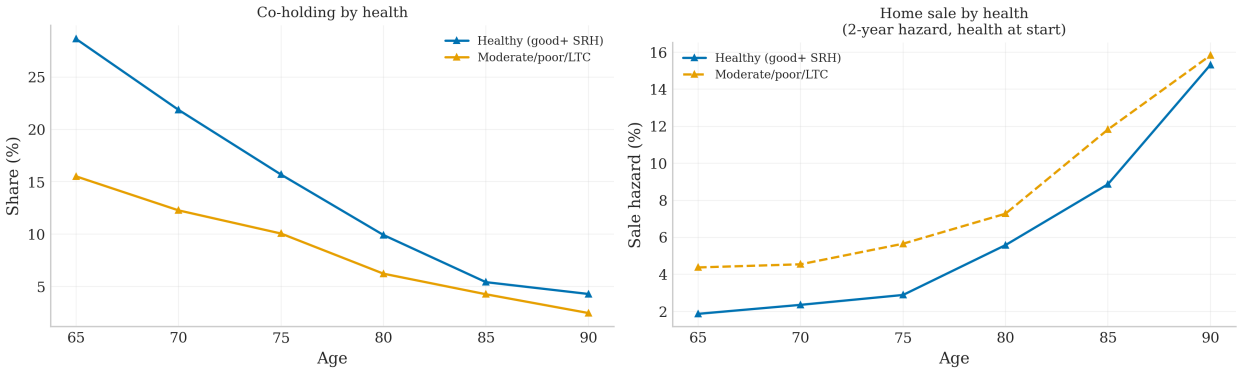


Figure 8: Co-holding by OOP medical expenses

Note: This figure plots the share of mortgage co-holders with liquid assets in excess of three months of income, across age for three out-of-pocket (OOP) medical expenses terciles in the HRS from 1998-2022.

Using de-identified Chase mortgage and deposit account data, Farrell, Bhagat, and Zhao (2018) show that borrowers with less than one mortgage payment equivalent in liquid reserves had default rates roughly seven times higher than borrowers with at least four mortgage payments in reserve. Additionally, 84 percent of defaults in their sample were from owners with less than four mortgage payments held as a buffer. Taken together, these facts suggest that retirement households are not simply drawing down wealth after paying off debt. Many continue to manage mortgage obligations alongside liquid savings and illiquid housing wealth. This makes expenditure shocks potentially important even when leverage is not especially high.

If this is true, then households who recently suffered from bad expense shocks should be observed with higher co-holding shares. Figure 8 provides evidence for just that phenomenon, though perhaps not as pronounced and perfectly monotonic as might be expected. This is because there is a second source of liquidity in the form of home sale. Additionally, the most likely source of expenses in retirement should come in the form of medical expenses, correlated to health. Figure 9 supports the story that expenses and health drive liquidity needs. The panel on the right suggests that homeowners in poor health are more likely to



(a) Co-holding Share

(b) Home Sale Hazard

Figure 9: Co-holding and Home Sale by Health

Note: This figure plots co-holding in panel (a) and home sale hazard in panel (b) by health status in the HRS from 1998-2022. Co-holders are defined as households with mortgage debt and liquid assets in excess of three months of income.

sell, which explains why they are less frequent co-holders: they don't own homes and thus no longer have a mortgage. While we'd expect to find the same kind of pattern with lagged OOP medical expenses and home sale, there is no clear pattern. The verdict: home owners may use loans to pay off medical expenses, and while poor health is correlated with expenses, poor health seems more likely to drive home sale.

The next section combines the salient facts documented here with the closed-end loan mechanism of Buck (2026) to build a retirement life-cycle model that explains co-holding. The model keeps the closed-end loan structure that generates co-holding, but extends it allowing re-financing. From here, the model focuses on OOP medical expenses correlated with health states to account for the specific heterogeneity in co-holding and home sale.

3 Model of Retirement

Buck (2026) develops a three-period model in which households choose liquid assets and closed-end debt optimally in the presence of future expenditure risk. The descriptive evidence above suggests that the same underlying tradeoff remains important in retirement, but in an environment where housing and mortgage debt are central. The model here is a natural extension of that framework: the core mechanism is preserved, while the model is expanded to incorporate a life-cycle, housing, mortgage debt, and retirement expenditure shocks.

3.1 Household Utility

Utility function. The model's environment is influenced by Nakajima and Telyukova (2020) and I use the same utility function. Each period, households earn utility from consumption c, housing h, and home-ownership o, and utility also depends on marital status s. The utility function is written as

$$u(c,h,s,o) = s\frac{\left(\frac{1}{\mu_s}c^{\eta}(\omega_o h)^{1-\eta}\right)^{1-\sigma}}{1-\sigma} \tag{1}$$

where $s=1$ denotes a single household and $s=2$ denotes coupled, $o=0$ renter and $o=1$ owner, μ_s is the household equivalence scale, and ω scales the value of housing utility for homeowners. Utility parameters are set such that $\omega_1 > \omega_0 = 1$ and $\frac{2}{\mu_2^{1-\sigma}} > \frac{1}{\mu_1^{1-\sigma}} = 1$.

Bequests. Households also have preferences over bequests. Following Nakajima and Telyukova (2020) the utility from bequests is given by

$$B(a,h,d) = \gamma\frac{\big((1+r)a + (1-\kappa)hp - d + \zeta\big)^{1-\sigma_B}}{1-\sigma_B} \tag{2}$$

where γ is a scale factor for bequest value, ζ is a value floor for bequests, p is the price of non-rental properties, d is the balance on mortgage debt, κ is cost of home sale, and σ_B is the bequest specific coefficient of relative risk aversion[2].

3.2 Dynamic Processes

Income. Households enter retirement with a fixed income type $i \in \{\underline{i}, \bar{i}\}$. In the model, period income aggregates pension income, social securities, and age-varying withdrawal profiles from retirement savings calibrated to match HRS. Income at age t takes the form

$$y_t = y_t(i, s_t)$$

so that income depends on age, marital status, and i. Treating retirement savings withdrawals as part of the income process allows the liquid assets state variable, a_t, to map directly onto liquid financial holdings in the data rather than just summarizing all non-housing wealth. Additionally, the income type index, i, is the main source of income heterogeneity. Section 4 specifies the calibration procedure.

Housing prices. The consumption value of housing is given by h, and the value of housing at any given time is p_t, such that the portfolio value of house size, h, is hp_t. Home

2. The addition of a separate bequest CRRA term, σ_B is a departure from (Nakajima and Telyukova 2020)

prices grow deterministically at the rate g, such that $p_{t+1} = (1+g)p_t$. Homeowners pay the depreciation rate on housing

$$\delta_t = \delta_0 + \delta_1 \max\{t - t^*, 0\} \tag{3}$$

to maintain the consumption value of their property. Depreciation takes on a form that weakly increases with time to reflect the fact that upkeep tends to become more difficult with age. This form is necessary to capture the shape of home sale and ownership over time. Households are endowed with ownership status and housing size at $t = 0$. Homeowners can decide to sell their properties at any time and either downsize to a smaller house, or become renters. Renters can also purchase homes at any time.

Renters pay the risk free interest rate r and the rental property specific depreciation rate δ_R on the value of the rental property of size h such that their rental payments are given by $(r + \delta_R)hp_{R,t}$. The value of rental properties $p_{R,t}$ follows the deterministic growth path $p_{R,t+1} = (1 + g_R)p_{R,t}$.

3.3 Idiosyncratic risks

Each period, households face three sources of idiosyncratic risk in the form of health risk, medical expense shocks, e, and a mortality risk[3]. In terms of health status, H, households can be indexed as healthy, $H = G$, moderate, $H = M$, bad, $H = B$, or long-term care, $H = L$. Health status transitions follow the Markov chain $\Gamma_{H_{t+1},H_t} = P(H_{t+1}|H_t)$. The primary role of health shocks in the model is to inform medical expense and mortality risk.

Medical expense shocks are discrete and follow a Markov process that depends on income, health, and age $\Psi(t, e_{t+1}, e_t, i, H_t) = P(e_{t+1}|t, e_t, i, H_t)$. Conditioning on income and health allows the model to match heterogeneity in the health-income gradient more closely (De Nardi, French, and Jones 2010).

In reality, coupled households may end up single in the future for a number of reasons. The model is agnostic to these reasons and treats both mortality and divorce risk as a single Markov chain. In particular, the mortality process is $\pi_{s_{t+1},s_t,i,H_t} = P(s_{t+1}|t, s_t, i, H_t)$ where $s = \{0, 1, 2\}$ represents deceased, single, and coupled households respectively. Each state where $s_{t+1} < s_t$ is treated as an absorbing state ($P(s_{t+1} > s_t|t, s_t, i, H_t) = 0$) suggesting singled households cannot re-marry. Just as with medical expense risk, mortality depends on the income index, health, and age to capture the gradient in income and mortality (De Nardi, French, and Jones 2010).

3. I don't explicitly model divorce, however, married couples can experience shocks where either only one or neither survives.

To reduce the notational burden, going forward I drop the t subscripts wherever possible and use a prime notation to denote future periods.

3.4 Medicaid

Medicaid represents the largest source of public insurance against medical expenses for retirees, and can affect liquid asset holdings, debt repayment, and co-holding. Medicaid is means tested across three relevant thresholds: liquid assets, income, and homestead exemption. The specific rules can vary by state so I focus on the federal floor to capture the greater picture. When a household qualifies for Medicaid, their medical expenses are covered which frees up financial resources for both consumption and saving. To capture the realism of Medicaid, the following eligibility criteria is imposed

$$EL = \begin{cases} 1 \quad \text{if all} \begin{cases} a & \leq a_M(s) \\ y(i,s) & \leq y_M(s) - e \\ hp - d & \leq HO \end{cases} \\ \\ 0 \quad \text{otherwise} \end{cases} \tag{4}$$

where EL is a binary denoting eligibility, $a_M(s)$ is the asset threshold, $y_M(s)$ is the income threshold, and HO is the homestead exemption. Notice that in order to qualify, the household must meet all three eligibility criteria simultaneously.

In the data, not all households eligible for Medicaid report receiving benefits (Borella, De Nardi, and French 2018). To account for this, households solve their problems expecting to receive Medicaid whenever they are eligible, and in the simulation a take-up probability $\xi \in (0, 1)$ determines which eligible households are recorded as recipients. Take-up therefore affects only the recorded recipiency moments: every eligible household behaves as insured[4]. This is a deliberate *ex post* simplification that lets the model match observed recipiency rates without requiring households to internalize information frictions or carry additional states. If a households is eligible in the model, they successfully receive the Medicaid benefit, their expenses e are covered for that period ($e = 0$) and they receive a Medicaid consumption floor $c > \underline{c}^{med}$.

4. The share of households effectively covered in the model is thus the eligible share, about $1/\xi$ times recorded recipiency. This is close to how Medicaid works for the elderly: eligibility rather than enrollment determines access, since coverage can be obtained, with limited retroactivity, when a large expense arises. To the extent that some eligible households would never enroll, the model overstates their insurance and understates their precautionary saving.

In addition to Medicaid, I also impose a smaller consumption floor $\underline{c}^{med} > \underline{c} > 0$ that applies in all states. This not only ensures the model can be solved with numerical stability, but also captures other social safety nets that are not modeled. Furthermore, since the value of insurance is so high for retired households, the consumption floor is also necessary to prevent wealthy households from aggressively holding liquid assets well beyond what is represented in the data.

Including Medicaid in the model not only adds realism to social insurance, but also has direct relevance for co-holding and can actually affect decisions two in different ways. First, the asset test creates a kink in the value function at $a = a_M$: households just above the liquid asset threshold have an incentive to decrease their liquid holding either through paying their mortgages or increased consumption. This can drive down co-holding at the intensive margin. Second, the homestead exemption protects housing wealth from the asset test which increases the value of holding onto a mortgage rather than paying it down. For liquid poor households, this can increase co-holding at the extensive margin. Therefore, Medicaid is not just an insurance mechanism in the model, but it plays a direct role in the tradeoff between debt obligations and liquidity.

3.5 Mortgages and Refinancing

In reality, households at the start of retirement hold mortgages with heterogeneous origination dates, terms, and remaining amortization schedules. Tracking this heterogeneity would require additional state variables and the imputation of contract terms not directly observed in HRS. I abstract from this by assuming that fixed minimum payments are determined from a standard 30 year mortgage with a 20% down-payment. Thus, minimum payments on mortgages take the form

$$\underline{m} = 0.8 \cdot phA(r_m, \bar{\tau}). \tag{5}$$

where

$$A(r_m, \bar{\tau}) = \frac{r_m(1+r_m)^{\bar{\tau}}}{(1+r_m)^{\bar{\tau}} - 1}^{5} \tag{6}$$

This structure retains the important closed-end features of mortgages while keeping the state space tractable. Mortgage debt evolves according to the standard law of motion

$$d' = (1+r_m)d - m \tag{7}$$

5. The $\bar{\tau}$ notation is used in general but I assume that $\bar{\tau} = 15$ to reflect a 30-year amortization.

where m is the payment made on debt. However, mortgagors must always make the minimum payment on debt.

Owners and home buyers are allowed to finance new loans provided the new loan satisfies the cap constraint, $\bar{D}$. The debt cap takes on multiple dimensions. First, households in the LTC health state are not given re-financing options. Second, new loans cannot exceed 80% of the home value, $0.80\bar{h}p$, equivalent to the typical minimum down payment on a new mortgage. Third, new loans must pass a debt-to-income (DTI) test $\bar{D}_{\text{DTI}}$. Specifically, the DTI test is

$$d' \leq \bar{D}_{\text{DTI}} = \frac{\theta_{\text{DTI}} \cdot y}{A(r_m, \bar{\tau})} \tag{8}$$

such that debt must be no greater than a fraction θ_{DTI} of income. Additionally, since new loans become rarer with age in the data, the DTI test becomes stricter with t after $t_R = 5$ (age 75) through

$$\chi(t) = \max\{0, 1 - \psi_D \max(t - t_R)\}^6. \tag{9}$$

Therefore, new loans are granted so long as the level of debt being chosen d' satisfies

$$d' \leq \bar{D} = \min\{0.8\bar{h}p,\ \chi(t)\bar{D}_{\text{DTI}}\}. \tag{10}$$

Lastly, households purchasing new loans are subject to fixed and proportional financing costs such that the cost of a new loan is

$$d' = (1 + f_1)d + f_0 \tag{11}$$

where d is the loan value prior to cost. This cost structure is specified capture both the fixed and proportional costs of financing new loans.

3.6 Renters

Renters make consumption, savings, and rental housing choices each period in addition to a home purchase choice. Non-home-buying renters face the Bellman problem

$$V_t^R(a : H, s, e, i) \equiv \max_{c,h,a'} \Big\{ u(c, h, s, 0) + \beta E\big[V_{t+1}^R(a' : H', s', e', i)\big|H, s, e\big] + \beta\pi_{s0} E\big[B(a', 0, 0)\big] \Big\} \tag{12}$$

6. This assumption is in the spirit of Nakajima and Telyukova (2020) who also tighten borrowing constraints with age. This is also a necessary condition to ensure mortgagor rates become extinct at a rate consistent with the data.

subject to

$$c + (\delta_R + r)hp_R + a' = y + (1+r)a - e_M \tag{13}$$

$$a' \geq 0 \tag{14}$$

$$e_M = \begin{cases} 0 & \text{if } EL = 1 \\ e & \text{otherwise} \end{cases} \tag{15}$$

$$V^R_{T+1} = B(a', 0, 0) \tag{16}$$

where t denotes age. Aside from the discrete housing choice, this problem amounts to a standard consumption-savings problem with incomplete markets.

The value function may not be differentiable for all a', resulting from the discrete housing choice. However, since the utility function remains concave in h, first-order conditions are still informative and useful for computations. The Lagrangian can be written as

$$\begin{aligned} V^R_t(a : H, s, e, i) \equiv \\ \max_{c,h,a'} \Big\{ & u(c, h, s, 0) + \beta E\big[V^R_{t+1}(a' : H', s', e', i)\big|H, s, e\big] + \beta \pi_{s0} E\big[B(a', 0, 0)\big] \\ & + \mu^R\big[y + (1+r)a - e_M - (\delta_R + r)hp_R - c - a'\big] + \lambda^R a' \Big\}. \end{aligned} \tag{17}$$

Home buying renters face a similar problem except for the additional home purchase decision. Home purchase for the buyer is made within the period so that the household must choose $\bar{h}$ and a grid of down-payments, $F \in \mathbb{F} = \{20\%, 50\%, 75\%, 100\%\}$, which will set the households debt $d' = (1 + f_1)(1 - F)h + f_0$. Additionally, the household can only qualify for a loan if the loan satisfies the DTI cap. Home buyers solve the problem

$$\begin{aligned} V^R_t(a : H, s, e, i) \equiv \\ \max_{c,h,a',d' \in d(\mathbb{F})} \Big\{ & u(c, h, s, 0) + \beta E\big[V^O_{t+1}(a'd', \bar{h} : H', s', e', i)\big|H, s, e\big] + \beta \pi_{s0} E\big[B(a', d', \bar{h})\big] \Big\} \end{aligned} \tag{18}$$

subject to 10 and

$$c + (\delta_R + r)hp_R + a' = y + (1+r)a - e_M - F\bar{h}p \tag{19}$$

$$a' \geq 0 \tag{20}$$

$$e_M = \begin{cases} 0 & \text{if } EL = 1 \\ e & \text{otherwise} \end{cases} \tag{21}$$

$$V^R_{T+1} = B(a', d', \bar{h}) \tag{22}$$

Similarly, the Lagrangian can be written as

$$V_t^R(a:H,s,e,i) \equiv \max_{c,h,a',d' \in d(\mathbb{F})} \Big\{ u(c,h,s,0) + \beta E\big[V_{t+1}^O(a'd', \bar{h}:H',s',e',i)\big|H,s,e\big] + \beta \pi_{s0} E\big[B(a',0,0)\big] + \mu^R\big[y+(1+r)a - e_M - F\bar{h}p - (\delta_R+r)hp_R - c - a'\big] + \lambda^R a' \}. \quad (23)$$

Since the home buyer only chooses loans discretely on a grid, the Euler equations for both Lagrangians can be written in general as

$$u_c(c,h,s,0) = \beta(1+r)E\big[u(c',\bar{h},s',0)\big|H,s,e\big] + \beta\pi_{s0}E\big[B_a(a',d',\bar{h})\big] + \lambda^R \quad (24)$$

$$\lambda^R a' = 0 \quad (25)$$

3.7 Sellers

Sellers are households whom in the previous period elected to liquidate their homes. At the beginning of the period, households receive the proceeds from their house sale $(1-\kappa)hp$ and must settle their debts. The term $1-\kappa$ captures the fact that house sale is costly. Otherwise, home owners effectively face the same problem as renters and it depends again on whether the sale is to downsize or to rent. Sellers becoming renters solve the problem

$$V_t^S(a,d,\bar{h}:H,s,e,i) \equiv \max_{c,h,a'} \Big\{ u(c,h,s,0) + \beta E\big[V_{t+1}^R(a':H',s',e',i)\big|H,s,e\big] + \pi_{s0}E\big[B(a',0,0)\big] \Big\} \quad (26)$$

subject to

$$c + (\delta_R+r)hp_R + a' = y + (1+r)a + (1-\kappa)\bar{h}p - d - e_M \quad (27)$$

$$a' \geq 0 \quad (28)$$

$$e_M = \begin{cases} 0 & \text{if } EL = 1 \\ e & \text{otherwise} \end{cases} \quad (29)$$

$$V_{T+1}^S = B(a_{T+1},0,0). \quad (30)$$

Equation 27 is the only effective change in the problem. As such, the FOCs remain the same for sellers as they do for renters. It is however still worth formulating the Lagrangian in this

case since it has implications for the Owner's problem. The Lagrangian can be written as

$$V_t^S(a, d, \bar{h} : H, s, e, i) \equiv$$
$$\max_{c,h,a'} \Big\{ u(c, h, s, 0) + \beta E\big[V_{t+1}^R(a' : H', s', e', i)\big|H, s, e\big] + \pi_{s0} E\big[B(a', 0, 0)|r\big]$$
$$+ \mu^R\big[y + (1+r)a - e_M - (\delta_R + r)hp_R + (1-\kappa)\bar{h}p_O - d - c - a'\big] + \lambda^R a' \Big\}. \quad (31)$$

For a seller downsizing, the household's problem can be written as

$$V_t^S(a, d, \bar{h} : H, s, e, i) \equiv$$
$$\max_{c,\bar{h}',a'} \Big\{ u(c, \bar{h}', s, 1) + \beta E\big[V_{t+1}^O(a'd', \bar{h}' : H, s', e', i)\big|H, s, e\big] + \pi_{s0} E\big[B(a', d', \bar{h}')\big] \Big\} \quad (32)$$

subject to equation 10 and

$$c + a' + (\delta + F)\bar{h}'p = y + (1+r)a + (1-\kappa)\bar{h}p - d - e_M \quad (33)$$
$$a' \geq 0 \quad (34)$$
$$e_M = \begin{cases} 0 & \text{if } EL = 1 \\ e & \text{otherwise} \end{cases} \quad (35)$$
$$V_{T+1}^S = B(a_{T+1}, d', \bar{h}'). \quad (36)$$

Notice how in the seller's case they move into the new home at the time of home sale. This means that at the same time households make their home sale, they can actually decide which home to purchase. This avoids needing additional decision nodes in the owner's problem while still allowing downsizing to happen seamlessly. Since the financing choice is discrete, the Euler equations for both problems can be written as

$$u_c(c, h, s, o) = \beta(1+r)E\big[u(c', h', s', o)\big|H, s, e\big] + \beta\pi_{s0}E\big[B_a(a', d', h')\big] + \lambda^S \quad (37)$$
$$\lambda^S a' = 0. \quad (38)$$

While this is effectively the same problem as the renter apart from the additional resources from house sale, the derivative of the value function with respect to debt is actually useful for the Owner's problem dynamics. In fact, this will be part of the Envelope theorem conditions used for the owner's problem. Assuming the seller previously maximized their value function

with respect to debt, this yields

$$V^S_{d,t}(a, d, \bar{h} : s, e, i) = -(1 + r_m)\mu^S. \tag{39}$$

In other words, the marginal value of decreasing debt for a seller is exactly the marginal value of consumption at the mortgage rate.

3.8 Owners

Owners are the focus of this paper. Since mortgages are closed-end contracts, costly refinancing frictions prevent households from freely switching between both objects which generates the co-holding behaviour the paper seeks to explain.

Each period, an owner makes a discrete choice among their three options: continue with existing mortgage (owner-keepers OK), refinance into a new contract (owner-refinancers OR), or sell next period (owner-sellers OS). Each decision has its own specific tradeoff. Keeping the mortgage preserves the existing contract but limits the household's ability to access housing wealth. Refinancing allows the household to tap extra liquidity but comes at a cost. Selling liquidates the entire housing position.

The owner's problem retains the same set of state variables as in the seller's problem, but adds debt, d, to the choice set. Owners either selling or continuing with the mortgage contract are required to make at least a minimum payment on debt, $m > \underline{m}$, derived from the schedule 6. Thus, mortgage balances evolve according to $d' = (1 + r_m)d - m$. Refinancers do not face the minimum payment constraint however, they pay a the refinancing costs. Debt evolution and refinancing costs embed the most important friction in the model: paying down debt requires an outflow of liquid assets while tapping home value requires either costly refinancing or sale, while access to debt may be constrained by LTV, DTI, or health. Therefore, liquid assets and mortgage debt are not perfect substitutes.

At the beginning of a period, the household must decide whether they would like to make mortgage payments on their existing mortgage, negotiate a new loan, or sell their home. The owner's value function can be written as

$$V^O(\boldsymbol{x}) \equiv \max\left\{V^{OS}(\boldsymbol{x}), V^{OK}(\boldsymbol{x}), V^{ON}(\boldsymbol{x})\right\} \tag{40}$$

where $\boldsymbol{x} = (a, d, \bar{h} : H, s, e, i)$. The set $\left\{V^{OS}(\boldsymbol{x}), V^{OK}(\boldsymbol{x}), V^{ON}(\boldsymbol{x})\right\}$ refers to each different problem structure that depends on the household's tenure choice. Each problem is defined explicitly below.

3.8.1 Owner-Sellers

Owner-sellers are home owners deciding to put their homes up for sale. Since home sale occurs in the period following the sale decision, owner-sellers are still owners within the current period, but the future value of this decision is the seller's value function. The value function for this choice can then be written as

$$V_t^{OS}(a, d, \bar{h}, : H, s, e, i) \equiv \max_{c,m,a',d'} \Big\{ u(c, h, s, 1) + \beta E\big[V_{t+1}^{S}(a', d', \bar{h} : H', s', e', i)\big|H, s, e\big] + \beta \pi_{s0} E\big[B(a', d', \bar{h})\big] \Big\} \quad (41)$$

subject to

$$c + \delta hp + m + a' = y + (1 + r)a - e_M \quad (42)$$

$$d' = (1 + r_m)d - m \quad (43)$$

$$a' \geq 0 \quad (44)$$

$$e_M = \begin{cases} 0 & \text{if } EL = 1 \\ e & \text{otherwise} \end{cases} \quad (45)$$

$$m > \underline{m} \quad (46)$$

$$d' \geq 0 \quad (47)$$

$$V_{T+1} = B(a', d', \bar{h}). \quad (48)$$

As in the simple model of Buck (2026), since d' and m are co-dependent, the Lagrangian can be simplified and written in terms of d and d' instead. The term δhp shows up in the budget constraint because households are required to perform maintenance to offset depreciation, δ. Again, keep in mind that the value functions will not be differentiable everywhere, so first-order conditions will only be valid in smooth neighborhoods. Since the value function is concave, the intuition is is still valid. The Lagrangian can be written as

$$V_t^{OS}(a, d, \bar{h} : H, s, e, i) \equiv \max_{c,a',d'} \Bigg\{ u(c, h, s, 1) + \beta E\big[V_{t+1}^{S}(a', d', \bar{h} : H', s', e', i)\big|H, s, e, r\big] + \beta \pi_{s0} E\big[B(a', d', \bar{h})\big] + \mu^O\big[y + (1 + r)a - (1 + r_m)d - e_M - \delta hp - c - a' + d'\big] + \lambda^O a' + \theta d' + \phi\big[\bar{d} - d'\big] \Bigg\}. \quad (49)$$

where $\bar{d} = (1 + r_m)d - \underline{m}$. For sellers, the Euler equations can be written as

$$u_c(c, \bar{h}, s, 1) = \beta(1 + r)E\big[u_c(c', h', s', 0), \big|H, s, e\big] + \beta\pi_{s0}E\big[B_a\big] + \lambda^O \tag{50}$$

$$\theta + u_c(c, \bar{h}, s, 1) = \beta(1 + r_m)E\big[u_c(c', h', s', 0)\big|H, s, e\big] - \beta\pi_{s0}E\big[B_d\big] + \phi \tag{51}$$

$$\lambda^O a' = 0 \tag{52}$$

$$\phi\big[\bar{d} - d'\big] = 0 \tag{53}$$

$$\theta d' = 0 \tag{54}$$

$$\theta\phi = 0. \tag{55}$$

The term $\beta E\big[u_c(c', h', s', 0)\big]$ from equation 51 derives directly from 39. This is due to the fact that mortgage debt is compounded within period based on outstanding debt at the beginning of the period. Furthermore, since $r_m > r$ and $B_a < B_d$ it will always be the case that $\theta > 0$ and $d' = 0$ for a seller if feasible. This implies that $\phi = 0$ and sellers make pay down as much debt as possible.

3.9 Owner-Keepers

Owner-keepers are households who decide to keep their existing mortgages. The only variation from the owner-seller problem is that now the future value function is equal to the owner's value function. Additionally, with some probability $\rho(H, s)$, households health and age forces the household into the sale branch. This is baked into the expectation by virtue of the $\mathcal{O}$ superscript: with probability $\rho(H, s)$, $\mathcal{O} = O$ and with probability $1 - \rho(H, s)$ $\mathcal{O} = S$. The problem can then be written as

$$\begin{aligned} V_t^{OK}(a, d, \bar{h}, : H, s, e, i) \equiv \max_{c,m,a',d'} \Big\{ & u(c, h, s, 1) \\ & + \beta E\big[V_{t+1}^{\mathcal{O}}(a', d', \bar{h} : H', s', e', i)\big|H, s, e\big] + \beta\pi_{s0}E\big[B(a', d', \bar{h})\big]\Big\} \end{aligned} \tag{56}$$

subject to

$$c + \delta hp + m + a' = y + (1+r)a - e_M \tag{57}$$

$$d' = (1+r_m)d - m \tag{58}$$

$$a' \geq 0 \tag{59}$$

$$e_M = \begin{cases} 0 & \text{if } EL = 1 \\ e & \text{otherwise} \end{cases} \tag{60}$$

$$m \geq \underline{m} \tag{61}$$

$$d' \geq 0 \tag{62}$$

$$V_{T+1} = B(a', d', \bar{h}). \tag{63}$$

The associated Lagrangian is

$$\begin{aligned} V_t^{OK}(a, d, \bar{h} : H', s, e, i) \equiv \max_{c,a',d'} \Bigg\{ & u(c, \bar{h}, s, 1) \\ & + \beta E\big[V_{t+1}^{OK}(a', d', \bar{h} : H', s', e', i)\big|H, s, e\big] + \beta \pi_{s0} E\big[B(a', d', \bar{h})\big] \\ & + \mu^O\big[y + (1+r)a - (1+r_m)d - e_M - \delta \bar{h} p - m - c - a' + d'\big] \\ & + \lambda^O a' + \theta d' + \phi\big[\bar{d} - d'\big] \Bigg\}. \end{aligned} \tag{64}$$

Before computing the Euler equations for future owners, it is worth discussing an envelope theorem condition. At the optimum, the derivative of the value function with respect to debt is

$$-V_{d,t}^O(\boldsymbol{x}) = (\mu^O - \phi)(1+r_m)^7 \tag{65}$$

implying there are some shadow values to paying off debt. This result refers to the same mechanism that leads to co-holding in the simple model of Buck (2026). The marginal value to paying down debt in the previous period is equal to the additional consumption gained from owing less future interest minus the shadow value of lowering future mortgage headroom. In simple terms, paying down mortgage debt saves interest payments however, it comes at the cost of lost liquidity flexibility in the future.

Equation 65 is the derivative of the value function for an owner-keeper but in reality, depending on the future draws on expense shocks and interest rates, the future owner could be either an owner-keeper or an owner-refinancer. Therefore, the value of future mortgage

7. Note that here I am using the negative derivative to infer the change in value from a decrease in mortgage debt.

headroom will only exist when households do not anticipate refinancing in the next period. Let $N(\boldsymbol{x}')$ refer to the future decision to sell, keep, or refinance the mortgage loan where $N(\boldsymbol{x}') = 0$ if re-financing and $N(\boldsymbol{x}') = 1$ otherwise. Then, the envelope theorem condition as

$$-V^{O}_{d,t+1}(\boldsymbol{x}') = (\mu' - \phi')(1 + r_m)N(\boldsymbol{x}') \tag{66}$$

The Euler equations for an owner who does not negotiate a new loan can then be written as

$$u_c(c, \bar{h}, s, 1) = \beta(1 + r)E\big[u_c(c', \bar{h}, s', 1)|H, s, e\big] + \beta\pi_{s0}E[B_a] + \lambda^O \tag{67}$$

$$\begin{aligned}\theta + \beta(1 + r_m)E[\phi'|H, s, e] + u_c(c, \bar{h}, s, 1) \\ = \beta(1 + r_m)E\big[u_c(c', \bar{h}, s', 1)|H, s, e\big] - \beta\pi_{s0}E[B_d]\end{aligned} \tag{68}$$

$$\phi\big[\nu_\tau d - d'\big] = 0 \tag{69}$$

$$\theta d' = 0 \tag{70}$$

$$\theta\phi = 0. \tag{71}$$

In terms of co-holding, these Euler equations imply the same behaviour as in the three-period model of Buck (2026). The difference here is that while mortgage contracts are closed-end, $E[\phi'|H, s, e]$ takes into account the flexibility offered by refinancing. In a sense, mortgage contracts are pseudo-closed-ended because while a single mortgage contract is closed-ended, it can also be resolved by taking out a new loan. In contrast to the simple model, here, mortgage contracts now acts as a friction that leads to co-holding rather than a strict barrier. This will imply a weaker but still important incentive to co-hold.

3.9.1 Owner-Refinancers

Owner-refinancers face the same value functions as owner-keepers with a slightly different set of constraints. Instead of a mortgage headroom constraint, since they are refinancing, these households choose a level of future debt no higher than 60% of their current home value. However, the lender takes a fee from their new loan at the price $df + f_0$, and the household is then required to absolve any outstanding debt. Since the household refinances, they are not required to make the minimum payment on debt. The problem can be written as

$$\begin{aligned}V^{OR}_t(a, d, \bar{h} : H, s, e, i) \equiv \max_{c,a',d'} \Big\{ u(c, h, s, 1) \\ + \beta E\big[V^{\mathcal{O}}_{t+1}(a', d', \bar{h} : H, s', e', i)\big|H, s, e\big] + \beta\pi_{s0}E\big[B(a', d', \bar{h})\big]\Big\}\end{aligned} \tag{72}$$

subject to

$$c + \delta hp + a' = y + (1+r)a - (1+r_M)d - e_M + (1-f)d' - f_0 \tag{73}$$

$$d' \leq \bar{D} \tag{74}$$

$$a' \geq 0 \tag{75}$$

$$e_M = \begin{cases} 0 & \text{if } EL = 1 \\ e & \text{otherwise} \end{cases} \tag{76}$$

$$d' \geq 0 \tag{77}$$

$$V_{T+1} = B(a', \bar{h}, d'). \tag{78}$$

The associated Lagrangian becomes

$$\begin{aligned} V_t^{OR}(a, d, \bar{h} : H, s, e, i) \equiv \max_{c,a',d'} \Big\{ & u(c, \bar{h}, s, 1) \\ & + \beta E\big[V_{t+1}^{\mathcal{O}}(a', d', \bar{h} : H', s', e', i)\big|H, s, e\big] + \beta \pi_{s0} E\big[B(a', d', \bar{h})\big] \\ & + \mu^O\big[y + (1+r)a - (1+r_m)d - e_M - \delta\bar{h}p - c - a' + d'(1-f) - f_0\big] \\ & + \lambda^O a' + \theta d' + \phi\big[0.6\bar{h}p - d'\big] \Big\}. \end{aligned} \tag{79}$$

This latest problem will modify the debt envelope theorem condition where instead,

$$-V_{d,t+1}^{O}(\boldsymbol{x}') = \mu'(1+r_m) \tag{80}$$

$$u_c(c, \bar{h}, s, 1) = \beta(1+r)E\big[u_c(c', \bar{h}, s', 1)\big|H, s, e\big] + \beta\pi_{s0}E[B_a] + \lambda^O \tag{81}$$

$$\theta + (1-f)u_c(c, \bar{h}, s, 1) = \beta(1+r_m)E\big[u_c(c', \bar{h}, s', 1)\big|H, s, e\big] - \beta\pi_{s0}E[B_d] + \phi \tag{82}$$

$$\lambda^O a' = 0 \tag{83}$$

$$\phi\big[\nu_\tau d - d'\big] = 0 \tag{84}$$

$$\theta d' = 0 \tag{85}$$

$$\theta\phi = 0. \tag{86}$$

The main difference in the Euler equations is that $(1-f)$ appears on the left-hand-side instead the ϕ' term. While not explicitly part of the choice set, since the refinancing fee applies to households on the entire debt position d, an owner will never refinance if $d' < (1+r_m)d - \underline{m}$, since this would be costlier than paying down debt.

4 Calibration

The model is calibrated in two stages. The first stage pins down parameters that can be directly inferred from the data or institutionally set parameters. The second stage estimates the remaining unidentified parameters using the simulated method of moments (SMM). Each model time period spans two years, and all dollar values are computed in terms of real 2022 USD. The data for calibration is sourced from the Rand HRS longitudinal files from 1998-2022 and FRED time series for various aggregate variables. A summary table of parameters and their treatment appears at the end of each subsection. Appendix A.1 provides clarifications on the methodologies and is referenced throughout.

4.1 Externally Calibrated Parameters

Income process. The model's income process is constructed from HRS data as a deterministic flow y indexed by entry quintile i, marital status s, and age t. Raw HRS total income includes both labour earnings and income on liquid wealth. Including labour income would be inconsistent with permanent retirement income and including liquid asset earnings would be a double counting of the return on a. Therefore, I define $y^{model} = \max\{0, y^{total} - ra_{+} - y^{earn}\}$ which preserves social security, pension income, and retirement-account withdrawals. While an endogenous draw-down of retirement savings would be preferred, this still preserves some of the value of retirement savings. Each household is then assigned a permanent income type by ranking y^{model} at ages 65-67 by quintiles using HRS sampling weights and excluding non-retired households[8]. Then, for each age, t, quintile-type, i, and marital status, s, each cell in y is generated as the weighted median within a year of the target age. Each age profile conditional on i and s is then smoothed using an isotonic regression to enforce monotone-decreasing paths[9]. The isotonic regression is also weighted by cell sample size so that noisy thin cells don't drag the fitted path. Under the same logic as the isotonic regression, I enforce that coupled income is at least as high as singled income by setting coupled equal to the maximum of both coupled and single income. This restriction ensures that widowhood never leads to an increase in income. Lastly, because households have expenses that are not modeled here, the income in y^{model} is not necessarily the model's

8. Because earnings are subtracted in y^{model}, a working 65-year-old with substantial labour income has $y^{model} \approx 0$ and gets bucketed into Q1 despite being middle class. This exclusion prevents this and keeps the meaning of i consistent.

9. Isotonic regressions removes non-economic variations across age resulting from attrition, cohort effects, and other noise that could lead to income fluctuations across the life-cycle. To avoid small/missing samples, every cell with fewer than 30 observations falls back to the median over marital status across age and quintile. Any remaining missing cell is then imputed from the nearest non-missing age. See appendix A.1 for the specification.

Table 1: Income quintiles by age and marital status.

		Age 65		Age 75		Age 85	
Status	Quintile	Gross	Disp.	Gross	Disp.	Gross	Disp.
Single	Q1	25.8	18.7	25.8	18.7	25.8	18.7
	Q2	39.6	31.1	39.1	30.7	35.2	27.1
	Q3	62.9	52.1	54.6	44.7	45.3	36.2
	Q4	106.8	91.6	83.8	70.9	63.0	52.2
	Q5	181.0	158.4	114.0	98.1	98.2	83.9
Couple	Q1	45.4	32.4	45.4	32.4	45.4	32.4
	Q2	53.4	39.6	53.4	39.6	53.4	39.6
	Q3	76.3	60.2	76.3	60.2	76.3	60.2
	Q4	112.1	92.4	108.9	89.5	101.2	82.6
	Q5	209.7	180.2	174.0	148.1	162.3	137.6

Note: Values are reported in 2022 $1,000 values.

Table 2: Prices and rates.

Parameter	Symbol	Annual	Model (2-yr)	Source
Risk-free rate	r	-0.23%	-0.45%	FRED TB3MS
Mortgage rate	r_m	3.17%	6.43%	FRED MORTGAGE30US
House-price growth	g_H	2.16%	4.37%	FRED CSUSHPINSA
Rent growth	g_R	1.02%	2.05%	FRED rent index
Renter depreciation	δ_R	2.00%	4.04%	Harding et al. (2007)

Note: Prices and rates are reported in real values deflated by CPIAUCSL.

disposable income, and mapping y^{model} directly onto $y_t(i,s)$ gives households far too much liquidity. To reflect disposable income, I subtract a 10% tax on income above $25k biannual, Medicare Part B + Medigap premiums of $4k single/$8k couple biannual, and a $3k property tax and insurance applied uniformly. Table 1 reports both gross and disposable income quintiles.

Prices and rates. The mortgage rate, r_m, and the housing/rent price growth rate, g, are both constructed as 1998-2022 averages from FRED. The mortgage rate uses 30-year mortgage rates (MORTGAGE30US), and house price growth uses the S&P Cotality Case-Shiller U.S. National Home Price Index (CSUSHPINSA) respectively. Real rates are deflated to 2022 USD from the FRED consumer price index for all urban consumers (CPIAUCSL). The renter depreciation rate δ_R is set to 2% annually, a standard value from the literature. The externally calibrated prices and rates are reported in table 2.

Because household liquid holdings are comprised of different portfolio objects (e.g., checking and savings accounts, bonds, stocks, etc.) the risk free interest rate is computed conditional on income type, i[10]. To compute conditional liquid asset returns, first I sample each household with at least \$1,000 in gross value across checking accounts, CDs, bond, and stock holdings. Then, across qualifying observations, I compute within income type weighted averages of portfolio holdings. For example, a household with \$1,000 in checking accounts and \$1,000 in stocks would be coded as $0.5, 0.0, 0.0, 0.5$. The aggregate results from the portfolio composition exercise are reported in table 3 In addition to computing holdings,

Table 3: Liquid-portfolio composition by income quintile (HRS).

	Checking	CDs	Bonds	Stocks	Holding liquid
Real return (annual)	−2.44%	0.07%	1.78%	6.82%	
Q1	78%	7%	2%	13%	36%
Q2	77%	9%	1%	13%	56%
Q3	71%	9%	2%	18%	76%
Q4	66%	10%	2%	22%	87%
Q5	52%	10%	4%	34%	94%

Note: Component shares are household-average portfolio weights among households holding a liquid bundle above \$1,000 (2022 USD), weighted by HRS sampling weights, pooled over 1998–2022, ages 65–95. "Holding liquid" is the weighted share of all households in the quintile above that threshold. Component returns are real 1998–2022 arithmetic annual means, deflated by CPIAUCSL, following the conventions of the prices-and-rates table. Income quintiles use the same cuts as the income process.

I use combinations of external data sources to infer returns on each type of holding. First, since checking accounts typically have no return, the nominal return on checking accounts is set to 0%. Second, since the Rand HRS file bundles CDs with government bonds and T-bills, I use the FRED TB3MS to infer returns. Both bond and stock returns are calculated by averaging yearly returns from Damodaran (2023). Returns are also reported in table 3 and are deflated to real terms using FRED CPIAUCSL.

Liquid returns conditional on income quintile are then computed using the average of each asset return weighted by the average holdings of each asset by income quintile. However, not all these assets are risk free (in particular stock holdings), but the model only allows households to store a risk free asset. To account for this, I measure certainty equivalents for each income quintile portfolio rather than just its time series mean. For a CRRA household with curvature σ, a second-order approximation to expected utility yields the Arrow–Pratt

10. Ideally, households would be allowed to hold different portfolio objects, but this would be too computationally burdensome.

risk premium (Pratt 1964; Arrow 1971). Therefore, the calibrated return is

$$r(i,\sigma) = \sum_c s_c(i)\, r_c \;-\; \frac{1}{2}\,\sigma\big(s_{stk}(i)\,\sigma_{EQ}\big)^2 \tag{87}$$

where s_c is the share of each asset type, r_c is the return for each asset type, s_{stk} is the share of stocks, σ is the coefficient of relative risk aversion (calibrated in the second stage), and σ_{EQ} is the standard deviation of stock returns[11]. In other words, the return on liquid assets measures the household-average portfolio return less a risk premium haircut of stock holdings. The final results of the exercise are compiled in 4 . Because the model reports using real

Table 4: First-stage calibrated liquid returns $r_a(i)$ by income quintile.

	Portfolio blend	CE haircut	$r_a(i)$ annual	$r_a(i)$ biannual
Q1	−0.99%	0.08%	−1.07%	−2.12%
Q2	−0.94%	0.08%	−1.02%	−2.03%
Q3	−0.48%	0.14%	−0.63%	−1.25%
Q4	−0.10%	0.22%	−0.31%	−0.62%
Q5	1.18%	0.54%	0.64%	1.27%

Note: The portfolio blend is $\sum_c s_c(i)\, r_c$ over the components and returns of the composition table. The certainty-equivalent haircut is $\frac{1}{2}\sigma\big(s_{stk}(i)\,\sigma_{EQ}\big)^2$ with $\sigma = 2.69$ and $\sigma_{EQ} = 18.4\%$, deducting the risk premium the model's riskless liquid asset cannot carry. Biannual values convert as $(1 + r_a)^2 - 1$. $r_a(i)$ is pinned in the first stage from portfolio composition and market returns only; it does not enter the SMM parameter vector.

terms, households with portfolios dominated by checking account holdings inherit much more negative real returns.

Health process. In the model, households are indexed as either healthy, $H = G$, moderate, $H = M$, bad health, $H = B$, or long-term care, $H = L$. In the data, a household in a given year is considered healthy if they report excellent/very good health, moderate if they report good/fair health, and poor health if they report poor health. In the event that they spend over 90 nights in a nursing home within a year, their health status is overwritten to $H = L$. Health transitions are then computed from the age-indexed transition kernel

$$\Gamma(H_{t+1}|H_t) \equiv \Pr(H_{t+1}|H_t,\ \text{age}_t,\ \text{alive at } t \text{ and } t+1) \tag{88}$$

so that health is computed out of the set of survivors in each survey year. Let the sample in

11. While this technically makes this object part of the second stage calibration, since s_{stk} and σ_{EQ} are both much lower than 1, deviations in σ have little affect on the total figure. See appendix A.1 for clarification.

Table 5: Health transition probabilities.

From \ To	Healthy	Moderate	Poor	LTC
Age 65				
Healthy	0.750	0.242	0.008	0.000
Moderate	0.183	0.746	0.063	0.007
Poor	0.025	0.441	0.512	0.022
LTC	0.161	0.645	0.018	0.176
Age 75				
Healthy	0.682	0.297	0.014	0.006
Moderate	0.165	0.744	0.080	0.012
Poor	0.033	0.351	0.576	0.041
LTC	0.046	0.498	0.306	0.150
Age 85				
Healthy	0.583	0.364	0.020	0.033
Moderate	0.149	0.703	0.098	0.050
Poor	0.049	0.356	0.458	0.137
LTC	0.078	0.465	0.194	0.263

Note: Calibrated from RAND HRS 1998-2022.

a given survey year, τ, be written such that

$$S_\tau = \{(l,t) \mid \text{alive}_{l,t+1} = \text{alive}_{l,t} = 1,\ H_{l,t+1},\ H_{l,t},\ \text{age}_\tau - 1 \leq \text{age}_{l,t} < \text{age}_\tau + 1\}. \tag{89}$$

Transition probabilities are then estimated from the data using the pseudo-maximum-likelihood estimator

$$\hat{\Gamma}(H_{t+1}|H_t) \equiv \frac{\sum_{j,t\in S_\tau} w_j \mathbf{1}\{H_{j,t+1}\}\mathbf{1}\{H_{j,t}\}}{\sum_{j,t\in S_\tau} w_j \mathbf{1}\{H_{j,t}\}}. \tag{90}$$

where w_j is the HRS household-level sampling weight for household j. Table 5 reports the Markov chain for health status.

Additionally, the model takes advantage of health states to capture the risk of forced home sale. In the data, I find highly liquid home owners who would not have financial incentives to move and difference them across health states to identify relocation hazard from health. The results for relocation hazard are in table 6 Interestingly, households in long term care are not highly likely to relocate. This may be because high-liquid LTC owners benefit from the homestead exemption which protects them from the asset test.

Demographics. Marital and mortality transitions follow a Markov chain across the states $s \in \{0,1,2\}$ for deceased, single, and coupled respectively. Both deceased and single

Table 6: Health-driven forced-relocation hazard $\rho_t(H, s)$ by age band.

	Single				Couple			
Health state	Pooled	65–74	75–84	85+	Pooled	65–74	75–84	85+
Healthy (baseline)	0.042	0.024	0.047	0.087	0.018	0.010	0.027	0.079
Moderate	0.033	0.007	0.024	0.081	0.009	0.000	0.007	0.033
Poor	0.039	0.025	0.061	0.081‡	0.027	0.014	0.013	0.044
LTC	0.060	0.060†	0.080*	0.081‡	0.034*	0.034†	0.034†	0.044†‡
Owner-wave pairs	30,399	12,794	12,793	4,812	34,097	21,299	11,113	1,685

Note: Per-two-year probability that a surviving owner is forced to sell for health reasons, applied to the owner-keeper branch on top of endogenous selling. Estimated on owner-wave pairs alive at both waves, among owners above the 75th percentile of liquid wealth ($159k, 2022 USD), so that sales are not financially motivated, and differenced against healthy owners of the same marital status in the same age band. The healthy row reports the raw hazard of healthy high-liquid owners, which is subtracted from every other row, so $\rho = 0$ for healthy owners by construction. The 85+ band applies to all model ages from 85 to 99. *Cell below 30 observations; uses the 50th-percentile cut ($27k) with the healthy baseline at the same cut. †Band cell below 20 observations; uses the pooled estimate. ‡Raised to the value of the next-healthier state so that ρ is non-decreasing in severity. Calibrated from RAND HRS 1998–2022.

are one-directional. Deceased is of course an absorbing state, but I do not model marriage decisions so single households never re-marry. Transition probabilities are estimated from the rand HRS longitudinal file and are computed conditional on age, t, income quintile-type, i, and health status H, to capture the empirical income-mortality gradient (Chetty et al. 2016) following the conditional structure of De Nardi, French, and Jones (2010). For couples, I estimate sex specific mortality hazards $\mu_m(t, i, H)$ and $\mu_f(t, i, H)$ and combine them by imposing a conditional independence assumption such that

$$\pi_{20} = \mu_m \mu_f \tag{91}$$

$$\pi_{21} = \mu_m + \mu_f - 2\pi_{20} \tag{92}$$

$$\pi_{22} = 1 - \pi_{20} - \pi_{21}. \tag{93}$$

Single household mortality, $\pi_{10}(t, i, H)$, is estimated directly from the pooled sample of single households[12].

Table 7 reports the calibrated life expectancy of households aged 65 across the health-income gradient. The gap in life expectancy across health status is relatively constant conditional on each income quintile. This is expected because health should be the biggest indicator of life expectancy, which is supported by the income gap in life expectancy conditional

12. Conditional independence simply means the probability of survival is independent of spousal survival, a typical modeling assumption for mortality risk with coupled households (Braun, Kopecky, and Koreshkova 2017)

Table 7: Life expectancy at 65 by income quintile and health state

	Health state				
Quintile	Healthy	Moderate	Poor	LTC	Health gap
Q1	88.7	82.1	73.3	71.7	+17.0
Q2	89.2	82.7	73.7	72.0	+17.2
Q3	89.4	82.9	73.9	72.1	+17.2
Q4	89.8	83.4	74.3	72.4	+17.4
Q5	90.8	84.5	75.4	73.2	+17.6
Income gap	+2.1	+2.4	+2.2	+1.5	

Note: Calibrated from RAND HRS 1998-2022.

on health. Healthy-wealthy individuals cannot significantly raise their life expectancies by virtue of being higher income compared to healthy poor households. That the gradient in income is stronger across moderate and poor health is indicative of the fact that wealthy households are more likely to move out of their poor health, or preserve their status quo longer.

Medical expense process. Medical expense risk is calibrated following De Nardi, French, and Jones (2010), extended so that both the level and the dispersion of expenses are conditioned on health. For the three community-dwelling health states, $H \in \{G, M, B\}$, I estimate

$$\ln e^m_{j,t} = \alpha(i_j) + \beta(i_j) \cdot (\text{age}_{j,t} - 60) + \gamma(H_{j,t}) + \zeta_{j,t}, \tag{94}$$

where $e^m_{j,t}$ is household j's biannual out-of-pocket (OOP) medical spending, $\alpha(i)$ and $\beta(i)$ are an intercept and per-year age slope specific to the household's permanent income quintile, and $\gamma(H)$ is a health-specific level shift with $\gamma(G) = 0$ as the normalization. Equation (94) is estimated in a single weighted least-squares regression on roughly 47,000 positive-OOP retiree household-waves.[13] Table 8 reports the estimates.

The residual $\zeta_{j,t}$ follows an AR(1),

$$\zeta_{j,t} = \rho\, \zeta_{j,t-1} + \eta_{j,t}, \tag{95}$$

with a common persistence parameter across health states and a health-specific stationary dispersion $\sigma_\zeta(H)$. I estimate ρ in a second step as the weighted through-origin regression

13. The sample restricts to retired households aged 60–95 with observed health and OOP spending of at least \$1,000 weighted by HRS sampling weights. Households in long-term care are excluded from the regression and handled separately below. Income quintiles are the permanent types assigned at ages 65–67, so households first observed after age 67 are excluded, consistent with the income and demographic processes.

Table 8: The estimated medical-expense process

Panel A: variation by health state				
Health state	$\gamma(h)$	$\sigma(h)$	Median ($)	Catastr. ($)
Healthy	0.00	1.31	2,167	29,755
Moderate	+0.20	1.35	2,644	38,977
Poor	+0.44	1.43	3,366	59,188
LTC	–	–	–	138,286
Panel B: process parameters				
Shock persistence ρ (biannual)				0.50
Income intercepts e^{α_q} (base)				\$1,443–\$3,303
Income age slopes β_q				-0.015 to -0.001 (near-flat)
LTC anchor (custodial)				461 nights × \$300/day
Estimation sample				47,728 person-waves

Note: Calibrated from RAND HRS 1998-2022. Monetary value reported in 2022 USD.

of residuals on their own lag, using only consecutive-wave pairs, which yields $\hat{\rho} \approx 0.50$ at the biannual frequency (an annual equivalent of $\sqrt{\hat{\rho}} \approx 0.71$, below the 0.922 of De Nardi, French, and Jones 2010). The stationary dispersion is the weighted standard deviation of the level-purged residuals within each health state, $\sigma_\zeta(G) = 1.31$, $\sigma_\zeta(M) = 1.35$, $\sigma_\zeta(B) = 1.43$: worse health widens the distribution of expense shocks more than it raises their mean. This asymmetry matters for the model's mechanics: a common dispersion would leave the poor-health right tail, which drives catastrophic-expense risk, far too thin.

The shock is discretized into five states with the method of Rouwenhorst (1995), which Kopecky and Suen (2010) show dominates alternative discretizations for persistent processes. Because the Rouwenhorst transition matrix depends only on (ρ, n), and ρ is common across health states, all households share a single transition kernel; health enters through the bin values, which span $\pm\sqrt{n-1}\,\sigma_\zeta(H)$ in logs around the deterministic component. Two properties of this construction are worth emphasizing. First, a household's shock position persists across health transitions: a household in the top expense bin that moves from moderate to bad health retains its place in the distribution, and its expense draw jumps because the bad-health bins are wider. Second, combining equation (94) with the discretized shock yields the model's expense object $e(t, e', i, H)$: a grid over age, shock state, income quintile, and health, whose transition kernel is the shared Rouwenhorst matrix. Table 9 reports the discretized distribution.

Households in long-term care ($H = L$) are treated separately, because self-reported OOP spending for institutionalized households is net of Medicaid payments and severely understates

Table 9: The discretized medical-expense distribution

Shock state	Stationary prob.	Healthy	Moderate	Poor	LTC
-2σ	0.06	158	179	191	138,286
-1σ	0.25	585	688	803	138,286
0 (median)	0.37	2,167	2,644	3,366	138,286
$+1\sigma$	0.25	8,030	10,151	14,114	138,286
$+2\sigma$	0.06	29,755	38,977	59,188	138,286

Note: Calibrated from RAND HRS 1998-2022. Monetary value reported in 2022 USD.

Table 10: Medicaid institutional parameters.

Parameter	Symbol	Single	Couple
Asset test (liquid wealth)	$\bar{a}_M(s)$	2	3
Income test (net of medical)	$\bar{y}_M(s)$	27	40
Homestead exemption	$\bar{H}$	750	750

Note: The asset test and homestead exemption are stocks; the income test is in biannual flows. Values in thousands of 2022 USD, set to 2025 U.S. federal floors.

the gross cost of care that the model requires. I instead anchor LTC costs with an external price and an internal quantity: the daily private nursing-home rate is set to \$300, consistent with the median private-room rate in the Genworth Financial (2023) Cost of Care Survey, and the length of stay is estimated from the HRS sample as the weighted mean of reported custodial nights among LTC households, 461 nights per biannual period (capped at 730). Together these imply a LTC cost of \$138,286 per biannual period, which replaces the expense process entirely in the LTC state: it is constant across shock states, income quintiles, and age.

Medicaid institutional parameters. Medicaid is modeled as a means-tested program that both covers medical expenses and guarantees a minimum level of consumption. A household with marital status, s, qualifies if it simultaneously satisfies an asset test, $a \leq \bar{a}_M(s)$, an income test on income net of medical spending, $y - m \leq \bar{y}_M(s)$, and, for homeowners, a home-equity test, $h - d \leq \bar{h}$ where $\bar{h}$ is the homestead exemption. Thresholds are set to 2025 U.S. federal floors: the asset test to SSI resource limits (\$2,000 single, \$3,000 couple), the income test to SSI limits scaled to medically-needy levels (\$27,000 single and \$40,000 couple biannual), and the homestead exemption to \$750,000 (within the federal \$730,000–\$1,140,000 range).

Table 11: Initial Conditions

Panel A. Aggregate entry distribution					
Households (N)	11,891				
Couple (%)	56.1				
Homeowner (%)	77.8				
Median liquid wealth (\$k)	9.1				
Mean liquid wealth (\$k)	140.2				
Owners with a mortgage (%)	49.6				
Median house value, owners (\$k)	227.7				
Median mortgage, mortgagors (\$k)	99.2				
Health: Healthy/Moderate/Poor/Long-term care (%)	43 / 50 / 7 / 1				
Panel B. By entry income quintile					
	Q1	Q2	Q3	Q4	Q5
Homeowner (%)	61	65	85	90	95
Median liquid (\$k)	0.1	0.7	8.3	25.0	128.7
Mortgagor (%)	32	29	39	46	51

Note: Calibrated from RAND HRS 1998-2022. Monetary value reported in 2022 USD.

4.2 Parameters Estimated by SMM

The remaining model parameters are estimated by the SMM, with a simulated sample of $12,000$ households drawn from the initial distribution reported in table 11. I then compute simulated model moments $\boldsymbol{M}_s(\boldsymbol{\theta})$, and choose θ to minimize the objective function

$$\hat{\boldsymbol{\theta}} = \arg\min_{\boldsymbol{\theta}} \big(\boldsymbol{M}_d - \boldsymbol{M}_s(\boldsymbol{\theta})\big)' \boldsymbol{W} \big(\boldsymbol{M}_d - \boldsymbol{M}_s(\boldsymbol{\theta})\big) \tag{96}$$

where $\boldsymbol{M}_d$ are the moments in the data and $\boldsymbol{W} = \text{diag}(1/\hat{se}_k^2)$ is the diagonal inverse-variance weight[14]. The reported loss value of the objective function is therefore equal to the sum of the square of simulated t statistics, $\sum_k t_k^2$. The weighting matrix $\boldsymbol{W}$ is obtained by bootstrapping over the sample of HRS households, see appendix A.1.

Table 12 summarizes the targeted moments in the data and the parameters they help to identify. In total, there are 17 structural parameters estimated by SMM and 36 moments in the objective. However, since the model is initiated at age 65, the initial distribution matches some moments by construction, so these moments are not targeted by the SMM but are listed for completeness[15]. Co-holding is the central object in the paper and targeting it

14. The objective function is optimized using a differential evolution algorithm followed by a Nelder-Mead polish.

15. Discrepancies in initial moment parameters exists either because the grid is too coarse, or the statistic is computed based off age 65 decisions.

Table 12: Targeted Moments

Targeted moment	Cells (age)	Primarily identifies
Median & 75th-pct. liquid wealth	65, 70, 75, 80, 85	β, σ, κ
Homeownership rate	65, 75, 85	ω_o, κ
Median home value, owners	70	η
Sale rate (overall, by health)	65–75, 75–85	κ, δ_0, δ_1
Mortgagor rate	65, 75, 85	f, θ_{DTI}, ψ_{refi}
Median mortgage balance	70, 80	f, θ_{DTI}, ψ_{refi}
Co-holding rate, by age	65, 75, 85	σ, κ (joint)
Co-holding rate, by income (pooled)	Q5, Q1–2; 65–85	β, $\underline{c}$, $\underline{c}_M(s)$
Weak-buffer mortgagors	70	σ, $\underline{c}$, $\underline{c}_M(s)$
Mean wealth at exit	at exit	γ, ζ, σ_B
Medicaid recipiency (low/high income)	65, 75, 85	ξ, $\underline{c}_M(s)$

Note: The liquid wealth moments at the 75th percentile are computed only at ages 65,75, and 85, giving 8 total liquid wealth moments. The sale rate is targeted overall, by healthy status, and by poor health status, which gives a total of 6 moments. Medicaid is separated by both low and high income status, giving a total of 6 moments.

ensures that the model maps to co-holding at the extensive margin. Targeting liquid wealth, mortgages, and weak buffers helps to ensure the model also captures the important features at the intensive margin as well. The homeownership and home sale moments then discipline the extensive housing margin.

In a nonlinear model, all moments jointly identify all parameters. However, specifically targeted moments help to primarily identify specific parameters. The liquid wealth moments in particular inform the model on how much households self-insure and on how quickly they pay down their wealth. This primarily helps to identify the discount factor, β, the coefficient of relative risk aversion, σ, and the consumption floors, $\{\underline{c}, \underline{c}_M(s)\}$, as it acts as an insurance. The homeownership rate and sale rate targets help the model to pin down the utility premium to owning, ω_o, and the cost of home sale, κ. The mortgage moments primarily pin down f, θ_{DTI}, and ψ_{refi}, because these parameters jointly determine the frictions related to refinancing. If mortgages persist longer and at higher rates, it must be the case that households refinance more. Wealth at exit helps to pin down the relative value of bequests, which is controlled by γ, σ_B, and ζ. However, since the model also targets median wealth at the 75th percentile, this this helps the model identify the luxury parameters, ζ and σ_B. If there is a very large mean wealth at exit, this would primarily inform the SMM that γ should be large. However, the wider the gap between median and 75th percentile liquid wealth holdings, the more bequests become a luxury good.

Table 13 reports the SMM estimates. β in particular deserves some discussion because it

Table 13: SMM-estimated structural parameters

Parameter	Description	Estimate
Preferences		
β	Discount factor (annual)	0.8544
σ	Relative risk aversion (CRRA)	2.545
ω_o	Ownership utility premium	66.06
η	Consumption weight (Cobb–Douglas)	0.7239
γ	Bequest intensity	11.00
ζ	Bequest shifter (luxury)	35.12
σ_B	Bequest curvature	2.341
Housing and mortgage frictions		
f	Proportional refinancing fee	0.1485
κ	Proportional selling cost	0.2301
δ_0	Owner depreciation, baseline (annual)	0.0020
δ_1	Owner depreciation, age slope (annual)	0.0077
θ_{DTI}	DTI refinancing cap	0.4786
ψ_{refi}	Refi shut-down rate (age)	0.2710
Medicaid and the safety net		
ξ	Medicaid take-up rate	0.4686
$\underline{c}$	Universal consumption floor (biannual, \$1,000s)	4.950
$\underline{c}^M$	Medicaid floor, singles (biannual, \$1,000s)	12.99
m_c	Medicaid floor, couples / singles	1.125

Note: All variables are reported annually unless otherwise stated.

implies a roughly 15% annual discount rate, well below the standard in the literature. The value I estimate here is identified in the data by liquid asset percentiles at different ages but households still accumulate too much wealth[16].

Despite including θ_{DTI}, the SMM estimation finds a proportional refinancing fee of about 15%, which is well above real closing costs that can range from around 2% - 5%. However, f is an effective refinancing cost that bundles together other frictions not modeled here[17]. A sensitivity analysis is included in table A.5 in appendix A.4.4 to help discuss identification.

Table 14 reports the model's fit against the targeted moments. Co-holding is somewhat

16. The estimation results for $\beta = 0.8544$ is much lower than the standard in the literature. However, increasing β incrementally towards a more typical 0.96 estimate, liquid asset holdings explode upward and median mortgage balances plummet, however, co-holding moments move as well but substantially less.

17. One explanation could be that the income I compute in the model incorporates savings withdrawals. These households may not meet standard income tests for qualifying for loans. Relatedly, under-qualified mortgagors receive risk premiums that can push interest rates up. Since interest rates are homogeneous, the refinancing fee may also absorb this margin. Beyond just market frictions, the inflated value might be capturing psychological costs associated with the stress of taking on more debt.

Table 14: Fitted Moments

Moment	Data	Model	t
Liquid wealth (2022 $k)			
Median (65)	8.7	8.9	+0.3
Median (70)	13.3	10.6	−3.4
Median (75)	16.3	7.9	−8.4
Median (80)	18.6	10.7	−7.9
Median (85)	23.3	40.9	+10.0
75th pct. (65)	96.2	108.9	+2.7
75th pct. (75)	129.6	109.3	−4.1
75th pct. (85)	159.8	187.2	+3.9
Housing & ownership (%; value $k)			
Ownership rate (65)	77.9	77.8	−0.2
Ownership rate (75)	78.1	78.4	+0.6
Ownership rate (85)	65.2	67.4	+3.2
House value, owners (70)	218.0	227.4	+2.8
Sale rate (65–75)	9.3	5.5	−6.9
healthy (65–75)	7.2	4.7	−3.7
sick (65–75)	11.5	6.4	−6.1
Sale rate (75–85)	20.0	19.7	−0.4
healthy (75–85)	17.8	17.7	−0.1
sick (75–85)	22.1	21.0	−0.9
Mortgage debt (%; balances $k)			
Mortgagor rate (65)	38.9	37.8	−1.7
Mortgagor rate (75)	23.3	19.6	−7.2
Mortgagor rate (85)	8.4	8.8	+0.9
Median balance (70)	90.1	96.3	+2.4
Median balance (80)	68.1	69.7	+0.5
Co-holding (mechanism, %)			
Co-holding rate (65)	21.0	19.7	−2.6
Co-holding rate (75)	12.0	9.8	−5.6
Co-holding rate (85)	4.3	2.5	−5.9
Q5, pooled	23.0	21.4	−2.2
Q1–2, pooled	7.2	5.0	−7.1
Weak-buffer mortg. (70)	49.8	49.5	−0.4
Bequest (2022 $k)			
Mean wealth at exit	264.9	255.1	−0.7
Medicaid recipiency (%)			
Low income (65)	21.6	23.9	+3.0
High income (65)	3.4	0.6	−11.0
Low income (75)	23.2	20.6	−3.2
High income (75)	3.5	1.2	−8.9
Low income (85)	23.5	20.5	−3.1
High income (85)	4.3	2.6	−4.7

Note: Fit at the SMM estimate. Dollar moments in 2022 $k; rates in percent. $t = (\text{model} - \text{data})/\widehat{\text{SE}}$, with clustered-bootstrap standard errors. With the large HRS sample the SEs are small, so $|t|$ measures statistical precision and overstates economic misfit; the Data–Model level gap is the economic magnitude. The sum of squared t-statistics equals the SMM objective (843.8); 25 of 36 targeted moments have $|t| > 2$.

below the data at every age, but before age 85, the shortfall is from mortgage shares rather than buffers. Conditional on being a mortgagor, the co-holding rate almost perfectly matches the data at ages 65 and 75: at 65, 54% in the data compared to 52%, and at 75, 52% compared to 50% in the model[18]. This is unsurprising because at age 70, the weak-buffer mortgagor share is almost exact[19]. At age 85, the model undershoots on both across the entire cross-section and conditional on being a mortgagor. For both misses, the limitation most likely responsible is that the model assumes full rationality and information, but behavioural motives may also play a role (Bertaut, Haliassos, and Reiter 2009; Gathergood and Weber 2014). Despite the limitation, co-holding has the right trend both in terms of the income gradient and shares of co-holders.

The results are mixed at the intensive margin. Mortgagor rates are mostly close to the true data moments and median holdings also have a reasonable fit. However, both median and 75th percentile liquid balances miss the mark, and the shortfall is mostly explained by the estimator trading in low liquidity from 70-80 to decrease excess liquidity at 85. In the data, following home sale, households typically see a very modest increase in liquid assets (see figure A.1 in appendix A.5). While it is not possible to directly track where the wealth from home sale ends up in the HRS, there are three main candidates. First, the wealth simply ends up in other types of savings the model ignores such as other real-estate, annuities, and trusts. Second, while the model does have a LTC state and forced relocation hazards, in the data, many elderly likely sell to accommodate and pay for deteriorating costs. While the model would see an accumulation of liquid assets in the following period, within the two-year periods between survey waves, care costs could easily eat up most of the housing wealth. Lastly are home transfers. Some exits from ownership involve no sale at all: federal Medicaid rules allow the home to be transferred without penalty to a spouse, a dependent or disabled child, or a child who lived in the home and provided care for the two years before the parent's admission to a nursing home (42 U.S.C. S1396p(c)(2)(A)), so equity can leave the household without ever appearing as liquid wealth.

Regardless of where homesale proceeds end up in the data, model households only have two choices: consume or save through liquid assets. Therefore, as the renting share grows around age 85, the model struggles to keep liquid assets low enough to match the data. The estimator's main lever to control this is to keep liquid balances low for as long as possible, which results in undershooting between ages 70 and 80, and then overshooting at age 85.

The remaining moments are home ownership and Medicaid. Ownership has a near perfect

18. Since co-holders are always mortgagors, the share of co-holders conditional on being a mortgagor is just the ratio $\frac{\text{co-hold share}}{\text{mortgage share}}$.

19. Weak buffer share implies mortgagor status without enough liquid assets to qualify as a co-holder.

fit at ages 65 to 75 and slightly overshoots at age 85, despite significantly underselling before 75. Simply, the data has more homeownership churn than in the model: more people exiting ownership and more people entering ownership. Medicaid recipiency is within a few points for the low income households but misses at the upper tail. Medicaid is important later on for policy experiments later but act mostly on the low income households where the fit is closer.

Taken together, the fit suggests that the model should be a credible laboratory for counterfactuals that work through owners' mortgage-versus-liquidity margin in early retirement, particularly from ages 65-80. However, from ages 85 and up, the model struggles to match liquidity moments so beyond this point the quantitative results diverge from the data.

5 Results

The estimation results in the previous section establishes how I discipline the model to reproduce a set of targeted moments. In this section, I ask what the estimated model implies beyond them. While the targeted moments only capture a few snapshots of life cycle portfolios, figure 10 shows the model's performance here from ages 65 to 95. The first five plots in figure 10 support conclusions discussed in the previous section. The most glaring model deviation from the data is that median liquid wealth appears to increase sharply after age 85, and this is primarily driven by an artifact from the data that the model cannot explain. Table 15 reports different median portfolio holdings prior to and after home sale. In the data, the median change in liquid assets among home sellers is zero, compared to almost $240,000 in the model after age 85. This is despite the fact that in the data, the same aged

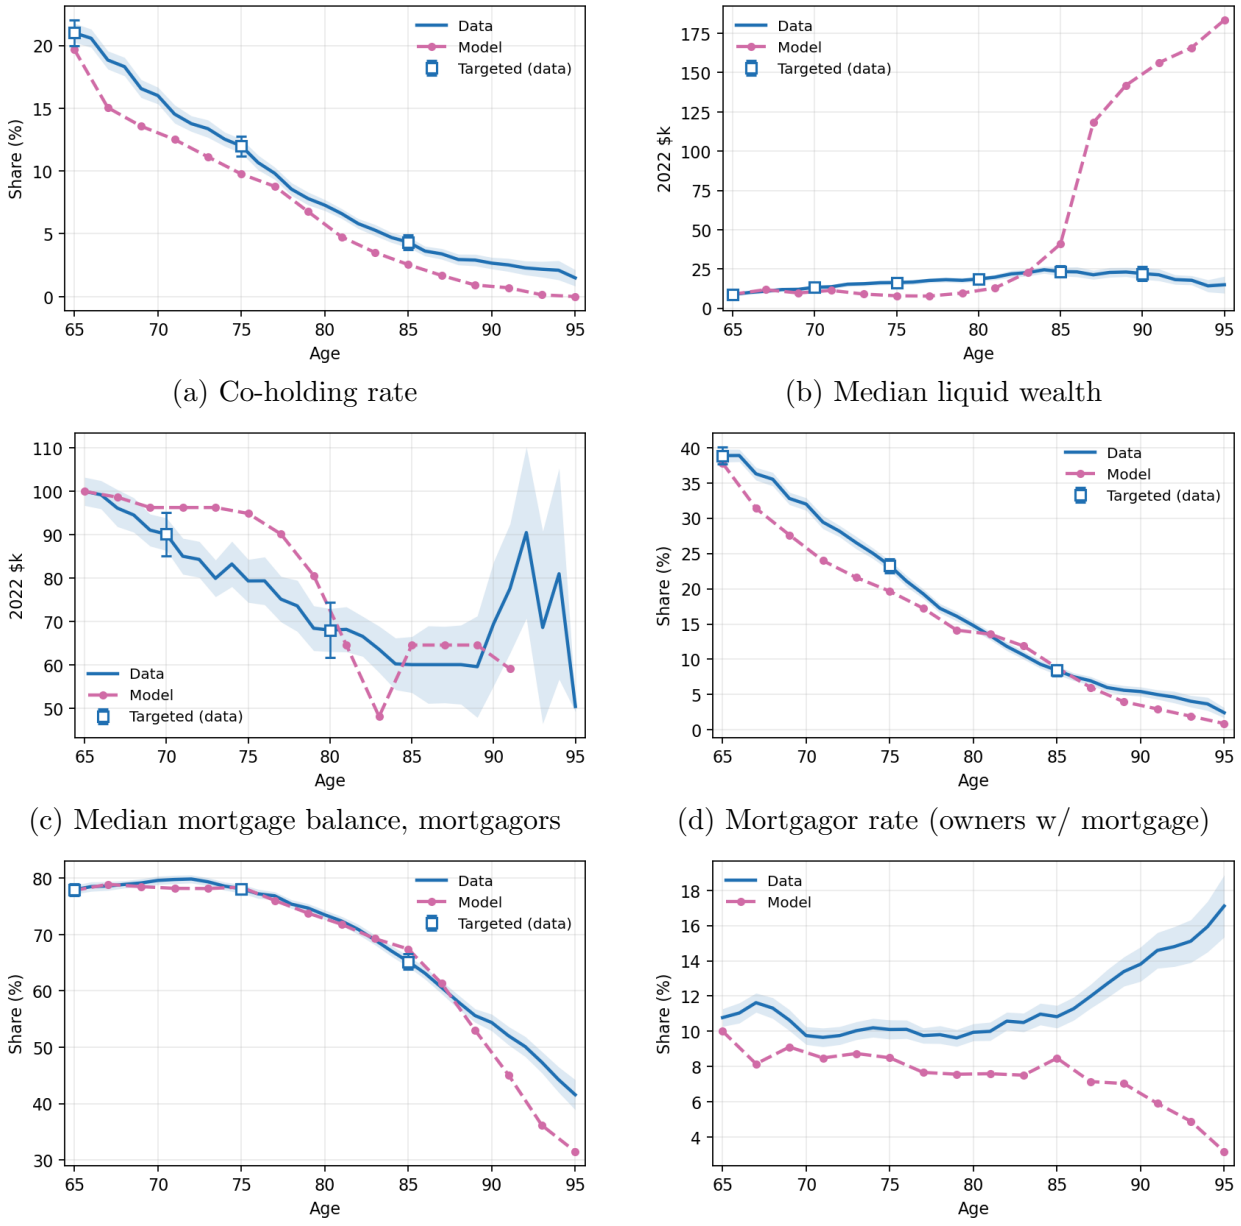
Table 15: What happens when older owners sell and rent: model and HRS, by age at the sale.

Statistic	Model 65-74	HRS 65-74	Model 75-84	HRS 75-84	Model 85-95	HRS 85-95
Households	788	1,197	1,229	1,383	645	804
Home value before	209	137	270	144	306	142
Home equity before	158	91	258	127	273	133
Liquid wealth before	20.7	1.4	18.3	11.9	142.6	32.3
Liquid wealth after	141.3	2.2	289.9	10.8	350.2	33.3
Change in liquid wealth, median	+91.0	+0.0	+195.8	+0.0	+239.2	+0.0
Change in liquid wealth, 75th pct.	+172.0	+17.8	+330.9	+28.4	+347.9	+37.4
Keep half of equity as liquid wealth	64%	17%	91%	18%	99%	21%
On Medicaid before	12.8%	12.8%	9.0%	9.0%	11.0%	9.7%
On Medicaid after	2.3%	14.0%	0.1%	12.3%	0.0%	15.4%

Note: Model: owners who sell and rent rather than buy another home, alive after the sale; before is the stock held at the sale, after is the first period as a renter. Home values are market values at the time of the sale in both model and HRS (model values lie on the house-size grid). HRS: owners who no longer own at the next wave and are alive at it; before and after are the two waves. Medians in 2022 $k except shares. Keep half: change in liquid wealth at least half of home equity before the sale. Some HRS exits are tenure misreports with no sale behind them, which pulls the median change toward zero; the 75th percentile is shown for that reason. HRS waves 4–16.

median sellers have $133,000 in home equity prior to sale[20]. The HRS does not track where the proceeds from home sale end, so this is still a puzzle, but there are a few hypotheses worth discussing. Table A.7 in appendix A.5 reports data on transfers made to children and grandchildren. Many households generally make *inter vivos* transfers, and from table A.9, sellers do appear to give more after home sale. However, transfers only account for about 8-10% of home equity, so this is likely not the source.

20. Sellers in the model also have higher median home value and home equities compared to the data. Since cash poor sellers are also likely to own smaller homes, this is again a consequence of the model not generating enough cash-poor sellers.



(a) Co-holding rate

(b) Median liquid wealth

(c) Median mortgage balance, mortgagors

(d) Mortgagor rate (owners w/ mortgage)

(e) Homeownership rate

(f) Medicaid recipiency rate (untargeted)

Figure 10: Model fit, life-cycle profiles, model vs. HRS data.

Note: Solid lines are HRS age profiles (three-year centered windows) with shaded 95% bootstrap bands; dashed lines with markers are model profiles from the simulated panel; open squares with whiskers are the targeted estimation moments ±1.96 standard errors. Monetary panels are in thousands of 2022 dollars. Panel f is untargeted: the estimation targets Medicaid recipiency for the low-income subgroup only, while the panel shown is the population rate.

From table A.8 in appendix A.5, it seems more likely that proceeds from home-sale are used to pay for entry into care like nursing homes. Among owners aged 85 and up who sell, 27.8% of sellers households report living in a nursing home, compared to just 4.6% who remain owners. Furthermore, 14.7% of sellers report household size increase after sale, which would be consistent with moving in with family following sale, or simply transferring ownership of the house to a "caregiver child."

While the model does capture relocation hazard, it does not capture the costs. Since model households receive the sale benefits from equity but have no other place to store the proceeds, there is a late-life surge in liquid assets compared to the data. Consequently, this also explains why in panel 10f, there is a divergence between Medicaid recipiency in the data and the model. Since more and more households downsize or transition to renting, their liquid balances are pushed well beyond the Medicaid asset test. While this model is not equipped to handle these questions, the results do point to some interesting behaviour around how retirees may balance their portfolios to receive Medicaid benefits.

5.1 Net-Worth vs. Portfolio Composition

The advantage of modeling liquid assets and debt as separate objects is to study how portfolio composition differs across net-worth. Table 16 reports statistics within net-worth quintiles. While the model matches the net-worth distribution in the data quite well, the

Table 16: Balance-sheet composition within net-worth quintiles at age 70, model and HRS, homeowners.

Quintile	Net worth model	Net worth HRS	Liquid model	Liquid HRS	Co-holding model	Co-holding HRS	HtM model	HtM HRS
1	52	47	0.0	0.0	15%	8%	80%	83%
2	148	132	1.9	7.6	20%	18%	66%	56%
3	214	247	2.0	27.4	21%	21%	54%	37%
4	388	450	35.6	113.4	14%	25%	46%	19%
5	842	1016	301.5	510.6	13%	28%	6%	7%

Note: Quintiles of net worth (liquid wealth plus home value minus mortgage) among living homeowners, formed separately in the model and the HRS. Net worth and liquid wealth are medians in 2022 $k. Co-holding: mortgage and liquid wealth above 3 months of income. Hand-to-mouth (HtM): liquid wealth below 3 months of income. Shares are of all homeowners in the quintile.

model stores much less liquid wealth within each quintile and is explained by two main reasons. First, in the bottom three quintiles, median liquid wealth in the model plateaus at about $2,000 because households there keep cash at or below the Medicaid asset limit to stay eligible. Second, at the fifth quintile, since model households are limited to the liquid asset for savings, at a certain point, the financial return to paying off debt outweighs

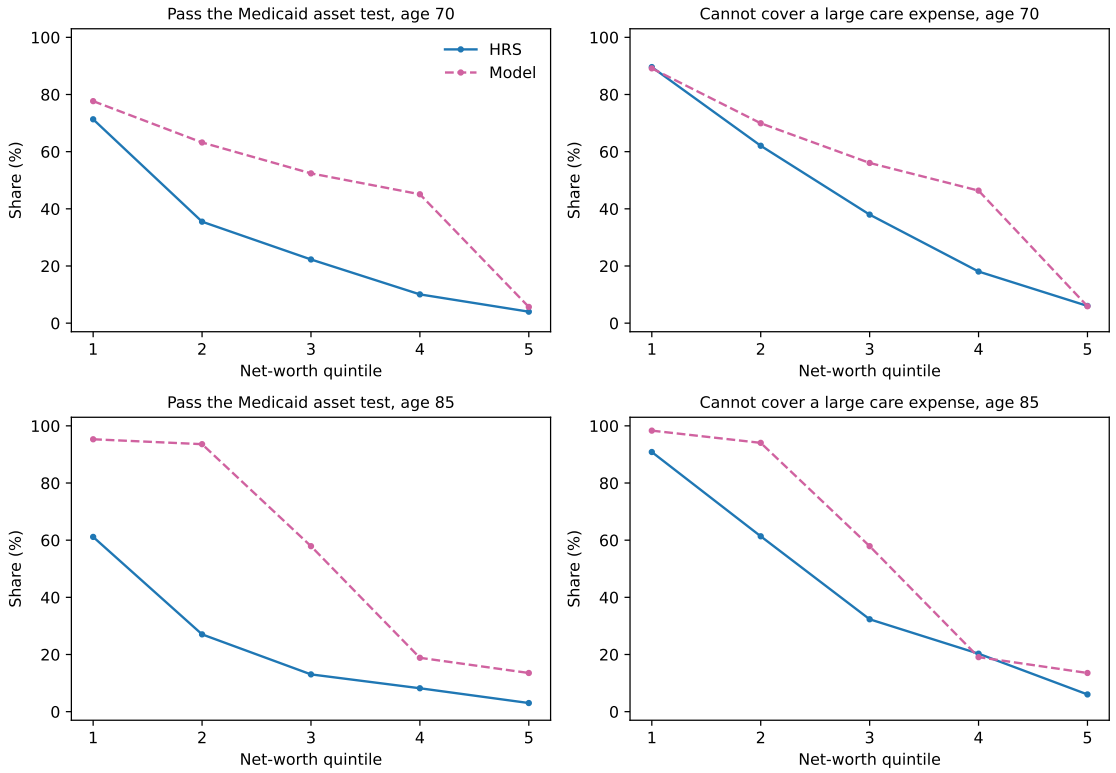


Figure 11: Asset thresholds by net-worth quintiles

Note: This figure plots the shares of households that would pass the Medicaid asset test (left) and do not have enough assets to cover a large medical shock (right) for ages 75 (top) and 85 (bottom) by net-worth quintiles. Large health shock is defined here as the 90th-percentile medical expense for that age.

precautionary motives, while in the data, wealthy households may have higher return options to save. Consequently, this also explains the mismatch in the co-holding distribution: since wealthier households pay off more debt in the model, they are less likely to co-hold.

Nevertheless, the co-holding and HtM columns of table 16 reveal that people with similar net-worth differ quite a lot in portfolio composition. At the 3rd quintile for example, the data reveals that 37% of households live with liquid assets lower than 3 months income (HtM), and another 21% have mortgage debt and liquid assets in excess of three months income (co-holders). These imply substantially different portfolio positions for households around $250,000 in net-worth.

Figure 11 plots the shares of households that would not be able to cover a large expense shock with liquid assets alone, and the share of households that would pass the Medicaid asset test. A model that represents household wealth only by net worth cannot determine liquid-asset eligibility from households' endogenous portfolio choices. It could reproduce the observed eligibility gradient only by imposing an auxiliary mapping from net worth into

liquid wealth or eligibility. Such a mapping may fit the cross section, but it would not provide a structural prediction when policies or incentives change.

Taken together, table 16 and figure 11 illustrate why net worth alone may be an insufficient statistic to study expense risk and Medicaid. Households with similar net worth can differ substantially in their holdings of liquid wealth and mortgage debt, creating heterogeneity in exposure to expense shocks and policy incentives that condition on liquid assets. Because a net-worth model would require an exogenous process to match the profiles in figure 11, it would provide little discipline for counterfactuals. Modeling liquid wealth and mortgage debt separately allows those outcomes to respond endogenously.

5.2 Decomposing Co-holding

Co-holding is optimal for some households in the model mainly through four interrelated channels. First, is the refinancing cost of debt. Because access to a new loan is not free, there is a tradeoff: paying down debt means lower cost of future mortgage interest payments, but because future costs are uncertain, paying down the mortgage also means loss of future borrowing capacity. In addition, this cost acts like a tax to anyone entering the mortgage market, which may deter new homeowners from co-holding to begin with. Second are mortgage access frictions. In the model, drawing the LTC state, or decreasing income from age or loss of spouse may result is losing access to the loan market partially or entirely. Since future expenses are uncertain, this again can increase co-holding because households might self-insure against bad states now, lest they be unable to take out a loan when they need it. Similarly to cost frictions, this can also decrease co-holding because it further restricts mortgage access. Both mechanisms can create precautionary borrowing motives to co-hold and access frictions to would be co-holders.

Third is expense shocks. Because the mortgage market is always costly to enter and access can become restricted in the future, the more a household is exposed to risk the more incentive they have to hold precautionary savings. Lastly are health shocks. Since the health process in part determines expense shocks, health risk can amplify precautionary savings motives. However, since health risk can also land a household in the LTC state, health risk can also amplify precautionary debt motives.

5.2.1 Decomposition

Because the mechanisms that generate co-holding in the model likely interact with one-another, a simple ablation of the model's frictions related to co-holding will produce results that depend on the order of ablation. Instead, I run a Shapley (1953) decomposition following

Shorrocks (2013), which does not depend on order. To keep the analysis tractable, I group the model's mechanisms into a set, $\mathcal{M}$, of the $K = 4$ groups outlined above: health risk, medical-expense risk, refinancing risk, and refinancing costs. For any subset $\mathcal{S} \subseteq \mathcal{M}$, let $\vartheta(\mathcal{S})$ be the share of co-holders when the mechanisms in $\mathcal{S}$ are on and all others are off. Every $v(\mathcal{S})$ is computed by re-solving the household problem with the same baseline parameters and shock draws. The $2^K = 16$ subsets then span every possible combination and therefore the contribution of group k is

$$\varphi_k = \sum_{\mathcal{S} \subseteq \mathcal{M} \setminus \{k\}} \frac{|\mathcal{S}|!(K - |\mathcal{S}| - 1)!}{K!} \big[\vartheta(\mathcal{S} \cup \{k\}) - \vartheta(\mathcal{S})\big]. \tag{97}$$

where φ_k is the Shapley value of mechanism k.

Alongside φ_k, I report two other statistics that create a bound for the entire effect. The first is the effect of group k alone, $v(\{k\}) - v(\varnothing)$: the contribution of the mechanisms in k when no other mechanisms are present. The second is a standard ablation, $v(\mathcal{M}) - v(\mathcal{M} \setminus \{k\})$: the contribution of the mechanisms in k when they are the only ones ablated. When the two differ, part of what the group does works through the other mechanisms in the model.

Table 17: What creates co-holding: Shapley contributions, co-holding rate.

	Ages 65-95			Ages 65-75			Ages 77-87		
Mechanism	Alone	Last	Shapley	Alone	Last	Shapley	Alone	Last	Shapley
Health risk	-0.2	-0.3	+1.9	-0.1	+0.5	+3.1	-0.4	-1.1	+0.7
Expense risk	+12.1	+1.4	+7.6	+12.0	+2.4	+9.4	+13.7	+0.3	+6.3
Refinancing risk	-1.1	+0.3	-1.5	+0.1	+0.6	-0.5	-2.7	+0.0	-2.8
Refinancing costs	+2.9	-8.7	-3.1	+6.3	-11.1	-1.9	-1.0	-6.9	-5.2
baseline−floor			4.9			10.1			-1.0
Floor			1.7			0.1			4.0
Baseline			6.6			10.3			3.0

Note: Every cell re-solves the model at the estimated parameters with a subset of the 4 mechanisms switched off, on the same households and shock draws (2^4 cells); medicaid and bequest motive are on in every cell. *Alone*: co-holding with only this mechanism on, minus the floor. *Last*: the baseline, minus co-holding with only this mechanism off. *Shapley*: the mechanism's marginal contribution averaged over all orders in which the mechanisms are introduced; the column sums to the explained amount. In every column a positive number means the mechanism raises co-holding. Co-holding: mortgagor with liquid assets above three months of income, share of living households, pooled over the ages of the band (weighted by the living at each age) and measured on the decisions made at each age.

Table 17 reports the main results of the exercise. The first thing to mention is that the four ablations groups together do not completely decompose co-holding in the quantitative model: across the entire panel, a 1.7% share of co-holding still persists after all four groups are ablated, with the residual coming from households aged 77-87, who see an increase in

co-holding under the ablations. Lumpy grids from the quantitative specification are the main culprit. In the baseline, most older households do not qualify for refinancing, or simply find it too costly to do so, so co-holders in the baseline are mostly tenured mortgagors. When access and refinancing costs are ablated, refinancing becomes optimal for home owners that want to use their house as a line of credit. However, since debt grids are discrete, it's nearly impossible to refinance small amounts, so the model chooses debt refinance positions that are too high and the agents hold the excess in liquid[21].

From ages 65-75, where co-holding is more pronounced, the floor lands at 0.1 percentage points, so the decomposition explains practically all co-holding here. The main driver is the expense shocks, landing at a Shapley value of 9.4 percentage points. Health risk also matters, and while substantially less important than expense shocks, they play an interesting role. Since health risk sorts households into different expense risk categories, adding health risk alone, or adding health risk as the last mechanism has only very subtle effects on co-holding overall. Therefore, the Shapley value at 3.1 percentage points shows that health risk moves co-holding through interaction with the other mechanisms.

Refinancing costs and access frictions appear to defy intuition from their negative Shapley values, but this just reflects the fact these frictions affect different populations differently. For example, renters and HtM home owners would benefit from an accessible and cheap mortgage market, but rarely become co-holders in the baseline. This is why turning every mechanism on except refinancing increases the co-holding share relative to the baseline.

Table 18 reports the Shapley decomposition among mortgaged homeowners, and shows how tenured co-holders respond to loan market frictions differently. Again, since co-holding is concentrated more before age 75 and the floor is better behaved, the focus here is on ages 65-75, but the headlines are the same. Again, expense shocks are the main drivers, but for mortgagors, both loan market mechanisms have the intuitive effects, and refinancing costs have the second highest Shapley value at 15.1%. This makes sense because for homeowners who already have mortgages, when refinancing is costly, it increases the value of maintaining debt headroom. By contrast, access risk tells two stories. On it's own, access to the loan market has very little impact on co-holding because under full certainty, homeowners can fully anticipate borrowing needs due to DTI tests, and the risk of ending in LTC is now zero. For this reason, the Shapley value is only 1%. However, adding access risk to the other three mechanisms increases co-holding by 8.5pp. This supports the hypothesis that for tenured mortgagors, restricting access to the loan market increases co-holding in the presence of

21. From ages 65-95, in the baseline model only 6.7% of co-holders are cash-outs compared to 97.3% at the Shapley floor. Of the remaining co-holders, 1.9% are tenured mortgagors and 0.7% are positions created from home purchase. The discrete grid is a limitation for obtaining a full 0% Shapley floor however these co-holders are different households from in the baseline.

Table 18: What creates co-holding: Shapley contributions, co-holding rate among the baseline's mortgagors.

	Ages 65-95			Ages 65-75			Ages 77-87		
Mechanism	Alone	Last	Shapley	Alone	Last	Shapley	Alone	Last	Shapley
Health risk	+0.0	+17.8	+9.1	-0.0	+18.0	+9.9	-0.2	+19.1	+7.7
Expense risk	+15.1	+29.0	+21.4	+15.5	+31.7	+24.0	+14.0	+21.9	+14.6
Refinance risk	-1.2	+8.4	+0.2	+0.1	+8.5	+1.0	-5.1	+8.2	-1.9
Refinance cost	+11.2	+18.9	+12.6	+14.6	+19.7	+15.1	+1.2	+17.0	+5.4
baseline−floor			43.3			50.0			25.7
Floor			1.5			0.1			5.8
Baseline			44.8			50.0			31.5

Note: Every cell re-solves the model at the estimated parameters with a subset of the 4 mechanisms switched off, on the same households and shock draws (2^4 cells); medicaid and bequest motive are on in every cell. *Alone*: co-holding with only this mechanism on, minus the floor. *Last*: the baseline, minus co-holding with only this mechanism off. *Shapley*: the mechanism's marginal contribution averaged over all orders in which the mechanisms are introduced; the column sums to the explained amount. In every column a positive number means the mechanism raises co-holding. Co-holding: mortgagor with liquid assets above three months of income, share of the households that hold a mortgage in the baseline at that age (the same households in every cell, so the statistic moves only through what they hold beside the balance), pooled over the ages of the band (weighted by the living at each age) and measured on the decisions made at each age.

higher risk.

6 Experiments

Frictions that prevent access to the mortgage market either directly by excluding household types, or indirectly by charging fees create margins of HtM, co-holders, and non-HtM. This section evaluates two specific counterfactuals the model can produce through exogenous factors that affect household decisions. The first is Medicaid. Given that Medicaid means tests are specifically defined on gross portfolio positions, counterfactual eligibility thresholds additionally offer a clear way to test the importance of portfolio composition on welfare.

The second is a rare event study. Under perfect information and perfect rationality, households make decisions that benefit their expected lifetime utilities the most. When it comes to debt, if households are overly optimistic about future outcomes, they may not be sufficiently prepared for future financial needs. Since mortgages are secured against the home, this model serves as a good workhorse to study how retired-aged households respond to unexpected changes in house prices.

6.1 Medicaid

As discussed previously, Medicaid works on co-holding in a few different directions. The asset test drives households to live HtM because asset thresholds are so low. The homestead exemption gives households an incentive to make sure their home equity does not exceed the limit, however, since the thresholds are so high, few households face this incentive. Therefore, Medicaid should drive more households to live HtM and should, if anything, make households less likely to co-hold.

In the baseline, to be eligible for Medicaid, single households need to keep assets below \$2k for singles and \$3k for couples, while the homestead exemption is held at \$750k. The income test is \$27,000 for singles and \$40,000 net of medical expenses, but since income is permanent and exogenous in the model, I focus just on asset and homestead reforms. The experiments in this section keep the calibrated parameters fixed and test two classes of Medicaid reforms: increases in the asset test, and decreases in the homestead exemption.

6.1.1 Behavioural Effects

Table 19 reports the results from counterfactuals run by tightening the homestead exemption. Because the homestead exemption is so large in the baseline, the threshold is

Table 19: Shrinking the Medicaid homestead exemption. Ages 65-75.

Policy	Co-holding	HtM	Mortgagor	Elig.	Med. liquid	Med. equity	Liquid
Baseline	10.2%	13.4%	20.4%	17.6%	9.1	235.9	2.8%
$\bar{H} = 250$	12.0%	11.7%	20.8%	16.3%	25.0	179.5	14.6%
$\bar{H} = 100$	16.4%	10.0%	23.7%	10.3%	53.9	179.0	23.1%
$\bar{H} = 50$	17.1%	10.4%	24.7%	5.6%	59.4	183.1	24.1%
$\bar{H} = 0$	17.0%	7.8%	22.7%	3.9%	60.4	185.6	24.4%

Note: Co-holding and the mortgagor rate are shares of all living households; HtM is the share of *owners* with a mortgage and less than three months of income in liquid assets. Elig. is the share meeting all three Medicaid tests (asset, home equity, income); realised take-up applies an exogenous participation rate to this group and so is proportional to it. Median liquid and median home equity are in thousands of dollars, the latter among owners. Liquid is the median household ratio of liquid assets to liquid assets plus home equity, among owners. All cells hold the estimated parameters fixed and share the same initial distribution and shock draws.

met at all equity levels for most home owners, and wealthy home owners likely fail the other Medicaid tests anyways. Even lowering the threshold to \$250,000 leads only to a decrease in Medicaid eligibility by about 1.3%. As the homestead exemption decreases further, co-holding increases, while conversely, the HtM share decreases. As evidenced by median liquid asset holdings, this is because for home owners, as qualifications on home equity tighten, they need

to rely more and more on self-insurance, driving up both debt and liquid savings, pushing people out of HtM. Despite this, the share of mortgagors can actually decrease because as homestead exemptions tighten, renting becomes the only way to qualify for Medicaid, so some home owners sell to rent.

The asset test produces a particularly interesting result in table 20. As the threshold

Table 20: Expanding the Medicaid asset test. Ages 65-75.

Policy	Co-holding	HtM	Mortgagor	Elig.	Med. liquid	Med. equity
Baseline	10.2%	13.4%	20.4%	17.6%	9.1	235.9
$\bar{a} = (6, 9)$	8.0%	17.5%	21.4%	19.3%	5.9	235.9
$\bar{a} = (10, 15)$	7.6%	19.9%	23.0%	21.6%	9.8	235.9
$\bar{a} = (20, 30)$	7.2%	20.9%	23.3%	26.3%	9.8	235.9
$\bar{a} = (60, 90)$	7.7%	21.4%	24.1%	28.8%	15.4	235.9

Note: Co-holding and the mortgagor rate are shares of all living households; HtM is the share of *owners* with a mortgage and less than three months of income in liquid assets. Elig. is the share meeting all three Medicaid tests (asset, home equity, income); realised take-up applies an exogenous participation rate to this group and so is proportional to it. Median liquid and median home equity are in thousands of dollars, the latter among owners. All cells hold the estimated parameters fixed and share the same initial distribution and shock draws.

for the asset test widens, the expectation might be that more home owners will opt to hold larger volumes of liquid assets. However, this ignores the fact that cash-rich households now have an incentive to hold less in liquid assets. They therefore pay down their mortgages and deplete their liquid assets holdings to qualify. This is evidenced by the decrease at $\bar{a} = (6, 9)$. It turns out that because insurance is so valuable for households, the effect of the cash-rich dominates any extraction of home equity by the cash-poor, and we observe a decrease in the share of co-holders, and an increase in HtM households. For the same reason that the asset test can go in different directions, widening from $\bar{a} = (20, 30)$ to $\bar{a} = (60, 90)$ actually sees a slight increase in co-holding because now home-owners can qualify for Medicaid while holding substantially high liquid assets.

6.1.2 Is Net-Worth a Sufficient Statistic?

The next set of experiments is designed to examine the extent to which portfolio composition contributes to welfare effects. Here, each household's welfare effect is computed as a liquid-wealth-equivalent measured in dollars directly from the value function.

Households are simulated from the calibrated model and two Medicaid reforms are considered: relaxing the liquid asset test from \$2k and \$3k to \$10k and \$15k for singles and couples respectively (a 5-fold increase), and a removal of the homestead exemption. Each reform is evaluated in two scenarios: the reform happens immediately at age 65, and the

reform is unexpectedly implemented at age 75. The liquid-wealth-equivalent, for a given household is the value Δ that solves

$$\text{Gainers, } V^R \geq V^B: \quad \Delta^* = \arg\min_{\Delta} \left\{ V^B(a+\Delta, \boldsymbol{z}) - V^R(a, \boldsymbol{z}) | V^B \geq V^R \right\} \tag{98}$$

$$\text{Losers, } V^R < V^B: \quad \Delta^* = \arg\min_{\Delta} \left\{ V^R(a+\Delta, \boldsymbol{z}) - V^B(a, \boldsymbol{z}) | V^R \geq V^B \right\} \tag{99}$$

where R and B superscripts refer to the reform and baseline regimes respectively, $\boldsymbol{z}$ is the vector of all state variables other than a, and Δ is the wealth-equivalent variation for gainers, and it is a compensating variation for losers.

Figure 12 shows the distribution of welfare gains across net-worth quintiles. Welfare changes are concentrated across the middle of the net-worth distribution, and they vary widely among households with the same net worth. Loosening the asset test benefits the

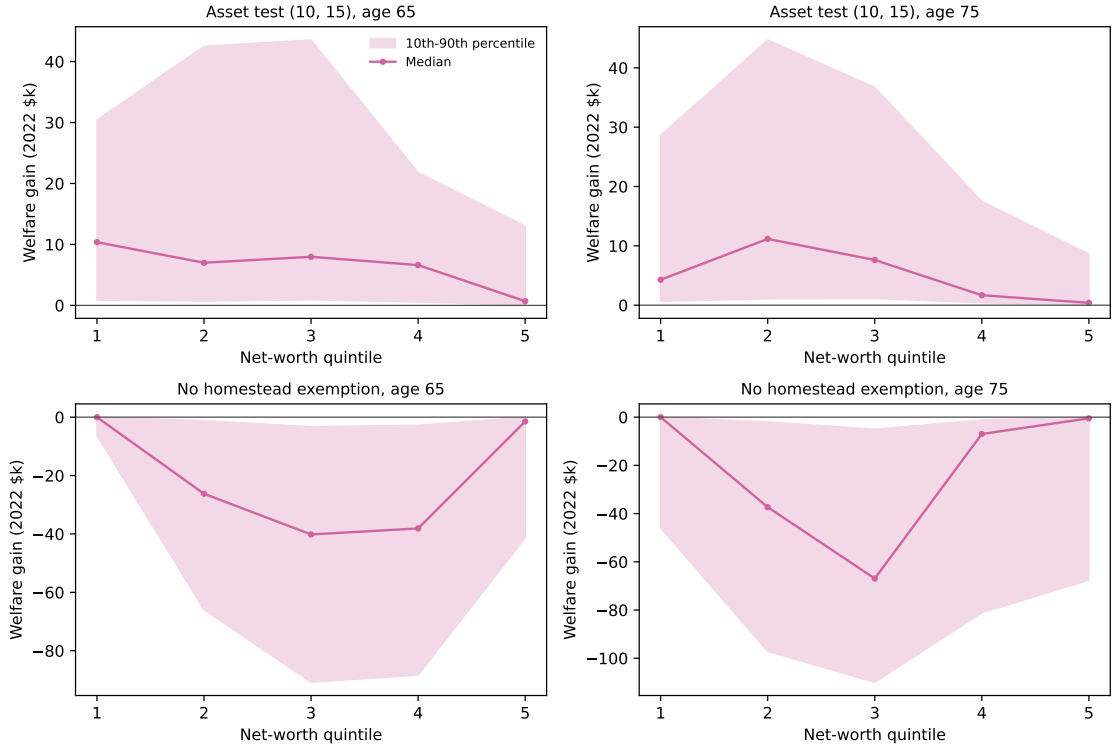


Figure 12: Distribution of households' welfare gains from Medicaid reforms.

Note: Each panel plots the liquid-wealth equivalent distribution across net-worth for Medicaid reforms: softer asset tests (top) and removing the homestead exemption (bottom). Reforms are implemented under two regimes: immediately at age 65 (left) and unexpectedly at age 75 (right). Solid line plot the Median welfare gains, and bands are the 10-90 percentile spreads. see tables A.11 and A.12 in appendix A.5

bottom three quintiles: the median household gains about \$10k in the bottom quintile at 65, and \$7k to \$8k in the next two, and the top quintile, rarely eligible, gains nothing. Removing

the homestead exemption costs the middle three quintiles the most: the median loss at 65 is about \$26k in the second quintile and \$40k in the third and fourth, and at 75 it reaches \$67k in the third. Within each quintile, dispersions are enormous. For the homestead exemption, at the third quintile, losses vary from near \$0 all the way up to around \$110k between the 10th and 90th percentiles.

The dispersion within quintiles follows the composition of the balance sheet, more so for the homestead exemption. At age 75, in each quintile, the 10% of households harmed the most are almost all owners with no liquid wealth. In the third quintile, these households lose \$131k on average, against \$55k on average for the rest of the quintile. This is not surprising, since by holding no liquid assets, they are entirely reliant on the homestead exemption. Without it, they have no insurance to manage shocks. The pattern for the asset test is more age dependent. In the second quintile at age 75, the 10% who gain most are mostly owners with cash that lies between the old and new asset limits, they gain on average \$72k compared to roughly \$14k for the rest of the quintile.

Since the spreads in figure 12 are substantial within net-worth deciles, I use the household level results to answer two questions, namely: given non-balance sheet household characteristics in the model, is net-worth a sufficient statistic to evaluate welfare? and, does separating debt from a net liquid asset position matter? To answer these questions I run a weighted least squares regression on the results with the following specification

$$\Delta_j^* = \alpha_0 + \sum_{k=2}^{10} \alpha_k \mathbf{1}\{\text{NW}_j \in \text{D}_k\} + \boldsymbol{b}'\boldsymbol{C}_j + \boldsymbol{\varrho}'\boldsymbol{\mathcal{B}}_j + \varepsilon_j \tag{100}$$

where D_k is the net-worth decile k with $\text{NW} = a + ph - d$, $\boldsymbol{C}_j$ is the vector of household characteristics (income type, marital status, and health), and $\boldsymbol{B}_j$ is the vector of portfolio objects: tenure status (owners, seller, renter), liquid wealth deciles, debt quintiles among mortgagors plus a zero debt state, and home equity quintiles among owners. In addition to the full specification, I run three restricted versions:

$$\text{NW } (\boldsymbol{b} = \boldsymbol{\varrho} = \boldsymbol{0}): \ \Delta_j^* = \alpha_0 + \sum_{k=2}^{10} \alpha_k \mathbf{1}\{\text{NW}_j \in \text{D}_k\} + \varepsilon_j \tag{101}$$

$$\text{NW + Characteristics } (\boldsymbol{\varrho} = \boldsymbol{0}): \ \Delta_j^* = \alpha_0 + \sum_{k=2}^{10} \alpha_k \mathbf{1}\{\text{NW}_j \in \text{D}_k\} + \boldsymbol{b}'\boldsymbol{C}_j + \varepsilon_j \tag{102}$$

$$\text{Characteristics + Composition } (\alpha_k = 0 \ \forall \ k > 0): \ \Delta_j^* = \alpha_0 + \boldsymbol{b}'\boldsymbol{C}_j + \boldsymbol{\varrho}'\boldsymbol{\mathcal{B}}_j + \varepsilon_j \tag{103}$$

Comparing the adjusted R^2 across specifications shows how much of the variation in welfare gains (or losses), Δ_j^*, are explained by net worth, household characteristics, and portfolio

composition. Under the null hypothesis, net-worth is a sufficient statistic given household characteristics. In other words, once net-worth and characteristics are known, composition adds no further information ($\boldsymbol{\varrho} = 0$). To test this, I use a heteroskedasticity-robust Wald test, and measure its economic importance by the increase in adjusted R^2 from adding composition, which under the null would be close to zero. The results of this exercise are reported in table 21.

Table 21: Net worth is not a sufficient statistic for the welfare effects of Medicaid reforms.

Reform	Age	R2 net worth	R2 NW + characteristics	R2 NW + char. + composition	R2 char. + composition	Composition adds	p-value
Asset test (10, 15)	65	0.08	0.29	0.34	0.33	+0.04	<0.001
Asset test (10, 15)	75	0.11	0.32	0.46	0.45	+0.14	<0.001
No homestead exemption	65	0.29	0.54	0.61	0.57	+0.07	<0.001
No homestead exemption	75	0.29	0.56	0.77	0.72	+0.21	<0.001

Note: R2: weighted R-squared of the gain on net-worth deciles (NW), characteristics (income type, marital status, health) and balance-sheet composition (tenure, liquid wealth, debt, home equity). Composition adds: increase in adjusted R2 from adding that block last. p-value: robust test that composition is jointly zero.

Portfolio composition adds between 4–14 percentage points of adjusted R^2 for the asset test reform, between 7–21 percentage points of adjusted R^2 for the homestead experiment, and the null is rejected at the 0.1% level for each reform.

6.2 Housing Price Crash

Taking on precautionary debt means purposely carrying excess debt to weather bad states, but this is only when the shocks are anticipated. In a model with closed-end debt, households facing more risk choose higher leverage (Buck 2026). This section experiments on an alternate housing price path, using a one-time permanent 30% drop in house prices at period $t = 1$, followed by a flattening period at $t = 2$, where housing prices are held constant. After the second period, housing prices then return to the original growth path. The experiment compares three counterfactuals to the calibrated baseline to answer a few questions, namely, how do household portfolios respond to a housing price crash? and is net worth a sufficient statistic to study the welfare effects of a wealth shock? The three counterfactual regimes studied in this section are described below.

Counterfactual 1: unanticipated shock. Let $\mathcal{V}_t^b$ denote the value function at time t for the baseline calibration and $\mathcal{V}_t^1$ denote the value function at time t for the unanticipated shock. Under this regime, households at period $t = 0$ solve the baseline problem

$$\mathcal{V}_0^1(\boldsymbol{x}) \equiv \max_{c,m,l',a',d'} \Big\{ u(c,h,s,o) + \beta E\big[\mathcal{V}_1^b(\boldsymbol{x}')\big|\boldsymbol{x}\big] + \beta\pi_{s0} E\big[B(a',d',\bar{h})\big] \Big\} \tag{104}$$

subject to the relevant constraints set at $t = 0$. However, once period $t = 1$ arrives, the new price path is realized and all households immediately adjust their beliefs such that the new regime is deterministic. Thus, for periods $t > 0$, they solve the problem

$$\mathcal{V}_t^1(\boldsymbol{x}) \equiv \max_{c,m,l',a',d'} \Big\{ u(c,h,s,o) + \beta E\big[\mathcal{V}_{t+1}^1(\boldsymbol{x}')\big|\boldsymbol{x}\big] + \beta\pi_{s0} E\big[B(a',d',\bar{h})\big] \Big\} \tag{105}$$

subject to the relevant constraint set.

Counterfactual 2: fully anticipated shock. Let $\mathcal{V}_t^2$ denote the value function at time t for the fully anticipated shock. Under this regime households solve the problem

$$\mathcal{V}_t^2(\boldsymbol{x}) \equiv \max_{c,m,l',a',d'} \Big\{ u(c,h,s,o) + \beta E\big[\mathcal{V}_{t+1}^1(\boldsymbol{x}')\big|\boldsymbol{x}\big] + \beta\pi_{s0} E\big[B(a',d',\bar{h})\big] \Big\} \tag{106}$$

subject to the relevant constraint set for all t.

Counterfactual 3: partially anticipated shock. This counterfactual sets a probability p hard-coded at $p = 0.25$ to the shock arriving at time $t = 1$. Thus, households weigh the probability of either of the previous counterfactuals arriving. At time period $t = 0$, households solve

$$\mathcal{V}_0^3(\boldsymbol{x}) \equiv \max_{c,m,l',a',d'} \Big\{ u(c,h,s,o) + (1-p)\cdot\beta E\big[\mathcal{V}_1^b(\boldsymbol{x}')\big|\boldsymbol{x}\big] + p\cdot\beta E\big[\mathcal{V}_1^1(\boldsymbol{x}')\big|\boldsymbol{x}\big] + \beta\pi_{s0} E\big[B(a',d',\bar{h})\big] \Big\} \tag{107}$$

subject to the relevant constraint set. Then, at $t = 1$, the shock materializes and they solve

$$\mathcal{V}_t^3(\boldsymbol{x}) \equiv \max_{c,m,l',a',d'} \Big\{ u(c,h,s,o) + \beta E\big[\mathcal{V}_{t+1}^1(\boldsymbol{x}')\big|\boldsymbol{x}\big] + \beta\pi_{s0} E\big[B(a',d',\bar{h})\big] \Big\} \tag{108}$$

subject to the relevant constraint set for all $t > 0$.

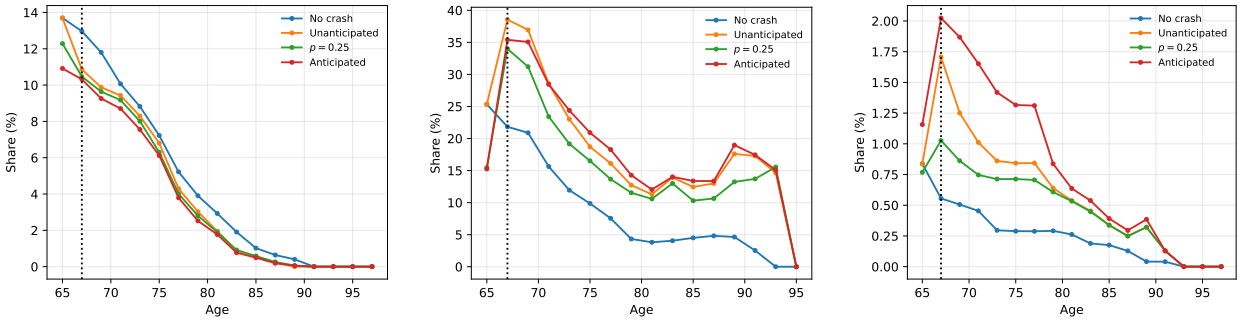
6.2.1 Behavioural Effects

The findings are summarized in table 22 and figure 13 The price drop has very little effect on co-holding at the extensive margin. The interpretation is simple, during a housing crash, co-holding should not move much as a result of the wealth shock alone. The results do however echo the prediction of the companion model: when households anticipate higher risk they choose higher leverage, but it seems to depend on leverage here. Median LTV is about the same in both the anticipated and unanticipated cases but increases at the upper tail. Interestingly, in the fully anticipated counterfactual, homeowners are the most likely to be underwater reaching a 1.7% share, compared to just 1.1% in the unanticipated version. Panel 13b in figure 13 shows the share of mortgagors above the refinancing cap. Across

Table 22: A 30% house-price decline at age 67, one flat period, then trend growth, under different beliefs. Ages 65-75.

Scenario	Co-holding	Mortgagor	Med. LTV	p90 LTV	Underwater
No crash	10.2%	20.4%	0.40	0.65	0.4%
Unanticipated	9.1%	20.4%	0.45	0.82	1.1%
$p = 0.25$	8.7%	19.1%	0.43	0.76	0.8%
Anticipated	8.4%	19.7%	0.44	0.88	1.7%

Note: Co-holding and the mortgagor rate are shares of all living households; LTV statistics are among mortgagors at decision-date prices. Underwater is the share of all households whose balance exceeds the house value. The unanticipated scenario holds age-65 decisions at baseline beliefs; the partial scenario assigns the crash its stated probability at age 65, with the crash then materializing; the anticipated scenario prices the crash path with certainty from age 65. All scenarios share the initial distribution and idiosyncratic shock draws.



(a) Co-holding rate (%)

(b) Mortgagors above refinancing cap

(c) Underwater Share

Figure 13: House-price crash experiment.

Note: a 30% price drop at the dotted line (age 67), one flat period, then trend growth resumes. Lines compare the no-crash control with crash paths under full, partial ($p = 0.25$) and no anticipation at the initial decision date.

each counterfactual, the crash pushes this share in the ballpark of 10-15 percentage points higher, and the gap persists all the way to age 95. This panel would seem to indicate that the housing crash doesn't just manifest as a wealth shock, but also as a housing access crash. As I'll show next, despite there being more homeowners losing access to refinancing, this doesn't really matter for welfare.

6.2.2 Welfare and Variations Across Portfolio

It wasn't surprising that portfolio composition was important for Medicaid reform since means testing is actually evaluated on the gross portfolio (gross assets and home equity). However, since a housing price shock is in essence a wealth shock, one might conjecture that net-worth could be a sufficient statistic, so evaluating the importance of portfolio composition here provides a different form of the same test. Since this is a wealth shock, this experiment also offers a natural way to study the explanatory power of net-worth: if net-worth is a sufficient statistic, then the liquid-wealth equivalent variation, Δ, should be equivalent to the loss in net-worth from the shock. Let Υ denote the wealth lost from the shock, then, define υ as the ratio $\upsilon = \Delta/\Upsilon$, then the following must be true

$\upsilon = 1$: the welfare loss from the crash is the same as the loss in net-worth

$\upsilon < 1$: housing wealth is worth less than cash to the household

$\upsilon > 1$: the crash costs more than the wealth lost.

If υ is significantly close to 1, then net-worth is a sufficient statistic for a housing crash. Experiments in this section are conducted only from counterfactual 1.

Wealth equivalent variations are computed the same way as they were for the Medicaid reforms, but the definition is re-stated here for convenience: for a given household the wealth equivalent variation is the value Δ that solves

$$\text{Gainers, } V^C \geq V^B: \quad \Delta^* = \arg\min_{\Delta} \left\{ V^B(a+\Delta, \boldsymbol{z}) - V^C(a, \boldsymbol{z}) | V^B \geq V^C \right\} \tag{109}$$

$$\text{Losers, } V^B < V^C: \quad \Delta^* = \arg\min_{\Delta} \left\{ V^C(a+\Delta, \boldsymbol{z}) - V^B(a, \boldsymbol{z}) | V^C \geq V^B \right\} \tag{110}$$

where C and B superscripts refer to the crash and baseline regimes respectively, $\boldsymbol{z}$ is the vector of all state variables other than a, and Δ is the wealth-equivalent variation for gainers, and it is a compensating variation for losers.

Table 23 reports the results from the welfare exercise. The average homeowner loses about \$0.70 per dollar of wealth lost in the crash, compared to \$0.62 for sellers. Since sellers already made the decision to sell prior to the crash and most of them are buying a different

sized home and they are are now paying less for the new home. This is the same reason

Table 23: The welfare cost of a 30% house-price crash per dollar of net worth lost, by balance sheet at age 67.

Group	Share	Households	Net worth	Net-worth loss	Welfare effect	Loss per dollar lost	Median per dollar lost
Renters	20.0%	2,593	8	0.0	+3.0	–	–
Sellers (moving this period)	5.5%	586	275	140.1	-87.0	0.62	0.63
Owners, no mortgage, liquid above the limit	27.4%	2,731	518	121.5	-81.8	0.67	0.79
Owners, no mortgage, liquid at or below the limit	20.3%	2,666	157	57.2	-57.8	1.01	1.28
Mortgagors, liquid above the limit	15.2%	1,495	228	117.0	-67.1	0.57	0.62
Mortgagors, liquid at or below the limit	11.7%	1,282	112	74.8	-54.4	0.73	0.83
All homeowners	80.0%	8,760	244	98.8	-69.3	0.70	0.77

Note: Welfare effect: liquid wealth (2022 $k) equivalent of an unanticipated 30% fall in house prices at 67 (one flat period, then trend growth) for each household in its pre-crash state; negative = loss. Net-worth loss: the fall in the market value of the house on impact (zero for renters). Loss per dollar lost: the group's total welfare loss over its total net-worth loss; median per dollar lost: the median of the household ratio. Limit: the Medicaid asset limit ($2k single, $3k couple), on liquid wealth. Sellers sell this period, at the crash price. Net worth: median, $k; net-worth loss and welfare effect: means.

that renters actually see a small gain in welfare, because their net-worth is unaffected and they can now purchase homes for a lower price.

There is a fair amount of variation across types, and the biggest differentiator is mortgagor status. For the average homeowner without a mortgage and under the Medicaid asset threshold, the loss is welfare is practically equal to the wealth lost with $v = 1.01$. By contrast, the least harmed by the wealth shock are mortgagors that have liquid assets in excess of the Medicaid liquid asset threshold, with average loss per dollar $v = 0.57$. These results reveal how Medicaid and mortgagor status are really the two margins that move welfare loss. Because welfare is measured by a wealth equivalent variation, giving liquidity to households that currently qualify for Medicaid means disqualifying them for insurance. Therefore, they need higher liquid injections to be made whole following the wealth shock. Despite mortgagors losing less than a dollar in welfare per dollar in wealth, the results don't necessarily mean that a mortgage is an insurance against a wealth shock, mortgagors may also be more constrained and thus value liquid cash more.

The results make it clear that despite being a wealth shock, the welfare loss differs across households depending on what their portfolio positions are. To study how important gross positions are to the welfare analysis, let M_1, M_2, and M_3 denote loan-to-value (LTV) bands,

and $\bar{M}$ denote mortgagor status such that

$$\begin{cases} M_1 = 1, M_2 = M_3 = 0 & \text{if LTV} \in (0, 0.4) \\ M_2 = 1, M_1 = M_3 = 0 & \text{if LTV} \in [0.4, 0.6) \\ M_3 = 1, M_1 = M_2 = 0 & \text{if LTV} \geq 0.6 \\ \bar{M} = M_1 + M_2 + M_3. \end{cases} \tag{111}$$

I then run the following weighted-least-squares regressions

$$\upsilon_j^S = \alpha_0^S + \beta^S \bar{M}_j^S + \mathcal{C}_j^S + \varepsilon_j \tag{112}$$

$$\upsilon_j^S = \tilde{\alpha}_0^S + \tilde{\beta}_1^S M_{1,j} + \tilde{\beta}_2^S M_{2,j} + \tilde{\beta}_3^S M_{3,j} + \mathcal{C}_j^S + \varepsilon_j \tag{113}$$

where

$$\mathcal{C}_i^S = \begin{cases} 0 & \text{(raw)} \\ \sum_k \alpha_k^S \mathbf{1}\{NW_j \in D_k\} + \sum_l \phi_l^S \mathbf{1}\{P_{t^*} h_j = v_l\} + X_j' \gamma^S & \text{(deciles)} \\ \sum_{b=1}^{50} \left(\mu_b^S + \lambda_b^S z_{jb}\right) \mathbf{1}\{NW_j \in B_b\} + \sum_l \phi_l^S \mathbf{1}\{P_{t^*} h_j = v_l\} + X_j' \gamma^S & \text{(fine)} \end{cases} \tag{114}$$

D_k are weighted net-worth deciles and B_b weighted net-worth bins, z_{ib} is net worth standardized within bin, b, v_l are the house values on the model's grid, and X_i contains income-type, marital-status and health indicators, and $S \in \{\text{below}, \text{above}\}$ denotes whether the household is above or below the Medicaid asset test limit so the regressions are run on both groups separately (the below group also includes households at the limit).

The regression results are reported in table 24. Holding net worth, house value, and household characteristics fixed, owners with a mortgage and liquid wealth above the Medicaid limit lose \$0.08 less welfare per dollar of net worth lost than owners without a mortgage. At the same net-worth, more debt means more cash, and cash is protected from the shock. However, this is only true to an extent. If the home owner's LTV is too high, the housing crash can push home value to a point where selling would actually require the seller to pay cash to settle the loan.

6.3 Do Gross Debt Positions Matter?

The two experiments showed that portfolio positions reveal additional welfare effects that net-worth alone cannot capture. The result echoes the HtM literature which argues that conditional on net worth, households behave differently depending on their liquidity (Kaplan,

Table 24: Welfare loss per dollar of net worth lost in the 30% crash: mortgagors relative to owners without a mortgage, on each side of the Medicaid asset limit (age 67).

Term	Above, raw	Above, deciles	Above, fine	At/below, raw	At/below, deciles	At/below, fine
Any mortgage	-0.10 (0.01)	-0.09 (0.01)	-0.08 (0.01)	-0.28 (0.03)	-0.11 (0.02)	-0.24 (0.06)
LTV below 40%	-0.14 (0.01)	-0.10 (0.01)	-0.08 (0.01)	-0.36 (0.04)	-0.11 (0.02)	-0.24 (0.06)
LTV 40% to 60%	-0.06 (0.01)	-0.06 (0.02)	-0.03 (0.02)	-0.12 (0.04)	-0.21 (0.07)	-0.33 (0.12)
LTV 60% or more	+0.02 (0.01)	+0.03 (0.03)	+0.07 (0.03)	-0.17 (0.04)	-0.14 (0.11)	-0.20 (0.16)
40% to 60% minus below 40%	+0.08 (0.01)	+0.05 (0.01)	+0.05 (0.02)	+0.24 (0.04)	-0.10 (0.07)	-0.10 (0.09)
60% or more minus below 40%	+0.16 (0.01)	+0.14 (0.02)	+0.15 (0.02)	+0.19 (0.04)	-0.03 (0.10)	+0.04 (0.13)
60% or more minus 40% to 60%	+0.08 (0.01)	+0.09 (0.02)	+0.10 (0.02)	-0.05 (0.04)	+0.07 (0.08)	+0.14 (0.09)
Households	4,226	4,226	4,226	3,948	3,948	3,948

Note: Weighted least squares of each owner's welfare loss per dollar of net worth lost, owners without a mortgage the base, on a mortgage indicator (row 1) or, in a separate regression, on indicators for the mortgagor's loan-to-value band (rows 2 to 4), with weights equal to the survey weight times the net-worth loss. LTV: the mortgage balance over the house's market value before the crash. Deciles: also net-worth decile indicators, house-value indicators (one per value) and characteristics (income type, marital status, health). Fine: the deciles replaced by 50 weighted net-worth bins with a separate linear term within each bin. At the same net worth and house value, a household with more debt holds as much more liquid wealth, so the controlled gaps compare gross positions at a given net position. Rows 5 to 7: differences between the band coefficients. Above: liquid wealth above the Medicaid asset limit ($2k single, $3k couple); at/below: at or below it. Heteroskedasticity-robust (HC1) standard errors in parentheses; –: fewer than 20 households in the group.

Violante, and Weidner 2014). However, that literature also nets debt into assets. With this model, I can go a step further and ask what a model that nets debt against liquid assets would get wrong.

To isolate the effect of gross positions on welfare loss, I use the model to compare households with debt to a counterfactual twin where their asset positions are netted against debt, and $\tilde{a} = a - d$[22]. The twin is identical in all states including net liquid asset position, the only difference between them becomes gross portfolio positions. Any net-asset model would not see any differences between these households. Since the twin study also applies equally to both experiments, I perform the twin study across both the Medicaid and unanticipated shock counterfactuals. The comparison covers 26 per cent of households at 65 and 12 per cent at 75, those holding any liquid wealth alongside a mortgage. Because errors of opposite sign cancel, the mean gross-debt effect says little about individual households. Table 25 therefore reports, for co-holders, the mean absolute gross-debt effect (the netted household average measurement error), its size relative to the mean absolute welfare effect, the misstated share by more than $1k, and the share of variation in welfare that the netted-twins account for. On average twins report similar effects to their gross counterparts, with mean welfare effects differing by between $1k and $8k. Netting the balance sheet would only change total welfare

22. In the event that $a - d < 0$, I set $\tilde{a} = 0$ such $\tilde{d} = d - a$ and the state for the twin is evaluated at $(\tilde{a}, \tilde{d})$. This would be equivalent to a negative position in a net-asset model.

Table 25: Co-holders against their netted twins.

Reform	Age	Co-holders	Effect	Effect, netted twin	Gross-debt effect	Mean absolute gross-debt effect	Share of effect size	Share above $1k	Variation captured by twin
Asset test (10, 15)	65	26%	+13.77	+12.63	+1.13	6.39	46%	53%	0.40
Asset test (10, 15)	75	12%	+13.90	+17.77	-3.88	7.79	55%	61%	0.42
$\bar{H} = 0$	65	26%	-23.93	-29.61	+5.68	9.14	38%	58%	0.43
$\bar{H} = 0$	75	12%	-21.39	-29.61	+8.22	10.18	48%	69%	0.53
Crash	67	19%	-66.40	-63.47	-2.93	16.83	25%	82%	0.54

Note: Netted twin: the same household with the cash it holds alongside its mortgage (the smaller of liquid wealth and the balance) used to pay the mortgage down, so it has the same net worth, net liquid wealth, house, income, health and marital status, but does not co-hold. Both are evaluated with the same value functions. Co-holders: share of households holding both liquid wealth and a mortgage. Effect: mean welfare effect of the experiment, 2022 $k. Gross-debt effect: household minus twin, the error a netted balance sheet makes in the household's welfare effect; it is zero for households that do not co-hold. Share of effect size: mean absolute gross-debt effect over mean absolute welfare effect. Variation captured by twin: one minus the weighted sum of squared gross-debt effects over the weighted variation of co-holders' effects (one if netting were harmless).

effects across all households by less than about 6%. However, for individual households, the measurement errors are large. At the household level, the welfare effects differ by about $6k to $10k in the Medicaid reforms which amounts to between 38% to 55% of the size of the welfare effect, and by almost $17k, a quarter of the effect, in the house-price crash. Furthermore, between 53% to 82% of co-holders are misstated by over $1k, and netted twins account for only 40% to 54% of the variation in welfare effects among co-holders.

The results from the experiments reveal two important lessons. First, Net worth on its own is not a sufficient statistic to study the welfare effects of Medicaid reform or a housing price crash because welfare effects vary substantially within households at the same net-worth, and portfolio composition captures a large share in the variation. Second, since co-holders make up at most about a quarter of all households, net asset positions capture most of the average welfare effects across all households. However, netting misstates welfare effects at the individual household level and therefore is a poor guide to how much each household gains or loses.

7 Conclusion

The share of US households retiring with debt has risen sharply from 1998-2022, and this increase has been mostly driven by mortgage debt. At the same time, nearly half of these households carry liquid assets in excess of three months income. Since the rate of return on liquid assets is low relative to mortgage interest rates, this suggests they would prefer to pay a premium for precautionary buffers, rather than pay down debt. Building on the theory of co-holding with closed-end loans in Buck (2026), this paper builds a quantitative model to

explain why co-holding is rational in retirement.

The model reproduces important facts around household portfolios. Between ages 65-75, among mortgagors, co-holding is very closely matched to the data: 54% against 52% at age 65 and 52% against 50% at age 75. Because the model understates the share of mortgagors around age 75, the total co-holding share is understated across the population. However, since the proceeds from homesale can only be stored as savings in the liquid asset, the model overstates liquid wealth accumulation late in life, and the calibration responds by undershooting liquid asset moments between age 70 and 80.

The calibrated model is used to decompose the mechanisms that drive co-holding by a Shapley method. The largest contributor is medical expense risk, which accounts for an average of 9.4 percentage points across all households from ages 65-75. This is the precautionary savings motive to co-hold. By contrast, refinancing costs, DTI/LTV constraints, and mortgage access risk introduce loan market frictions that affect different households in different ways. Refinancing costs in particular introduce a tradeoff: paying down debt now can save on future interest payments, but paying down too fast risks paying re-access costs. For mortgaged home-owners, this is the second largest driver of co-holding. At the same time, while re-access frictions don't contribute much directly, ablating access risk alone decreases co-holding among mortgagors by 8.5 percentage points between ages 65-75. Thus, these frictions amplify co-holding indirectly through the other mechanisms. These are the precautionary debt motives to co-hold. Meanwhile, because loan market frictions also restrict access to new loans, these mechanisms decrease co-holding among debt free households.

The main advantage of this model is that it creates gross liquid, illiquid, and debt positions, which makes it a good laboratory to study what netted models miss. Net-worth is a reasonable guide to a policy's total cost but a poor guide to its incidence. Under a reform that eliminates the Medicaid homestead exemption, cash-poor owners are hit the hardest because they use their house as a shield against the asset test. A second experiment considers the welfare losses from a 30% crash in housing prices. Liquid-wealth equivalent measures of welfare run from 0.57 to 1.01 dollars per dollar of wealth lost, depending primarily on mortgage and whether liquid savings pass the Medicaid asset test. While the experiments reveal welfare heterogeneity within net-worth quintiles, I use twin households created by cloning the baseline co-holders and forcing their liquid assets to a net-of-debt position to study the extent to which net-worth misses welfare changes. The netted twins misstate each co-holders welfare effect by about one quarter to one half while the aggregate effect only modestly changes. So while netted positions might be sufficient to capture averages, they substantially miss the variation across households.

The model still has some important limitations. First, because there is only a single

liquid asset position, and renters have no other savings options, liquid wealth erupts past its target at age 85. By comparison in the data, liquid positions only increase modestly following home sale. Additionally, HRS households are seen to make significant transfers to relatives (children and grandchildren) over the course of their lifecycles. Since the model only has bequests, there is no place for households to make *inter vivos* transfers. Second, the model only matches Medicaid in an *ad-hoc* sense, and data limitations make it hard to compare true means tests. When households report assets to the HRS, they report current asset balances from their accounts. In reality, Medicaid only considers savings in the account that remain at the end of the month, so Medicaid eligible households in the data may report higher asset holdings to HRS than they keep at the end of the month. Furthermore, different states have different asset thresholds, so forcing everyone to meet the Federal floor is at best a one-size-fits-all approach. Lastly, since co-holding is measured by simultaneous asset and debt positions, and the model is solved numerically on a grid, the Shapley decomposition cannot reach a true zero co-holding floor.

One set of extensions follows directly from these limitations. The model accumulates too much liquid wealth after age 85 because sellers keep their proceeds in the liquid account, whereas in the HRS, proceeds either go into care, transferred *inter vivos*, or moved into assets outside the liquid portfolio. The model has long-term care costs and a health-dependent sale hazard, but draws them separately, so a seller faces the same expense distribution as a non-seller in the same health state. A care state that triggers the sale and carries a persistent expense would close that gap and would likely raise Medicaid recipiency among higher-income households, where the model falls short. A second step is to let the Medicaid means test affect the portfolio households bring into retirement, rather than taking the age 65 distribution as given. More broadly, the results suggest that the question for retirement policy is not only how much retirees have saved but in what form: how wealth is held determines how households fare when their health, their house price, or the rules around them change.

A Appendix

A.1 Data

This appendix details the methodologies for collecting and presenting the data in sections 2, 4, and 5.

A.1.1 Source and Sample

This paper takes micro-data from the Health and Retirement Study (HRS) and complements it with data from FRED and (Damodaran 2023). Data is collected from the RAND HRS files taking 13 waves from 1998-2022, and the sample is restricted to households aged 65-95. Earlier waves are excluded to maintain consistency in the construction and interpretation of household balance sheet measures across survey years. The sample is also restricted to ages 65–95 because the number of households above age 95 declines quickly. Additionally, households who are not retired when observed between ages 65 and 67 are dropped because the income process in the paper would misallocate their types (see variable and process definitions below).

Statistics reported throughout the paper are weighted using the HRS respondent-level analysis weight (RwWTRESP) with one adjustment to account for long-term care (LTC). The RwWTRESP weight is calibrated to the community-dwelling population and therefore omits nursing-home residents. To capture their contributions to aggregate statistics, I carry forward each household's last observed positive weight which effectively holds representation of households in LTC at a constant from entry to exit. Carrying these weights forward effectively increases the weighted LTC share of households from 0.5% to 3.0%.

Dollar values are reported in 2022 dollars deflated using the FRED CPIAUCSL series. A summary of the cleaned HRS data is reported in table A.1.

Table A.1: HRS analysis panel

Panel A: living household-waves (ages 65+ statistics)						
			Owners (%)		Medians (2022 $k)	
Year	N	N_{65+}	Own	Mortgagor	Income	Liquid
1998	14,395	7,830	72.6	16.3	34.7	20.0
2000	13,213	7,811	72.6	17.5	35.3	22.1
2002	12,349	7,966	72.5	18.3	36.1	23.8
2004	13,645	8,080	72.4	20.0	35.8	22.5
2006	12,605	8,351	72.4	21.6	36.3	17.0
2008	11,897	8,315	73.3	23.4	37.0	18.6
2010	15,280	8,006	73.7	23.8	37.4	14.3
2012	14,316	7,882	74.3	24.6	38.5	12.4
2014	13,241	7,709	74.4	25.9	41.7	12.0
2016	15,023	7,476	74.3	25.9	42.4	11.9
2018	12,467	6,779	75.1	27.5	43.1	11.3
2020	11,490	6,780	76.8	29.2	45.1	13.1
2022	11,924	6,910	75.9	29.2	44.1	15.0

Panel B: pooled ages 65+, weighted	
Couple household (%)	43.0
Health: excellent/very good (%)	36.0
Health: good/fair (%)	53.0
Health: poor (%)	8.0
Health: LTC ($\geq$90 NH nights) (%)	3.0
Medicaid receipt, past 2y (%)	11.2
Any OOP medical, past 2y (%)	89.5
Median OOP | Positive ($2022, bian)	3,269

Notes: RAND HRS Longitudinal File 1992–2022 (v1), one row per household-wave (financial respondent's record), living rows; counts are unweighted. Statistics use the HRS respondent weight with the carry-forward adjustment for institutionalized rows described in the text; dollar values are deflated to 2022 dollars by CPI-U (FRED CPIAUCSL). Liquid assets are checking, CDs, bonds, and stocks net of non-housing debt (IRA balances excluded); mortgage debt includes home-equity borrowing. OOP medical spending is measured over the two-year wave window.

A.1.2 Variable Definitions

Mortgage debt and home ownership. To estimate household mortgage debt, I use the HRS variable for all reported mortgages a and land contracts on the primary residence (HwAMORT) in addition to the variable for home-equity loans and lines of credit on the primary residence (HwAHMLN). Primary residences are inferred from the primary residence value variable (HwAHOUS) and secondary residencies are excluded from both the value and debt side of the analysis. Loan-to-value (LTV) is therefore the ratio of mortgage debt to home value.

Liquid assets and co-holding. As in Buck (2026), co-holding is defined as carrying more than three months income in liquid assets, and in this case, a positive mortgage balance from HwAMORT + HwAHMLN. By contrast to the SCF measure in Buck (2026), liquid assets are computed as the sum of HwACHCK (checking, savings, and money market account balances), HwABOND (corporate, municipal, foreign bonds, and bond funds), and HwASTCK (directly held stocks and stock mutual funds) minus HwADEBT (non mortgage debt). This definition is broader than the SCF measure because it includes stock and bond holdings.

Health states and OOP medical expenses. The HRS asks about health state (RwSHLT) where households report excellent, very good, good, fair, and poor health. Additionally, I add that a household is considered in long-term care (LTC) for any 2-year period in which they report 90 or more nights spent in nursing home. The model's four states are encoded as good (excellent or very good), $H = G$, moderate (good or fair), $H = M$, poor, $H = B$, or LTC, $H = L$, where LTC overrides the self-reporting. Out-of-pocket (OOP) medical expenses are computed as the sum of respondent (RwOOPMD) and spouse (SwOOPMD) OOP medical expenses.

A.2 Model Solution Routine

The model is solved recursively through backwards induction using a combination of grid search, and derivative methods. The renter, seller, and owner problems are all solved along grid approximations of the state space. Starting at age $t = T$, each point inside the grid approximation of the state space is solved backwards from $t = T$ to $t = 0$. Future value functions are approximated using linear interpolations where needed.

Medicaid. If a household qualifies for Medicaid then their medical expenses are covered and their new consumption floor is set to $\underline{c}^{\text{med}} > \underline{c}$. In the decision rules, if they qualify if they pass all the relevant tests. In the event that a household qualifies, they solve their decision problem under the assumption they will receive the benefit with certainty. The simulation applies a probability ξ that the benefit is received. This is an assumption applied

after the fact to simplify the number of iterations. Without it, each subproblem would have to be performed twice for anyone who could potentially qualify. For each subproblem listed below, if the household is Medicaid eligible, switch $\underline{c}$ for $\underline{c}^{\text{med}}$.

The Renter's Problem. The renter's problem can be solved by choice of a' over a continuous range and h over a grid. Since h is a discrete choice, a candidate a' is selected for each possible discrete h choice, and the highest (a', h) pair is used as the solution. For each grid point, the feasible set is $a' \in [0, I - \underline{c})$, where I denotes the net available resources. If $\underline{c}$ is not affordable, then $\underline{c}$ is selected as the default consumption level. If there exists an $a' \in [0, I - \underline{c})$, then the solver find the optimal choice of a' using a Brent-Dekker method. Once the convergence criteria is met, the candidate solution a' is tested against the second-order condition (SOC) using a second derivative test. If the candidate solution fails the SOC, the solution falls back to a coarse grid global search with refinements.

The Seller's Problem. The seller's problem diverges depending if the household transitions to renting or is downsizing. The method solves both branches as selects the decision that maximizes utility. The renting branch conveniently is the same problem as the renter, except that the seller has home sale proceeds as part of their gross resources. Otherwise, the solution routine follows the same method as the renter. For the purchasing branch, the seller has two discrete choices: $\bar{h}$ and d'. The solution method for this branch selects the optimal a' for each $(\bar{h}, d')$ pair following the same method as the Renter's problem.

The Owner's Problem. The owner's problem simultaneously solves a', d' using the first order conditions, with a coarse 2D grid search adding two refinements. First, the method attempts to solve the first-order condition using a nested bisection. If it is successful within the convergence criteria, the candidate (a', d') pair is tested against the SOC through a bordered Hessian test. If the first-order condition fails to converge, or the candidate fails the SOC, the algorithm passes onto a coarse grid search over d' with a nested golden-section search for a'. The coarse grid for d' is then refined twice around the candidate. Since owner's can also transition to sellers, this method is applied to both branches, with the utility maximizing branches and their respective (a', d') pairs kept. For owners in the keep or refinancing branch who are Medicaid eligible, if there is no feasible combination of (a', d') such that $c > \underline{c}_M(s)$, they are rescued by Medicaid, however, if they are not eligible they are forced to sell. For Owner-sellers, sellers, and renters, if there is no feasible combination of (a', d') such that $c > \underline{c}$, then $\underline{c}$ is stored as the solution with $a' = 0$ and $d' = \bar{d}$.

A.3 First-Stage Processes

Income. The model's income process combines total household income and retirement saving account draw-downs. The process is computed from the following relation

$$y^{\text{model}} = \max\left\{0, y^{\text{total}} - r \cdot \max(a, 0) - y^{\text{earn}}\right\} \tag{A.1}$$

where y^{total} is build from the RAND HRS HwITOT variable and y^{earn} is constructed from both RwIEARN (respondent earnings) and SwIEARN (spousal earnings). Since y^{total} includes capital income, I remove $r \cdot \max(a, 0)$ since this is already from returns on liquid assets. y^{earn} is also netted out to prevent earnings from labour income, and households who fully retire after age 65 would then see a sudden drop in model income despite receiving a higher permanent income-quintile draw. This also preserves IRA/401(k) draw downs, and other capital income not accounted for in the model.

Each household is assigned a permanent income type based on its first observation in the longitudinal file between ages 65-67 conditional on being retired. Since y^{earn} is netted out, keeping unretired households would misallocate their income type which is a large source of the heterogeneity.

y^{model} is also indexed by age/time, t, marital status s, and initial income quintile i, such that $y^{\text{model}} = y_j^{\text{model}}(t, i, s)$, where the j subscript refers to the household ID number in HRS. I then take the weighted household median across each cell in $y_j^{\text{model}}(t, i, s)$, using age ± 1 year around the target age to obtain $\tilde{y}_t = \tilde{y}^{\text{model}}(t, i, s)$. Cells with fewer than 30 observations are imputed by pooling across marital status. Cells that are still empty then take on the nearest non-missing age. Then, to ensure income is weakly decreasing with age, I run the entire profile through isotonic regression constrained monotone-decreasing using the following procedure

$$\hat{y} = \arg\min_{y \in \mathbb{R}^T} \sum_{t=0}^{T-1} n_t (y_t - \tilde{y}_t)^2 \quad \text{subject to} \quad y_0 \geq \cdots \geq y_{T-1} \tag{A.2}$$

where n_t is the number of observations in cell (s, i) at age t. The results from the isotonic regression are reported in table A.2

Finally, because income at this point is still gross and not disposable, using $\hat{y}$ as the income process would overestimate resources. To capture disposable income, I apply a 10% tax on income above \$25,000 biannual, flat rates for medicare, \$4,000 for singles and \$8,000 for couples biannual, and a flat \$3,000 biannual fee for property tax and insurance, applied uniformly. The disposable income results are reported in table A.2. The disposable income process is used for the calibrated model.

Portfolio composition and returns. The mortgage rate, r_m, house price growth, g_H, and

Table A.2: Construction of the income grid $y_t(i, s)$.

		Age 65			Age 75			Age 85		
Status	Quin.	Med.	Isotonic	Disp.	Med.	Isotonic	Disp.	Med.	Isotonic	Disp.
Single	Q1	21.9	25.8	18.7	27.0	25.8	18.7	28.3	25.8	18.7
	Q2	38.4	39.6	31.1	39.1	39.1	30.7	34.6	35.2	27.1
	Q3	62.2	62.9	52.1	54.6	54.6	44.7	44.4	45.3	36.2
	Q4	106.8	106.8	91.6	83.9	83.8	70.9	61.0	63.0	52.2
	Q5	181.0	181.0	158.4	114.0	114.0	98.1	106.1	98.2	83.9
Couple	Q1	17.8	45.4	32.4	71.0	45.4	32.4	63.8	45.4	32.4
	Q2	41.8	53.4	39.6	65.4	53.4	39.6	71.5	53.4	39.6
	Q3	68.6	76.3	60.2	82.8	76.3	60.2	79.6	76.3	60.2
	Q4	109.4	112.1	92.4	108.9	108.9	89.5	101.2	101.2	82.6
	Q5	209.7	209.7	180.2	174.7	174.0	148.1	170.1	162.3	137.6

Note: Quin. denotes quintile and med. denotes median. Values are biannual 2022 dollars, thousands (model units). *Median* is the HRS sampling-weighted median of model-consistent income $y^{model} = \max\{0, y^{total} - r_F a^+ - y^{earn}\}$ among households within ±1 year of the reported age, in the given entry-income quintile and current marital status. *Isotonic* is the cell-size-weighted monotone-decreasing least-squares projection of the median profile over model periods. *Disp.* subtracts the disposable-income wedge: a 10% rate on income above \$25k, Medicare Part B and Medigap premiums of \$4k (single) and \$8k (couple), and \$3k of property tax and insurance. The disposable grid is the input to the solver.

rental price growth, g_R, are all averaged from FRED time series data (see table 2). Returns to liquid assets depend on income type, i, such that $r(i)$ is computed as the within income-quintile average portfolio return.

Household-level portfolio composition is estimated using the following equation

$$p_{j,h} = \frac{b_{j,h}}{b_{\text{CHECKING},h} + b_{\text{CDs},h} + b_{\text{BONDS},h} + b_{\text{STOCKS},h}} \tag{A.3}$$

where $b_{j,h}$ is household h's balance for $j \in \{\text{CHECKING}_h, \text{CDs}_h, \text{BONDS}_h, \text{STOCKS}_h\}$. Then, let $\bar{p}_j(i)$ denote the average portfolio share of type j for income quintile i. Balances for each account type are computed from the RAND HRS HwACHCK, HwACD, HwABOND, HwASTCK variables for checking, CDs, bonds, and stock accounts respectively.

Returns to account type j are inferred from the Damodaran (2023) annual returns dataset using the 10-year Treasury total return for bonds and the S&P 500 total return including dividends for stocks. CD accounts are estimated using the FDIC national CD rates. Checking accounts are assumed to have zero return. Each return is computed as the average from 1998-2022 and deflated using the FRED CPIAUCSL series.

Within-income quintile returns, $r(i)$, are then estimated using the Arrow-Pratt risk

premium (Pratt 1964; Arrow 1971) to account for risky returns to stocks

$$r(i,\sigma) = \sum_j \bar{p}_j(i) r_j - \frac{1}{2}\sigma \left(\bar{p}_{\text{STOCKS}}(i)\sigma_{EQ}\right)^2 \tag{A.4}$$

where $r(i) = r(i,\sigma)$ depends on the CRRA parameter σ, $\bar{p}_{\text{STOCKS}}$ is the share of stocks, and σ_{EQ} is the standard deviation of real S&P total returns from 1998-2022. It should be noted that this technically makes $r(i,\sigma)$ a 2-stage parameter, but σ gets hard-coded at $\tilde{\sigma} = 2.69$ in the first stage before $\hat{\sigma} = 2.545$ is used in the final calibration. This is done because with $\sigma_{EQ} = \%18.42$ and the highest stock share $\bar{p}_{\text{STOCKS}}(5) = 34\%$ means the square term dominates the haircut expression and deviations in σ barely moves $r(i,\sigma)$.

Health Process. Let households be indexed by j and waves by ω. For each household-wave, let $\delta_{j,\omega} \in \{0,1\}$ be the deceased indicator with 1 indicating deceased, $w_{j,\omega}$ the HRS weight, $\text{SRH}_{j,\omega} \in \{G, M, B, L\}$ be the self-reported health, and $N_{j,\omega}$, the nursing-home nights over the two-year wave window. The household's health state is given by

$$H_{j,\omega} = \begin{cases} G & \text{if } N_{j,\omega} < 90 \text{ and } \text{SRH}_{j,\omega} \in \{\text{excellent, very good}\} \\ M & \text{if } N_{j,\omega} < 90 \text{ and } \text{SRH}_{j,\omega} \in \{\text{good, moderate}\} \\ B & \text{if } N_{j,\omega} < 90 \text{ and } \text{SRH}_{j,\omega} \in \{\text{bad}\} \\ L & \text{if } N_{j,\omega} \geq 90 \end{cases} \tag{A.5}$$

The sample is taken from the set of survivors with the following restrictions

$$\mathcal{S} = \big\{(j,\omega) \;\big|\; \text{retired}_{j,\omega},\ \delta_{j,\omega} = \delta_{j,\omega+1} = 0, \\ H_{j,\omega} \text{ and } H_{j,\omega+1} \text{ are observed}, \text{age}_{j,\omega} \in [65, 94]\big\}. \tag{A.6}$$

Each age cell is in the set

$$\mathcal{S}_t = \left\{(j,\omega) \in \mathcal{S} |\ \text{age}_{j,\omega} \in [\text{age}_t - 1, \text{age}_t + 1)\right\}. \tag{A.7}$$

For each t and $(x,y) \in \{G,M,B,L\}^2$, the transition probability matrix is estimated as the weighted share of households transitioning from x to y in subsequent waves

$$\hat{\Gamma}(y|x) = \frac{\sum_{(j,\omega)\in\mathcal{S}_\sqcup} w_{j,\omega} \mathbf{1}\{H_{j,\omega} = x\}\mathbf{1}\{H_{j,\omega+1} = y\}}{\sum_{(j,\omega)\in\mathcal{S}_\sqcup} w_{j,\omega}\mathbf{1}\{H_{j,\omega} = x\}}. \tag{A.8}$$

Relocation Hazard. The sample for relocation is restricted to homeowners surviving to

the next period, formally

$$\mathcal{O} = \big\{(j,\omega) \;\big|\; \text{retired}_{j,\omega},\ \delta_{j,\omega} = \delta_{j,\omega+1} = 0, \\ \text{own}_{j,\omega} = 1, H_{j,\omega} \text{ observed}, \text{age}_{j,\omega} \in [65,94]\big\} \quad \text{(A.9)}$$

and further restricted to unconstrained households

$$\underline{\mathcal{O}} = \big\{(j,\omega) \in \mathcal{O} \big| a_{j,\omega} > \kappa\big\} \quad \text{(A.10)}$$

where $a_{j,\omega}$ is liquid assets and κ is the 75th percentile of liquid assets. The estimator is given as

$$\text{Hazard}(x,s) = \frac{\sum_{(j,\omega)\in\underline{\mathcal{O}}} w_{j,\omega}\mathbf{1}\{H_{j,\omega} = x, s\}\mathbf{1}\{\text{own}_{j,\omega+1} = 0\}}{\sum_{(j,\omega)\in\underline{\mathcal{O}}} w_{j,\omega}\mathbf{1}\{H_{j,\omega} = x, s\}}. \quad \text{(A.11)}$$

***Mortality* and Marital Transitions**. Mortality risk and marital transitions are part of the same process. The data for this process is described by the set

$$\mathcal{M} = \big\{(j,\omega) \;\big|\; \text{retired}_{j,\omega},\ \delta_{j,\omega} = \delta_{j,\omega+1} = 0, \text{age}_{j,\omega} \in [65,94], i \text{ and } H_{j,\omega} \text{ observed}\big\}. \quad \text{(A.12)}$$

where $\delta_{j,\omega} = 0$ denotes survival, i denotes income quintile, and $H_{j,\omega}$ denotes health state. Mortality and divorce hazard are then estimated from a weighted binomial GLM with logit link from the following model

$$\text{P}(\delta_{j,\omega} = 1|\boldsymbol{z}_{j,\omega}) = \Lambda\Bigg(\boldsymbol{z}_{j,\omega}\boldsymbol{\beta} + \\ \sum_{i=2}^{5}\gamma_i\mathbf{1}\{i_j = i\} + \sum_{l\in\{G,M,B,L\}}\alpha_l\mathbf{1}\{H_{j,\omega} = l\} + \sum_{l\in\{G,M,B,L\}}\eta_l\mathbf{1}\{H_{j,\omega} = l\}\tilde{a}\Bigg) \quad \text{(A.13)}$$

where

$$\Lambda(v) = \frac{1}{1+e^{-v}}, \quad \text{(A.14)}$$

$\boldsymbol{z}_{j,\omega} = (1, \tilde{a}_{j,\omega}, \tilde{a}^2_{j,\omega}, F_h, C_{j,\omega})$, $\tilde{a} = \text{age} - 75$, F denotes female, and C denotes couple. γ_i, α_l, and η_l are all dummy variables capturing the households income quintile at age 65, the household's health status, and an interaction between health status and age. This is a standard process following (De Nardi, French, and Jones 2010).

The fitted model is used to predict mortality and divorce hazard jointly as follows. For each (t,i,H), let

$$\mu^m_{t,i,H} = \Lambda(\hat{v}|F=0,C=1), \quad \mu^f_{t,i,H} = \Lambda(\hat{v}|F=1,C=1) \quad \text{(A.15)}$$

where m denotes male and f denotes female. Assuming that couple transitions are independent of spousal death, then

$$\pi_{20} = \mu^m \mu^f, \quad \mu_{21} = \mu^m + \mu^f - 2\mu^m \mu^f, \quad \mu_{22} = (1-\mu^m)(1-\mu^f). \tag{A.16}$$

For singles, let ι denote the share of female singles, then

$$\pi_{10} = (1-\iota)\Lambda(\hat{v}|F=0, C=0) + \iota\Lambda(\hat{v}|F=1, C=0). \tag{A.17}$$

Medical Expense Process. The expense process follows De Nardi, French, and Jones (2010) (DFJ)with a few notable deviations discussed in sequence. For households in health states $H \in \{G, M, B\}$ and income-quintile i, biannual out-of-pocket (OOP) medical spending satisfies

$$\ln e^m_{j,\omega} = \alpha^m(i_j) + \beta^m(i_j) \cdot (\text{age}_{j,\omega} - 60) + \gamma^m(H_{j,\omega}) + \zeta_{j,\omega}, \quad \gamma^m(G) = 0 \tag{A.18}$$

$$\zeta_{j,\omega} = \rho\zeta_{j,\omega-1} + \varepsilon_{j,\omega}, \quad \text{sd}(\zeta|H) = \sigma_H \tag{A.19}$$

where e^m denotes OOP medical expenses, α^m is an income type dependent intercept, and β^m is an age slope conditional on income type, and γ^m is a health dependent intercept. The error term $\zeta_{j,\omega}$ is a single AR(1) process estimated on the total residual. This is the first departure from (DFJ) who decompose the term into a persistent AR(1) and a transitory term. The biggest implication here is lower long-horizon expense risk than (DFJ), implying weaker precautionary motive, which translates to more conservative co-holding results. The estimated sample for estimation uses the set

$$\mathcal{E} = \big\{(j,\omega) \;\big|\; \text{retired}_{j,\omega},\ \delta_{j,\omega} = \delta_{j,\omega+1} = 0, \\ \text{age}_{j,\omega} \in [60, 94], i \text{ and } H_{j,\omega} \text{ observed}, \text{OOP}_{j,\omega} \geq \$1\big\}. \tag{A.20}$$

The second departure from DFJ here is that I do not condition on gender because expenses are measured at the household-level. The

Equation A.19 is estimated using weighted OLS on RAND HRS weights. Persistence is computed using

$$\hat{\rho} = \frac{\sum w_{j,\omega}\zeta_{j,\omega}\zeta_{j,\omega-1}}{\sum w_{j,\omega}\zeta^2_{j,\omega-1}} \tag{A.21}$$

and $\hat{\sigma}_H$ is the weighted mean of squared residuals within each health state.

I use a Rouwenhorst (1995) discretization to reduce the continuous AR(1) shock the finite

Markov chain $\mathrm{P}(e'|e)$. The procedure is recursive. Start with $p = (1+\rho)/2$ and let

$$\boldsymbol{P}_2 = \begin{bmatrix} p & 1-p \\ 1-p & p \end{bmatrix}. \tag{A.22}$$

Then, define $\boldsymbol{P}_n$ recursively as

$$\boldsymbol{P}_n = p\begin{bmatrix} \boldsymbol{P}_{n-1} & \mathbf{0} \\ \mathbf{0}' & 0 \end{bmatrix} + (1-p)\begin{bmatrix} \mathbf{0} & \boldsymbol{P}_{n-1} \\ 0 & \mathbf{0}' \end{bmatrix} + (1-p)\begin{bmatrix} \mathbf{0}' & 0 \\ \boldsymbol{P}_{n-1} & \mathbf{0} \end{bmatrix} + p\begin{bmatrix} 0 & \mathbf{0}' \\ \mathbf{0} & \boldsymbol{P}_{n-1} \end{bmatrix} \tag{A.23}$$

for $n \geq 3$ with the interior rows $2, ..., n-1$ divided by 2. This defines the transition matrix.

Nodes are located symmetrically and equidistantly spanning $\pm\sqrt{n-1}$ standard deviations following

$$e_k = \frac{\sigma_H(2k-n-1)}{\sqrt{n-1}}, \quad k = 1, ..., n. \tag{A.24}$$

For 5 possible states, this leads to the grid $\boldsymbol{e} = (-2\sigma_H, \sigma_H, 0, \sigma_H, 2\sigma)$.

The third departure from DFJ is that long-term care (LTC) expenses are handled separatly to avoid double counting in the calculation of expenses, following the treatment of nursing-home costs as a distinct risk with its own price (Kopecky and Koreshkova 2014). To compute the cost of LTC, I take nursing home nights (NrSNIT in RAND HRS), clipped at 730 for LTC households and compute the weighted mean, 461 nights. I assume the price is roughly \$300 per day, roughly the 2022 national private room rate, and bundle the costs together (Genworth Financial 2023)[23]. The average cost is then computed as

$$461 \text{ nights} \times \$300/\text{night} = \$138,286 \tag{A.25}$$

biannually.

A.4 Simulated Method of Moments Solver

A.4.1 Initial Distribution

The solver initializes each simulated household at $t = 0$ with the state vector $\boldsymbol{x}_0 = (a, d, \bar{h} : H, s, i)$: liquid assets, debt, home value, health, marital status, OOP medical expenses, and initial income quintile drawn exogenously from the HRS. The sample is drawn from all observed alive households at age 65-67. If a household appears twice in that age range, the

23. While Genworth never published a standalone 2022 or 2023 survey, their March 12 2024 press release stated the price of a private room is a nursing home had increased to \$116,800 in 2023, 4.9% up from 2022. From this estimate, I estimate the private cost as roughly $\frac{\$116,800}{1.049} = \$111,344/\text{yr} \approx \$300/\text{day}$

earlier wave is used as their initial observation. The sample is not restricted to only strictly retired households despite the income quintiles cutting working households. This is done to preserve a larger sample. To remedy the income-quintile inconsistency, If a household is not retired, I instead find the first observed retired wave for that household and assign their income quintile from that wave. If retirement is never observed then the observation is dropped. Each variable used from the RAND HRS files and a description of their construction is summarized in table A.3.

Table A.3: Initial-distribution states and their sources.

State	RAND source	Construction
a_0	`HwACHCK`, `HwACD`, `HwABOND`, `HwASTCK`, `HwADEBT`	Checking + CDs + bonds + stocks − other debt; IRA excluded; clipped at zero
h_0	`HwAHOUS`	Primary residence value
d_0	`HwAMORT` + `HwAHMLN`	First mortgage plus home-equity loans (HELOC-inclusive)
s_0	`RwMSTAT`	1 = single (`MSTAT` $\in$ 4–8), 2 = couple ($\in$ 1–3)
i_0	`HwITOT`, `RwIEARN`, `SwIEARN`	Income process explained in appendix A.3
H_0	`RwSHLT`, `RwNRSNIT`	Concept explain in appendix A.1.2
weight	`RwWTRESP`	Carried forward from last positive value for LTC households

A.4.2 Moments

All moments in the SMM routine are computed using the RAN HRS files and come from the sample of all living household-waves from 1992-2022[24]. Age moments are pooled where age $\in \{\text{age}_{\text{moment}} - 1, \text{age}_{\text{moment}} + 1\}$. For example, if a household is observed at age 66 and the computation is for median liquid wealth at age 65, then that household is counted. All monetary values are deflated using 2022 as the base year from the FRED `CPIAUCSL` series. Moments are weighted using the HRS `HwWTRESP` variable, with unweighted LTC observations using weights carried forward from the last positive weight. Monetary variables are also capped at the 99th percentile within age. Table A.4 reports the entire list of moments, their sources, and a description of their construction.

24. Bequest moments are the exception as they use the sample of first deceased observation.

Table A.4: Targeted moments: source variables and construction.

Moment	RAND source	Construction
Panel A. Liquid wealth decumulation		
Median Liquid Assets $a \in \{65, 70, 75, 80, 85\}$	`HwACHCK`, `HwACD`, `HwABOND`, `HwASTCK`, `HwADEBT`	Weighted median of checking + CDs + bonds + stocks − other debt, IRA excluded.
p75 Liquid Assets $a \in \{65, 75, 85\}$	as above	Weighted 75th percentile of the same variable. Disciplines the upper tail that the median cannot, and is where co-holding capacity lives.
Panel B. Housing		
Ownership Share a $a \in \{65, 75, 85\}$	`HwAHOUS`	Weighted mean of $\mathbf{1}[\texttt{housing} > 0]$.
Median housing Wealth age 70	`HwAHOUS`	Weighted median among owners only.
Sale rate ages 65-75, Sale rate ages 75-85	`HwAHOUS`	Ten-year transition: of owners observed at $\text{age}_0 \pm 1$, the weighted share observed as renters at their first wave in $\text{age}_1 \pm 1$. Requires the household to appear at both endpoints, so attrition and mortality select the sample.
Sell Rate Healthy, Sell Rate Sick $w \in \{65\text{-}75,\ 75\text{-}85\}$	`HwAHOUS`, `RwSHLT`, `RwNRSNIT`	Same transition, split on health at the *left* endpoint: healthy = $\texttt{health} = G$, sick = $\texttt{health} \geq 1$. Because health is a fixed first-stage input, the healthy moment identifies the voluntary selling margin and the sick moment overidentifies the health channel.
Panel C. Mortgage debt		

Table A.4 continued

Moment	RAND source	Construction
Mortgagor rate by age $\in \{65, 75, 85\}$	`HwAMORT`, `HwAHMLN`	Weighted mean of $\mathbf{1}[\text{mortgage} > 0]$, where mortgage is first mortgage + home-equity loans. HELOC-inclusive, matching the SCF `MRTHEL` concept used in Buck (2026).
Median Mortgage age $\in \{70, 80\}$	as above	Weighted median among mortgagors.
Panel D. Co-holding (the mechanism)		
Co-holding rate age $\in \{65, 75, 85\}$	mortgage and liquid components above; `HwITOT`, `RwIEARN`, `SwIEARN`	Weighted share of all households with mortgagor = 1 *and* liquid assets greater than 3 months income, i.e. a mortgage held alongside more than three months of deterministic income in liquid form.
Weak buffer Mortgagors age 70	as above	Weighted share *of mortgagors* with liquid *below* the same co-hold. The complement of co-holding within the mortgagor population; separates "few mortgagors" from "mortgagors without buffers" as explanations of a co-holding miss.
Panel E. Bequests		
Mean wealth exit	prior-wave `HwATOTB`	Weighted mean of $\max(0, \text{wealth at exit})$ over death rows aged 65+, using the last wave's reported wealth as the bequest proxy. Raw, not winsorized. HRS weights are missing on death rows, so the carried-forward weight is essential here rather than merely corrective.
Panel F. Medicaid recipiency		

Table A.4 continued

Moment	RAND source	Construction
Medicaid low income, Medicaid high income age $\in \{65, 75, 85\}$	`RwGOVMD`, `SwGOVMD`	Weighted mean of an indicator for respondent or spouse receiving Medicaid in the past two years. "Low" pools entry quintiles Q1–Q2, "high" pools Q3–Q5; the pooling keeps cells adequate at 85 while preserving the gradient that identifies the asset test. Waves 1–2 (pre-1996) drop out because the question was not asked. Cells with fewer than 30 observations are set missing.

A.4.3 Simulated Method of Moments

Let $\boldsymbol{\Theta} = \{\beta, \sigma, \omega_o, \eta, \gamma, \zeta, f, \kappa, \sigma_B, \delta_0, \delta_1, \xi, \theta_{\text{DTI}}, \psi_{\text{refi}}, \underline{c}, \underline{c}_M(s=1), \underline{c}_M(s=2)\}$ denote the vector of parameters targeted in the SMM. The estimator is then the vector $\hat{\boldsymbol{\Theta}}$ such that

$$\hat{\boldsymbol{\Theta}} = \arg\min_{\boldsymbol{\Theta}} \mathcal{L}(\boldsymbol{\Theta}), \quad \mathcal{L}(\boldsymbol{\Theta}) = \big(\boldsymbol{m}(\boldsymbol{\Theta}) - \hat{\boldsymbol{m}}\big)' \boldsymbol{W} \big(\boldsymbol{m}(\boldsymbol{\Theta}) - \hat{\boldsymbol{m}}\big), \quad \boldsymbol{W} = \text{diag}\big(1/\hat{\sigma}_k^2\big) \tag{A.26}$$

where $\mathcal{L}$ is the loss function, $\boldsymbol{m}$ is the vector of simulated moments, $\hat{\boldsymbol{m}}$ is the vector of data moments, and $\boldsymbol{W}$ a weighting matrix. The moments in $\hat{\boldsymbol{m}}$ are computed from the data following table A.4.

Given a value for $\boldsymbol{\Theta}^n$, the solution method simulates $N = 11,891$ household rows used once from the initial distribution at $t = 0$, and $e = 0$ drawn from the stationary distribution at entry. The simulated moments are computed the same way as the data moments to yield $\boldsymbol{m}(\boldsymbol{\Theta}^n)$. The loss value, $\mathcal{L}$ is then updated and a new guess is initialized. The optimal choice $\hat{\boldsymbol{\Theta}}$ is solved from a two stage algorithm: first, an exploratory differential evolution algorithm is applied to find a sufficient basin for a Nelder-Mead polish.

The standard errors $\hat{\sigma_k}$ are computed by the following bootstrapping method:

1. Let $\mathcal{F}$ denote the distribution estimated from the HRS sample, then the data moments can be written as $\hat{\boldsymbol{m}}(\mathcal{F})$. Let Q be the total number of independent households in the HRS sample, the initial distribution is resampled Q times with replacement to obtain the new distribution $\mathcal{F}_l$. Each moment is then re-calculated to obtain $\hat{\boldsymbol{m}}_l(\mathcal{F}_l)$.

Table A.5: Sensitivity Analysis

Parameter	wealth	housing	debt	mechanism	bequest	rate
β	**43**	30	6	9	1	11
σ	28	**33**	10	5	4	20
ω_o	24	**31**	4	18	2	21
η	14	**37**	9	17	2	21
γ	14	**30**	13	16	3	23
ζ	**28**	20	15	12	3	23
f	21	16	11	**28**	1	22
κ	23	**35**	3	10	0	29
σ_B	29	**30**	10	9	5	17
δ_0	**33**	16	10	20	0	20
δ_1	22	**34**	9	15	1	18
ξ	20	27	4	7	1	**41**
θ_{DTI}	8	**34**	21	18	0	20
ψ_{refi}	17	18	14	15	3	**33**

Note: The sensitivity analysis follows from Andrews, Gentzkow, and Shapiro (2017). Each row refers to the specific parameter, each column summarizes a group of moments. Wealth aggregates all median and p75 liquid asset moments, housing aggregates all ownership and home sale moments, debt aggregates mortgagor rate and medians, mechanism aggregates co-holding and weak buffer moments, bequest is the mean wealth at exit moment, Medicaid take-up aggregates all Medicaid moments. Each cell is then constructed as the sum of the absolute value of each component within the column group.

2. After repeating 500 times to obtain $(\hat{\boldsymbol{m}}_1, ...\hat{\boldsymbol{m}}_{500})$ the bootstrap mean is

$$\bar{\boldsymbol{m}} = \frac{1}{500}\sum_{l=1}^{500}\hat{\boldsymbol{m}}_l. \tag{A.27}$$

The bootstrap variance is then approximately

$$\hat{\sigma}_k^2 = \frac{1}{499}\sum_{l=1}^{500}\left(\hat{m}_l^k - \bar{m}^k\right)^2 \tag{A.28}$$

for each data moment k.

A.4.4 Sensitivity Analysis

To validate the identification strategy, I use a Andrews, Gentzkow, and Shapiro (2017) sensitivity analysis, the results of which are reported in table A.5 With more moments than parameters, every estimate is a weighted combination of all of them. Let $\boldsymbol{\Lambda}$ denote the weight

where

$$\mathbf{\Lambda}[j,k] = \frac{\partial \hat{\Theta}_j}{\partial m_k} \tag{A.29}$$

where $\hat{\Theta}_j$ denotes the estimated parameter j, and m_k denotes the moment k. Thus, $\mathbf{\Lambda}$ measures how much each parameter moves in response to a change in data moments. Solving for $\mathbf{\Lambda}$ directly would be costly, so Instead, I compute

$$\mathbf{G} = \frac{\partial \tilde{m}_k}{\partial \hat{\Theta}_j} \tag{A.30}$$

by perturbing parameters by a fraction, δ_j, such that

$$\Theta_j^+ = \hat{\Theta}_j + \delta_j|\hat{\Theta}_j|, \quad \Theta_j^- = \hat{\Theta}_j + \delta_j|\hat{\Theta}_j| \tag{A.31}$$

and

$$\frac{\partial \tilde{m}_k}{\partial \hat{\Theta}_j} \approx \frac{\tilde{m}_k^+ - \tilde{m}_k^-}{\Theta_j^+ - \Theta_j^-}. \tag{A.32}$$

Then, since from the first order condition, the estimator solves

$$\mathbf{G}(\hat{\mathbf{\Theta}})'\mathbf{W}[\tilde{\boldsymbol{m}}(\hat{\mathbf{\Theta}}) - \boldsymbol{m}] = 0. \tag{A.33}$$

total differentiating with respect to $\boldsymbol{m}$ yields

$$\mathbf{G}'\mathbf{W}\left(\frac{\partial \tilde{\boldsymbol{m}}}{\partial \hat{\mathbf{\Theta}}}\right)\left(\frac{\partial \hat{\mathbf{\Theta}}}{\partial \tilde{\boldsymbol{m}}}\right) - \mathbf{G}'\mathbf{W} = 0 \tag{A.34}$$

and thus by rearranging

$$\frac{\partial \hat{\mathbf{\Theta}}}{\partial \boldsymbol{m}} = (\mathbf{G}'\mathbf{W}\mathbf{G})^{-1}\mathbf{G}'\mathbf{W} = \mathbf{\Lambda}. \tag{A.35}$$

Since Andrews, Gentzkow, and Shapiro (2017) define $\mathbf{G}$ as the Jacobian of the moment moment condition, $\boldsymbol{m} - \tilde{\boldsymbol{m}}$, while I use the Jacobian of the simulated moments, the two conditions differ by a minus sign, so to give the same interpretable conditions I use the matrix

$$\mathbf{\Lambda} = -(\mathbf{G}'\mathbf{W}\mathbf{G})^{-1}\mathbf{G}'\mathbf{W}. \tag{A.36}$$

Table A.5 reports the normalized sensitivity measure of $\mathbf{\Lambda}$ for each group of parameters. For a given parameter, the table sums the absolute value measures of $\mathbf{\Lambda}$ for the entire row, then reports the percentage that each parameter group contributes to the measurement of $\mathbf{\Lambda}$.

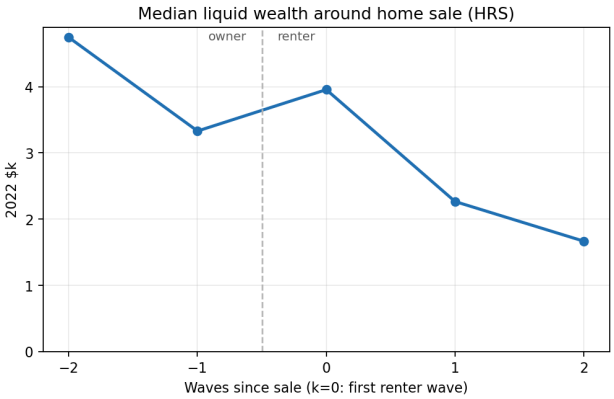


Figure A.1: Median liquid Wealth Prior to and After Sale

Note: Solid line plots median liquid wealth of sellers in the HRS for the last two waves of ownership (-2 & -1), the first wave following home sale (0), and two more waves after that (1 & 2). Vertical dashed line indicates transition from ownership status.

A.5 Estimation and Results Validation Tables

This appendix lists additional HRS data results in support of the discussions.

Table A.6: Financial transfers to children and grandchildren, HRS, by permanent-income quintile (2022 $k per two years).

Age	Quintile	Obs.	Share giving	Mean gift	Median gift, givers	Amount missing
65-74	Q1	7,833	15.7%	1.3	2.4	2%
65-74	Q2	8,351	23.8%	2.3	3.4	1%
65-74	Q3	8,930	33.4%	3.5	4.0	1%
65-74	Q4	9,323	42.1%	5.4	5.2	1%
65-74	Q5	8,649	53.5%	15.6	10.0	1%
75-84	Q1	5,030	8.9%	0.7	2.3	2%
75-84	Q2	6,519	17.5%	1.6	3.2	2%
75-84	Q3	7,430	28.1%	3.1	3.8	2%
75-84	Q4	7,524	37.1%	6.8	6.0	2%
75-84	Q5	6,077	52.4%	19.5	11.3	1%
85-95	Q1	2,574	6.7%	0.7	4.0	6%
85-95	Q2	3,065	15.2%	3.1	3.7	3%
85-95	Q3	2,955	22.7%	3.1	4.2	2%
85-95	Q4	2,904	31.1%	6.7	8.0	1%
85-95	Q5	2,056	47.0%	26.2	14.0	2%

Note: HRS waves 4–16 (1998–2022). Quintiles of household mean real income over ages 65+.

Table A.7: Financial transfers to children and grandchildren, HRS, by liquid wealth (2022 $k per two years).

Age	Group	Obs.	Share giving	Mean gift	Median gift, givers	Amount missing
65-74	Rest	33,641	31.2%	3.6	4.2	1%
65-74	Top liquid quartile	9,445	49.6%	14.5	10.0	1%
75-84	Rest	24,919	24.8%	3.3	4.3	2%
75-84	Top liquid quartile	7,661	48.1%	16.9	10.6	1%
85-95	Rest	10,261	18.6%	3.3	5.0	3%
85-95	Top liquid quartile	3,293	40.6%	18.8	13.6	1%

Note: Living households aged 65–95, HRS waves 4–16 (1998–2022). Transfers are household help to children or grandchildren since the previous interview (RAND HRS Family Data), deflated to 2022 dollars. Top liquid quartile: above the weighted 75th percentile of liquid wealth at the same age. Mean gift includes non-givers as zero. Amount missing: share of givers whose amount is not reported. For comparison, the model's 75th-percentile liquid gap widens by $10.0k per two years between 75 and 85.

Table A.8: Care spending, nursing-home residence and household composition around a home sale, HRS.

Age	Group	Nursing home before	Nursing home after	NH stay before	NH stay, sale	Household grew
65-74	Sell and rent	0.5%	4.4%	3.6%	8.9%	19.0%
65-74	Keep owning	0.1%	0.3%	1.5%	2.1%	6.9%
65-74	Sell, equity >= 150k	0.6%	5.4%	3.0%	10.5%	16.6%
65-74	Keep, equity >= 150k	0.1%	0.2%	1.4%	1.7%	5.8%
75-84	Sell and rent	4.4%	16.4%	9.6%	25.0%	14.2%
75-84	Keep owning	0.4%	1.3%	3.6%	5.9%	5.7%
75-84	Sell, equity >= 150k	4.2%	15.4%	10.9%	25.0%	13.5%
75-84	Keep, equity >= 150k	0.3%	0.9%	3.5%	5.5%	5.4%
85-95	Sell and rent	7.9%	27.8%	16.1%	42.0%	14.7%
85-95	Keep owning	1.7%	4.6%	7.4%	11.9%	6.7%
85-95	Sell, equity >= 150k	6.7%	20.5%	14.5%	33.8%	14.8%
85-95	Keep, equity >= 150k	1.2%	3.7%	7.1%	11.9%	7.1%

Note: Means. Nursing home: financial respondent living in a nursing home at this wave (before) and the next (after). NH stay: any nursing-home stay in the two years before the sale (before) or the two years of the sale (sale). Household grew: more people in the household at the next wave. Respondent-level items refer to the household's financial respondent.

Table A.9: Transfers to children and grandchildren around a home sale, HRS: owners who sell and rent versus owners who keep their home.

Age	Group	Households	Equity (median)	Give, sale window	Mean gift before	Mean gift, sale	Mean gift after	Gifts / equity
65-74	Sell and rent	1,197	91	32.5%	4.6	6.9	4.7	8.2%
65-74	Keep owning	29,203	174	37.9%	7.6	7.1	7.4	5.9%
65-74	Sell, equity >= 150k	340	273	47.2%	9.5	15.6	10.6	8.1%
65-74	Keep, equity >= 150k	15,332	285	44.4%	10.9	10.3	10.8	5.5%
75-84	Sell and rent	1,383	127	27.1%	5.7	9.5	7.3	10.2%
75-84	Keep owning	19,267	180	33.7%	8.3	7.2	7.4	6.2%
75-84	Sell, equity >= 150k	556	250	33.2%	9.1	14.0	13.9	9.5%
75-84	Keep, equity >= 150k	10,678	272	40.0%	11.8	10.1	10.0	5.9%
85-95	Sell and rent	804	133	24.1%	19.1	6.6	6.4	6.9%
85-95	Keep owning	4,824	168	25.7%	7.7	8.2	8.4	7.8%
85-95	Sell, equity >= 150k	337	240	32.5%	37.4	10.4	7.2	5.0%
85-95	Keep, equity >= 150k	2,513	259	30.3%	11.0	11.3	12.6	7.8%

Note: Owners at one wave who are alive at the next; sellers no longer own at the next wave. Transfers are reported for the period since the previous interview: before covers the two years before the sale, sale the two years in which it happened, after the following two years. Means include non-givers as zero (2022 $k). Give: share making any transfer in the sale window. Gifts / equity: mean transfers over the sale and after windows divided by mean home equity before the sale. Descriptive: sellers and keepers differ in health and family circumstances. HRS waves 4–16; RAND HRS Family Data.

Table A.10: Exits from ownership over the sale-rate windows, by household characteristics at the start of the window, model and HRS.

Window	Group	Exit model	Exit data	Model / data	Share of owners model	Share of owners data
65-75	All owners	5.6%	9.3%	0.60	–	–
65-75	Liquid Q1 (lowest)	8.5%	16.8%	0.51	–	–
65-75	Liquid Q2	7.5%	11.1%	0.67	–	–
65-75	Liquid Q3	3.9%	6.3%	0.61	–	–
65-75	Liquid Q4	2.5%	2.9%	0.87	–	–
65-75	Liquid < \$5k	8.2%	16.0%	0.51	33.6%	31.6%
65-75	Widowed in window	6.8%	16.0%	0.42	19.4%	11.2%
75-85	All owners	14.5%	20.2%	0.72	–	–
75-85	Liquid Q1 (lowest)	18.1%	31.0%	0.58	–	–
75-85	Liquid Q2	17.0%	17.5%	0.97	–	–
75-85	Liquid Q3	12.9%	17.3%	0.74	–	–
75-85	Liquid Q4	10.1%	14.9%	0.68	–	–
75-85	Liquid < \$5k	16.9%	31.1%	0.54	44.2%	24.7%
75-85	Widowed in window	13.7%	23.8%	0.58	29.8%	18.1%

Note: Owners at the first age who are alive at the second; exit: no longer an owner at the second age (the sale-rate moment definition). Model / data: ratio of the two exit rates. Share of owners: the group's share of the window's owners, shown where it is informative (liquid quartiles are 25% by construction). Liquid quartiles are weighted quartiles within the window's owners. Model sellers count as owners when they buy another home. HRS waves 4–16, balanced over the window.

Table A.11: Who drives the dispersion of welfare effects within net-worth quintiles (No homestead exemption, age 75): the 10% most affected households against the rest of their quintile.

Quintile	Effect, top	Effect, rest	Liquid, top	Liquid, rest	Between limits, top	Between limits, rest	Cash-poor owner, top	Cash-poor owner, rest
1	-78.1	-3.3	0.0	13.8	1%	21%	95%	9%
2	-119.8	-34.1	0.0	2.0	1%	5%	99%	64%
3	-130.8	-54.8	0.0	0.0	1%	3%	98%	73%
4	-97.0	-24.4	0.0	108.7	0%	1%	97%	38%
5	-76.3	-8.0	0.0	200.7	0%	0%	100%	18%

Note: Top: households in the top 10% of the absolute welfare effect within their net-worth quintile. Effect: mean welfare effect, 2022 \$k. Liquid: medians. Between limits: liquid wealth above the current Medicaid limit (\$2k single, \$3k couple) but within the expanded one (\$10k, \$15k), the households the asset-test reform makes eligible. Cash-poor owner: owner with under \$5k of liquid wealth.

Table A.12: Who drives the dispersion of welfare effects within net-worth quintiles (Asset test (10, 15), age 75): the 10% most affected households against the rest of their quintile.

Quintile	Effect, top	Effect, rest	Liquid, top	Liquid, rest	Between limits, top	Between limits, rest	Cash-poor owner, top	Cash-poor owner, rest
1	+48.7	+5.9	0.0	13.7	27%	18%	30%	16%
2	+72.2	+14.2	1.9	1.9	7%	4%	80%	66%
3	+57.1	+9.2	0.0	0.0	5%	2%	90%	74%
4	+39.2	+3.4	1.8	107.0	1%	1%	84%	40%
5	+21.0	+1.0	1.3	201.9	2%	0%	77%	20%

Note: Top: households in the top 10% of the absolute welfare effect within their net-worth quintile. Effect: mean welfare effect, 2022 $k. Liquid: medians. Between limits: liquid wealth above the current Medicaid limit ($2k single, $3k couple) but within the expanded one ($10k, $15k), the households the asset-test reform makes eligible. Cash-poor owner: owner with under $5k of liquid wealth.

References

Agarwal, Sumit, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, Tomasz Piskorski, and Amit Seru. 2017. "Policy intervention in debt renegotiation: Evidence from the home affordable modification program." *Journal of Political Economy* 125 (3): 654–712.

Aguiar, Mark, Mark Bils, and Corina Boar. 2025. "Who are the Hand-to-Mouth?" *Review of Economic Studies* 92 (3): 1293–1340.

Amromin, Gene, Jennifer Huang, and Clemens Sialm. 2007. "The tradeoff between mortgage prepayments and tax-deferred retirement savings." *Journal of Public Economics* 91 (10): 2014–2040.

Andrews, Isaiah, Matthew Gentzkow, and Jesse M. Shapiro. 2017. "Measuring the Sensitivity of Parameter Estimates to Estimation Moments." *The Quarterly Journal of Economics* 132 (4): 1553–1592. https://doi.org/10.1093/qje/qjx023.

Arrow, Kenneth J. 1971. *Essays in the Theory of Risk-Bearing.* Chicago: Markham.

Bédard, Nicolas, and Pierre-Carl Michaud. 2021. "Playing with fire? Household debt near retirement in Canada." *Canadian Public Policy* 47 (1): 56–71.

Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra. 2018. "House prices and consumer spending." *The Review of Economic Studies* 85 (3): 1502–1542.

Bertaut, Carol C, Michael Haliassos, and Michael Reiter. 2009. "Credit card debt puzzles and debt revolvers for self control." *Review of Finance* 13 (4): 657–692.

Borella, Margherita, Mariacristina De Nardi, and Eric French. 2018. "Who receives Medicaid in old age? Rules and reality." *Fiscal Studies* 39 (1): 65–93.

Braun, R Anton, Karen A Kopecky, and Tatyana Koreshkova. 2017. "Old, sick, alone, and poor: A welfare analysis of old-age social insurance programmes." *The Review of Economic Studies* 84 (2): 580–612.

Browning, Martin, and Thomas F Crossley. 2001. "The life-cycle model of consumption and saving." *Journal of economic perspectives* 15 (3): 3–22.

Buck, Jared. 2026. "Co-Holding with Closed-End Loans." Working paper, Concordia University.

Campbell, John Y, and Joao F Cocco. 2003. "Household risk management and optimal mortgage choice." *The Quarterly Journal of Economics* 118 (4): 1449–1494.

Chen, Anqi, Siyan Liu, and Alicia H Munnell. 2023. "What are the implications of rising debt for older Americans?" *Center for Retirement Research. Boston College. CRR WP* 11.

Chetty, Raj, László Sándor, and Adam Szeidl. 2017. "The effect of housing on portfolio choice." *The Journal of Finance* 72 (3): 1171–1212.

Chetty, Raj, Michael Stepner, Sarah Abraham, Shelby Lin, Benjamin Scuderi, Nicholas Turner, Augustin Bergeron, and David Cutler. 2016. "The association between income and life expectancy in the United States, 2001-2014." *Jama* 315 (16): 1750–1766.

Cocco, Joao F. 2005. "Portfolio choice in the presence of housing." *The Review of Financial Studies* 18 (2): 535–567.

Damodaran, Aswath. 2023. *Historical Returns on Stocks, Bonds and Bills: 1928–Current.* Dataset, NYU Stern School of Business. Accessed 2026/09/23. https://pages.stern.nyu.edu/~adamodar/New_Home_Page/datafile/histretSP.html.

De Nardi, Mariacristina, Eric French, and John B Jones. 2010. "Why do the elderly save? The role of medical expenses." *Journal of political economy* 118 (1): 39–75.

Druedahl, Jeppe, and Casper Nordal Jørgensen. 2018. "Precautionary borrowing and the credit card debt puzzle." *Quantitative Economics* 9 (2): 785–823.

Foote, Christopher L, Kristopher Gerardi, and Paul S Willen. 2008. "Negative equity and foreclosure: Theory and evidence." *Journal of urban economics* 64 (2): 234–245.

Friedberg, Leora, Wei Sun, and Anthony Webb. 2023. "The House: Is it an Asset or a Liability?" *Ann Arbor, MI. University of Michigan Retirement and Disability Research Center (MRDRC) Working Paper; MRDRC WP 2023-473. https://mrdrc.isr.umich.edu/publications/papers/pdf/wp473.pdf.*

Fulford, Scott L. 2015. "How important is variability in consumer credit limits?" *Journal of Monetary Economics* 72:42–63.

Ganong, Peter, and Pascal Noel. 2020. "Liquidity versus wealth in household debt obligations: Evidence from housing policy in the great recession." *American Economic Review* 110 (10): 3100–3138.

Gathergood, John, and Arna Olafsson. 2026. "The coholding puzzle: New evidence from transaction-level data." *The Review of Financial Studies* 39 (6): 1877–1908.

Gathergood, John, and Jörg Weber. 2014. "Self-control, financial literacy & the co-holding puzzle." *Journal of Economic Behavior & Organization* 107:455–469.

Genworth Financial. 2023. *Cost of Care Survey.* https://investor.genworth.com/news-events/press-releases/detail/972/genworth-releases-cost-of-care-survey-results-for-2023. Accessed 03/09/2026.

Greene, Claire, and Joanna Stavins. 2023. "Credit card debt puzzle: liquid assets to pay household bills." *International Review of Economics* 70 (4): 503–535.

Gross, David B, and Nicholas S Souleles. 2002. "Do liquidity constraints and interest rates matter for consumer behavior? Evidence from credit card data." *The Quarterly journal of economics* 117 (1): 149–185.

Hubbard, R Glenn, Jonathan Skinner, and Stephen P Zeldes. 1995. "Precautionary saving and social insurance." *Journal of political Economy* 103 (2): 360–399.

Islam, Khan, and Julien Picault. 2026. "Unveiling the co-holding puzzle: insights from two decades of Canadian household data." *International Journal of Bank Marketing,* 1–22.

Kaplan, Greg, and Giovanni L Violante. 2014. "A model of the consumption response to fiscal stimulus payments." *Econometrica* 82 (4): 1199–1239.

Kaplan, Greg, Giovanni L Violante, and Justin Weidner. 2014. *The wealthy hand-to-mouth.* Technical report. National Bureau of Economic Research.

Kopecky, Karen A, and Tatyana Koreshkova. 2014. "The impact of medical and nursing home expenses on savings." *American Economic Journal: Macroeconomics* 6 (3): 29–72.

Kopecky, Karen A., and Richard M. H. Suen. 2010. "Finite State Markov-Chain Approximations to Highly Persistent Processes." *Review of Economic Dynamics* 13 (3): 701–714.

Ludvigson, Sydney. 1999. "Consumption and credit: a model of time-varying liquidity constraints." *Review of economics and statistics* 81 (3): 434–447.

Lusardi, Annamaria, Olivia S Mitchell, and Noemi Oggero. 2020. "Debt and financial vulnerability on the verge of retirement." *Journal of Money, Credit and Banking* 52 (5): 1005–1034.

Medina, Paolina C, and Michaela Pagel. 2025. "Does Saving Cause Borrowing? Implications for the Coholding Puzzle." *The Journal of Finance* 80 (5): 2689–2738.

Meier, Stephan, and Charles Sprenger. 2010. "Present-biased preferences and credit card borrowing." *American Economic Journal: Applied Economics* 2 (1): 193–210.

Mian, Atif, Kamalesh Rao, and Amir Sufi. 2013. "Household balance sheets, consumption, and the economic slump." *The Quarterly Journal of Economics* 128 (4): 1687–1726.

Nakajima, Makoto, and Irina A Telyukova. 2020. "Home equity in retirement." *International Economic Review* 61 (2): 573–616.

———. 2017. "Reverse mortgage loans: A quantitative analysis." *The Journal of Finance* 72 (2): 911–950.

Pratt, John W. 1964. "Risk Aversion in the Small and in the Large." *Econometrica* 32 (1–2): 122–136.

Rouwenhorst, K. Geert. 1995. "Asset Pricing Implications of Equilibrium Business Cycle Models." In *Frontiers of Business Cycle Research,* edited by Thomas F. Cooley, 294–330. Princeton, NJ: Princeton University Press.

Shapley, Lloyd S. 1953. "A Value for n-Person Games." In *Contributions to the Theory of Games, Volume II,* edited by Harold W. Kuhn and Albert W. Tucker, 307–317. Annals of Mathematics Studies 28. Princeton University Press.

Shefrin, Hersh M, and Richard H Thaler. 1988. "The behavioral life-cycle hypothesis." *Economic inquiry* 26 (4): 609–643.

Shorrocks, Anthony F. 2013. "Decomposition Procedures for Distributional Analysis: A Unified Framework Based on the Shapley Value." *Journal of Economic Inequality* 11 (1): 99–126.

Stango, Victor, and Jonathan Zinman. 2009. "Exponential growth bias and household finance." *the Journal of Finance* 64 (6): 2807–2849.

Sussman, Abigail B, and Rourke L O'brien. 2016. "Knowing when to spend: Unintended financial consequences of earmarking to encourage savings." *Journal of Marketing Research* 53 (5): 790–803.

Telyukova, Irina A. 2013. "Household need for liquidity and the credit card debt puzzle." *Review of Economic Studies* 80 (3): 1148–1177.

Telyukova, Irina A, and Randall Wright. 2008. "A model of money and credit, with application to the credit card debt puzzle." *The Review of Economic Studies* 75 (2): 629–647.

Thaler, Richard. 1985. "Mental accounting and consumer choice." *Marketing science* 4 (3): 199–214.

Tobin, James. 1957. *Consumer Debt and Spending: Some Evidence from Analysis of a Survey.* Technical report. Cowles Foundation for Research in Economics, Yale University.

Vihriälä, Erkki. 2025. "Intra-household frictions, anchoring, and the credit card debt puzzle." *Review of Economics and Statistics* 107 (2): 510–522.

Yao, Rui, and Harold H Zhang. 2005. "Optimal consumption and portfolio choices with risky housing and borrowing constraints." *The Review of Financial Studies* 18 (1): 197–239.