Housing and Portfolio Choice over the Wealth Distribution

Zoltán Rácz *
Stockholm School of Economics

February 28, 2025

Abstract

Why do the rich take more financial risk and hence earn a higher return on their portfolios on average? In this paper, I argue that understanding the interdependence of optimal housing decisions, debt taking and portfolio allocation over the wealth distribution is key to explaining this robust empirical pattern. As apart from being a means of investment, housing also serves as a consumption good, households with a lower financial wealth to human capital ratio optimally choose a higher share of housing out of wealth. On the one hand, this implies that for relatively wealth-poor homeowners risky liquid assets are mechanically crowded out from their portfolio. Second, since this mechanism also makes poorer households optimally more leveraged, the effects are magnified by the wedge between borrowing and lending rates: if the interest rate on debt is higher, indebted households effectively face a lower risk premium, and thus are provided with lower incentives to hold risky assets. Calibrating a rich life-cycle model to the saving and home ownership profiles over age in Swedish administrative data I find that these mechanisms enable matching the increasing risky share pattern over the wealth distribution. I decompose the effect of different channels and also show that the model predicts a higher marginal propensity of stock investments for the rich.

1 Introduction

A long-standing challenge in the field of household finance concerns understanding the robust finding that wealthier households choose to hold portfolios offering

*SSE Department of Finance, Zoltan.Racz@hhs.se

†I am grateful to Lars Ljungqvist and Paolo Sodini for their advice, feedback and support. I also thank Gualtiero Azzalini, Laurent Bach, Corina Boar, Jaroslav Borovička, David Domeij, Tore Ellingsen, Mark Huggett, Paul Klein, Arash Nekoei, Robert Östling, Lars E.O. Svensson and seminar participants at the Stockholm School of Economics, New York University and conference participants at ESEM 2024 and EFA 2024 for useful discussions and comments.

higher risk and higher expected returns. An example of this pattern from Swedish administrative data is displayed in Figure 1. With the exception of the first two wealth deciles (in which even households with negative wealth are present), the share of cash decreases, the share of risky assets increases monotonically, while housing and debt display a hump-shaped pattern.

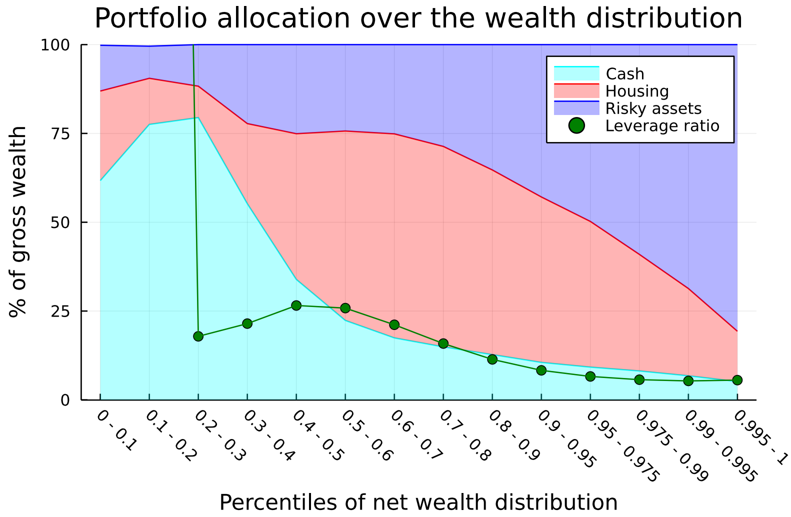

Figure 1: Average portfolio allocation of households belonging to different quantiles of the wealth distribution. The leverage ratio is measured as debt to gross wealth and is over 100% for the bottom two wealth deciles. Details on the data and variable construction can be found in Section 6.

Understanding the origins of these findings is crucial for two reasons: First, the higher risky share of the rich implies higher capital income on average, which can substantially increase wealth inequality as demonstrated by Hubmer et al. (2021). Therefore explaining portfolio choice is important to understand the causes and possible remedies of wealth inequality. In addition, heterogeneous portfolio composition patterns also mean heterogeneity in terms of exposure to any shock that is specific to an asset class, such as changes in house price, debt costs or stock market returns. This implies that to understand the non-linear effects of these shocks, the heterogeneous portfolio allocations need to be taken into account.

Unfortunately, an optimally increasing risky share in wealth is difficult to generate in models assuming standard preferences. In the baseline model by Merton (1969) and Samuelson (1969) with constant relative risk aversion preferences the optimal risky share is constant in the absence of labor income. After adding a realistic income process to the model, the optimal risky share becomes a decreasing function of wealth, keeping the income process fixed. This is so because as idiosyncratic labor income and the stock market are very weakly correlated, regarding the portfolio choice decision, human capital behaves as a near substitute of the risk-less asset.

A large literature has extended these standard portfolio choice models to generate an increasing risky share over wealth. One approach lies in choosing more sophisticated ways of modeling labor income risk: Chang et al. (2018) explore a framework where individuals have noisy expectations about their abilities at the beginning of their careers. This generates additional income risk for the young, which is resolved gradually over their lifetime. Depressing risk-taking for the young contributes to an increasing risky share over wealth, since young individuals are poorer than the average. Catherine (2021) introduces cyclical skewness in the labor income risk: in his paper disastrous labor market outcomes are more likely to occur in times of bad stock market performance, hence individuals with little wealth are reluctant to invest in stocks. While these papers can generate low and increasing risk-taking at the bottom of the wealth distribution, the effect of these channels fades away for the rich, as their human capital is relatively less important. There are also several papers generating increasing risky share even on the top of the wealth distribution, all of which involve some kind of non-homothetic preferences. Carroll (2000) assumes holding wealth provides utility directly, besides through financing consumption, while Wachter and Yogo (2010) include a luxury good in their model, generating decreasing relative risk aversion as agents get richer and optimally spend relatively more on luxuries. Finally, Cioffi (2021) achieves the same effect by modeling housing as a necessity.

In this paper, I argue that the portfolio allocation patterns illustrated by Figure 1 can be largely explained by taking into account the rich interactions of housing, debt and portfolio allocation choices. In particular, I show that the optimal housing to total wealth ratio is decreasing over the wealth distribution even with standard homothetic preferences. To understand this result, it is essential to explore the dual role of housing in consumption and saving decisions. Namely, while real estate is clearly a device of saving and investment, home ownership is also beneficial through obtaining housing services. Looking at housing chiefly as an instrument of investment it does not seem obvious at all why the ratio of housing to other components of wealth should depend on net wealth, assuming preferences are roughly homothetic. Indeed, in classical portfolio choice models the amount of wealth is irrelevant to determine portfolio shares. Even preference parameters, such as risk aversion, affect only the amount of safe assets in the portfolio, while the optimal share of different assets in the risky portfolio depends only on the joint distribution of the assets' returns.

To see why an asset providing a consumption flow changes this conclusion, it is simplest to consider a setting where all returns are deterministic and housing gives a lower return than alternative assets. In this case, if real estate provided no housing services, buying a house would always be a suboptimal decision. Therefore, abstracting away from transaction costs, homeownership can be interpreted directly as a way of consuming housing services, such that its cost equals the foregone capital income from choosing to save into the house instead of a superior vehicle of investment. As long as house transactions are not too large, this observation remains largely true in a more realistic model as well.

Once housing is primarily thought of as an object of consumption, expecting a non-constant share of housing in wealth is natural. Indeed, consumption and end-of-period savings depend on income and initial wealth in a very different manner: First, raising one's income in the current and all future periods by the same amount would induce an approximately identical change in consumption and leave savings intact. Second, due to consumption smoothing, increasing initial wealth has a much stronger effect on end-of-period savings than on consumption. As long as the importance of housing services in total consumption is roughly constant, this implies that the optimal share of housing from total wealth should decrease in the wealth-to-income ratio. In particular, I show that in a frictionless setting, there is a linear relationship between the share of human capital out of net worth (defined as the sum of human capital and net wealth) and the share of housing out of net wealth. I also provide descriptive evidence that such a pattern can actually be observed among Swedish households.

As both in data and typical heterogeneous agent models income is a lot less unequally distributed than wealth, there exists a negative correlation between wealth and the importance of human capital in one's net worth, which translates to the observation that the housing ratio appears to be a decreasing function of wealth among homeowners. A declining housing share among owners contributes to the increase in portfolio share of risky liquid assets over the wealth distribution in two ways: First, through the mechanism described above, housing mechanically crowds out risky assets in the case of poorer home-owners. Second, when the interest paid on debt is higher than that of the risk-free deposit rate, stock holding is less advantageous to indebted households than the rest of the population, since the effective risk premium is smaller. As the higher housing shares of the poor can only be supported by relatively higher debt, the poor are disproportionately more leveraged, and this latter channel again contributes to an increasing risky share over the wealth distribution. Besides the increasing average risky share, this model can also match the high leverage ratio of the poor and the hump-shaped share of housing. On the bottom of the wealth distribution housing share increases due to the transition from renting, after which it starts decreasing as the wealth to human capital ratio of households increases.

There is already a large literature investigating the relationship between housing, debt and portfolio choice decisions, and hence many ingredients of this model can be found in earlier papers in some form. Already Yao and Zhang (2004), Cocco (2005) and Flavin and Yamashita (2011) pointed out that optimal housing choice is increasing in human capital and that housing can crowd out risky investments. One contribution of this paper is determining in a simple frictionless model exactly in what way optimal housing level is determined by wealth and human capital. Namely, it is shown that the optimal share of housing in wealth is linear in the reciprocal of the share of wealth in net worth, defined as the sum of wealth and human capital. Furthermore, it is demonstrated that the same functional form provides a fair approximation of housing decisions even in a standard life-cycle model. My analysis of the interactions of portfolio choice decisions and interest rate wedges in the risk-free rate is related to Davis et al. (2006) and Willen and Kubler (2006), who show that a higher borrowing rate than the risk-free rate can help explain empirically plausible age profiles for stock holdings without relying on large values of risk aversion. However, their analysis lacked housing: while due to the corner solutions caused by the interest rate wedge, the authors found solution methods relying on first-order conditions more desirable, at that time there was no such method to solve models featuring discrete decisions. I overcome this difficulty by applying a modified EGM method following Iskhakov et al. (2017) and Fella (2014) and hence in spite of the presence of housing I obtain precise and easy to interpret policy functions. In fact, solving a standard life-cycle model with housing and portfolio choice by a first-order condition-based method is the second contribution of this paper. A third contribution is that by combining the housing and higher borrowing rate channels, I build a model that is able to replicate portfolio choice patterns for homeowners over the whole wealth distribution. By presenting appropriate counterfactuals I also demonstrate that the fit is indeed due to the mechanisms discussed above and explain in which part of the wealth distribution they matter. Finally, I show that in this model the marginal propensity of stock investments is roughly increasing in wealth. This finding is of interest due to the recent policy debate following the stimulus packages related to the Covid crisis and the subsequent flow of retail investments to the stock market.

The most relevant paper to the current one is probably Cioffi (2021), who also explains the increasing risky share schedule over wealth through housing. However, he does so by assuming non-homothetic preferences: since housing in his model is a necessity, poor agents will optimally own a relatively larger home and consume relatively more housing services. As they get richer, the curvature of the utility function changes, increasing the optimal risky share at the same time as housing becomes less important. The intriguing question is, given the mechanisms in the current paper, why does Cioffi (2021) need to resort to non-homothetic preferences to generate these patterns? The answer most likely lies in debt: As Cioffi (2021) only models unbacked loans of quite limited size, but his model lacks mortgages, most households cannot choose the optimal housing level postulated by the ratio of their wealth and human capital. This also diminishes the crowding out effect of housing, crucial to how the current paper fits the data.

In addition, it is also worth mentioning Chetty et al. (2017), who estimate the effect of housing on portfolio choice and find that while additional home equity causes owners to choose a higher risky share, increasing both the house value and mortgage by the same amount depresses risk-taking. These findings are in line with my model, since as it is discussed in Section 7.3, from the point of portfolio choice decisions, the difference between the level of risk-free assets and the borrowing limit is the equivalent of cash in models with no borrowing. Since this paper allows borrowing against one's home, house value changes correspond to varying borrowing limits, directly affecting the implicit cash to consider when choosing the optimal amount of risky assets.

The rest of the paper is organized as follows: first in Section 2 I study optimal housing choice in a simple toy model. Then in Section 4 a rich life-cycle model is introduced with housing and portfolio choice with the aim of inspecting the effects of housing in a more realistic setup. Next, in Section 5 and 6 I discuss the solution and calibration of the model. Section 7 present the result and discusses the mechanisms of the model in detail. Finally, Section 8 concludes.

2 Toy model

The aim of this section is to illustrate how one can interpret housing choices as optimal consumption decisions when frictions are absent. In this section therefore there are no borrowing limits (apart from the last period) or transaction costs. For the sake of obtaining closed form solutions, I also assume that there is no uncertainty.

The household lives for T periods and gains utility from housing services and non-durable consumption. Housing is expressed in units of non-durable consumption equivalents and hence its relative price is 1. Saving is possible into two assets: bonds and houses, with constant gross returns R and R^h. For simplicity assume that no rental market exists for housing and hence the amount of housing owned determines how much housing services one consumes in a given period.

The household maximizes

$$\sum_{t=0}^{T} \beta^t U(c_t, h_t)$$

such that

$$b_t + h_t = Rb_{t-1} + R^h h_{t-1} + y_t - c_t \qquad \forall t \tag{1}$$

$$Rb_T + R^h h_T \geq 0$$

The first constrain is the budget constraint, while the second constraint ensures that the agent does not die in debt.

Absent a borrowing limit on bonds, the problem would not possess a well-defined solution when $R^h \geq R$ holds, assuming that utility is strictly increasing in the housing services. Indeed, there would be no obstacle of consuming an arbitrary amount of housing services, financed by debt.

From now on assume that $R^h < R$, i.e. the financial return on housing is strictly lower than that of the alternative asset. The economic meaning of this relation is that ignoring the services it provides, housing is always a bad investment. The only way to make housing a profitable investment is to make use of the housing services it generates. $R^h < R$ creates an interesting trade-off in the portfolio choice decision: the bond pays a higher return, but housing provides services contributing to current utility. When consuming housing instead of investing more into bonds, the agent has to sacrifice some interest income as a price. This intuition is formalized next. First note that $\frac{R-R^h}{R}h_t$ equals the foregone capital income from consuming housing instead of saving in bonds, discounted to present terms[1]. By subtracting this term from both sides of equation (1),

$$b_t + h_t - \frac{R - R^h}{R}h_t = Rb_{t-1} + R^h h_{t-1} + y_t - c_t - \frac{R - R^h}{R}h_t$$

is obtained, which can be significantly simplified by introducing two economicaly meaningful new variables: Define the bond equivalent of total savings a_t and total expenditure x_t as

$$\begin{aligned} a_t &= b_t + \frac{R^h}{R}h_t \\ x_t &= c_t + \frac{R - R^h}{R}h_t \end{aligned} \tag{2}$$

Then the optimization problem can be rewritten as

$$\begin{aligned} &\sum_{t=0}^{T} \beta^t v(x_t) \\ \text{s.t.} \quad & a_t = Ra_{t-1} + y_t - x_t \qquad \forall t \end{aligned}$$

where v is the indirect utility function where u is maximized taking equation (2) as a constraint. This problem is a standard consumption/saving problem, suggesting that in the absence of transaction costs, borrowing limits and other frictions, housing does not fundamentally change the nature of optimal consumption decisions. In this framework, housing is a consumption good that is paid for by sacrificing some interest income.

To make this point more concrete, assume that $R\beta = 1$ and the utility function is Cobb-Douglas over housing and non-durable consumption, i.e. $u(c_t, h_t) = \tilde{u}\big(c_t^{1-\omega} \cdot$

[1] In an alternative model where housing services are obtained on a frictionless rental market instead of home ownership, $\frac{R-R^h}{R}$ would be replaced by the rental rate, but all the rest of the analysis would stay identical.

h_t^ω). In this case it is easy to show that optimal housing choice is given by

$$h_t = A_t\Big(w_t + \sum_{s=0}^{T-t} \frac{y_{t+1}}{R^s}\Big) = A_t\Big(w_t + HC_t\Big),$$

showing that optimal housing level equals an age-dependent fraction one's total net worth, which is the sum of current wealth and human capital. Naturally, as housing a kind of consumption, it turns out to be optimal to consume housing services out of future income, and hence higher human capital implies a higher optimal housing level. Note that relationship is not due to the investment good role of housing and in fact, it goes against that: indeed, higher future incomes depress the level of optimal savings. Intuitively, this already implies that the importance of human capital relative to total wealth must be an important determinant of the optimal share of housing out of total wealth, which formally follows from the above equation:

$$\frac{h_t}{w_t} = A_t\Big(1 + \frac{HC_t}{w_t}\Big) \tag{3}$$

I have shown that in this simple framework with unlimited borrowing, the optimal share of housing to wealth is decreasing in the wealth to human capital ratio. In the rest of the paper, a full-fledged life-cycle model is considered to investigate the magnitude and consequences of this effect in a more realistic setting.

3 Empirics (draft)

After deriving a simple equation determining the optimal housing share, it is of course an interesting question whether a similar relationship exists in the data. In the following, I provide some suggestive correlational evidence that it is indeed so. As many households hold very little net wealth, directly testing equation (3) would lead to imprecise estimates. This can be avoided by considering the sub-sample of home-owners and rearranging the equation as follows:

$$\frac{w_t}{h_t} = \frac{1}{A_t}\frac{w_t}{w_t + HC_t} \tag{4}$$

After estimating households human capital in Swedish administrative data, and plotting the quantities present on the two sides of equation (4) for the home-owner sample, I obtained the following figure.

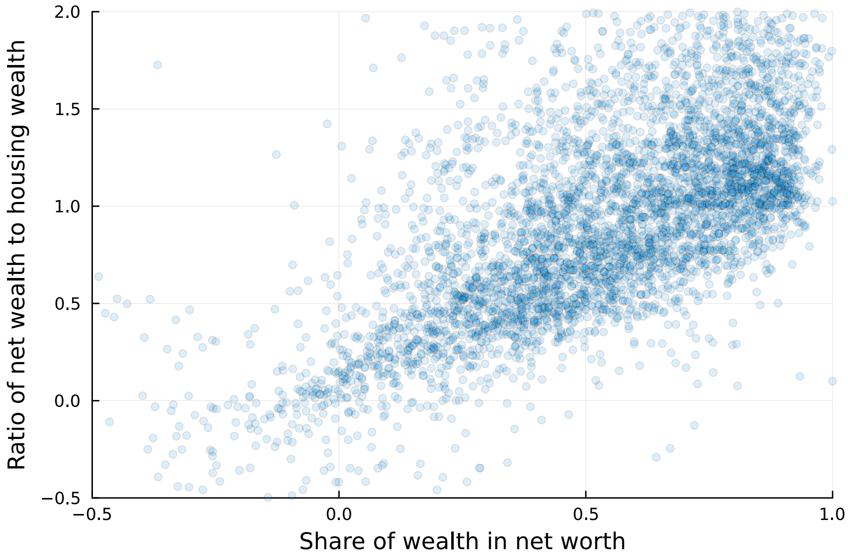

The strong positive correlation provides cross-sectional evidence that the share of human capital out of net wealth is a crucual determinant of housing choices.

4 Full Model

4.1 Demographics and intratemporal utility

I analyze a partial equilibrium overlapping generation economy populated by a continuum of finitely lived households indexed with i, which are ex-ante identical except for their age (j), starting wealth and starting persistent income level ($z_{i,25}$). Each period a measure one of 25 years old households are born. Survival is stochastic, but 100 years old households die for certainty. In general, the probability that a household of age j will be alive next period is denoted by q_{j+1}. One period in the model corresponds to one year. In each period, households derive instantaneous utility from non-durable consumption (c) and housing services (h) according to Cobb-Douglas preferences:

$$U(c,h) = h^{\omega} c^{1-\omega}$$

where ω is the optimal share of housing in expenditure in the absence of frictions.[2]

[2] In an earlier version of this paper, a CES-utility function with subsistence levels was used, and the main findings presented in this paper were found to be robust in that more general setup.

4.2 Income process

Working-age households supply labor inelastically and are subject to exogenous labor income shocks. Their log labor income is composed of a deterministic secular growth term (gt), a deterministic age term (f), a persistent (z) part following an AR(1) process and a transitory (ν) stochastic part.

$$y_{ij} = gt + f_j + z_{ij} + \nu_{ij}$$
$$z_{ij} = \rho z_{ij-1} + \varepsilon_{ij}$$

The two shocks ε and ν are iid normal random variables with zero mean

$$\varepsilon_{ij} \sim N(0, \sigma_\varepsilon^2)$$
$$\nu_{ij} \sim N(0, \sigma_\nu^2)$$

and $z_{i,25}$ is drawn from the implied stationary distribution of z. Note that all shocks to labor income are idiosyncratic and the sole aggregate change is through the deterministic, exponential growth at gross rate $G = \exp(g)$. Households at age 65 retire and afterwards their persistent income level is deterministically pinned down by their persistent income state at the time of retirement. However, they are still subject to transitory shocks which have smaller variance than during working time, i.e. for $j > 65$ we have

$$y_{ij} = gt + f_j + z_{i,65} + \nu_{ij}^r$$
$$\nu_{ij}^r \sim N(0, \sigma_{\nu^r}^2)$$

The presence of income risk is set to match the observed cross-sectional variance of disposable income residuals of retired persons, but it is also instrumental for the technical purpose of smoothing out jumps arising from non-convexities in integrands when computing expectations. Since the model does not feature risks related to unexpected medical expenses, the true total uncertainty faced by the elderly is probably still underestimated.

4.3 Housing

Housing services (h) can be obtained either by renting or by owning a house (H). Homes of any size H are in infinite supply at market value $P_t^h H$ where P^h is an aggregate price level subject to stochastic deviations around a deterministic trend as follows:

$$P_t^h = G_h^t \exp(\tilde{p}_t^h)$$
$$\tilde{p}_t^h = \rho^h \tilde{p}_{t-1}^h + \varepsilon_t^h \qquad \varepsilon_t^h \sim N(0, \sigma_h^2)$$

Owners consume housing services corresponding to their housing level, i.e. it is not possible to rent another apartment nor to rent out part of one's own home. In particular, housing services equal to the size of the home multiplied with G_h^t, which represents the improving efficiency of utility provided by a dwelling of constant size. Thus

$$h_{it} = G_h^t H_{it} \qquad \text{if } H_{it} > 0.$$

where $H = 0$ corresponds to renting. Therefore it is assumed that the secular trend in house prices is fully justified by home improvements leading to homes providing higher utility, and hence the aggregate price changes do not cause trends in the size of demanded houses. Housing consumption for renters is flexible but entails paying a rental cost amounting to $\tau P_t^h h/G_h^t$, where τ is a real rental rate including operational costs. Rescaling by G^t is necessary, since h/G_h^t expresses the current home size capable of producing home services h in time t and P_t^h is the price of homes expressed in size terms.

Selling or buying real estate involves transaction costs: When purchasing a house one has to pay $\Phi > 1$ times the house value, but selling results in income corresponding to only $\alpha < 1$ times house value. In addition, home owners must pay a maintenance cost amounting to χ times the home value. Apart from repair activities counteracting any depreciation of real estate, χ also includes operational and home improvements costs. To sum up, the total net costs originating from discrete home ownership decisions sum up as

$$D(H_{it-1}, H_{it}, P_t^h) = P_t^h H_{it} \chi \mathbb{1}_{H_{it}>0} + \big(H_{it}\Phi - H_{it-1}\alpha\big) P_t^h \mathbb{1}_{H_{it-1} \neq H_{it}} \tag{5}$$

4.4 Asset Choice

In addition to housing, agents also have access to several liquid means of saving (s): They can invest into bonds (B) offering a risk-free gross rate R^f and stocks (ξ) with stochastic risky gross rate R_t with $\log(R_t) \sim N(\mu_M, \sigma_M^2)$. When participating in the stock market ($\xi > 0$), a yearly participation cost F has to be payed. In addition, it is possible to take on debt in the form of mortgage (M) and consumption loans (L) with constant rates R^m and R^l, respectively. To sum up, liquid savings can be allocated as below:

$$s_{it} = B_{it} + \xi_{it} + F\mathbb{1}_{\xi_{it}>0} - M_{it} - L_{it}$$

All liquid investments are subject to a no short-position constraint. Furthermore, one cannot take a mortgage larger than a δ fraction of her current house value and consumption loan has an exogenous upper limit $\overline{L}$. In addition, both kinds of debt are bounded from above by a modified natural debt limit, which is a function of age and the persistent income state and is set such that a household's expected lifetime income HC must be sufficient to cover repayment even without spending more than a prespecified η fraction of yearly income on debt costs. This means that as households age, borrowing constraints get tighter, but interest rates are not affected.

$$0 \leq B_{it} \tag{6}$$
$$0 \leq \xi_{it} \tag{7}$$
$$0 \leq M_{it} \leq \min\left\{\eta^m HC(z_{it}, j), \delta P^h_t H_{it}\right\} \tag{8}$$
$$0 \leq L_{it} \leq \min\left\{\eta^c HC(z_{it}, j), \overline{L}\right\} \tag{9}$$

It is assumed that mortgages are renegotiated each year costlessly according to current conditions. This is a strong and unrealistic assumption, but is necessary to avoid adding one more state variable to the model and it is line with most papers in the literature featuring both housing and portfolio choice. It is assumed that

$$R^f < R^m < R^l$$

As a consequence, households hold debt only if they do not hold bonds and consumption loan is chosen only when the total mortgage capacity is exhausted. This also means that one could equivalently treat the three risk-free assets as one, with a risk-free but state-dependent interest rate.

4.5 Bankruptcy

Due to a combination of low savings, low income outcome and tightening credit constraints caused by falling house prices, it can happen that even after selling its home, a household cannot finance positive consumption without taking on more debt than allowed. In this case bankruptcy occurs: any home owned is lost and savings are set to the minimum. In other words, debts are renegotiated such that the household is moved to the borrowing constraint, but all further debts are canceled. In addition, expenditure is exogeneously set to ζ, which plays the role of a consumption floor as in Hubbard et al. (1995) and Ameriks et al. (2011). ζ represents unmodeled insurance mechanisms in the form of transfers either from the government or from relatives. Intuitively, this parameter will determine how hard households try to avoid staying too close to their respective borrowing constraint.

4.6 Timing and Intertemporal Decisions

At the beginning of each period t, the agent is given her idiosyncratic cash on hand a_{it}, persistent income and housing level, and the aggregate house price and wage level. For new-born households cash on hand is composed of bequests and labor income, while for older households bequests are replaced by their own after return liquid savings ($\hat{s}_{it}$) from the previous period. She then chooses her current housing level, how much non-durable goods and housing services to consume in the current period, how much to save for the next period s_{it}, and how to allocate those savings across the liquid instruments. Home owners must consume housing services equaling to their home size, while renters may adjust housing consumption freely in each period. Utility is aggregated over time according to Epstein-Zin preferences.

To summarize, the maximization problem of agent i is

$$\begin{aligned} V_j(G^t, P_t^h, a_{it}, z_{it}, H_{it-1}) = \max_{\{c,B,L,M,\xi,H,h\}} \Big\{ &(1-\beta)U\big(c_{it}, h_{it}\big)^{1-\psi} \\ &+ \beta\Big(q_{j+1}\mathbb{E}_t\Big[V_{j+1}(G^{t+1}, P_{t+1}^H, a_{it+1}, z_{it+1}, H_{it})^{1-\gamma}\Big] \\ &+ (1-q_{j+1})\mathbb{E}_t\Big[B(G^{t+1}, P_{t+1}^H, \hat{s}_{it+1} + \alpha P_{t+1}^h h_{it})^{1-\gamma}\Big]\Big)^{\frac{1-\psi}{1-\gamma}} \Big\}^{\frac{1}{1-\psi}} \end{aligned}$$

subject to

$$a_{it} = c_{it} + s_{it} + \tau \frac{h_{it}}{G_h^t} P_t^h \mathbb{1}_{H_{it}=0} + D(H_{it-1}, H_{it}, P_t^h) \tag{10}$$

$$s_{it} = B_{it} + \xi_{it} + F\mathbb{1}_{\xi_{it}>0} - M_{it} - L_{it} \tag{11}$$

$$\hat{s}_{it+1} = R_{t+1}\xi_{it} + R^f B_{it} - R^m M_{it} - R^l L_{it} \tag{12}$$

$$a_{it+1} = \hat{s}_{it+1} + \exp(y_{it+1}) \tag{13}$$

$$h_{it} = G_h^t H_{it} \qquad \text{when } H_{it} > 0 \tag{14}$$

$$H_t \in \{0, H_1, \dots H_l\} \tag{15}$$

and the no short position and borrowing constraints from section 4.4. Households have a bequest motive for saving. It is well known that the curvature of the bequest function can vastly affect savings for the old Nardi and Fella (2017). In this model non-trivial curvatures are already present due to the non-convexities caused by housing decisions, therefore for the sake of expositional clarity I avoid "warm-glow" type bequest utilities and choose a more neutral approach. It is assumed that the bequest function coincides with a scaled version of newborn households' value function. Hence,

$$B(G^{t+1}, P_{t+1}^H, \hat{s}_{it+1} + \alpha P_{t+1}^h h_{it}) = \kappa \mathbb{E}_t\Big[V_{25}(G^{t+1}, P_{t+1}^h, \hat{s}_{it+1} + \alpha P_{t+1}^h H_{it} + Y_1, z_{i25}, 0)\Big]$$

where κ determines the strength of bequest motive and z_{i25} is assigned from the cross-sectional stationary distribution of z. Note that when simulating the model, only a constant fraction θ of the bequest is given to an offspring, the rest is distributed evenly across the whole population, but this is not taken into account in the optimization problem. This is a necessary simplification: while in a more consistent setting the households would take into account how bequest are distributed over the whole population and how the marginal utility of wealth depends on the age and wealth of the inheritor, this dimension is abstracted away from in the current formulation. Instead, the optimal saving decision problem is solved by assuming that a young household will receive the bequests.

5 Solution of the model

5.1 Optimal policies

By taking advantage of the recursive structure of the problem, the optimal policies can be found backwards: On one level this means that the algorithm starts at age 100 and progresses towards the young. Furthermore, each period can be divided sequentially into three sub-periods: After learning the value of current shocks the agent decides on her housing level, then decides how much of the remaining wealth she should consume or save, and lastly decides how her liquid savings are allocated among risky and safe assets. This implies an appropriate order of solution: First the optimal risky asset allocation is found conditionally on the housing decision and the amount of total liquid savings. Once the optimal risky share as a function of liquid savings and housing is known, the marginal value from liquid savings is determined for each housing choice: This allows for solving for the optimal consumption-saving decision as a function of the discrete housing choice and the cash on hand remaining after potential costs involving home owning. Finally, the optimal discrete housing decision is solved by simply comparing values resulting from the potential alternatives.

On the technical side, the main challenge is caused by the fact that the combination of discrete and continuous decisions gives rise to non-concave value functions. Inconveniently, this means that first order conditions are in general only necessary but not sufficient conditions for optimality. To counter this problem I follow Fella (2014) and Iskhakov et al. (2017) who develop a version of the endogenous grid point method by Carroll (2006) with discrete choices. The key idea is that even though optimal consumption may be non-monotonic and the optimal saving policy may be discontinuous, in a large class of models optimal savings is still a monotonic function of cash-on-hand. This observation allows for efficiently identifying the optimal sections of the candidate consumption correspondence pinned down by the first order conditions, resulting in the optimal consumption policy. Also due to non-convexities, integrands are typically not smooth, and hence extra care needs to be taken when computing expectations. All further details on the numerical procedures are relegated to Appendix B. In the end, computation of all risky share and consumption decisions are based on first order conditions, which results in higher precision and speed than what value function iteration would allow for.

Since the bequest function is linked to the value function of the youngest agents, an extra outer loop is needed when finding optimal policies. Namely, first all policies are computed using a guessed bequest function. Then the bequest function is updated using the computed value function of 25 years old households and this process is repeated until the change in the bequest function is small enough. Fortunately, using a sensible guess in practice already after two iterations sufficient precision is achieved. This happens so, because new-born households are so far from death in expectation that their value function is barely influenced by the bequest function.

5.2 Wealth distribution

After solving for the policy functions, the steady state distribution is obtained by simulating several generations over time. The first cohort starts with no wealth. I assume that a θ fraction of dying agents' wealth is given to a newborn household, while the remainder is evenly distributed among all survivors in the economy. Receiving bequests from the previous cohort, new and new cohorts are simulated until the distribution of state variables does not change substantially anymore across cohorts. When solving for the wealth distribution it is assumed that aggregate shocks realize at their mean values. Of course, since this was not assumed when solving for optimal policies, the presence of aggregate uncertainty still affects the steady state distribution.

6 Calibration

The parameters used when solving the model are chosen to provide a suitable basis for comparing policy functions over the wealth distribution obtained from the model with empirical estimates using Swedish administrative data. Most parameters describing the economic environment of households are exogenously set, either based on own computations or other sources. The rest of the parameters, in particular those determining preferences are estimated to match aggregate saving and portfolio allocation patterns over the life cycle.

6.1 Exogenously set parameters

Throughout the paper Swedish administrative data is used between 1994-2015, including a rich set of yearly variables on demographics and income on the individual level for the whole universe of Swedish population. Since household identifiers are included, any quantities can be computed both at the household or individual level. In addition, for years 1999-2007 a rich data on wealth is available on the level of individual securities, bank accounts or other assets, collected due to a wealth tax. A detailed description of this dataset is available in Bach et al. (2020).

Income process parameters and age effects are estimated on disposable income net of capital income of 25-60 years old men. The parameters of the stochastic process are estimated by GMM to match variances and one year autocovariances of log real income residuals. The latter are obtained by projecting out year, age, cohort, education group and occupation dummies from log real income. Using year, age and cohort fixed effects at the same time involves a multicollinearity problem, which is solved by assuming that cohort effects are orthogonal to a trend and sum to 0. I chose this normalization to provide a close match to the model, which also abstracts away from cohort effects. [3] The age effects supplied by the same regression are used as the age profile figures f_j. I also tried estimating the income process parameters by separating the sample according to education level. As moments of income residuals were in each case fairly similar to that of the pooled sample, the results from the pooled regression are used when solving the model. However, apart from affecting stochastic properties of the income residuals, fixed effects also influence the level of average earnings. By abstracting away from these when solving the model, I most likely underestimate the true level of inequality in terms of income, even if the level of income risk for each household is roughly correct.

The parameters of the housing price process are estimated from the yearly price index of detached houses provided by Statistics Sweden, which is available since 1981. The series is made nominal using SCB's CPI data after which a trend is removed from the log real values. The resulting trend coincides with the growth term is real wage up to 3 digits, hence it is assumed that the two growth rates coincide which greatly simplifies the solution of the household's problem. ρ_h and σ_h are set to match variance and auto-correlation of the resulting residual series. It should be noted that while assuming that all involved stochastic processes are Markov is essential to solve to model without additional state variables, this assumption might be a strong one regarding house prices as the series appears to have substantial deviations on longer frequencies (5-8 years). During the solution of the model, all continuous Markov chains are replaced with a discretized version obtained by the Rouwenhorst method.

Housing share ω is set to match the share of housing expenditures out of disposable income in the case of renters in Sweden in 2007 (the earliest available figure), see Statistics Sweden (2022a). This figure is suitable benchmark as long as disposable income and total expenditure of renters is not drastically different. Transaction cost for buying a house is set to 3.5%, approximately representing a 1.5% stamp duty on the real estate value and an extra 2% for a typical mortgage initiation cost. The transaction cost for selling is set to 4%, since a typical commission for real estate agents is reported at $3-5\%$. Maintenance cost χ is set to 4.5%, based on calculations of Svensson (2023). Approximately 1.5% form this figure should be interpreted as depreciation, while the rest represents operational and maintenance costs.

[3] An alternative, and perhaps conceptually preferable approach would be taking into account cohort effects both for earnings and wealth patterns. Unfortunately, due to the shortness of the panel of wealth variables, cohort effects for the latter are difficult to disentangle from age effects without making strong assumptions.

The risky investment return parameters are computed from monthly data on the Swedish stock market index between 1983-2016, taking the 30% capital gains tax into account. This results in a mean log return of 0.0646, with a standard deviation of 0.12. Interest rates take their average value between October 2005 and the December 2007, which is the intersection of the time periods when these numbers are reported by SCB and when the wealth data is available. The risk free rate is set to 1.3%, the average interest rate on bank accounts provided to households, which is the primary vehicle of risk-free investments in Sweden. The average mortgage rate equals 4%, while the interest on consumption loan is 7.5%.

The maximal loan-to-value ratio for mortgages is set to 0.85, a typical value in the literature. Note that currently this maximum is enforced by law in Sweden, but it was not yet the case in the years, from which detailed wealth data is available. The payment-to-income ratio is set to $\eta_m = 0.18$ for mortgages, which would be implied by the chosen mortgage rate and a recent regulation limiting the debt to income ratio at 4.5 in Sweden. Finally, $\eta_c = 0.2$ and the maximal level of consumption loan $\overline{L}$ is set twice the value of average income in the model. The latter two values are based on Finansinspektionen (2019), a report on consumption loans by the Financial Supervisory Authority.

6.2 Estimated parameters

The rest of the parameters (β, κ, θ, γ, ψ, ζ, F and τ) are estimated by simulated method of moments in a way that certain patterns generated by the model match their empirical counterparts. The reason τ is included among the parameters to estimate together with preference parameters is that without having a free parameter influencing directly the decision between renting and owning, home-ownership patterns are very difficult to match. Since the mechanisms discussed in this paper work through housing decisions, representing the latter in a realistic way in the model crucial to claim that findings of this paper are relevant to understand the effects of home-ownership in real life economies. In this section I am first describing the estimation targets, the method of estimation and the match provided by the model.

As the model's performance is judged based on how well it can match portfolio allocation patterns over the wealth distribution, it is reasonable to calibrate the above parameters such that the shares of different asset classes in the aggregate is close to the data. In particular, having a life-cycle model at hand, I am matching the paths of average net wealth, risky assets, debt, housing value, stock market participation and home-ownership ratio over age. Note that in order to avoid hard-wiring portfolio choice policies already at the calibration stage, I am targeting the average value of wealth components for each age instead of targeting average portfolio shares, which would give less freedom for the model to generate different portfolio choice patterns. In addition, to make comparisons of patterns over the wealth distribution meaningful, it is crucial that wealth distribution induced by the model is also a close match to the empirical one. Therefore I am also targeting the shares of wealth held by households in selected quantiles of the wealth distribution.

Targeted moments in the data are computed as follows: For all empirical patterns only data from year 2000 is used, since during the subsequent years an exceptional housing price boom took place and due to the way how home value data is collected, filtering out year effects without heterogeneous distortions over the wealth distribution would be very difficult. The unit of analysis is the household, which is represented by the total income and wealth of its members divided by the number of adults in the household.[4] The age of the household is defined as the age of the household's oldest member. Risky liquid assets include stocks, funds (except money market funds), commercial real estates, derivatives and private equity. Housing includes residential real estate and vacation homes. Home-owners are households with positive housing, while participants are those with positive risky liquid asset holdings. Cash (or liquid safe assets) includes cash, bank deposits and money market funds. The debt variable includes both mortgages and unbacked credits. Using these raw values of cash and debt would be problematic to compare with the model however, since while on one hand a significant share of the population holds both cash and debt in the data, this can never happen in the model. For the sake of more fair comparisons, I use netted out versions of these two quantities when computing averages for each age group as follows:

$$cash_{net} = \max\{0, cash - debt\}$$
$$debt_{net} = \max\{0, debt - cash\}$$

Note that from now on empirical cash and debt values will always refer to their netted out versions. To make the results comparable with the model, all obtained averages over age are normalized by the average total income in the economy.

Measures of the wealth distribution are all computed based on net wealth, defined as the sum of risky and safe liquid assets and housing, minus debt. The quantiles considered for matching the wealth distributions are the following ones: 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 0.95, 0.975, 0.99, 0.995. This comparison is performed as follows: First, I compute the net wealth level corresponding to the above quantiles and normalize it with average income. Then I compute how big share of the population in the model is in the wealth groups defined by the same cutoffs. Then I take differences of the resulting population shares from those of the data.

The estimation is performed as follows. For a given guess for the parameter vector, the policy functions and the wealth distribution are solved for. Next the averages of net wealth, housing wealth, risky wealth, cash, debt, participation rate

[4]Statistics Sweden classifies two adults as one household if (1) they are married or (2) they have common children and are registered at the same address. According to this definition, never more than two adults are present in one household.

and home ownership rate are computed for ages 25-100 years in the model and I take their differences with the empirical values. Monetary values all rescaled by average income income. Then I add the difference in wealth shares as described above. In total, this process results in $76 \times 7 + 14 = 544$ moments, which are organized in a column vector z. Parameters are chosen to minimize

$$z'Wz$$

where W is a diagonal weight matrix. Weights are adjusted such that the importance of the wealth distribution equals the other patterns in spite of consisting of fewer numbers. In addition, life cycle patterns are weighted such that weights correspond to the age distribution of the population. First global optimization is performed by evaluating the objective function at 5000 Sobol points in the parameter space. Then the solution is found by performing local optimization using the Nelder-Mead method started from the best point.

The values of the resulting parameters are summarized in Table 1. The estimated value for $\tau = 7.1\%$ is in the range of realistic values: the most important component of rental costs is rental rate itself, for which an upper bound is 5.7%, the ratio of rents and prices of newly built rental dwellings in 2000. Note that in Sweden the majority of rental apartments are affected by rent control, so most households face a lower rental rate than this value. In addition, renters also have to pay operational costs, approximately increasing the rental cost by 2.5 percentage points.

Next, the matched patterns are reported in Figure 2. In general, most life-cycle patterns are matched reasonably precisely until retirement. Moreover, the paths of net wealth, housing wealth, risky assets, debt, home ownership and participation all replicate the hump-shaped patterns over the whole life cycle found in the data. However, the speed of dissaving generated by the model is significantly faster than its empirical counterpart. In addition, the little saving that is predicted by the model for the old is chosen in a far too risk seeking manner. In particular savings of the old are allocated almost exclusively in risky assets, while housing and cash are avoided, in a stark contrast with data. These findings might be a consequence of not modeling health risks for the elderly as De Nardi (2004). Of course, in general a worse fit for the policies of the elderly is to be expected, since during estimation the corresponding moments got a smaller weight when computing the objective function. A further deviation from the data is the insignificant amount of cash holdings.

Preference parameters			
β	time preference rate	0.938	estimated
κ	Bequest strength	0.932	estimated
θ	Bequest share to offspring	0.473	estimated
γ	risk aversion	8.81	estimated
ψ	inverse EIS	0.761	estimated
ω	housing share	0.276	SCB
ζ	consumption insurance	0.045%*	estimated
Returns and participation cost			
R^f	deposit rate	1.013	SCB
μ_M	expected log stock market return	0.0646	SIXRX
σ_M	s.d. of log stock market return	0.14	SIXRX
R^m	interest rate - mortgage	1.04	SCB
R^c	interest rate - consumption loan	1.075	SCB
F	fixed participation cost	1.8%*	estimated
Income			
g	drift of aggregate wage growth	0.0213	data
ρ	auto-correlation of persistent component	0.924	data
σ_ε	s.d. of shocks to persistent income	0.171	data
σ_ν	s.d. of shocks to transitory income	0.356	data
$\sigma_{\nu pen}$	transitory pension	0.094	data
Housing			
ρ_h	autocorrelation of housing prices	0.9334	data
σ_h	s.d. of housing price shocks	0.0836	data
$\min_h$	minimal housing size	1*	preset
Φ	buying costs	1.035	preset
α	selling costs	0.96	preset
τ	rental costs to price ratio	0.071	estimated
η_m	PTI mortgage	0.3	preset
η_c	PTI consumption loan	0.2	FI
$\overline{L}$	maximal consumption loan	2*	FI
δ	mortgage max LTV	0.85	preset
χ	maintenance cost	0.04	Svensson (2023)

Table 1: Calibrated values for model parameters. Quantities marked with an asterisk * are expressed relative to average yearly income.

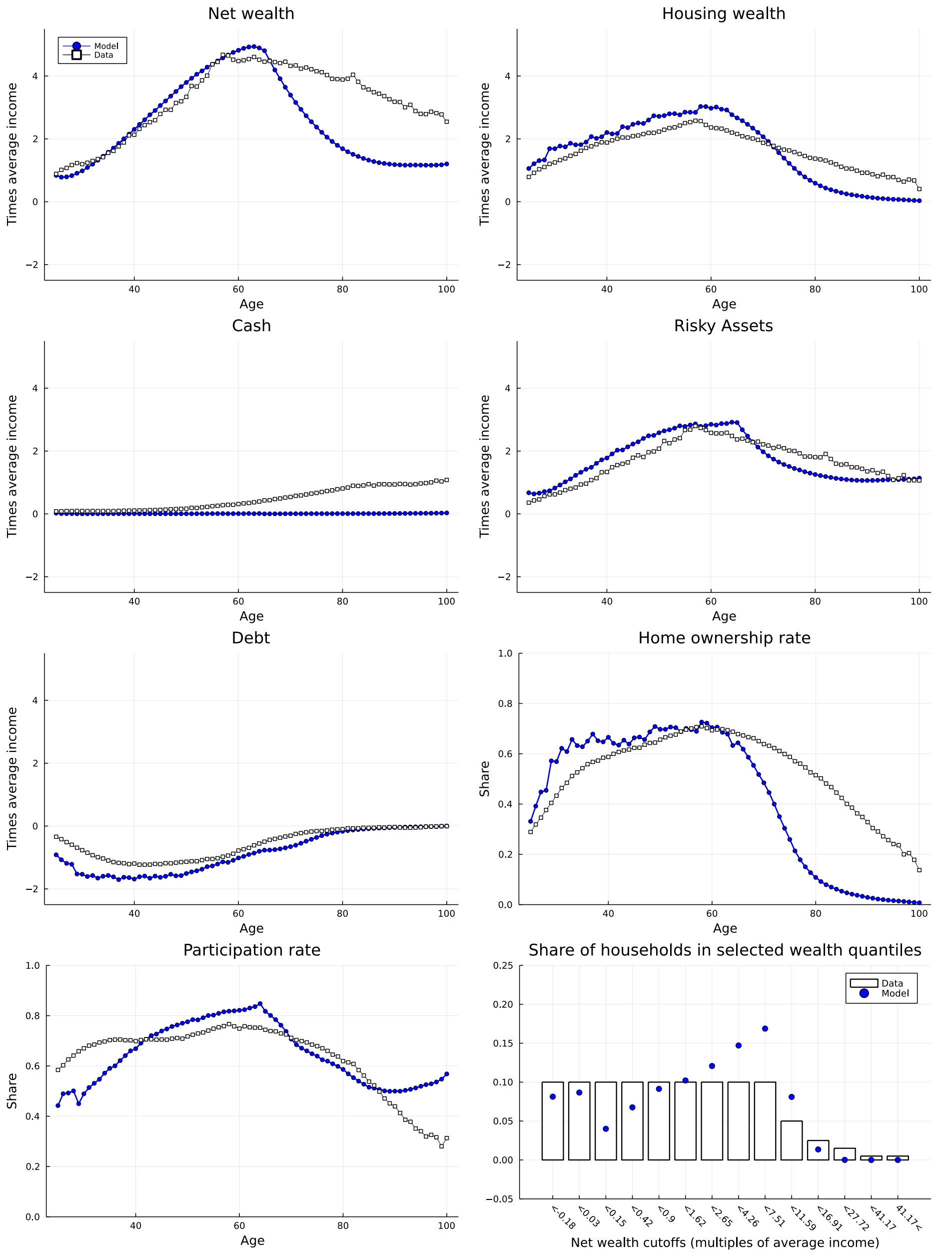

Figure 2: Targeted moments in the model and in the data

This is partly driven by the slight overshooting of debt - since debt and cash are never held at once - and also the fact that over a significant range of wealth levels it is optimal to hold no debt and no cash, before hoarding cash becomes optimal (see section 7.3 for a discussion). Finally, note that even though the wealth distribution is relatively well matched, the model predicts very little mass of households owning more wealth then 16.91 times the average yearly income, which corresponds to the richest 2.5% in the data. While this is not surprising given that this model does not contain any common ingredients of models generating realistic wealth inequality on the right tail (such as preference heterogeneity, skewed income shocks, entrepreneurship, or idiosyncratic returns), this warns us that the current model might be most useful to interpret portfolio choice patterns over the wealth distribution after discarding the super rich.

7 Results

7.1 Policy Functions

In the following, portfolio policies are presented over the wealth distribution, where the share of each portfolio component is expressed as the ratio of the corresponding component with gross wealth, defined as the sum of safe and risky liquid wealth and housing wealth. To compare policies over the wealth distribution in the data versus the model, the following procedure is used: After creating wealth bins in the data using the same percentiles as in section 6, I normalize the thresholds by the average yearly income of the economy. Then I build the corresponding wealth bins in the simulated wealth distribution using the same multiples of average income in the model and compute average policies over the resulting wealth quantiles. Therefore when comparing policies in the model and data, we know that the compared households are similarly wealthy, but they do not necessarily sit at the same point at their respective wealth distribution. This procedure has the advantage that comparisons are not distorted by differences in the exact shape of the wealth distributions in the model and the data. Therefore the following figures present portfolio choice patterns generated by the model and their empirical counterparts over the empirical wealth distribution.

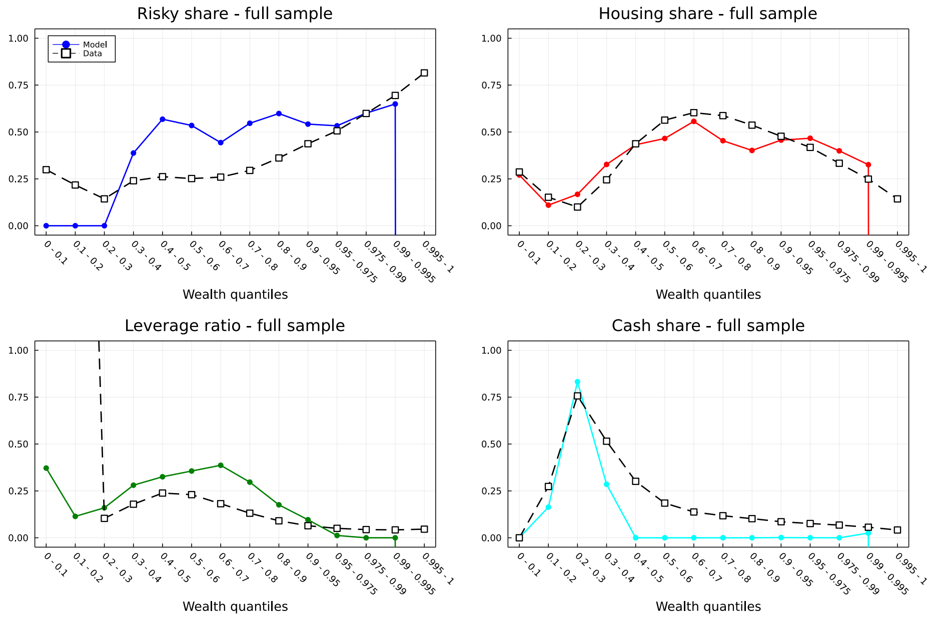

Figure 3: Average portfolio allocation patterns of households over the empirical wealth distribution. Model versus data. Note that households with zero gross wealth are assigned share zero for all components.

It turns out that most portfolio choice patterns are matched surprisingly well. One apparent shortcoming of the model fit is that it significantly over-predicts risk-taking for the middle of the wealth distribution. This generates a too high risky share and leverage ratio, while a practically zero cash ratio. However, for the richest 10% the model fit again becomes better, and in particular, the risky share is not counterfactually decreasing for the more wealthy part of the wealth distribution, which is a common feature of models without non-homotheticities. Note that the last data point in the model results is missing, since the model cannot generate agents with wealth corresponding to the richest half percent of the empirical wealth distribution. To understand what drives the deviations from empirical patterns, it is useful to consider home-owners and renters separately. Before that however, it is important to check that home-ownership varies over the wealth distribution similarly in the model and data. As seen in Figure 4, the model correctly replicates the increasing share of home owners along the wealth distribution.

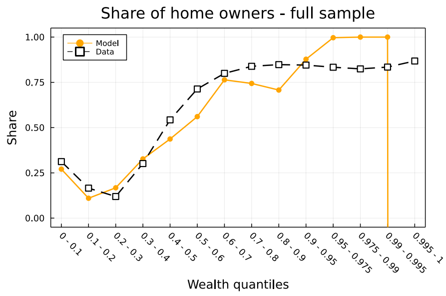

Figure 4: Home-ownership ratio over the empirical wealth distribution. Model versus data. Note that households with zero gross wealth are assigned share zero for all components.

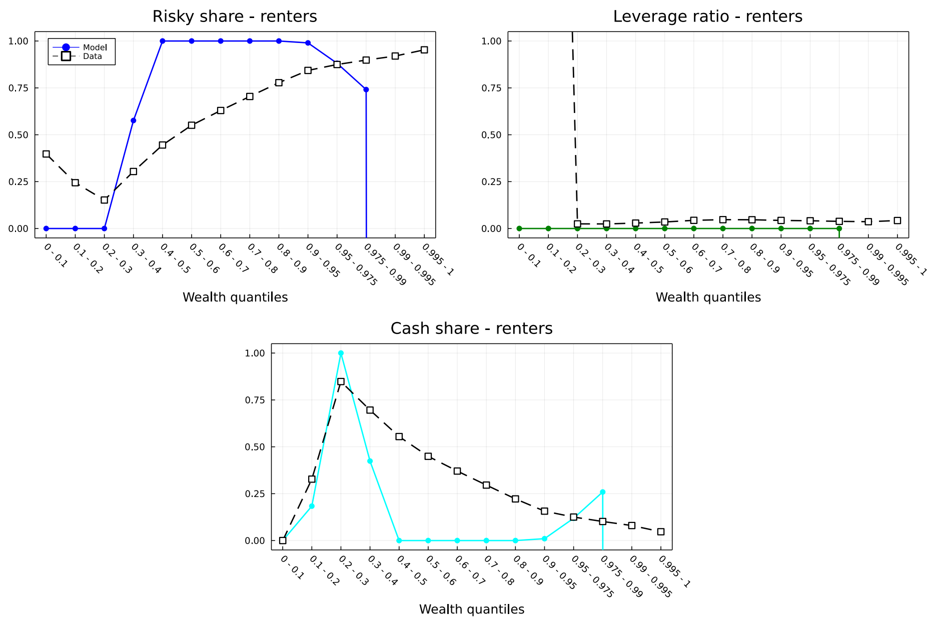

Figure 5: Average portfolio allocation patterns of renters over the empirical wealth distribution. Model versus data. Note that households with zero gross wealth are assigned share zero for all components.

Note that both the data and the model features a significant share of heavily indebted homeowners: More people in the bottom decile own housing than any of the following two deciles. After understanding how households with different home-ownership status are spread out over the stationary distribution, next we examine portfolio choice decision of renters.

The obtained patterns for renters accurately represent the typical difficulties models face when trying to match portfolio choice over the wealth distribution. Poor households choose not to participate in the stock market due to high fixed costs. On the other hand, once they are past the wealth levels where the participation cost has a sufficient deterring effect, their risky share immediately jumps to 100%. This is due to the fact that their human capital is still large relative to financial wealth. Since in this model labor income is not correlated with the stock market, human capital is a close substitute of cash when thinking about optimal portfolio choice, driving up the optimal share of risky assets out of gross wealth. For the richest renters however, the weight of human capital is getting smaller, leading to decreasing risky share and increasing cash share on top of the wealth distribution, which is counterfactual. The inability of the model to match portfolio choice of renters accurately is not very surprising as all mechanisms in the model that might matter for renters have a bite only for the poorest segment, namely precautionary motives amplified by housing price risk and the fixed participation cost. Let us consider the corresponding figures for home-owners next.

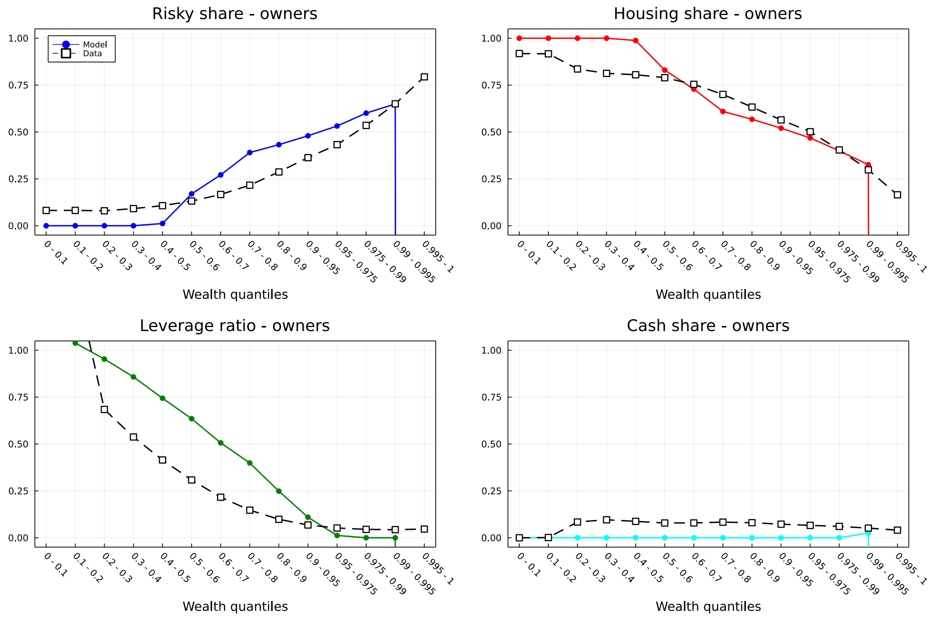

Figure 6: Average portfolio allocation patterns of home-owners over the empirical wealth distribution. Model versus data

In contrast to renters, the model performs very well to reproduce portfolio choice patterns for home-owners, especially considering these moments were not targeted during the calibration. In particular, both housing and risky share patterns are precisely matched by the model. Debt holdings are significantly overshot in the model, even though the decreasing pattern and flattening out around 0 for the richest segment is well captured. The share of safe liquid assets is low and flat in data, while the same quantity is practically flat 0 in the model. In the following subsections the mechanisms behind these patterns are examined in detail.

7.2 Housing choice

It is clear from Figure 4 that to understand how the model generates an increasing risky share for home-owners, it is necessary to find out what drives the decreasing optimal housing share over the wealth distribution. The toy model in Section 2 offered a simple explanation: since housing decisions are made to smooth net worth - including human capital - over one's lifetime, while saving decisions are to smooth only wealth over time, households with relatively more human capital should optimally choose higher housing ratios. So far we have seen that the optimal housing share monotonically decreases over wealth, so it is of interest to see, if the same is true for the share of human capital out of total net worth.[5] As shown by Figure 7, this is in indeed the case, suggesting that the mechanism uncovered in the toy model might be relevant in the full life-cycle model as well.

However, we can perform a somewhat more direct test to demonstrate the connection between human capital and optimal housing consumption. As shown in the toy model, there exists a hyperbolic relationship between the optimal share of housing in net wealth and the share of financial wealth in total net worth. Both these quantities can be computed in the model for all states (determined by age, persistent income state, housing level and liquid wealth) after which the resulting pairs of values can be plotted as shown on Figure 8.

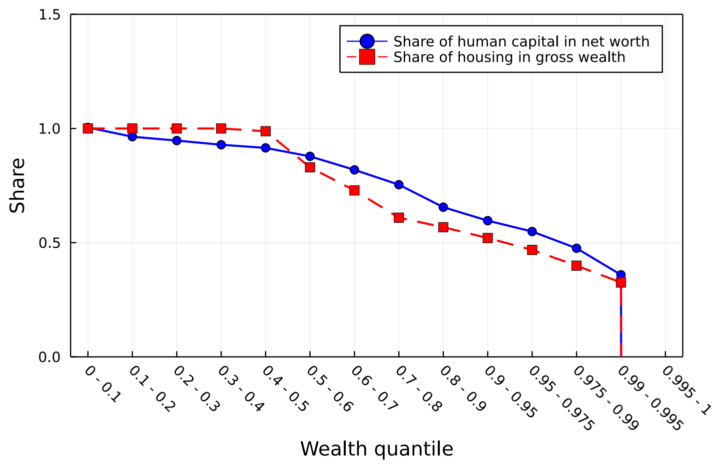

Figure 7: Shares of human capital out of net worth and share of housing in gross wealth over the wealth distribution.

[5]Human capital in the model is computed as the discounted expected value of one's lifetime income, taking survival probabilities into account, discounted by R^m.

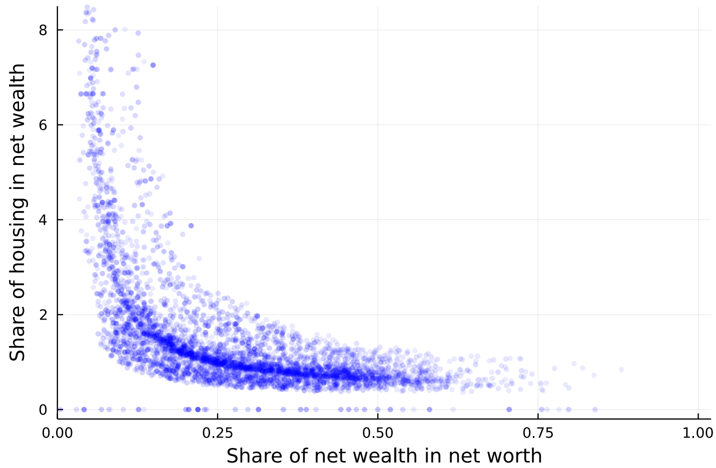

Figure 8: The relationship of the share of housing wealth in net wealth and the share of wealth in net worth in the model. Each dot represents one state, and the color intensity is a monotonic function of the mass of households in the given state in the stationary distribution. For computational reasons only a random sample of all states is shown on this graph.

Figure 8 clearly depicts a hyperbolic pattern for home-owners, conforming the findings in the analytically tractable model of Section 2. Note that the horizontal line at level 0 corresponds to renters. We can therefore conclude that households with less wealth relative to their human capital will optimally choose a high housing share, which mechanically crowds out risky assets from these households' portfolios. The size of this effect will be discussed in Section 7.4, after a discussion on how the interest rate wedge shapes optimal risk taking.

7.3 Asset choice and the interest rate wedge

After understanding what induces the crowding out effect of housing on risky investment for poor households, we focus on an orthogonal effect. Even considering housing choice fixed, interesting patterns in the risky share are generated by the differences among interest rates of different risk-free investment vehicles. For most home-owner states, due to the house transaction costs, for a fairly large interval of wealth it is optimal to stay in the current housing state. Outside of this interval, on the lower end the household optimally sells its house and buys a smaller one or becomes a renter. On the other extreme, it switches to a larger home often by taking on more debt. To understand optimal policies in the region with no change in housing, it is most useful to examine the optimal portfolio choice sub-problem of the household's full optimization problem, taking given current housing and expenditure. For the rest of this subsection, assume that $\gamma < 1$ for expositional clarity.[6] Fixing the current state, with some abuse of notation define

$$\overline{V}_{j+1}(\hat{s}') = \beta\mathbb{E}\Big[q_{j+1}V_{j+1}(G', P_H', \hat{s}' + \exp(y'), z', H)^{1-\gamma} \\ + (1 - q_{j+1})B(G', P_H', \hat{s}' + \alpha P_H' H)^{1-\gamma} \mid \hat{s}'\Big] \tag{16}$$

where time and household indices are omitted and values with prime denote quantities determined in the next period. Then, taking current savings s as given, the portfolio allocation sub-problem can be formulated as

$$\max_{B,\xi,M,L} \mathbb{E}\big[\overline{V}_{j+1}(\hat{s}')\big] \tag{17}$$

subject to (6)-(9) and (12). Denote the Lagrange multiplier to the last constraint with λ. Furthermore, for a generic component of saving X, let λ_X be the Lagrangian multiplier corresponding to the lower bound for X, while let λ^X denote the same for the upper bound, whenever applicable. Then in the optimum the following first order conditions will hold:

$$\mathbb{E}\big[R'\overline{V}'_{j+1}(\hat{s}')\big] = \lambda - \lambda_\xi \tag{18}$$

$$R^f\mathbb{E}\big[\overline{V}'_{j+1}(\hat{s}')\big] = \lambda - \lambda_B \tag{19}$$

$$R^m\mathbb{E}\big[\overline{V}'_{j+1}(\hat{s}')\big] = \lambda + \lambda_M - \lambda^M \tag{20}$$

$$R^l\mathbb{E}\big[\overline{V}'_{j+1}(\hat{s}')\big] = \lambda + \lambda_L - \lambda^L \tag{21}$$

Note that all Lagrange-multipliers are non-negative and the corresponding complementary slackness conditions hold. Before starting to interpret the policy functions in a concrete case to gain intuition of the optimal solutions, it is worth making two points. First, as discussed in the model description, in optimum only one of the risk-free saving vehicles are in the interior of their respective feasible regions. Indeed, as the left hand side of equations (19)-(21) all differ, at most one of their right hand sides can equal λ. Second, if we have an interior solution both for a generic risk-free asset X and ξ, then a first order condition of the form

$$\mathbb{E}\big[R'\overline{V}'_{j+1}(\hat{s}')\big] = R^X\mathbb{E}\big[\overline{V}'_{j+1}(\hat{s}')\big] \tag{22}$$

holds. Let us assume that the conditional expectation in (16) smoothed out any potential non-concavities and hence $\overline{V}$ is strictly concave. [7] In this case it easy to

[6] The other case with $\gamma > 1$ can be dealt with identically apart from a few sign changes.

[7] This is the case if the shocks integrated out by the conditional expectation operator represented a large enough uncertainty.

show that the optimal ξ is a increasing function of the risk premium $\mathbb{E}\big[R'\big] - R^X$. After making these observations we are ready to examine the intuition behind optimal portfolio choice policies in a given state, shown in Figure 9. Note that the considered state is not completely representative, and is chosen to give a full overview of the possible configurations of binding asset pricing equations. In particular, regions A and E are typically missing for younger or poorer households, since they have higher affinity to move into a home of the most efficient size, instead of sticking to what they have.

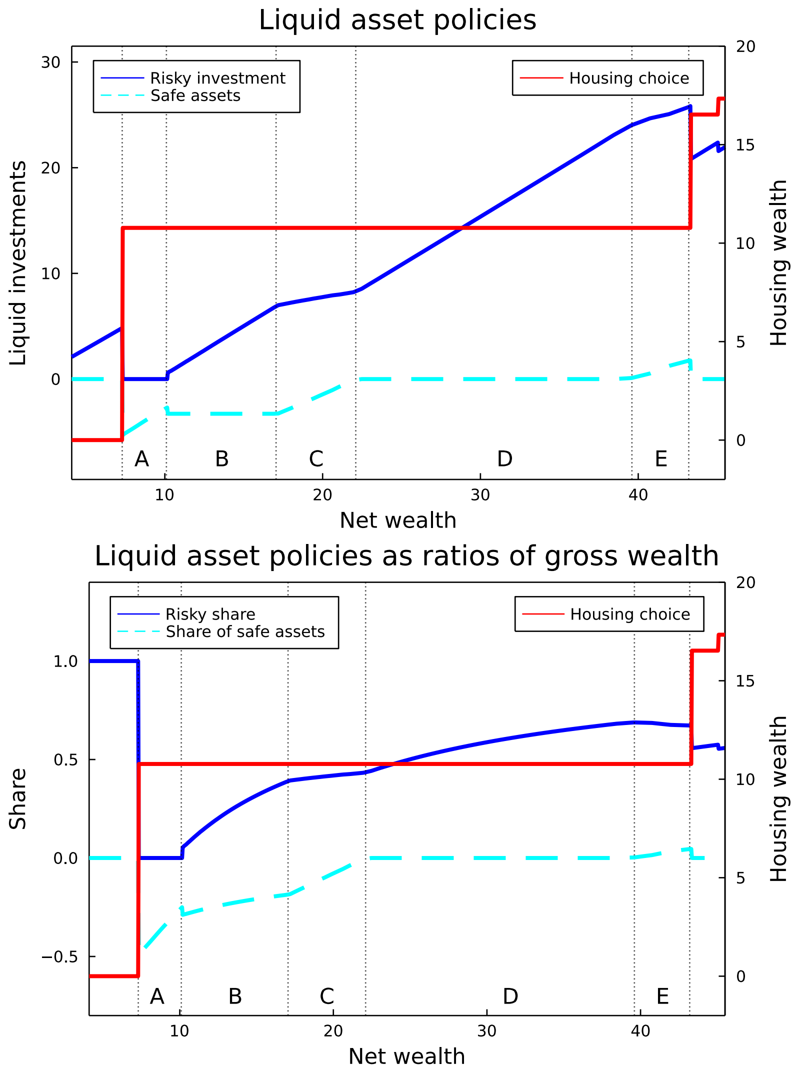

Figure 9: Optimal risky and risk-free savings and housing choices as a function of wealth for 62 years old households with house of value 10.77 times average income, and in the average persistent income state. The top figure displays saving decisions directly, while the bottom panel shows their value as a ratio of gross wealth. The inaction region with respect to housing choice is divided into regions A to E, discussed in the text.

In region A the household does not participate in the stock market, but holds only debt as liquid asset. If the household decided to save more but still remained in A, then the only adjustment to optimal policies is decreasing her debt. Next, in zone B the household starts holding stocks, which is financed by completely exhausting the borrowing capacity corresponding to mortgage. The reason for this choice being optimal is twofold: First, due to the high interest rate on consumption loan, holding stocks is never optimal while exceeding the borrowing limit for mortgage. Second, the optimal risky investment level implied by ignoring the constraints and solving (22) with $X = m$ is too high to be financed without relying on consumption loans. Hence in optimum, until the latter statement seizes to be true, the household keeps her mortgage level at the corresponding borrowing limit and responds to changes in saving one-to-one with increasing risky investments.

Next, in region C optimal portfolio choice is determined by (22) with $X = m$. As a result, both risky and safe investment increase over wealth here, in a roughly linear manner. In fact, the ratio of these two slopes corresponds to the optimal risky share in standard portfolio choice models such as Merton (1969), as holding one unit less debt is analogous to one unit more of bonds. The subsequent region D is characterized by no holding of any risky-free assets. This happens so, as at these wealth levels, the unconstrained optimum assuming R^m as alternative risk-free rate would imply negative mortgage holdings, while the corresponding optimum with R^f would imply negative bond holdings, both of which are ruled out in our setup. Therefore the optimal risky investment is in between these two benchmarks, also illustrated by

$$R^m \mathbb{E}\big[\overline{V}'_{j+1}(\hat{s}')\big] - \lambda_M = \mathbb{E}\big[R' \overline{V}'_{j+1}(\hat{s}')\big] = R^f \mathbb{E}\big[\overline{V}'_{j+1}(\hat{s}')\big] + \lambda_f$$

which holds here. Finally, zone E is the analogue of C, but with the bond serving as the locally unique relevant risk-free alternative of risky assets. It is worth noting, that optimal policies in region E appear less linear than those in other zones. Specifically, on the right edge of the inaction region optimal risky investment is increasing faster in wealth than in the rest of region E. This is due to the mechanism explored in the second chapter of this thesis, implying higher risk taking close to a wealth boundary representing a saving target.[8]

It is easy to appreciate the effects of wedges between interest rates on portfolio choice patterns through the following thought experiment: By allowing negative bond holdings in the current period, but keeping housing decisions as fixed, optimal policies could be obtained by prolonging the optimal policies of region E over all regions A to E. In that case both the amount of risky investments and debt would be higher, which has two important implications. First, risky share would be higher

[8]The discrete time nature of this model makes the size of this effect small relative to an otherwise equivalent model in continuous time. This is because in a discrete time setup being close to a switching boundary in the current period does not necessarily imply being close to it again next period with a high enough probability.

already for lower wealth levels and second, its growth over over wealth would be smaller. A key how wedges between interest rates cause an increasing risky share is the presence of the risk-free asset transition regions B and D, where the risky share always monotonically grows. Note that this finding is not specific to the chosen measure of the risky share, namely the ratio of risky investments over gross wealth. As in these regions the level of risk-free liquid assets is constant, while risky investment increases in wealth, the same conclusion would hold for any natural alternative measure of risk taking in saving.

Apart from exploring the connections of interest rate wedges and optimal portfolio allocation, this section clarifies another feature of the model, namely that in spite of the high estimated risk-aversion coefficient, very little cash is held in this economy. Indeed, as debt is adjustable without paying any transactional costs, holding less debt than its maximal level is conceptually equivalent to holding cash in a model with no borrowing. This would imply that the formal equivalent of risky share in a model with debt is the ratio of risky assets and the sum of risky assets and the excess amount of risk-less assets relative to the borrowing capacity (unfortunately and unobservable quantity). In fact, this is a possible interpretation of the optimal investment rule in Section 7 of Merton (1971) where the agent is allowed to borrow up to the natural debt constraint implied by her income. To sum up, in this model risk aversion is not identified simply by the relation between stock and cash holdings, but it is also relevant how much less debt agents hold relative to their borrowing limits.

7.4 Counterfactuals

In the previous two subsections two mechanisms were discussed, both capable of contributing to an increasing risky share over wealth. To explore to what extent these channels are able to generate empirically plausible patterns over the wealth distribution, I construct two corresponding counterfactuals with the aim of turning off these two channels one-by-one.

In the first counterfactual the possibility of endogenous housing choice is turned off: all agents are forced to choose a house which brings them closest to the aggregate housing wealth to net wealth ratio in the benchmark model, among the alternatives not bringing them to immediate bankruptcy. After this hard-wired housing decision, the households choose expenditure and portfolio choice optimally. Note that it is not the ratio of housing wealth to gross wealth, but to net wealth which is considered here. This choice provides a more natural comparison, since the mechanism discussed in Section 2 implies a pattern on the share of housing out of net wealth.

The second counterfactual concerns the channel through interest rate wedges. To consider the least extreme alternative, I turn off this effect by resolving the model assuming that households wrongly think that $R^f = R^m = R^c = 1.04$ - just as mortgage rate in the original setup - in the stage when they decide about the risky share. Note that for the sake of comparison I make sure that apart from lowering the utility from saving due to suboptimal portfolio choice, the savings and housing decisions are not affected directly. In practice, under this scenario when a household gets rich enough not to hold debt in optimum, they face a higher risk-free rate when allocating their savings across different assets than in the benchmark specification. On the other hand, when crossing under the borrowing limit corresponding to mortgage, the household does not face a higher interest rate, increasing debt taking.

As I would like to isolate the direct effect of these modified policies from effects through generating a different wealth distribution, I perform the following exercise: I take the stationary distribution of the benchmark model and assume that suddenly all households find themselves in the counterfactual in question and choose their current policies accordingly. Therefore average policies are compared across the benchmark specification and the two counterfactuals over an identical distribution of states. As both channels affect home-owners only, for clarity, only results for home owners are reported in Figure 10.

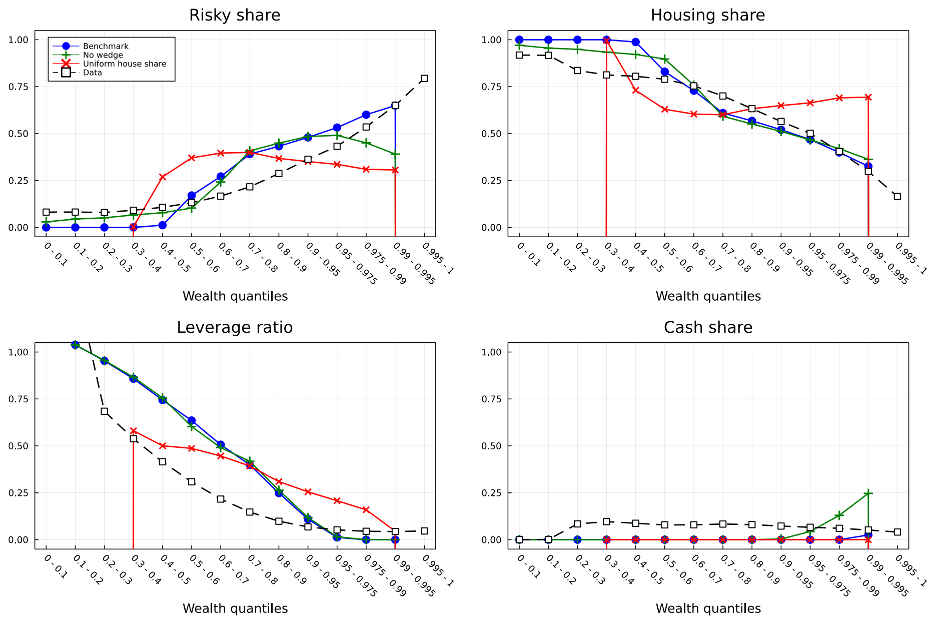

Figure 10: Average portfolio allocation patterns of home-owners over the empirical wealth distribution. Counterfactuals versus benchmark and data

Forcing a uniform choice of housing share results in a risky share pattern not unlike typical portfolio choice models predict. After an initial increase over the bottom half of the wealth distribution, risky share slowly decreases in wealth, as human capital gradually becomes less important. The increasing segment for the poor is partly caused by growing participation, and partly driven by agents who cannot afford switching to the house level prescribed by the counterfactual due to the presence of a minimal housing size. Therefore by breaking the connection between wealth and housing decisions, the model loses its ability to generate an increasing risky share for the rich. Furthermore, forcing the rich to choose larger houses makes them much more leveraged compared to the data and the benchmark model. By construction, the decreasing pattern in housing share disappears as well. Note however, that it is still not exactly constant due to the difference between gross and net wealth.

Turning off the interest rate wedge between the mortgage rate and the two other risk-free rates also destroys the monotonic increasing nature of optimal risky share over wealth. Especially the disappearance of the difference between borrowing and lending rates is a problem: Under this counterfactual households without debt will start hoarding cash, which depresses the risky share for the richest. This highlights the importance of a decreasing risk-free rate schedule in effectively making risk-premium to be an increasing function of wealth and hence generating the increasing risky share all along the wealth distribution. It is worth noting that the calibration overshoots the amount of debt in the benchmark economy. Hence, with a more precise fit, where more households are affected by the interest rate of the bond, the effect of this counterfactual would be presumably even larger.

From the above counterfactuals it becomes apparent that the baseline model can generate an increasing risky share for the wealthy home-owners due to the presence of housing and debt decisions and their interactions.

7.5 Marginal propensity of risky investments

To estimate the effect of almost any economic policy measure, a crucial question is how an average household would spend an unexpected additional dollar. In particular, an object of central importance in macroeconomic models is the marginal propensity of consumption, due to the general equilibrium effects caused by the aggregate consumption response. Another interesting question is what happens to the rest of the money. Will it be invested in the stock market, is it spent on paying back debt or it stays on a bank account untouched? Since the model considered in this paper is able to generate realistic portfolio compositions over an empirically plausible wealth distribution - abstracted away from the very top - it is of interest to investigate what drives the marginal propensities of different ways to spend an additional dollar over the wealth distribution. Defining the corresponding marginal propensities as derivatives of non-durable consumption, housing expenditure, risky investment and safe asset holdings with respect to cash-on-hand, we obtain the decomposition reported in Figure 7.5.

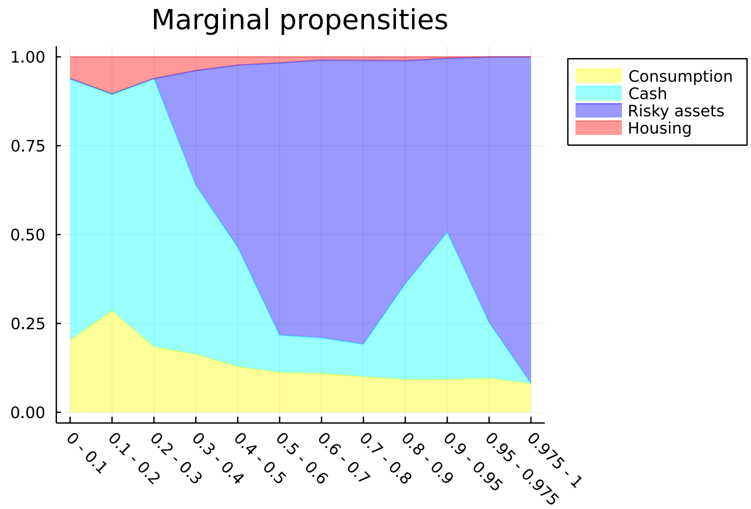

Figure 11: Marginal propensities of different means of saving and expenditure.

On surprising finding is the unusually large magnitude of marginal propensity of consumption. This is a consequence of the estimated preference parameters: While the high risk aversion and low discount factor are not unusual in the household finance literature, they are far from typical in the macroeconomic literature on MPCs.[9] Focusing on the part of money not spent immediately, we can conclude that in the bottom deciles all stimulus not consumed would be saved in the risk-less asset, or more precisely it would be used to pay back debt. On the other hand, the richest 2.5% would invest all their left over money into risky investments. The transition between these two extremes is not exactly monotonic however. To understand the forces shaping marginal propensities of different means of saving, we need to return to the classification of net wealth regions based on the optimal qualitative nature of saving policies, introduced in Figure 9. By grouping all households into categories A to E and plotting their share over wealth quantiles, Figure 7.5 is obtained, providing an explanation for the patterns of marginal propensities.

[9]By examining the equivalent graphs separately for owners and renters, it is apparent that while the illiquidity of housing plays a role generating higher MPCs for the rich, its maximal size is even higher for renters.

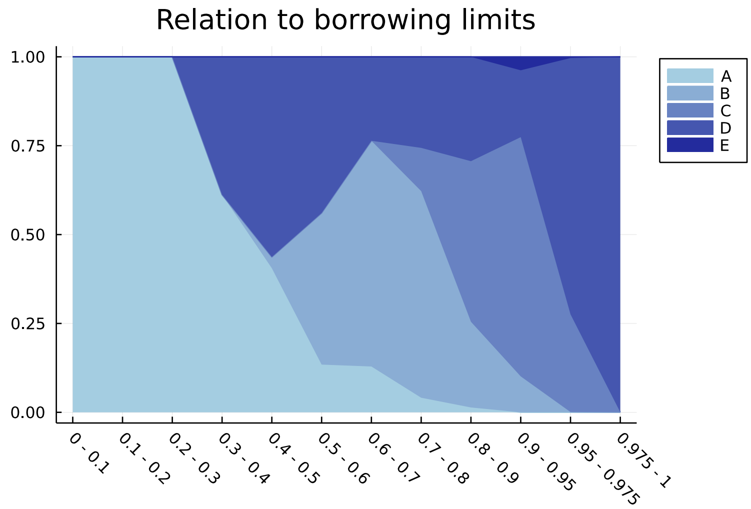

Figure 12: Distributions of risk-free asset positions relative to borrowing limit. A - non-participants; B - participants, maximal mortgage; C - participants, interior solution with positive mortgage; D - participants, zero risk-free assets; E - participants, positive bond holding.

As it was discussed in Section 7.3, regions B and D are characterized by a flat risk-free asset policy function and hence households belonging to these groups invest all their additional savings into risky assets. By comparing Figures 7.5 and 7.5 it is clear that the wealth quantiles with high marginal propensity of risky assets are exactly those with a high share of group B or D households. Intuitively, the marginal propensity of risky investments is driven up by households, which do not need additional risk-free assets for dealing with financial risk, thanks to their ample unused borrowing capacity.

8 Conclusion

This paper argues that modeling housing and debt-taking decisions can go a long way in matching the portfolio allocation of homeowners. In particular, I showed that since housing choice is conducted similarly to consumption decisions, households with a higher share of human capital in their net worth will optimally choose a higher housing share relative to their wealth. This means that housing mechanically crowds out risky investments to a larger degree for less wealthy households, keeping human capital fixed. Since in the model the average human capital to net worth ratio decreases in wealth in the stationary distribution, this leads to a decreasing housing share and increasing risky share in wealth for home-owners. Indeed, when in a counterfactual households are not allowed to adjust their housing ratio, but instead are forced to set it at the cross-sectional average, then the risky share is decreasing over the top half of the wealth distribution.

Another contribution of this paper is linking the above channel with the effects of debt holding. Indeed, for wealth-poor households the ideal house size can be financed only by a relatively high reliance on debt, which affects the average risk premium they face for two reasons: First, the interest rate of debt is higher than that of bonds. Second, if several types of debts are available with different interest rates, the agent optimally chooses the more expensive alternatives for financing her house purchase only after exhausting all other possibilities. Therefore, since on average poorer households have typically a lower wealth-to-net worth ratio, they optimally choose a higher housing share, will take relatively more debt and as a consequence, will face lower risk premia relative to more wealthy households. This channel endogenously results in a wealth-dependent effective risk premium, contributing to the increasing risky share over the wealth distribution. Lastly, to illustrate the importance of the interplay between debt and portfolio allocation decisions, the marginal propensities of all uses of additional liquidity were computed. I showed that a larger share of helicopter money would be spent on risky assets by the rich, largely due to the presence of a kind of household which can exist only in a framework where there is an interest rate wedge between bonds and debt: when a sufficiently large amount of unexhausted borrowing capacity is available to provide safety in rainy days, a household might optimally allocate all additional savings into the risky asset.

Apart from explaining portfolio choice patterns over the wealth distribution in aggregate data, it would also be interesting to test several implications of this model in micro-data, which is something I am planning to work on next. The first of these concerns the relationship between the share of human capital in net worth and the share of housing in wealth. Indeed, Swedish administrative data contains detailed information on the income and demographic characteristics of each individual over a long time period, which would allow constructing suitable estimates of households' human capital. Combining this with the available data on wealth including housing, exploring how the two are related is a feasible task. Second, as discussed in Section 7.3 a key determinant of optimal portfolio allocation is the difference between the amount of risk-free assets and the borrowing limit, which acts as a substitute for cash in a framework without transaction costs related to debt adjustment decisions. I am planning to test this implication by investigating portfolio choice changes around house purchase transactions, when all involved quantities change by construction, allowing us to separate the true effects from differences among individual fixed effects.

One obvious limitation of the model studied in this paper is pooling all liquid safe vehicles of saving, such as cash and various forms of debt into one composite asset in the optimal solution. In particular, it is the lack of any costs related to debt initiation and renegotiation which makes it sub-optimal to hold a positive amount of cash and debt at the same time. This simplification enabled the introduction of debt in the model without the need to add an extra state variable. However, this approach has two significant drawbacks: First, as a large majority of households do not keep a positive amount of cash, this model is not a suitable tool to investigate patterns of the risky share out of liquid assets, which is a representative measure of risk-taking in portfolios, besides the ratio of risky assets to gross wealth used in this paper. Second, the model understates the disadvantages of debt holding [10] and as a consequence, has difficulties matching saving patterns over the life-cycle without overshooting the amount of debt held in the economy. It is a natural question of how the mechanisms presented in this paper would be affected under a more realistic representation of debt. When taking or adjusting debt involves costs, households would presumably find it optimal to hold some cash even if their borrowing capacity is large. It also seems plausible that this precautionary cash holding would be a concave function of wealth, taking the housing and debt state given, in which case the points made in this paper could still apply and possibly be extended to the risky share out of liquid assets as well. Whether or not the housing and debt channel could still produce empirically plausible risky share patterns over the life-cycle and the wealth distribution is a quantitative question, which I am planning to address in future research.

[10] Boar et al. (2021) investigates how debt initiation costs can significantly reduce the effective liquidity of households, since taking debt is avoided when facing smaller adverse shocks.

References

John Ameriks, Andrew Caplin, Steven Laufer, and Stijn Van Nieuwerburgh. The joy of giving or assisted living? using strategic surveys to separate public care aversion from bequest motives. *The Journal of Finance*, 66(2):519–561, 2011. doi: https://doi.org/10.1111/j.1540-6261.2010.01641.x. URL `https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6261.2010.01641.x`.

Laurent Bach, Laurent E. Calvet, and Paolo Sodini. Rich pickings? risk, return, and skill in household wealth. *American Economic Review*, 110(9):2703–47, September 2020. doi: 10.1257/aer.20170666. URL `https://www.aeaweb.org/articles?id=10.1257/aer.20170666`.

Corina Boar, Denis Gorea, and Virgiliu Midrigan. Liquidity Constraints in the U.S. Housing Market. *The Review of Economic Studies*, 89(3):1120–1154, 09 2021. ISSN 0034-6527. doi: 10.1093/restud/rdab063. URL `https://doi.org/10.1093/restud/rdab063`.

Christopher D. Carroll. Portfolios of the rich. NBER Working Papers 7826, National Bureau of Economic Research, Inc, Aug 2000. URL `https://ideas.repec.org/p/nbr/nberwo/7826.html`.

Christopher D. Carroll. The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91(3):312–320, 2006. ISSN 0165-1765. doi: https://doi.org/10.1016/j.econlet.2005.09.013. URL `https://www.sciencedirect.com/science/article/pii/S0165176505003368`.

Sylvain Catherine. Countercyclical Labor Income Risk and Portfolio Choices over the Life Cycle. *The Review of Financial Studies*, 12 2021. ISSN 0893-9454. doi: 10.1093/rfs/hhab136. URL `https://doi.org/10.1093/rfs/hhab136`.

Yongsung Chang, Jay H. Hong, and Marios Karabarbounis. Labor market uncertainty and portfolio choice puzzles. *American Economic Journal: Macroeconomics*, 10 (2):222–62, April 2018. doi: 10.1257/mac.20160207. URL `https://www.aeaweb.org/articles?id=10.1257/mac.20160207`.

Raj Chetty, László Sándor, and Adam Szeidl. The effect of housing on portfolio choice. *The Journal of Finance*, 72(3):1171–1212, 2017. doi: https://doi.org/10.1111/jofi.12500. URL `https://onlinelibrary.wiley.com/doi/abs/10.1111/jofi.12500`.

Riccardo A. Cioffi. Heterogeneous risk exposure and the dynamics of wealth inequality. 2021. URL `https://www.rcioffi.com/files/jmp/cioffi_jmp2021_princeton.pdf`.

João F. Cocco. Portfolio choice in the presence of housing. *The Review of Financial Studies*, 18(2):535–567, 2005. ISSN 08939454, 14657368. URL `http://www.jstor.org/stable/3598045`.

Steven J. Davis, Felix Kubler, and Paul Willen. Borrowing costs and the demand for equity over the life cycle. *The Review of Economics and Statistics*, 88(2): 348–362, 2006. ISSN 00346535, 15309142. URL `http://www.jstor.org/stable/40043000`.

Mariacristina De Nardi. Wealth inequality and intergenerational links. *Review of Economic Studies*, 71(3):743–768, 2004. URL `https://EconPapers.repec.org/RePEc:oup:restud:v:71:y:2004:i:3:p:743-768`.

Giulio Fella. A generalized endogenous grid method for non-smooth and non-concave problems. *Review of Economic Dynamics*, 17(2):329–344, 2014. ISSN 1094-2025. doi: https://doi.org/10.1016/j.red.2013.07.001. URL `https://www.sciencedirect.com/science/article/pii/S1094202513000392`.

Finansinspektionen. Swedish consumption loans, 2019. URL `https://www.fi.se/en/published/reports/reports/2019/swedish-non-mortgage-lending/`.

Marjorie Flavin and Takashi Yamashita. Owner-occupied housing: Life-cycle implications for the household portfolio. *American Economic Review*, 101(3): 609–14, May 2011. doi: 10.1257/aer.101.3.609. URL `https://www.aeaweb.org/articles?id=10.1257/aer.101.3.609`.

Robert Hubbard, Jonathan Skinner, and Stephen Zeldes. Precautionary saving and social insurance. *Journal of Political Economy*, 103(2):360–99, 1995. URL `https://EconPapers.repec.org/RePEc:ucp:jpolec:v:103:y:1995:i:2:p:360-99`.

Joachim Hubmer, Per Krusell, and Anthony A. Smith. Sources of us wealth inequality: Past, present, and future. *NBER Macroeconomics Annual*, 35:391–455, 2021. doi: 10.1086/712332. URL `https://doi.org/10.1086/712332`.

Fedor Iskhakov, Thomas H. Jørgensen, John Rust, and Bertel Schjerning. The endogenous grid method for discrete-continuous dynamic choice models with (or without) taste shocks. *Quantitative Economics*, 8(2):317–365, 2017. doi: https://doi.org/10.3982/QE643. URL `https://onlinelibrary.wiley.com/doi/abs/10.3982/QE643`.

Robert C. Merton. Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics*, 51(3):247–257, 1969. ISSN 00346535, 15309142. URL `http://www.jstor.org/stable/1926560`.

Robert C Merton. Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4):373 – 413, 1971. ISSN 0022-0531. doi: https://doi.org/10.1016/0022-0531(71)90038-X. URL http://www.sciencedirect.com/science/article/pii/002205317190038X.

Mariacristina De Nardi and Giulio Fella. Saving and Wealth Inequality. *Review of Economic Dynamics*, 26:280–300, October 2017. doi: 10.1016/j.red.2017.06.002. URL https://ideas.repec.org/a/red/issued/16-340.html.

Paul A. Samuelson. Lifetime portfolio selection by dynamic stochastic programming. *The Review of Economics and Statistics*, 51(3):239–246, 1969. ISSN 00346535, 15309142. URL http://www.jstor.org/stable/1926559.

Statistics Sweden. Housing expenditures, housing expenditures as a percentage of disposable income, average disposable income minus net housing expenditures by tenure, 2022a. URL https://www.scb.se/en/finding-statistics/statistics-by-subject-area/household-finances/income-and-income-distribution/households-finances/pong/tables-and-graphs/housing-costs/housing-expenditures-housing-expenditures-as-a-percentage-of-disposable-income-average-di

Lars E. O. Svensson. Are swedish house prices too high? why the price-to-income ratio is a misleading indicator. 2023. URL https://larseosvensson.se/files/papers/are-swedish-house-prices-too-high.pdf.

Jessica Wachter and Motohiro Yogo. Why do household portfolio shares rise in wealth? *Review of Financial Studies*, 23(11):3929–3965, 2010. URL https://EconPapers.repec.org/RePEc:oup:rfinst:v:23:y:2010:i:11:p:3929-3965.

Paul Willen and Felix Kubler. Collateralized borrowing and life-cycle portfolio choice. 2006 Meeting Papers 578, Society for Economic Dynamics, 2006. URL https://ideas.repec.org/p/red/sed006/578.html.

Rui Yao and Harold H. Zhang. Optimal Consumption and Portfolio Choices with Risky Housing and Borrowing Constraints. *The Review of Financial Studies*, 18 (1):197–239, 01 2004. ISSN 0893-9454. doi: 10.1093/rfs/hhh007. URL https://doi.org/10.1093/rfs/hhh007.

A Additional Figures

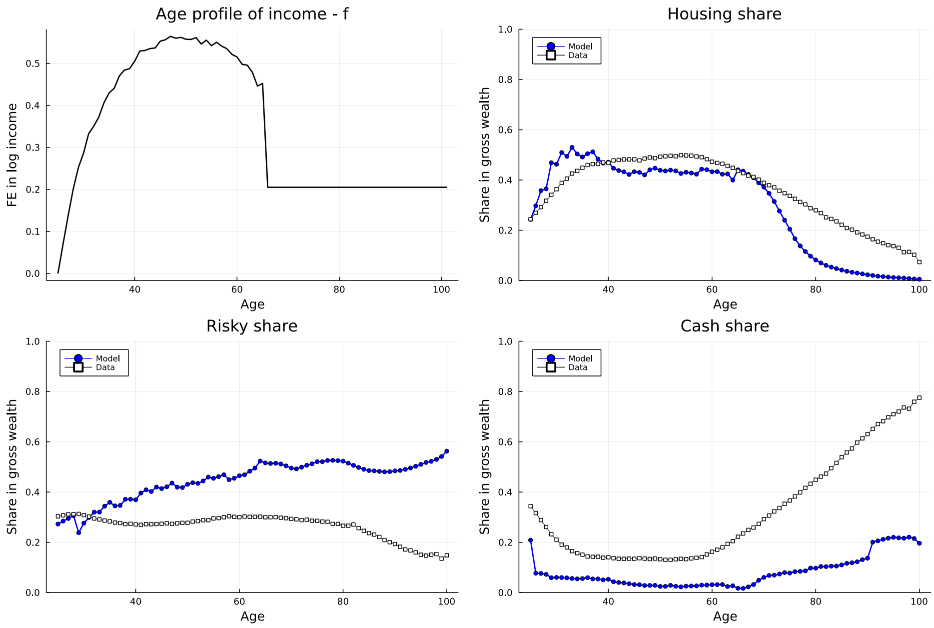

Figure 13: Income profile, i.e. the values f_j on upper left panel. The rest three figures compare untargeted moments across data and the model, in particular, the share of the three components of gross wealth.

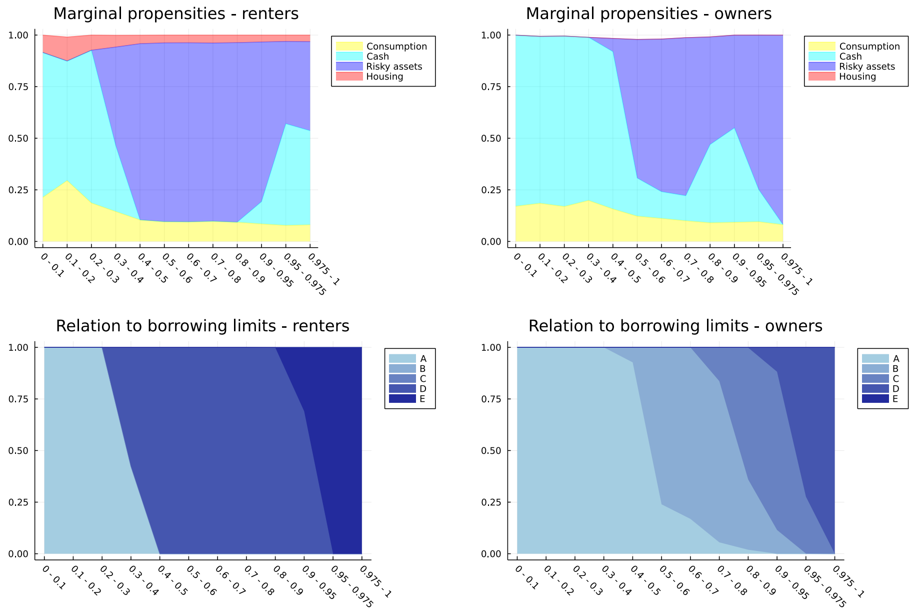

Figure 14: Marginal propensities and borrower status separately for renters and home-owners.

B Numerical Appendix

B.1 Solution of the household's problem

As it was mentioned in the main text, it is tractable to solve the household's problem by considering individual decisions sequentially in the following order:

1. Home ownership state H
2. Saving s versus expenditure x
3. Stock market investments ξ versus vehicles of risk-free decisions.

Here, expenditure is meant to consist of non-durable consumption and rental costs when applicable. Since we consider the expenditure decision after choosing the discrete housing state, it is worthwhile to define indirect utility as a function of expenditure and the housing state as follows:

$$\begin{aligned} u(x,H) \equiv \quad & \max_{c,h} \quad U(c,h) \\ & \text{subject to} \quad x = c + \tau h \widetilde{P} \mathbb{1}_{H=0} \\ & \qquad\qquad\quad H = h \quad \text{if } H > 0 \end{aligned}$$

It is easy to show that $u(x,H)$ satisfies

$$u(x,H) = \begin{cases} H^{\omega} x^{1-\omega} & \text{if } H > 0 \\ \left(\frac{\omega}{\tau \widetilde{P}}\right)^{\omega} (1-\omega)^{1-\omega} x & \text{if } H = 0. \end{cases} \tag{23}$$

Another source of simplification is through normalizing the problem with deterministic trends in aggregate variables. As mentioned earlier, the empirical estimates for secular wage growth G and the house price time trend G_h are nearly identical, and in particular their difference is statistically not significant. Based on this, from now on I will assume $G = G_h$. Using this equality and the homogeneity properties of the preferences, aggregate states except the deviation of housing prices from trend can be eliminated from the recursive problem,[11] which simplifies as shown below:

$$\begin{aligned} V_j(\tilde{p}_h, a, z, H) = \max_{\{x,B,L,M,\xi,H'\}} & \Big\{ (1-\beta) u\big(x, H'\big)^{1-\psi} \\ & + \beta G^{1-\psi} \Big(q_{j+1} \mathbb{E}\Big[V_{j+1}(\tilde{p}'_h, a', z', H')^{1-\gamma} \Big] \\ & + (1-q_{j+1}) \mathbb{E}\Big[B(\tilde{p}'_h, \hat{s}' + \alpha \tilde{p}'_h H')^{1-\gamma} \Big] \Big)^{\frac{1-\psi}{1-\gamma}} \Big\}^{\frac{1}{1-\psi}} \end{aligned} \tag{24}$$

[11] With $G \neq G_h$, one aggregate state representing the current wedge between the aggregate wage and house price index would be needed to keep track of.

subject to

$$a = x + s + D(H, H', \tilde{p}_h) \tag{25}$$
$$s = B + \xi + F\mathbb{1}_{\xi>0} - M - L \tag{26}$$
$$\hat{s}'G = R'\xi + R^f B - R^m M - R^l L \tag{27}$$
$$a' = \hat{s}' + \exp(y') \tag{28}$$
$$H \in \{0, H_1, \dots H_l\} \tag{29}$$

Note that these equations are obtained through some abuse of notation, as variables a, s, $\hat{s}$, x, h, B, ξ, F, M, L and $\exp(y)$ and functions V and B are all rescaled versions of themselves by G^t. Furthermore, for the sake of increased readability of lengthy formulas, in this appendix all time indices are eliminated. Instead, values corresponding to the next period are denoted with the prime sign, as customary.

We first divide the above problem into two a discrete housing choice problem and the rest:

$$V_j(\tilde{p}_h, a, z, H) = \max_{H'} \widehat{V}_j(\tilde{p}_h, \hat{a}, z, H') \tag{30}$$

subject to

$$\hat{a} = a - D(H, H', \tilde{p}_h)$$
$$\hat{a} \geq -\min\left\{\eta^m HC(z,j), \delta \tilde{P}^h H'\right\} - \eta^c HC(z,j)$$
$$H' \in \{0, H_1, \dots H_l\}$$

where $\widehat{V}$ denotes the post housing-choice value. $\hat{a}$ is cash-on-hand available after potential housing transactions. Constraint (31) states that housing choices leading to immediate bankruptcy are ruled out. The post housing-choice subproblem is defined by

$$\begin{aligned}\widehat{V}_j(\tilde{p}_h, a, z, H') = \max_{\{x,B,L,M,\xi\}} &\Big\{(1-\beta)u\big(x, H'\big)^{1-\psi} + \\ &+ \beta G^{1-\psi}\Big(q_{j+1}\mathbb{E}\Big[V_{j+1}(\tilde{p}'_h, a', z', H')^{1-\gamma}\Big] + \\ &+ (1-q_{j+1})\mathbb{E}\Big[B(\tilde{p}'_h, \hat{s}' + \alpha\tilde{p}'_h H')^{1-\gamma}\Big]\Big)^{\frac{1-\psi}{1-\gamma}}\Big\}^{\frac{1}{1-\psi}}\end{aligned} \tag{31}$$

subject to

$$\hat{a} = x + s \tag{32}$$
$$s = B + \xi + F\mathbb{1}_{\xi>0} - M - L \tag{33}$$
$$\hat{s}'G = R'\xi - R^f B - R^m M + R^l L \tag{34}$$
$$a' = \hat{s}' + \exp(y') \tag{35}$$

Finally, we need to solve the consumption-saving problem and the portfolio allocation problem. The latter is partly trivial since as the households aim to minimize interest expenses on debt, we have

$$\begin{cases} B = s - \xi - F\mathbb{1}_{\xi>0},\ M = 0,\ L = 0 & \text{if } s - \xi - F\mathbb{1}_{\xi>0} \geq 0 \\ B = 0,\ M = -(s - \xi - F\mathbb{1}_{\xi>0}),\ L = 0 & \text{if } 0 \geq s - \xi - F\mathbb{1}_{\xi>0} \geq -M_{max} \\ B = 0,\ M = M_{max},\ L = -(s - \xi - F\mathbb{1}_{\xi>0}) - M_{max} & \text{if } -M_{max} \geq s - \xi - F\mathbb{1}_{\xi>0} \end{cases} \tag{36}$$

where M_{max} is a temporary abbreviation for $\min\left\{\eta^m HC(z,j), \delta\tilde{P}^h H'\right\}$.

To simplify ideas and notation, let us introduce the expected continuation value, given today's non-trivial saving policies

$$\begin{aligned} \widetilde{V}_j(\tilde{p}_h, z, s, \xi, H') = &\Big(q_{j+1}\mathbb{E}\Big[V_{j+1}(\tilde{p}'_h, a', z', H')^{1-\gamma}\Big] \\ &+ (1-q_{j+1})\mathbb{E}\Big[B(\tilde{p}'_h, \hat{s}' + \alpha\tilde{p}'_h H')^{1-\gamma}\Big]\Big)^{\frac{1-\psi}{1-\gamma}} \end{aligned}$$

To be more concrete, $\widetilde{V}$ takes s and ξ as given and the other portfolio allocation decisions are determined by (36). Then we can write

$$\widehat{V}_j(\tilde{p}_h, a, z, H') = \max_{\{x,\xi\}}\left\{(1-\beta)u\big(x, H'\big)^{1-\psi} + \beta G^{1-\psi}\widetilde{V}_j(\tilde{p}_h, z, a - x, \xi, H')\right\}^{\frac{1}{1-\psi}} \tag{37}$$

subject to

$$x \leq a + \min\left\{\eta^m HC(z,j), \delta\tilde{P}^h H'\right\} + \eta^c HC(z,j) \tag{38}$$

$$0 \leq \xi \leq a - x + \min\left\{\eta^m HC(z,j), \delta\tilde{P}^h H'\right\} + \eta^c HC(z,j) \tag{39}$$

Algorithm to solve for value and policy functions Fix discrete grids for savings $\{\hat{s}_1 = \bar{s}, \hat{s}_2, \ldots, \hat{s}_i, \ldots, \hat{s}_N\}$, for the possible values of permanent income $\{z_1, \ldots, z_{i_z}, \ldots, z_{N_z}\}$ and the housing price deviations $\{\tilde{p}^h_1, \ldots, \tilde{p}^h_{i_p}, \ldots, \tilde{p}^h_{N_p}\}$. In step j assume either that $q_{j+1} = 0$ or that we already have value and policies for period $j+1$. Consider a given housing level H and an exogenous state $(z_{i_z}, \tilde{p}^h_{i_p})$. Next it is described how to compute optimal policies and the value function of a household starting in this state at age j.

1. For all i (savings values) compute

   (a) the optimal risky share ξ. Given s_i, the optimal risky share is found by globally optimizing

   $$\xi = \arg\max_{\xi} \widetilde{V}_t(\tilde{p}^h_{i_p}, z_{i_z}, s_i, \xi, H')$$

In particular, I evaluate the derivative $\partial \widetilde{V}_j / \partial \xi$ for 100 equidistant points ranging from 0 to $\min\left\{\eta^m HC(z_{i_z}, j), \delta \exp(\tilde{p}^h_{i_p}) H'\right\} + \min\left\{\eta^c HC(z_{i_z}, j), \bar{L}\right\}$. In addition, I add grid points slightly higher and lower than the cutoffs corresponding to the boundaries appearing in (36). This is necessary, as at these points $\partial \widetilde{V}_j / \partial \xi$ is not continuous, since there is a jump in the interest rate of the effective risk-free investment alternative. In each segment without such a jump, where the sign switch of the derivative is consistent with an interior optimum (the exact direction of the right sign switch depends on parameters), I solve for the candidate by the secant method. Among the candidates I also include endpoints and regime switch points where the sign switch is consistent with being an optimum. Then the candidate with the highest $\widetilde{V}_j$ is chosen as optimal. Note that in the absence of discrete housing decisions, this would be a concave problem and relying on the first-order condition would be sufficient.

2. When saving is strictly larger than the borrowing limit, candidates for consumption are pinned down by first order condition

$$(1-\beta)\frac{\partial u\big(x, H'\big)^{1-\psi}}{\partial x} = \beta G^{1-\psi} \frac{\partial \widetilde{V}_j(\tilde{p}_h, z, s, \xi, H')}{\partial s}$$

Since u is still concave, for each grid point of savings, there is exactly one suitable value for consumption which also implies a value for the start of the period cash-on-hand through the budget constraint. However, due to the non-concavities in $\widetilde{V}$ it might happen that the correspondence between the savings grid points and cash-on-hand is not one-to-one, therefore for some segments of cash-on-hand more saving candidates are available, only one of which is actually optimal. Since optimal saving is still an increasing function of cash-on-hand (this already follows from u being concave and $\widetilde{V}$ being increasing, see Iskhakov et al. (2017)), there is a trick to find to correct one.

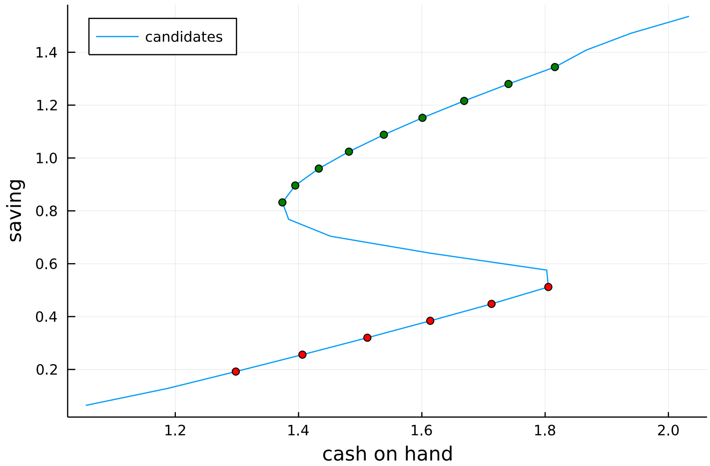

Figure 15: The implied red and green values of cash-on-hand overlap. We need to identify where the optimal saving function jumps from the red segment to the green one.

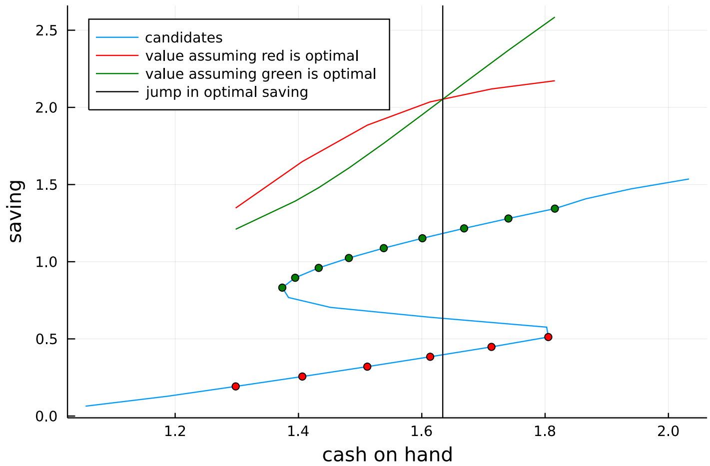

Figure 16: Conditional value functions are computed over both the red and green segments. Their intersection shows where the discontinuity in saving (and the other policy functions) will be. Around the switch extra grid points are added (not indicated on the figure)

Saving grid points (and corresponding values of ξ) belonging to suboptimal regions are deleted. Furthermore, an extra grid point added to capture constrained agents as in Carroll (2006): The cash-on-hand value implied by minimal saving is the lowest point where the first order condition holds with equality. Below that level, the agent would optimally consume all their assets. After this the conditional optimal consumption, saving and investment policies are obtained through linear interpolation.

3. The above step is repeated with setting $\xi = 0$. Then the threshold for participation is found by comparing implied values with and without partification at the gridpoints, after which I solve for the participation threshold $\bar{a}$ by the secant method on the appropriate interval. I add extra nodes around $\bar{a}$, and build the final interpolated state-conditional optimal policies and value function by combining the two subcases according to the participation decision.

4. As above functions $\widehat{V}$ were interpolated, we can choose optimal housing by solving (30), which is a simple discrete optimization problem. Thresholds for optimal housing choices are again computed via the secant method. Additional grid points are added just below and above the regime switch points, after which the final policy and value functions can be interpolated. The grid used for this is obtained by combining the cash-on-hand grids obtained through the endogenous grid point step in each case. Thus the desirable property of EGM placing more points on regions with higher curvature is maintained.

B.2 Integration

In this section it is discussed how to efficiently and precisely compute expectations with discontinuous and divergent integrands, as in the current model. During the solution procedure described above, several conditional expectations of form

$$\mathbb{E}_t \left[F_{t+1}(a', z', H', p'_H)\right] = \mathbb{E}_t \left[F_{t+1}(a', z', H', p'_H) | s, \xi, z, H', p_H\right]$$

need to be evaluated. In most economic models f would be a smooth function in which case computing expectations via discretizing all shocks and state variables can give an adequate approximation. In particular, when shocks follow a normal distribution, the Gauss-Hermite quadrature rule is known to give very precise estimates even with relatively few grid points. However in our case due to the non-convexities from discrete decisions we have to be more careful: Since f can be discontinuous in a', unwanted jumps would be introduced in the result when slightly changing s would move a discretized a' value from one side of a discontinuity point to the other. This means that exact relation of the jumps in f (which is a deep property of the model) and the location of grid points for shocks (which is arbitrary from the theoretical point of view) can affect the results. In this particular model, this effect unfortunately significantly affects the estimated policy functions, even with fairly large grids for the discretized shocks. To avoid this issue, special care is needed when computing expectations.

Note that the above problem actually appears for one shock only. When integrating over multiple shocks over time, if the innermost expectation is computed correctly, it smooths out the discontinuities in the integrand assuming that the innermost shock has a non-degenerate distribution. I give this role to the transitory income shock, since all agents are subject to it (unlike market return, which would smooth results only for participants). Hence,

$$\mathbb{E}_t\Big[F_{t+1}(a', z', H', p^{H\prime})|s, \xi, z, H', p^H\Big] = \sum_i \sum_j \sum_k P(r_i^M)P(z'|z)P(p_k^{H\prime}|p^H)\cdot$$
$$\cdot \int_{-\infty}^{\infty} F_{t+1}\big(\hat{s}'(s, r_i^M) + \exp(z' + \nu + f_{j+1}), z_j', H', p_k^{H\prime}\big)H(\nu)\,\mathrm{d}\nu$$

Now let us take all the other shocks as given and thus consider the integrand only a function of

$$a' = \hat{s}' + \exp(z' + \nu + f_{j+1})$$

where $\hat{s}'$ is savings including realized returns. With some abuse of notation, we hence are interested in:

$$\int_{-\infty}^{\infty} F(a'(\nu))H(\nu)\,\mathrm{d}\nu$$

Sin ce ν follows a normal distribution, if F were well approximated by polynomials, the most efficient way of computing this integral would be by Gauss-Hermite quadrature. There are however two problems with this:

1. F is only piece-wise continuous creating issue described above. This can be solved by integrating piece-wise between the regime changes corresponding to future optimal housing transactions. This means that it is necessary to keep track of the locations of all jumps in policy functions to compute expectations precisely.

2. F might diverge to infinity close to bankruptcy, i.e. when a' is close to $\bar{a} = s_{min}(z', H', p^{h\prime}, j)$. This commonly happens if F is a function of marginal utility.

   In our case we can assume that $G(a') = F(a')(a' - \bar{a})^p$ is smooth and bounded for an adequate p. As the Epstein-Zin time aggregator is composed of power functions, this turns out to be a reasonable assumption. Let

   $$y = \Big(\hat{s}' - \bar{a} + \exp(y_0 + \nu)\Big)^{1/q}$$

where $y_0 = z' + f_{j+1}$ and q is a non-zero real number. This implies

$$\nu = \phi(y) = \log(y^q + \bar{a} - \hat{s}') - y_0$$

and

$$\phi'(y) = \frac{qy^{q-1}}{y^q + \bar{a} - \hat{s}'}.$$

Then

$$\begin{aligned}
\int_c^d F(a'(\nu))H(\nu)\,\mathrm{d}\nu &= \int_c^d G(a'(\nu))(a'(\nu) - \bar{a})^{-p} H(\nu)\,\mathrm{d}\nu \\
&= \int_{\left(\exp(y_0+c)+\hat{s}'-\bar{a}\right)^{1/q}}^{\left(\exp(y_0+d)+\hat{s}'-\bar{a}\right)^{1/q}} G(y^q + \bar{a}) y^{-pq} H(\phi(y)) \frac{qy^{q-1}}{y^q + \bar{a} - \hat{s}'}\,\mathrm{d}y \\
&= \int_{\left(\exp(y_0+c)+\hat{s}'-\bar{a}\right)^{1/q}}^{\left(\exp(y_0+d)+\hat{s}'-\bar{a}\right)^{1/q}} F(y^q + \bar{a}) H(\phi(y)) \frac{qy^{q-1}}{y^q + \bar{a} - \hat{s}'}\,\mathrm{d}y
\end{aligned}$$

which can be integrated by Gauss-Legendre quadrature since the integrand is well-behaved by the second line of the above equation, when $pq < 0$. Intuitively, this transformation moves all divergence from the integrand to the integrations limits, which is less problematic, since the limit points are not meant to approximate nearby points.