The Composition and Distribution of Wealth and Aggregate Consumption Dynamics*

Bulent Guler Burhan Kuruscu Baxter Robinson

May 30, 2022

Abstract

We study how the composition and distribution of household wealth affects the average marginal propensity to consume (MPC) and the distribution of MPC's. We document several facts in the Survey of Consumer Finances about the composition of household portfolios between housing, mortgage debt, and liquid financial assets. We then build a rich quantitative lifecycle model with heterogeneous returns that matches both the composition and concentration of wealth as well as key aggregate facts about housing and mortgage markets in the US. We use the model to decompose the importance of return heterogeneity, long-term fixed rate mortgages and refinancing, as well as owner-occupied housing on the average MPC. The decomposition exercise shows that return heterogeneity is the most important factor in generating a high average MPC. Illiquid owner-occupied housing also plays a substantial role, while the presence of illiquid mortgage debt reduces the average MPC, as it allows households to borrow against their illiquid assets. We use our model to compare the aggregate and distributional effects of a one-time stimulus payment to the effects of a mortgage debt relief program. We show that mortgage debt relief that targets high loan-to-value households is both less effective at boosting aggregate consumption and more regressive than equal stimulus payments to all households, and that a careful consideration of household's portfolios is vital for policy analysis.

*Guler: Indiana University, bguler@indiana.edu; Kuruscu: University of Toronto, burhan.kuruscu@utoronto.ca; Robinson: University of Western Ontario, brobin63@uwo.ca
For generous financial support, Kuruscu thanks the Social Sciences and Humanities Research Council of Canada and Robinson thanks the Stephen A. Jarislowsky Chair in Central Banking.

1 Introduction

Households' portfolios of assets and debts play a crucial role in determining how they respond to economic shocks. Households that own more liquid financial assets are much less likely to change their consumption patterns due to small fluctuations in their income (Ganong et al., 2020). By contrast, households that own substantial illiquid assets may have large consumption responses to changes in income (Kaplan and Violante, 2014). Understanding both how economic shocks will impact the aggregate economy as well as how fiscal and monetary policy can be used to improve economic outcomes therefore requires the study of both the composition and the distribution of household wealth.

In this paper, we study how both the composition and distribution of household wealth affects the average marginal propensity to consume (MPC) and the distribution of MPC's. We focus on housing and mortgage markets because houses are the largest asset that most households own and because mortgage debt is the largest source of debt for most households. For example, households in the fifth decile of net worth on average own \$52,779 worth of housing and only \$11,046 worth of financial and business assets. These households also have an average of \$33,803 in mortgage debt and only \$3,875 in other debts. We contrast the illiquid nature of housing wealth and mortgage debt with more liquid financial assets. In doing so, we study how the composition of housing assets, financial assets and mortgage debt across the distribution of wealth impacts households' willingness to spend out of a small temporary income shock, their marginal propensity to consume (MPC).

To study how the composition and distribution of wealth impact MPCs, we build a quantitative life cycle model. In the model, each household decides each period whether to rent or own housing, and how large a house they wish to rent or own. Households that own illiquid houses can take out long-term defaultable mortgage debt collateralized by the value of their house. All households can save in liquid financial assets. Households receive uninsurable stochastic labor income shocks and entrepreneurial ability shocks. Different levels of entrepreneurial ability generate heterogeneity in rates of return. We calibrate this model to match key aggregate facts about wealth, housing, mortgage markets, and entrepreneurship as well as the concentration of wealth (net worth) at the top and in the hands of entrepreneurs.

A key feature of the calibrated model is the lower discount factor needed to generate the aggregate wealth-to-income ratio, compared to a model without rate-of-return heterogeneity. With return heterogeneity, households with higher returns accumulate substantial levels of wealth despite having a low discount factor. The majority of households hold little wealth, allowing the model to match both the aggregate amount of wealth as well as the concentration at the top.

Although not targeted, the model also matches the composition of household portfolios between housing and liquid financial wealth. While there is a substantial amount of liquid financial assets in both the model and the data, the median household owns very few of them. The mean liquid wealth-to-income ratio is 2.43 in our model, close to the data's value of 2.62. The median liquid wealth-to-income ratio in the model is 0.06 where the data counterpart is 0.16. Furthermore, 26% of households in the model are wealthy hand-to-mouth (HtM) consumers, in the sense that they own no liquid financial wealth, which falls within the empirical range reported in Kaplan and Violante (2014). Finally, the mean housing wealth-to-income ratio is 1.98 whereas it is 1.97 in the data.

Since many households own relatively few liquid assets, they have high marginal propensities to consume. Overall, the model generates an average annual MPC of 0.29 from a one-time unanticipated transfer of $500. We use our model to decompose the importance of return heterogeneity, long-term fixed mortgages, and owner-occupied housing on the average MPC, sequentially shutting down each of these channels and recalibrating the model to the remaining relevant moments. These decomposition exercises show that return heterogeneity is the most important factor (generating 14 percentage points of the average MPC) and illiquid owner-occupied housing is the second most important factor (7 percentage points). Illiquid mortgage debt reduces the average marginal propensity to consume by 2 percentage points, as it allows households to borrow against their illiquid assets.

Under rate of return heterogeneity, high-return households have strong incentives to save, which means that they accumulate a substantial fraction of aggregate wealth. This generates both realistic wealth inequality and higher aggregate savings for a given discount factor. Shutting down rate of return heterogeneity eliminates these high-return households, reducing wealth inequality, and requires a much larger discount factor in order to match the same wealth-to-income ratio. This higher discount factor encourages saving across the distribution, and thus the median liquid wealth-to-income ratio increases from 0.06 to 0.48 despite the fact that the mean liquid wealth-to-income ratio declines from 2.43 to 2.09, due to reduced financial savings of the wealthy.[1] The fraction of hand-to-mouth

[1] In this case, the homogeneous returns model cannot generate the wealth concentration observed in the data. We recalibrate the model by targeting only the aggregate facts about household wealth, housing, and

households decreases substantially as a result of shutting down rate-of-return-heterogeneity, as households across the distribution save more due to their higher discount factor.

Next, we eliminate the ability of households to borrow against their house in long-term mortgage debt. When we remove mortgage debt, homeowners lose the ability to self-insure their consumption by extracting housing equity. Households accumulate more liquid assets in response, which increases the mean liquid wealth-to-income ratio from 2.09 to 2.23, and increases the average MPC to 0.17. Thus, by studying household's gross holdings of housing assets and mortgage debt rather than net housing wealth, we show that the capacity to borrow against illiquid assets slightly reduces the scope of illiquid assets to explain the high average MPCs found in empirical work. Finally, when we exclude owner-occupied housing, the model is reduced to a standard incomplete-markets lifecycle model with non-durable consumption and rental housing consumption. The median liquid wealth-to-income ratio increases to 1.27 since liquid wealth is the only type of wealth in the economy and the average MPC falls to 0.10.

Finally, motivated by the fact that our model matches both the concentration of wealth and the composition of wealth between housing, mortgage debt, and financial assets, we use our model to compare the aggregate and distributional effects of a one-time stimulus payment to the effects of a mortgage debt relief program. The mortgage relief policy specifically targets households with loan-to-value ratios exceeding 90%, while the stimulus payments accrue equally to all households.[2] The mortgage relief policy is far less effective at boosting aggregate consumption, with only 22% of the fiscal cost spent on consumption in the current year compared to 43% for the stimulus payment. The mortgage relief policy is also far more regressive, with 57% of the funds flowing to households in the top 20% of the wealth distribution, since it is wealthier households that own larger homes and correspondingly have larger mortgage debts. On the other hand, when evaluated using the homogeneous returns model, these two policies result in a much more similar similar aggregate consumption response, and the mortgage relief program looks misleadingly progressive. We conclude that a careful consideration of household's portfolios is vital for conducting this type of policy analysis.

mortgage markets ignoring the distribution of this wealth.

[2]Households receive a stimulus payment of $260, which is exactly the same total fiscal cost as the mortgage debt relief program.

Related Literature

This paper contributes to a large literature studying the aggregate consumption response to fiscal stimulus. A great deal of empirical work has been done estimating the average marginal propensity to consume (see for example among others, Johnson et al. (2006), Parker et al. (2013), Jappelli and Pistaferri (2014), Fagereng et al. (2021), and Lewis et al. (2019)).

In standard incomplete market models with one asset, the average marginal propensity to consume is very low and similar across households, except for those constrained by a borrowing limit. These constrained households act in a hand-to-mouth fashion, where they consume all of their income each period. However, when these models are calibrated to match the level of aggregate wealth in the data, very few households end up being borrowing constrained. Thus, the challenge for the quantitative macroeconomics literature has been to generate a sizable hand-to-mouth consumer base consistent with aggregate wealth accumulation in the data.[3]

In an influential paper, Kaplan and Violante (2014) report that a substantial fraction of households in the data are "wealthy hand-to-mouth", that is they have substantial net worth held in illiquid assets but they lack liquid resources. They contrast these "wealth hand-to-mouth" households with "poor hand-to-mouth" households who do not have liquid or illiquid assets. They show that a calibrated two-asset model, with both liquid and illiquid assets, can generate the fraction of both poor and wealthy hand-to-mouth households in the data and produce a high average MPC, consistent with the data. By contrast, our main contribution is to study how return heterogeneity, which generates empirically realistic wealth inequality, influences the economy-wide average MPC and show that it is at least as important as owner-occupied housing.

Our paper is also related to other quantitative papers studying the average MPC. Attanasio et al. (2020) study a life-cycle model of housing and liquid asset accumulation when households face a behavioural temptation for immediate consumption. They argue that the two-asset model without temptation requires an unreasonably large difference between returns on housing and on liquid assets in order to generate a high fraction of wealthy hand-to-mouth households and that the two-asset model has difficulty generating a large share of wealthy hand-to-mouth households when transaction costs are as high as in the data. Our model, on the other hand, generates a large fraction of wealthy

[3] See Kaplan and Violante (2021) for an excellent review.

hand-to-mouth households with reasonable transaction costs and without requiring an unreasonably higher return on the housing relative to the liquid asset return. Thus, we provide a complementary mechanism that can generate a large fraction of wealthy hand-to-mouth households under reasonable return differentials and transaction costs.

Carroll et al. (2017) study how the distribution of marginal propensities to consume is influenced by preference heterogeneity, which also generates some wealth inequality. They find that discount rate heterogeneity can boost the model-generated marginal propensity to consume substantially. In appendix 11 we contrast our results with rate of return heterogeneity to models of wealth inequality generated with preference heterogeneity and with highly skewed income shocks in the spirit of Castañeda et al. (2003). Finally, Kaplan and Violante (2021) compare four separate models, each containing only one mechanism to generate a high MPC: discount rate heterogeneity as in Carroll et al. (2017), temptation preferences as in Attanasio et al. (2020), return heterogeneity, and the two-asset model as in Kaplan and Violante (2014). They argue that the first three models can generate high MPC's only with unrealistically low median wealth while the two-asset model in Kaplan and Violante (2014) achieves it without this caveat. In this paper, however, we show that the return heterogeneity plays at least as important a role as having two-assets. Our benchmark model with both housing and return heterogeneity not only generates reasonable median wealth but also matches the mean and median liquid wealth in the data, both of which are critical for determining the distribution of MPCs.

Our paper draws heavily from the literature on wealth inequality. As shown by Quadrini (1999), Guvenen et al. (2019), and Cagetti and De Nardi (2006) models of rate of return heterogeneity can generate empirically relevant wealth inequality and match the concentration of wealth. More recent work show that persistent return heterogeneity is key for generating several features of the wealth data that are challenging for models of wealth inequality (See Benhabib, Bisin and Luo (2017); Benhabib and Bisin (2018) for the Pareto tail of the wealth distribution, Gabaix, Lasry, Lions and Moll, 2016; Jones and Kim, 2018 for the dynamics of wealth inequality, and Guvenen et al. (2019) for the high fraction of self-made billionaires among the very rich). On the empirical front, a growing number of papers document large and persistent differences in rates of return across individuals even after adjusting for risk and size (see for example Fagereng, Guiso, Malacrino and Pistaferri (2020), Bach et al. (2020), Campbell et al. (2019) and Smith et al. (2019)). Motivated by this work, our model includes rate of return heterogeneity through an entrepreneurship mechanism.

In the next section we document several facts from the Survey of Consumer Finances about household portfolios that is relevant to our analysis. We then present our life cycle model in section 3. In section 4, we detail how we calibrate our model to the data. We compare the composition and distribution of wealth between the data and our model in section 5. Our main findings about how the aggregate consumption response depends on the composition and distribution of wealth are presented in section 6 and we study the policy implications in section 7. Finally, we conclude in section 8.

2 Empirical Facts

We start by documenting several facts about the composition and distribution of US household wealth using the 1995 Survey of Consumer Finances (SCF). We focus on the year 1995 in order to avoid the instability of housing prices in the subsequent decades. These facts remain relatively constant over time, as can been seen in appendix 9, which document the same facts in the other Survey of Consumer Finances waves from 1989-2019.

We broadly categorize assets into three different categories; housing, financial assets, and private businesses. The largest illiquid asset class that most households own is housing, which we define as the sum of the value of primary residences and other residential real estate, which includes vacation homes and time shares.

We contrast the ownership of illiquid housing with relatively more liquid financial assets, which we define as the sum of liquid financial assets, certificates of deposit, mutual funds, bonds, the cash value of life insurance, stocks, non-residential real estate and other financial assets. We include retirement accounts only for households where the head of the household is older than 65, as early-withdrawal penalties typically make it costly to access before retirement age. Finally, we include the value of private business wealth in a separate category.[4]

We then define mortgage debt as all debt held against the principal residence or other residential real estate. Other debts include credit card debts and miscellaneous debts, but exclude student loans. Net worth is the sum of housing, financial assets, and private businesses less the value of mortgage debt and other debts.

Table 1: Wealth, Asset and Debt Shares

	Top 1%	Next 4%	Next 5%	Deciles 9th	8th	7th	6th	5th	4th	3th	2nd	1st
Net Worth	39	21	11	12	7	5	3	2	1	0	-0	-1
Financial Assets	40	24	12	11	5	3	2	1	1	0	0	0
Housing	7	14	12	18	14	12	10	7	4	1	0	1
Private Businesses	71	17	4	4	2	1	0	0	0	0	0	0
Mortgage Debt	5	9	8	12	12	12	14	12	9	2	1	5

The first row shows the share of total net worth held by different percentiles of the distribution of net worth. Each of the other rows shows the share of the total of one asset or debt category held by different percentiles of the distribution of net worth.

2.1 Financial Wealth

Wealth is highly unequally distributed across different US households. Table 1 shows that the top 1% of US households own 39% of aggregate wealth and the top 5% own 60% of aggregate wealth. Financial assets are slightly more unequally distributed than net worth. The wealthiest 1% of US households own 40% of aggregate financial assets and the wealthiest 5% own 64% of financial assets. To the extent that the ownership of liquid financial assets allows individuals to smooth consumption across time, the concentration of financial assets at the top of the wealth distribution has important implications for how households will be able to respond to both idiosyncratic and aggregate shocks.

Figure 1 shows the composition of household assets over the wealth distribution. The share of assets invested in housing declines over the wealth distribution, with some non-monotonicity around zero net worth. The share of private business increases starkly over the wealth distribution, which is unsurprising given that the wealthiest 1% of households own 71% of aggregate private business wealth.

2.2 Mortgage Debt

[4]We exclude from our analysis both vehicles and non-residential real estate. Both of these asset categories are relatively illiquid, but are considerably smaller than the value of residential housing.

Figure 1: Composition of Assets over the Wealth Distribution

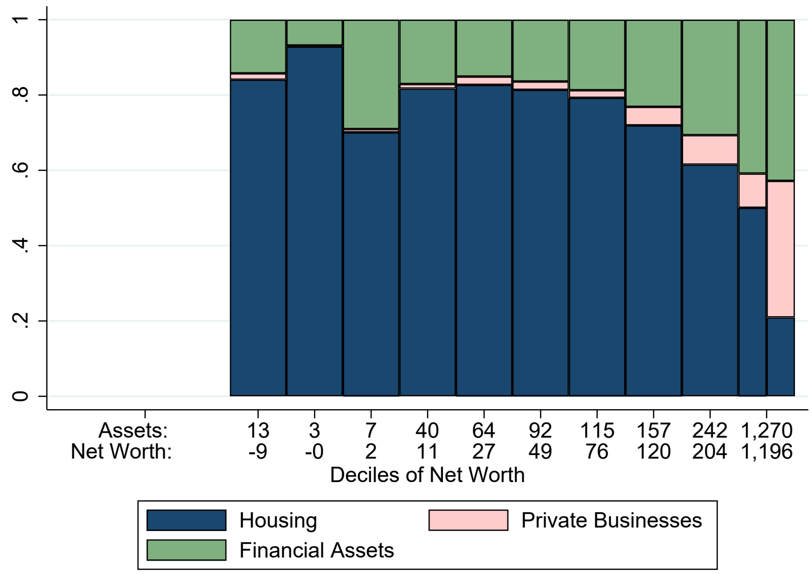

Asset composition by net worth decile. The top decile is split in two to better show the asset composition at the top of the distribution. The x-axis is labeled with the average level of assets and the average net worth within each decile.

Wealthy households hold a substantial amount of mortgage debt. Figure 2 shows the average amount of mortgage debt over the wealth distribution. Despite their wealth, the wealthiest 10% of households own 22% of aggregate mortgage debt.

Many of these wealthy households could pay off the entire balance of their mortgage if they sold their financial assets. Figure 3 shows what proportion of mortgage owning households could fully repay their mortgages if they sold various categories of financial assets and used the proceeds to pay off their mortgage debts. 60% of the wealthiest 5% of householdscould sell their financial assets and fully pay off their mortgage balance. Among the same group, 50% have enough stocks and mutual funds to fully pay off their mortgage balances. The fact that they choose not to is consistent with the idea that they earn higher rates of return on their financial assets than the mortgage interest rates that they pay.

3 Lifecycle Model

In order to study how the composition and distribution of household wealth, we build a quantitative life-cycle model with rate of return heterogeneity, idiosyncratic income

Figure 2: Mortgage Debt over the Wealth Distribution

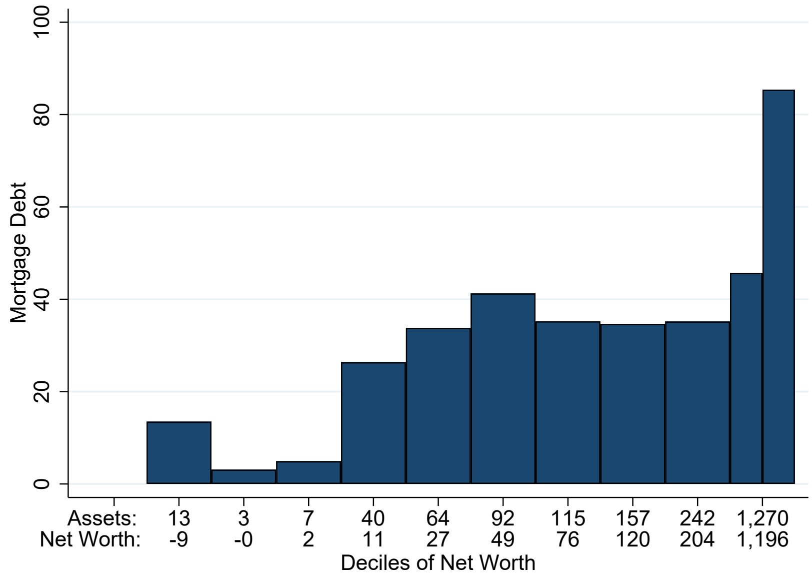

Average mortgage debt by net worth decile. The top decile is split in two to better show the increase in mortgage debt at the top of the distribution. The x-axis is labeled with the average level of assets and the average net worth within each decile.

uncertainty, and incomplete markets. We distinguish liquid financial assets from illiquid housing. Housing is illiquid in the sense that households have to pay a transaction cost whenever it is sold. Housing also generates a consumption flow and can be financed by long-term defaultable mortgages. Households in the model choose how much to invest in both liquid financial assets and illiquid housing. The model therefore generates an endogenous distribution of financial assets, housing, and mortgage debt.

3.1 Environment

Demographics and Preferences: The economy is populated with a continuum of households indexed by i. Age is indexed by j. Households are born at age 1 and start working in the labor market, retire at age J_r, and die with certainty at age J. At each age j, households face a probability of death given by ψ_j.

Households enjoy utility from non-durable consumption goods c and from housing services s. They maximize expected lifetime utility:

$$\mathbb{E}_0 \sum_{j=1}^{J} \beta^{j-1} \Psi_j u(c_{ij}, s_{ij}), \tag{1}$$

Figure 3: Payable Mortgage Debt over the Wealth Distribution

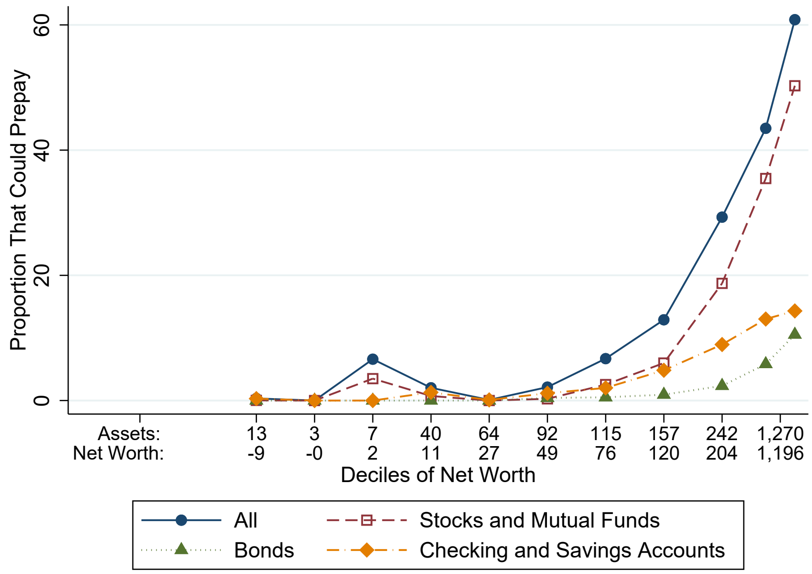

Proportion that could prepay by net worth decile. A household is classified as able to prepay if the value of their stocks and mutual funds, theirs bonds, their checking and savings accounts, or all of these assets together exceeds the total amount of outstanding mortgage debt. The x-axis is labeled with the average level of assets and the average net worth within each decile.

where β is the time discount factor and $\Psi_j = \Pi_{k=1}^{j}(1-\psi_k)$ is the cumulative survival probability until the age of j.

Stochastic Labor Earnings: Working-age households are subject to idiosyncratic income uncertainty: before retirement, log labor income consists of a deterministic component $h(j)$, which only depends on age, and a stochastic component ν_{ij}, which is assumed to follow an AR(1) process. After retirement, the household receives constant social security benefits which only depends on the labor income of the last period working age. Thus, a household's income process y_{ij} is given by

$$y_{ij} = \begin{cases} \exp(h(j) + \nu_{ij}), & \text{if } j \leq J_r \\ y_R(z_{J_r}), & \text{if } j > J_r \end{cases} \tag{2}$$

$$\nu_{ij} = \rho_\nu \nu_{ij-1} + \varepsilon_{ij}, \quad \varepsilon_{ij} \sim i.i.d. \quad N(0, \sigma_\varepsilon^2),$$

where ρ_ν is the persistence of the stochastic component of income and ε is the i.i.d labor income shock which is normally distributed with mean 0 and standard deviation σ_ε.

Financial Assets: All households can save in risk-free, liquid, one period bonds a_{ij}, which delivers a deterministic and exogenous return of r^a each period. There is no unsecured borrowing, so $a_{ij} \geq 0$. In addition to the risk-free bonds, households may have access to higher return investment via entrepreneurship. Each household has a stochastic entrepreneurial ability z_{ij}, which allows the household to operate a project. The project will generate profits according to:

$$\pi_{ij} = \max_{k_{ij}} z_{ij} k_{ij}^{\gamma_e} - r^k k_{ij} - f \tag{3}$$

where r^k is the interest cost of capital and f is a common fixed cost all entrepreneurs face. Entrepreneurial capital is limited by a collateral constraint. Based on the amount of risk-free bonds a_{ij} the household owns their capital stock is limited by:

$$k_{ij} \leq \phi(z_{ij}) \times a_{ij} \tag{4}$$

where $\phi(z_{ij}) > 1$ represents the looseness of the borrowing constraint depending on an entrepreneur's ability z_{ij}. The entrepreneurial ability follows a first-order Markov process with associated transition matrix $\Pi_z \left(z_{ij+1} | z_{ij}\right)$.

Housing Choices: At each age j, a household can choose between renting or owning a house. If they choose to rent they must pay p^r for each unit of housing services s_{ij} they rent. If a household chooses to become a homeowner, they can purchase a house of quality h for a price $p^h h$. A house of quality h delivers housing services of $s_{ij} = h$ each period it is owned. Houses are illiquid, so that the sale of a house is subject to transaction costs, which are assumed to be λ fraction of the house's value. There is a minimum house size $\underline{h}$ below which a purchased house will produce zero housing services, though renters are able to rent houses of lower size than $\underline{h}$. Each period a house is owned, the owner must pay δh in maintenance costs in order to prevent depreciation.

A homeowner can take out a long-term mortgage loan against their house. To ensure computational tractability, we assume that mortgages are amortized over the maximum lifespan of a household J. As death is stochastic, all households that die before period J will have the balance of their mortgage paid out of the value of their remaining assets at the end of life.

When a household takes out a loan, they receive an amount m of consumption good after paying a fixed cost φ_f and a variable cost φ_v proportional to the originated mortgage balance. Households can refinance their mortgage at any time by paying the same fixed and variable costs as when purchasing a house. Households can also default on mortgages at any time period. Upon default the household is free of mortgage debt, but the lender seizes the house and the household loses all the equity in the house.

Mortgage Pricing Each period that the household has a mortgage balance, they make a payment d_{ij} based on the standard amortization formula:

$$\begin{aligned} m_{ij} =& d_{ij}\left[1+\frac{1}{(1+r^m)}+\frac{1}{(1+r^m)^2}+...+\frac{1}{(1+r^m)^{J-j}}\right] \\ d_{ij} =& m_{ij}\frac{r^m(1+r^m)^{J-j}}{(1+r^m)^{J-j+1}-1}=m_{ij}\eta_j \\ m_{ij+1} =& (m_{ij}-d_{ij})*(1+r^m) \end{aligned} \tag{5}$$

Following Hatchondo et al. (2015), Kaplan et al. (2020), and Arslan et al. (2020), we assume that mortgage amortization is calculated using the risk-free mortgage rate. We then price the default risk of each household by having them receive a lower amount of mortgage debt m based on their default probability.

Mortgages are priced endogenously by a set of competitive banks.[5] Since this is a partial equilibrium model, the only purpose of the banks is to pin down the mortgage interest rate pricing based on each household's likelihood of default.

Banks lend to a borrower at a mortgage interest rate based on the borrower's assets (a), the house size they are purchasing (h), their entrepreneurial ability (z), their age (j) and their current labour income shock (ν). Define a mortgage borrower's state as $\theta_{ij}=(a_{ij},h_{ij},z_{ij},j,\nu_{ij})$. Once the mortgage is originated, for a mortgage with outstanding balance m_{ij}, in every period the lender will receive a mortgage payment $d_{ij}(m_{ij})$ if the household stays in the current house without refinancing, the outstanding mortgage balance m_{ij} if the household chooses to refinance the mortgage or sell the house, or the foreclosure value if the household defaults on the mortgage. Foreclosure costs are κ fraction of the value of the house.

The bank's value of a mortgage contract is therefore given by:

$$\bar{v}^l(\theta_{ij})=\frac{1}{1+r^m}\int_{\theta_{ij+1}}v^l(\theta_{ij+1})\Pi(\theta_{ij+1}|\theta_{ij}) \tag{6}$$

[5] Arslan et al. (2015), Chatterjee and Eyigungor (2015), and Guler (2015) use similar models of mortgage pricing.

where

$$v^l(\theta_{ij}) = \begin{cases} d_{ij}\left(m_{ij}\right) + \overline{v}^l(\theta_{ij}) & \text{if} \quad \text{Stay} \\ m_{ij} & \text{if} \quad \text{Refinance or Sell} \\ (1-\kappa)p^h h_{ij} & \text{if} \quad \text{Default} \end{cases}$$

d_{ij} is given by 5 and Π is derived from the policy function of the household.

In equilibrium, these competitive banks must make zero profits. This should imply that the price of a mortgage at the time of origination should be equal to the present discounted value of payments:

$$b\left(\theta_{ij}\right) m_{ij} = \overline{v}^l(\theta_{ij})$$

In practice, default rates are extremely low or zero in the steady state of the model. Because houses do not experience negative house-price shocks, all home owning households have positive net equity in their house, and so will only choose to default if the transactions costs of selling exceeds their equity. This is true for very few households, as the endogenously priced mortgage interest rate would be extremely high if households wanted to take out such a large mortgage that this would be likely.

3.2 Recursive Problem

Every period a household begins the period as an active renter, inactive renter or a homeowner. An active renter can choose to continue to rent or to buy a house. An inactive renter has a default flag in their history and is excluded from owning a house in that period. They regain access to the housing market in the next period with probability ζ. A homeowner can choose to stay in the same sized house, sell their current house to buy a new house, sell their current house to become a renter, refinance or default on their mortgage debt if they have any. Then, the beginning of period value associated with being an active renter is given by:

$$V^R(\theta_R) = \max\left\{V^{Rent}(\theta_R), V^{Buy}(\theta_R)\right\}$$

where $\theta_R = (a, 0, 0, \nu, z, j)$ is the vector of state variables for a renter. Similarly, for a homeowner the beginning of period value function is given by:

$$V^H\left(\theta_H\right) = \max\left\{V^{Stay}(\theta_H), V^R(\tilde{\theta}_R), V^{Refi}(\theta_H), V^{Default}\left(\theta_R\right)\right\}$$

where $\theta_H = (a, h, m, \nu, z, j)$ is the vector of state variables for a homeowner and $\tilde{\theta}_R = \left(a + p^h h\left(1-\lambda\right) - m, 0, 0, \nu, z, j\right)$ is the vector of state variables for a homeowner selling their house. λ is the transaction cost of selling a house. Next, we define each of these value functions.

Inactive Renter

$$V^{Default}(\theta_R) = \max_{c,s,a',k} u(c,s) + \beta(1-\psi_j)\left[\zeta \mathbb{E}V^R(\theta'_R) + (1-\zeta)\,\mathbb{E}V^{Default}(\theta'_R)\right] \tag{7}$$

$$c + a' + p^r s = y(j,\nu) + \pi(z,a,k) + a(1+r^a)$$

Inactive renters choose the level of consumption (c), the amount of housing services to rent (s), the amount to save in the risk-free bond (a'), and the amount of capital (k) to use in the entrepreneurial project.

Active Renter

$$V^{Rent}(\theta_R) = \max_{c,s,a',k} u(c,s) + \beta(1-\psi_j)\mathbb{E}V^R(\theta'_R) \tag{8}$$

$$c + a' + p^r s = y(j,\nu) + \pi(z,a,k) + a(1+r^a)$$

If a household chooses to rent, they choose the level of consumption (c), the amount of housing services to rent (s), the amount to save in the risk-free bond (a'), and the amount of capital (k) to use in the entrepreneurial project.

Buyer

$$V^{Buy}(\theta_R) = \max_{c,a',h',m',k} u(c,h') + \beta(1-\psi_j)\mathbb{E}V^H(\theta'_H) \tag{9}$$

$$c + a' + p^h h' = y(j,\nu) + \pi(z,a,k) + a(1+r^a) + q\left(\theta'_H\right) m' - \left(\varphi_f + \varphi_\nu m'\right)\mathbb{I}_{(m'>0)}$$

If a household purchases a new house, they choose the level of consumption (c), the size of the house to buy (h'), the amount of mortgage debt (m'), the amount to save in the risk-free bond (a') and how much capital (k) to use in the entrepreneurial project.

Homeowner Staying in Current House

$$V^{Stay}(\theta_H) = \max_{c,a',k} u(c,h) + \beta(1-\psi_j)\mathbb{E}V^H(\theta'_H) \tag{10}$$

$$c + a' + \eta_j m + \delta h = y(j,\nu) + \pi(z,a,k) + a(1+r^a)$$
$$m' = m\left(1-\eta_j\right)\left(1+r^m\right)$$

where η_j is defined in equation 5. If a homeowner chooses to stay in their home , they make a simple consumption (c) savings (a') decision and choose how much capital (k) to use in the entrepreneurial project.

Homeowner Refinances their Mortgage

$$V^{Refi}(\theta_H) = \max_{c,a',m',k} u(c,s) + \beta(1-\psi_j)\mathbb{E}V^H(\theta'_H) \tag{11}$$

$$c + a' + m = y(j,\nu) + \pi(z,a,k) + a(1+r^a) + q\left(\theta'_H\right)m' - \varphi_f - \varphi_\nu m'$$

Ifa homeowner chooses to refinance their mortgage, they choose the level of consumption (c), the amount of mortgage debt (m'), the amount to save in the risk-free bond (a') and how much capital (k) to use in the the entrepreneurial project.

4 Calibration

We calibrate the model to match micro data on household portfolios from the 1995 Survey of Consumer Finances. We focus on the year 1995 in order to avoid the instability of housing prices in the subsquent decades. For a set of parameter that are difficult to identify using the data, we rely on commonly used values within the literature. We then jointly calibrate the other parameters to match moments from the data.

A model period corresponds to one year. Households are born at age 25 and may live for 60 years. The stochastic mortality rates $\{\psi_j\}$ are taken from the Center for Disease Control and Prevention. After age 85 ($J = 60$) households die with certainty.

Preferences

For preferences, we use an additive CRRA utility function:

$$u(c,s) = \frac{c^{1-\sigma}}{1-\sigma} + \gamma_s \frac{s^{1-\sigma_s}}{1-\sigma_s}$$

where σ and σ_s capture risk aversion and intertemporal elasticity of substitution for non-durable consumption and housing services respectively. We set $\sigma = 2$, a standard value in the literature, and calibrate σ_s to capture the slightly declining share of housing wealth in total wealth as wealth increases.

Labour Income Process

In order to capture the life-cycle profile of earnings, we include a life-cycle component of earnings ($h(j)$), which is assumed to be a quadratic function in age.[6] From age 25-65, households also face a stochastic labour income process $\nu_{i,j}$ that follows an AR(1) process, which we discretize using the Tauchen method:

$$\nu_{ij} = \rho_\nu \nu_{i,j-1} + \varepsilon_{ij} \qquad\qquad \varepsilon_{ij} \sim N(0, \sigma_\varepsilon)$$

We choose $\rho_\nu = 0.96$ and $\sigma_\varepsilon = 0.16$, which matches the estimation results in Storesletten et al. (2004).

After age 65, we assume households retire and so receive a constant amount of pension income, which is calculated according to the method in Guvenen and Smith (2014). Household labour income in a retirement period is then given by:

$$y_{ij} = g\left(y_{ij=40}\right)$$

where g represents the predicted average lifetime income conditional on the last period's labour income before retirement. Forecasting coefficients are obtained by running an OLS regression of average income on a constant and the last labour income realization before retirement in a simulation.

Entrepreneurship

Entrepreneurial ability is also comprised of the same lifecycle component as labour earnings ($q(j)$) plus an AR(1) process, which we discretize using the Rowenhorst method:

[6] Coefficients are chosen to match doubling of income between age 25 and 55, 10% drop in income between age 55 and 65, and average income of 1.

$$z_{ij} = \rho_z z_{i,j-1} + \chi_{ij} \qquad \chi_{ij} \sim N(0, \sigma)$$

The entrepreneur's borrowing constraint is parameterized as in Guvenen et al. (2019) as:

$$k_{ij} \leq \phi \times (z_{ij} - \underline{z}) \times a_{ij}$$

The entrepreneur's returns to scale (γ_e) is set to 0.85 as in Midrigan and Xu (2014).

Financial Markets and Retirement Accounts

Households can save through risk-free financial asset and we set the annual interest rate for financial assets to 2%. There is no unsecured borrowing, so $a \geq 0$.

Retirement accounts are a non-trivial source of financial wealth for many households. However, because they are typically costly to access before retirement, early withdrawals from retirement accounts are rare and account for only a small fraction of retirement pension wealth (Engelhardt (2002) and Poterba et al. (2007)). As a consequence, retirement account balances are poorly suited to help households smooth out consumption. In order to keep our model computationally tractable, yet still recognize that the accumulation of financial wealth within retirement accounts is a key channel through which people save for retirement, we borrow an approach for dealing with retirement accounts found in Berger et al. (2018). In the model, we do not keep track of retirement account balances. Instead, when households retire at age 65, they receive a one-time lump-sum payment that represents gaining access to the balances of their retirement accounts. This one-time payment is a multiple (ξ) of their their income at age 65. We set $\xi = 1.02$, which corresponds to the ratio of average retirement account balances to average income for households with a head aged 63-67 in the Survey of Consumer Finances.

Housing

For the housing related parameters we follow Arslan et al. (2020). We set the depreciation rate of houses, δ, to 1.5%, consistent with estimates provided in **?**. The transaction cost of selling a house, λ, is set to 7% capturing the realtor fee, sales costs and local taxes and is consistent with estimates provided in **?**. Following**?**, we set the cost of selling a foreclosed house, κ, to 25%. Following the estimated provided by the Consumer

Financial Protection Bureau, we set the fixed cost of originating a mortgage, φ_f, to 2% of average labor income, variable cost of mortgage origination, φ_ν, to 1%.[7]. Finally, we set the probability of removal of default flag, ζ, to 14% capturing the fact that foreclosure flag is removed from credit history after 7 years.[8]

4.1 Externally Chosen Parameters

Table 2 summarizes all of the parameters we choose according to commonly defined values in the literature.

Table 2: Exogenously Chosen Parameters

Interest rate	r^A	2.00 %
Transaction cost of selling housing	λ	0.07
Depreciation of housing	δ	0.015
Variable mortgage origination cost	φ_v	0.01
Fixed cost of refinancing	φ_f	0.02
Bank loss on foreclosure	κ	0.25
Probability the default flag is removed	ζ	0.14
Coefficient of relative risk aversion	σ	2.00
Dispersion of labour ability	σ_ε	0.16
Persistence of labour ability	ρ_ν	0.96
Entrepreneur's returns to scale	γ_E	0.85

4.2 Internally Calibrated Parameters

Table 3 summarizes the internally calibrated parameters and table 4 summarizes the targeted moments. We calibrate the discount factor β to match an aggregate wealth to income ratio of 3.84.[9]

We discipline the curvature on housing utility (σ_s) using the housing share of gross assets owned by the wealthiest 10% of households (0.28). The minimum house size ($\underline{h}$), the price of housing (p^h), the price of renting (p^r), and the household's preference for housing services (γ_s) jointly govern the housing and rental markets in the economy. These are jointly used to match the proportion of homeowners (0.60), the proportion of homeowners younger than 40 (0.42), the share of housing expenditure out of total income (0.25), and the average owner-to-renter house size ratio (2). The risk-free mortgage interest rate (r^m) is used to match the ratio of mortgage debt to income (0.67).

[7]See https://www.federalreserve.gov/pubs/refinancings/default.htm

[8]https://www.experian.com/blogs/ask-experian/how-does-a-foreclosure-affect-credit/

[9]Data moments are computed using 1995 Survey of Consumer Finance whereever possible.

Table 3: Internally Calibrated Parameters

Parameters	Heterogeneous Returns	Homogeneous Return Models: Mortgages	Housing	One-Asset
β	0.87	0.93	0.95	0.95
γ_s	2.14	1.88	0.30	2.93
σ_s	3.28	4.66	1.36	-
$\underline{h}$	1.16	1.55	0.70	-
p^h	1.76	1.58	0.76	-
p^r	0.18	0.13	0.10	-
r^M (%)	5.34	2.86	-	-
$\bar{z}$	-2.15	-	-	-
σ_z	0.96	-	-	-
ρ_z	0.94	-	-	-
f	1.83	-	-	-
ϕ	65.52	-	-	-

Table 4: Targeted Model Moments

Moments	Targets	Heterogeneous Returns	Homogeneous Return Models: Mortgages	Housing	One-Asset
Wealth to Income	3.84	3.68	3.84	3.71	3.84
Housing Expenditure Share	0.25	0.25	0.25	0.25	0.25
Housing Share of Wealthy 10%	0.28	0.24	0.28	0.25	-
Prop. of Homeowners	0.60	0.64	0.66	0.64	-
Prop. of Homeowners < Age 40	0.42	0.41	0.40	0.32	-
Owner-Rental House Size Ratio	2.00	2.16	2.00	1.85	-
Mortgage Debt to Income	0.67	0.69	0.67	-	-
Entrepreneurs' Wealth Share	0.42	0.42	-	-	-
Top 1% Wealth Share	0.39	0.37	-	-	-
Top 10% Wealth Share	0.72	0.79	-	-	-
Proportion of Entrepreneurs	0.10	0.09	-	-	-
Entrepreneurial Debt to GDP	1.50	1.57	-	-	-

Entrepreneurial ability (z) follows an AR(1) process that is discretized using the Rowenhorst technique. All households start life with the lowest entrepreneurial ability shock, and over time they receive stochastic shocks. As a cohort ages, the distribution of entrepreneurial ability therefore converges toward the invariant distribution of the AR(1) process. We calibrate the mean, standard deviation and auto-correlation of the discretized AR(1) process to match the share of wealth owned by the wealthiest 1% and 10% of households (0.39 and 0.72 respectively) as well as the share of wealth held by entrepreneurs (0.42). We calibrate the fixed cost of entrepreneurship (f) to match the proportion of households that own and actively manage a business (0.10). We calibrate the leverage constraint of the entrepreneurs (ϕ) to match the ratio of entrepreneurial debt to income (1.50). Finally, we set the entrepreneur's cost of borrowing (r^k) equal to the risk-free mortgage interest rate (r^m).

In addition to our full model, we also calibrate three other versions of our benchmark model without rate of return heterogeneity. These versions allow us to illustrate the relative impact of return heterogeneity, long-term mortgage debt, and illiquid housing for on the distribution of MPCs. We call these alternative models the homogeneous returns models. In these homogeneous returns models, we do not target the moments regarding the wealth concentration, and we set every household's entrepreneurial ability equal to zero. As a result, these versions grossly understate the wealth concentration observed in the data. Nevertheless, these models are useful benchmarks to compare since they are similar to models of housing and mortgage markets used in most recent work (see for example, Kaplan et al. (2020); Garriga and Hedlund (2020); Arslan et al. (2015); Guler (2015)).

We contrast three versions of the homogeneous returns models. The first has both illiquid housing and long-term mortgage debt, which we label either the full homogeneous returns more or the "Mortgages" model. The second has illiquid housing, but households cannot borrow against the value of their homes, which we label the "Housing" model. Finally, our third version has no illiquid housing assets or long-term mortgage debt, in this model all households must rent housing services, which we call the "One-Asset" model.

5 Consistency between the Model and the Data

Our full model, with heterogeneous returns, matches the distribution of wealth much better than the homogeneous returns mode. This is obvious given that the heterogeneous return process is disciplined directly with statistics characterizing the concentration of wealth at the top of the distribution. In this section, we show that in addition to matching the distribution of wealth, including heterogeneous returns also improves the model's ability to match the composition of the wealth over the wealth distribution.

Figure 4 shows the composition of assets over the wealth distribution in the data and in both the heterogeneous returns and full homogeneous returns models. Both models do a good job at predicting the decline in the share of wealth invested in housing at the top of the wealth distribution. Given that households in the models cannot engage in unsecured borrowing and their houses cannot lose value after they purchase them, no households have negative net worth nor negative equity in their homes. As a result, the models have some

Figure 4: Asset Composition

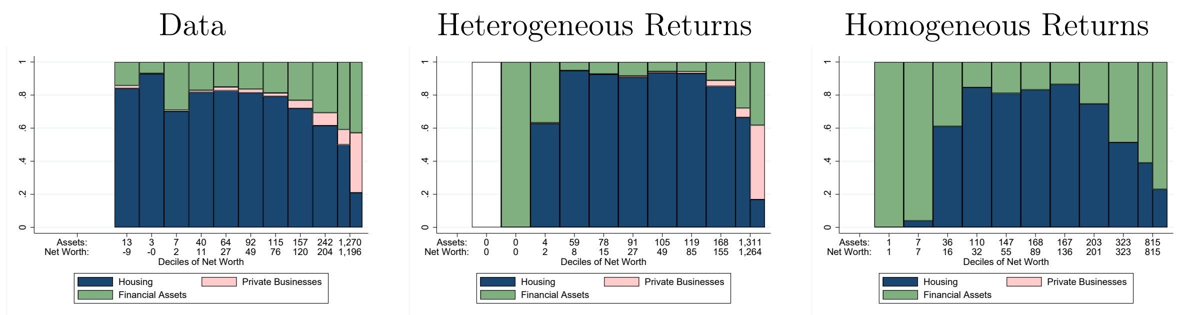

Asset composition by net worth decile. The top decile is split in two to better show the asset composition at the top of the distribution. The x-axis is labeled with the average level of assets and the average net worth within each decile.

difficulty predicting the high share of housing at the bottom of the distribution.

In the following subsections we focus on the distribution of financial wealth and mortgage debt. The former matters for the average MPC since it is liquid and the latter matters for the policy analysis we conduct in section 7. However, the distributions of other assets including housing are reported in Table 10 in Appendix 10.

5.1 Financial Assets

The heterogeneous and homogeneous returns models are calibrated to match the same wealth-to-income ratios. However, given the the two models have different savings mechanisms, the composition of that wealth is very different. Particularly important for the distribution of MPCs, is the distribution of liquid financial assets, which allow households to smooth their consumption over time.

Table 5: Share of Financial Assets Held over the Wealth Distribution

		Next		Deciles								
	Top 1%	4%	5%	9th	8th	7th	6th	5th	4th	3th	2nd	1st
SCF (1995)	40	24	12	11	5	3	2	1	1	0	0	0
Het. Returns	47	36	8	3	1	1	1	1	1	0	0	0
Homo. Returns	18	29	17	17	6	2	3	3	2	2	1	0

The distribution of these financial assets in the heterogeneous returns model comes much closer to matching the distribution seen in the data. In the data, 40% of all financial assets are held by the wealthiest 1% of households. Table 5 shows that the heterogeneous returns model slightly over-predicts this, while the homogeneous returns model has far too few financial assets held by the wealthiest 1% households. As a result, the homogeneous returns model has far too many assets held by the bottom half of households. This is vitally important for the average MPC, as households without much liquid financial wealth will be the ones with high MPCs. In the data, less than 3% of financial assets are owned by the bottom half of households. The heterogeneous returns model is able to match this exactly, while the corresponding number is 9% in the homogeneous returns model.

5.2 Mortgage Debt

Figure 5 shows the level of mortgage debt over the wealth distribution. In the data, the top decile has the largest average amount of mortgage debt. The homogeneous returns model entirely fails to generate this debt, instead having zero mortgage debt at the top of the distribution. Without any source of rate-of-return heterogeneity, wealthy households in the homogeneous returns model choose to fully pay off their mortgage debt. By contrast in the heterogeneous returns model, many wealthy households choose to take out mortgage debt in order to invest additional financial resources in their businesses, which earn higher rates of return.

However, the model fails to generate enough mortgage debt for households between the 70th and 95th percentiles. One caveat of the current heterogeneous returns calibration is that higher returns are only attainable through entrepreneurship. Since only 9 percent of the population in the model are entrepreneurs and the calibrated entrepreneurial productivity process is very persistent, the model generates a large amount of mortgage debt at the very top of the wealth distribution. In reality, people can earn higher returns due to investing in publicly-traded stocks, mutual funds or other financial assets. We are currently considering a broader definition of high-return earners that includes both entrepreneurs and stock market investors, which should generate larger mortgage holdings not just for the very top but also for those close to the top of the wealth distribution.

6 Quantitative Results

In this section, we use our model to decompose the importance of return heterogeneity, long-term fixed mortgages and refinancing, and owner-occupied housing on the average MPC, sequentially shutting down each of these channels and recalibrating the model to the remaining relevant moments. We calculate the marginal propensity to consume in all of these models by giving households a $500 tax rebate and comparing their level of

Figure 5: Mortgage Debt over the Wealth Distribution

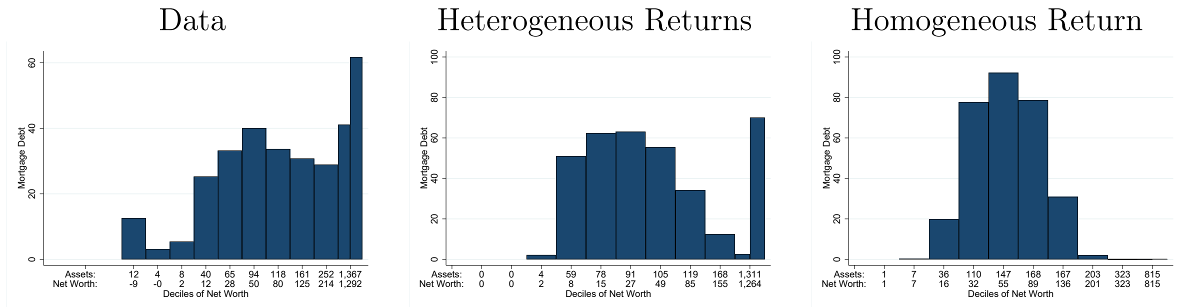

Average Mortgage Debt by net worth decile. The top decile is split in two to better show the asset composition at the top of the distribution. The x-axis is labeled with the average level of assets and the average net worth within each decile.

consumption with and without the tax rebate. Table 6 reports the average marginal propensity to consume across all households in the stationary distribution of each model, as well as the proportion of hand-to-mouth households and the average MPC across different subgroups in the economy

6.1 The Role of Inequality

Including heterogeneous returns in the model allows us to match the high degree of wealth inequality we observe in the data, and also generates a larger aggregate marginal propensity to consume (MPC). The average MPC in the heterogeneous returns model is 0.29, which is almost twice as large as the average MPC in the homogeneous returns model 0.15. This large difference is due to the proportion of the population who choose to be hand-to-mouth, in the sense that they save no liquid wealth.

As can be seen in table 6 in the heterogeneous returns model one quarter of all households choose to be hand-to-mouth, while in the homogeneous returns model it is only 6% of households. Households that can earn a high rate of return on their assets have substantially stronger savings motives. Not only is their rate of return higher, but accumulating additional financial assets allows them to expand the scale of their business and therefore increase the amount of profit they can earn. This mechanism generates the high concentration of financial assets that we see in the data, where the top 10% wealthiest households own 76% of financial wealth.

The strong savings motive mechanically reduces the average MPC for these high-return households. However, it also gives rise to a more realistic concentration of

Table 6: The Average Marginal Propensity to Consume

\centering

	Average MPC	Proportion of Hand-To-Mouth			Average MPC		
					Hand-To-Mouth		>0 Liquid
		All	Rich	Poor	Rich	Poor	Wealth
SCF Data		0.26	0.10	0.16			
Het. Returns	0.29	0.25	0.12	0.13	0.89	0.54	0.15
Homo. Returns							
Mortgages	0.15	0.06	0.02	0.03	0.92	0.66	0.11
Housing	0.17	0.08	0.05	0.03	0.91	0.76	0.12
One-Asset	0.10	0.03	0.00	0.03	-	0.62	0.09

We calculate the proportion of hand-to-mouth households in the data as the proportion of households with liquid financial assets less than half of their monthly income.

Table 7: Components of Wealth

	Mean			Median		
	$\frac{\text{Wealth}}{\text{Income}}$	$\frac{\text{Fin Wealth}}{\text{Income}}$	$\frac{\text{Housing}}{\text{Income}}$	$\frac{\text{Wealth}}{\text{Income}}$	$\frac{\text{Fin Wealth}}{\text{Income}}$	$\frac{\text{Housing}}{\text{Income}}$
Data	3.84	2.62	1.97	0.86	0.16	1.22
Het. Returns	3.68	2.43	1.98	0.46	0.06	1.63
Homo. Returns:						
Mortgages	3.84	2.09	2.44	1.57	0.48	3.30
Housing	3.71	2.23	1.47	1.55	0.47	0.78
One-Asset	3.84	3.84	-	1.27	1.27	-

wealth, and substantially boosts the aggregate wealth-to-income ratio. As a consequence, the heterogeneous returns model needs a much lower discount factor in order to generate the same wealth-to-income ratio as in the data. The lower discount factor boosts the marginal propensity to consume of the entire distribution of households, both because it mechanically increases the share of temporary income that households want to consume, and also because it reduces the desired liquid savings of low-return households generating more hand-to-mouth households.

By contrast, in the homogeneous returns model, the absence of high-return high-wealth households mean that a higher discount factor is needed to match the same wealth-to-income ratio. This higher discount factor raises all households' savings motives equally, leading to more liquid wealth throughout the distribution. As a consequence, far fewer households choose to hold no liquid financial assets, making them far less responsive to a $500 cash rebate.

6.2 The Role of Mortgage Debt

Removing long-term mortgage debt from the full homogeneous returns model increases the aggregate marginal propensity to consume by 2 percentage points. Intuitively, preventing households from borrowing against their illiquid housing wealth makes the housing wealth less liquid and makes it more difficult for households to smooth their consumption over time, raising the aggregate marginal propensity to consume. While refinancing mortgage debt has associated transaction costs, these costs are far less than the transaction costs associated with buying or selling a new house and so preventing households from borrowing against their illiquid assets removes one margin of adjustment. Overall though, the effect of removing mortgage debt is small, as homeowners who know that they now cannot borrow against their illiquid assets, endogenously choose to save more in the liquid asset. Table 7 reports that the financial wealth to income ratio in the economy rises from 2.09 to 2.23.

6.3 The Role of Housing

Removing illiquid housing reduces the average marginal propensity to consume from 0.17 down to 0.10. This mechanism is similar to the one in Kaplan and Violante (2014). In the absence of housing there are no wealthy hand-to-mouth households, who have substantial net worth held in illiquid housing but no liquid wealth. Due to the presence of transaction costs, these wealthy hand-to-mouth household are unwilling to adjust the size of their home in order to respond to small income shocks. Due to their substantial net worth, when

Table 8: Mortgage Relief vs. Equal Cash Transfers: Heterogeneous Returns Model

		$\frac{\Delta C}{\text{Fiscal Cost}}$	Benefits By Wealth Quintile				
	Fiscal Cost		0-20%	20-40%	40-60%	60-80%	80-100%
Mortgage Relief	0.6%	22.0%	0.00	0.26	0.17	0.00	0.57
Equal Cash Transfers	0.6%	42.9%	0.20	0.20	0.20	0.20	0.20

they receive a cash transfer they are willing to spend much more of it than a relatively poor hand-to-mouth household, who does not have any illiquid wealth, would be willing to. Across the models we consider, these wealthy hand-to-mouth households have the highest average MPCs. However, the presence of housing as an illiquid asset in the absence of return heterogeneity is not as strong as the one reported in Kaplan and Violante (2014). This difference is due to the fact that the share of wealthy hand-to-mouth consumers in the housing model is 0.05, which is half of its data counterpart of 0.10.

7 Policy Implications

In this section, we study the policy implications of these mechanisms. In response to recent aggregate economic shocks, like the 2008 Financial Crisis and the COVID-19 pandemic, governments have sought to support aggregate consumption demand and to provide some poorly-targeted insurance with direct cash transfers to households. Using our calibrated model as a laboratory, we are able to consider multiple counterfactual policies and evaluate both their efficacy at boosting aggregate consumption demand and their distributional consequences.

We compare the efficacy of mortgage relief payments to households with high LTVs to a policy where all households are given an equal cash transfer. In the mortgage relief policy, all mortgage debtors with a Loan-To-Value (LTV) ratio exceeding 90% are given a one-time cash transfer sufficient to reduce their LTV down to exactly 90%. This policy is unanticipated, and so it does not create an incentive to take out a large amount of mortgage debt or to refinance in anticipation of the policy. The fiscal cost of this policy is 0.6% of GDP in the economy with heterogeneous returns. We contrast this mortgage relief policy with an equal cash transfer. We set the size of the equal cash transfer such that its aggregate fiscal cost is exactly equal to the mortgage relief policy.

Table 7 shows the efficacy and distributional consequences of these policies. The first column indicates the fiscal cost as a share of GDP. The second column shows what

Table 9: Mortgage Relief vs. Equal Cash Transfers: Homogeneous Returns Model

		ΔC	Benefits By Wealth Quintile				
	Fiscal Cost	Fiscal Cost	0-20%	20-40%	40-60%	60-80%	80-100%
Mortgage Relief	0.0%	17.3%	0.00	0.89	0.10	0.00	0.00
Equal Cash Transfers	0.0%	21.7%	0.20	0.20	0.20	0.20	0.20

proportion of the total fiscal cost is spend on consumption within the year that it is dispersed. While only 22.0% of the mortgage relief ends up consumed, 42.9% of the cash transfers are consumed within the year.

The last five columns show the proportion of the fiscal stimulus is received by households across the distribution of wealth. 57% of the mortgage relief goes to households in the top 20% of the wealth distribution. While a small proportion of households in the top 20% of the wealth distribution have high LTV ratios, those that do own disproportionately large houses, and so have very large mortgage debts. Reducing their mortgage debts down until they have a LTV of 90% is thus very expensive.

When we conduct the same exercise in the homogeneous returns model, we get very different results. First, since many fewer people have LTVs below 90%, the mortgage relief policy is much cheaper to implement. Table 7 shows that the total fiscal cost in the homogeneous returns model is less than 0.1% of GDP. The beneficiaries of this policy are also much poorer in this model, falling mostly between the 20th and 40th percentiles of the wealth distribution. As a consequence the difference in aggregate consumption responses of the mortgage relief and cash transfers is far smaller than in the heterogeneous returns model. Overall, these results suggest that policy conclusions drawn from models that fail to generate the concentration as well as the composition of wealth observed in the data can be misleading.

8 Conclusion

This paper studies the importance of the composition and distribution of wealth for the economy's average marginal propensity to consume. To do so, we build a quantitative lifecycle model with rate of return heterogeneity, illiquid housing and long-term mortgage debt. We calibrate this model to key aggregate facts from the Survey of Consumer Finances, and show how this model can closely match both the distribution of wealth, as well as the composition of that wealth over the distribution. In the calibrated model, the average annual marginal propensity to consume is 0.29. We then study the importance of both the composition and distribution of wealth for the aggregate consumption response of the economy. Eliminating the rate of return heterogeneity mechanism from the model and recalibrating reduces the aggregate marginal propensity to consume to 0.15. Removing mortgage debt has a small impact on increasing the average MPC to 0.17. Also removing the illiquid owner-occupied housing from the model reduces the aggregate marginal propensity to consume to 0.10. Overall, return heterogeneity appears to be the most important mechanism for generating a high average MPC.

We also demonstrate the policy relevance of our results by studying the efficiency and distributional consequences of mortgage relief and equal cash transfers. We find that equal cash transfers are both much more effective at boosting aggregate consumption and that they are far more equitable. We conclude that both the composition of wealth of liquid and illiquid assets, as well as the distribution of those assets, can play an important role in determining the aggregate marginal propensity to consume. As policy makers attempt to evaluate the likely effect of different policies, the results in this paper suggest that a careful consideration of household's portfolios is vital.

References

Arslan, Yavuz, Bulent Guler, and Burhan Kuruscu, "Credit Supply Driven Boom-Bust Cycles," *Working Paper*, 2020.

_ , _ , **and Temel Taskin**, "Joint Dynamics of House Prices and Foreclosures," *Journal of Money, Credit and Banking*, March 2015, *47* (S1), 133–169.

Attanasio, Orazio, Agnes Kovacs, and Patrick Moran, "Temptation and Incentives to Wealth Accumulation," *National Bureau of Economic Research Working Paper Series*, 2020, (28938).

Bach, Laurent, Laurent E. Calvet, and Paolo Sodini, "Rich Pickings? Risk, Return, and Skill in Household Wealth," *American Economic Review*, September 2020, *110* (9), 2703–2747.

Benhabib, Jess, Alberto Bisin, and Mi Luo, "Earnings Inequality and Other Determinants of Wealth Inequality," *American Economic Review*, May 2017, *107* (5), 593–597.

_ **and** _ , "Skewed Wealth Distributions: Theory and Empirics," *Journal of Economic Literature*, 2018, *56* (4), 1261–1291.

Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra, "House Prices and Consumer Spending," *The Review of Economic Studies*, July 2018, *85* (3), 1502–1542.

Cagetti, Marco and Mariacristina De Nardi, "Entrepreneurship, Frictions, and Wealth," *Journal of Political Economy*, 2006, *114* (5), 835–870.

Campbell, John Y., Tarun Ramadorai, and Benjamin Ranish, "Do the Rich Get Richer in the Stock Market? Evidence from India," *AER: Insights*, 2019, *1* (2), 225–240.

Carroll, Christopher, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White, "The Distribution of Wealth and the Marginal Propensity to Consume," *Quantitative Economics*, November 2017, *8* (3), 977–1020.

Castañeda, Ana, Javier Díaz-Giménez, and José-Víctor Ríos-Rull, "Accounting for the U.S. Earnings and Wealth Inequality," *Journal of Political Economy*, August 2003, *111* (4), 818–857.

Chatterjee, Satyajit and Burcu Eyigungor, "A Quantitative Analysis of the U.S. Housing and Mortgage Markets and the Foreclosure Crisis," *Review of Economic Dynamics*, April 2015, *18* (2), 165–184.

Engelhardt, Gary V., "Pre-Retirement Lump-Sum Pension Distributions and Retirement Income Security: Evidence from the Health and Retirement Study," *National Tax Journal*, December 2002, *55* (4), 665–685.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri, "Heterogeneity and Persistence in Returns to Wealth," *Econometrica*, 2020, *88* (1), 115–170.

_ **, Martin B. Holm, and Gisle J. Natvik**, "MPC Heterogeneity and Household Balance Sheets," *American Economic Journal: Macroeconomics*, 2021, *13* (4), 1–54.

Gabaix, Xavier, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll, "The Dynamics of Inequality," *Econometrica*, November 2016, *84* (6), 2071–2111.

Ganong, Peter, Damon Jones, Pascal J Noel, Fiona E Greig, Diana Farrell, and Chris Wheat, "Wealth, Race, and Consumption Smoothing of Typical Income Shocks," *National Bureau of Economic Research Working Paper Series*, 2020, (27552).

Garriga, Carlos and Aaron Hedlund, "Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession," *American Economic Review*, June 2020, *110* (6), 1603–1634.

Guler, Bulent, "Innovations in Information Technology and the Mortgage Market," *Review of Economic Dynamics*, July 2015, *18* (3), 456–483.

Guvenen, Fatih and Anthony A. Smith, "Inferring Labor Income Risk and Partial Insurance From Economic Choices: Inferring Labor Income Risk and Partial Insurance," *Econometrica*, November 2014, *82* (6), 2085–2129.

_ **, Gueorgui Kambourov, Burhan Kuruscu, Sergio Ocampo, and Daphne Chen**, "Use It or Lose It: Efficiency Gains from Wealth Taxation," *National Bureau of Economic Research Working Paper Series*, 2019, (26284).

Hatchondo, Juan Carlos, Leonardo Martinez, and Juan M. Sánchez, "Mortgage Defaults," *Journal of Monetary Economics*, November 2015, *76*, 173–190.

Jappelli, Tullio and Luigi Pistaferri, "Fiscal Policy and MPC Heterogeneity," *American Economic Journal: Macroeconomics*, October 2014, *6* (4), 107–136.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles, "Household Expenditure and the Income Tax Rebates of 2001," *The American Economic Review*, 2006, *96* (5), 1589–1610.

Jones, Charles I and Jihee Kim, "A Schumpeterian Model of Top Income Inequality," *Journal of Political Economy*, 2018, *126* (5), 42.

Kaplan, Greg and Giovanni L. Violante, "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 2014, *82* (4), 1199–1239.

_ and _ , "The Marginal Propensity to Consume in Heterogeneous Agent Models," *Working Paper*, 2021, pp. 1–52.

_ , Kurt Mitman, and Giovanni L Violante, "The Housing Boom and Bust: Model Meets Evidence," *Journal of Political Economy*, 2020, *128* (9), 3285–3345.

Krusell, Per and Anthony A. Smith Jr., "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, October 1998, *106* (5), 867–896.

Lewis, Daniel J., Davide Melcangi, and Laura Pilossoph, "Latent Heterogeneity in the Marginal Propensity to Consume," *SSRN Electronic Journal*, 2019.

Midrigan, Virgiliu and Daniel Yi Xu, "Finance and Misallocation: Evidence from Plant-Level Data," *American Economic Review*, 2014, *104* (2), 422–458.

Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland, "Consumer Spending and the Economic Stimulus Payments of 2008," *The American Economic Review*, 2013, *103* (6), 2530–2553.

Poterba, James M, Steven F Venti, and David A Wise, "Preretirement Cashouts and Foregone Retirement Saving: Implications for 401(k) Asset Accumulation," in "Themes in the Economics of Aging," University of Chicago Press, 2007, pp. 23–58.

Quadrini, Vincenzo, "The Importance of Entrepreneurship for Wealth Concentration and Mobility," *Review of Income and Wealth*, 1999, *45* (1), 1–19.

Smith, Matthew, Danny Yagan, Owen M Zidar, and Eric Zwick, "Capitalists in the Twenty-First Century," *The Quarterly Journal of Economics*, 2019, *134* (4), 1675–1745.

Storesletten, Kjetil, Christopher I. Telmer, and Amir Yaron, "Consumption and Risk Sharing over the Life Cycle," *Journal of Monetary Economics*, April 2004, *51* (3), 609–633.

9 Empirical Facts over Time

In this section, we show that our empirical facts from section 2 hold across different waves of the Survey of Consumer Finances.

Figure 6: Composition of Assets over the Wealth Distribution

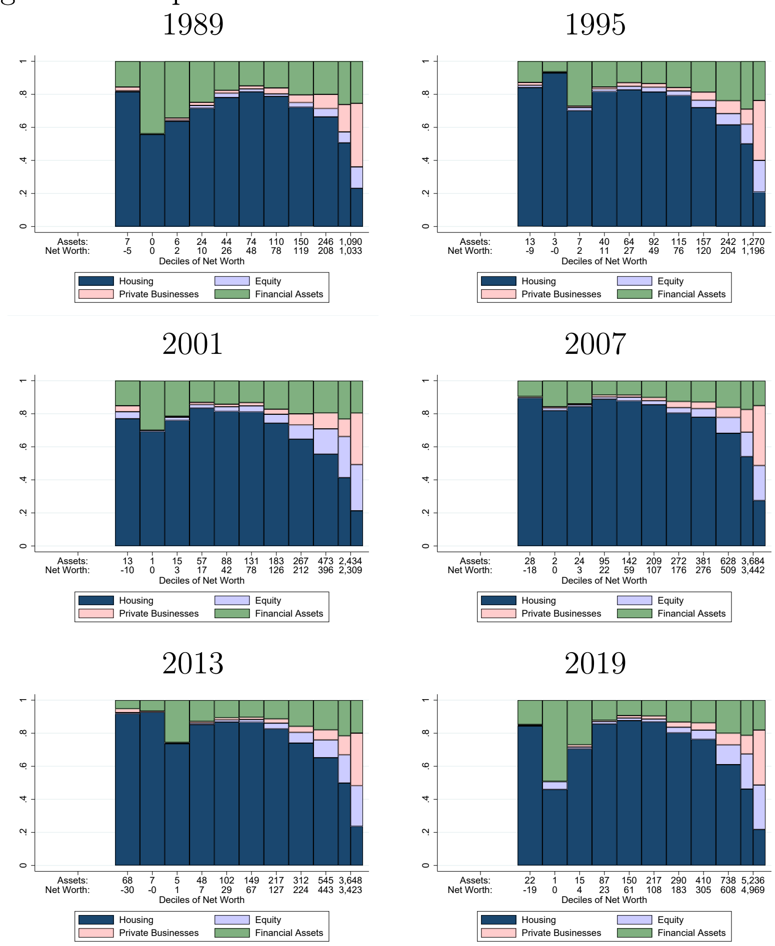

Figure 6 compares the asset composition over the wealth distribution across different waves of the Survey of Consumer Finances. Across all of these waves, the housing share of assets declines starkly over the wealth distribution. By contrast, the private business and public equity share of assets increases over the wealth distribution. Variability in the house share just above a net worth of zero is often driven by a small number of households and as such is likely subject to sampling variability.

Figure 7 shows the average level of mortgage debt over the wealth distribution. In all years, the largest amount of mortgage debt is held by the wealthiest households.

Figure 7: Mortgage Debt over the Wealth Distribution

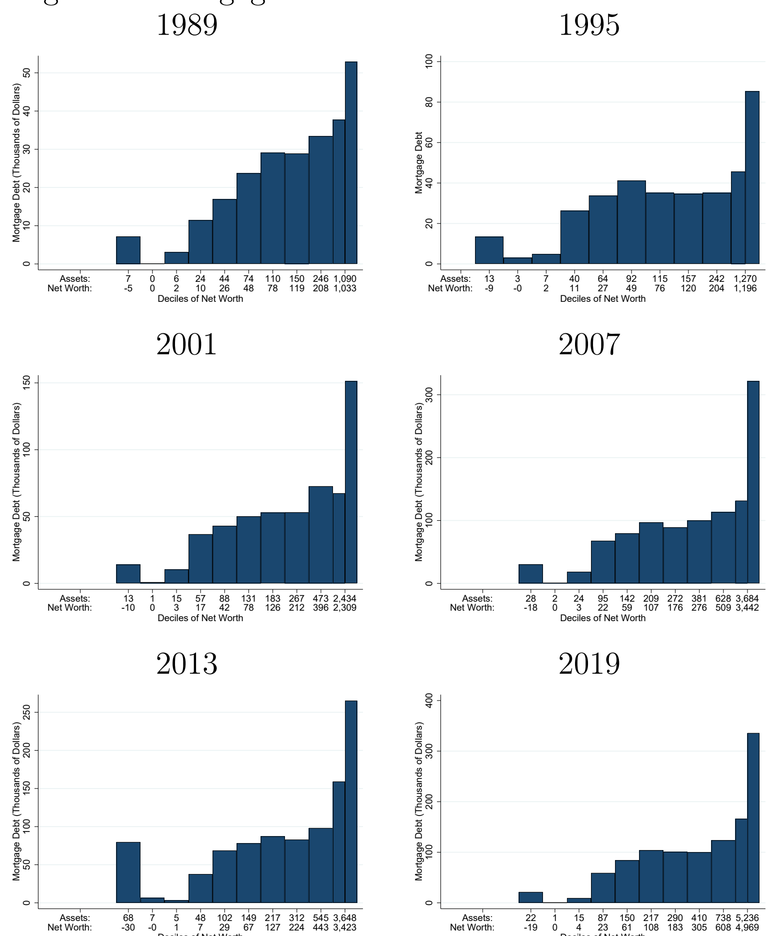

Figure 8 shows what proportion of households could liquidate financial assets and fully pay off their mortgage debt. Of households within the wealthiest 5%, around 60% of them could fully pay off their mortgages by liquidating financial assets.

Figure 8: Payable Mortgage Debt over the Wealth Distribution

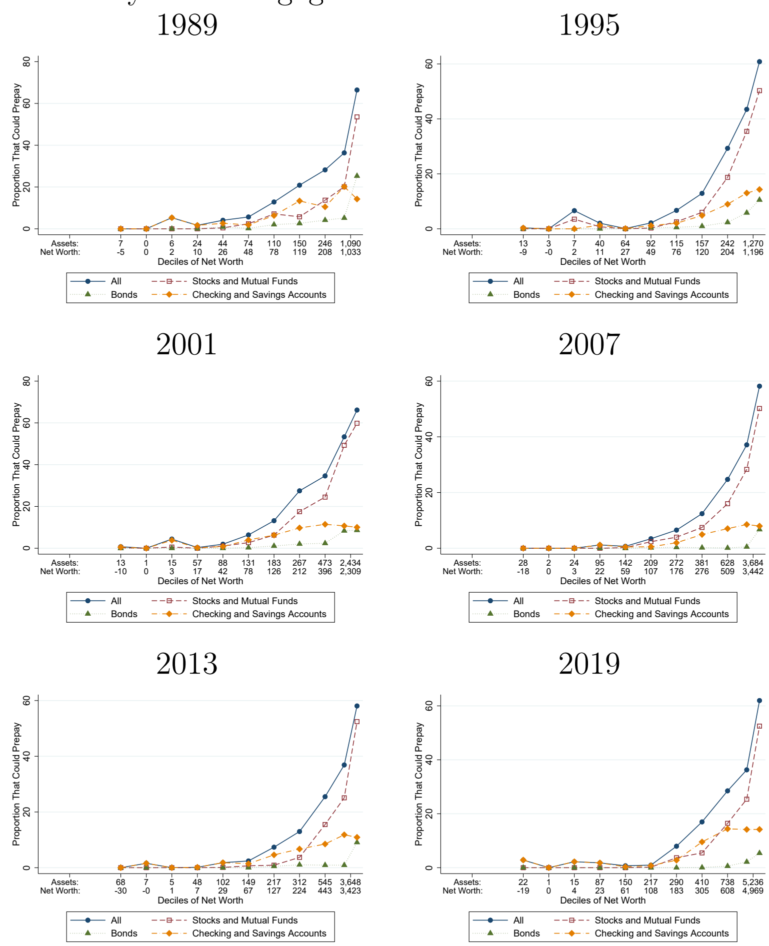

10 Additional Comparisons between Model and Data

Table 10: Share of Different Assets Held over the Wealth Distribution

		Next		Deciles									
	Top 1%	4%	5%	9^{th}	8^{th}	7^{th}	6^{th}	5^{th}	4^{th}	3^{th}	2^{nd}	1^{st}
Net Worth												
SCF (1995)	39	21	11	12	7	5	3	2	1	0	-0	-1
Het. Returns	37	32	10	10	5	3	2	1	1	0	0	0
Homo. Returns	12	22	15	19	12	8	5	3	2	1	0	0
Financial Assets												
SCF (1995)	40	24	12	11	5	3	2	1	1	0	0	0
Het. Returns	47	36	8	3	1	1	1	1	1	0	0	0
Homo. Returns	18	29	17	17	6	2	3	3	2	2	1	0
Business Assets												
SCF (1995)	71	17	4	4	2	1	0	0	0	0	0	0
Het. Returns	59	38	2	1	0	0	0	0	0	0	0	0
Homo. Returns	0	0	0	0	0	0	0	0	0	0	0	0
Gross Housing												
SCF (1995)	7	14	12	18	14	12	10	7	4	1	0	1
Het. Returns	6	17	12	17	13	11	10	8	6	0	0	0
Homo. Returns	3	9	9	16	14	14	13	11	9	2	0	0
Mortgage Debt												
SCF (1995)	5	9	8	12	12	12	14	12	9	2	1	5
Het. Returns	4	7	0	4	11	17	20	20	16	1	0	0
Homo. Returns	0	0	0	0	1	10	26	31	26	7	0	0

11 Alternative Wealth Inequality Mechanisms

In the main body of our paper we contrast our heterogeneous returns model against a homogeneous returns model. In this section, we consider two other models that are commonly used to generate wealth inequality. In particular, we calibrate a "superstar" income shock version of the model, following Castañeda et al. (2003), and a preference heterogeneity version, where agents have different discount factors (β), as in Krusell and Smith (1998).

11.1 Calibration of Superstar Income Shock Model

Table 11: Superstar Shock Model Internally Calibrated Parameters

Moment	Data	Model	Parameter	Value
Wealth to Income Ratio	3.84	3.82	β	0.90
Proportion of Homeowners	0.60	0.63	$\underline{h}$	0.88
Debt to income Ratio	0.67	0.69	r^m	4.50 %
Share of housing in total income	0.25	0.27	p^h	1.61
Home Ownership rate Age 40 and less	0.42	0.41	p^r	0.14
Owner to Renter House Size Ratio	2.00	1.98	γ_s	2.20
Housing Share of Wealth for Top 10%	0.28	0.27	σ_s	3.94
Wealth Gini	0.77	0.81	ν_{n_ν}	5.79
Top 1% Wealth Share	0.39	0.39	$Prob(\nu_{n_\nu}\|\cdot)$	0.06 %
Top 10% Wealth Share	0.72	0.71	$Prob(\nu_{n_\nu}\|\nu_{n_\nu})$	15.61 %

In the superstar income shock model, we add an additional "superstar" level of labour income, which is a high level of income that is both rare and impersistent. The impersistence is important as it provides a strong savings motive for individuals. Castañeda et al. (2003) demonstrate that this type of process can be calibrated in order to match top wealth shares in the US, though **?** show that empirically relevant higher order moments of income shocks does not generate the high top wealth shares observed in US data.

We follow a similar calibration strategy for the superstar income shock model as in 4.We target the same housing and mortgage market moments in both the superstar income shock model and the heterogeneous returns model. In order to calibrate the superstar income shock, we start with the same labour income process as in the heterogeneous returns model, but add one additional superstar state. This state is governed by three additional parameters, the probability of attaining this state, the probability of leaving this state, and the income level of the state. We then discipline these parameters using three moments of the wealth distribution, the top 1% wealth share, the top 10% wealth share, and the wealth Gini.

11.2 Calibration of Preference Heterogeneity Model

Table 12: Preference Heterogeneity Model Internally Calibrated Parameters

Moment	Data	Model	Parameter	Value
Wealth to Income Ratio	3.84	3.84	β	0.75
Proportion of Homeowners	0.60	0.65	$\underline{h}$	1.25
Debt to income Ratio	0.67	0.68	r^m	2.80 %
Share of housing in total income	0.25	0.26	p^h	1.99
Home Ownership rate Age 40 and less	0.42	0.40	p^r	0.17
Owner to Renter House Size Ratio	2.00	2.03	γ_s	0.82
Housing Share of Wealth for Top 10%	0.28	0.28	σ_s	4.74
Top 20% Wealth Share	0.84	0.80	σ_β	0.09
Top 50% Wealth Share	0.98	0.93	ρ_β	0.76

In the preference heterogeneity model, agents have heterogeneous discount factors (β). We calibrate an AR(1) process for β by targeting the top 20% wealth share and the top 50% wealth share. Attempts to calibrate the preference heterogeneity model to match the top 1% and top 10% wealth shares were unsuccessful while restricting $\beta \in [0, 1]$.

11.3 Consistency between the Models and the Data

Table 13: Share of Financial Assets Held over the Wealth Distribution

				Deciles								
	Top 1%	Next 4%	Next 5%	9^{th}	8^{th}	7^{th}	6^{th}	5^{th}	4^{th}	3^{th}	2^{nd}	1^{st}
SCF (1995)	40	24	12	11	5	3	2	1	1	0	0	0
Het. Returns	47	36	8	3	1	1	1	1	1	0	0	0
Superstar Shocks	59	23	7	4	2	1	2	1	1	1	0	0
Beta Heterogeneity	19	29	17	17	6	3	3	3	2	2	1	0

Figure 9: Asset Composition

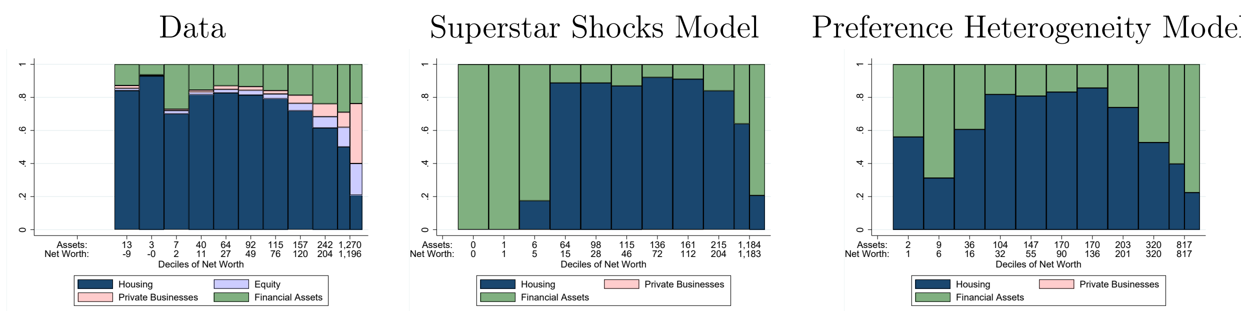

Figure 9 shows the composition of assets over the wealth distribution in the data, the superstar shocks model and the preference heterogeneity model. Both models again do a good job at predicting the decline in the share of wealth invested in housing at the top of the wealth distribution. Both of these models are also able to match the very low levels of wealth at the bottom of the distribution. The preference heterogeneity model also does a better predicting a high housing share at the bottom of the distribution.

Figure 10 shows the level of mortgage debt over the wealth distribution. In the data, the top decile has the largest average amount of mortgage debt. Both the superstar shock model and the preference heterogeneity model entirely fail to generate this debt, instead having little or no mortgage debt at the top of the distribution.

11.4 Implications of Inequality

The average MPC in the heterogeneous returns model is 0.29. The superstar shock model has a very similar MPC of 0.28, despite a lower share of poor hand-to-mouth households. Households that receive superstar income shocks have very strong savings motives. As a

Figure 10: Mortgage Debt over the Wealth Distribution

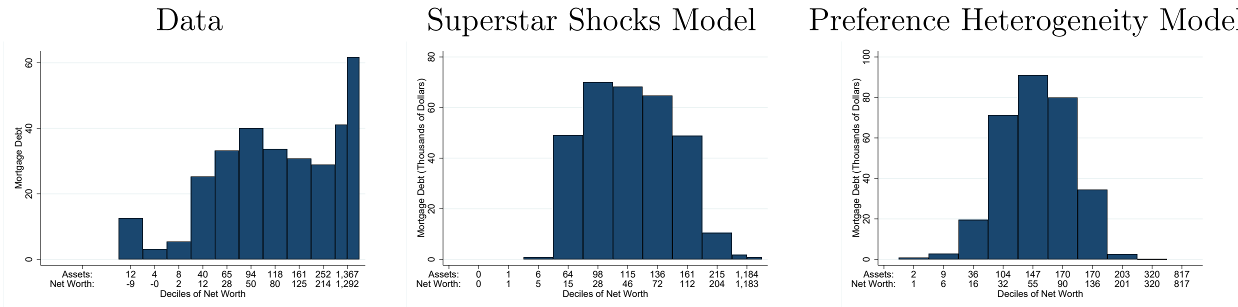

Table 14: The Average Marginal Propensity to Consume

	Average MPC	Proportion of Hand-To-Mouth: All	Rich	Poor	Average MPC: Hand-To-Mouth Rich	Hand-To-Mouth Poor	>0 Liquid Wealth
SCF Data		0.26	0.10	0.16			
Het. Returns	0.29	0.25	0.12	0.13	0.89	0.54	0.15
Superstar Shocks	0.28	0.19	0.12	0.07	0.85	0.66	0.17
Beta Heterogeneity	0.14	0.07	0.02	0.05	0.93	0.65	0.10

consequence, the discount factor necessary to match the same wealth-to-income ratio is lower than the homogeneous returns model. As a result we get both more hand-to-mouth households and higher average MPCs.

By contrast, in the preference heterogeneity model, the average MPC is much lower at 0.14. Since our implementation of this preference heterogeneity model is unable to generate the top tail wealth inequality, financial wealth is too dispersed across the population. As a result, only 7% of households have no financial wealth and so the average MPC is much lower.

12 Computational Solution

MPC Calculation

We calculate marginal propensities to consume by comparing a household's consumption to their consumption if they were given an unanticipated \$500. Technically, this is the average propensity to consume out of additional \$500, though we follow the empirical literature in referring to this consumption response as the marginal propensity to consume.

For some households on the margin between different discrete choice, for example to buy a new house or to refinance a mortgage, an additional \$500 may lead them to choose a different discrete choice. Changes in these discrete choices can sometimes lead to calculated marginal propensities to consumer that are either negative or larger than one. An alternative measure that is easy to compute in our model is to directly back out the marginal propensities to consume from the derivative of the consumption function. There are noticeable discrepancies between the two calculated methods, with the derivative based method generally resulting in higher measured marginal propensities to consume. We prefer to stick the method using an unanticipated \$500 transfer , as this provides the best analogue to empirical estimates of the marginal propensity to consume.