The Secular Shift in the MPC Distribution

Pablo Sánchez*

October 30, 2026

Click here for the latest version

Abstract

This paper studies how and why the distribution of marginal propensities to consume (MPCs) has changed over time, and what this implies for fiscal policy. I document a secular shift: MPCs have risen among low-wealth households and fallen for mid-wealth households. I argue that this is caused by rising house prices and rents, and the secular decline in interest rates. To understand and quantify the drivers, I build a quantitative life-cycle model with housing tenure and portfolio choice. Higher house prices put homeownership out of reach for low-wealth households, while larger rents push more of them into liquidity constraints. Among mid-wealth households, greater down payments delay homeownership and encourage saving, while lower interest rates shift these savings toward liquid buffers. These findings underscore the importance of targeting the right people: transfers targeted toward low-wealth households have become increasingly powerful.

*Northwestern University; pablo.sanchez@u.northwestern.edu; https://pablosanchezf32.github.io/
I am extremely grateful to my advisors Marty Eichenbaum, Giorgio Primiceri, and Matt Rognlie for their invaluable guidance and support. For helpful comments and suggestions, I thank Marios Angeletos, Sergio Rebelo, Diego Känzig, Kunal Sangani, Rodolfo Oviedo, Junko Oguri, Cristoforo Pizzimenti, Jana Obradovic, Tim Seida. I also thank the participants at the Macroeconomic Reading Group and 501 Seminar at Northwestern and the summer research program seminar at the Bank of Mexico for their valuable feedback.

1 Introduction

The marginal propensity to consume (MPC) is a central determinant of macroeconomic policy transmission. Fiscal transfers are a powerful tool to stimulate the economy. Yet, their effect depends not only on their size, but also on who receives them. Transfers are more effective when targeted toward high-MPC households, who spend a larger share of the additional income. Reaching high-MPC households generates a larger and faster boost to aggregate demand per dollar of government spending. Accordingly, fiscal policy depends critically on the distribution of MPCs and on policymakers' ability to identify high-MPC households.

A large literature documents substantial heterogeneity in consumption responses across income, wealth, and household balance sheets (Fagereng *et al.*, 2021; Ganong *et al.*, 2025). These studies consistently show a correlation between MPC and liquid wealth: MPC is higher for liquidity-constrained households. Conversely, total wealth is not correlated with MPC given the existence of wealthy liquidity-constrained households (Kaplan *et al.*, 2014). These findings discipline the MPC distribution in modern macroeconomic models used for policy analysis. However, that distribution is typically treated as time invariant (e.g., Auclert *et al.*, 2018, 2025; Kaplan *et al.*, 2018). Yet household balance sheets are not static: the level, distribution, and composition of wealth change over time. Whether the MPC distribution, and therefore the effectiveness of fiscal targeting and overall policy transmission, has changed with them remains largely unknown. Understanding the drivers of any change in the MPC distribution allows policymakers to anticipate who these high-MPC households are, rather than relying on past estimates. This paper asks how the MPC has evolved across the wealth distribution in the United States, what forces have driven these changes, and how they have altered the effectiveness of fiscal policy.

The paper has three main findings. First, I document a substantial shift in the MPC across the wealth distribution from 1999 to 2019 in the United States. The MPC of low-wealth households increased from 0.13 to 0.26, while that of mid-wealth households declined from 0.38 to 0.16. Second, I argue that this is caused by a combination of rising house prices, higher rents and the secular decline in interest rates. These changes affect savings behavior and the distribution of liquidity-constrained high-MPC households. Third, I do a counterfactual experiment and show that fiscal policy targeted toward low-wealth households has become increasingly powerful. Targeting the bottom of the distribution has become the most effective tool to stimulate the economy through transfers.

Taken together, these three key findings underscore the importance of targeting the right people to stimulate the economy. Knowing who these high-MPC households are and why they are liquidity-constrained is central for policymakers to make effective, real time

policy decisions. I discuss my findings in more detail below.

Empirics. I estimate the MPCs over time using household panel data from the Panel Study of Income Dynamics (PSID) between 1999 and 2019 for the working-age population. Building on Blundell *et al.* (2008), I identify the (non-housing) consumption response to transitory income innovations by instrumenting current income growth with future income growth. I trace the MPC over time using rolling windows of four consecutive PSID waves and apply the same procedure separately to wealth groups.

I show that the MPC has moved in opposite directions across the wealth distribution. The MPC of low-wealth households increased from 0.13 to 0.26, while that of mid-wealth households declined from 0.38 to 0.16. The MPC estimate for high-wealth households has remained relatively low, around 0.1. These trends contrast with the aggregate MPC, which remained roughly stable at about 0.2.

In the early periods of the sample, MPC estimates for mid-wealth households are larger than those for the low-wealth group. This is at odds with standard one-asset heterogeneous-agent macroeconomic models, in which liquidity-constrained households are predominantly at the bottom of the distribution. However, I show that there are two forces behind this. First, the mid-wealth group had a substantial amount of wealthy liquidity-constrained households (as in Kaplan *et al.* (2014)) in the earlier periods of the sample. These are predominantly recent homeowners whose wealth is mainly illiquid. Despite being wealthy, they are liquidity-constrained.

Second, low-wealth households, who are predominantly renters, can respond to large income losses by cutting their housing expenditure. Renters can move to a cheaper unit relatively easily. This buffers their consumption against large transitory shocks. Homeowners lack this margin, since selling a house is costly and slow. Recent owners in the middle of the distribution therefore cut consumption more sharply after a large income loss.

I reinforce the main result by looking at the evolution of the share of liquidity-constrained, high-MPC households. I use the hand-to-mouth[1] definition from Kaplan *et al.* (2014) as a proxy for these households. I show that the share of liquidity-constrained low-wealth households increased, while it declined for mid-wealth. Additionally, in line with Fagereng *et al.* (2021), Ganong *et al.* (2025), and others, I find that household liquidity moved opposite to MPCs: it declined for low-wealth and increased for the mid-wealth

[1]Hand-to-mouth (HtM) households consume all of their current income and carry little to no liquid wealth from one period to the next. Kaplan *et al.* (2014) define poor HtM households as those who live paycheck to paycheck with zero asset accumulation, whereas wealthy HtM households are those that hold significant illiquid assets, such as housing or retirement accounts, but remain highly sensitive to income shocks due to a lack of readily available cash.

group. These findings together suggest an important shift in the distribution of liquidity-constrained households.

To understand what drove this shift, I document three facts related to household saving behavior. First, I show that homeownership declined over the sample period for the working-age population from 59% to 48%. Despite lower mortgage rates during this period, ownership became harder to attain as increasing house prices translated into larger down payments.

Second, housing-cost-to-income ratios increased substantially for the low-wealth group from 0.39 to 0.49. These households, predominantly renters (around 85%), were largely hit by higher rents. On the other hand, lower mortgage rates benefited mid- and high-wealth households, who are mostly owners. Despite higher prices, mortgage payments fell after the global financial crisis (GFC).

Third, household portfolios have become more liquid. The share of liquid assets out of total assets increased across all groups. Using the Survey of Consumer Finances (SCF), I show that households have substituted away from fixed-income assets[2] toward other assets, such as liquid accounts (e.g., checking accounts). With liquid assets yielding close to zero during the sample period, the secular decline in interest rates narrowed the return gap between the fixed-income and liquid assets. The (opportunity) cost of holding liquid assets went down, resulting in an asset substitution from fixed-income towards liquid assets.

Model. I propose a theory in which the MPC evolution is driven by three secular trends related to these facts: rising house prices, larger rents, and the decline in interest rates.[3] I build a heterogeneous-agent life-cycle model with income risk, portfolio choice and housing tenure. Households choose between liquid buffers or a higher return illiquid asset. The former provides insurance against income shocks and the latter provides a savings vehicle, mainly used to buy a house. They also choose between renting and owning. Renters choose from a rental menu and pay a proportional rental cost q every period. The rental menu serves as a buffer in case households are hit by a large shock. Owners make a down payment $(1-\theta)p$ when buying and pay their mortgage in subsequent periods.

I calibrate the baseline model to match wealth, portfolio-allocation, and homeownership targets at the beginning of the sample. The model generates two groups of liquidity-constrained, high-MPC households: poor renters at the bottom and recent homeowners with large mortgages in the middle. The former are standard in heterogeneous-agent models; the latter arise from the presence of illiquid assets, as in Kaplan *et al.* (2014).

[2] Assets with fixed rates of return linked to interest rates.

[3] The price-to-rent ratio also increased over the same period.

The two groups differ in their ability to adjust housing. Renters can respond to large income losses by moving to a smaller unit, which cushions their consumption. Owners lack this margin since selling a house is costly and slow, so they cut consumption more aggressively. As a result, the model reproduces the initial MPC ordering in the data, with mid-wealth households having the highest MPC.

I use the model to quantify how higher house prices, higher rents, and lower interest rates have shaped MPCs across the wealth distribution. I take these prices as exogenous. They are linked in equilibrium, and rising house prices and rents are likely driven in part by the secular decline in interest rates, but I abstract from these links and ask what their joint change implies for household behavior. I feed the price changes from the data into the model to generate a counterfactual economy. It reproduces the three empirical facts and the shift in the MPC distribution. It accounts for the entire rise in the low-wealth MPC and about three quarters of the decline in the mid-wealth MPC. Decomposing the effect of each price, I find that house prices and rents drive most of the rise at the bottom, while house prices and interest rates drive most of the decline in the middle. I discuss each channel in turn.

First, higher house prices alone reproduce the decline in homeownership and move MPCs in opposite directions across wealth groups. Poor renters are priced out of homeownership and stop saving toward a down payment. Thus, the amount they use for consumption (rather than saving) out of every extra dollar increases, raising MPC at the bottom. On the other hand, would-be buyers must save longer for a larger down payment. This shrinks the group of recent homeowners with little liquidity left after buying, lowering the MPC at the middle.

Second, higher rents alone reproduce the rising rent burden at the bottom and tighten constraints there through two channels. Some households move down the rental ladder. Their rent payments fall, but they use up the buffer that the rental menu provides. With less room left to downsize, the next large income loss falls more heavily on their consumption. Others stay in their unit and pay more for the same housing. Rent absorbs a larger share of their income and crowds out liquid savings, leaving them with thinner buffers against income shocks. Both channels increase the share of liquidity-constrained low-wealth households and raise the MPC at the bottom.

Third, lower interest rates alone reproduce the shift toward liquid portfolios. With a narrower spread between illiquid and liquid returns, households wealthy enough to save in illiquid assets build liquid buffers instead. Would-be buyers take longer to buy a house as their savings vehicle disappears, reducing the share of constrained new owners.

Taken together, the three price changes reshape household saving at both ends of the

wealth distribution. At the bottom, higher house prices put homeownership out of reach and higher rents erode buffers, pushing more low-wealth households into liquidity constraints. In the middle, larger down payments delay home purchases and encourage saving, while lower interest rates tilt these savings toward liquid assets, leaving fewer recent buyers without a buffer. The model needs all three price changes to match the empirical facts quantitatively and qualitatively.

Fiscal policy. Finally, I ask how the effectiveness of targeted fiscal policy has changed. I measure effectiveness as the difference in the aggregate consumption response between a transfer targeted to low-wealth households and a uniform transfer of the same total size. Choosing between a targeted or uniform transfer of the same size, however, is not straightforward. The MPC is nonlinear as the effect of a transfer on consumption declines with transfer size. Thus, targeting a fixed budget on fewer households raises the amount each recipient gets and lowers their response per dollar. Targeting therefore involves a trade-off: it reaches households with higher MPCs, but at a size where those MPCs are lower.

I show that this trade-off has shifted sharply over the sample. In the baseline economy, targeting low-wealth households is less effective than a uniform transfer. The highest-MPC households are recent homeowners with large mortgages in the middle of the distribution, and a transfer restricted to the bottom misses them. This ranking reverses in the counterfactual economy. A transfer of 1% of annual income, comparable to the 2008 stimulus payments, raises output by 50% more when targeted to low-wealth households than when distributed uniformly. The same reversal holds when targeting by income, a variable much easier to observe in practice. Targeting households at the bottom of the distribution has become a powerful tool to stimulate the economy.

Related literature and contributions. This paper contributes to three strands of the literature. First, it relates to the empirical literature measuring marginal propensities to consume and their heterogeneity across households. A large body of work estimates consumption responses using tax rebates, stimulus payments, surveys, lottery prizes, and semistructural income-process methods (Johnson *et al.*, 2006; Parker *et al.*, 2013; Jappelli and Pistaferri, 2014; Misra and Surico, 2014; Blundell *et al.*, 2008; Fagereng *et al.*, 2021; Commault, 2022; Ganong *et al.*, 2025). More recent work emphasizes both the substantial dispersion of household-level MPCs and the limited ability of standard observable characteristics to explain it (Lewis *et al.*, 2024). Most of this literature estimates MPC heterogeneity within a particular episode or characterizes an average cross-sectional MPC schedule.

My contribution is to introduce a time dimension. Building on Blundell *et al.* (2008), I mea-

sure MPC over time by estimating it in rolling windows of the panel. I show that MPCs rose among low-wealth households and fell among mid-wealth households between 1999 and 2019, even as the aggregate MPC remained approximately constant. These findings imply that a cross-sectional MPC schedule estimated in one period need not remain a stable input into policy analysis.

Second, the paper contributes to work connecting consumption sensitivity to household liquidity and portfolio composition in heterogeneous-agent models. Kaplan and Violante (2014) and Kaplan *et al.* (2014) show that households may hold substantial illiquid wealth while remaining liquidity-constrained, and that a two-asset structure can generate large consumption responses to transitory transfers. Subsequent work studies how liquid wealth, illiquid wealth, indebtedness, and household-specific characteristics shape MPCs and monetary-policy transmission (Auclert *et al.*, 2018; Auclert, 2019; Auclert *et al.*, 2019, 2025; Kaplan *et al.*, 2018; Aguiar *et al.*, 2025; Crawley and Kuchler, 2023; Kaplan and Violante, 2022; Luetticke, 2021).

This paper studies how the joint distribution of liquidity and MPCs changes as households adjust their saving, portfolio composition, and housing tenure to the macroeconomic environment. I show how house prices, rents and interest rates change the distribution of liquidity-constrained households.

Finally, the paper relates to research on housing, household saving, and the heterogeneous transmission of interest rates. Existing work studies consumption responses to house prices, mortgage debt, debt-service payments, and mortgage refinancing (Berger *et al.*, 2018; Cloyne *et al.*, 2020; Flodén *et al.*, 2021; Eichenbaum *et al.*, 2022; Anenberg *et al.*, 2025; Foulis *et al.*, 2026). A complementary literature examines how house prices and down-payment requirements shape saving for home purchase, and how financial constraints and housing illiquidity influence tenure choices over the life cycle (Engelhardt, 1994; Halket and Vasudev, 2014). Homeownership also provides a hedge against fluctuations in rents (Sinai and Souleles, 2005), while monetary policy can affect both housing tenure and rental costs (Dias and Duarte, 2019, 2022; Abramson *et al.*, 2026).

My contribution is to connect rising house prices and rents, together with declining interest rates, to the evolution of MPCs across the wealth distribution. The mechanism operates through access to homeownership, housing payments, and liquid asset accumulation. Higher house prices put homeownership out of reach for low-wealth households, leaving them exposed to rising rents that tighten their financial constraints. Among mid-wealth households, larger down payments delay homeownership and encourage saving, while lower interest rates shift these savings toward liquid assets. These adjustments increase liquidity among mid-wealth households and reduce the prevalence of liquidity-

constrained behavior, generating a divergence in MPCs across the wealth distribution.

Paper structure. The remainder of the paper is organized as follows. Section 2 describes the data. Section 3 presents the empirical framework and documents the evolution of MPCs across the wealth distribution. Section 3.4 studies the evolution of liquidity-constrained households and liquid wealth. Section 4 documents the three empirical facts used to study savings behavior and motivate the model. Section 5 develops the model, while Section 6 presents the mechanism behind the MPC evolution. Section 7 presents the fiscal policy experiments. Section 8 concludes.

2 Data

My main data source is the Panel Study of Income Dynamics (PSID), a household panel survey that runs from 1968 to 2023. The PSID was conducted annually until 1997 and biennially from 1999 onward. It contains information on income and employment throughout, and comprehensive measures of consumption and wealth from 1999 onward. I describe it in section 2.1. Additionally, I complement my main analysis with data from the Survey of Consumer Finances (SCF), which I present in section 2.2.

2.1 Panel Study of Income Dynamics (PSID)

The PSID is the main dataset in this paper. The raw data require substantial harmonization as the survey methodology changes across years. I therefore use a publicly available harmonized version, the PSID-SHELF (Social, Health, and Economic Longitudinal File) of Daumler *et al.* (2025). The PSID-SHELF tracks over 8,000 families and 53,000 individuals across 42 waves from 1968 to 2021.[4] I exclude the 2021 wave from the baseline analysis as it is distorted by the COVID-19 pandemic and the subsequent government stimulus payments.[5] Additionally, I only consider data starting from 1999 as it contains comprehensive data on consumption, which is needed to estimate the MPC in Section 3. These waves also contain wealth data.

The baseline sample includes households whose head is aged 25 to 55, to focus on the working-age population, consistent with Kaplan *et al.* (2014). I exclude households that are out of the labor force (e.g., retired) and, within each wave, drop observations in the top and bottom percentile of income growth. Income is total household labor income, the sum of the labor earnings of the head and the spouse. The baseline consumption measure is total consumption excluding housing expenditure.[6]

[4] This excludes the latest PSID wave, from 2023.

[5] Appendix B.6 reports the main results including this wave.

[6] Total consumption includes food at home, food delivered, food eaten out, hospital and doctor bills,

I use net wealth to classify households into wealth groups. Net wealth includes housing and other real estate, savings, investments, vehicles, and other assets, net of mortgages, student loans, and other debt. Following the literature, I define liquid wealth as liquid assets (checking and savings accounts) minus liquid debt (credit card debt). Credit card debt, however, is reported as a separate item only in recent waves; before that, it is included in other debt. I therefore compute the share of credit card debt in other debt in the waves where both are available, and apply this share to impute liquid debt in all waves from 1999 onward.

I construct wealth groups as follows. I compute the average net wealth of each household across all the waves in which it is observed, and then split households into three equally sized groups based on this average.[7] This makes households remain in the same group throughout the sample. Appendix B.5 shows that the main results hold when I define the groups within each period.

Table 2.1 reports descriptive statistics for the whole sample and by wealth group, covering income, expenditure, wealth, liquid wealth, and demographic characteristics. As expected, income, expenditure, liquid wealth, age, homeownership, and education all increase with wealth. Importantly, low-wealth households are mostly renters (only 16 percent own their home), whereas most mid- and high-wealth households are homeowners.

Lower-wealth households also spend a larger share of their income and save less: total expenditure is 87 percent of labor income for the low-wealth group, compared with 69 and 57 percent for the mid- and high-wealth groups. Similarly, housing expenditure absorbs a larger share of income for this group (35 percent, compared with 28 and 24 percent). The low-wealth group has, on average, negative net wealth and negative liquid wealth. Finally, portfolios become less liquid as wealth increases.

2.2 Survey of Consumer Finances (SCF)

I use the Survey of Consumer Finances (SCF) to complement the PSID. The SCF is a triennial survey conducted by the Federal Reserve Board between 1989 and 2022. It collects detailed information on the balance sheets, income, and demographic characteristics of U.S. households. It provides a comprehensive view of household finances, covering both assets and liabilities, as well as the distribution of wealth across the population. I use it to

prescriptions, health insurance, mortgage payments, rent, utilities, homeowners insurance, property taxes, vehicle loans, leases and down payments, auto insurance, additional vehicle expenses, vehicle repairs, gasoline, parking, bus, taxi, other transportation expenses, education, and childcare. The baseline measure excludes the housing items: mortgage payments, rent, utilities, homeowners insurance, and property taxes.

[7] The thresholds between groups are US$14,984 (low to mid) and US$103,062 (mid to high).

Table 2.1: Descriptive Statistics by Wealth Group (PSID)

		By wealth group		
	Overall	Low-wealth	Mid-wealth	High-wealth
Income				
Labor income	60,162	32,018	53,729	94,784
Expenditure				
Total expenditure	39,647	27,799	37,091	54,072
Housing expenditure	16,336	11,351	14,898	22,767
Expenditure excl. housing	23,312	16,448	22,193	31,305
Wealth				
Net wealth	190,376	-10,495	48,603	533,263
Total assets	267,755	28,147	111,284	664,136
Total debt	77,380	38,643	62,681	130,873
Liquid wealth				
Liquid wealth	14,068	-2,132	6,019	38,340
Liquid assets	16,837	2,468	7,930	40,130
Liquid debt	2,768	4,600	1,911	1,790
Household characteristics				
Age	40	38	40	43
Homeownership (%)	51.0%	16.3%	56.5%	80.4%
College education (%)	36.3%	25.3%	31.4%	52.3%

Notes: Entries are weighted means using cross-sectional weights. Monetary variables are reported in 2024 dollars. Homeownership and college education are reported in percent. College education: the head has at least 16 years of schooling.

study how the composition of household wealth has changed over time.

All monetary variables are in 2024 dollars, as in the PSID. The baseline sample includes households whose head is aged 25 to 55. I also use a period comparable to the PSID sample, 1998 to 2019.[8] I divide households into three groups according to their net wealth, using the PSID thresholds. Appendix A.2 shows data after breaking the sample into three equally sized groups across the whole sample and three equally sized groups every year.

Table A.1 in Appendix A.2 reports descriptive statistics for the whole SCF sample and by wealth group. The qualitative patterns are consistent across the two surveys. However, I use the SCF primarily to study the composition of household portfolios. Table A.2 in Appendix A.2 decomposes total assets by wealth group. Financial assets become less liquid as wealth increases, and the primary residence is the largest asset of low- and mid-wealth households.

[8]Because the SCF is triennial, I cannot exactly replicate the PSID time frame.

Finally, I use the Consumer Expenditure Survey (CEX) to measure housing expenditure more precisely than in the PSID, in particular rent and mortgage payments. Appendix A.3 describes this survey.

3 Empirics

This section presents the empirical framework to estimate the marginal propensity to consume (MPC) over time. Section 3.1 describes the methodology of Blundell *et al.* (2008) and how I adapt it to measure the MPC over time. Section 3.2 presents the main results: the MPC evolution across the wealth distribution. Section 3.3 compares the MPC at the beginning and at the end of the sample for other groups: housing tenure, age and income. Finally, section 3.4 explores the evolution of liquidity-constrained households and shows that it is consistent with the trends for MPC.

3.1 MPC Estimation

This section presents a framework to estimate the MPC over time. A large empirical literature measures MPCs using a variety of approaches. Because of data limitations, however, most of these studies estimate the MPC at a given point in time (e.g., Fagereng *et al.* (2021), Jappelli and Pistaferri (2014), Johnson *et al.* (2006), Parker *et al.* (2013)). Measuring how the MPC evolves requires long panel data on both income and consumption, which are rarely available.

The PSID is one of the few public datasets in the United States with a panel structure, enough detail, and a long enough time span to measure the MPC over time. I build on the methodology of Blundell *et al.* (2008) to obtain these estimates. This is standard in the literature and has been used by Kaplan *et al.* (2014) and Berger *et al.* (2018), among others. However, these studies estimate the MPC for a single cross-section rather than tracking it over time. I present a simple extension that measures the MPC over time, where the idea is to estimate the MPC in rolling windows of the panel.

Blundell *et al.* (2008) estimate the response of consumption to transitory income shocks, instrumenting current income growth with future income growth. The method only requires a panel that follows the income and consumption of the same households over time. It needs at least three consecutive waves; I use four to reduce noise. I start by residualizing income and consumption on age, birth year, race, sex, family size, region, employment status, and housing and mortgage status. This income measure $y_{i,t}$ is the sum of a permanent component $z_{i,t}$, which follows a random walk, and a transitory component $\varepsilon_{i,t}$:

$$y_{i,t} = z_{i,t} + \varepsilon_{i,t}, \qquad z_{i,t} = z_{i,t-1} + \eta_{i,t}.$$

The permanent innovation $\eta_{i,t}$ and the transitory shock $\varepsilon_{i,t}$ are i.i.d. with mean zero and variances σ_η^2 and σ_ε^2, and are independent of each other. Consumption responds to the permanent and transitory components of income, and has an additional idiosyncratic component $u_{i,t}$: $c_{i,t} = \psi z_{i,t} + \phi\varepsilon_{i,t} + u_{i,t}$. The object of interest is ϕ, the marginal propensity to consume out of a transitory shock. The challenge is that neither $z_{i,t}$ nor $\varepsilon_{i,t}$ is observed in the data. To solve this, we can start by taking first differences,

$$\Delta y_{i,t} = \eta_{i,t} + \Delta\varepsilon_{i,t}, \qquad \Delta c_{i,t} = \psi\eta_{i,t} + \phi\Delta\varepsilon_{i,t} + \Delta u_{i,t}. \tag{1}$$

Following Blundell *et al.* (2008), I estimate ϕ from an IV regression of $\Delta c_{i,t}$ on $\Delta y_{i,t}$, instrumenting $\Delta y_{i,t}$ with next wave's income growth, $\Delta y_{i,t+1}$:

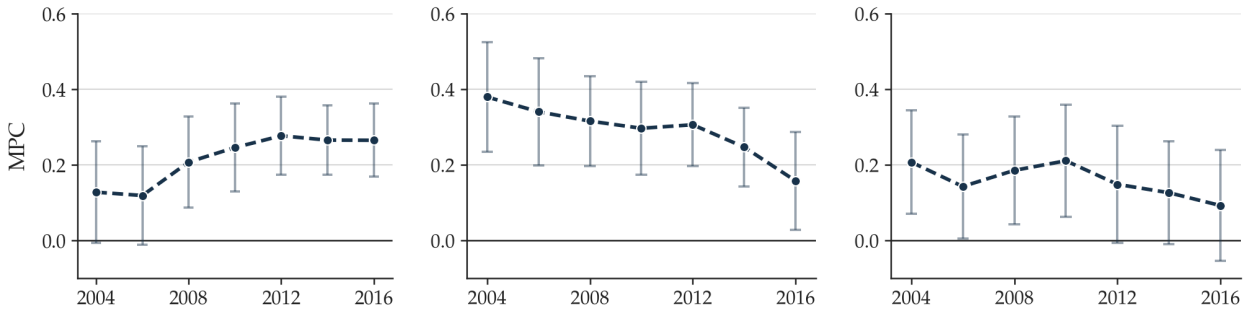
$$\phi^{IV} = \mathrm{Cov}(\Delta c_{i,t}, \Delta y_{i,t+1}) \,/\, \mathrm{Cov}(\Delta y_{i,t}, \Delta y_{i,t+1}). \tag{2}$$

Appendix B.1 shows that the instrument is valid and that $\phi^{IV} = \phi$ under the assumptions of the model, and Appendix B.2 shows that these assumptions are consistent with the data. I can use this methodology to estimate group-specific MPCs as Kaplan *et al.* (2014) and Berger *et al.* (2018) do.

Finally, to trace the evolution of the MPC over time, I use a rolling-window approach. I estimate the MPC using four consecutive PSID waves, and then move the estimation window forward one wave at a time, dropping the earliest wave and adding the next one. This procedure generates a sequence of overlapping MPC estimates that traces the evolution of household consumption responses over time, resembling a moving average.

3.2 Main Results

I present the main result in Figure 3.1. It shows the MPC evolution across the wealth distribution. The MPC of low-wealth households has increased over time, while the MPC of mid-wealth households has declined. The MPC of high-wealth households has remained low and noisy throughout.



(a) Low Wealth (b) Mid Wealth (c) High Wealth

Figure 3.1: MPC Evolution Across Wealth Groups

Notes: PSID. MPC estimated on four consecutive waves (an eight-year window), plotted at the midpoint of the window. Wealth groups: terciles of household-average net worth. 95% confidence intervals.

In contrast, I show that the aggregate MPC remained stable at around 0.2–0.25 over the sample period (Figure B.10 in Appendix B.5), consistent with other estimates in the literature (e.g., Kaplan *et al.*, 2014). Despite some cyclicality (as in Gross *et al.* (2020)), looking at this estimate alone might lead to the wrong policy conclusions.

Standard heterogeneous-agent models predict that the poorest households respond the most to unexpected income shocks. Figure 3.1 shows that this is not always the case. At the beginning of the sample, the MPC estimate for mid-wealth households was higher than that for low-wealth households. I argue that this is driven by two factors. First, there is an important presence of wealthy liquidity-constrained households (the wealthy hand-to-mouth, or HtM, households of Kaplan *et al.* (2014)). Second, low-wealth households are predominantly renters who absorb part of a large income shock by adjusting their housing.

Wealthy HtM households hold substantial illiquid wealth, mostly home equity, but little liquid wealth, and they have high MPCs. Figure B.6 in Appendix B.3 shows that they are concentrated in the mid-wealth group: 28 percent of mid-wealth households are wealthy HtM, compared with 24 percent of low-wealth households and 17 percent of high-wealth households. Many of them are recent homeowners who have used their savings for a down payment and owe a fixed mortgage payment every month. They are liquidity-constrained and respond strongly to income shocks.

The second feature concerns how households respond to unexpected income shocks. Apart from adjusting consumption, households can also change their housing spending by moving to a cheaper place when facing an income loss. This is a common response to job loss. Chetty and Szeidl (2007) show in the PSID that households rarely adjust their housing, but that unemployment shocks make them more likely to move,

especially renters, and that households that keep their housing cut their other consumption more. Other households move in with their parents or with others (Kaplan, 2012; Wiemers, 2014), or fall behind on their rent (Pattison, 2024). This margin is easier to use for renters than for owners. An owner who wants to spend less on housing has to sell the house, which is costly and takes time.[9] Hence, this adjustment margin serves as a buffer and mutes the response of low-wealth households (predominantly renters) to large income losses relative to mid-wealth households (mainly owners).

I study this margin in two ways. First, I do an event study where I compare households with large transitory income losses[10] with households whose income is stable. Figure 3.2a shows that low-wealth households adjust their housing much more often. After a large transitory loss, the share that adjusts rises by 19 percentage points, from a base of 8 percent, compared with 10 percentage points for mid-wealth households and 4 percentage points for high-wealth households. It also shows that mid-wealth households cut their non-housing consumption more than low-wealth households. Their non-housing consumption falls by 10 percent, compared with 5 percent for low-wealth households, even though the income loss of low-wealth households is larger. Table B.1 in Appendix B.4 reports all the estimates.

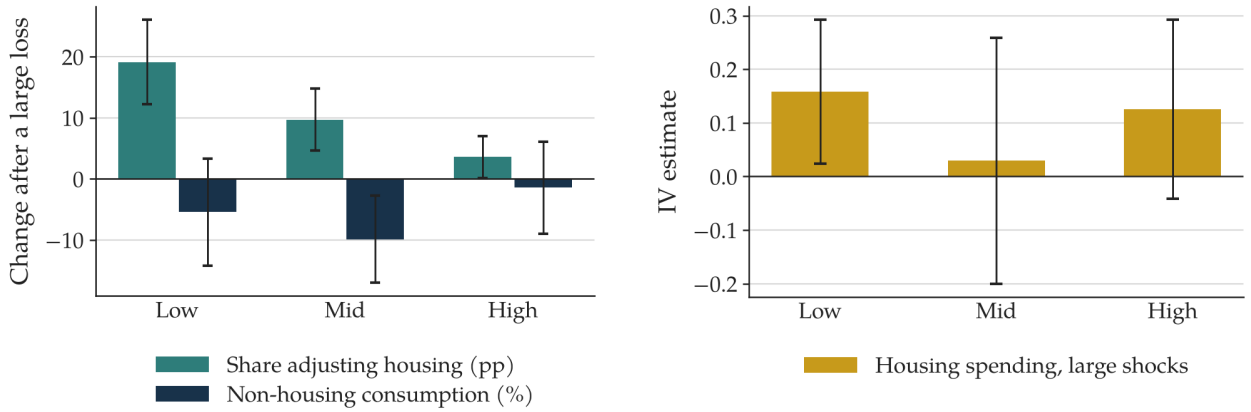
The second approach uses the estimator of Section 3.1, with housing spending and total consumption in place of non-housing consumption. I focus on the response of housing spending to large shocks.[11] Figure 3.2b shows that, for large shocks, the housing spending of low-wealth households responds strongly (0.16), while mid-wealth households do not respond (0.03). Appendix B.4 shows that these responses come from spells of unemployment. Low-wealth households absorb a large part of the income loss through housing adjustments, while mid-wealth households cut their non-housing consumption almost twice as much.

These findings suggest that housing works as insurance for low-wealth households. When their income falls sharply, many of them move to cheaper housing or move in with others, and they cut their non-housing consumption by less. Mid-wealth households have little liquid wealth and their housing is costly to adjust, so they strongly adjust their non-housing consumption after a large income loss. Both features make the MPC of mid-wealth households higher than that of low-wealth households.

[9] Owners can always extract equity through a home equity line of credit (HELOC) or a similar instrument. However, this instrument is less available for recent owners as their home equity is small.

[10] A household has a large transitory income loss when its labor income falls by more than 20 percent between two waves and recovers by more than 15 percent at the next wave.

[11] Large shocks are the quarter of observations with the largest absolute income growth in the next wave, in each year.



(a) Housing Adjustment and Non-Housing Consumption

(b) Housing Spending Response to Large Shocks

Figure 3.2: Housing Adjustment and Consumption Responses to Large Transitory Income Shocks

Notes: PSID, 1999–2019. Panel (a): change in the share of households that adjust their housing (percentage points) and in non-housing consumption (percent) after a large transitory income loss, relative to households with stable income. Panel (b): response of housing spending to the quarter of largest income shocks, estimated as in Section 3.1. Appendix B.4 gives the definitions and the other responses. 95% confidence intervals.

3.3 Additional Results

Table 3.1 compares the MPC in the initial period (1999–2007) and in the final period (2011–2019) when households are grouped by other characteristics. The first rows repeat the main result for comparison. Because the two windows do not overlap, I treat the two estimates as independent when I compute the standard error of the change. The fall in the MPC of mid-wealth households is significant at the 5 percent level. The rise for low-wealth households falls just short of the 10 percent level ($p = 0.10$). The change for high-wealth households is not significant. Importantly, the difference between the two groups is significant at the 1 percent level (Table B.3 in Appendix B.5).

There are three main takeaways from the table. First, the MPC of renters increased, from 0.17 to 0.22, while the MPC of owners declined, from 0.27 to 0.14. This is in line with Figure 3.1: most low-wealth households are renters, while most mid- and high-wealth households are owners (Table 2.1). Figure B.11 in Appendix B.5 shows the evolution over the whole sample, and Figure B.12 splits renters and owners by wealth group.

Second, the MPC of younger households (25 to 39) increased slightly, while the MPC of older households (40 to 55) declined. Again, this is consistent with the main result, as younger households are more likely to be renters with little wealth, and older households are more likely to be owners. Third, grouping households by income does not reproduce

the divergence. The changes are small, are not ordered by income, and none of them is statistically different from zero.

Table 3.1: MPC in the Initial and Final Period by Group

	Initial, 1999–2007		Final, 2011–2019		Change	
	MPC	(s.e.)	MPC	(s.e.)	Δ MPC	(s.e.)
A. Wealth						
Low wealth	0.13	(0.07)	0.26	(0.05)	+0.14	(0.08)
Mid wealth	0.38	(0.07)	0.16	(0.07)	−0.22	(0.10)
High wealth	0.21	(0.07)	0.09	(0.08)	−0.11	(0.10)
B. Housing tenure						
Renters	0.17	(0.05)	0.22	(0.05)	+0.05	(0.07)
Owners	0.27	(0.06)	0.14	(0.06)	−0.13	(0.09)
C. Age of the head						
25–39	0.15	(0.05)	0.20	(0.04)	+0.05	(0.07)
40–55	0.32	(0.06)	0.22	(0.06)	−0.10	(0.09)
D. Income						
Low income	0.27	(0.07)	0.23	(0.05)	−0.04	(0.08)
Mid income	0.21	(0.06)	0.12	(0.07)	−0.09	(0.10)
High income	0.15	(0.08)	0.21	(0.07)	+0.06	(0.11)

Notes: PSID. MPC estimated as in Section 3.1 on four waves. Standard errors of the change treat the two periods as independent samples. Wealth and income groups: terciles of household-average net worth and labor income. Tenure and age are measured in each wave.

3.4 Liquidity-Constrained Households

The MPC trends suggest there was an important shift in the distribution of liquidity-constrained households. To analyze this, I proxy for liquidity-constrained households using the hand-to-mouth (HtM) definition of Kaplan *et al.* (2014).[12] Figure B.5 in Appendix B.3 shows the evolution of HtM households and of liquid wealth for each wealth group. Panel (a) of Figure B.5 shows the share of HtM households in each wealth group, relative to the initial period. The share of HtM households declined by about a fifth in the mid-wealth group, from 36 to 30 percent, and by about a sixth in the high-wealth group, from 20 to 17 percent. In the low-wealth group, it remained stable, at about 53 percent. Fewer mid-wealth households are constrained, and more low-wealth are.

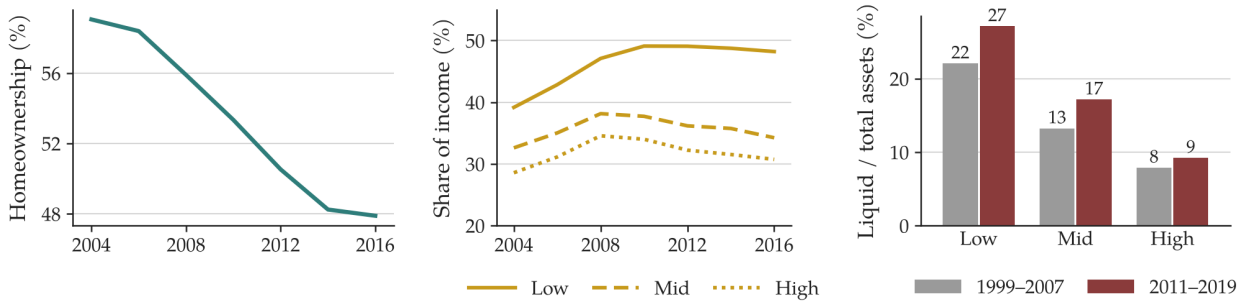
[12] Kaplan *et al.* (2014) classify a household as HtM if it consumes all of its available liquid resources each pay period, so that its consumption tracks its income. Concretely, a household is HtM if (i) its liquid wealth is non-negative but no larger than half of its income per pay period, or (ii) its liquid wealth is negative and within half a pay period's income of its credit limit. HtM households are then split by their illiquid wealth, such as housing net of mortgages and retirement accounts. *Wealthy* HtM households hold positive illiquid wealth. *Poor* HtM households hold zero or negative illiquid wealth. Appendix A.1 describes how I implement this classification in the PSID and the SCF.

Panel (b) shows the change in liquid wealth relative to income for each group. Consistent with panel (a), liquidity is a mirror image of the MPC. Liquid wealth fell by about 5 percent of income for low-wealth households, and increased by about 4 percent of income for mid- and high-wealth households.

Taken together, the evidence of this section shows that the divergence in MPCs reflects a change in who is liquidity-constrained. At the bottom of the distribution, liquid wealth fell, and more households are poor HtM, without a house or other illiquid wealth to fall back on (Appendix B.3). In the middle, liquid wealth rose, and fewer households are HtM. Housing plays a central role in both groups. It is the main asset of mid-wealth households, and it is the margin that low-wealth renters use to absorb large income losses. This suggests looking at how housing and household portfolios changed over the period. The next section documents three facts: homeownership declined, housing costs rose at the bottom of the distribution, and portfolios became more liquid.

4 Stylized Facts

This section examines household balance sheets and housing costs to understand the forces behind the changes in MPCs and liquidity documented in Section 3. I document three stylized facts related to savings behavior. First, Section 4.1 shows that homeownership declined among the working-age population, as did the likelihood that renters become owners. Second, Section 4.2 documents that housing costs relative to income increased, especially for low-wealth households, driven by rising rents. Third, Section 4.3 shows that household portfolios became more liquid, as households shifted away from fixed-income assets.



(a) Homeownership Rate (b) Housing Costs-to-Income (c) Liquid Asset Share

Figure 4.1: Homeownership, Housing Costs, and Portfolio Liquidity

Notes: PSID. Panels (a) and (b): four-wave windows plotted at their midpoint. Panel (b): rent for renters and mortgage payments for owners, relative to income, averaged across households. Panel (c): liquid assets relative to total assets, averaged across households, in the initial (1999–2007) and final (2011–2019) periods. Wealth groups: terciles of household-average net worth.

4.1 Homeownership

This section shows that homeownership declined among the working-age population over the sample period. The homeownership rate declined steadily, from about 59 to 48 percent (panel (a) of Figure 4.1).[13] The panel structure of the PSID also allows me to study transitions into homeownership. Figure C.3 in Appendix C.2 reports the likelihood that a renter becomes an owner within the next three waves. It gives a similar picture: the likelihood fell from almost 60 percent to about 40 percent, so fewer renters became owners over the sample period.

Figure C.4 in Appendix C.2 shows the homeownership rate for each wealth group. Two patterns stand out. First, most low-wealth households are renters, while most mid- and high-wealth households are owners. Second, homeownership declined in all groups. Since mid- and high-wealth households are mainly owners, the decline is more pronounced for them.

These facts are consistent with the substantial increase in house prices and in the price-to-rent ratio over the sample period (Figure C.6 in Appendix C.2). House prices roughly doubled between 1999 and 2019 and, although they fell by about 20 percent after the GFC, they never returned to their late-1990s level. Rents also increased (Section 4.2), but by less, so the price-to-rent ratio ended the sample about 25 percent above its 1999 level. Higher prices raise the down payment required to buy a house, which makes homeownership harder to reach even when mortgage rates are low.

4.2 Housing Costs

I now show that housing costs relative to income increased, especially for low-wealth households. Panel (b) of Figure 4.1 presents total housing costs relative to income for each wealth group. Housing costs consist of rent for renters and mortgage payments for owners. I compute the ratio for each household and then average across households. Because the average is sensitive to households with very large ratios, I also report the median in Figure C.8 in Appendix C.3.

There are two main patterns. First, low-wealth households spend a larger share of their income on housing. Second, this share increased the most for them, from about 39 to 49 percent of income. For mid- and high-wealth households, housing costs relative to

[13] Recall that each point averages four waves (an eight-year window). Figure C.5 in Appendix C.2 reports the homeownership rate wave by wave, and it tells a consistent story: homeownership peaks in the early 2000s at about 60 percent, falls sharply after the GFC, reaches a trough of about 47 percent in 2015, and recovers slightly in the last waves. This increase continued after the COVID-19 pandemic. Although the homeownership rate has not returned to its pre-GFC level, it has kept rising since the pandemic, at least as of this writing.

income increased until the GFC and have declined since then.

Figure C.7 in Appendix C.3 splits housing costs by tenure. The large increase for low-wealth households is driven by substantially higher rents, as rent-to-income for low-wealth renters rose from about 39 to 50 percent. Rent-to-income also increased, though by less, for mid- and high-wealth renters. In contrast, owners in all wealth groups follow the same pattern: mortgage costs rose until the GFC and declined afterward, as mortgage rates fell.

Appendix C.3 confirms these patterns with additional data. Figure C.9 shows that national rents for new tenants rose steadily, by about 60 percent between 1999 and 2019, and, unlike house prices, did not fall after the GFC. Figure C.10 reports rent-to-income and mortgage payments-to-income from the CEX, which measures housing payments more precisely than the PSID. The two series diverge sharply after 2008: by the end of the sample, rent-to-income is about 20 percent above its initial level, while mortgage payments-to-income is about a third below it. Both sources therefore tell a consistent story: rent burdens increased, while mortgage burdens declined after the GFC.

4.3 Portfolio Liquidity

Finally, I show that household portfolios became more liquid across the wealth distribution. I focus on the asset side of the portfolio. Panel (c) of Figure 4.1 shows the average share of liquid assets in total assets in the PSID, in the initial and in the final period. The share declines with wealth: in the initial period, liquid assets represent about 22 percent of total assets for low-wealth households, 13 percent for mid-wealth households, and 8 percent for high-wealth households. Moreover, the liquid share increased in all groups, by about 5 percentage points for low-wealth households, 4 percentage points for mid-wealth households, and 1 percentage point for high-wealth households.

The PSID, however, does not provide much detail on the composition of household portfolios. I therefore turn to the SCF and focus on financial assets. Table C.1 in Appendix C.4 shows the share of the four main components of financial assets in 1998 and 2019: liquid assets, fixed-income assets, retirement assets, and equities.[14] Between 1998 and 2019, the share of liquid assets increased by about 14 percentage points and the share of retirement assets by about 3 percentage points, while the shares of fixed-income assets and equities declined by about 10 and 7 percentage points, respectively. Note that these shares are

[14]The categories follow the decomposition in Table A.2 in Appendix A.2. Liquid assets include checking and savings accounts, as well as money market, call, and prepaid card accounts. Fixed-income assets include certificates of deposit, directly held bonds, savings bonds, and the cash value of whole-life insurance. Retirement assets include both equity and non-equity holdings in retirement accounts. Equities include pooled investment funds, directly held stocks, other managed assets, and other financial assets.

measured out of financial assets, while panel (c) of Figure 4.1 measures the liquid share out of total assets. The same substitution appears when shares are measured relative to total assets, which also include housing (Figure C.11 in Appendix C.4).

The decline in fixed-income assets is particularly informative. By definition, fixed-income assets pay returns linked to interest rates, so they became less attractive as interest rates declined over the sample period. With liquid assets yielding close to zero, falling fixed-income returns narrowed the return gap between the two, consequently reducing the (opportunity) cost of liquidity. The share of every fixed-income asset in financial assets fell substantially over the sample period (Table C.1).

Taken together, these facts are informative about households' saving behavior and, ultimately, about MPCs. Low-wealth households, who are mostly renters, faced rising rents and fewer opportunities to become owners. Mid- and high-wealth households, who are mostly owners, faced lower mortgage costs after the GFC and shifted their portfolios toward liquid assets as interest rates declined. Behind these household-level patterns are three aggregate price changes: house prices roughly doubled, rents rose steadily, and interest rates declined. I propose a model to explore the role of these prices in the MPC in section 5.

5 Model

This section develops a quantitative life-cycle model of housing tenure and portfolio choice. The goal is to understand and quantify the divergence in MPCs documented in Section 3, using the findings of Section 4. I use the model to ask how much of that divergence can be accounted for by the changes in house prices, rents and interest rates between the early 2000s and the late 2010s, and through which channels. I abstract away from general equilibrium forces and take prices as exogenous. This keeps the model simple enough to show how prices change household balance sheets and, through them, who are the high-MPC households.

The model produces two groups of liquidity-constrained, high-MPC households. The first group is formed by poor renters at the bottom of the wealth distribution, who hold almost no wealth of any kind. The second group is formed by households in the middle of the distribution who hold most of their wealth in illiquid assets. They are recent homeowners, who have turned their savings into a down payment and owe a fixed mortgage payment. These are the wealthy HtM households of Kaplan *et al.* (2014).

I first describe the environment in Section 5.1. Then, I turn to the household's problem in Section 5.2. Section 5.3 describes how I measure the MPC in the model. Section 5.4 shows

the model calibration, and Section 5.5 analyzes the baseline economy.

5.1 Environment

I start by describing the model environment. This is a life-cycle model with income risk, housing tenure and portfolio choice. I first describe the life-cycle component. Then I expand on the income process and household preferences. Finally, I introduce the housing and financial sectors.

Demography. Time is discrete and annual, and age is denoted by t. Households enter the economy at age 25 as renters with no wealth, work until 64 and retire at 65. As is standard in life-cycle models, the probability of dying rises with age after retirement. I use the Gompertz form

$$\pi_t = 1 - m_0\, e^{m_1(t-65)}, \qquad 65 \leq t < 85, \tag{3}$$

where π_t is the probability of surviving to the next year, m_0 is the probability of dying at 65 and m_1 is the rate at which it rises with age. The probability of dying before 65 is negligible, and no household lives beyond 85.

Income. During working life, log income is the sum of a deterministic age profile, a persistent component and a transitory shock:

$$\log y_t = \log g_t + z_t + \log(1 + \varepsilon_t), \tag{4}$$

$$z_{t+1} = \rho z_t + \eta_{t+1}, \qquad \varepsilon_t \sim U[-\bar{\varepsilon}, \bar{\varepsilon}] \text{ i.i.d.} \tag{5}$$

The persistent innovation η is normally distributed with mean zero and variance σ_η^2, and $\sigma_z = \sigma_\eta/\sqrt{1-\rho^2}$ is the stationary standard deviation of z. The persistent component is not a random walk, as in Section 3. However, as described in Section 5.4, the income process is very persistent in the data (ρ close to 1). At retirement, income becomes a pension that rises less than proportionally with the last persistent state, $\log y_t = \log \mu_{ss} + \lambda z_{64}$. This is a reduced form of the progressive Social Security benefit formula (Kaplan and Violante, 2014; De Nardi *et al.*, 2010).

Preferences. Households value non-housing consumption c and housing services h. A renter's period utility is $u(c) + v_R(h)$, and an owner's is $u(c) + v_O(h) + \zeta$, with

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}, \qquad v_k(h) = \omega\, \frac{(h - \bar{h}_k)^{1-\gamma_h}}{1-\gamma_h}, \quad k \in \{R, O\}. \tag{6}$$

Housing utility has a Stone–Geary form. The floor $\bar{h}_k$ makes housing demand inelastic for poor households, so that housing takes a larger share of their budget (Albouy *et al.*, 2016). Utility is separable between consumption and housing. This implies that housing affects the marginal utility of consumption only through the budget constraint, which is the main channel of the paper. The parameter ζ is a utility premium from owning, which stands in for the non-financial benefits of ownership. Households discount the future at rate β and value the wealth they leave at death through a warm-glow bequest motive:

$$B(w) = \nu \frac{(w + \bar{w})^{1-\gamma}}{1-\gamma}, \tag{7}$$

where w is the estate: the household's liquid and illiquid savings plus its home equity, net of the cost of selling the house.[15]

Housing. A household either rents or owns. Renters choose among a small menu of options. They can rent one of two units, a small one and a standard one, of size $h_1 < h_2$ at an annual rent qh_j, or live with relatives or others without paying rent, in a smaller space $h_0 < h_1$. A renter pays a small moving cost F when it changes units. The menu lets renters adjust their housing costs after an income shock, by moving to a cheaper unit or in with others. The rental menu acts as a buffer against large shocks, consistent with the data (Section 3.2).

Owners occupy a house of value pH. A buyer pays a down payment and a closing fee, $(1-\theta)\,pH + \kappa$, and finances the rest with a fixed-rate mortgage of T years at rate r_m. The annual mortgage payment is

$$M = \frac{r_m\, \theta pH}{1 - (1 + r_m)^{-T}}, \tag{8}$$

and the outstanding balance after n payments, d_n, follows from the usual amortization schedule. Owners pay a maintenance cost $\chi\, pH$ every year. An owner can sell the house at any time and receives its home equity, $(1-\delta)\,pH - d_n$, net of the selling cost δ and the fee κ. In the baseline, owners cannot refinance their mortgage or extract equity without selling. Appendix D.13 studies a version of the model in which they can.

Financial assets. Households hold a liquid asset a with gross return R and an illiquid asset $b \geq 0$ with a higher return $R_b > R$. The illiquid asset stands for the fixed-income assets defined in Section 2. These are certificates of deposit, savings bonds and other bonds, and the cash value of life insurance. Following Kaplan and Violante (2014), it is costly to extract money out of the illiquid asset. Households can always add at no cost,

[15] The bequest motive slows down the decumulation of wealth after retirement (De Nardi, 2004; Lockwood, 2018). It matters little for working-age households, for whom death is a distant event.

but reducing the balance from b to $b' < b$ costs $k_0 + k_1(b - b')$. k_0 is a fixed cost and k_1 is a proportional cost, such as an early-withdrawal penalty. A buyer can use its illiquid savings toward the down payment without paying the withdrawal cost.[16] The illiquid asset is therefore a good vehicle to save toward a house and a poor buffer against income shocks.

Households can also borrow a small amount through unsecured credit, such as a credit card. Liquid wealth can be negative down to a limit $\underline{a}$, and debt pays a gross rate $R_d > R$. Note that $R_d > R_b$ is also needed. Otherwise, there would be an arbitrage opportunity and households would borrow as much as possible to invest in the illiquid asset. The return on liquid wealth is therefore $R(a) = R$ if $a \geq 0$ and $R(a) = R_d$ if $a < 0$.

5.2 The Household's Problem

Households start the year in state $\Omega \equiv (z, \varepsilon, a, b, h, n)$. It is composed of persistent income z, the transitory shock ε, liquid wealth a, illiquid wealth b, housing h (a rental option or an owned house) and, for an owner, the number of years since purchase n, which determines its outstanding mortgage. The illiquid balance b includes the return earned over the previous year: a household that keeps b' in the illiquid asset this year starts next year with $R_b b'$. I can reduce the number of states without loss of generality. First, liquid wealth and current income enter the problem only through their sum, cash on hand $x = R(a)\, a + y$, because both are funds available to spend during the year. The transitory shock is i.i.d., so once it is included in x it carries no information about the future (as opposed to the persistent component z). Second, the mortgage balance is a deterministic function of the years since purchase. Housing and the mortgage can therefore be described by a single housing state s, which is either a rental option j or an owned house of vintage n. The state can therefore be reduced to (z, b, x, s).

Each year the household first chooses its housing for the year, s', and then its consumption c, liquid saving a' and new illiquid balance b'. Moving from s to s' requires a housing payment $\mathcal{C}(s, s')$:

$$\mathcal{C}(s, s') = \begin{cases} q h_{j'} \mathbf{1}\{j' \neq 0\} + F\, \mathbf{1}\{j' \neq j\} & \text{renter in } j \text{ moves to } j', \\ (1 - \theta)\, pH + \kappa + \chi pH & \text{renter buys a house}, \\ M\, \mathbf{1}\{n < T\} + \chi pH & \text{owner of vintage } n \text{ stays}, \\ -\big[(1 - \delta)\, pH - d_n\big] + \kappa + q h_{j'} & \text{owner of vintage } n \text{ sells and rents } j'. \end{cases} \tag{9}$$

Let $\Psi(b, b') = \mathbf{1}\{b' < b\}\, [k_1(b - b') + k_0]$ be the cost of reducing the illiquid balance from b to

[16] Appendix D.13 shows the model without this assumption.

b'. A renter who buys a house pays no withdrawal cost on the part of its illiquid savings that goes to the down payment and the fee. For each housing choice s', the household solves

$$\mathcal{V}_t(z, b, x, s, s') = \max_{c,\, a',\, b'} u(c) + v(s') + \beta\Big\{\pi_t\, \mathbb{E}\big[V_{t+1}(z', R_b b', x', s') \,\big|\, z\big] + (1 - \pi_t)\, B(w')\Big\} \tag{10}$$

$$\begin{aligned} \text{s.t.} \quad & c + a' + b' + \Psi(b,\, b') = x + b - \mathcal{C}(s, s'), \\ & x' = R(a')\, a' + y_{t+1}, \\ & w' = R(a')\, a' + R_b b' + \mathbf{1}\{\text{owns}\}\big[(1 - \delta)\, pH - d_{n'}\big], \\ & a' \geq -\underline{a}, \qquad b' \geq 0, \end{aligned}$$

where $v(s')$ is the utility of housing s', including the ownership premium for an owner, w' is the estate that the household leaves if it dies at the end of the year, and n' is the number of years since purchase of the house in s'.

The choice between renting and owning is subject to idiosyncratic taste shocks. They capture reasons to move that the model does not describe, such as a job relocation or a change in family size. Because of these shocks, the household does not choose its tenure deterministically. Given its state, it owns with some probability, and this probability rises smoothly with the value of owning relative to renting. Let $\mathcal{V}_t^R$ be the value of the best renting option and $\mathcal{V}_t^O$ the value of the best owning option. Each year the household draws i.i.d. type-I extreme value shocks for the two options, with scale σ. The probability that it owns and its expected value are then

$$\Pr(\text{own}) = \frac{1}{1 + e^{(\mathcal{V}_t^R - \mathcal{V}_t^O)/\sigma}}, \qquad V_t(z, b, x, s) = \sigma \log\left(e^{\mathcal{V}_t^R/\sigma} + e^{\mathcal{V}_t^O/\sigma}\right). \tag{11}$$

These shocks generate realistic flows between renting and owning. Appendix D.13 shows that the results are the same when each rental unit also has its own taste shock. Additionally, they smooth the value function, which makes the endogenous grid method easier to apply (Rust, 1987; Iskhakov *et al.*, 2017). I solve the model by backward induction, with the endogenous grid method (Carroll, 2006) and an upper envelope step (Iskhakov *et al.*, 2017). I provide full details in Appendix D.1.

5.3 MPC in the Model

As in Section 3, I estimate the MPC as the response of consumption to transitory income shocks, building on Blundell *et al.* (2008). Within each wealth group, the MPC is measured

as in (2)

$$\phi^{IV} = \frac{\text{Cov}(\Delta c_{i,t}, \Delta y_{i,t+1})}{\text{Cov}(\Delta y_{i,t}, \Delta y_{i,t+1})}, \tag{12}$$

where $\Delta c_{i,t}$ and $\Delta y_{i,t}$ are the changes in residual log consumption and log income between two waves. Future income growth isolates the transitory part of current income growth. A transitory shock today lowers income growth tomorrow, while a persistent shock does not. Appendix D.5 tests the restrictions behind the estimator in the simulated panel, and shows that it recovers the true response of consumption to transitory shocks, which the model observes directly.

I compute this object in the model exactly as in the data. From each steady state, I simulate a panel of 40,000 households and draw, for each of them, a window of four biennial waves, as in the PSID.[17] As in the data, I use changes between waves two years apart, although the model would allow annual changes.[18] As Blundell *et al.* (2008) do, I residualize log income and log non-housing consumption before estimating (12).[19] I estimate (12) by wealth tercile, ranking households by their average net worth over the window, as in the data.[20] Throughout the paper, the MPC of the model refers to this estimate.

I consider an alternative MPC measure. This is the response to a small, unexpected cash transfer τ: the share of it that a household spends within the year,

$$\text{MPC}_i^{\tau} = \frac{c_i(x_i + \tau) - c_i(x_i)}{\tau}. \tag{13}$$

I call it the *transfer MPC*. Households are allowed to re-optimize their housing choice after the transfer. This is the relevant measure for the fiscal transfer experiments in Section 7. It differs from the baseline MPC for two reasons. First, consumption is a nonlinear function of cash on hand, so households respond differently to a small transfer than to the large and often negative transitory shocks that identify (12). Second, a household hit by a large negative shock can adjust its housing, by moving to a cheaper unit or in with others. This buffer mutes the consumption response against a large negative shock. A small transfer does not trigger this adjustment. Appendix D.6 compares the two measures. The baseline MPC measures the response to the income shocks that households actually face, which is what the data identify. The transfer MPC measures the response to a policy that does not exist in the data, a small unexpected transfer. In the model, the two measures only differ

[17] Income in the PSID is measured with error, so I add classical measurement error to log income in the simulated panel. This attenuates the estimate of the MPC but does not affect the results. I use a conservative measure (a variance of 0.02 of log income) relative to the PSID Validation Study (Bound *et al.*, 1994).

[18] The MPC estimates are practically the same with annual changes (Appendix D.5).

[19] In the model, age is the only control.

[20] See Appendix D.10 for alternative groupings.

in magnitude but change in the same direction over time, which supports the evidence of Section 3.

5.4 Calibration

A model period is a year, and money amounts are in units of mean annual labor income. Table 5.1 lists the parameters along with the target or sources. I set most of them from the literature or directly from the data, and I choose eight of them jointly so that the model matches eight moments of the PSID and the SCF in 1999–2007, the initial period of the empirical analysis.

Targeted moments. I calibrate eight parameters by targeting eight moments. The ownership premium determines the homeownership rate, and the discount factor the holdings of liquid assets and fixed income. The space of a household living with others sets how many households choose this option. The rent per unit and the renters' housing floor determine the housing costs of low- and mid-wealth renters. The scale of the taste shock controls how often owners sell, and the bequest weight the decline of wealth in old age. Finally, the proportional withdrawal cost determines how much fixed income high-wealth households hold, because a larger cost makes the illiquid asset less attractive. Table D.1 in Appendix D.2 shows that the model matches the eight moments closely.

A discount factor of 0.852 is low compared with models calibrated to total wealth. However, it is common for models that match liquid wealth moments. With a larger β, households would accumulate much more wealth. Because the illiquid asset pays a higher return, they would keep most of the additional saving in it, and they would hold almost no liquid buffer. Matching the distribution of liquid wealth requires low effective discount factors (Carroll *et al.*, 2017). Moreover, liquidity-constrained households are associated with impatience (Aguiar *et al.*, 2025).

Income. I estimate the income process on labor earnings in the PSID by minimum distance on the variance of residual earnings by age and its autocovariances. The persistent component is very close to a random walk: its annual persistence is $\rho = 0.964$. Households therefore treat persistent shocks almost as if they were permanent. This supports the random walk assumption behind the method of Blundell *et al.* (2008). Appendix B.2 shows that the variance of long differences of income in the PSID rises linearly with the horizon, consistent with a random walk (Figure B.3).

On the other hand, I account for unemployment insurance to measure transitory risk. Most of the transitory volatility comes from unemployment spells, which are accompanied by this type of insurance. Since the model does not have unemployment, and thus unemployment insurance, I interpret income in the model as income after unemployment

Table 5.1: Parameters

Parameter	Description	Value	Target or source
Chosen jointly to match moments of the data			
ζ	Ownership premium	1.48	Homeownership, 56.5% (PSID)
β	Discount factor	0.852	Liquid assets and fixed income, 0.44 (SCF)
h_0	Space when living with others	0.38	Neither own nor pay rent, 6.0% (PSID)
q	Rent per unit of housing	0.133	Housing cost / income, low-wealth renters, 0.27 (PSID)
$\bar{h}_R$	Housing floor, renters	0.033	Housing cost / income, mid-wealth renters, 0.21 (PSID)
σ	Scale of the tenure taste shock	0.38	Owners who sell each year, 3.5% (PSID)
ν	Bequest weight	11.0	Net worth at 75+ relative to 65–74, 0.77 (PSID)
k_1	Withdrawal cost, proportional	0.33	Fixed income, top third of net worth, 0.59 (SCF)
Set from data or the literature			
$\rho,\ \sigma_z$	Persistent income	0.964, 0.78	PSID labor earnings, minimum distance
$\bar{\varepsilon}$	Transitory income, half-width	0.298	Standard deviation 0.17 (see text)
g_t	Age profile of earnings	cubic	PSID earnings by age
$\mu_{ss},\ \lambda$	Pension level, elasticity to z_{64}	0.78, 0.6	Social Security (Kaplan and Violante, 2014)
$m_0,\ m_1$	Mortality after 65	0.013, 0.095	SSA period life table
γ	Curvature of utility from consumption	2	Berger *et al.* (2018)
pH	House value	3.3	Rent-to-price ratio (Davis *et al.*, 2008)
$\theta,\ T,\ r_m$	Loan-to-value, mortgage term, rate	0.8, 30, 4.5%	Conventional fixed-rate mortgage, real rate
$\chi,\ \delta,\ \kappa$	Maintenance, selling cost, closing fee	4.5%, 8%, 0.15	Maintenance and property taxes; broker fees; closing costs
$R,\ R_b$	Gross returns, liquid and illiquid	1.00, 1.06	See text
k_0	Withdrawal cost, fixed	0.015	Kaplan and Violante (2014)
$\underline{a},\ R_d$	Credit limit, rate on debt	0.083, 1.12	One month of income; credit card rates (Federal Reserve G.19)
$h_1,\ h_2$	Rental units	0.45, 1	Rents of the cheapest quarter of units (PSID)
$\omega,\ \gamma_h,\ \bar{h}_O$	Housing utility	0.55, 4, 0.125	Normalizations
$F,\ \bar{w}$	Moving cost; bequest shifter	0.02; 1	Small; bequests are a luxury good (De Nardi, 2004)

Notes: Money amounts in units of mean annual labor income. Returns and rates are real. The eight parameters in the upper panel are chosen jointly; each is listed next to the moment that mainly identifies it.

insurance. This measure gives a standard deviation of 0.17 (Appendix D.4). Without trimming the largest income changes, labor earnings give a standard deviation of 0.32. Appendix D.13 shows how the results change with this measure.

Financial assets. The fixed withdrawal cost is 1.5% of mean income. This is close to the fixed cost of Kaplan and Violante (2014), \$1,000 against mean labor income of about \$53,000 in the 2001 SCF. With a fixed cost of this size, households do not use their illiquid

savings to absorb ordinary income fluctuations. This is what makes the asset illiquid for the MPC. The illiquid asset earns a premium of 6% over the liquid asset in the baseline economy, about the return on long-term saving instruments relative to deposits, and the liquid asset earns no real return. Households can borrow up to one month of mean income at a real rate of 12%, consistent with the average interest rate on credit card accounts between 1999 and 2007 (Federal Reserve, G.19).

Housing. The value of the house is 3.3 years of mean income. The ratio of the rent per unit of housing to the price per unit is 4.4%, within the range estimated by Davis *et al.* (2008) for this period. A buyer needs 0.81 years of income for the down payment and fees, and then pays 0.31 a year in mortgage payments and maintenance, against 0.13 to rent the standard unit. The size of the cheapest rental unit is set from the dispersion of rents in the PSID: renters in the cheapest quarter of the rental market, relative to the median in their state and year, pay 0.40–0.48 times the rent of the others.

5.5 The Baseline Economy

Figure 5.1 shows some descriptive statistics for the baseline economy by wealth group. Low-wealth households are mostly renters with almost no wealth, and many of them live in the cheapest rental units or with others. Mid-wealth households are mostly owners, many of them recent buyers. Their home equity is small relative to their income, about half a year of mean income.[21] However, this represents most of their net worth, and they hold little liquid wealth. High-wealth households are owners with substantial home equity, liquid wealth and fixed-income savings. Despite being untargeted, the model reproduces the homeownership rate of each group, and the ordering of net worth across groups. Home equity is the main asset of mid- and high-wealth households, as in the SCF. Appendix D.3 shows that the model also reproduces how homeownership, liquid wealth and net worth evolve with age.

Figure 5.2 compares the MPC of the model with the one estimated in data in the initial period.[22] Although MPCs are not a calibration target, the model closely matches both their magnitude and their ordering in the data. It is highest for mid-wealth households (0.38 in the model and 0.38 in the data) and lower for low-wealth households (0.20 in the model and 0.13 in the data) and high-wealth households (0.22 and 0.21). The estimate is also close to the true elasticity of consumption to transitory income shocks, which the

[21] Allowing them to borrow against their home would not change this much, because their home equity is small: most of them bought their house recently with a down payment of 20 percent, and selling costs absorb 8 percent of the value of the house. Appendix D.13 studies a version of the model in which owners can extract home equity at any time by paying a fee, as with a cash-out refinancing or a home equity line of credit. The MPC of mid-wealth households remains the highest, and the divergence in MPCs survives.

[22] The initial period is the window of four PSID waves centered on 2004, which covers 1999–2007.

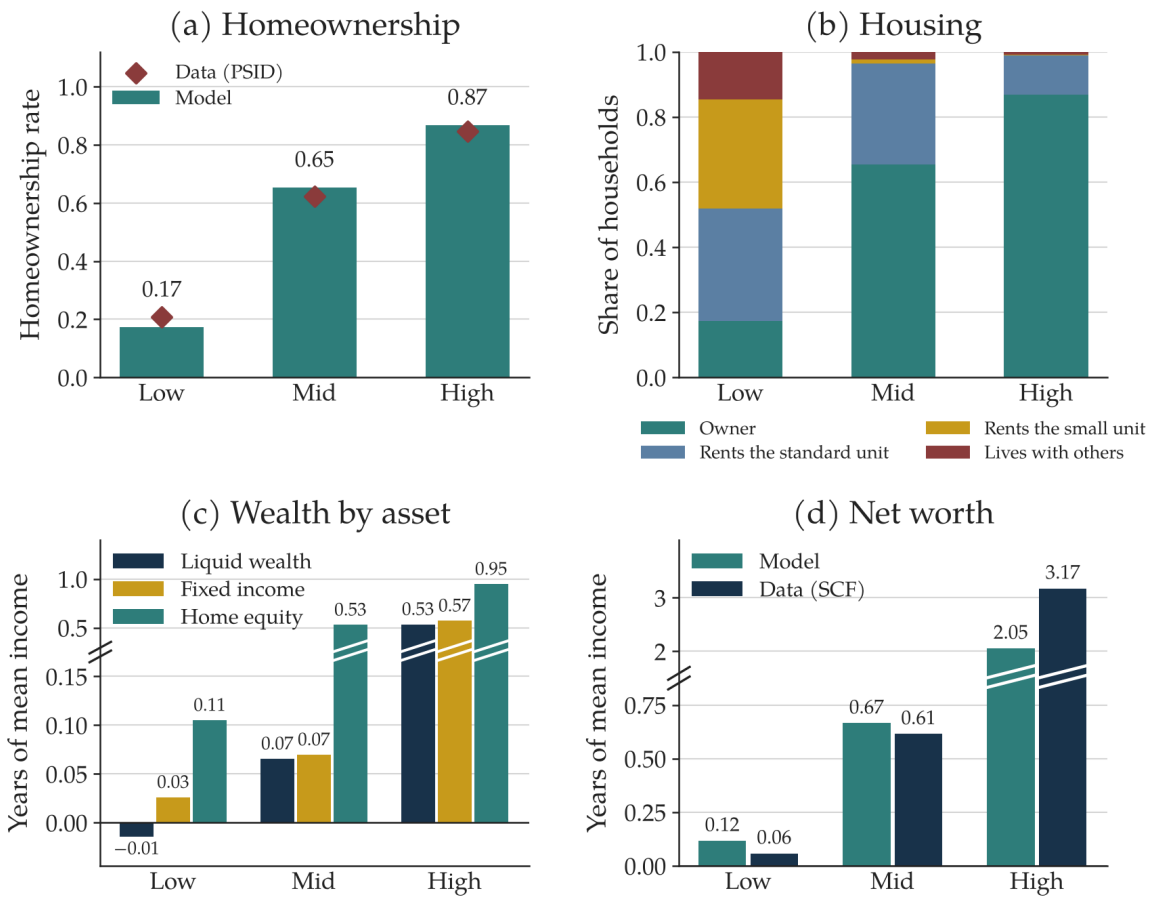


Figure 5.1: Housing and Wealth in the Baseline Economy

Notes: Baseline economy, heads aged 25–55. (a) Homeownership in the model (bars) and in the PSID, 1999–2007 (diamonds). (b) Housing of each group. (c) Liquid wealth, fixed income and home equity, in years of mean income. (d) Net worth in the model and in the SCF, 1998–2007. Panels (c) and (d) have a break in the vertical axis. Wealth groups: terciles of household-average net worth.

model observes directly: 0.26, 0.43 and 0.30 for low-, mid- and high-wealth households (Figure D.3 in Appendix D.5).

Mid-wealth households have the highest MPC because many of them hold most of their wealth in illiquid form. This group contains a large share of recent homeowners that used their savings for the down payment and owe a fixed mortgage payment. Despite being wealthy, they respond to an income shock by adjusting their consumption. Table D.1 in Appendix D.2 shows the share of HtM households by wealth group.

On the other hand, low-wealth households hold almost no wealth, but their MPC is lower than that of mid-wealth households. As in the data, housing plays a key role. Panel (a) of Figure D.4 provides more color on this. Holding housing fixed, the MPC is very similar for both groups. Low-wealth households are mainly renters that respond to large losses

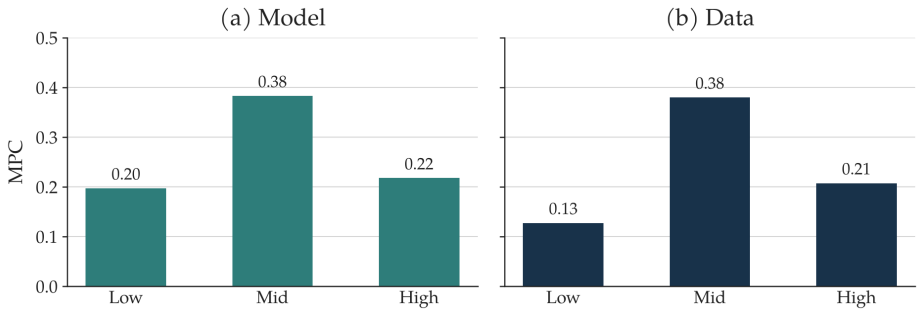


Figure 5.2: MPC by Wealth Group, Baseline Economy and Initial Period

Notes: (a) MPC estimated on the panel simulated from the baseline economy, as in Section 3.1. Figure D.3 compares it with the true elasticity of consumption to transitory shocks. (b) PSID, 1999–2007. Wealth groups: terciles of household-average net worth.

by moving down the rental ladder (Figure D.5). This adjustment mutes their consumption response to the large income losses that enter the MPC estimate. High-MPC mid-wealth households are predominantly new owners who cannot adjust their housing cheaply.[23] As a consequence, they respond more aggressively to income shocks.

Overall, this model reproduces the main features of the economy presented in Sections 3 and 4. It produces two groups of constrained households, one at the bottom composed of renters and one at the middle that contains recent homeowners. Renters at the bottom can use the rental menu as a buffer against larger shocks. I now turn to study how house prices, rents and interest rates affect the distribution of liquidity-constrained high-MPC households.

6 Mechanism

The model of Section 5 presents the baseline economy calibrated to household balance sheets and housing in the initial periods of the sample. Despite not targeting the MPC, it reproduces the ordering and magnitude of the data in that period. As in Section 3, the MPC is highest for mid-wealth households, who hold most of their wealth in illiquid form, and lower for low-wealth households, who can use the rental menu as insurance. In this section I study and quantify the effect of higher prices, larger rents and the decline of interest rates on the MPC across the wealth distribution.

I compare two steady states: the baseline economy, which represents the early 2000s, and a counterfactual economy with the prices of the late 2010s. I abstract away from

[23] Selling costs absorb 8 percent of the value of the house, and a seller who wants to own again later pays the closing fee and a new down payment. Borrowing against the house is not an alternative for most mid-wealth owners, because their home equity is small (footnote 21).

equilibrium analysis and I take all three prices as exogenous. These are tightly connected in equilibrium, and part of the rise in house prices and rents is itself a consequence of the fall in interest rates. I remain agnostic about these common drivers and ask only what their joint change implies for household behavior. Section 6.3 discusses how the analysis could be extended to general equilibrium.

The section has three parts. Section 6.1 introduces the three changes together and shows that the model reproduces the divergence in MPCs, the evolution of liquid wealth of Section 3.4 and the three stylized facts of Section 4. Section 6.2 introduces the changes one at a time and in pairs, shows that all three are needed, and explains how each of them moves the MPC. Section 6.3 discusses the general equilibrium links between the three prices.

6.1 Main Experiment

I change three prices between the baseline economy and the counterfactual economy. First, house prices rise by 28 percent. Second, rents per unit of housing rise by 22 percent. Both are measured relative to the mean income of the PSID sample, between 1999–2007 and 2011–2019 (Appendix C shows the aggregate trends in house prices and rents). Third, the premium of the illiquid asset over the liquid asset falls from 6 to 2 percent. This is the fall in long-term interest rates over the period. The 10-year Treasury yield fell from 6.0 percent in 2000 to 1.8 percent in 2012 and 2016. Short rates fell even more: the federal funds rate went from above 6 percent in 2000 to almost zero for most of the 2010s. A fall of four percentage points is in line with the fall in the 10-year yield between 2000 and the mid-2010s, and smaller than the fall in short rates. All other parameters stay at their baseline values.

Figure 6.1 compares the change in the MPC in the model with the change in the data. The model reproduces the distributional shift from the data. The MPC of low-wealth households rises by 0.14 in the model and by 0.14 in the data. The MPC of mid-wealth households falls by 0.17 in the model and by 0.22 in the data. The model therefore accounts for all of the rise at the bottom and for three quarters of the fall in the middle (Table D.5 in Appendix D.7). It does not account for the fall of the MPC of high-wealth households. This change, however, is not informative as the MPC of high-wealth households is noisy in the data, and its change is not statistically different from zero (Table 3.1). In the counterfactual economy low-wealth households have the highest MPC in the model, as in the data. The true elasticity of consumption to transitory income shocks, which the model observes directly, changes in the same way: it rises by 0.18 at the bottom and falls by 0.18 in the middle (Figure D.3 in Appendix D.5).

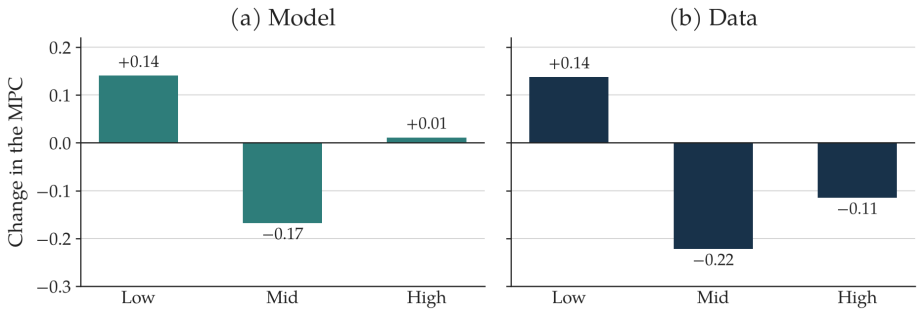


Figure 6.1: Change in the MPC by Wealth Group, Model and Data

Notes: (a) Change between the baseline and the counterfactual economy, estimated on simulated panels as in Section 3.1. (b) Change in the PSID between 1999–2007 and 2011–2019. Wealth groups: terciles of household-average net worth.

The model also reproduces the changes in liquidity behind the divergence in the data (Section 3.4). Fewer mid-wealth households are liquidity-constrained at the end of the sample, and liquid wealth rises for them. The opposite happens for low-wealth households, consistent with the data. The counterfactual economy also reproduces the three stylized facts. First, homeownership falls. Aggregate ownership falls by 12 percentage points in the model, from 56.5 to 44.5 percent, against 11 points in the PSID. The model also reproduces where the fall is concentrated (panel (a) of Figure D.8 in Appendix D.11): ownership falls by 23 percentage points among mid-wealth households in the model and by 21 in the data, against 9 and 5 among low-wealth households and 4 and 7 among high-wealth households. Second, the housing costs of renters rise (panel (b)). I report the median ratio of housing costs to income, because this ratio is very skewed. In the model, the median rises by 2 and 3 percentage points in the middle and at the top of the distribution, against 8 and 3 in the data. At the bottom, renters in the model respond to higher rents by moving to cheaper units or in with others. Because of this adjustment, the median housing cost of low-wealth households who still pay rent rises by only 0.5 percentage points in the model, against 6 in the data. Higher rents alone raise it in every group, but higher house prices lower it, because households who would have bought a house remain renters (Section 6.2).[24] Third, portfolios become more liquid in every group, as in the data (panel (c)). The liquid share of assets rises by 4 to 7 percentage points in the model and by 2 to 3 in the data.[25]

[24]The median housing cost of all households, owners included, falls in the middle of the distribution in the model, because fewer mid-wealth households pay a mortgage. In the data, it rises in every group. I therefore focus on renters, the group whose housing costs the model is designed to capture.

[25]Here the liquid share is the liquid assets of each group over its total assets, the measure I compute in the model, and groups are defined within each period. The average share across households of Section 4.3 rises by 1 to 5 percentage points.

6.2 Disentangling the Effect

The main experiment changes three prices at once. In this section I introduce them one at a time and in pairs, keeping the others at their baseline values. I explore which change drives the MPC of each group, and whether the model needs all three to match the data. Each change moves a different part of household balance sheets, and none of them reproduces the stylized facts alone.

Table 6.1 reports the change in the MPC, homeownership, housing costs of low-wealth renters and liquid share of assets, for each price change across each wealth group. There are three takeaways from this table. First, the rise of the MPC at the bottom comes from higher rents ($+0.06$) and higher house prices ($+0.09$). Second, almost all of the fall in the middle comes from higher house prices (-0.17). Third, lower interest rates barely move the MPC on their own (by at most 0.02 in any group), even though they substantially shift portfolios. The effects are close to additive. For example, higher rents and house prices together raise the MPC at the bottom by 0.14 and lower it in the middle by 0.16, against 0.15 and 0.17 for the sum of their separate effects.

Note that all three changes get substantially close to the data. Only larger rents and lower interest rates together match the three empirical moments qualitatively. However, the change in homeownership is close to zero. The only way to match the empirical patterns qualitatively and quantitatively is by having the three changes together. I discuss each change below.

House prices. Higher house prices raise the MPC at the bottom of the distribution and reduce it in the middle. They work mainly through the extensive margin: who becomes an owner, and when. Panel (a) of Figure 6.2 shows that homeownership falls the most in the middle of the distribution, from 65 to 45 percent, against 17 to 9 percent at the bottom and 87 to 82 percent at the top. Panel (b) shows why: higher prices make it much harder for renters to become owners. Among renters aged 25–40, the share who own a house ten years later falls from 42 to 18 percent at the bottom of the distribution, and from 85 to 69 percent in the middle. This changes how households respond to income shocks.

At the bottom, households give up saving for a house. In the baseline economy, some low-wealth renters save toward a down payment. When their income falls, they draw on these savings and barely cut their consumption (by 0.02 per dollar lost, against 0.35 for the other low-wealth households). With higher prices, the purchase is further away, and most of them stop saving for it. Without these savings, more of the low-wealth group is exposed to income shocks, and its MPC rises.

In the middle, the composition of the group changes. In the baseline economy, the middle

Table 6.1: The Three Changes, Separately, in Pairs and Together

	Change in the MPC			Ownership	Rent burden	Liquid share
	Low	Mid	High	(pp)	(pp)	(pp)
Higher rents	+0.06	0.00	0.00	−0.2	+1.5	−0.1
Higher house prices	+0.09	−0.17	+0.02	−11.1	−1.4	−0.1
Lower interest rates	+0.02	−0.02	0.00	−0.9	−0.1	+4.6
Rents and house prices	+0.14	−0.16	+0.01	−11.1	+0.6	−0.3
House prices and rates	+0.08	−0.19	+0.01	−12.0	−1.6	+5.5
Rents and rates	+0.08	−0.01	0.00	−0.9	+1.5	+4.6
All three	+0.14	−0.17	+0.01	−12.0	+0.5	+5.5
Data	+0.14	−0.22	−0.11	−11.1	+6.4	+2.3

Notes: Changes relative to the baseline economy. MPC estimated on simulated panels with the method of Section 3.1. Ownership: homeownership rate of heads aged 25–55. Rent burden: median housing costs relative to income of low-wealth renters who pay rent. Liquid share: liquid assets over total assets (liquid assets, fixed income and housing), average of the three wealth groups. Data: PSID, change between the initial (1999–2007) and the final (2011–2019) period. Wealth groups in the data: terciles of household-average net worth for the MPC, as in Table 3.1, and terciles of net worth within each period for the other columns, as in the model. In the data, the liquid share is liquid assets over total assets of each group.

of the distribution is full of recent buyers, who have used their savings for the down payment and owe a mortgage payment. The MPC of these liquidity-constrained owners is one of the highest in the model. Panel (c) shows that larger prices cut their share of mid-wealth households from 42 to 23 percent. The households who would have bought remain renters and keep their savings, and the share of mid-wealth households who are renters saving for a house rises. These renters respond little to income shocks, so the MPC of the middle group falls.

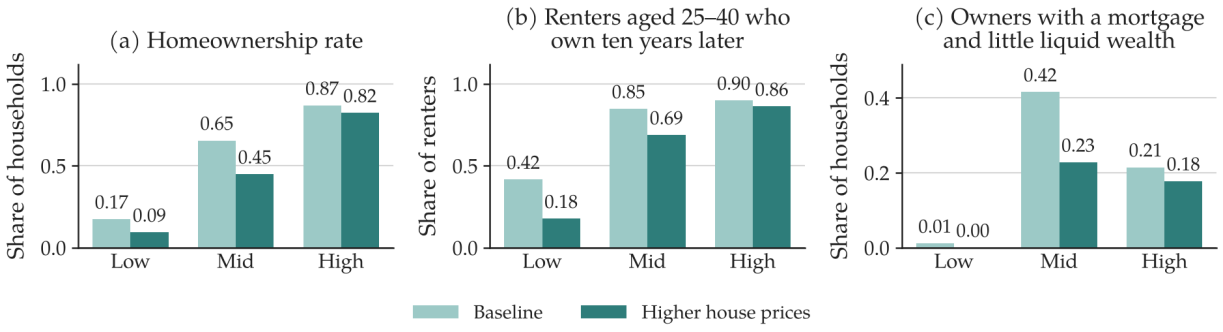


Figure 6.2: House Prices and Access to Ownership

Notes: Baseline economy and economy with higher house prices only. (c) Owners with a mortgage and less than half a month of income in liquid wealth. Wealth groups: terciles of household-average net worth, over the life cycle in panel (a) and over four biennial waves in panels (b) and (c), as in the estimates of the MPC.

Rents. Higher rents increase the housing costs of renters, most of whom are low-wealth households. They respond in two ways. Some stay in the same rental unit and pay more for it. Panel (a) of Figure 6.3 shows that they become more constrained. It compares the same renters, at the same ages, who live in the same unit in both economies. Their transfer MPC[26] of low-wealth renters who stay in their unit rises from 0.21 to 0.32. Others move down the rental ladder. Panel (b) shows how renters are distributed across the ladder. With higher rents, the share of renters who live with others rises from 14 to 23 percent, and fewer renters occupy the standard unit. This adjustment keeps the housing costs of low-wealth renters from rising much, but it also uses up part of the insurance that housing provided. After a fall in income of more than 30 percent, 12 percent of renters move to cheaper housing in the baseline economy, and 9 percent with higher rents. Households who already live in the cheapest option cannot move further down.

Panel (c) shows why this matters for the MPC. It compares the change in the transfer MPC caused by higher rents when households can adjust their housing and when they cannot. At the bottom, the MPC rises in both cases as higher rents leave less room in the budget. However, it rises more when housing can adjust (0.07 against 0.05): the housing margin is a less effective source of insurance with higher rents. Hence, the MPC for low-wealth households increases through these two forces. First, households become more constrained

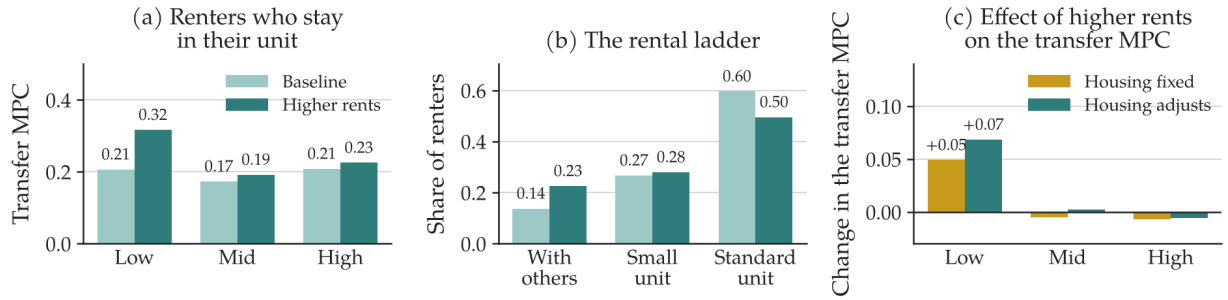


Figure 6.3: Rents and the Rental Ladder

Notes: Baseline economy and economy with higher rents only. Transfers of 0.5% of mean income. (a) Renters who pay rent and live in the same unit as the year before, in both economies: the same households at the same ages. Wealth groups of the baseline economy. (b) Share of renters in each housing option. (c) Change in the transfer MPC caused by higher rents. Wealth groups: terciles of household-average net worth.

Interest rates. Lower interest rates make the illiquid asset less attractive. Panel (a) of Figure 6.4 shows that, with a premium of 2 percent, almost no household holds fixed income. The share of high-wealth households who hold it falls from 60 to 3 percent, and

[26] I use the transfer MPC of Section 5.3 because, unlike the MPC estimated on four waves, it can be computed for a given set of households.

the share of mid-wealth households from 17 to 6 percent. Households move these savings into liquid form, and the liquid share of assets rises by 5 percentage points.

This large change in portfolios barely moves the MPC, because the households who held fixed income did not have a high MPC to begin with. Panel (b) compares, in the baseline economy, the transfer MPC of households who hold fixed income with that of households who do not. Holders spend 0.10 of a small transfer in the middle of the distribution and 0.21 at the top, against 0.46 and 0.38 for the others. Most of them are high-wealth owners with a liquid buffer, or renters saving for a house, who would save a windfall anyway. For these households, moving savings from fixed income to liquid assets does not change how much they consume out of a transfer or out of an income shock. The households for whom liquidity matters, recent buyers with no liquid buffer, hold no fixed income, and they are affected mainly by house prices. This is why lower rates barely move the MPC, alone or together with the other two changes (Table 6.1).

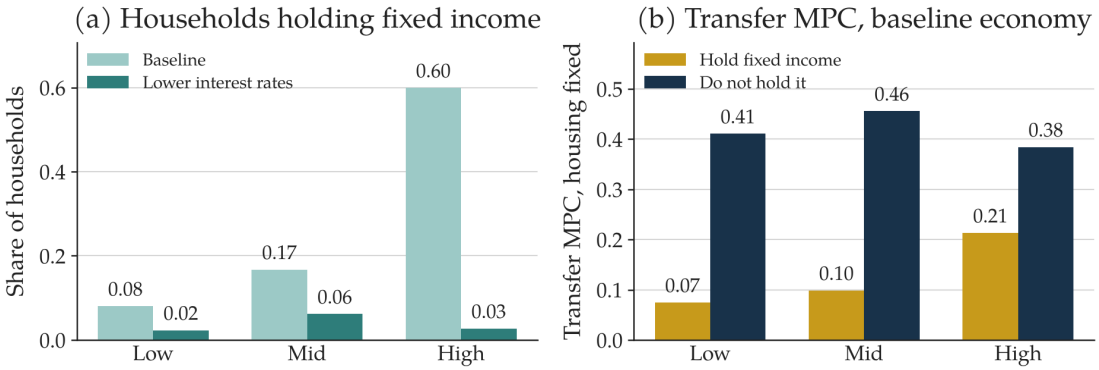


Figure 6.4: Lower Interest Rates and Fixed-Income Holdings

Notes: (a) Baseline economy and economy with lower interest rates only. (b) Baseline economy, transfer of 0.5% of mean income, housing held fixed. Wealth groups: terciles of household-average net worth.

Each change moves a different part of household balance sheets. Higher rents act on households who rent today: they become more constrained, and the housing margin that insured them is partly used up. Higher house prices act on households who would like to own in the future: at the bottom, many give up saving for a house and respond more to shocks; in the middle, fewer households become liquidity-constrained owners, and more remain renters with savings. Together, these two changes account for the rise of the MPC at the bottom, and higher house prices for most of its fall in the middle. Lower interest rates make portfolios more liquid, but they barely move the MPC.

6.3 Discussion: Prices and General Equilibrium

Throughout the paper, I take house prices, rents and interest rates as given. These three prices are tightly linked in equilibrium, and part of the rise in house prices and rents is likely a consequence of the fall in interest rates. This section explains how they are connected, why their joint behavior is hard to pin down in a model, and why I prefer to take their changes from the data.

A house is a long-lived asset, and its price reflects the value of the housing services that it provides over many years. The user cost approach makes this explicit (Poterba, 1984; Himmelberg *et al.*, 2005). A household is indifferent between owning and renting when the rent equals the user cost of owning: the interest forgone on the value of the house, plus maintenance, taxes and depreciation, minus the expected appreciation. When the interest rate falls, the user cost falls, and the price at which owning and renting are equally attractive rises relative to the rent. The price-to-rent ratio is the inverse of the user cost, so it rises when rates fall, as it did in the sample period (Appendix C.2).

The effect of lower rates on rents is ambiguous. On the one hand, lower rates make owning cheaper. Some renters become owners, the demand for rental housing falls, and rents fall. On the other hand, higher prices raise the down payment needed to buy a house. Households who cannot save enough remain renters for longer, and the demand for rental housing rises (Dias and Duarte, 2019; Abramson *et al.*, 2026). This second force is stronger when credit is tight, as it was after the financial crisis. On the supply side, lower rates reduce the cost of financing new housing, which pushes rents down, but only where land and regulation allow the housing stock to grow. Which force dominates is an empirical question (Sommer *et al.*, 2013). In the PSID, rents rose faster than income over the sample (Section 4.2), which suggests that the second force dominated.

Closing the model would require three additional blocks. The first is the supply of housing, whose response to prices depends on land and regulation and differs a lot across local markets. The second is a rental sector, in which the financing costs of landlords and the tenure choices of households determine rents. The third is the interest rate itself. The secular decline in interest rates is a global phenomenon, driven by forces such as demographics, inequality and the demand for safe assets (Mian *et al.*, 2021).

For these reasons, I take the changes in the three prices from the data and remain agnostic about their common drivers. The effect of each price on the MPC goes through household balance sheets, which the model matches in the baseline economy and whose changes it reproduces. A general equilibrium model would change why prices moved, but, for a given change in prices, the response of households would be the same.

7 Fiscal Policy

The divergence in MPCs changes who responds the most to a fiscal transfer. In this section I ask what it implies for the design of one-off transfers, such as the stimulus payments of 2001, 2008 and 2020–2021. I study the effectiveness of targeted fiscal policy relative to uniform transfers and how it has changed over time.

The relevant object for this question is the transfer MPC of Section 5.3: the share of a one-off transfer that a household spends within the year, letting it re-optimize its housing. The transfer MPC and the MPC move in the same direction between the two periods (Appendix D.6), so the divergence documented in Section 3 carries over to the response to transfers. However, the transfer MPC depends on the size of the transfer. A liquidity-constrained household without liquid wealth spends most of a small transfer, but beyond some amount it saves the rest. Hence, targeting transfers to high-MPC households is not always more effective than transferring the same amount uniformly. I first explore the nonlinearity of MPCs in Section 7.1. Then, I present the main fiscal experiment, which shows how the effectiveness of targeted transfers has changed, in Section 7.2. Finally, in Section 7.3, I compare the cost of raising spending by a given amount with targeted and with uniform transfers.

7.1 MPC Nonlinearity

The MPC out of a transfer is nonlinear relative to the size of the transfer. As the size of the transfer increases, the consumption response goes down. Households are going to start saving their transfer as soon as the transfer is so large that they are no longer constrained. A key concept to understand this is the consumption gap. This is the difference between the consumption a constrained household would choose without its constraint and its current consumption. A transfer smaller than the gap relaxes the constraint, and the household spends most of it. A larger transfer lifts the household off its constraint, and beyond that point it saves most of each additional dollar. The MPC is therefore roughly constant for transfers below the gap and falls for transfers above it. A household with a small gap spends most of a small transfer but little of a large one. A household with a large gap spends a large share of both. Moreover, the same monetary transfer affects households differently depending on the consumption gap.

Panel (a) of Figure 7.1 shows the transfer MPC of all households by the size of the transfer. In the baseline economy, households spend 0.28 of a transfer of 0.5% of mean income and 0.18 of a full year of income. For small and moderate transfers, the curve barely moves between the two economies, consistent with the stable aggregate MPC of Section 3.2. Panel (b) shows the same curve by wealth group in the baseline economy, and Table E.1 in

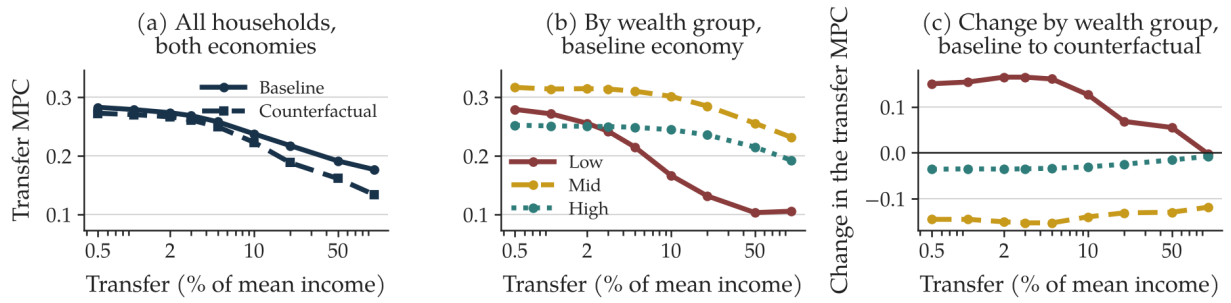


Figure 7.1: Transfer MPC by the Size of the Transfer

Notes: Share of a one-off transfer spent within the year, with housing re-optimized. (b) Baseline economy. (c) Change between the two economies. Wealth groups: terciles of current net worth. Transfers in percent of mean income, logarithmic scale.

Appendix E shows which households drive each group. There are three main takeaways. First, the MPC of low-wealth households falls steeply with the size of the transfer: they spend 0.28 of a transfer of 0.5% of mean income, 0.13 of a transfer of 20%, and 0.11 of a full year of income. The households that respond are the renters at or near their credit limit, 36% of the group. They spend half of a small transfer, but they are poor in every sense, so their gap is small and a transfer of 20% of mean income already lifts most of them off their constraint. The other low-wealth renters spend little of any transfer (0.15), because repaying unsecured debt at 12% is a better use of a windfall.

Second, the MPC of mid-wealth households barely depends on the size of the transfer (0.32 of a transfer of 0.5%, 0.31 of a transfer of 5%, and 0.29 of a transfer of 20%). Most of them (72%) are owners with a mortgage, with a net worth of 0.62 years of mean income, almost all in their home, and liquid wealth of only 0.04. Given their total wealth, they would like to consume much more than their liquid resources allow, so their gap is large and they spend about half of any transfer up to 20% of mean income. The other mid-wealth households are renters saving for a house. They save a transfer toward the down payment, and for some of them the transfer brings the purchase forward. The down payment then absorbs the transfer and part of their consumption, which is why their average transfer MPC is negative. Third, high-wealth households are not constrained. They spend about a quarter of any transfer, which reflects their impatience and their remaining lifetime rather than a lack of liquidity.

Panel (c) of Figure 7.1 shows that the divergence shifts these curves, mostly for small and moderate transfers. Low-wealth households spend 0.43 of a small transfer in the counterfactual economy, against 0.28 in the baseline economy: more of them are at or near their credit limit (44% against 36%), and the others are closer to it. Their gap grows

but remains small, so their curve falls even more steeply and meets the baseline one for large transfers. In the middle, owners with a mortgage still spend about half of a transfer, but they are now only 38% of the group. The middle is now mostly made of renters with liquid savings, who spend almost nothing of a transfer, and the MPC of mid-wealth households falls to about 0.17 at every size up to 20%.

7.2 Targeted Transfers Effectiveness

I now turn to the main fiscal experiment. Consider a budget of Z per household aged 25–55, in units of mean annual income. A *uniform* transfer pays Z to every household. A *targeted* transfer pays Z/f to each household in the bottom f of the distribution of net worth, and nothing to the others, so that both cost the same. The spending that each generates within the year, per dollar of budget, is the average transfer MPC of the recipients, each evaluated at the transfer that it actually receives:

$$S(f;Z) = \mathbb{E}\big[\,\mathrm{MPC}_i(Z/f) \;\big|\; i \in \text{bottom } f\,\big], \qquad G(f;Z) = \frac{S(f;Z)}{S(1;Z)} - 1, \tag{14}$$

where G is the gain from targeting: the extra spending that the targeted transfer generates relative to the uniform one. I compute $\mathrm{MPC}_i(\tau)$ in the simulated panel for transfers between 0.5% and 100% of mean income, and interpolate between them. I focus on a transfer targeted to low-wealth households, the bottom third of the distribution of net worth, and consider two budgets, 1% and 3% of mean income per household. For comparison, the 2008 stimulus payments amounted to roughly 1% of mean household income, and the 2021 payments to roughly 3%.

Panel (a) of Figure 7.2 shows the gain from targeting low-wealth households. There are two main takeaways. First, the divergence turned targeting from counterproductive into effective. With a budget of 1%, a transfer to low-wealth households generates 13% less spending than a uniform transfer in the baseline economy, and 50% more in the counterfactual economy. Two forces add up. The MPC of low-wealth households rose, so a targeted transfer generates more spending. And the MPC of mid-wealth households fell, so a uniform transfer, much of which goes to the middle, generates less. In the baseline economy, a uniform transfer reached the liquidity-constrained owners of the middle, while a transfer targeted to low-wealth households missed them.

Second, the size of the budget matters because of the nonlinearity of Section 7.1. With a budget of 3%, each low-wealth household receives 9% of mean income, a size at which many of them are already lifted off their constraint. Targeting them changes spending by -35% in the baseline economy and by $+17\%$ in the counterfactual economy. The gain

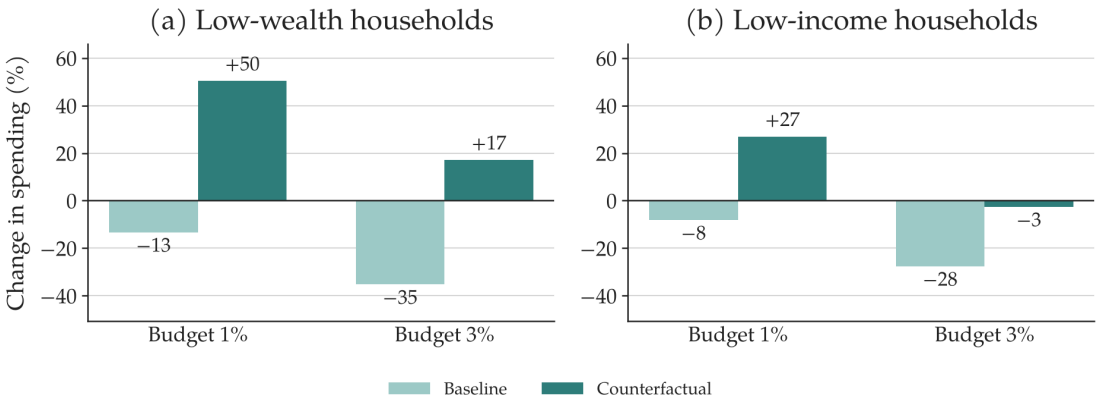


Figure 7.2: Gain from Targeting Transfers

Notes: Gain $G(1/3; Z)$ from equation (14): the change in spending within the year from a transfer to the bottom third, relative to a uniform transfer with the same budget Z. Targeting by net worth (a) and by income (b).

is smaller than with a budget of 1%, but the divergence moves it in the same direction. The design of a transfer program should therefore depend on its size, and the divergence widened the range of budgets for which targeting pays.

The targeting variable also matters. Stimulus payments are usually phased out with income, so a natural alternative is to target households with low income. Panel (b) of Figure 7.2 shows the gain from targeting the bottom third of the distribution of income. Income is a worse target than net worth. With a budget of 1%, targeting low-income households changes spending by -8% in the baseline economy and by $+27\%$ in the counterfactual economy, against $+50\%$ when targeting by net worth. With a budget of 3%, the change is -3% in the counterfactual economy. The divergence also raised the gain from targeting by income, but by less. Households with low income in a given year include many households with a temporary fall in income who hold liquid savings, and many high-MPC households in the middle of the wealth distribution have average incomes. Appendix E shows the gain from targeting any share of the distribution of net worth, income and liquid wealth. Liquid wealth, which identifies liquidity-constrained households directly, is the best target in both economies.

7.3 Targeting Makes Stimulus Cheaper

The previous exercise holds the budget fixed and compares spending. An alternative question is how much a government needs to spend to obtain a given increase in spending. Suppose that a uniform transfer of Z_u per household generates a total increase in spending of $Z_u\, S(1; Z_u)$. I look for the budget Z_f such that a transfer targeted to the bot-

tom f of the distribution of net worth generates the same increase:

$$Z_f \, S(f; Z_f) = Z_u \, S(1; Z_u). \tag{15}$$

If $Z_f < Z_u$, targeting reaches the same stimulus at a lower cost. Figure E.1 in Appendix E shows Z_f/Z_u for a transfer to low-wealth households and for a transfer to low-income households, for uniform transfers of 1% and 3% of mean income.

The divergence made stimulus through transfers to low-wealth households almost twice as cheap. To match the spending generated by a uniform transfer of 1%, a transfer to low-wealth households costs 120% of the uniform transfer in the baseline economy, and only 64% in the counterfactual economy. For a uniform transfer of 3%, it costs 194% of the uniform transfer in the baseline economy and 78% in the counterfactual economy. In the baseline economy, targeting low-wealth households was an expensive way to stimulate spending. In the counterfactual economy, it saves between 22 and 36 percent of the cost. Targeting low-income households also became cheaper, but by less: it costs 111% and 76% of a uniform transfer of 1% in the two economies, and 154% and 104% of a uniform transfer of 3%.

The aggregate response to a uniform transfer, in contrast, barely changed. A uniform transfer of 1% generates spending of 0.28 per dollar in the baseline economy and 0.27 in the counterfactual economy, consistent with the stable aggregate MPC of Section 3.2. A policymaker who looks only at the aggregate MPC would conclude that nothing has changed. The distribution of MPCs shows that the effectiveness of targeted transfers has changed a lot.

8 Conclusions

This paper studies how the distribution of marginal propensities to consume across the wealth distribution has changed over time, the drivers behind these changes and what they imply for fiscal policy. Using the PSID between 1999 and 2019, I show that the MPC of low-wealth households increased, while that of mid-wealth households declined. At the beginning of the period, the MPC was higher for the mid-wealth households than for the low-wealth group. This was driven by two forces. First, the mid-wealth group had a large share of liquidity-constrained wealthy households, mainly owners with a mortgage. Second, low-wealth households used the rental market as a buffer against large shocks.

Across time, the distribution of liquidity-constrained households changed. This share increased at the bottom and reduced at the middle of the distribution. The shift came with changes in household balance sheets. Liquid wealth fell for low-wealth households and

rose for mid- and high-wealth households. Homeownership declined, and fewer renters became owners. Housing costs rose, especially for low-wealth renters. Portfolios became more liquid, as households moved away from fixed-income assets. These changes are connected to three secular trends: rising house prices, larger rents and the secular decline in interest rates.

I build a life-cycle model with housing tenure, a rental ladder, and liquid and illiquid assets to explain and quantify the role of rising house prices, rising rents and falling interest rates on the MPC across the wealth distribution. Calibrated to the early 2000s, the model generates two groups of high-MPC households. These are poor renters at the bottom and owners with a large mortgage in the middle. Additionally, given the existence of illiquid assets and a rental market, it reproduces the MPC ordering of the data without targeting it.

I abstract away from general equilibrium forces and take prices as exogenous. These are tightly connected in equilibrium, and it is likely that the secular decline in interest rates is the main driver. However, I proceed by computing a counterfactual economy where I change prices exogenously to their level in the late 2010s. The model accounts for all of the rise of the MPC at the bottom and for three quarters of its fall in the middle. Each price change moves a different part of household balance sheets. Higher house prices delay or prevent home purchases and remove high-MPC households from the middle. Higher rents push low-wealth renters down the rental ladder and use up part of the insurance that housing provided. Lower interest rates make portfolios more liquid. Only the three together reproduce the stylized facts and the full divergence in MPCs.

The policy implication is that the gains from fiscal targeting change over time. The high-MPC households are not always the same in the economy. Changes in the macroeconomic environment can alter both who they are and the gains from targeting them. Given the strong presence of liquidity-constrained wealthy households in the mid-wealth group, a uniform transfer is more effective in the baseline economy than a transfer targeted to low-wealth households. However, this reverses in the counterfactual economy, where the gain of targeting these households can go up to 50% relative to the uniform case. Hence, targeted transfers towards low-wealth households have become more powerful.

This paper opens several avenues for future research. First, the secular decline in interest rates has reversed, and rates appear likely to remain higher for longer. The mechanisms in this paper suggest that this reversal need not undo the shift in MPCs. A wider spread between illiquid and liquid returns may draw savings back out of liquid buffers, while house prices remain high and higher mortgage rates make homeownership even harder to reach. How the MPC distribution responds when interest rates rise, but housing costs

do not fall is an open question with direct implications for policy transmission.

Second, this paper takes house prices, rents, and interest rates as exogenous. A natural next step is to link them in equilibrium, for instance by letting lower interest rates raise house prices and rents through housing demand and constrained supply, and to build a full theory of how these forces jointly reshape the distribution of MPCs.

Third, the shift in MPCs across the wealth distribution matters for monetary policy as well as fiscal policy. Monetary policy works in part through household cash flows, mortgage payments, and rents, and its transmission depends on which households are liquidity-constrained and how their balance sheets are exposed to interest rates. If constrained households are now mostly renters rather than recent mortgagors, the channels through which monetary policy reaches spending may have changed substantially.

References

ABRAMSON, B., DE LLANOS, P. and HAN, L. (2026). Monetary policy and rents, working paper.

AGARWAL, S., CHOMSISENGPHET, S., MAHONEY, N. and STROEBEL, J. (2015). Regulating consumer financial products: Evidence from credit cards. *The Quarterly Journal of Economics*, **130** (1), 111–164.

AGUIAR, M., BILS, M. and BOAR, C. (2025). Who are the hand-to-mouth? *The Review of Economic Studies*, **92** (3), 1293–1340.

ALBOUY, D., EHRLICH, G. and LIU, Y. (2016). *Housing Demand, Cost-of-Living Inequality, and the Affordability Crisis*. Working Paper 22816, National Bureau of Economic Research.

ANENBERG, E., SCHARLEMANN, T. and VAN STRAELEN, E. (2025). Borrowing and spending in the money: Debt substitution and the cash-out refinance channel of monetary policy. *American Economic Review*, **115** (11), 3909–3940.

AUCLERT, A. (2019). Monetary policy and the redistribution channel. *American Economic Review*, **109** (6), 2333–2367.

—, BARDÓCZY, B., ROGNLIE, M. and STRAUB, L. (2019). Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. *NBER Working Paper*, (w26123).

—, ROGNLIE, M. and STRAUB, L. (2018). *The Intertemporal Keynesian Cross*. NBER Working Paper 25020, National Bureau of Economic Research, revised April 2024.

—, — and — (2025). Fiscal and monetary policy with heterogeneous agents. *Annual Review of Economics*, **17**.

BERGER, D., GUERRIERI, V., LORENZONI, G. and VAVRA, J. (2018). House prices and consumer spending. *The Review of Economic Studies*, **85** (3), 1502–1542.

BLUNDELL, R., PISTAFERRI, L. and PRESTON, I. (2008). Consumption inequality and partial insurance. *American Economic Review*, **98** (5), 1887–1921.

BOUND, J., BROWN, C., DUNCAN, G. J. and RODGERS, W. L. (1994). Evidence on the validity of cross-sectional and longitudinal labor market data. *Journal of Labor Economics*, **12** (3), 345–368.

CARROLL, C., SLACALEK, J., TOKUOKA, K. and WHITE, M. N. (2017). The distribution of wealth and the marginal propensity to consume. *Quantitative Economics*, **8** (3), 977–1020.

CARROLL, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, **91** (3), 312–320.

CHETTY, R. and SZEIDL, A. (2007). Consumption commitments and risk preferences. *Quarterly Journal of Economics*, **122** (2), 831–877.

CLOYNE, J., FERREIRA, C. and SURICO, P. (2020). Monetary policy when households have debt: New evidence on the transmission mechanism. *The Review of Economic Studies*, **87** (1), 102–129.

COMMAULT, J. (2022). Does consumption respond to transitory shocks? Reconciling natural experiments and semistructural methods. *American Economic Journal: Macroeconomics*, **14** (2), 96–122.

CRAWLEY, E. and KUCHLER, A. (2023). Consumption heterogeneity: Micro drivers and macro implications. *American Economic Journal: Macroeconomics*, **15** (1), 314–341.

DAUMLER, D., FRIEDMAN, E. and PFEFFER, F. T. (2025). *PSID-SHELF User Guide and Codebook, 1968–2021, Beta Release*. Tech. rep., Survey Research Center, Institute for Social Research, University of Michigan.

DAVIS, M. A., LEHNERT, A. and MARTIN, R. F. (2008). The rent-price ratio for the aggregate stock of owner-occupied housing. *Review of Income and Wealth*, **54** (2), 279–284.

DE NARDI, M. (2004). Wealth inequality and intergenerational links. *Review of Economic Studies*, **71** (3), 743–768.

—, FRENCH, E. and JONES, J. B. (2010). Why do the elderly save? The role of medical expenses. *Journal of Political Economy*, **118** (1), 39–75.

DIAS, D. A. and DUARTE, J. B. (2019). *Monetary Policy, Housing Rents, and Inflation Dynamics*. International Finance Discussion Paper 1248, Board of Governors of the Federal Reserve System.

— and — (2022). *Monetary Policy and Homeownership: Empirical Evidence, Theory, and Policy Implications*. International Finance Discussion Paper 1344, Board of Governors of the Federal Reserve System.

EICHENBAUM, M., REBELO, S. and WONG, A. (2022). State-dependent effects of monetary policy: The refinancing channel. *American Economic Review*, **112** (3), 721–761.

ENGELHARDT, G. V. (1994). House prices and the decision to save for down payments. *Journal of Urban Economics*, **36** (2), 209–237.

FAGERENG, A., HOLM, M. B. and NATVIK, G. J. (2021). MPC heterogeneity and household balance sheets. *American Economic Journal: Macroeconomics*, **13** (4), 1–54.

FLODÉN, M., KILSTRÖM, M., SIGURDSSON, J. and VESTMAN, R. (2021). Household debt and monetary policy: Revealing the cash-flow channel. *The Economic Journal*, **131** (636), 1742–1771.

FOULIS, A. K., HAZELL, J., MIAN, A. R. and TRACEY, B. (2026). *How Do Interest Rates Affect Consumption? Household Debt and the Role of Asset Prices*. NBER Working Paper 34911, National Bureau of Economic Research.

GANONG, P., JONES, D., NOEL, P., FARRELL, D., GREIG, F. and WHEAT, C. (2025). *Liquid Wealth and Consumption Smoothing of Typical Labor Income Shocks*. Tech. rep., Working paper.

GROSS, T., NOTOWIDIGDO, M. J. and WANG, J. (2020). The marginal propensity to consume over the business cycle. *American Economic Journal: Macroeconomics*, **12** (2), 351–384.

HALKET, J. and VASUDEV, S. (2014). Saving up or settling down: Home ownership over the life cycle. *Review of Economic Dynamics*, **17** (2), 345–366.

HIMMELBERG, C., MAYER, C. and SINAI, T. (2005). Assessing High House Prices: Bubbles, Fundamentals and Misperceptions. *Journal of Economic Perspectives*, **19** (4), 67–92.

ISKHAKOV, F., JØRGENSEN, T. H., RUST, J. and SCHJERNING, B. (2017). The endogenous grid method for discrete-continuous dynamic choice models with (or without) taste shocks. *Quantitative Economics*, **8** (2), 317–365.

JAPPELLI, T. and PISTAFERRI, L. (2014). Fiscal policy and MPC heterogeneity. *American Economic Journal: Macroeconomics*, **6** (4), 107–136.

JOHNSON, D. S., PARKER, J. A. and SOULELES, N. S. (2006). Household expenditure and the income tax rebates of 2001. *American Economic Review*, **96** (5), 1589–1610.

KAPLAN, G. (2012). Moving back home: Insurance against labor market risk. *Journal of Political Economy*, **120** (3), 446–512.

—, MOLL, B. and VIOLANTE, G. L. (2018). Monetary policy according to HANK. *American Economic Review*, **108** (3), 697–743.

— and VIOLANTE, G. L. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, **82** (4), 1199–1239.

— and — (2022). The marginal propensity to consume in heterogeneous agent models. *Annual Review of Economics*, **14**, 747–775.

—, — and WEIDNER, J. (2014). The wealthy hand-to-mouth. *Brookings Papers on Economic Activity*, **2014** (1), 77–138.

KOPECKY, K. A. and SUEN, R. M. H. (2010). Finite state Markov-chain approximations to highly persistent processes. *Review of Economic Dynamics*, **13** (3), 701–714.

LEWIS, D., MELCANGI, D. and PILOSSOPH, L. (2024). *Latent Heterogeneity in the Marginal Propensity to Consume*. NBER Working Paper 32523, National Bureau of Economic Research.

LOCKWOOD, L. M. (2018). Incidental bequests and the choice to self-insure late-life risks. *American Economic Review*, **108** (9), 2513–2550.

LUETTICKE, R. (2021). Transmission of monetary policy with heterogeneity in household portfolios. *American Economic Journal: Macroeconomics*, **13** (2), 1–25.

MIAN, A., STRAUB, L. and SUFI, A. (2021). Indebted demand. *The Quarterly Journal of Economics*, **136** (4), 2243–2307.

MISRA, K. and SURICO, P. (2014). Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs. *American Economic Journal: Macroeconomics*, **6** (4), 84–106.

PARKER, J. A., SOULELES, N. S., JOHNSON, D. S. and MCCLELLAND, R. (2013). Consumer spending and the economic stimulus payments of 2008. *American Economic Review*, **103** (6), 2530–2553.

PATTISON, N. (2024). *Landlords as Lenders of Last Resort? Late Housing Payments During Unemployment*. Working Paper 2401, Southern Methodist University, Department of Economics.

POTERBA, J. M. (1984). Tax Subsidies to Owner-Occupied Housing: An Asset-Market Approach. *The Quarterly Journal of Economics*, **99** (4), 729–752.

ROUWENHORST, K. G. (1995). Asset pricing implications of equilibrium business cycle models. In T. F. Cooley (ed.), *Frontiers of Business Cycle Research*, Princeton, NJ: Princeton University Press, pp. 294–330.

RUST, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica*, **55** (5), 999–1033.

SINAI, T. and SOULELES, N. S. (2005). Owner-occupied housing as a hedge against rent risk. *The Quarterly Journal of Economics*, **120** (2), 763–789.

SOMMER, K., SULLIVAN, P. and VERBRUGGE, R. (2013). The equilibrium effect of fundamentals on house prices and rents. *Journal of Monetary Economics*, **60** (7), 854–870.

WIEMERS, E. E. (2014). The effect of unemployment on household composition and doubling up. *Demography*, **51** (6), 2155–2178.

YOUNG, E. R. (2010). Solving the incomplete markets model with aggregate uncertainty using the Krusell–Smith algorithm and non-stochastic simulations. *Journal of Economic Dynamics and Control*, **34** (1), 36–41.

A Data Appendix

This appendix provides additional details on the data. I first describe the construction and cleaning of the three datasets I use: the PSID, the SCF, and the CEX. I then report additional descriptive statistics from the SCF. Results including the 2021 PSID wave are in Appendix B.6.

A.1 Dataset Cleaning

This section describes how I construct and clean each of the three surveys: the PSID, the SCF, and the CEX. For each one, I describe the dataset, the cleaning process, and the main definitions I use throughout the paper (e.g., HtM households).

Panel Study of Income Dynamics (PSID). I use the PSID-SHELF harmonized extract of the PSID family and individual files, 1968–2021, and keep the biennial waves from 1999 to 2021. The PSID started as an annual survey and moved to biennial interviews after 1997. Comprehensive expenditure questions and a wealth module in every wave were introduced in 1999, so 1999 is the first wave in which income, consumption, and wealth are all observed for the same households. Each wave of the PSID-SHELF contains one record per individual. I keep the reference person of each family unit, so that each household contributes one observation per wave. I restrict the sample to the nationally representative core sample (the original 1968 SRC sample and its descendants), and drop the low-income SEO oversample and the 1997/1999 immigrant refresher families, consistent with Blundell *et al.* (2008). All statistics are weighted with the PSID cross-sectional individual weight of the reference person. Monetary variables are deflated using the PSID-SHELF deflator and expressed in constant 2024 dollars.

I focus on the working-age population (ages 25 to 55) and drop households whose reference person is permanently out of the labor force. Within each wave, I drop observations whose income growth lies outside the 1st–99th percentile range, as well as observations in the top percentile of the ratios of liquid wealth, cash on hand (liquid wealth net of one month of housing expenditure), and housing expenditure to labor income.

Labor income is the sum of the labor earnings of the reference person and of the spouse or partner. Consumption is total family expenditure. The consumption measure used in the MPC regressions excludes housing expenditure (rent, mortgage payments, property taxes, insurance, and utilities), so that the estimated response is not mechanically affected by housing payments that are predetermined within the year. This measure includes expenditure on food, transportation, education, child care, health, clothing, recreation, and trips.

I use net worth, including home equity, to group households. Liquid assets are the sum of checking and savings accounts and other liquid assets. Liquid debt is only reported in the last few waves, while non-mortgage debt is reported in every wave. I therefore proceed in two steps. First, I compute the share of liquid debt in non-mortgage debt (credit card, student, medical, legal, and family loans) in the waves in which both are available. Second, I impute liquid debt in every wave by multiplying non-mortgage debt by this share. Liquid wealth is then liquid assets net of liquid debt.

Illiquid wealth is net worth minus liquid wealth, that is, home equity, other real estate, business and vehicle equity, IRAs and annuities, and directly held stocks and mutual funds, net of all debt other than liquid debt. A household is a homeowner if it owns the dwelling it occupies; all other households are classified as renters.

Following Kaplan *et al.* (2014), a household is hand-to-mouth (HtM) if its liquid wealth m is small relative to its per-period income. With a two-week pay period, per-period income is $y = Y/26$, where Y is annual labor income, and the household is HtM if

$$0 \leq m \leq \tfrac{1}{2}\, y \qquad \text{or} \qquad m < 0 \;\text{ and }\; m \leq \tfrac{1}{2}\, y - \underline{m},$$

where the credit limit $\underline{m}$ is 18.5 percent of annual income, the median credit limit relative to income reported by Kaplan *et al.* (2014). Poor HtM households are HtM households with non-positive illiquid wealth; wealthy HtM households are HtM households with positive illiquid wealth.

I divide households into three wealth groups using the weighted distribution of average net worth, where the average is taken over all waves in which the household is observed. A household's group is therefore fixed over time. In the descriptive sample (reference persons aged 25 to 55, all waves), the boundaries between groups are \$14,984 and \$103,062.

Survey of Consumer Finances (SCF). I use the twelve triennial waves of the SCF from 1989 to 2022. The SCF releases five imputation replicates per household. I average the replicates within each household, so that each household contributes one observation, and use the analysis weights of the summary extract. The extract reports monetary variables in constant 2022 dollars. I convert them to 2024 dollars with the CPI-U, as in the PSID. The wealth groups use the PSID thresholds, so both surveys use the same dollars.

Consistent with the PSID, I keep households whose reference person is aged 25 to 55. I drop observations with missing education, marital status, housing, financial assets, or net worth, and observations with negative total income. Within each survey year, I then drop the bottom and top percentile of three ratios: liquid wealth to income, net worth to income, and liquid wealth to net worth.

I use the SCF to study the composition of household balance sheets. I first divide financial assets into four categories: liquid, fixed-income, retirement, and equities. Liquid assets consist of checking, savings, money market, and call accounts, and prepaid cards (which are only recorded from 2016 onward). Fixed-income assets consist of certificates of deposit, directly held bonds, savings bonds, and the cash value of whole-life insurance. Retirement assets consist of IRAs, Keogh plans, and account-type pensions such as 401(k)s. Finally, equities consist of directly held stocks, pooled investment funds, other managed assets, and other financial assets.

Real assets consist of the primary residence and other non-financial assets, which include vehicles, other real estate, businesses, and other assets. Portfolio shares are computed household by household and then averaged using the survey weights. For robustness, I also use three alternative classifications that move certificates of deposit into liquid assets, the cash value of life insurance into retirement assets, or both, and a fourth that drops life insurance altogether. The trends reported in the paper are not sensitive to these choices.

The SCF HtM classification mirrors the one in the PSID. Consistent with Kaplan *et al.* (2014), liquid wealth is liquid assets, scaled up by 5.5 percent to account for cash holdings not recorded in the survey, net of credit card balances. Illiquid wealth is the primary residence and other real estate net of real-estate debt, plus certificates of deposit, the cash value of life insurance, and retirement accounts. A household is HtM if its liquid wealth is between zero and one week of labor and transfer income, or if it is negative and below one week of income minus a credit limit of one month of income. Poor and wealthy HtM households are distinguished by the sign of their illiquid wealth.

Finally, I assign households to wealth groups by applying the PSID thresholds ($14,984 and $103,062) to their net worth in the survey year.

Consumer Expenditure Survey (CEX). I use the family-level (FMLI) files of the CEX Interview Survey, quarterly from the first quarter of 1990 to the fourth quarter of 2023. Each consumer unit is interviewed up to four times at three-month intervals and reports its expenditures for the preceding three months. Because variable names change across survey redesigns, I harmonize the income, asset, tax, and expenditure variables across years.

Financial assets are the sum of the balances reported in the asset questions available in each year (checking, savings, bonds, securities, and other money owed before 2013, and a single liquid-asset question from 2013 onward). Their value one year earlier is recovered from the reported changes. Reported income taxes are no longer released after 2015, so after-tax income uses the CEX's estimated taxes when reported taxes are missing.

Expenditure categories follow the published CEX aggregates. Nondurable consumption is the sum of food, alcohol, apparel, entertainment, personal care, reading, tobacco, and cash contributions; a broader measure adds education and health. Housing expenditure is the published housing aggregate (shelter, utilities, household operations, and furnishings and equipment), and the shelter component is further split into rent, mortgage interest and charges, and other owned-dwelling expenses. Income is family income before taxes. All nominal variables are deflated with the CPI.

I drop consumer units that report no food-at-home or no nondurable expenditure in the quarter, that have missing or negative income before taxes, or that report zero income in every interview. I keep units whose head is aged 25 to 55. Within each tenure group and quarter, I also drop the bottom and top percentile of nondurable and of durable expenditure. Because the housing-payment series are meant to track the cost of a given tenure status, I drop units that change tenure between interviews.

A.2 Additional Data from the SCF

This section reports additional descriptive statistics from the SCF. Table A.1 is the SCF analog of Table 2.1. The SCF is qualitatively similar to the PSID: income, liquid wealth, and homeownership all increase with wealth. Quantitatively, however, the two surveys differ. The SCF oversamples wealthy households and captures the upper tail of the wealth distribution more fully, so average income and wealth are higher than in the PSID. Given these differences, I use the PSID whenever possible and use the SCF only to complement the analysis.

Table A.2 decomposes total assets for the whole sample and for each wealth group. There are two key takeaways. First, financial assets become less liquid as wealth increases: liquid assets represent 37 percent of financial assets for low-wealth households, but only about 12 percent for high-wealth households. Second, financial assets account for a larger share of total assets as wealth increases, rising from 13 percent for low-wealth households to 34 percent for high-wealth households. The primary residence is the largest asset of low- and mid-wealth households (57 and 63 percent of total assets, respectively), compared with 37 percent for high-wealth households.

A.3 Consumer Expenditure Survey (CEX)

I also use the Consumer Expenditure Survey (CEX), a nationally representative survey conducted by the U.S. Bureau of Labor Statistics that collects detailed information on household expenditures, income, and demographic characteristics. Unlike the SCF, which is designed to measure household balance sheets, the CEX is designed to measure con-

Table A.1: Descriptive Statistics by Wealth Group (SCF)

		By wealth group		
	Overall	Low-wealth	Mid-wealth	High-wealth
Income				
Labor income	123,620	49,082	75,702	185,007
Wealth				
Net wealth	418,653	-6,699	51,150	817,029
Total assets	561,197	38,902	136,744	1,036,667
Total debt	142,544	45,601	85,595	219,638
Liquid wealth				
Liquid wealth	20,580	-19	2,989	39,767
Liquid assets	24,028	1,918	6,012	44,180
Liquid debt	3,448	1,937	3,023	4,413
Household characteristics				
Age	40	37	38	43
Homeownership (%)	62.7%	14.2%	56.4%	89.9%
College education (%)	33.2%	17.2%	23.2%	46.3%

Notes: Entries are weighted means using cross-sectional weights. Monetary variables are reported in 2024 dollars. Homeownership and college education are reported in percent.

sumption and expenditure patterns. The CEX Interview Survey is quarterly and has a rotating panel design: each household is interviewed for up to four consecutive quarters and then rotates out of the sample. The CEX started in 1972, but its current structure was first implemented in 1984.

I use the CEX to construct measures of housing expenditure, focusing on mortgage payments for homeowners and rent payments for renters. Its detailed expenditure data allow a more precise measure of housing expenditure relative to consumption and income than the PSID. However, the CEX does not contain comprehensive wealth data, which is why I use it only as a secondary data source. Appendix A.1 describes the cleaning of the CEX.

B Empirics Appendix

This appendix complements Section 3. I first derive the MPC estimator of Blundell *et al.* (2008) and then test the assumptions of the underlying model, following the original paper. I then show how poor and wealthy HtM households are distributed across wealth groups. Next, I report additional results on how households adjust after large transitory income losses, the role of unemployment, and how this adjustment changed after 2008. I then report the MPC under alternative definitions of the wealth groups, and the MPC of renters and owners, overall and within each wealth group. Finally, I report the main

Table A.2: Asset Decomposition by Wealth Group (SCF)

		By wealth group		
	Overall	Low-wealth	Mid-wealth	High-wealth
Total assets	561,197	38,902	136,744	1,036,667
Financial assets	182,772	5,123	22,546	352,598
Liquid assets	24,028	1,918	6,012	44,180
Checking and savings accounts	16,122	1,740	5,333	28,755
Other liquid assets	7,906	178	678	15,426
Fixed-income assets	18,510	524	3,226	35,224
CDs and bonds	9,558	136	1,014	18,588
Cash value of life insurance	8,952	389	2,212	16,637
Retirement assets	80,295	2,205	10,712	154,513
Equities in retirement accounts	47,794	1,137	5,584	92,463
Non-equity retirement assets	32,501	1,068	5,129	62,050
Equities and other investments	59,939	476	2,595	118,680
Pooled investment funds	25,076	78	601	49,959
Directly held stocks and other	34,863	399	1,994	68,721
Real assets	378,425	33,778	114,198	684,069
Real estate (primary residence)	218,420	22,123	86,222	383,154
Vehicles and other real assets	160,004	11,655	27,976	300,915

Notes: Entries are weighted means using cross-sectional weights and are reported in 2024 dollars. Financial assets are decomposed into liquid assets, fixed-income assets, retirement assets, and equities and other investments. Checking and savings accounts combine checking and savings balances. Other liquid assets include prepaid card, call, and money market accounts. CDs and bonds include certificates of deposit, directly held bonds, and savings bonds. Retirement assets are decomposed into equities held in retirement accounts and non-equity retirement assets. Directly held stocks and other include directly held stocks, other managed assets, and other financial assets. Vehicles and other real assets combine vehicles with nonfinancial assets other than the primary residence. Prepaid card balances are only available beginning in 2016, and unavailable values are treated as zero.

results including the 2021 wave, which covers the COVID-19 period.

B.1 MPC Model Derivations

This section derives the results for the model in Section 3.1. I first show that $\Delta y_{i,t+1}$ is a valid instrument for $\Delta y_{i,t}$ under the assumptions of the model. To simplify notation, I drop the household subscript i.

Consider first instrument relevance:

$$\begin{aligned}\operatorname{Cov}(\Delta y_{t+1}, \Delta y_t) &= \operatorname{Cov}\left(\eta_{t+1} + \varepsilon_{t+1} - \varepsilon_t, \eta_t + \varepsilon_t - \varepsilon_{t-1}\right) \\ &= -\operatorname{Var}(\varepsilon_t) = -\sigma_\varepsilon^2 \neq 0.\end{aligned}$$

Thus, Δy_{t+1} is relevant for Δy_t whenever the variance of transitory income shocks is positive. Next, consider the exclusion restriction. Equation (1) can be rewritten as

$$\Delta c_t = \phi \Delta y_t + v_t, \qquad v_t \equiv (\psi - \phi)\eta_t + \Delta u_t.$$

Under the assumptions of the model,

$$\mathrm{Cov}(\Delta y_{t+1}, v_t) = 0,$$

because future income growth is orthogonal to the current permanent innovation η_t and to the consumption innovation Δu_t. Hence, Δy_{t+1} satisfies the exclusion restriction.

The IV estimand is therefore

$$\phi^{IV} = \frac{\mathrm{Cov}(\Delta c_t, \Delta y_{t+1})}{\mathrm{Cov}(\Delta y_t, \Delta y_{t+1})}.$$

Using equation (1),

$$\begin{aligned}\mathrm{Cov}(\Delta y_{t+1}, \Delta c_t) &= \mathrm{Cov}\left(\eta_{t+1} + \varepsilon_{t+1} - \varepsilon_t, \psi\eta_t + \phi(\varepsilon_t - \varepsilon_{t-1}) + \Delta u_t\right) \\ &= -\phi\sigma_\varepsilon^2.\end{aligned}$$

Consequently,

$$\phi^{IV} = \frac{-\phi\sigma_\varepsilon^2}{-\sigma_\varepsilon^2} = \phi.$$

B.2 Testing Empirical Model Assumptions

The MPC estimates in Section 3 follow Blundell *et al.* (2008). Recall that residual labor income is the sum of a permanent and a transitory component, and that consumption growth responds to each with its own coefficient:

$$\begin{aligned}y_{i,t} &= z_{i,t} + \varepsilon_{i,t}, \\ z_{i,t} &= z_{i,t-1} + \eta_{i,t}, \\ \Delta c_{i,t} &= \psi\eta_{i,t} + \phi\Delta\varepsilon_{i,t} + \Delta u_{i,t},\end{aligned}$$

where ε and η are serially uncorrelated and mutually independent, and u captures consumption innovations, including classical measurement error in consumption. The MPC out of transitory income is estimated as

$$\hat{\phi} = \frac{\mathrm{Cov}(\Delta c_{i,t}, \Delta y_{i,t+1})}{\mathrm{Cov}(\Delta y_{i,t}, \Delta y_{i,t+1})},$$

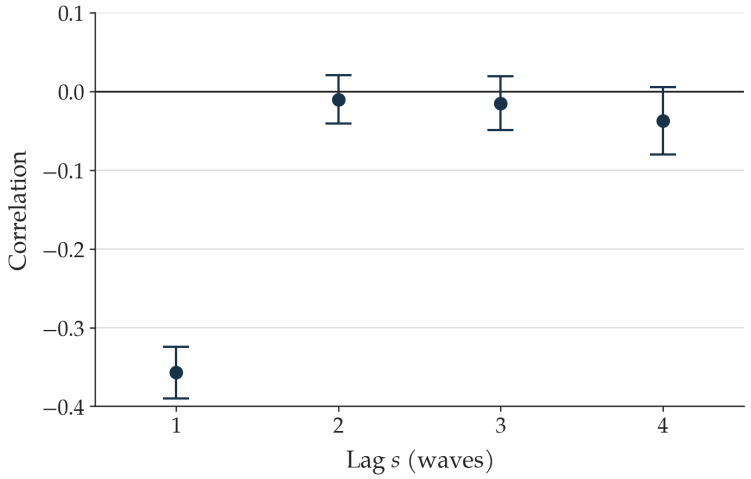
that is, from an IV regression of $\Delta c_{i,t}$ on $\Delta y_{i,t}$ with next wave's income growth as the instrument. With biennial data, Δ is a two-year difference, $t{+}1$ denotes the next wave, and $\eta_{i,t}$ accumulates two years of permanent shocks. The estimator rests on three restrictions that can be checked in the data:

1. The transitory component has no memory beyond one wave (two years), so that $\Delta y_{i,t+1}$ is correlated with $\Delta y_{i,t}$ only through $\varepsilon_{i,t}$.
2. Households have no advance information about future shocks, so that $\Delta c_{i,t}$ is uncorrelated with income growth beyond the next wave.
3. The permanent component of income is close to a random walk, so that $\Delta y_{i,t+1}$ carries no information about $\eta_{i,t}$.

This section examines each restriction in the estimation sample, pooling all years. The data support all three restrictions.

Transitory income has no memory beyond two years. Under the income process, two-year income growth $\Delta y_{i,t} = \eta_{i,t} + \varepsilon_{i,t} - \varepsilon_{i,t-1}$ shares a component with the previous wave's income growth only through the transitory shock $\varepsilon_{i,t-1}$. Its autocorrelation is therefore negative at one wave and zero at every longer lag. Figure B.1 shows exactly this pattern. The first-order autocorrelation is -0.36 and precisely estimated, while the autocorrelations at two, three, and four waves (four, six, and eight years) are small (-0.01, -0.02, and -0.04) and statistically indistinguishable from zero at the 5 percent level (p-values of 0.52, 0.39, and 0.09). Transitory shocks therefore wash out within two years. The first-order autocovariance is also the first stage of the estimator, and it is strong: the first-stage F-statistic is 833 in the full sample and between 236 and 352 within wealth groups, far above conventional weak-instrument thresholds.

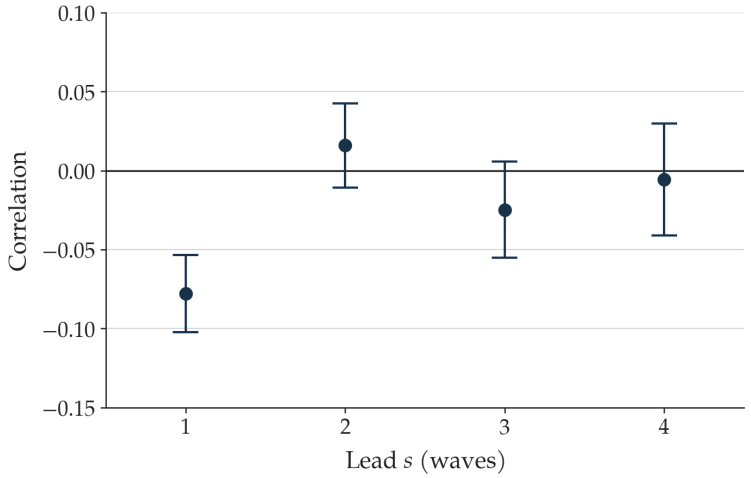
Figure B.1: Autocorrelation of income growth



Notes: Correlation between two-year growth in log residual labor income, $\Delta y_{i,t}$, and its value s waves ($2s$ years) earlier, $\Delta y_{i,t-s}$. Both observations belong to the MPC estimation sample, extended to include the 2021 wave (1999–2021). Moments are pooled across years after removing year means. Bars are 95% confidence intervals based on standard errors clustered by household.

No advance information. If households learn about shocks only when they occur, consumption growth at t is correlated with next wave's income growth only through the transitory shock $\varepsilon_{i,t}$, which it shares with $\Delta y_{i,t+1}$, and it is uncorrelated with income growth further ahead. If instead households anticipated future shocks and acted on that knowledge, consumption today would move with income changes several waves ahead. Figure B.2 shows the correlation between $\Delta c_{i,t}$ and $\Delta y_{i,t+s}$. At one wave, the correlation is -0.08: this is the identifying moment in the numerator of the IV estimate. It is negative and precisely estimated, as the model predicts when $\phi > 0$. At two, three, and four waves, the correlations (0.02, -0.02, and -0.01) are not statistically significant (p-values of 0.24, 0.11, and 0.76). Consumption does not anticipate income changes four, six, or eight years ahead.

Figure B.2: Consumption growth and future income growth



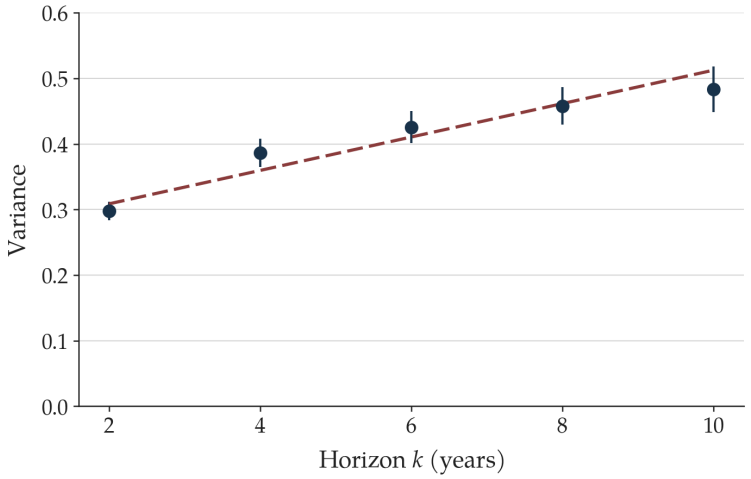
Notes: Correlation between two-year growth in log residual consumption, $\Delta c_{i,t}$, and two-year growth in log residual labor income s waves ($2s$ years) later, $\Delta y_{i,t+s}$. The value at $s = 1$ is the identifying moment of the estimator. Both observations belong to the MPC estimation sample, extended to include the 2021 wave (1999–2021), so that $\Delta c_{i,t}$ covers the waves from 2001 to $2021 - 2s$. Moments are pooled across years after removing year means. Bars are 95% confidence intervals based on standard errors clustered by household.

The permanent component is a random walk. The variance of the k-year difference in income is

$$\mathrm{Var}(y_{i,t} - y_{i,t-k}) = k\,\sigma_\eta^2 + \mathrm{Var}(\varepsilon_{i,t} - \varepsilon_{i,t-k}), \tag{16}$$

where σ_η^2 now denotes the annual variance of permanent shocks. Once k exceeds the memory of the transitory component, the second term is constant, so the variance rises linearly with the horizon, with slope σ_η^2. Figure B.3 shows the variance of long differences in the estimation sample, for horizons of two to ten years. The variance rises steadily with the horizon, and a straight line fits the profile closely: the slope is 0.025 per year, in the range of the permanent variances estimated by Blundell *et al.* (2008), and the R^2 is 0.91. Permanent shocks therefore accumulate as a random walk predicts, so next wave's income growth carries no information about past permanent shocks.

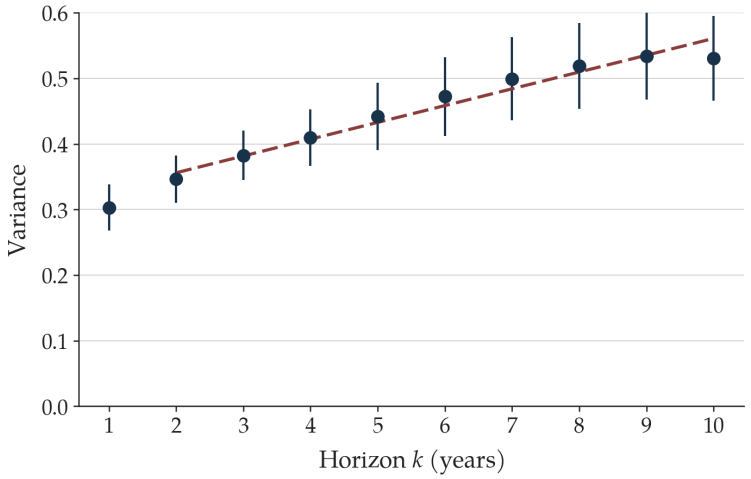
Figure B.3: Variance of long differences of income



Notes: Variance of $y_{i,t} - y_{i,t-k}$, the k-year difference in log residual labor income, for $k = 2, 4, \ldots, 10$. Both observations belong to the MPC estimation sample, extended to include the 2021 wave (1999–2021), and income is demeaned by year. Bars are 95% confidence intervals based on standard errors clustered by household. The dashed line is a weighted least-squares fit, with weights equal to the inverse sampling variance.

The annual waves of the PSID (1970–1997) provide a robustness check with horizons of every length between one and ten years. Applying the same income definition, residualization, and sample restrictions as in the estimation sample, Figure B.4 shows the same pattern. From two years onward, the variance rises linearly, with a slope of 0.026 per year, very close to that of the estimation sample, and an R^2 of 0.96. The one-year difference lies below the line, as an annual MA(1) transitory component implies. The random-walk behavior of the permanent component is therefore a stable feature of the PSID across three decades, rather than a property of the recent waves.

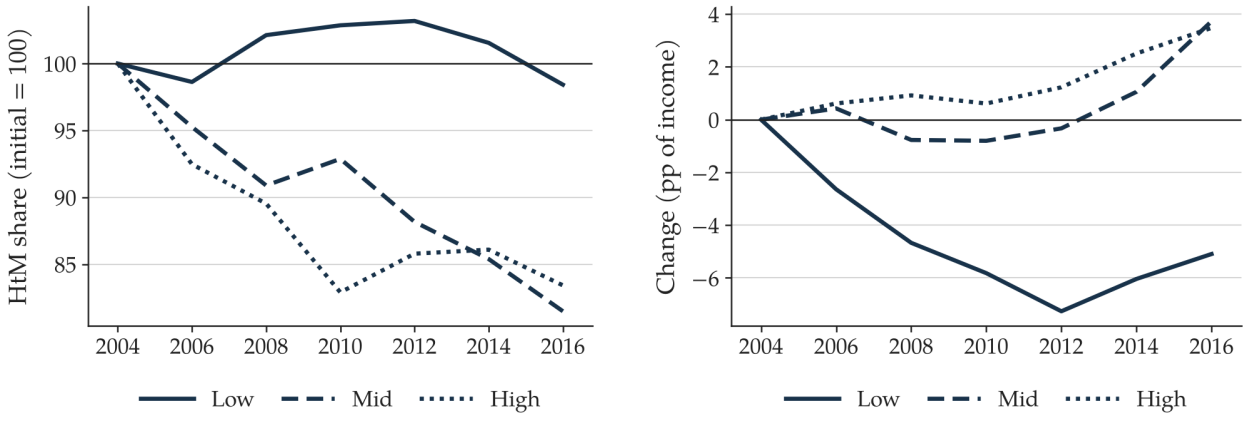
Figure B.4: Variance of long differences of income: annual waves, 1970–1997



Notes: Variance of $y_{i,t} - y_{i,t-k}$, the k-year difference in log residual labor income, for $k = 1, \ldots, 10$, in the annual waves of the PSID (1970–1997), with the income definition, residualization, and sample restrictions of the MPC estimation sample. Both observations of every pair belong to the sample, and income is demeaned by year. Bars are 95% confidence intervals based on standard errors clustered by household. The dashed line is a weighted least-squares fit for $k \geq 2$, with weights equal to the inverse sampling variance. The horizon $k = 1$ is excluded from the fit because an annual MA(1) transitory component makes $\mathrm{Var}(\varepsilon_{i,t} - \varepsilon_{i,t-k})$ constant only from $k = 2$ onward.

B.3 HtM Evolution

Figure B.5 shows the evolution of HtM households and of liquid wealth for each wealth group, discussed in Section 3.4.



(a) Share of HtM Households

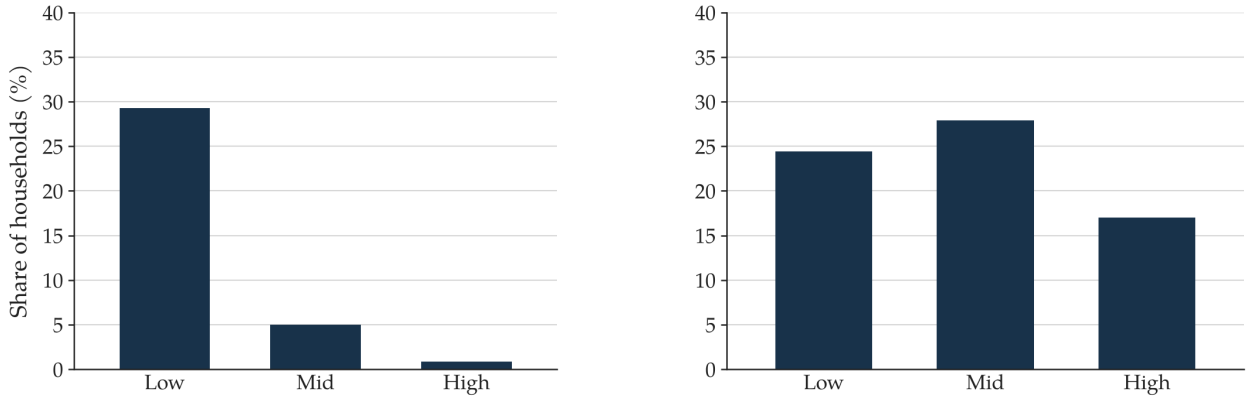
(b) Liquid Wealth

Figure B.5: HtM Households and Liquid Wealth Across Wealth Groups

Notes: PSID. Panel (a): share of HtM households (Kaplan *et al.*, 2014) relative to the initial window (= 100). Panel (b): change in liquid wealth relative to income since the initial window, in percentage points of income. Four-wave windows plotted at their midpoint.

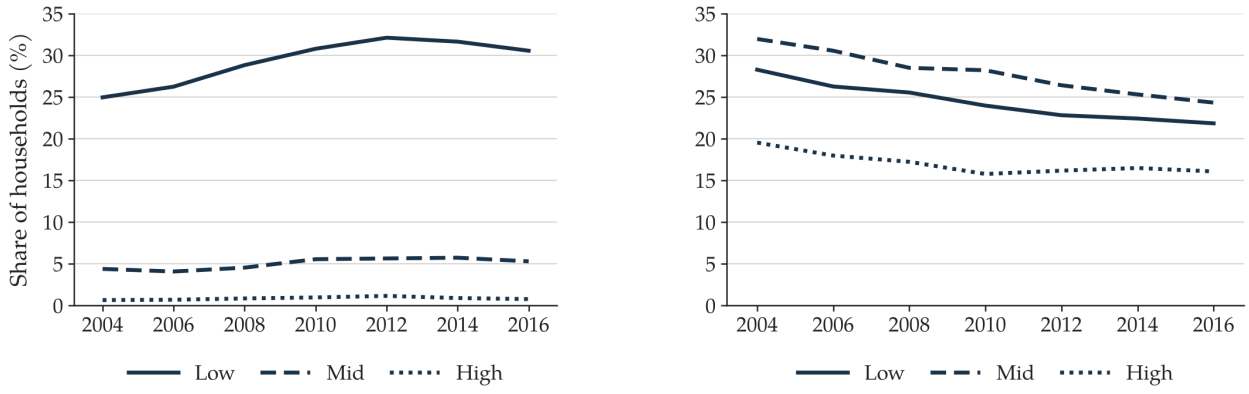
The rest of this section shows how poor and wealthy HtM households are distributed across wealth groups. Figure B.6 shows the share of poor and wealthy HtM households in each wealth group, and Figure B.7 shows how these shares evolved over time. There are two main takeaways. First, the share of wealthy HtM households is larger in the mid-wealth group than in the low-wealth group, which is consistent with the higher MPC of mid-wealth households at the beginning of the sample. Second, consistent with Section 3.4, the share of poor HtM households has increased, while the share of wealthy HtM households has decreased.

Figure B.6 is key to interpreting Figure 3.1: fiscal transfers targeted toward low-wealth households are not always the most effective tool, because these households are not necessarily the ones with the highest MPC. Moreover, Kaplan *et al.* (2014) show that wealthy HtM households typically have high MPCs, which makes this fact even more relevant.



(a) Poor HtM (b) Wealthy HtM

Figure B.6: Poor and Wealthy HtM by Wealth Groups



(a) Poor HtM (b) Wealthy HtM

Figure B.7: Poor and Wealthy HtM Evolution Across Wealth Groups

B.4 Housing Adjustment After Large Transitory Losses

This section reports the change in spending after a large transitory income loss, which complements Section 3.2, and the role of unemployment. Table B.1 reports the estimates: the change in labor income, non-housing consumption, housing spending and housing adjustment after a large transitory loss (Panel A), and the response of spending to large transitory shocks (Panel B). Low-wealth households cut their housing spending the most, while mid-wealth households cut their non-housing consumption the most.

Table B.1: Responses to Large Transitory Income Losses

	Low-wealth	Mid-wealth	High-wealth
Panel A. Large drop in labor income that recovers, relative to stable income			
Labor income at t (%)	-74.5	-60.1	-65.8
Labor income at $t+1$ (%)	65.5	54.2	63.9
Non-housing consumption (%)	-5.4	-9.9	-1.4
	(4.5)	(3.6)	(3.8)
Housing spending (%)	-5.5	-3.4	-2.9
	(5.2)	(4.7)	(4.1)
Adjusts housing (pp)	19.1	9.7	3.6
	(3.5)	(2.6)	(1.8)
Households with the event	201	192	200
Panel B. Response to large transitory income shocks (2SLS)			
Non-housing consumption	0.22	0.20	0.10
	(0.06)	(0.08)	(0.07)
Housing spending	0.16	0.03	0.13
	(0.07)	(0.12)	(0.09)
Total consumption	0.20	0.16	0.10
	(0.05)	(0.07)	(0.06)
Adjusts housing	0.00	0.06	-0.02
	(0.05)	(0.05)	(0.03)

Notes: PSID, 1999–2019, MPC sample, pairs of consecutive waves. Panel A compares households whose labor income falls by more than 20% between $t-1$ and t and recovers by more than 15% between t and $t+1$ with households whose labor income changes by less than 10% in both waves; family size is unchanged. Entries are coefficients on the event indicator in a regression with year fixed effects (percent changes in log points, percentage points for housing adjustment). A household adjusts its housing if it rented at $t-1$ and at t pays at least 20% less in rent or no longer pays rent, or if it owned at $t-1$ and no longer owns at t. Panel B reports the response to large transitory shocks from the estimator of Section 3.1, allowing the coefficient to differ for the quarter of observations with the largest absolute income growth in the next wave, with year fixed effects. Income and consumption are residualized without controls for housing tenure, mortgage status and family size. Standard errors clustered by household in parentheses.

Most large transitory losses come with a spell of unemployment. I say that a pair of waves has a spell of unemployment if the head or the spouse is temporarily laid off or

looking for work at the interview of the first wave, the second wave, or the following wave. A spell of unemployment is part of 31 percent of the large transitory losses of low-wealth households, 24 percent of those of mid-wealth households, and 16 percent of those of high-wealth households. Figure B.8 reports the response of non-housing consumption, housing spending and total consumption to large transitory shocks (Panel B of Table B.1), separately for pairs of waves with and without a spell of unemployment. There are two main takeaways. First, large changes in income outside spells of unemployment do not move consumption or housing in any group. These changes are likely a mix of measurement error and changes that households expect or that are insured. Second, in spells of unemployment the pattern of Section 3.2 is stronger. Low-wealth households absorb a large part of the loss through housing (0.23 out of a total response of 0.31), while mid-wealth households cut their non-housing consumption almost twice as much as low-wealth households (0.58 against 0.31). The share of low-wealth households that adjust their housing after a large transitory loss rises by 30 percentage points when the loss comes with a spell of unemployment, and by 15 percentage points otherwise, against 13 and 9 percentage points for mid-wealth households.

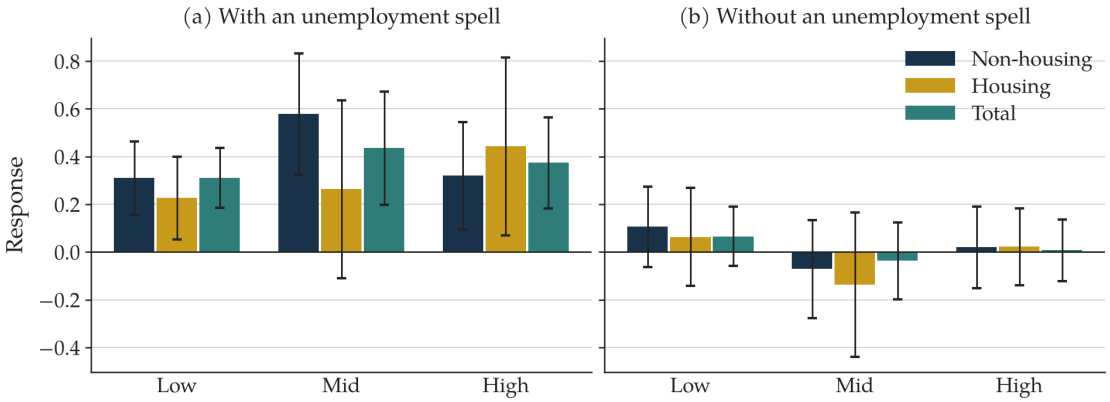
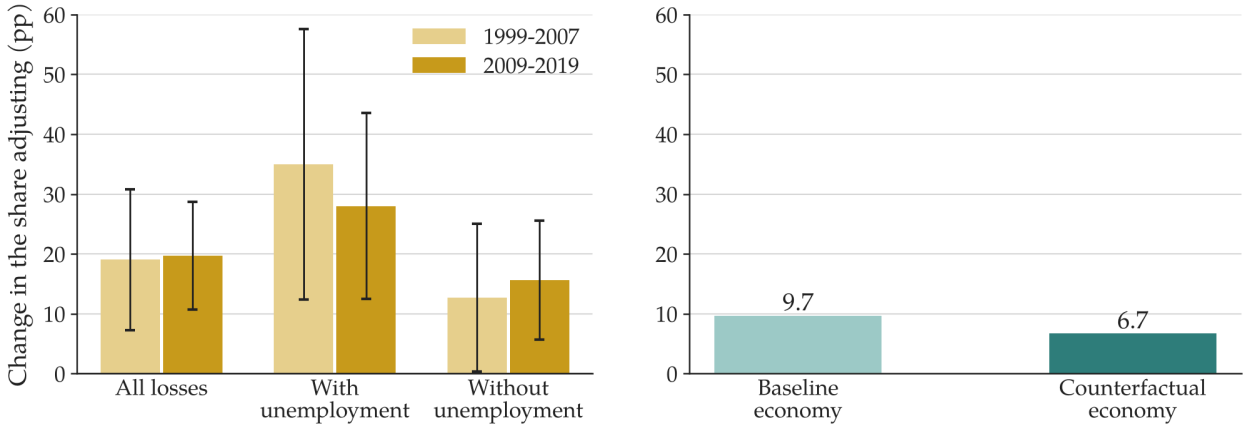


Figure B.8: Responses to Large Transitory Shocks, With and Without Unemployment

Notes: The figure reports the response of non-housing consumption, housing spending, and total consumption to large transitory income shocks, estimated as in Panel B of Table B.1, separately for pairs of waves in which the head or the spouse is temporarily laid off or looking for work at the interview of either wave or of the following wave (panel a), and for the other pairs (panel b). Sample: 1999–2019. Error bars denote 95% confidence intervals computed from standard errors clustered by household.

Finally, I study how the housing margin of low-wealth households changed over the sample. Figure B.9 repeats the comparison of Panel A of Table B.1 for low-wealth households, before the Great Recession (1999–2007) and after it (2009–2019), and compares it with the model. The estimates are noisy, because each period has few large transitory losses, but they support the mechanism of Section 6.2. There are two main takeaways. First, losses that come with a spell of unemployment became more common: they are 29 percent of

the large transitory losses of low-wealth households before 2008 and 33 percent after. Second, in spells of unemployment, which are the losses that move consumption, low-wealth households became less likely to adjust their housing. After such a loss, the share of low-wealth households that adjust their housing rises by 35 percentage points before 2008 and by 28 percentage points after. Outside spells of unemployment, the adjustment does not fall, and, because more losses come with unemployment, the adjustment after any large loss is about the same in both periods (19 and 20 percentage points). The model predicts the same decline (panel (b)). In the counterfactual economy, higher rents push renters to the bottom of the rental ladder, where they have less room to move down, and the increase in the share of low-wealth households that adjust their housing after a large transitory loss falls from 9.7 to 6.7 percentage points. In the model, as in the data, housing provides less insurance at the bottom after the rise in rents.



(a) Data (b) Model

Figure B.9: Housing Adjustment of Low-Wealth Households After a Large Transitory Loss, Before and After 2008

Notes: Panel (a) reports the change in the share of low-wealth households that adjust their housing for households whose labor income falls by more than 20 percent between two consecutive PSID waves and recovers by more than 15 percent at the next wave, relative to households whose labor income changes by less than 10 percent in both waves, as in Panel A of Table B.1. Initial period: pairs of waves starting in 1999–2007. Final period: pairs of waves starting in 2009–2019. A loss comes with unemployment if the head or the spouse is temporarily laid off or looking for work at the interview of either wave or of the following wave; the control group excludes these households. Estimates control for year fixed effects. Error bars denote 95% confidence intervals computed from standard errors clustered by household. Panel (b) reports the same comparison in the simulated panels of the baseline and the counterfactual economy, for all large transitory losses. Wealth groups: terciles of household-average net worth.

B.5 Additional Results and Robustness

This section reports the aggregate MPC, the MPC under alternative definitions of the wealth groups, and the MPC by housing tenure.

Aggregate MPC. Figure B.10 shows the evolution of the aggregate MPC. Each point estimate represents four waves (eight years). There is some cyclicality, consistent with Gross *et al.* (2020), but the aggregate MPC remained stable at around 0.2–0.25 over the sample period.

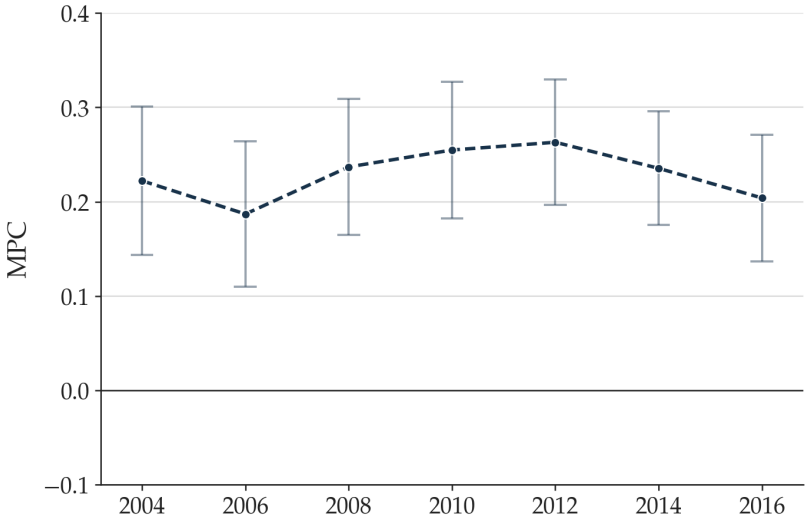


Figure B.10: Aggregate MPC over time

Notes: PSID. Aggregate MPC estimated on four consecutive waves (an eight-year window), plotted at the midpoint of the window. 95% confidence intervals.

Definition of the wealth groups. In the main text, each household belongs to the same wealth group in every wave, based on its average net worth over the whole sample (Section 2). Table B.2 reports the MPC in the initial and final windows when I instead split households into terciles within each window (Panel B), and within each window and five-year age group (Panel C). The last definition compares households of the same age, so it is not affected by the aging of the sample. The main result does not depend on how I define the groups. Under the three definitions, the MPC of mid-wealth households falls by between 0.15 and 0.22, the MPC of low-wealth households rises, and low-wealth households have the highest MPC in the final window. The rise at the bottom is smaller when groups are defined within each window, but the ordering of the MPC reverses in all three cases.

Differences between groups. Table B.3 tests whether the change in the MPC differs between groups. The divergence between low- and mid-wealth households is the most precise result: the difference between their changes is 0.36, with a standard error of 0.13, significant at the 1 percent level. The difference between low- and high-wealth households is significant at the 10 percent level, and so is the difference between renters and owners. The differences between age groups and between income groups are not statisti-

Table B.2: MPC by Wealth Group Under Alternative Definitions of the Groups

	Low-wealth	Mid-wealth	High-wealth
Panel A. Terciles over the whole sample (baseline)			
Initial window, 1999–2007	0.13	0.38	0.21
Final window, 2011–2019	0.26	0.16	0.09
Change	+0.14	−0.22	−0.11
	(0.08)	(0.10)	(0.10)
Panel B. Terciles within each window			
Initial window, 1999–2007	0.19	0.32	0.16
Final window, 2011–2019	0.25	0.17	0.14
Change	+0.06	−0.15	−0.03
	(0.08)	(0.09)	(0.10)
Panel C. Terciles within each window and age group			
Initial window, 1999–2007	0.21	0.31	0.14
Final window, 2011–2019	0.29	0.17	0.11
Change	+0.08	−0.15	−0.04
	(0.09)	(0.09)	(0.10)

Notes: PSID. MPC estimated as in Section 3.1 in the initial window (four waves, 1999–2007) and in the final window (four waves, 2011–2019). Panel A uses the groups of the main text. Panel B splits households into terciles of their average net worth within each window. Panel C splits households into terciles of their average net worth within each window and five-year age group of the head. Standard errors of the change in parentheses, treating the two windows as independent samples.

cally different from zero.

Housing tenure. Figure B.11 shows the evolution of the MPC for renters and owners. The MPC of renters increased over time, while the MPC of owners declined.

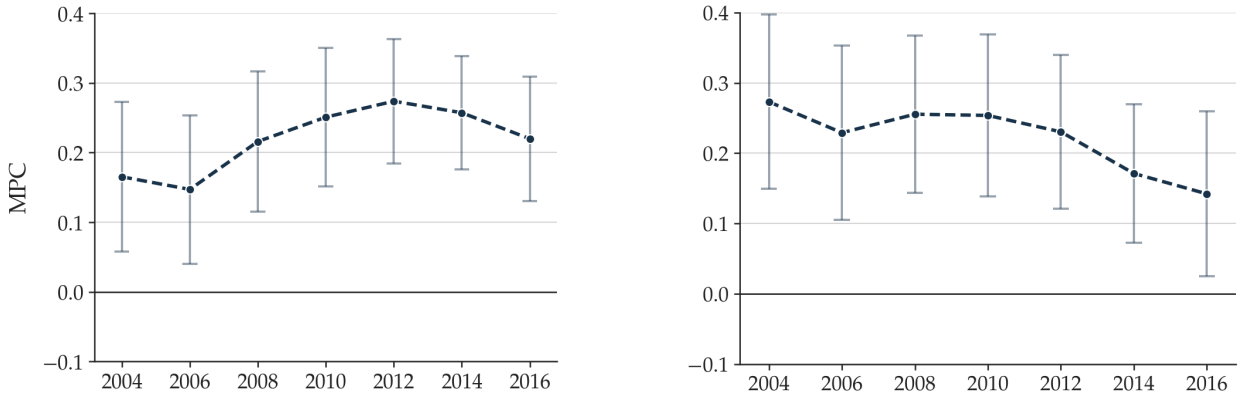
Figure B.12 reports the MPC of renters and owners within each wealth group. It is consistent with Figure 3.1, but noisier, because splitting the sample into more groups reduces the number of observations in each. The clearest trends are for low-wealth renters and for mid-wealth renters and owners; the estimates for the remaining groups are relatively noisy.

In the model, low-wealth renters and mid-wealth owners are the most relevant groups for the mechanism. The trends for these groups are in line with the model: the MPC of low-wealth renters increases, while the MPC of mid-wealth households declines, for both renters and owners.

Table B.3: Differences in the Change in the MPC Between Groups

	Change in the MPC		Difference	(s.e.)
	First group	Second group		
A. Wealth				
Low minus mid wealth	+0.14	−0.22	+0.36***	(0.13)
Low minus high wealth	+0.14	−0.11	+0.25*	(0.13)
Mid minus high wealth	−0.22	−0.11	−0.11	(0.14)
B. Housing tenure				
Renters minus owners	+0.05	−0.13	+0.19*	(0.11)
C. Age of the head				
25–39 minus 40–55	+0.05	−0.10	+0.15	(0.11)
D. Income				
Low minus mid income	−0.04	−0.09	+0.05	(0.13)
Low minus high income	−0.04	+0.06	−0.10	(0.14)

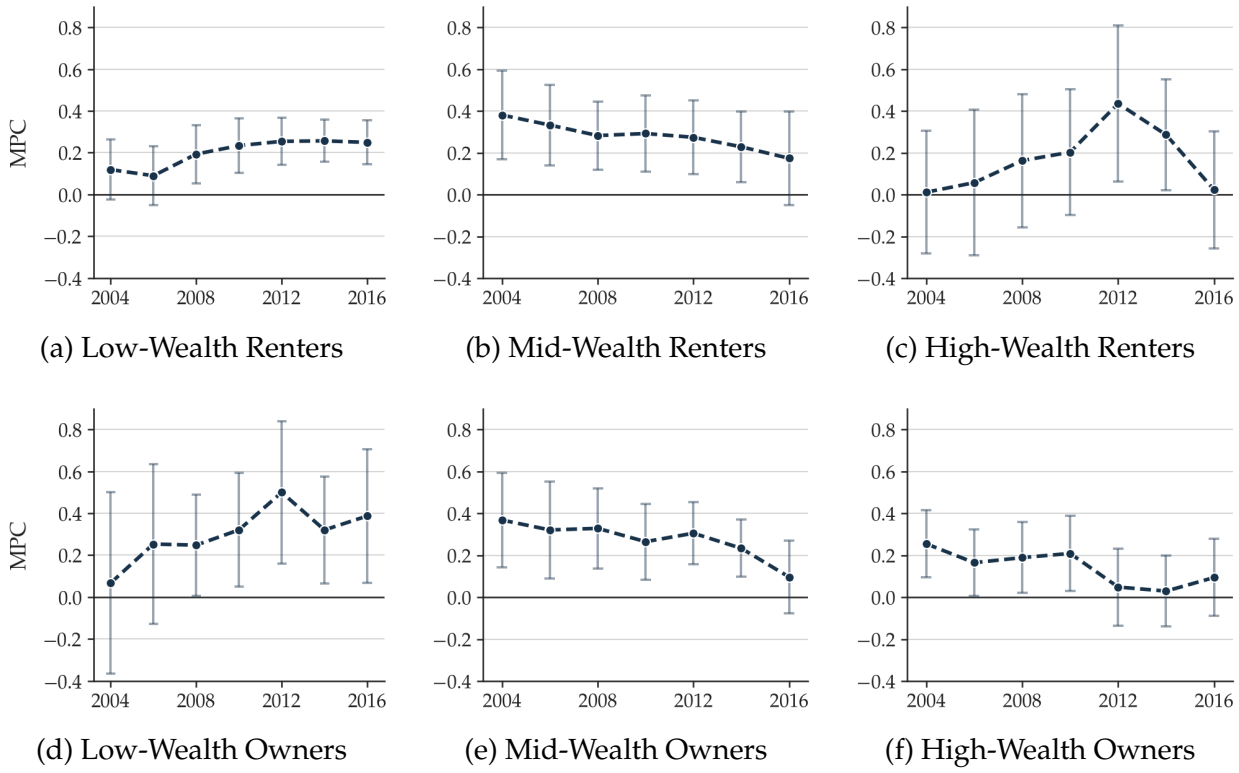
Notes: PSID. Change in the MPC between the initial period (1999–2007) and the final period (2011–2019), estimated as in Section 3.1, and the difference between the changes of two groups. The standard error of the difference treats the four estimates as independent. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.



(a) Renters (b) Owners

Figure B.11: MPC Evolution Across Housing Tenure

Notes: The figure reports estimates of the marginal propensity to consume (MPC) for renters and owners by survey year. Each point estimate is obtained using four consecutive PSID waves (an eight-year window) and is plotted at the midpoint of the corresponding sample period. Error bars denote 95% confidence intervals.



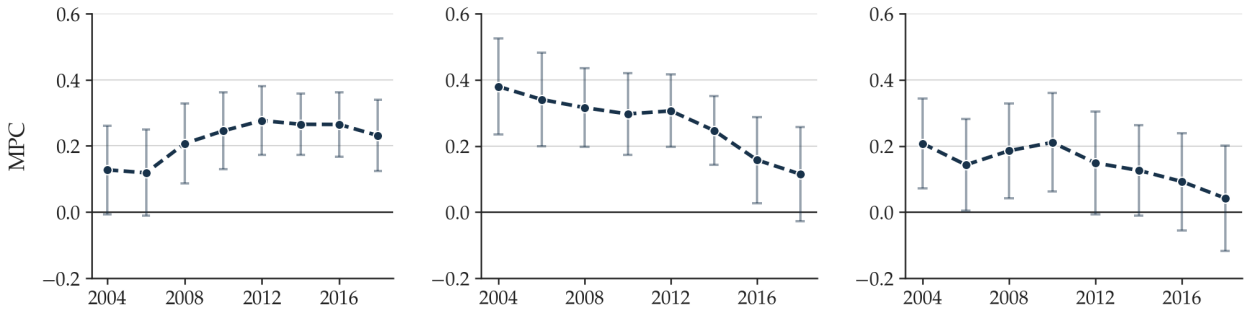
(a) Low-Wealth Renters (b) Mid-Wealth Renters (c) High-Wealth Renters

(d) Low-Wealth Owners (e) Mid-Wealth Owners (f) High-Wealth Owners

Figure B.12: MPC Evolution for Renters and Owners Across Wealth Groups

B.6 PSID with COVID Sample

This section reports the main results including the last available wave, 2021. Figure B.13 shows that the MPC results are robust to including the COVID-19 period: the trends in Figure 3.1 continue in the last observation. The only minor difference is that the MPC of low-wealth households decreases slightly in the last observation, although it remains above its initial value.

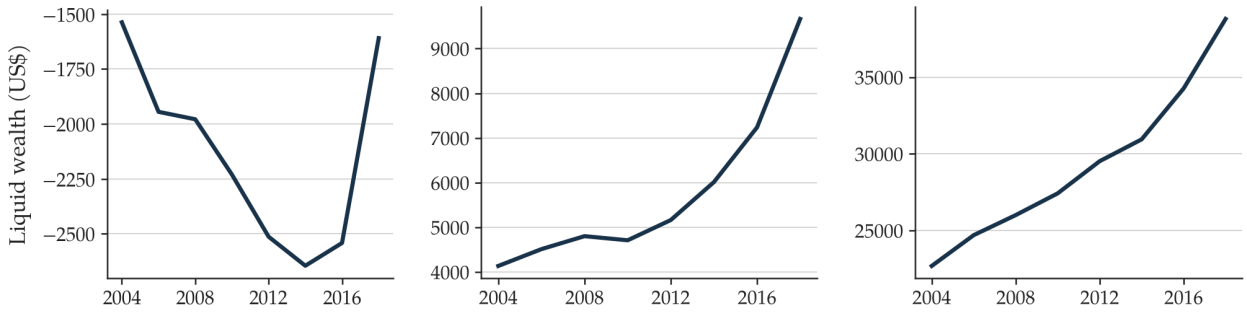


(a) Low Wealth (b) Mid Wealth (c) High Wealth

Figure B.13: MPC Evolution Across Wealth Groups Including COVID

Notes: The figure reports estimates of the marginal propensity to consume (MPC) for each wealth group by survey year, including the 2021 wave. Each point estimate is obtained using four consecutive PSID waves (an eight-year window) and is plotted at the midpoint of the corresponding sample period. Error bars denote 95% confidence intervals.

Figure B.14 shows the evolution of household liquid wealth once the 2021 wave is included. In contrast to the MPC, adding this wave materially changes the picture: the downward trend in the liquid wealth of low-wealth households reverses. The explanation is simple: government transfers. In 2020 and 2021, the federal government distributed several rounds of stimulus payments that temporarily increased household liquidity. This is consistent with evidence from the SCF and other sources, which show that household checking account balances increased substantially during this period. As Figure B.13 shows, however, this temporary increase in liquidity does not change how households respond to income changes.



(a) Low Wealth (b) Mid Wealth (c) High Wealth

Figure B.14: Liquid Wealth Evolution Including COVID

Notes: The figure reports average liquid wealth for each wealth group by survey year, including the 2021 wave. Each point estimate is obtained using four consecutive PSID waves (an eight-year window) and is plotted at the midpoint of the corresponding sample period.

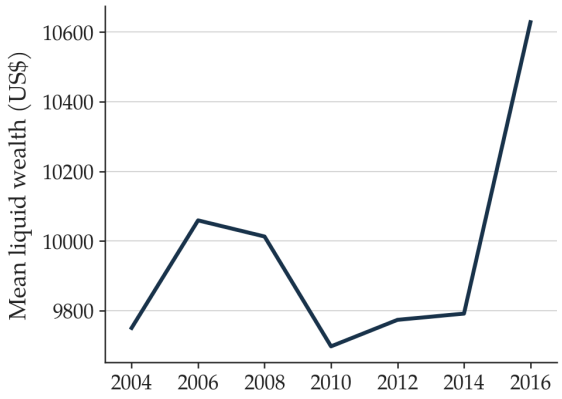
C Stylized Facts Appendix

This appendix provides additional evidence for the evolution of liquid wealth in Section 3.4 and for the stylized facts in Section 4. Appendix C.1 shows that the divergence in liquid wealth also appears in the distribution of the whole sample and across housing tenure, linking liquid wealth to the housing facts. Appendix C.2 documents the aggregate housing-market trends behind the decline in homeownership: rising house prices and a rising price-to-rent ratio. Appendix C.3 shows that the increase in housing costs is robust to using the median instead of the average and to measuring housing payments in the CEX, and documents the evolution of national rents. Finally, Appendix C.4 shows that the shift from fixed-income toward liquid assets in Section 4.3 is robust to measuring portfolio shares relative to total assets rather than financial assets. The increases in house prices and rents documented here, together with the secular decline in interest rates behind this portfolio shift, are the three aggregate price changes that the model takes as given.

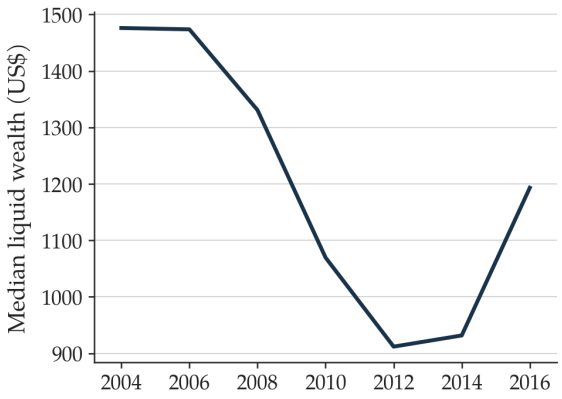
C.1 Liquidity

Figure C.1 shows average and median liquid wealth for the whole PSID sample. The two measures move in opposite directions. Average liquid wealth increased from about $9,750 to about $10,600, with most of the increase at the end of the sample. Median liquid wealth, in contrast, fell from about $1,480 to a trough of about $910 in the early 2010s, and only partially recovered, to about $1,200. The typical household therefore lost liquidity over the sample period, while the gains were concentrated among more liquid households. This is the aggregate counterpart of the divergence across wealth groups in Figure B.5: a stable or rising average masks a decline in liquidity at the bottom of the distribution and an increase further up, just as the stable aggregate MPC in Figure B.10 masks diverging MPCs across wealth groups.

Figure C.2 splits average liquid wealth by housing tenure. Owners hold substantially more liquid wealth than renters, and the gap widened over the sample period. The average liquid wealth of owners rose steadily, from about $13,900 to about $17,600, an increase of roughly 25 percent. The average liquid wealth of renters fluctuated between about $3,600 and $4,200 and increased only modestly, by roughly 15 percent. As a result, the gap between owners and renters widened from about $10,200 to about $13,400. This divergence is consistent with the housing-cost trends in Section 4.2: owners benefited from falling mortgage payments after the GFC, while renters faced rising rents. The modest increase for renters may also reflect a change in the composition of this group. As homeownership became harder to reach (Section 4.1), more mid-wealth households remained renters, and these households hold more liquid wealth than the typical low-wealth renter.



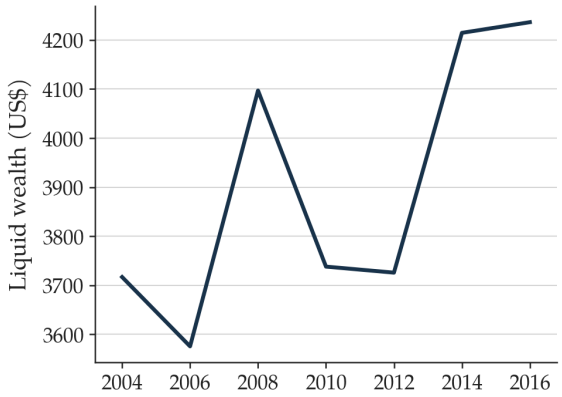
(a) Average



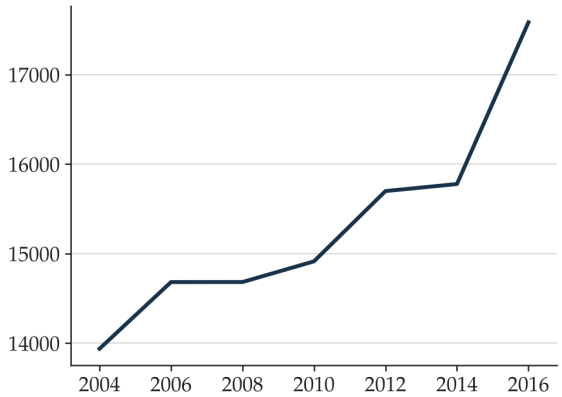
(b) Median

Figure C.1: Average and Median Liquid Wealth

Notes: Panel (a) reports average liquid wealth and panel (b) median liquid wealth for the whole sample by survey year. Each point is obtained using four consecutive PSID waves (an eight-year window) and is plotted at the midpoint of the corresponding sample period.



(a) Renters



(b) Owners

Figure C.2: Average Liquid Wealth of Renters and Owners

Notes: Panel (a) reports average liquid wealth for renters and panel (b) average liquid wealth for owners by survey year. Each point is obtained using four consecutive PSID waves (an eight-year window) and is plotted at the midpoint of the corresponding sample period.

C.2 Homeownership

Figure C.3 reports the probability that a household renting in a given wave owns its home in at least one of the next three waves. The probability fell from almost 60 percent to about 40 percent.

Figure C.4 reports the homeownership rate for each wealth group, discussed in Section 4.1.

Figure C.5 reports the homeownership rate in each PSID wave, without averaging across

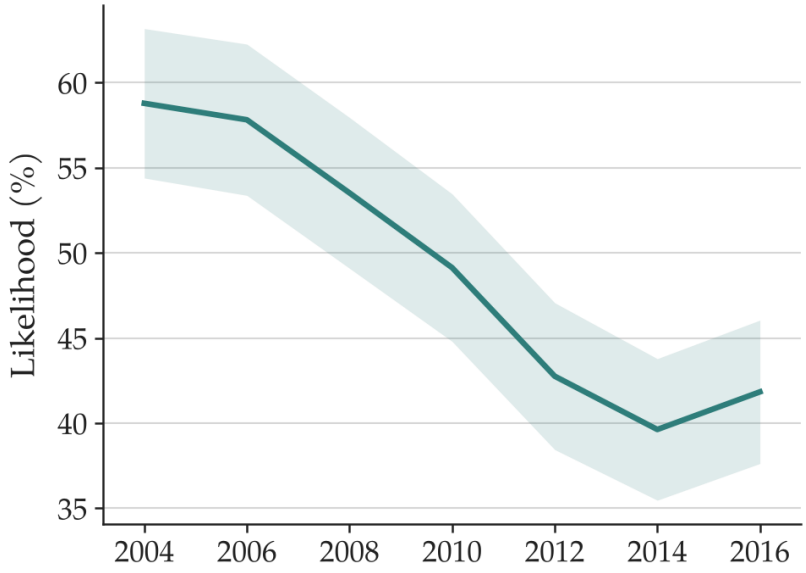
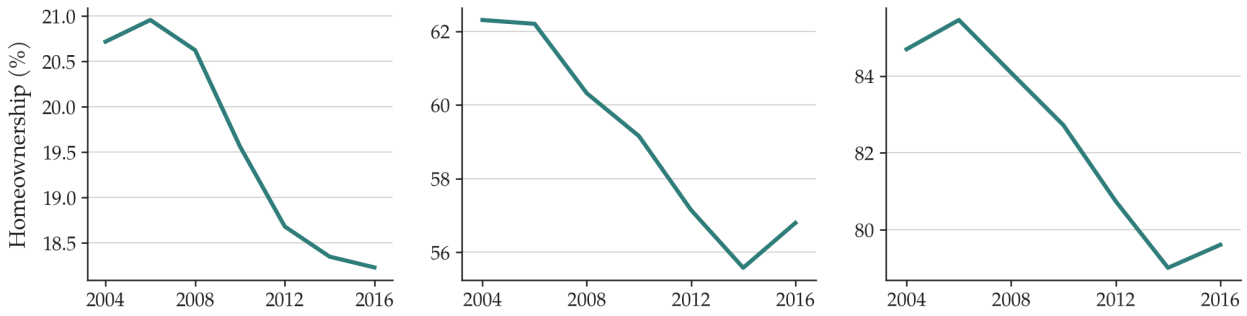


Figure C.3: Likelihood of Becoming an Owner

Notes: PSID. Predictive margins by year from a linear probability model of the six-year renter-to-owner transition on year fixed effects, controlling for observables. 95% confidence interval, standard errors clustered by household.



(a) Low Wealth (b) Mid Wealth (c) High Wealth

Figure C.4: Homeownership Evolution Across Wealth Groups

Notes: The figure reports the homeownership rate for each wealth group. Each point is obtained using four consecutive PSID waves (an eight-year window) and is plotted at the midpoint of the corresponding sample period.

waves. Homeownership was roughly stable at about 59 to 60 percent until 2003 and declined gradually until 2009. It then fell sharply after the GFC, from about 57 percent in 2009 to 51 percent in 2011, and reached a trough of about 47 percent in 2015. It recovered only partially afterward, to about 50 percent in 2019. The rolling-window estimates in Figure 4.1a smooth this profile, but capture the same decline.

Figure C.6 shows the aggregate housing-market trends behind this decline. House prices rose by about 75 percent between 1999 and 2007, fell by about 20 percent during the GFC,

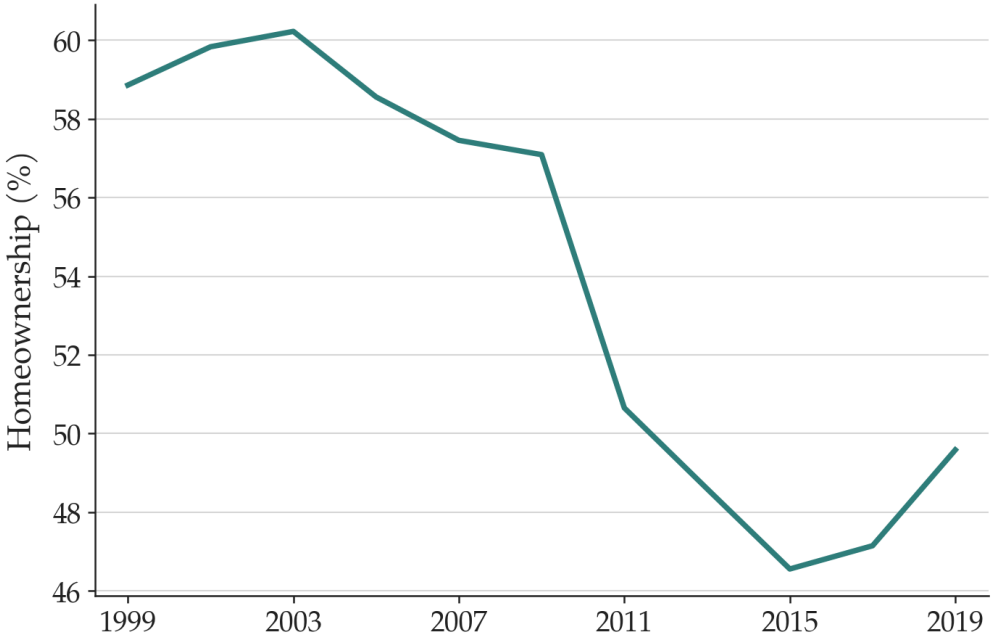


Figure C.5: Homeownership Rate by Wave

Notes: The figure reports the homeownership rate in each PSID wave over the sample period 1999–2019, without averaging across waves.

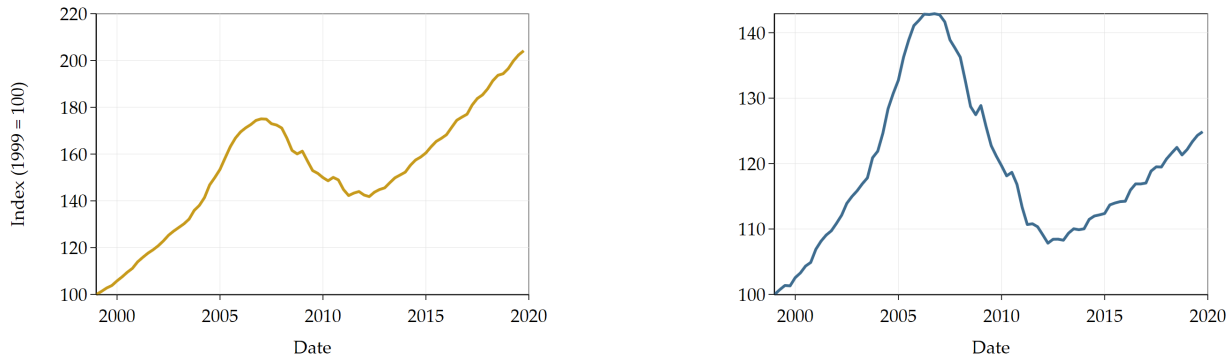
and then rose steadily, ending the sample at about twice their 1999 level. The price-to-rent ratio followed a similar boom-bust pattern: it peaked about 43 percent above its 1999 level in 2006, fell back to about 8 percent above it in 2012, and rose again to about 25 percent above it by 2019. House prices therefore increased faster than rents over the sample period, even though rents also rose (Appendix C.3).

Higher prices relative to rents raise the down payment required to buy a house relative to the cost of renting, which makes it harder for renters to accumulate enough wealth to become owners. This is consistent with the decline in the probability of becoming an owner in Figure C.3, and with the mechanism in the model, in which higher house prices keep potential buyers in the rental market for longer. Lower mortgage rates partially offset higher prices by reducing mortgage payments, but they do not reduce the down payment, which is often the binding constraint for households saving to buy a house.

C.3 Housing Expenditure

Figure C.7 splits housing costs by tenure, discussed in Section 4.2.

Figure C.8 reports the median ratio of housing costs to income for each wealth group. The median is less sensitive to households with very large ratios than the average in panel (b) of Figure 4.1. Its levels are lower, but it shows the same patterns. The median ratio of



(a) House Prices (b) Price-to-Rent Ratio

Figure C.6: House Prices and the Price-to-Rent Ratio

Notes: Panel (a) reports house prices and panel (b) the price-to-rent ratio, both indexed to 1999 = 100, the beginning of the sample period. House prices are measured by the U.S. Federal Housing Finance Agency All-Transactions House Price Index for the United States (USSTHPI), retrieved from FRED, Federal Reserve Bank of St. Louis. The price-to-rent ratio is constructed using the Consumer Price Index for new tenant rent (R-CPI-NTR) from the Bureau of Labor Statistics (BLS).



(a) Rent-to-Income, Renters (b) Mortgage Costs-to-Income, Owners

Figure C.7: Housing Costs-to-Income by Tenure

Notes: PSID. Average ratio of rent to income for renters (a) and of mortgage payments to income for owners (b), by wealth group. Four-wave windows plotted at their midpoint.

low-wealth households rose from about 30 to 34 percent of income and remained high throughout the sample. For mid- and high-wealth households, the median ratio rose until the GFC, to about 29 and 27 percent of income, respectively, and declined afterward.

Figure C.9 shows the evolution of national rents, measured by the BLS new tenant rent index. Rents increased steadily over the sample period, by about 60 percent between 1999 and 2019. Unlike house prices, rents did not fall after the GFC; they only flattened briefly between 2008 and 2010. This steady increase is the aggregate counterpart of the rising rent burden of low-wealth renters in Figure C.7a.

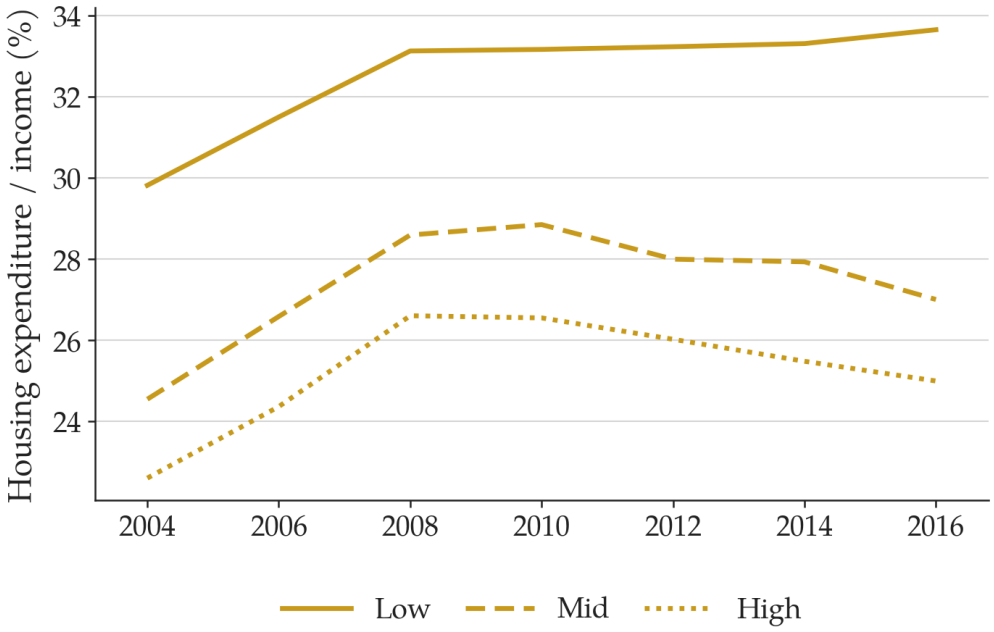


Figure C.8: Median Housing Costs-to-Income Across Wealth Groups

Notes: The figure reports the median ratio of housing costs to income for each wealth group. Housing costs consist of rent for renters and mortgage payments for owners. Each point is obtained using four consecutive PSID waves (an eight-year window) and is plotted at the midpoint of the corresponding sample period.

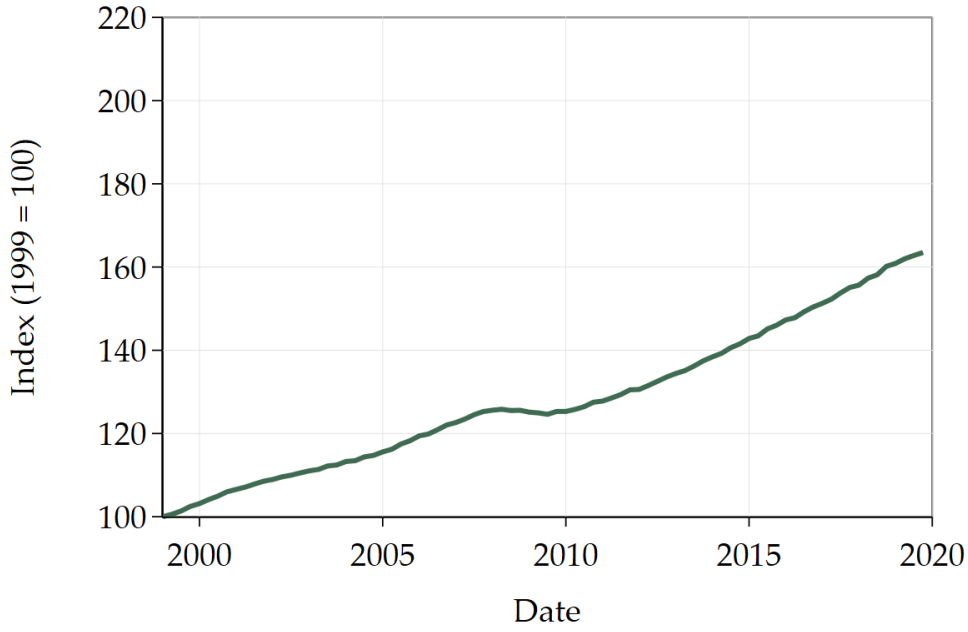


Figure C.9: Rents over Time

Notes: The figure reports the Consumer Price Index for new tenant rent (R-CPI-NTR) from the Bureau of Labor Statistics (BLS), indexed to 1999 = 100, the beginning of the sample period.

Finally, Figure C.10 reports rent-to-income and mortgage payments-to-income in the CEX, indexed to the beginning of the sample. The CEX measures housing payments more precisely than the PSID, and at a quarterly frequency. The two series moved together until the housing boom, when both rose. After 2008, however, they diverged sharply. Rent-to-income kept rising, peaked about 30 percent above its initial level around 2011, and remained about 20 percent above it at the end of the sample. Mortgage payments-to-income, in contrast, fell steadily after 2008 and ended the sample about a third below its initial level. These patterns confirm the PSID evidence in Figures C.7a and C.7b: after the GFC, renters faced rising housing costs relative to income, while owners benefited from lower mortgage payments as mortgage rates declined.

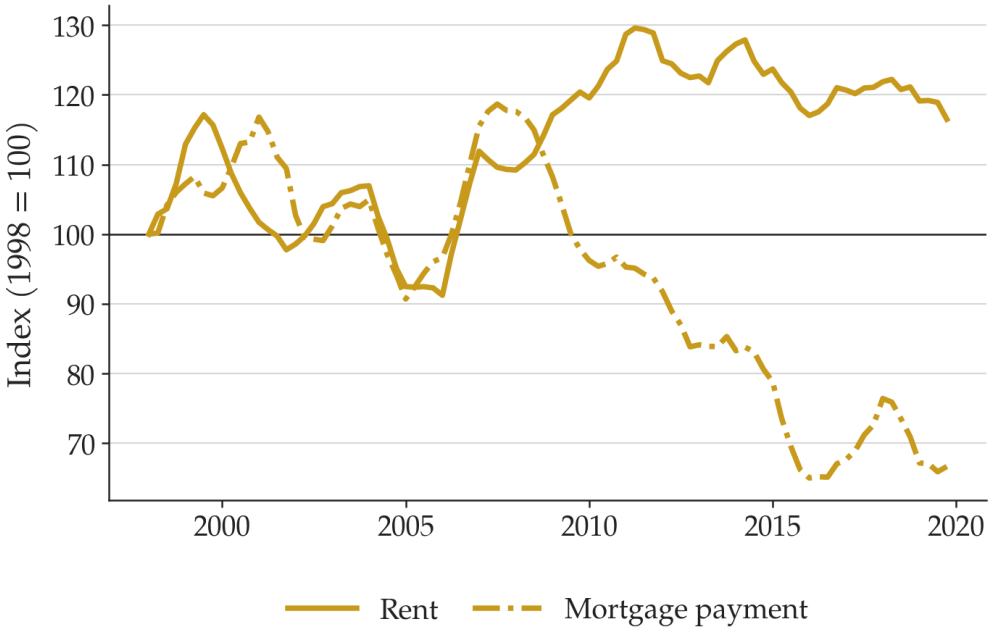


Figure C.10: Rent and Mortgage Payments-to-Income in the CEX

Notes: The figure reports the ratio of rent to income for renters and the ratio of mortgage payments to income for owners in the CEX, both indexed to 1999 = 100, the beginning of the sample period.

C.4 Portfolio Composition

Table C.1, discussed in Section 4.3, measures portfolio shares relative to financial assets.

A potential concern is that these shares reflect changes in the size of financial assets relative to the rest of the balance sheet, rather than substitution across financial assets. For example, rising house prices increase the weight of housing in household portfolios, which could mechanically change the composition of the remaining financial portfolio. Figure C.11 addresses this concern by measuring the share of each financial asset in total assets, which also include real assets such as housing and vehicles.

Table C.1: Composition of Financial Assets, 1998 and 2019 (SCF)

	Share of financial assets (%)			Share of fixed-income assets (%)
	1998	2019	Change (pp)	
Asset classes				
Liquid assets	36.4	50.4	+14.0	
Fixed-income assets	15.3	4.7	−10.5	
Retirement assets	33.5	36.8	+3.3	
Equities	14.9	8.1	−6.8	
Fixed-income assets				
Cash value of life insurance	10.0	3.6	−6.4	52.4
Savings bonds	2.1	0.4	−1.7	28.0
Certificates of deposit	2.7	0.7	−2.0	17.2
Directly held bonds	0.4	0.0	−0.4	2.4

Notes: Average share of each asset in financial assets, heads aged 25–55. Shares are computed for each household and then averaged across households with cross-sectional weights. The last column reports the average share of each component in fixed-income assets among households that hold fixed-income assets, pooling all waves. 1998 is the SCF wave closest to the beginning of the PSID sample.

The same substitution appears. Between 1998 and 2019, the share of liquid assets in total assets increased by about 4 percentage points and the share of retirement assets by about 1.5 percentage points, while the shares of fixed-income assets and equities declined by about 4 and 2.5 percentage points, respectively. These changes are smaller than in Table C.1 because financial assets represent only part of total assets (Table A.2). The shift from fixed-income to liquid assets, however, remains clear: the gain in the liquid share roughly mirrors the loss in the fixed-income share.

The timing is also informative. The liquid share of total assets fell slightly during the housing boom, as rising house prices increased the weight of real assets in household portfolios. It then rose sharply between 2010 and 2016, when interest rates were close to zero, and remained high afterward. The decline in fixed-income assets, in contrast, was concentrated in the early 2000s and continued gradually afterward. Taken together, the four series imply that the share of financial assets in total assets fell by about 6 percentage points by 2004 and had almost fully recovered by 2016. The increase in portfolio liquidity is therefore not an artifact of the denominator: at the end of the sample, households held a larger share of their total assets in liquid form than at the beginning.

D Model Appendix

This appendix complements Sections 5 and 6. I first describe how I solve the model, compare it with the data, and explain how I measure the transitory income shock. I then

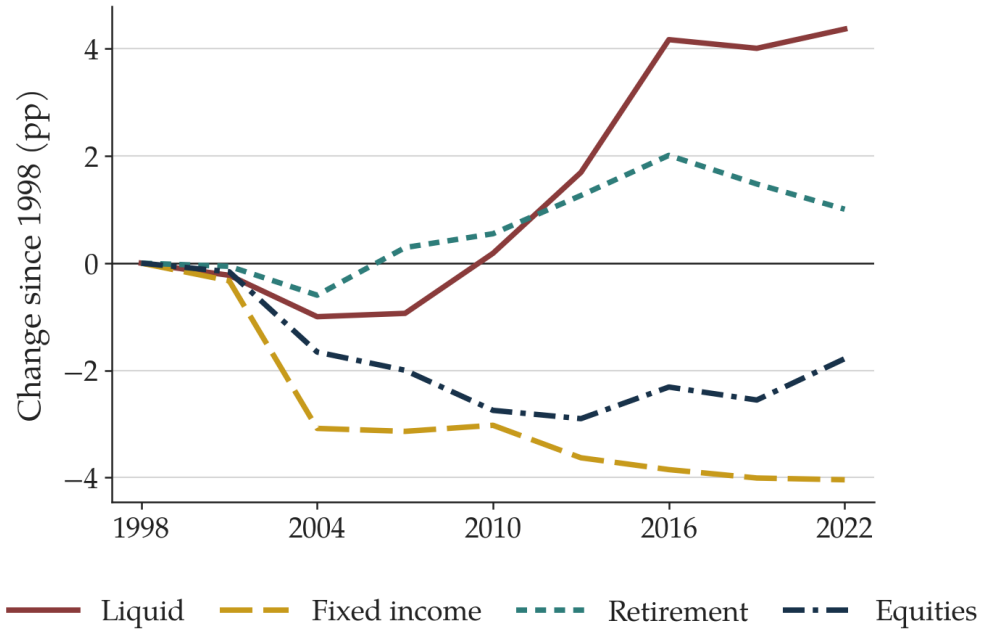


Figure C.11: Change in the Composition of Financial Assets Relative to Total Assets

Notes: The figure reports the change in the shares of liquid, fixed-income, retirement, and equity assets out of total assets in the SCF, relative to 1998. Shares are computed for each household and then averaged across households. The base year, 1998, is the SCF wave closest to the beginning of the PSID sample.

test the estimator of the MPC in the model, compare the MPC with the true elasticity of consumption to transitory shocks and with direct measures of the response to a transfer, and show how households use the housing margin. Finally, I report the MPC under other definitions of the wealth groups, the stylized facts in the model, the effect of each price change separately, and the robustness of the main experiment. Appendix E reports additional results on fiscal policy.

D.1 Computation

This section describes how I solve the model, compute its stationary distribution, simulate the panel used to estimate the MPC, and calibrate the parameters. All the code is written in Python and compiled with Numba. Solving one economy, computing its stationary distribution and simulating its panel takes about seven minutes on a computer with two cores.

States and grids. The state of a household of age t is (z, b, x, s) (Section 5.2). I discretize the persistent component z with the method of Rouwenhorst (1995) into three states. This method is accurate for very persistent processes (Kopecky and Suen, 2010). Entrants draw z from a mix of the stationary distribution and the middle state, so that income is less dispersed at 25 than later in life. The transitory shock ε is uniform, so I do not need a

quadrature rule for it. The value function is linear in cash on hand between grid points, and next year's cash on hand is linear in ε, so I integrate over ε exactly. The housing state s is one of the three rental options or an owned house of vintage $n = 0, \ldots, T$, where the vintage determines the outstanding mortgage. Cash on hand has 150 points for renters and 120 for owners. The grid goes from the lowest possible income to the largest cash on hand that a household can reach, and the points are closer to each other at the bottom, where the consumption function has the most curvature. The illiquid balance has 16 points, also closer to each other near zero. End-of-year liquid saving a' has 93 points, including the region of unsecured debt down to $-\underline{a}$.

Backward induction. I solve the problem backward, from age 85 to age 25. At 85, the household dies at the end of the year for sure, so the continuation value is the bequest alone. At each age t, given the value function V_{t+1} and its derivatives, I do four steps.

1. *End-of-year value.* For each housing choice s', persistent state z, illiquid balance b' and liquid saving a' on their grids, I compute

$$W_t(a', b'; z, s') = \pi_t \, \mathbb{E}\big[V_{t+1}(z', R_b b', R(a')\, a' + y_{t+1}, s') \,\big|\, z\big] + (1 - \pi_t)\, B(w'),$$

   and its derivatives with respect to a' and b'. The expectation sums over the three values of z' and integrates over ε'.

2. *Consumption with a fixed balance.* A household that keeps its illiquid balance at b' and has m to spend after its housing payment solves $\max_c u(c) + \beta W_t(m - c, b'; z, s')$. I solve it with the endogenous grid method (Carroll, 2006). For each a' on the grid, the Euler equation $u'(c) = \beta \, \partial W_t / \partial a'$ gives consumption, and $m = c + a'$ gives the cash that makes this choice optimal. The function W_t is not concave, because the housing choices of later years and the withdrawal cost create kinks in it. The Euler equation can then have several solutions for the same m. The upper envelope step of Iskhakov *et al.* (2017) keeps, for each m, the solution with the highest value. Below the first endogenous point, the borrowing limit binds and the household consumes $m + \underline{a}$. I store the consumption function on a common grid of 900 values of m.

3. *Consumption with a deposit.* A household that deposits chooses its total saving $S = a' + b'$ and how to split it. For each S on the grid, the best split solves $\widehat{W}_t(S) = \max_{0 \le b' \le S} W_t(S - b', b'; z, s')$. I find it with a grid search over b', followed by a bisection on the first-order condition $\partial W_t / \partial a' = \partial W_t / \partial b'$ between the neighbors of the best grid point. I then apply step 2 to $\widehat{W}_t$, with the cash $m + b$ that includes the current balance. This option is valid only if the new balance is at least b; otherwise, the household would be withdrawing without paying Ψ.

4. *Housing and withdrawals.* For each state (z, b, x, s) on the grid and each housing choice s', I compute the cash left after the housing payment, $m = x - \mathcal{C}(s, s')$. The household compares three options: keeping its balance (step 2), depositing (step 3), and withdrawing to a lower balance $b' < b$ on the grid, which leaves $m + (1 - k_1)(b - b') - k_0$ to spend (step 2). A renter who buys a house uses its illiquid balance toward the down payment and the fee without paying the withdrawal cost, and starts the year as an owner with a zero balance. The best option gives $\mathcal{V}_t(z, b, x, s, s')$. I then combine the best renting and the best owning option with equation (11), which gives $V_t(z, b, x, s)$. Its derivatives with respect to x and b are the averages of the derivatives of each option, weighted by the logit probabilities.

Stationary distribution. I compute the stationary distribution with a forward sweep over ages, using a non-stochastic simulation in the spirit of Young (2010). At each age, the distribution is a set of masses over the housing state, the persistent state, the illiquid balance on its grid, and a fine grid of cash on hand with steps of 0.002 times mean income (coarser above six times mean income). At 25, every household is a renter in the standard unit with no wealth, and its cash on hand is its first income. For each mass, I compute the probability of each tenure choice with equation (11) and the policies of each branch, which give the end-of-year balances (a', b'). A balance b' between two grid points is split between them so that the average balance is preserved. I then move each mass to next year's persistent state with the transition matrix and spread it over the fine grid of cash on hand, integrating over the uniform transitory shock exactly. The cross-section weights each age by the probability of surviving to it. All the moments of the stationary distribution are computed from these masses.

Simulated panel. The MPC of the model is estimated on a simulated panel (Section 5.3). I simulate the panel in lockstep with the forward sweep: at each age, the 40,000 households use the same policy functions as the distribution, interpolated in cash on hand, so that I do not need to solve the model a second time. Each household is followed from 25 to 55. The tenure choice is drawn from the logit probabilities, and income from the income process. I use the same random draws in every economy, so that the difference between two economies does not come from simulation noise.

Calibration. The eight calibrated parameters solve $m(\Theta) = \hat{m}$ (Section 5.4). I solve this system with a damped Newton method. The Jacobian is computed by finite differences at the starting point and updated with Broyden's method after each step. Each moment is weighted by its tolerance, and the step is shortened when it would move a parameter too far from the current point. The algorithm stops when every moment is within its tolerance of its target.

D.2 Targeted and Untargeted Moments

Table D.1 compares the moments of the baseline economy with the data. The model matches the eight targeted moments of Section 5.4. It also reproduces several moments that I do not target. Homeownership rises with wealth as in the data, from about a fifth of low-wealth households to more than four fifths of high-wealth households. Mid-wealth households are as likely to be HtM in the model as in the data (33 and 36 percent). Liquid wealth and fixed income are small for low- and mid-wealth households, in the model and in the data. The model has fewer HtM households at the bottom than the data (35 against 53 percent), and high-wealth households hold more liquid wealth and fixed income than in the SCF.

Table D.1: Targeted and Untargeted Moments, Heads Aged 25–55, 1999–2007

	Data	Model	Source
Targeted			
Homeownership (%)	56.5	56.5	PSID
Liquid assets and fixed income	0.44	0.44	SCF
Neither own nor pay rent (%)	6.0	6.0	PSID
Housing cost / income, low-wealth renters	0.27	0.26	PSID
Housing cost / income, mid-wealth renters	0.21	0.21	PSID
Owners who sell each year (%)	3.5	3.5	PSID
Net worth, 75+ relative to 65–74	0.77	0.77	PSID
Fixed income, top third of net worth	0.59	0.59	SCF
Untargeted			
By wealth group (low / mid / high)			
Homeownership (%)	21 / 62 / 85	17 / 65 / 87	PSID
HtM (%)	53 / 36 / 20	35 / 33 / 15	PSID
Liquid wealth	−0.01 / 0.02 / 0.33	−0.01 / 0.07 / 0.53	SCF
Fixed income	0.00 / 0.03 / 0.40	0.03 / 0.07 / 0.57	SCF

Notes: Model moments from the baseline economy, heads aged 25–55 unless stated. Housing cost relative to income: total housing costs over total income of renters who pay rent. Liquid wealth: liquid assets net of credit card debt. Liquid assets, liquid wealth and fixed income in years of mean labor income. Wealth groups: in the PSID and in the model, terciles of household-average net worth; in the SCF, the same dollar thresholds as in the PSID ($14,984 and $103,062 of 2024), as in Table A.1. Top third of net worth: terciles of net worth within each SCF wave and in the model. SCF: 1998–2007 waves.

D.3 Life-Cycle Profiles

The model also reproduces how ownership and wealth evolve with age. Figure D.1 compares homeownership, liquid wealth and net worth by age in the model and in the PSID over 1999–2007. Ownership rises with age at about the rate observed in the data, from about a quarter of households in their late twenties to more than two thirds after 40. Median liquid wealth is small at every age, in the model and in the data, and median net worth rises steadily over working life.

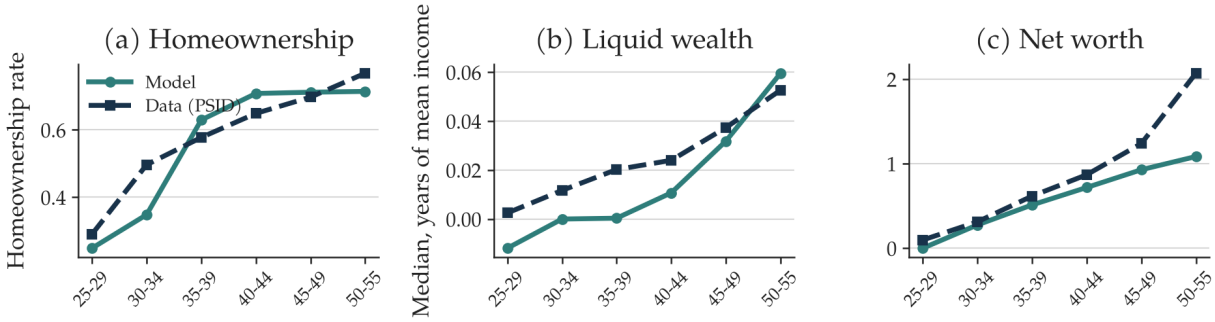


Figure D.1: Life-Cycle Profiles, Model and Data

Notes: Homeownership, median liquid wealth and median net worth by age group, heads aged 25–55. Data: PSID, 1999–2007. Liquid wealth and net worth in years of mean income.

D.4 The Transitory Income Shock

The model has no unemployment and no unemployment insurance (UI). I interpret income in the model as labor earnings after UI, and I measure the transitory risk of this measure of income in the PSID. This section describes how.

I use the PSID family files for 1999–2019. They report the UI received by the head and by the spouse in the previous year. From 2005 on, the PSID reports the annual amount. For 1999–2003, I build it from the amount, the period to which it refers, and the months in which it was received.[27] I add UI to the labor earnings of the head and the spouse. I also build two broader measures: earnings, UI and other public transfers (workers' compensation, TANF and other welfare, and SNAP), and total family income. All of them are in real terms and divided by the square root of family size. I keep heads aged 25–55 of the SRC sample who are in the labor force, and I residualize log income by wave on age, education, race, sex, family size, number of children, region, employment status and marital status, close to the controls of Section 3.1. For each measure, I estimate the variance of the transitory component as Blundell *et al.* (2008) do: minus the covariance of consecutive two-year changes in residual log income. I then subtract a measurement error variance of 0.02, the one I add to the simulated panel (footnote 17).

Table D.2 reports the standard deviation of the transitory component. A few very large income changes have a large effect on this covariance, so I report the estimates under four treatments of extreme changes. With the rule of Blundell *et al.* (2008), which drops households with an income change above 500 percent or below −80 percent, the standard deviation of earnings and UI is 0.16. Dropping the top and bottom percentile of income

[27] This rule reproduces the annual amounts that the PSID reports from 2005 on in 99.9 percent of the cases in which both are available.

changes in each wave gives 0.18 to 0.19. The model uses 0.17. Without any trimming, the standard deviation is 0.30, because a small number of very large changes dominate the covariance. Appendix D.13 shows the model with a transitory shock of 0.32, the estimate for labor earnings without trimming.

UI lowers the standard deviation by about 0.01 under every treatment, and the broader measures give about the same numbers. UI is received in 6.5 percent of household-years. In the years in which earnings fall below half of their usual level, only one household in seven receives UI, and it replaces about a fifth of the lost earnings of those who do.

As a check, I apply the same code to the sample of Blundell *et al.* (2008): married couples in the annual waves of 1978–1991, with their measure of family income after taxes and transfers. It gives a transitory standard deviation of 0.17, including measurement error, against about 0.175 implied by their estimates.

Table D.2: Transitory Income Risk in the PSID, 1999–2019

Extreme income changes removed	Labor earnings	Earnings and UI	Earnings, UI and transfers	Total family income
Rule of Blundell *et al.* (2008)	0.16	0.16	0.17	0.18
Top and bottom 1% of log changes	0.19	0.18	0.18	0.18
Top and bottom 1% of percent changes	0.20	0.19	0.19	0.19
No trimming	0.32	0.30	0.30	0.30

Notes: Standard deviation of the transitory component of log income, net of measurement error. PSID family files, biennial waves 1999–2019, SRC sample, heads aged 25–55 in the labor force. Income of the head and the spouse, in real terms and divided by the square root of family size. UI: unemployment insurance. Transfers: workers' compensation, TANF and other welfare, and SNAP. Log income is residualized by wave on the controls of Section 3.1. The variance of the transitory component is minus the covariance of consecutive two-year changes, as in Blundell *et al.* (2008), minus a measurement error variance of 0.02. Rows: treatment of extreme income changes, applied to each measure. Rule of Blundell *et al.* (2008): drop households with income below \$100 or a change above 500% or below −80% in any wave. Top and bottom 1% of percent changes: the rule of the MPC sample (Section 2.1).

D.5 The MPC Estimator in the Model

The model allows me to test the method of Blundell *et al.* (2008) in a setting where the truth is known. Figure D.2 applies the tests of Appendix B.2 to the simulated panel of the baseline economy, with the same residualization as the estimator. Panel (a) shows that, as in the PSID, income growth is strongly negatively correlated with income growth one wave earlier (-0.30 in the model and -0.36 in the data), and essentially uncorrelated at longer lags. The transitory component has no memory beyond one wave. Panel (b) shows that consumption growth is correlated with income growth one wave ahead, which is the source of identification of the estimator, but not with income growth further ahead.

Households do not respond to income changes before they happen. Panel (c) shows that the variance of income changes rises almost linearly with the horizon, as a random walk implies. The persistent component of the model is not a random walk, but it is persistent enough that the estimator treats it as one.

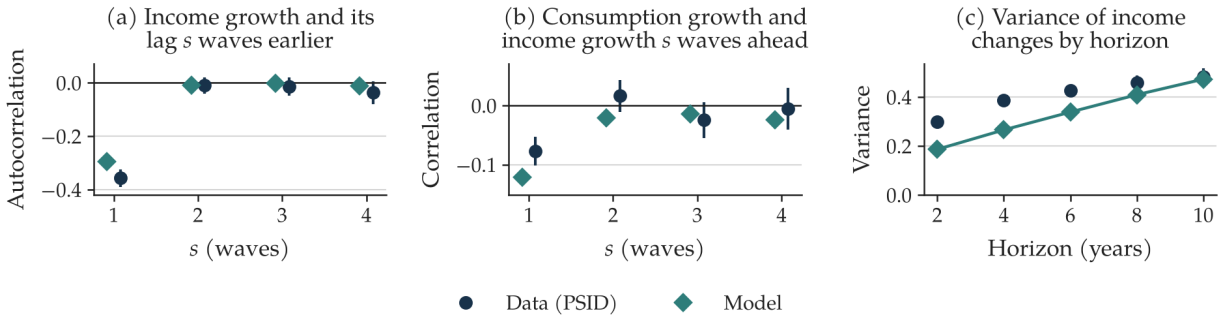


Figure D.2: The Restrictions of the Estimator in the Model

Notes: Simulated panel of the baseline economy, biennial waves, log income and log non-housing consumption residualized on age. Data: PSID, as in Appendix B.2, with 95% confidence intervals. (a) Correlation of two-year income growth with income growth s waves earlier. (b) Correlation of consumption growth with income growth s waves ahead. (c) Variance of income changes by horizon.

The model also observes the object that the estimator targets. For each household, I know the transitory shock that hits its income. I can therefore compute the true elasticity of consumption to transitory shocks: the coefficient of a regression of the change in log non-housing consumption on the transitory component of log income, on the same households and waves as the estimate. Figure D.3 compares the two. The estimate recovers both the ordering of the groups and its change. In the baseline economy, the true elasticity is highest in the middle (0.43, against 0.26 at the bottom and 0.30 at the top), as is the estimate. Between the two economies, the true elasticity rises by 0.18 at the bottom and falls by 0.18 in the middle, against $+0.14$ and -0.17 for the estimate. The estimate is below the true elasticity in every group (footnote 17), but it measures the same divergence. The divergence in the MPC is a divergence in how much households respond to transitory income shocks.

Table D.3 reports the estimate with annual changes instead of two-year changes, which the model allows but the PSID does not. The estimates are very close to the baseline ones, in levels and in changes.

D.6 Measures of the MPC

The model also allows me to compare the MPC with direct measures of the response to a transfer, computed on the same households. The transfer MPC holds housing fixed:

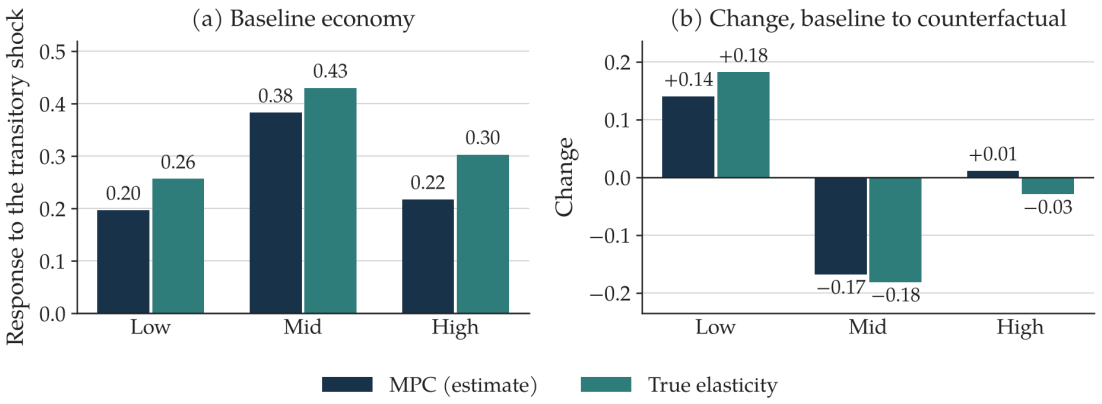


Figure D.3: The MPC and the True Elasticity of Consumption to Transitory Shocks

Notes: MPC: estimate of equation (12) on the simulated panel, as in Section 5.3. True elasticity: coefficient of a regression of the change in log non-housing consumption on the transitory component of log income, on the same households and waves. (a) Baseline economy. (b) Change between the baseline and the counterfactual economy. Wealth groups: terciles of household-average net worth.

Table D.3: The MPC of the Model with Other Choices of the Estimator

	Baseline economy			Change		
	Low	Mid	High	Low	Mid	High
Baseline	0.20	0.38	0.22	+0.14	−0.17	+0.01
Annual changes	0.21	0.35	0.20	+0.12	−0.14	−0.01
True elasticity	0.26	0.43	0.30	+0.18	−0.18	−0.03

Notes: Estimates of equation (12) on the simulated panels of the baseline and the counterfactual economy. Baseline: four biennial waves and two-year changes, as in Section 5.3. Annual changes: seven annual waves over the same six years. True elasticity: as in Figure D.3. Wealth groups: terciles of household-average net worth.

the household receives a transfer of 0.5% of mean income after it has chosen its housing for the year. Figure D.4 compares it with the MPC, in the baseline economy and their change between the two economies. The levels differ, but the changes are very similar. Between the two economies, the transfer MPC rises at the bottom of the distribution by 0.11, against 0.14 for the MPC, and falls in the middle by 0.13, against 0.17. The method of Blundell *et al.* (2008) recovers the change in the response of households to transfers, even if the level of the MPC differs from the level of the transfer MPC. The divergence in the data is therefore informative about the effectiveness of fiscal transfers.

Table D.4 explains why the levels differ. It goes from the transfer MPC with housing fixed to the MPC, one step at a time, for the baseline economy. Letting households re-optimize their housing lowers the response to a small transfer at the bottom, because some renters use it to stay in a better unit instead of consuming it. The response to large transfers de-

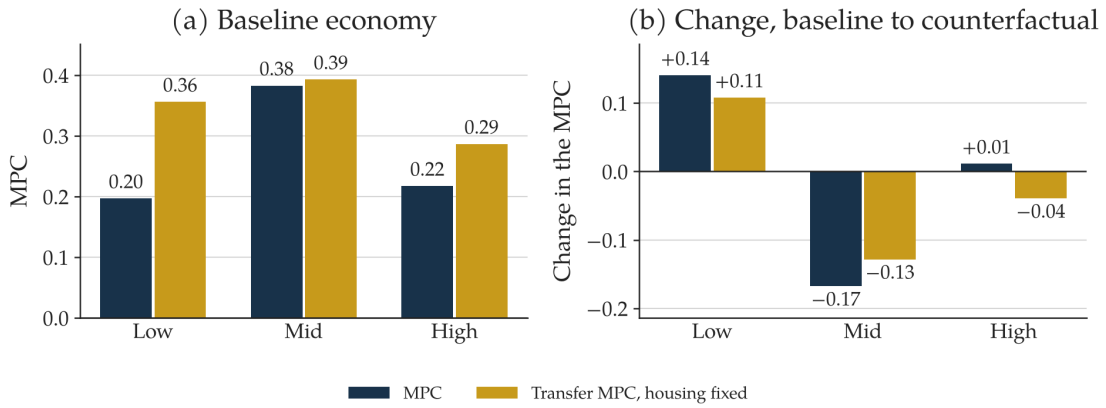


Figure D.4: The MPC and the Transfer MPC in the Model

Notes: Panel (a) shows the MPC and the transfer MPC with housing held fixed by wealth group in the baseline economy, for a transfer of 0.5% of mean income. Panel (b) shows the change in each measure between the baseline and the counterfactual economy. Wealth groups: terciles of household-average net worth.

pends on their sign: households spend a smaller share of a large positive transfer and cut consumption by a larger share of a large negative one, because consumption is concave in cash on hand. The MPC is identified from transitory shocks of this size, both positive and negative. Low-wealth households respond to large losses partly by moving down the rental ladder, which lowers their consumption response. The gap between the transfer MPC with housing fixed and the MPC is therefore largest at the bottom, where the housing margin is used, and small in the middle, where it is not.

Table D.4: From the Transfer MPC to the MPC, Baseline Economy

	Low	Mid	High
Transfer MPC, housing fixed, transfer of 0.5%	0.36	0.39	0.29
Transfer MPC, housing adjusts, transfer of 0.5%	0.25	0.31	0.26
Transfer MPC, housing adjusts, transfer of −0.5%	0.28	0.33	0.25
Transfer MPC, housing adjusts, transfer of 5%	0.19	0.32	0.26
Transfer MPC, housing adjusts, transfer of −5%	0.32	0.34	0.25
True elasticity to transitory shocks	0.26	0.43	0.30
MPC	0.20	0.38	0.22

Notes: Baseline economy. Transfer MPCs: change in non-housing consumption within the year divided by the transfer, in percent of mean income, with housing held fixed or re-optimized. True elasticity: as in Figure D.3. MPC: estimate of equation (12). Wealth groups: terciles of household-average net worth.

Figure D.5 shows how households use the housing margin. Panel (a) repeats the event study of Section 3.2 on the simulated panel. After a large transitory loss, the share of low-wealth households who adjust their housing rises by 10 percentage points relative

to households with stable income, against 19 in the PSID. As in the data, the response is much smaller in the middle and at the top of the distribution. Panels (b) and (c) show the share of households who choose cheaper housing after a one-off loss of cash, by the size of the loss. Small losses do not trigger a move. Above 5 percent of income, low-wealth households start to move down the rental ladder or in with others, and after a loss of 30 percent of income, 27 percent of them do. Mid- and high-wealth households, most of whom are owners, adjust much less, because selling a house is costly. After a loss of 30 percent of income, 22 percent of renters move to cheaper housing, against 9 percent of owners. The housing margin therefore mutes the response of low-wealth households to the large shocks that identify the MPC, but not their response to a small transfer.

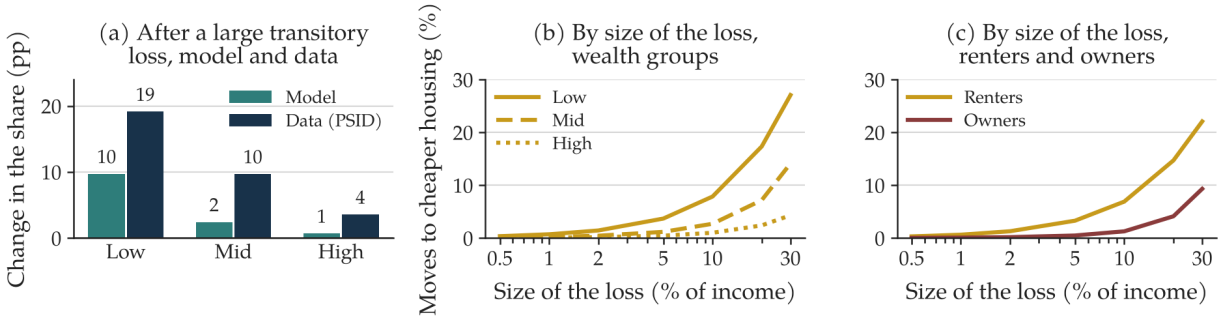


Figure D.5: Housing Adjustment After Income Losses in the Model

Notes: Baseline economy. (a) Change in the share of households that adjust their housing when income falls by more than 20 percent and recovers by more than 15 percent at the next wave, relative to households with stable income, in the model and in the PSID (Section 3.2). (b) and (c) Share of households, aged 25–55, that choose cheaper housing after a one-off loss of cash on hand, in percent of their own income, than they would have chosen without it: a cheaper rental unit, living with others, selling the house, or not buying it. Tenure at the beginning of the year. Wealth groups: terciles of household-average net worth.

D.7 MPC by Wealth Group in the Two Economies

Table D.5 reports the MPC of each wealth group in the baseline and in the counterfactual economy, and in the initial and final periods of the PSID (Section 6.1). In the model, the MPC of low-wealth households rises from 0.20 to 0.34, and that of mid-wealth households falls from 0.38 to 0.22. The ordering of the two groups reverses, as in the data, where the MPC of low-wealth households rises from 0.13 to 0.26 and that of mid-wealth households falls from 0.38 to 0.16.

D.8 HtM Households

Table D.1 shows the share of HtM households by wealth group in the baseline economy and in the PSID at the beginning of the sample. Mid-wealth households are about as

Table D.5: MPC by Wealth Group, Model and Data

	Low-wealth	Mid-wealth	High-wealth
Model			
Baseline economy	0.20	0.38	0.22
Counterfactual economy	0.34	0.22	0.23
Change	+0.14	−0.17	+0.01
Data			
Initial period	0.13	0.38	0.21
Final period	0.26	0.16	0.09
Change	+0.14	−0.22	−0.11

Notes: Model: baseline economy and counterfactual economy with the house prices, rents and interest rates of the 2010s. Data: PSID, initial period (window centered on 2004, 1999–2007) and final period (window centered on 2016, 2011–2019). Wealth groups: terciles of household-average net worth. The high-wealth group is the least informative: its MPC is the least precisely estimated in the data, and its change is not statistically different from zero (Table 3.1).

likely to be HtM in the model as in the data (33 and 36 percent). Figure D.6 shows how the share of HtM households changed between the two economies. It falls the most in the middle of the distribution, where fewer households hold most of their wealth in a house or in fixed income.

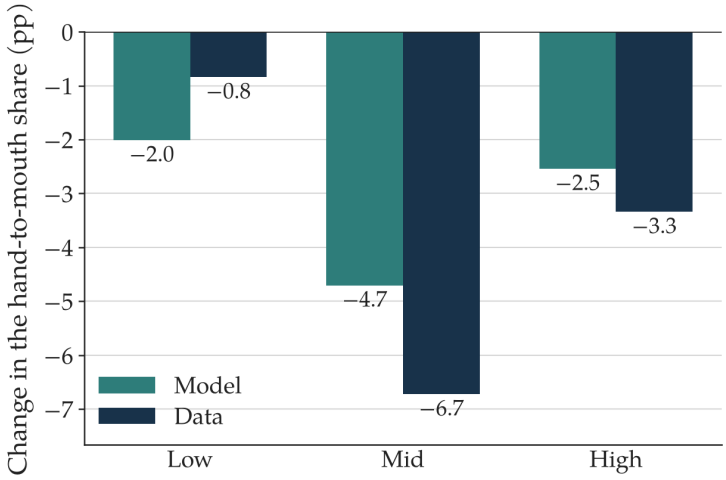


Figure D.6: Change in the Share of HtM Households, Model and Data

Notes: Change in the share of HtM households, defined as in Kaplan *et al.* (2014), between the baseline and the counterfactual economy (model) and between the initial and the final period in the PSID, windows centered on 2004 and 2016 (data), in percentage points. Wealth groups: terciles of household-average net worth.

D.9 Results by Wealth Decile

Figure D.7 shows the MPC and the transfer MPCs by decile of household-average net worth. In the baseline economy, the MPC is highest between the fifth and the seventh deciles, where recent buyers and households saving for a house are concentrated. These are the wealthy HtM households of Section 5.5. The transfer MPC with housing fixed is highest in the first decile, and letting housing adjust lowers it the most in the bottom half of the distribution, where most households rent. Panel (b) shows the change between the two economies. The MPC rises in the bottom four deciles and falls sharply in the fifth and sixth deciles, where recent buyers were concentrated in the baseline economy.

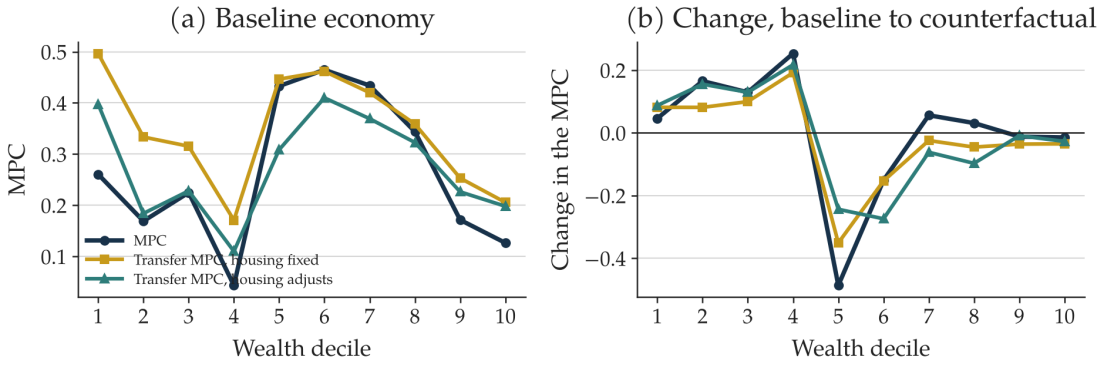


Figure D.7: MPC by Wealth Decile in the Model

Notes: Panel (a) shows the MPC and the transfer MPCs with housing held fixed and with housing re-optimized, by decile of household-average net worth, in the baseline economy. Panel (b) shows the change between the baseline and the counterfactual economy.

D.10 Alternative Groupings

In the data and in the model, I group households by their net worth averaged over several waves, but the averages cover different periods. In the data, the main grouping averages net worth over all the waves in which the household is observed in 1999–2019 (Panel A of Table B.2). In the model, it averages net worth over the four waves used to estimate the MPC (Panel B). Table D.6 compares the data and the model under the three groupings of Table B.2. To build Panel A in the model, I follow each household for up to eleven biennial waves, the length of the PSID sample, and estimate the MPC on four-wave windows.

The MPC rises at the bottom and falls in the middle under every grouping, in the data and in the model. With the same grouping in both, the model accounts for about two thirds of the change in the data in Panel A ($+0.10$ and -0.12, against $+0.14$ and -0.22), and more than all of it in Panel B ($+0.14$ and -0.17, against $+0.06$ and -0.15). Within age groups (Panel C), the change in the model is much smaller ($+0.02$ and -0.04), while

in the data it is about the same as in Panel B. In the model, higher house prices mostly delay ownership: the average age at the first purchase rises from 33 to 37. Recent buyers, who have the highest MPC, are older in the counterfactual economy, but within their age group they are still in the middle of the distribution. In the data, the fall in the middle survives within age groups.

Table D.6: MPC by Wealth Group Under Alternative Definitions of the Groups, Data and Model

	Initial period / baseline economy			Change		
	Low	Mid	High	Low	Mid	High
Panel A. Terciles over the whole sample						
Data	0.13	0.38	0.21	+0.14	−0.22	−0.11
				(0.08)	(0.10)	(0.10)
Model	0.20	0.39	0.22	+0.10	−0.12	−0.01
Panel B. Terciles within each window (main text of the model)						
Data	0.19	0.32	0.16	+0.06	−0.15	−0.03
				(0.08)	(0.09)	(0.10)
Model	0.20	0.38	0.22	+0.14	−0.17	+0.01
Panel C. Terciles within each window and age group						
Data	0.21	0.31	0.14	+0.08	−0.15	−0.04
				(0.09)	(0.09)	(0.10)
Model	0.22	0.38	0.20	+0.02	−0.04	−0.01

Notes: MPC estimated as in Section 3.1. Data: PSID, initial period 1999–2007 and final period 2011–2019, as in Table B.2; standard errors of the change in parentheses. Model: baseline and counterfactual economy. Panel A: terciles of net worth averaged over all the waves in which the household is observed (in the model, up to eleven biennial waves, the length of the PSID sample). Panel B: terciles of net worth averaged over the four waves of each window. Panel C: as in Panel B, within five-year age groups of the head.

Grouping by current net worth gives the same picture. With current net worth, the MPC of the model is 0.20, 0.30 and 0.26 for low-, mid- and high-wealth households in the baseline economy, and 0.23, 0.20 and 0.25 in the counterfactual economy. The MPC rises at the bottom and falls in the middle, and by the end of the period low-wealth households have a higher MPC than mid-wealth households. The changes are smaller than with household-average net worth, because current net worth moves with the income shocks that identify the MPC: a household hit by a large negative shock draws down its wealth and moves to a lower group. The same happens in the data, where the estimates with current wealth groups are noisier but point in the same direction.

D.11 Stylized Facts in the Model

Figure D.8 compares the changes in homeownership, housing costs, portfolios and liquid wealth between the two economies with the changes in the data (Section 6.1). The model reproduces the fall in homeownership and where it is concentrated, and portfolios become more liquid in every group, as in the data. It understates the rise in the housing costs of low- and mid-wealth renters, because renters in the model respond to higher rents by moving to cheaper units or in with others. Liquid wealth rises for mid- and high-wealth households, as in the data, but much more at the top (by 0.20 of income, against 0.03). At the bottom, it barely changes in the model and falls in the data. Appendix D.12 explains why.

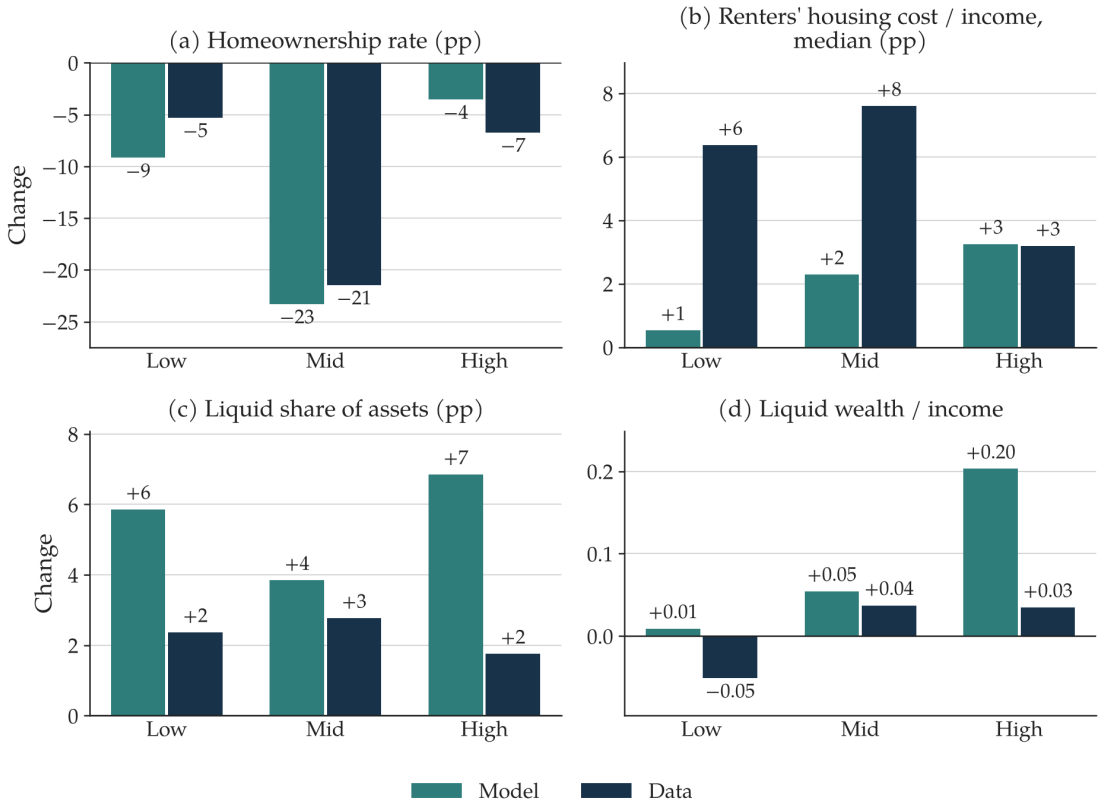


Figure D.8: Homeownership, Housing Costs, Portfolios and Liquid Wealth in the Model and in the Data

Notes: Changes between the two economies (model) and between the initial and the final period (PSID). (a) Homeownership rate. (b) Median housing costs relative to income of renters who pay rent. (c) Liquid assets over total assets. (a)–(c): terciles of net worth within each economy or period. (d) Average liquid wealth relative to income, terciles of household-average net worth. Data: initial (1999–2007) and final (2011–2019) period.

D.12 Each Price Change Separately

Table 6.1 in Section 6.2 reports the change in the MPC when each price changes alone. Figure D.9 shows the stylized facts and liquid wealth under each change. Higher rents raise the housing costs of renters. Higher house prices reduce access to ownership, most of all in the middle of the distribution. Lower interest rates raise liquid wealth and make portfolios more liquid. No single change reproduces all the facts, but the three together do.

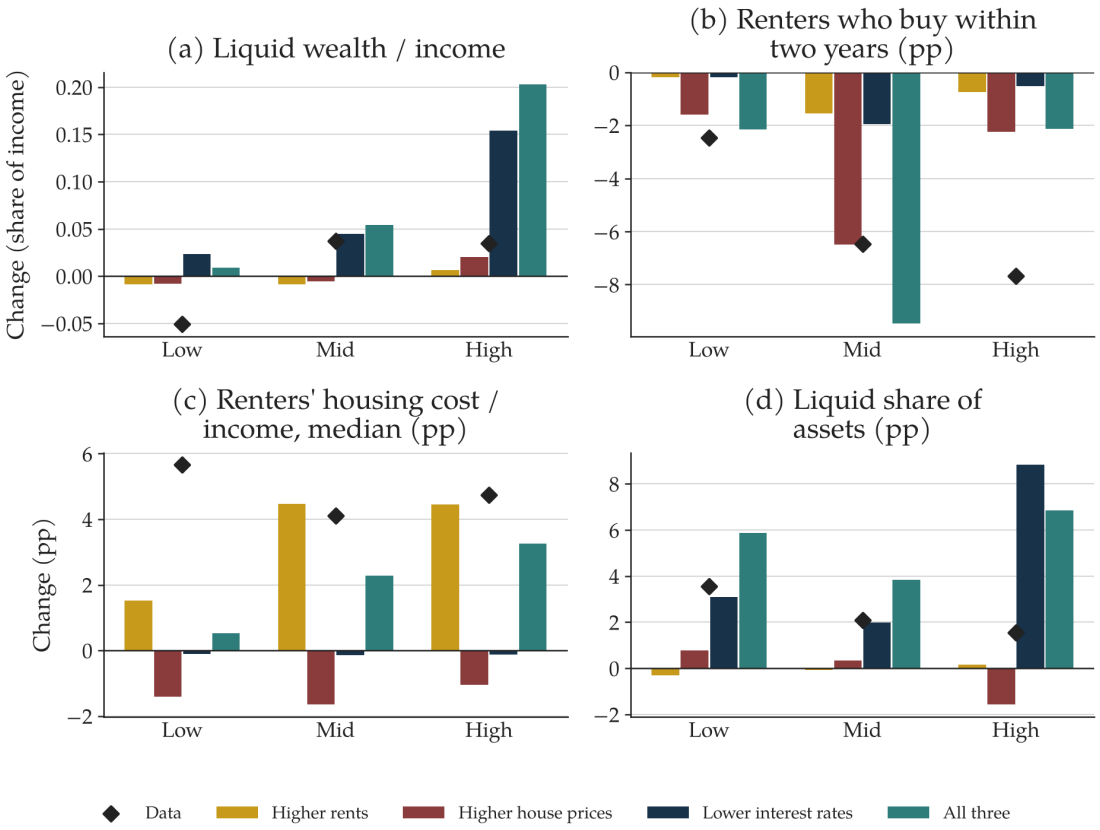


Figure D.9: The Stylized Facts under Each Price Change

Notes: Changes relative to the baseline economy when each price changes alone and when the three change together (bars), and changes in the PSID between the initial period, 1999–2007, and the final period, 2011–2019 (diamonds). (a) Average liquid wealth relative to income. (b) Share of renters who own a house two years later, in percentage points. (c) Median housing costs relative to income of renters who pay rent, in percentage points. (d) Liquid assets over total assets, in percentage points. Wealth groups: terciles of household-average net worth.

Liquid wealth at the bottom. Panel (a) of Figure D.9 shows that two opposing forces act on the liquid wealth of low-wealth households. Higher rents and higher house prices lower it: each reduces the average by about 0.01 of income, as renters pay more for housing and fewer of them accumulate savings toward a down payment. Lower interest rates

push in the opposite direction, by $+0.02$. Some low-wealth renters save for a house partly in fixed income, and when its return falls they hold those savings in liquid form instead. This raises liquid wealth at the top of the low-wealth group, but it does not affect the typical low-wealth household: its median liquid wealth barely changes, in the model and in the data. In the data, the fall in the average comes from liquid debt, which the model does not capture.

D.13 Robustness

This section studies how the main results change when I relax some of the assumptions of the model. In each case I change one feature of the baseline economy and repeat the main experiment of Section 6.1. Except in the last three cases, the other parameters stay at their baseline values. Table D.7 reports the MPC by wealth group in the baseline economy, its change between the baseline and the counterfactual economy, and the change in homeownership.

Table D.7: Robustness of the Main Experiment

	MPC, baseline			Change in the MPC			Ownership
	Low	Mid	High	Low	Mid	High	(change, pp)
Baseline	0.20	0.38	0.22	+0.14	−0.17	+0.01	−12.0
Alternative price changes	0.20	0.38	0.22	+0.14	−0.21	0.00	−17.0
No down-payment exemption	0.22	0.36	0.22	+0.12	−0.16	+0.01	−11.2
Refinancing allowed	0.17	0.41	0.18	+0.16	−0.12	+0.03	−10.2
Lower selling cost	0.20	0.38	0.21	+0.14	−0.18	+0.02	−11.8
Transitory s.d. 0.32	0.32	0.34	0.30	+0.07	−0.19	+0.01	−11.9
Unsecured credit at 20%	0.18	0.38	0.22	+0.12	−0.16	+0.01	−12.0
Taste shock on each rental unit	0.19	0.38	0.22	+0.12	−0.17	+0.01	−12.1
Data	0.13	0.38	0.21	+0.14	−0.22	−0.11	−11.1

Notes: MPC estimated on simulated panels with the method of Section 3.1. Wealth groups: terciles of household-average net worth. Ownership: change in the homeownership rate of heads aged 25–55 between the baseline and the counterfactual economy. All economies are simulated with the same random draws as in the main text, so the baseline row is the same as in Table D.5. Data: PSID, windows centered on 2004 and 2016, and 1999–2007 to 2011–2019 for ownership.

Alternative price changes. In the main experiment, I measure the rise in house prices and rents relative to the mean income of the PSID sample. Measuring it relative to median income gives larger changes: house prices rise by 38 percent and rents by 31 percent. With these changes, the MPC rises by 0.14 at the bottom and falls by 0.21 in the middle, closer to the data. Homeownership falls by 17 percentage points, more than in the data.

No down-payment exemption. In the baseline, a buyer can use its illiquid savings for the down payment without paying the withdrawal cost. When buyers pay the withdrawal cost like any other household, the illiquid asset becomes a worse vehicle to save for a house, and fewer households in the middle hold it. The MPC still rises at the bottom (by 0.12) and falls in the middle (by 0.16).

Refinancing and equity extraction. In the baseline, owners cannot refinance or extract equity without selling the house. In this version, an owner can reset its mortgage to the original loan-to-value ratio at any time and receive the difference between the new loan and its outstanding balance in cash, paying a fee of 2 percent of the house value. This works like a cash-out refinancing or a home equity line. Owners can now use their home equity as a buffer, but only at a cost and only for amounts large enough to justify the fee. The MPC of mid-wealth households in the baseline economy is even higher than in the baseline (0.41, against 0.38), and the divergence remains: the MPC rises by 0.16 at the bottom and falls by 0.12 in the middle. Access to home equity does not undo the mechanism, because the households who drive it are recent buyers, who have little equity to extract, and renters saving for a house.

Lower selling cost. The housing margin is harder to use for owners than for renters because selling a house is costly. Lowering the selling cost from 8 to 5 percent of the house value makes it easier for owners to adjust their housing. The results barely change: the MPC rises by 0.14 at the bottom and falls by 0.18 in the middle.

Unsecured credit at 20 percent. The baseline rate on unsecured debt is 12 percent in real terms, based on the interest rate on credit card accounts that were assessed interest. Including fees, the cost of credit card borrowing is higher (Agarwal *et al.*, 2015). With a real rate of 20 percent and the model recalibrated, the MPC rises by 0.12 at the bottom and falls by 0.16 in the middle.

A larger transitory shock. In the baseline, the transitory shock has a standard deviation of 0.17, the transitory risk of earnings after unemployment insurance (Section 5.4). Here I set it to 0.32, the transitory standard deviation of labor earnings over 1999–2019 when the largest income changes are not trimmed (Appendix D.4). I recalibrate the model to the same moments. The model needs a lower discount factor (0.814 instead of 0.852) and no bequest motive, and it misses the rate at which owners sell (4.4 percent a year against 3.5 percent): with shocks this large, some owners have to sell after a bad year. The MPC still falls in the middle (by 0.19) and rises at the bottom (by 0.07). However, the MPC in the baseline economy is almost the same in the three groups, and the mid-wealth households no longer have the highest MPC. A shock of this size gives every household the risk of the households most exposed to unemployment. With larger shocks, mid-

wealth households hold more liquid wealth, and fewer of them are wealthy HtM (19 percent instead of 27 percent). At the bottom, the savings from moving to a cheaper unit cover a smaller part of a bad year, so the housing margin insures low-wealth households less. In this version, higher rents alone barely change the MPC at the bottom (by 0.01, against 0.06 in the baseline). In the PSID, transitory risk is not the same for everyone: years of unemployment are three times as frequent for households with low earnings as for households with high earnings.

Taste shocks on each rental unit. In the baseline, the taste shocks apply only to the choice between renting and owning, and a renter chooses its unit without them. Here every option, each rental unit and the house, has its own taste shock, with the same scale σ. These shocks capture reasons to change units that the model does not describe, such as a change in family size or a new job in another part of the city. I recalibrate the model to the same moments, and the parameters barely change. Renters now change units more often (18 percent of them each year, against 11 percent in the baseline), and low-wealth households are more likely to adjust their housing after a large transitory loss (by 14 percentage points, against 10 in the baseline and 19 in the PSID). The main results do not change: the MPC rises by 0.12 at the bottom and falls by 0.17 in the middle, and homeownership falls by 12 percentage points. As in the baseline, higher rents push renters down the rental ladder and the housing margin provides less insurance: after a large transitory loss, the share of low-wealth households that adjust their housing rises by 12 percentage points in the counterfactual economy, against 14 in the baseline economy.

E Fiscal Policy Appendix

This appendix complements Section 7. I first show which households drive the change in the transfer MPC of each wealth group. I then show the cost of matching a uniform transfer with a targeted transfer, and the gain from targeting other shares of the distribution and other variables.

Transfer MPC by type of household. Table E.1 reports the transfer MPC by size and type of household in the two economies (Section 7.1). It shows that the change in the transfer MPC of each group comes from a change in its composition. At the bottom, renters at or near the credit limit spend about half of a small transfer in both economies, and they become more common, from 36 to 44 percent of low-wealth households. Other renters also spend more of a small transfer (0.30 against 0.15). In the middle, owners with a mortgage spend about half of a small transfer in both economies, but they become much less common, from 72 to 38 percent of mid-wealth households. They are replaced by renters saving for a house, who spend little or nothing of a transfer.

Table E.1: Transfer MPC by Size and Type of Household

	Baseline economy				Counterfactual economy			
		Transfer (% of income)				Transfer (% of income)		
	Share	0.5	5	20	Share	0.5	5	20
Low-wealth households	1.00	0.28	0.22	0.13	1.00	0.43	0.38	0.20
Renters at or near the credit limit	0.36	0.51	0.39	0.20	0.44	0.59	0.50	0.28
Other renters	0.64	0.15	0.12	0.09	0.56	0.30	0.28	0.13
Mid-wealth households	1.00	0.32	0.31	0.29	1.00	0.17	0.16	0.15
Owners with a mortgage	0.72	0.51	0.48	0.41	0.38	0.51	0.49	0.42
Renters saving for a house	0.28	−0.18	−0.14	−0.05	0.62	−0.04	−0.05	−0.01
High-wealth households	1.00	0.25	0.25	0.24	1.00	0.22	0.22	0.21
Owners	0.97	0.27	0.27	0.25	0.95	0.25	0.24	0.22

Notes: Share of a one-off transfer spent on non-housing consumption within the year, with housing re-optimized. Wealth groups: terciles of current net worth, heads aged 25–55. Renters at or near the credit limit: liquid wealth within 2% of mean income of the credit limit. Share: share of the households of each wealth group in each type; the types shown cover at least 97% of each group. A transfer brings the purchase of a house forward for some renters saving for a house; the down payment then absorbs the transfer and part of their consumption, so the average transfer MPC of this group is negative.

The cost of matching a uniform transfer. Figure E.1 shows the budget of a transfer targeted to low-wealth households, and to low-income households, that generates the same increase in spending as a uniform transfer (Section 7.3). In the baseline economy, both targeted transfers cost more than the uniform transfer. In the counterfactual economy, a transfer to low-wealth households costs between 64 and 78 percent of the uniform transfer.

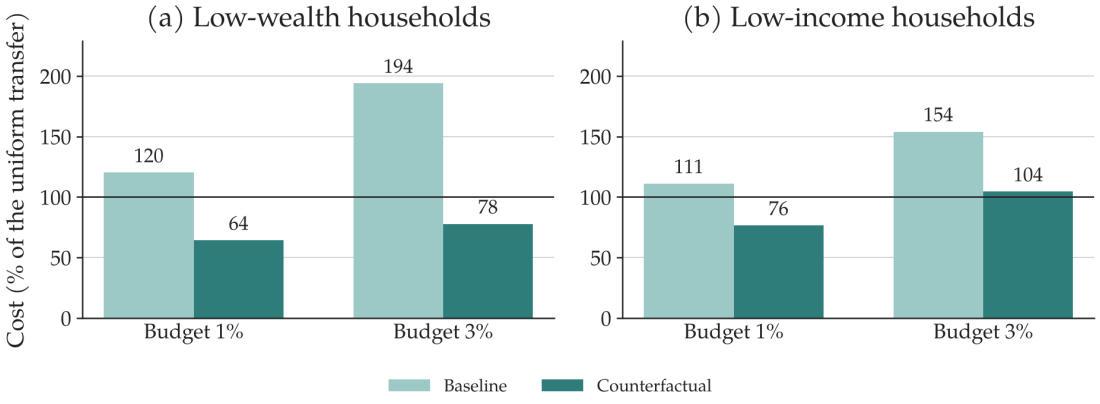


Figure E.1: Cost of Matching a Uniform Transfer

Notes: Budget of the targeted transfer that matches the spending of a uniform transfer of Z, in percent of Z (equation (15)). Targeting by net worth (a) and by income (b).

Other targets. Section 7.2 focuses on transfers targeted to low-wealth households, the bottom third of the distribution of net worth. Figure E.2 shows the gain from targeting any share f of the distribution of net worth (panel a) and of income (panel b), and Figure E.3 shows the cost of a transfer targeted to the bottom f of the distribution of net worth that matches a uniform transfer. With a budget of 1%, targeting the bottom 10% of the distribution of net worth raises spending by 12% relative to a uniform transfer in the baseline economy, and by 76% in the counterfactual economy, and it costs 80% and 44% of the uniform transfer, respectively. Targeting too narrowly does worse than targeting a little more broadly in both economies, because the transfer each household receives becomes too large. Targeting by income does worse than targeting by net worth for every share of households targeted.

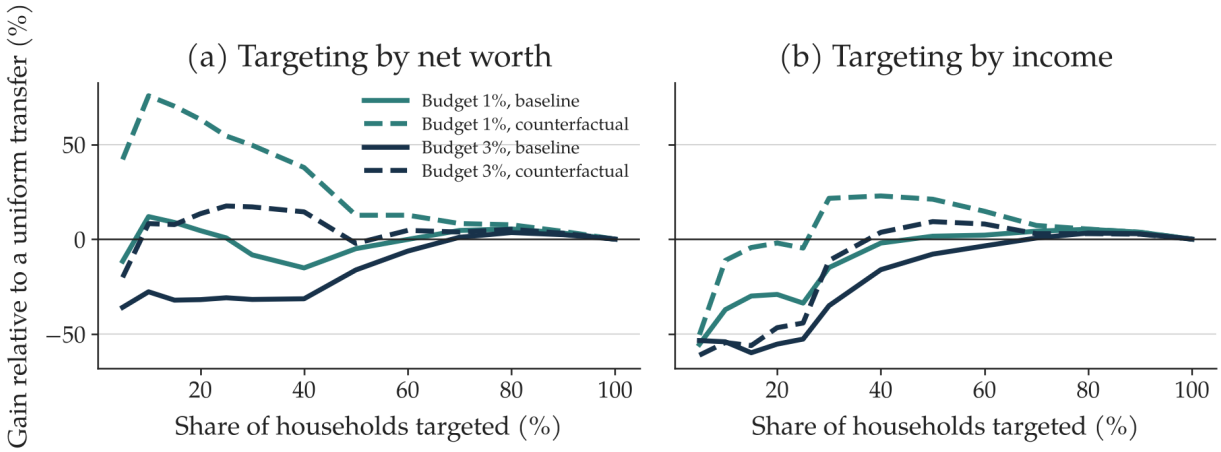


Figure E.2: Gain from Targeting Relative to a Uniform Transfer

Notes: Gain $G(f; Z)$ from equation (14) of targeting the bottom f of the distribution of net worth (a) and of income (b), relative to a uniform transfer with the same budget Z per household, in the baseline (solid) and the counterfactual (dashed) economy. Budgets as a percentage of mean annual income. Heads aged 25–55.

Figure E.4 shows the gain from targeting households with the lowest liquid wealth, relative to a uniform transfer with the same budget. Liquid wealth is the best target in both economies. It identifies HtM households directly, including recent buyers in the middle of the net worth distribution. With a budget of 1%, targeting the bottom 15% of the distribution of liquid wealth more than doubles spending relative to a uniform transfer, in the baseline and in the counterfactual economy. Liquid wealth is harder to observe than income, however, and the gain from targeting by liquid wealth changed less between the two economies than the gain from targeting by net worth.

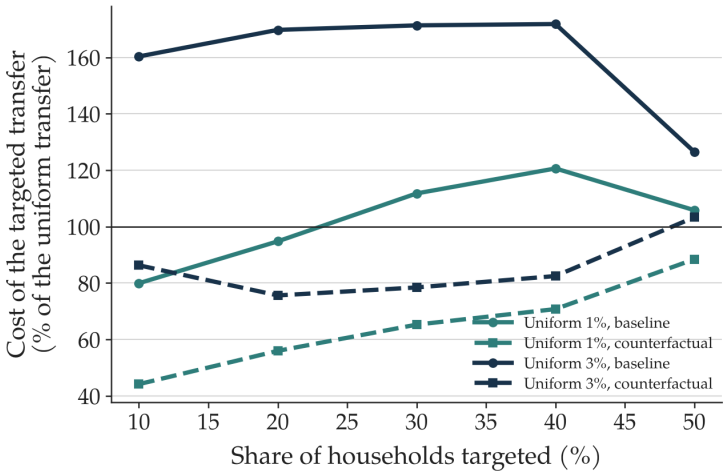


Figure E.3: Cost of a Targeted Transfer that Matches a Uniform Transfer

Notes: Budget Z_f of a transfer targeted to the bottom f of the distribution of net worth that generates the same increase in spending as a uniform transfer of Z_u, as a percentage of Z_u (equation (15)), in the baseline (solid) and the counterfactual (dashed) economy. Heads aged 25–55.

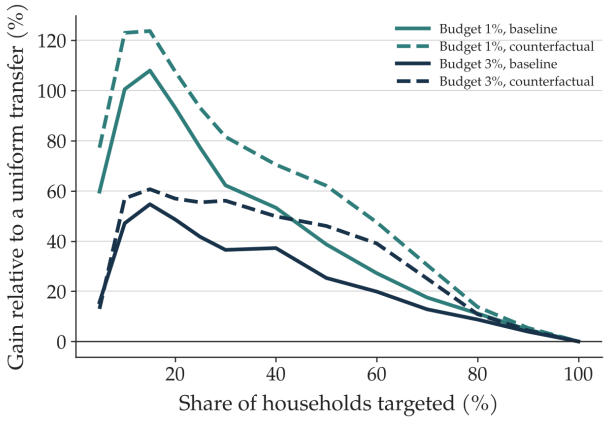


Figure E.4: Gain from Targeting by Liquid Wealth

Notes: Gain $G(f;Z)$ from equation (14) of targeting the bottom f of the distribution of liquid wealth, relative to a uniform transfer with the same budget Z per household, in the baseline (solid) and the counterfactual (dashed) economy. Heads aged 25–55.