To Have or Not to Have: Understanding Wealth Inequality*

Pavel Brendler†
University of Bonn

Moritz Kuhn
University of Mannheim,
CEPR, and IZA

Ulrike Steins
Evonauts GmbH

August 18, 2024

Abstract

Differences in household saving rates are a key driver of wealth inequality. But what determines these differences in saving rates and wealth accumulation? We provide a new answer to this long-standing question based on new empirical evidence and a new modeling framework. In the data, we decompose U.S. household wealth into its main portfolio components to document two new empirical facts. First, the variation in wealth by income is mainly driven by differences in participation in asset markets rather than by the amounts invested. Wealth differences are a matter of *to have or not to have*. Second, the large heterogeneity in asset market participation closely follows observed differences in access to asset markets. Combining these two facts, we develop a new model of life-cycle wealth accumulation in which income-dependent market access is the key driver of differences in asset market participation and saving rates by income. The calibrated model accurately captures the joint distribution of income and wealth. Eliminating heterogeneity in access to asset markets increases wealth accumulation in the bottom half of the income distribution by 32%. Facilitating access to employer-sponsored retirement accounts improves broad-based wealth accumulation in the U.S. economy. Historical data support the model's prediction.

Keywords: Wealth Inequality; Labor Market Heterogeneity; Household Portfolios.
JEL codes: D31; E21; H31.

*We are thankful for very useful comments to Árpád Ábrahám, Corina Boar, Jake Bradley, Chris Busch (discussant), Jan Eeckhout, Giulio Fella, Joachim Hubmer, Winfried Koeniger, Dirk Krueger, Felix Kubler, Burhan Kuruscu, Benjamin Moll, Simon Mongey, Ed Nosal, Krzysztof Pytka, José-Víctor Ríos-Rull, Florian Scheuer, Minchul Yum, and the seminar participants at the 2024 CEPR European Summer Symposium in International Macroeconomics in Paris, 2023 Macro Workshop in Lisbon, 2023 Workshop on "Macroeconomic implications of housing, household finances, and wealth dynamics" in St. Gallen, 2023 SED Meeting in Cartagena, 2022 Econometric Society Meeting in Berlin, Bank of Portugal, Danmarks Nationalbank, Federal Reserve Banks of Atlanta, St. Louis, and Cleveland, Toulouse School of Economics, and the Universities of Birmingham, Bologna, Bonn, Dortmund, Frankfurt, Lugano, and Nottingham. We thank Moritz Scheidenberger for excellent research assistance. Steins gratefully acknowledges financial support from the Science Foundation of Sparkassen-Finanzgruppe. Kuhn gratefully acknowledges financial support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/1–390838866.

†Corresponding author. Address: University of Bonn, Institute for macroeconomics, Adenaueralle 24-42, 53113 Bonn, Germany. Email: pavel.brendler at uni-bonn.de.

1 Introduction

Wealth inequality has always been one of the largest economic disparities (Saez and Zucman, 2014; Kuhn et al., 2020). The determinants of wealth accumulation have therefore been a long-standing topic of economic research. In theory, differences in saving rates by income are one of the potential determinants of differences in wealth accumulation. In the data, such differences find strong support (Dynan et al., 2004).[1] While the consequences of differences in saving rates for wealth inequality are understood, their causes remain a question that still requires answers.[2] The goal of this paper is to provide a new answer to this long-standing question.

We start from the observation that actual wealth accumulation occurs across multiple assets that make up the household portfolio, so that differences in saving rates and wealth accumulation may arise from differences in asset market participation (the extensive margin) or from differences in saving conditional on participation (the intensive margin). Focusing on the main components of the U.S. household portfolio, we find that heterogeneity in asset market participation is the dominant driver of wealth disparities and provide evidence that this heterogeneity closely relates to variation in reported asset market *access*. We thus confirm a conjecture in Dynan et al. (2004, p. 436) that differential access to asset markets is the cause of differences in saving rates.

Guided by this new empirical evidence, we develop a theory of life-cycle wealth accumulation with portfolio choice, asset market participation, and heterogeneous asset market access. We use the model to qualitatively and quantitatively assess the role of the empirically documented heterogeneity in asset market access for wealth accumulation across the income distribution. In the model, heterogeneous asset market access and the resulting differences in asset market participation are the key drivers of low saving rates among lower-income households. Eliminating heterogeneity in asset market access increases wealth accumulation in the bottom half of the income distribution by 32%, slightly reduces wealth accumulation in the top half of the income distribution (3%), and thus reduces overall wealth inequality.

Consistent with our results, we find that unequal access to asset markets is also an important issue in the current policy debate in the United States. We use the developed model to evaluate the recently proposed bipartisan *Retirement Saving for Americans Act*, which aims to improve access to retirement savings for low-income workers. Implementing such a policy in our model, we find that it successfully broadens wealth accumulation and helps low-income households to accumulate life-cycle savings. Thus, our empirical and theoretical results suggest that asset markets do not provide a level playing field with worse opportunities for low-income households to build wealth.

The main data source for our empirical analysis is the *Survey of Consumer Finances* (SCF). We

[1] Over the past decade, there has been renewed interest in the topic, with a rapidly growing literature studying differences in saving rates and their implications for the distribution of wealth, see for example Saez and Zucman (2014); Alan et al. (2015); Krueger et al. (2016); Straub (2019); Fagereng et al. (2021).

[2] There is a very large literature on wealth inequality, but work that explicitly examines differences in saving rates by income is scarce. One of the papers that examines the reasons for differences in saving rates by income is Hubbard et al. (1995). They point to means-tested social security to explain the low saving rates of income-poor households. De Nardi (2004) and Straub (2019) point to non-homothetic preferences as a driver of high saving rates at the top of the distribution. See the literature review for more details.

document that wealth-to-income ratios rise sharply with household income and that this is true for each of the three major portfolio components: home equity, retirement accounts, and business equity. Together, these three components account for about 70% of the (marketable) wealth of U.S. households. We refer to the sum of these components as *core wealth* and document two facts about its joint distribution with income.[3] First, we show that the overwhelming share of the variation in wealth holdings by income for each portfolio component is driven by extensive margin differences. Strikingly, we find that the intensive margin is constant along the income distribution or, in the case of home equity, actually declines with income. In essence, these results show that understanding wealth differences is a matter of *to have or not to have*.

Second, we examine the sources of heterogeneity in asset market participation and document large differences in asset market access by income. Approval of a mortgage application by a bank is the key condition for access to the housing market, and we document substantial heterogeneity in approval rates along the income distribution, increasing by a factor of five between the bottom and top income deciles. Quantitatively, this pattern is consistent with the documented rising trend in the extensive margin for home equity, suggesting that approval rates, and hence market access, are an important driver of participation rates. In contrast, the share of mortgage applications remains virtually constant along the income distribution, suggesting that differences in demand are unlikely to account for rising wealth-to-income ratios.

Similarly, we document large heterogeneity in access to the market for employer-sponsored retirement accounts, with offer rates increasing by a factor of four along the income distribution. This increase in access to the pension market is twice as large as the variation in the pick-up rates of these plans conditional on receiving an offer. Thus, differences in pick-up rates contribute in part to the documented variation in participation rates, but as in the case of housing, the main source of variation comes from offer rates and, hence, asset market access. We provide extensive sensitivity analysis for these facts. The relationship between mortgage approval rates, offer rates, and income remains statistically and economically significant even after controlling for a wide range of household characteristics. We corroborate our findings using *Home Mortgage Disclosure Act* (HMDA) data on mortgage approval rates, where we also control for debt-to-income ratios and down payments, and the *National Compensation Survey* for pension plans.

Guided by this new empirical evidence, we develop a life-cycle model of wealth accumulation with portfolio choice and heterogeneity in access to asset markets. Our framework differs from existing models of saving behavior and life-cycle wealth accumulation in two important ways. First, our model features a discrete choice of whether to invest in the three asset classes that comprise core wealth (the decision to participate in asset markets). Conditional on the decision to save, agents in the housing and pension markets invest in and commit to the fulfillment of predetermined financial contracts (fixed mortgage contracts and pension plans). Consistent with models of illiquid wealth (Kaplan and Violante, 2014; Bayer et al., 2019), households in our model are typically wealthy hand-to-mouth agents. We abstract from saving in liquid assets and short-run consumption smoothing since we show in the data that non-core wealth contributes insignificantly to life-cycle wealth accumulation. To account for the important role of business

[3] Core wealth overlaps to a large extent with the definition of illiquid assets in Kaplan and Violante (2014).

equity, we introduce a discrete occupational choice for entrepreneurial activity. As a second departure from the existing literature, we allow for income-dependent differences in access to mortgage and pension markets, in line with our empirical findings. In the tradition of the standard incomplete markets model (Bewley, 1986; Huggett, 1993; Aiyagari, 1994), our analysis abstracts from microfounding the reasons for the lack of access to asset markets by income and focuses instead on its consequences.[4] As a result, saving decisions are discrete decisions given existing opportunities over financial contracts, rather than driven by a marginal saving decision governed by an Euler equation.

We calibrate the model to the U.S. economy and show that it closely matches the joint distribution of income and wealth. While the calibration does not target heterogeneity along the extensive or intensive margin of any portfolio component, we still find that the model achieves an excellent fit of both margins along the income distribution. Thus, it also provides a very good fit to the empirically documented rising wealth-to-income ratios with income. In addition, the model provides a close fit to the untargeted increasing saving rates by income. This close fit offers a new interpretation of these differences as arising from differences in market access by income rather than differences in the demand for wealth accumulation. The idea that differences in asset market access are an important driver of the empirical differences in saving rates is already a conjecture in the discussion in Dynan et al. (2004, p. 436). Our results provide evidence for the validity of this conjecture.

To quantify the importance of income-dependent access to asset markets for wealth accumulation and inequality, we conduct a counterfactual experiment in which we remove income dependence of access to asset markets in the mortgage and pension markets by setting access rates to the economy-wide average. The extensive margins in the housing and pension markets increase substantially in the lower half of the income distribution, although participation in the pension market does not increase at the very bottom of the income distribution due to the availability of the social security system, consistent with Hubbard et al. (1995). Worse access to mortgages in the top half of the income distribution reduces the intensive margin of home equity accumulation by limiting the trading opportunities for upgrading to larger houses for high-income households. Thus, a model with homogeneous access to asset markets overstates the wealth-building opportunities of low-income households while it understates them for high-income households. As a consequence, wealth-to-income ratios increase by 32% in the bottom half of the income distribution, while they decrease by 3% in the top half.

We apply our framework to study how limited access to the private pension market affects wealth accumulation for low-income households. Unequal access to retirement savings is at the heart of the recently proposed bipartisan *Retirement Savings for Americans Act*. According to the Act, workers without access to employer-sponsored retirement plans should become eligible for government retirement accounts, with provisions to subsidize pensions for low-income workers. In the model, we introduce the program as an alternative freely accessible retirement saving plan with government contributions that phase out as income increases.

[4]For entrepreneurial activity, we focus on the friction of finding a valuable business idea at the occupational decision stage rather than on the financial constraint of scaling up a valuable idea. In the spirit of the housing and retirement markets, market entry (receiving an idea) is the key friction for entrepreneurs (Bloom et al., 2020).

The policy experiment produces three notable results. First, the increase in the share of workers enrolled in the government plan is large in the bottom half of the distribution and is largest, at 33%, in the third decile of the income distribution, leading to a significant increase in wealth accumulation among lower-income households. Averaged over all households in the bottom half of the income distribution, pension wealth increases by 69%, while the economy-wide average increase is much smaller (14%). This demonstrates the substantial impact of policies that increase the ability of lower-income households to access the pension market and build retirement savings. Second, the vast majority of workers enrolled in the newly available plans were excluded from employer-sponsored plans in the baseline model. While we find evidence of crowding out from employer plans to government plans, this effect is small (15% of total enrollment in government plans). Thus, our policy experiment highlights the important role of supply-side barriers in limiting wealth accumulation in the private pension market. Third, the cost of the reform is low, at less than 1% of total labor income. Enrollment in the newly available plans peaks around the third income decile, where government subsidies are limited due to their gradual phasing out. While enrollment is significant among higher income households, they do not qualify for subsidies. The main impact of the policy comes from improving access rather than providing subsidies.

To validate our findings, we conduct an experiment assuming an economy without pension plans and then introducing employer-sponsored retirement plans with observed access rates. We compare the model's prediction with the evolution of private retirement wealth during the 1980s when the supply of employer-sponsored retirement accounts grew rapidly following the passage of the *Economic Recovery Tax Act* of 1981. Quantitatively, the model explains almost two-thirds of the total growth in the extensive margin for pension wealth between 1983 and 1998 and replicates remarkably closely the observed trajectory of wealth-to-income ratios for pension wealth and the slowdown in home equity accumulation in the U.S. during this period.

1.1 Literature review

Our work lies at the intersection of three strands of the macroeconomic literature.

First, we contribute to the large literature studying the sources of wealth inequality. Existing research has extended the standard incomplete markets model à la Aiyagari (1994) through various channels to match the high degree of wealth concentration in the data. These channels include preference heterogeneity (Krusell and Smith, 1998), rich income dynamics (Castañeda et al., 2003), entrepreneurial activity (Quadrini, 2000; Cagetti and De Nardi, 2006), bequests (De Nardi, 2004), medical expense risk (De Nardi et al., 2010), capital return heterogeneity (Benhabib et al., 2019; Hubmer et al., 2021), or non-homothetic preferences (De Nardi, 2004; Straub, 2019).[5] The distinguishing feature of these studies is their focus on the very top of the wealth distribution. Our paper complements this work by accounting for wealth accumulation throughout the income distribution, and we, therefore, do not aim at accounting for wealth concentration at the very top of the distribution. A notable exception to the existing literature

[5] A comprehensive review of this literature is beyond the scope of this paper. We refer the reader to Cagetti and De Nardi (2008) and De Nardi and Fella (2017) for excellent reviews.

is the study by Hubbard et al. (1995), which seeks to explain the drivers of strongly increasing saving rates throughout the income distribution. They highlight the important role of the social security system for the low saving rates of households at the bottom of the income distribution. While the authors assume equal saving opportunities and market access for all households, we provide evidence of heterogeneous asset market access and show that this has implications for wealth accumulation even for households in the upper part of the income distribution.

Second, we contribute to the growing macroeconomic literature on saving behavior and portfolio choice that highlights the role of financial frictions in asset allocation (Kaplan and Violante, 2014; Kaplan et al., 2018; Luetticke, 2021; Bayer et al., 2024; Bilbiie, 2024). Our approach differs from this literature in two important respects. First, these studies focus on the illiquid asset as a single asset category, while we differentiate between its two major components – pension wealth and home equity. This detailed approach allows us to define the notion of asset market participation and study the role of frictional access to each market for wealth accumulation. Second, we introduce income dependence in access to asset markets, supported by our extensive empirical evidence. We demonstrate that a model with homogeneous access to asset markets substantially overstates the wealth accumulation of low-income households. In modeling the financial frictions, we follow the parts of this literature that assume stochastic access to asset markets but we allow for additional heterogeneity. We share with this strand of the literature that we take financial frictions as given, without microfounding them. Instead, we focus on studying their consequences for saving behavior and wealth accumulation.

Finally, we connect to the household finance literature that emphasizes heterogeneity in asset market participation as an important part of heterogeneity in investment behavior (see survey in Gomes et al., 2021). This literature mainly looks at the stock market and documents large variation in participation (Mankiw and Zeldes, 1991; Haliassos and Bertaut, 1995; Guiso et al., 2002). The proposed mechanisms have in common that they focus on the demand side, i.e. the investor's willingness to enter the market, assuming that all households have equal access to the market. Instead, we focus on missing asset market access from the supply side.

The rest of the paper is structured as follows. Section 2 conducts the empirical analysis. Section 3 introduces the model. The calibration strategy is discussed in Section 4. Our main findings are presented in Section 5. Section 6 conducts the policy experiment. Section 7 concludes.

2 Empirical analysis

2.1 The data

We rely on the Survey of Consumer Finances (SCF) data, which is the most widely used dataset to study the financial situation of U.S. households (see Bhutta et al., 2020; Kuhn and Rıos-Rull, 2016, and references therein). It provides detailed information on household income, wealth, and the household's portfolio composition.

Our empirical analysis focuses on two main variables: income and wealth. We follow Bhutta et al. (2020) for the definitions of these variables. We only deviate by excluding capital income (capital gains, dividends, and interest) from total income to avoid any mechanical relationship between income and wealth and to align the definitions in the data and the quantitative model.[6] Income is then the sum of pre-government wage and salary income, business income (income from a sole proprietorship or a farm), and transfer income. Wealth is the value of all financial and real assets net of all debts, i.e., marketable net worth.

We construct business wealth, home equity, and retirement accounts as key portfolio components on the household balance sheet. We refer to the sum of these three assets as *core wealth*. Business wealth is the sum of businesses in which the household has an active interest.[7] Home equity comprises both the primary residence and other residential real estate (e.g., vacation homes) minus housing debts, such as mortgages and home equity lines of credit. Retirement accounts include pension plans accumulated in individual retirement and thrift accounts, net of any loans taken against pensions. To cover the remainder of the household balance sheet, we construct the following three asset classes: (i) stock, bonds, and mutual funds, (ii) liquid assets, and (iii) other assets.[8]

We apply the following sample selection to align the data with our quantitative framework. We focus on household heads aged 25–65, where the head works in dependent employment or is self-employed. We drop households who report running a business but not actively managing it and set the household's business wealth or home equity to zero if the reported values are negative. To abstract from business-cycle variation, we pool the four survey waves from 2010, 2013, 2016, and 2019, covering data over the last decade, but the main empirical findings remain robust when pooling together all available data since 1989. Appendix A.1 describes the construction of the variables and the sample selection criteria in more detail. Appendix A.2 presents the main results when using all available data.

2.2 Total wealth and core wealth

Table 1 reports the shares of main asset components in total wealth held by three income groups (0–50%, 50–90%, and 90–100% of the income distribution). In each column, the portfolio shares sum up to 100% for a given income group.

The last column shows the economy-wide average shares of each asset class in total wealth. The portfolio components with the largest shares are business wealth (30.6%), home equity (21.4%), and retirement accounts (17.7%). Together, these three asset classes constitute core wealth and comprise 70% of household wealth. The remaining wealth components and their shares in total wealth are: Stocks, bonds, and mutual funds (15.6%), Liquid assets (6.1%), and Other assets (8.6%). Stocks, bonds, and mutual funds comprise the largest asset class, not

[6]Excluding capital income does not qualitatively change any of the empirical findings.

[7]Business wealth is the value of the business if sold net of all liabilities. Since our quantitative model will be populated with business owners, we restrict the focus to actively managed businesses to be consistent with the definition of entrepreneurs in the literature (Cagetti and De Nardi, 2006). However, this restriction has only a negligible impact on the empirical findings.

[8]Liquid assets comprise all types of transaction accounts (checking, savings, money market accounts, etc.) plus certificates of deposit net of credit lines, credit balances, etc. Other assets include remaining financial and real assets such as cash value of life insurance, vehicles, etc., net of any corresponding debts.

Table 1: Portfolio composition by income groups in the SCF, %.

	Wealth component	Share of asset held by income group:			
		$< 50\%$	$50-90\%$	$> 90\%$	All
	Business wealth	18.5	16.1	33.9	30.6
+	Home equity	41.2	31.5	18.5	21.4
+	Retirement accounts	17.2	28.2	15.9	17.7
=	**Core wealth**	**77.0**	**75.8**	**68.3**	**69.8**
+	Stocks, bonds, and mutual funds	8.8	8.6	17.3	15.6
+	Liquid assets	5.6	6.3	6.0	6.1
+	Other assets	8.7	9.3	8.4	8.6
=	**Total wealth**	**100.0**	**100.0**	**100.0**	**100.0**

included in core wealth, but these assets are heavily concentrated at the top of the distribution.[9]

Looking in detail at the portfolio shares along the income distribution displayed in the three middle columns of the table, we see that for each income group, the sum of business wealth, home equity, and pension wealth accounts for the largest share of their total wealth. For the bottom 90% of households, the share of these three portfolio components is even higher than the average and virtually the same for the first two income groups covering the bottom 90% (77.0% and 75.8%). For the top 10%, the share is only slightly lower than 70%. As one can see from the last rows of the table, the key difference between the top 10% and the bottom 90% is the higher share of stocks, bonds, and mutual funds, which is below 9% for the bottom 90% of the income distribution and almost twice as large (17.3%) for the top 10%. The largest share of *other assets* are cars. Together with liquid assets, which primarily consist of transaction accounts, they account for nearly two-thirds of the non-core wealth of the bottom 90%. These low/no-return asset classes are typically not used to accumulate life-cycle savings.

To characterize the joint distribution of income and wealth, we construct wealth-to-income ratios and study their variation across deciles of the income distribution. Figure 1 plots wealth-to-income ratios for total wealth (solid-circle line) and core wealth (solid-triangle lines) by income deciles. To construct the figure, we compute wealth-to-income ratios for each household and report the average ratio within each decile. To exclude outliers, we winsorize the top 1% of the wealth-to-income ratios for total wealth and each core wealth component and then construct wealth-to-income ratios for core wealth from the winsorized data.[10]

The figure shows that wealth-to-income ratios for *total wealth* rise with income. More specifically, the ratios rise linearly up to the 8^{th} income decile and exhibit a steeper increase towards the top of the income distribution. Looking at the level, we find that wealth corresponds to approximately annual income in the first income decile, whereas households in the top income decile hold wealth that is five times their annual income, so a top-income decile household accumulates four more dollars of wealth per dollar of income than a household in the bottom-

[9] Kuhn et al. (2020) document large and persistent differences in portfolio composition of U.S. households along the wealth distribution and report that the top 10% of households own more than 90% of all stocks.

[10] Figure 1 remains very similar if we plot the ratio of average wealth to average income by decile.

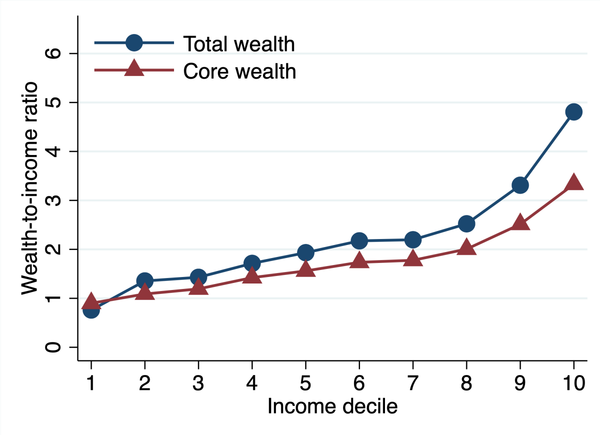

Figure 1: Wealth-to-income ratios for total wealth and core wealth in the SCF.

Notes: Core wealth is the sum of business wealth, home equity, and retirement accounts. *Total wealth* comprises core wealth plus stocks, bonds, mutual funds, liquid assets, and remaining assets. The value of each asset component is computed net of debts. See Appendix A.1 for the sample selection criteria.

income decile. A consequence of these differences is that household wealth is significantly more concentrated than income.[11]

Consider next the wealth-to-income ratios for *core wealth* by income deciles in Figure 1. First, we note that core wealth ratios typically lie below those for total wealth because core wealth is a subset of total wealth. Second, the two lines increase in parallel over most of the income distribution, indicating that both total wealth and core wealth follow the same pattern in their joint distribution with income. This proportionality between core wealth and total wealth suggests that core wealth captures most of the variation in wealth by income among U.S. households. Therefore, understanding the unequal distribution of core wealth components is crucial for identifying the drivers of total wealth accumulation and wealth inequality. In Appendix A.3, we show that non-core wealth is mainly concentrated at the top of the income distribution and households without core wealth have, on average, zero non-core wealth. Hence, there is no evidence of substitution from core to non-core wealth for households without core wealth.

2.3 Extensive and intensive margin effects

Figure 2 shows the wealth-to-income ratios for the three components of core wealth. We observe a consistent picture across all asset classes with the same qualitative pattern as for total wealth: An increase that is close to linear up to the 8^{th} decile and a steeper increase in the two top deciles. Hence, each asset class contributes to higher savings by income and a stronger concentration of wealth compared to income. The most notable difference is the very steep increase of business wealth in the top income decile, which makes business wealth a key driver of the wealth concentration at the top of the income distribution, as previously documented by Quadrini (2000) and Cagetti and De Nardi (2006).

Breaking total wealth into different components allows us to decompose the total variation

[11] As we confirm later, the relationship between income and wealth documented in this section remains economically and statistically significant even after controlling for a broad set of household demographics.

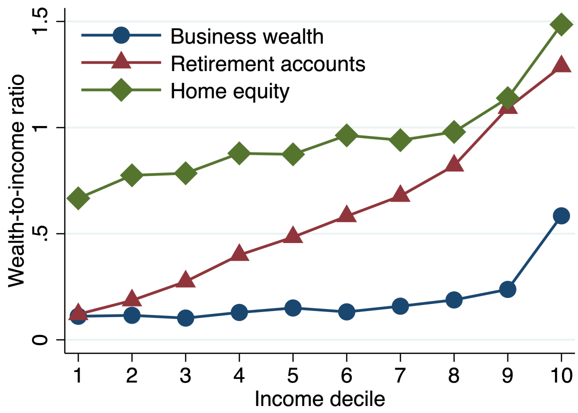

Figure 2: Wealth-to-income ratios for core wealth components in the SCF.

in wealth-to-income ratios into the share of households holding each asset (extensive margin) and the asset value held by those households who own positive amounts of the asset (intensive margin). Denoting by ω_a^d the wealth-to-income ratio in decile d for asset class a, we have:

$$\omega_a^d = \mu_d^a \times \tilde{\omega}_a^d,$$

where μ_d^a is the population share of households holding positive amounts of asset a in income decile d and $\tilde{\omega}_a^d$ is the wealth-to-income ratio, conditional on holding the asset. We refer to μ_d^a as the extensive margin for asset a in decile d, while $\tilde{\omega}_a^d$ denotes the intensive margin.

Figure 3 presents the decomposition results. Consider first the extensive margin variation μ_d^a in Panel (a). The figure shows that participation rates in each asset component rise strongly with income, meaning that the share of households who invest in home equity, retirement funds, or business wealth increases with income.[12] The figure reveals substantial variation in participation rates by income, with a 6-fold increase in participation for retirement accounts and business wealth and a 4-fold increase for home equity. Notably, these rising patterns at the extensive margin resemble the increasing pattern in the wealth-to-income ratio in Figure 2.

These results are in stark contrast to the patterns of the intensive margin variation $\tilde{\omega}_a^d$ in Figure 3(b). The intensive margin of business equity is approximately constant along the income distribution and it increases slightly for retirement accounts. Strikingly, the ratio of home equity to income even declines in income. These findings suggest that once a household has invested in these assets, asset accumulation relative to income varies little with income, indicating roughly constant saving rates among asset market participants. Regarding wealth inequality, the intensive margin variation suggests that wealth inequality would be similar to, or even less than, income inequality. Thus, the results underscore the crucial role of asset market participation in explaining differences in saving rates and wealth accumulation across the income distribution.

To quantify the contribution of the extensive and intensive margins for wealth accumulation, we construct wealth-to-income ratios for core wealth under two counterfactual scenarios. In the

[12] We assume that a household owns an asset if the value of the asset held is strictly positive, but our results are robust to using positive cutoff values close to zero.

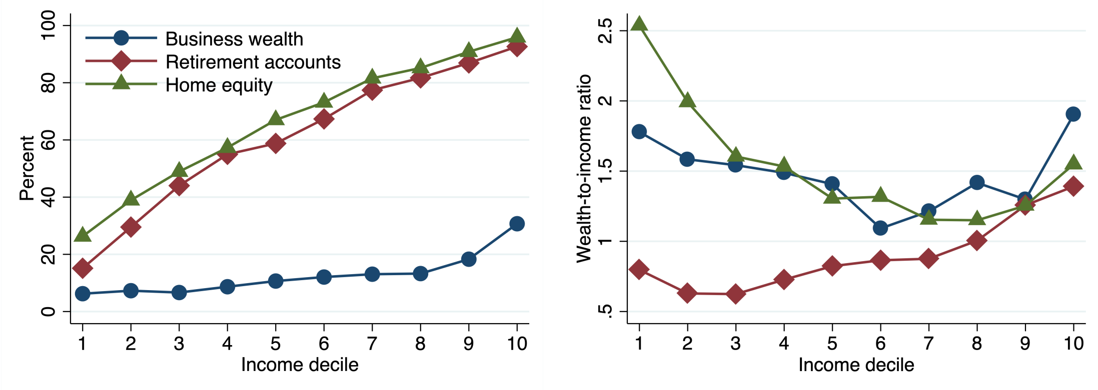

(a) Extensive margin

(b) Intensive margin

Figure 3: Extensive and intensive margin variation in core wealth components in the SCF.

Notes: The left panel shows the share of households (in percent) with a positive amount of a given asset by income deciles. The right panel shows wealth-to-income ratios for a given asset class by income deciles, conditional on households holding a positive amount of the asset.

fixed extensive margin scenario, we keep the extensive margins of the three core wealth components at their sample average and vary only the intensive margin from Figure 3(b). Reversely, the *fixed intensive margin* scenario sets the intensive margins constant at their sample averages and varies only the extensive margin from Figure 3(a).

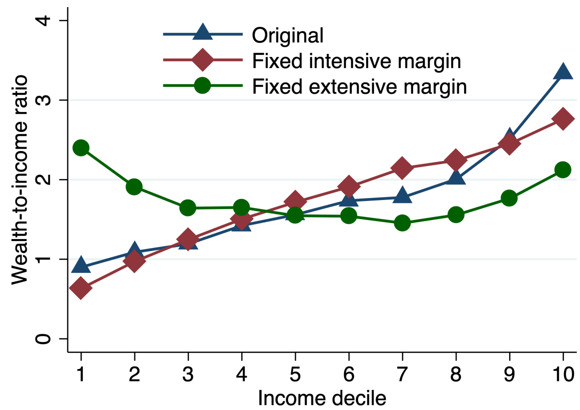

Figure 4: Extensive and intensive margin effects on wealth-to-income ratios.

Notes: *Fixed extensive margin* experiment assumes the extensive margins to be constant across the income distribution and equal to the average extensive margin for a given asset class, while the intensive margins correspond to the observed values in the data. The reverse applies to the *Fixed intensive margin* experiment.

Figure 4 presents the results of the two counterfactual scenarios together with the original wealth-to-income ratios for core wealth (solid-triangle line). Comparing the counterfactual cases to the observed data, we get striking results. Eliminating the variation in the extensive margin (*fixed extensive margin*), we obtain wealth-to-income ratios that decrease up to the 7^{th} income decile, in contrast to the increasing ratios in the data. What is more, wealth-to-income ratios in the top decile are below wealth-to-income ratios in the bottom decile. By contrast, the effect of removing the intensive margin variation (*fixed intensive margin*) on the variation in wealth by income is modest for most of the income distribution, and we only find notable variation in the top decile of the distribution. As in the original data, wealth-to-income ratios remain strongly increasing over the entire income distribution, suggesting that the variation in the extensive margin along the income distribution plays a pivotal role in explaining the variation in wealth by income.

In summary, this decomposition of U.S. wealth differences by income establishes that variation from the extensive margin accounts virtually for the entire differences in wealth accumulation along the income distribution. The intensive margin varies little and in a u-shaped pattern between the bottom and the top of the distribution. From a theory perspective, the importance of the extensive margin scrutinizes the focus of many macroeconomic models on marginal investment decisions, in which wealth accumulation is governed by the Euler equation.[13] While the marginal consideration captures well average wealth accumulation, it misses the substantial variation across the income distribution stemming from the asset market participation decision. Our empirical findings suggest that understanding why some households do not invest in certain asset classes is key to understanding the joint distribution of wealth and income.[14]

2.4 Income-dependent asset market access

The documented variation in the extensive margin can result from differences in the demand for and the supply of assets. When studying saving rate differences by income, Dynan et al. (2004) already conjecture that differences in market access (supply) could be a potential driver of differences in saving behavior. In this subsection, we provide empirical evidence for their conjecture supporting the idea that market access to the housing and retirement market differs across households with different incomes.

2.4.1 Housing market

Approval of a mortgage application is the key condition for accessing the housing market. To study access to the housing market, we rely again on SCF data that also provide information on mortgage approval rates by income. We define a household as an applicant if it has applied for a mortgage or home-based loan or requested to refinance an existing mortgage in the past 12 months. We assign households that did not apply because they anticipated being rejected as well as households who did not receive full funding to the group of households that applied and were rejected.[15]

In the SCF sample, 20% of households applied for a mortgage or refinancing during the past 12 months, and, on average, one in five households got rejected or received only partial funding.

[13]Such deviations from the Euler equation are in line with the class of HANK models (Kaplan and Violante, 2014; Bayer et al., 2019) where household consumption dynamics follow the consumption dynamics of hand-to-mouth consumers for extended periods.

[14]While highly instructive, this decomposition exercise is stylized. It assumes the extensive and intensive margins to be constant by income while both margins are equilibrium outcomes. In the quantitative part, we will examine the effects of a uniform asset market access, allowing extensive and intensive margins to respond endogenously. This approach will account for substitution across assets in response to changes in market access.

[15]In the SCF, 5% of households did not apply because they anticipated their application to be rejected.

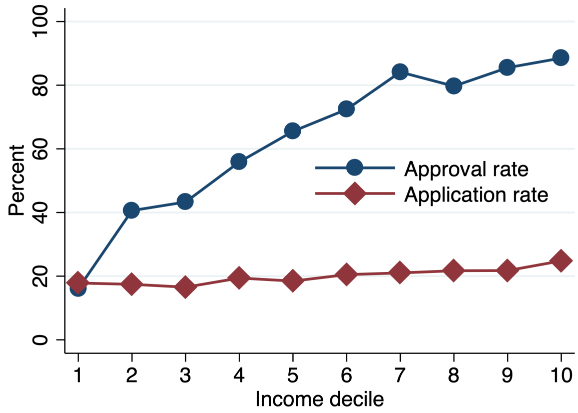

Figure 5: Application rate and approval rates for mortgage applications in the SCF.

Notes: An applicant is a household who submitted a loan application in the past 12 months or who did not submit an application in the past 12 months because he/she thought it would be denied. An application is treated as approved if the household receives full funding. An application is treated as rejected if the request for credit is turned down, or the household receives less credit than what was applied for, or the household did not apply for any credit because it thought it would be denied. Applications for a mortgage, home-based loan, and refinancing are considered.

Figure 5 shows that the average approval rate hides a lot of heterogeneity. The solid-circle line shows how approval rates vary with the household's income. To maintain consistency with the preceding analysis, the boundaries of the income deciles in the figure are kept the same. The approval rate at the median is 60% lower than the average approval rate of 80%. Along the income distribution, we observe substantial heterogeneity in approval rates increasing by a factor of five from 16% in the bottom income decile to roughly 80% in the 7^{th} income decile and 90% in the top decile. The figure, therefore, reveals that the increase by income is non-linear. Approval rates rise steeply up to the 7^{th} decile where they almost level off. This pattern aligns with the rising trend in the extensive margin for home equity documented in Figure 3(a), suggesting that approval rates and participation rates are strongly correlated.[16]

At the same time, the solid-diamond line in Figure 5 shows the share of applicants for a mortgage, which we take as a measure of demand. This share remains virtually constant along the income distribution at 20%. This suggests that the demand for housing is unlikely to be responsible for the strong rise in the extensive margin in the data. In our model, we will calibrate directly to these approval rates by income and allow for housing demand and mortgage applications to adjust endogenously in the presence of such financial frictions.

2.4.2 Pension market

Similar to the mortgage market, the SCF data also provide information on market access to employer-sponsored retirement plans. In the SCF sample, such employer-provided retirement

[16] Excluding mortgage refinancing applications reduces approval rates in the lower half of the income distribution on average by 5pp, thus strengthening our finding regarding the steep gradient in approval rates.

accounts comprise the largest share (64%) of total retirement wealth.[17] Given that we focus on the employer-provided pension plans in this subsection, we restrict the SCF sample to households whose head reports to be working in dependent employment.

Roughly 40% of household heads do not hold an active employer-sponsored pension account, either because their employer does not offer such a plan (28%) or because they are not eligible to be included in the plan (12%). The remaining 60% of workers have active employer-sponsored retirement plans, and we refer to the share of workers with an active retirement account as the *offer rate*. Looking at workers who received an offer for an employer plan, we find that 80% of those households are enrolled in the plan. We refer to this measure as the *pick-up rate*.[18]

The solid-circle line in Figure 6 uncovers the large variation in offer rates along the income distribution (as before, the boundaries of the income deciles are kept unchanged). The offer rate increases by a factor of four from approximately 20% in the bottom income decile to more than 80% in the top decile. This evidence supports the conjecture that households with different incomes do not face equal access to the pension market. The solid-triangle line depicts the variation in the pick-up rates by income. We observe a higher level and a substantially lower variation in pick-up rates than in offer rates. The pick-up rate increases only by a factor of 1.7 from the bottom to the top income decile. Thus, we find that for retirement accounts differences in pick-up rates (demand) contribute partly to the documented variation in the participation rates (Figure 3(a)), but as for housing, the major source of variation comes from offer rates and asset market access.

As in the case of housing, we will directly calibrate the model to be consistent with the empirical evidence on pension market access by income and will include an endogenous acceptance decision to account for the empirical variation in pick-up rates. Hence, in our model, we will follow the large literature on incomplete market models and take frictions as given without microfounding them when exploring their consequences for wealth accumulation.

2.5 Market access, contribution rates, and copay rates

We have shown that the observed differences in asset market participation are consistent with unequal access to asset markets by income. An alternative explanation for differences in asset market participation by income could be differences in interest rates for mortgages, contribution rates of workers, or copay rates of employers to retirement accounts. In Appendix A.4.4, we rely on the SCF data to explore differences by income along these dimensions and conclude that the available evidence does not support such alternative explanations.

For retirement plans, the median employer copay rate by income is constant up to the 6^{th} decile and declining for higher incomes. Declining copay rates tend to lead to lower rather than higher

[17]The second largest account types are individual retirement accounts and Keoghs plans (30% of total pension wealth), which are accounts workers typically opened with their previous employers. To be consistent with the SCF data, we will include rollover from previous employers in our model below. Consistent with the data, we find that about a third of wealth in retirement plans has been rolled over from previous employers.

[18]This naming constitutes a slight abuse of concepts as the offer rate in the SCF data is a stock of households having received an offer rather than the share of households receiving an offer per period of time (flow rate). The same applies to the pickup rates. We will take this into account when calibrating the model to these data.

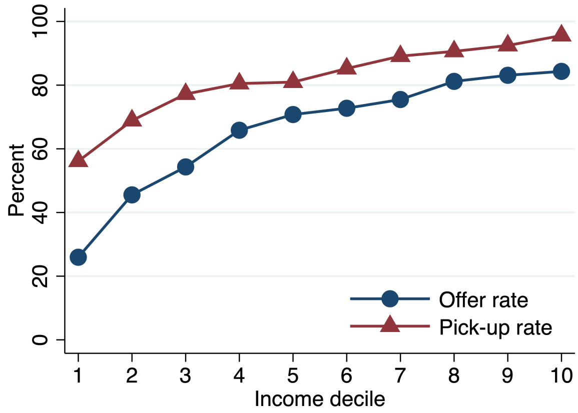

Figure 6: Offer rates and pick-up rates for employer-provided retirement accounts in the SCF.
Notes: The SCF sample is restricted to households whose head works in dependent employment. Income deciles are defined on a full sample to maintain consistency with the preceding analysis. *Offer rate* denotes the fraction of employees whose employer offers a retirement plan and who are eligible to participate. *Pick-up rate* is defined as the proportion of employees who enroll in the employer-sponsored retirement plan out of the total pool of employees who have access to such a plan and are eligible to enroll.

participation rates with income. In addition, we observe that the within-income-group variation is substantially larger than the between-income-group variation.

For worker contribution rates, the median rate varies little with income and shows only a slight increase for the top deciles of the income distribution. The increase is, however, small. As for copay rates, the within-income-group variation in contribution rates is almost an order of magnitude larger than the between-income-group variation. Thus, the available evidence for retirement plans does not support a systematic relationship between copay rates (or contribution rates) and participation behavior. These differences can, however, contribute to within-income-group variation of wealth holdings that our model does not fully account for.

For mortgage interest rates, there is virtually no variation in the median mortgage interest rate across income groups. Importantly, we also do not find variation in the 10^{th} percentile of mortgage interest rates across income groups. If low-income households were less likely to get mortgages with low interest rates approved, this would have resulted in selection and decreasing interest rates with systematically higher interest rates for approved mortgages at the bottom of the income distribution. Our results do not support such selection.

2.6 Sensitivity analysis

We have provided evidence that income, wealth accumulation, and asset market access are strongly correlated. However, several observable factors can potentially account for the relationship between income and wealth. In Appendix A.4.1, we conduct extensive sensitivity analyses controlling for household demographics (age, race, education, marital status, sex) in addition to household income. We find that the relationship between income and wealth remains economically and statistically significant in all cases. We also conduct sensitivity analyses for the income dependence of mortgage approval rates and offer rates for employer-sponsored retirement plans and demonstrate that conditional on a broad set of household characteristics, income deciles have a statistically significant and positive effect on both margins.

Our data source for the main empirical analysis is the SCF but we corroborate the results on income dependence of market access also based on other independent, high-quality data sources. In Appendix A.4.2, we use *Home Mortgage Disclosure Act* (HMDA) data to study the relationship between mortgage approval rates and income. The HMDA data allow us to control in addition to household demographics and income for mortgage characteristics. In particular, we control for the loan-to-value ratio at origination and debt-service-to-income ratios. When we do not control for mortgage characteristics, we find a positive relationship between mortgage approval rates and income that becomes even stronger when controlling for household characteristics, state, and year-fixed effects. Importantly, when we control for the loan-to-value ratio as a wealth-related variable in the regression for approval rates, we estimate a statistically insignificant coefficient. By contrast, the coefficient on the income-related debt-service-to-income ratio is statistically significant, supporting our focus on income for asset market access. These results do not change when we include each of the mortgage controls individually or jointly.

In Appendix A.4.3, we rely on the *National Compensation Survey* (NCS) as an alternative and independent data source to corroborate our results on market access by income for employer-provided pension plans. We rely on publicly available data and data from a special data request to the BLS to study offer and pick-up rates by income. Based on these data, we find that offer rates strongly increase in income from 30% in the bottom decile to 90% in the top decile. As in the SCF data, the pick-up rates vary less than offer rates along the income distribution.

3 The model

Guided by the empirical evidence, we develop a model of asset market participation to quantify the role of heterogeneous asset market access for differences in saving rates and wealth accumulation. In the model, we will focus on life-cycle wealth accumulation as the most important saving motive (e.g. Gourinchas and Parker, 2002; Cagetti, 2003). We will, therefore, abstract from low-return liquid assets that are used to accumulate precautionary savings to smooth out transitory income shocks (Kaplan and Violante, 2014; Bayer et al., 2019).[19] To account for the empirically important role of business wealth, we incorporate an occupational choice allowing us to jointly study the three components of core wealth (home equity, pension plans, and business wealth).

Our model departs from existing models of wealth accumulation with frictional asset markets in two important dimensions. First, we do not consider the saving decision at the intensive margin governed by an Euler equation but instead restrict the saving decision to a discrete choice over pre-determined financial contracts. We assume that pre-specified mortgage contracts and pension plans exist and are being offered to agents. Agents who sign a financial contract commit to executing it, and by executing the financial contract, they accumulate wealth in a portfolio

[19] In Appendix A.3, we demonstrate that the overwhelming share of life-cycle wealth accumulation is accounted for by core wealth.

with home equity and pension accounts. This choice of pre-specified contracts is motivated by our empirical evidence of a roughly constant intensive margin among asset market participants (Section 2.3).

Second, we allow access to financial markets to depend on the agent's labor market situation, thereby introducing income-dependent market access whose extent we estimated from the data (Section 2.4). Hence, our framework emphasizes asset market access as a source of variation in wealth. We demonstrate that this shift in focus from the *how much to save* to the *if to save* allows to closely account for the empirically observed joint distribution of wealth and income.

3.1 Overview

There is a continuum of finitely-lived agents. Agents enter the model at age $j = 1$, participate in the labor market during the first J periods, and retire at age $J+1$. During retirement, the agent may die with a fixed probability $1-\rho$ each period so that the retirement period is of stochastic length $1/(1-\rho)$.

Each model period is split into two stages: a labor market stage and a housing investment stage. At the labor market stage, households receive job offers and business opportunities. Job offers may come with a retirement plan, so the agent needs to decide whether to accept the offer and, conditional on accepting it, whether to enroll in a plan if offered. Business opportunities provide an alternative to dependent employment. Entrepreneurs can always set up a retirement plan on their behalf when starting a new business. At the housing investment stage, agents receive housing offers and decide whether to apply for a mortgage to buy a house.

Agents accumulate wealth in retirement funds p and home equity $h-m$, where h denotes the value of housing and m – the mortgage value. The dividend flow d of an active business is associated with business wealth b (details below). Hence, the agent's total wealth, denoted by w, is given by:

$$w = p + h - m + b \tag{1}$$

and, therefore, matches the definition of core wealth from the empirical analysis.

3.2 Preferences

During the working stage, the agent's instantaneous utility function is additively separable in non-durable consumption c and housing services h:

$$u(c,h) = \alpha \log(c) + (1-\alpha)\log(\underline{h} + h), \tag{2}$$

where α is the weight on non-durable consumption and $\underline{h}$ denotes a minimum level of housing received by each agent that we introduce to abstract from a separate renting market.

3.3 Retirement

We focus on life-cycle wealth accumulation during working life. We model retirement in a reduced form, assuming that agents derive utility from their accumulated wealth and, in addition, receive a lump-sum social security benefit with present value g (g for). We include social security benefits in our analysis, as previous studies have demonstrated that they can create negative incentives for wealth accumulation among low-income households (Hubbard et al., 1995). To derive the utility from wealth in retirement, we solve a deterministic consumption-saving problem where consumption choices depend on total wealth only, agents have per-period utility $u(c) = \log(c)$ and face a fixed probability ρ to survive to the next period. The solution to this problem is to consume a constant fraction $1 - \beta\rho$ of total wealth $w + g$ each period, where $w = p + h - m + b$ is accumulated lifetime wealth. This results in the value function of wealth at the entry of retirement

$$v(w) = \frac{\log(w + g)}{1 - \beta\rho} + q \tag{3}$$

with the constant $q = -\frac{\log(1-\beta\rho)}{1-\beta\rho}$. We will use this value function during retirement in the recursive representation of the model below.

3.4 Labor market stage

At the labor market stage, the agent enters as a worker or entrepreneur. The worker's labor-market state is denoted by l. A worker in a labor-market state l supplies labor inelastically and earns pre-tax labor income $l > 0$. If an agent enters as an entrepreneur, she owns a business whose market value is denoted by $b > 0$. We model business ideas as seeds to Lucas trees that, if realized, provide per-period dividends d. Motivated by the empirical evidence of a constant business wealth-to-income ratio at the intensive margin, we assume that business wealth b corresponds to a constant multiplier Ψ of business income d:[20]

$$b = \Psi d. \tag{4}$$

At the beginning of the labor market stage, agents may receive new labor market opportunities for the next period. We denote the agent's current labor-market state by (l, d), the job offers and business opportunities they receive by $(\tilde{l}, \tilde{d})$, and the chosen employment state for the next period by (l', d'), and we will follow the same convention for other model variables. To keep the model tractable, we assume that a worker never holds business wealth and a business owner does not receive labor income. Therefore, our convention is to set $b = d = 0$ if the agent is a worker ($l > 0$) and to set $l = 0$ if the agent is a business owner ($b = \Psi d > 0$).

In each period, a worker receives a job offer with probability P_W and a business opportunity with probability P_E. The probability of not receiving any job offer or business opportunity is

[20] Business wealth corresponds to the firm's equity, and we implicitly assume that all business equity is intangible capital. The simplifying assumption that agents do not invest physical capital in starting a firm can be viewed as an outcome of a well-functioning venture capital market in which an entrepreneur with an idea always receives financing.

$1 - P_W - P_E$, in which case the agent remains in her current labor market state. The worker's idiosyncratic job offer $\tilde{l} > 0$ follows an AR(1) process:

$$\log \tilde{l} \quad = \quad \rho_W \log l + \varepsilon \quad \text{with } \varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2). \tag{5}$$

Conditional on receiving an offer $\tilde{l}$, the worker faces an exogenous probability $\pi_P(\tilde{l})$ of being offered an employer-sponsored retirement account and $\tilde{\psi}$ is an indicator function equal to one if the job offer comes with an option to open a pension account and zero otherwise.

Workers decide whether to accept the job offer, internalizing that the chosen labor-market state will affect both the probability of being granted a mortgage during the housing investment stage later on and the labor market dynamics in the subsequent periods according to eq. (5). Conditional on accepting the job offer with an option to set up a pension account, a worker must also choose whether to enroll in the plan. If she rejects the job offer, she remains in her current labor-market state with $l' = l$ and $\psi' = \psi$, where ψ indicates whether she is enrolled in a plan with her current employer. We formulate the decision problem in Section 3.8 recursively.

With probability P_E, a worker receives a business opportunity with a dividend flow $\tilde{d}$. For simplicity, we assume that the business opportunity arriving to workers is an i.i.d. random variable drawn from a uniform distribution and is, therefore, uncorrelated with the worker's current job. Workers who receive a business opportunity must decide whether to realize it.

Entrepreneurs derive a dividend income d each period and consume the value of their business b during retirement according to eq. (3). We assume that once agents decide to become entrepreneurs, they cannot switch back to being workers for the remainder of their working lives. Entrepreneurs receive a new business opportunity $\tilde{d}$ each period, which follows a stochastic process equivalent to eq. (5). Upon receiving a new business opportunity, entrepreneurs must decide whether to realize or forgo it. The private pension market for entrepreneurs is assumed to be frictionless, i.e., all business ideas arrive with an opportunity to open a retirement account, and we allow business owners to choose whether they want to set up a private pension account whenever they accept a new business opportunity.

3.5 Retirement accounts

Workers and entrepreneurs receive opportunities to invest in private pension accounts. While business owners can always access the private pension market when receiving a new business opportunity, access to employer-provided retirement accounts is frictional. We assume that the probability of being offered an employer-sponsored retirement account, $\pi_P(\tilde{l})$, is exogenous and depends on the worker's labor-market state $\tilde{l}$.

While the supply of employer-sponsored plans is exogenous, our model generates an endogenous demand for these assets. Retirement accounts are attractive for life-cycle wealth accumulation for two reasons. First, they are subject to preferential tax treatment since contributions are made from pre-tax labor income. Second, the employer matches a portion of the employee's contribution rate. In the SCF data, we find that, on average, employers match on average 80% of employees' contributions. Our framework incorporates both of these features that make retirement plans an important means of wealth accumulation.

We model the retirement plan as a pre-specified financial product where workers commit to paying a fixed share χ_W of their labor income l into the account each period as long as they work for the same employer. The employer matches the employee's contribution at a fixed rate σ_P, and workers can deduct the amount $\chi_W l$ from their taxable income. Each period, the account holder receives a fixed rate of return r_P on her pension account balance p. Therefore, the law of motion for pension wealth reads:

$$p' = (1 + r_P)p + \psi'(1 + \sigma_P)\chi_W l'. \tag{6}$$

If a worker chooses to stay in her current job with an active pension plan ($\psi' = \psi = 1$ and $l' = l$), she continues to pay into the retirement account. Employer-provided pension plans are fully portable across jobs, so when workers move from one job to another, they carry their accumulated retirement savings with them and continue paying into the account as long as they again enroll in a pension plan with the new employer. If the new employer does not offer a plan or the worker does not enroll ($\psi' = 0$), the worker keeps the existing account so that the next period's balance equals $p' = (1 + r_P)p$. Similarly, if the worker becomes an entrepreneur, she keeps her account and may choose to continue paying into it as an entrepreneur.

Entrepreneurs can also accumulate savings in a private pension account, and we allow them to open an account whenever they realize a new business opportunity. Like employer-provided plans, private pension accounts are pre-specified financial products that require paying each period a fixed share χ_E of the dividend payments d. Similar to eq. (6), the law of motion for pension wealth for business owners reads:

$$p' = (1 + r_P)p + \psi'\chi_E d'. \tag{7}$$

3.6 Housing investment stage

At the housing investment stage, agents receive investment opportunities for housing. We model the housing market as a search market where agents receive each period a random opportunity to buy a house of value $\tilde{h} \in \mathcal{H}$. Housing offers follow a log-normal distribution with mean μ_h and variance σ_h^2. We introduce an additional state $\tilde{h} = 0$ in the support $\mathcal{H}$, which we interpret as an opportunity to move to rental housing and consume housing service $\underline{h}$. We denote the distribution of housing offers over the discretized support $\mathcal{H}$ by $\lambda(\tilde{h})$. A housing offer is an i.i.d. random variable across agents and time.

Agents can always reject a housing offer and remain in their current house. Agents who prefer to downsize and draw a corresponding offer ($\tilde{h} < h$) always get their mortgage application approved. If the agent is a first-time home buyer ($h = 0$) or decides to increase the value of existing housing and draws a corresponding offer ($\tilde{h} > h$), she must take out a mortgage whose value is denoted by $m \geq 0$. Consistent with our empirical analysis, we assume that the probability of being approved for a mortgage, denoted by $\pi_H(y')$, depends on the household's labor market situation $y' \in \{l', d'\}$.[21] Workers and entrepreneurs face the same frictions in accessing the housing market, conditional on their income. To parameterize the probability schedule $\pi_H(y')$, we will directly use the approval rates for mortgage applications from the empirical section (Figure 5). The probability of being granted a mortgage is independent of the housing value $\tilde{h}$.

While access to the mortgage market is frictional and exogenous, two forces generate endogenous demand for housing in our model. First, housing provides agents with a utility flow during the working stage according to eq. (2). Second, the accumulation of home equity $h - m$ contributes to life-cycle savings as agents consume their accumulated wealth during retirement (eq. 3). Hence, housing serves a dual role as a consumption and investment good.

Like retirement plans, mortgages arrive as pre-specified financial contracts. A contract has a fixed interest r_M and requires a fixed repayment share κ of the housing value each period. The law of motion for mortgages reads:

$$m' = (1 + r_M)m + h' - h - \mathbb{1}_{m>0}\kappa h, \tag{8}$$

where $\mathbb{1}_{m>0}$ is an indicator function that is one if the current mortgage balance m is positive. A mortgage balance of zero $(m = 0)$ applies to non-participants in the housing market and first-time home buyers, for whom it also holds that $h = 0$, and outright homeowners with $h > 0$ and $m = 0$. Although there is no downpayment for first-time home buyers, most home buyers in our model will upgrade their homes $(h' > h)$ and make downpayments by using their accumulated home equity $(h - m)$. For the new house only additional housing investments will be financed through a mortgage. Downsizing households $(h' < h)$ will repay their mortgage. All agents with outstanding mortgage debt $(m > 0)$ have to pay interest r_M on their debt balance m and make an amortization payment κh. As there are no unexpected income declines in the model, agents have no incentive to default on mortgage contracts and we therefore abstract from a default decision. The fact that mortgages as financial contracts constitute financial commitments limits the demand for housing.

3.7 Budget constraints

An agent's taxable income is her market income $y \in \{l, d\}$ net of mortgage interest payments and contributions to retirement plans that can be made out of pre-tax income and reduce taxable income. Taking into account our convention that whenever $l > 0$, $d = 0$, and vice versa, the taxable income $\bar{y}$ becomes

$$\bar{y} = \underbrace{(l + d)}_{\text{market income}} - \underbrace{\psi(\chi_W l + \chi_E d)}_{\text{retirement contributions}} - \underbrace{\mathbb{1}_{m>0}(r_M m)}_{\text{mortgage interest payments}} . \tag{9}$$

[21] Importantly, Table 11 in the appendix shows that there is no significant effect of loan-to-value ratios on approval rates in the HMDA data once we control for income and household characteristics. We, therefore, abstract from this constraint in our model.

Observe that only agents who have an active pension plan ($\psi = 1$) pay contributions into the retirement account, and only agents who have a positive mortgage balance make mortgage interest payments ($\mathbb{1}_{m>0} = 1$). Conditional on the current labor market situation and housing, the agent's non-durable consumption in the current period equals the after-tax income net of pension contributions and amortization payments κh in case of a positive mortgage balance:

$$c = y - T(\bar{y}) - \psi(\chi_W l + \chi_E d) - \mathbb{1}_{m>0}(\kappa h), \tag{10}$$

where $T(\bar{y})$ denotes the income tax function specified in the calibration section.

Hence, consumption corresponds to the income net of the financial contracts and taxes and follows hand-to-mouth dynamics, with few poor hand-to-mouth households and a majority of wealthy hand-to-mouth households, in line with the empirical evidence in Kaplan et al. (2014). Note that wealth dynamics of wealthy hand-to-mouth agents in our model differ from Kaplan and Violante (2014). Whereas agents in their model build up liquid wealth to deposit it infrequently into their illiquid account and follow hand-to-mouth behavior if their wealth is tied up in illiquid assets, wealth accumulation in our model happens constantly by executing long-term financial contracts. As financial contracts are commitments, consumption dynamics of households follow income dynamics one-for-one.

3.8 Dynamic programming problem

We summarize the agent's joint labor market and financial state at the beginning of a period by her age j and a state vector $\boldsymbol{x} \equiv (l, d, \psi, p, h, m)$, where (l, d) denotes the labor market state, ψ the indicator function whether the agent has an active pension plan, p the pension account balance, h the housing stock, and m the mortgage balance.

We describe the dynamic programming problem proceeding backward. Agents enter retirement at age $J+1$. Their value function is $v(w)$ derived in Section 3.3. During working life, each period is divided into two consecutive stages: the labor market stage and the housing investment stage. Proceeding backward, we start with the housing investment stage, where the agent receives housing offers $\tilde{h}$ and decides whether to apply for a mortgage or not.

To save on notation, let the state vector $\boldsymbol{z}' = (l', d', \psi', p')$ summarize the outcome of the labor market stage of the current period. The decision about moving to a smaller house $\tilde{h} \leq h$ does not involve risk about mortgage approval, so if the agent takes the offer for a smaller house, her value function is given by $V_{j+1}(\boldsymbol{z}', \tilde{h}, m')$, and if she declines it, the value is given by $V_{j+1}(\boldsymbol{z}', h, m')$. If the agent applies for a mortgage for the offer $\tilde{h} > h$ and receives approval for financing, which happens with probability $\pi_H(y')$, her value function is given by $V_{j+1}(\boldsymbol{z}', \tilde{h}, m')$. If she does not apply for a mortgage the offer or she is not approved, she remains in the current house, and her continuation value is $V_{j+1}(\boldsymbol{z}', h, m')$. In both cases, the mortgage balance m' follows the law of motion from eq. (8).

At the beginning of the housing investment stage, agents form expectations over housing offers, and their expected value can be summarized as follows:

$$\begin{aligned}\bar{V}_j(\boldsymbol{z}',h,m) &= \beta\sum_{\tilde{h}>h}\lambda(\tilde{h})\left(\pi_H(y')\max\left\{V_{j+1}(\boldsymbol{z}',\tilde{h},m'),V_{j+1}(\boldsymbol{z}',h,m')\right\}+(1-\pi_H(y'))V_{j+1}(\boldsymbol{z}',h,m')\right)\\ &\quad+\beta\sum_{\tilde{h}\leq h}\lambda(\tilde{h})\max\left\{V_{j+1}(\boldsymbol{z}',\tilde{h},m'),V_{j+1}(\boldsymbol{z}',h,m')\right\}. \end{aligned} \tag{11}$$

Observe that if the agent does not apply for a mortgage for a housing offer $\tilde{h} > h$, the expression in round brackets in the first line above boils down to $V_{j+1}(\boldsymbol{z}',h,m')$, which is the value of staying in house h.

At the labor market stage, the agent may receive a job offer or business opportunity $(\tilde{l},\tilde{d})$. If an agent receives a job offer, she also learns whether it comes with a retirement account, which is denoted by $\tilde{\psi}\in\{0,1\}$. For business owners, access to the pension market is frictionless, always leading to $\tilde{\psi}=1$ (so we will omit it to simplify notation). Note that agents can choose to enroll in a pension plan (if offered one), only when they switch jobs or start a new business.

The decision problem of the agent at the labor market stage for different combinations of offers can then be expressed as

$$V_j^{WP}(\boldsymbol{x};\tilde{l},1) = \max\{\bar{V}_j(\boldsymbol{x}),\max\{\bar{V}_j(\tilde{l},\tilde{d},0,p,h,m),\bar{V}_j(\tilde{l},\tilde{d},1,p,h,m)\}\}, \tag{12}$$

$$V_j^{WN}(\boldsymbol{x};\tilde{l},0) = \max\{\bar{V}_j(\boldsymbol{x}),\bar{V}_j(\tilde{l},\tilde{d},0,p,h,m)\}, \tag{13}$$

$$V_j^{E}(\boldsymbol{x};\tilde{d}) = \max\{\bar{V}_j(\boldsymbol{x}),\max\{\bar{V}_j(\tilde{l},\tilde{d},0,p,h,m),\bar{V}_j(\tilde{l},\tilde{d},1,p,h,m)\}\}, \tag{14}$$

where superscripts indicate the occupation (W for worker and E for entrepreneur) and the opportunity to open a pension account (P when a retirement plan is available and N if not). Specifically, the first line is the worker's decision problem with a job offer $\tilde{l}$ that includes a retirement account ($\tilde{\psi}=1$). $\bar{V}_j(\boldsymbol{x})$ stands for the agent's utility at the beginning of the housing stage, introduced above, if the agent does not change the labor market state, leading to $l=l'$, $d=d'$, $\psi=\psi'$, and p' follows from eq. (6)–(7). $\bar{V}_j(\tilde{l},\tilde{d},0,p,h,m)$ denotes the continuation value from accepting the job offer but rejecting the investment opportunity, while $\bar{V}_j(\tilde{l},\tilde{d},1,p,h,m)$ denotes the value from accepting the job offer and enrolling in the plan. The second line is the worker's decision problem without the opportunity to invest in a retirement plan ($\tilde{\psi}=0$). The third line corresponds to the case when the agent receives a business opportunity $\tilde{d}$ and can set up a retirement plan if preferred.

Combining eq. (12)– (13), we further define

$$V_j^{W}(\boldsymbol{x};\tilde{l}) = \pi_P(\tilde{l})V_j^{WP}(\boldsymbol{x};\tilde{l},1)+(1-\pi_P(\tilde{l}))V_j^{WN}(\boldsymbol{x};\tilde{l},0)$$

as the expected value of receiving a job offer $\tilde{l}$, where $\pi_P(\tilde{l})$ is the probability of being offered an employer-sponsored retirement account. Combining the case of receiving a job offer, a business opportunity, and the event of receiving no offer, occurring with probabilities P_W, P_E, and $1-P_W-P_E$, respectively, the value function at the beginning of the labor market stage is given by

$$V_j(\boldsymbol{x}) = u(c,h)+P_W\mathbb{E}[V_j^{W}(\boldsymbol{x};\tilde{l})\mid l]+P_E\mathbb{E}[V_j^{E}(\boldsymbol{x};\tilde{d})]+(1-P_W-P_E)\bar{V}_j(\boldsymbol{x}).$$

Since a business opportunity arriving to workers is an i.i.d. random variable, the second expectation operator is unconditional. Consumption follows from the budget constraint in eq. (10) as the residual of the decision from the previous period and the current labor market status.

Finally, since business owners do not return to dependent employment by assumption, their value function at the beginning of the period reads:

$$V_j(\boldsymbol{x}) = u(c, h) + \mathbb{E}\left[V_j^E(\boldsymbol{x}; \tilde{d}) \mid d\right], \tag{15}$$

where $V_j^E(\boldsymbol{x}; \tilde{d})$ is defined in eq. (14).

4 Calibration

We calibrate the model to the SCF data from the empirical analysis in Section 2. We set some parameters outside the model and calibrate the remaining parameters inside the model. Although all of the latter parameters must be calibrated jointly, we provide an intuitive discussion about the relationship of each parameter to a corresponding data moment. Table 2 lists all model parameters and reports the corresponding source or the empirical target.

To initialize the distribution of agents, we proceed as follows. All agents are assumed to enter the model at age 20. The initial conditions for these agents are parameterized based on the subsample of households aged 20–24 from the SCF data. When comparing the model to the data, we discard this age group to avoid the influence of the initial distribution, focusing instead on the age group 25–65, consistent with the sample selection in the empirical section.

At age 20, agents enter the model in the lowest labor-market state and without any pension wealth, but there is a share of agents who work for an employer offering a retirement account that we set to 30%. The share of entrants with positive home equity is 17%. We set their initial house value and the mortgage balance such that their home equity-to-income ratio, $(h - m)/y$, is equal to 3.9 and the mortgage-to-housing ratio, m/h, is 53%.

A subset of the model parameters, for example, those governing the stochastic process for labor income or business ideas, are occupation-specific. To calibrate these parameters, we define a household in the SCF data as a worker if the household head reports working for someone else and as a business owner if the household head reports actively managing a business.

Demographics and utility

The model period is one year. Agents enter at age 20 and retire at age 66. After dropping the first five years, the duration of the working life is $J = 41$ periods. The average duration of retirement is set to $1/(1 - \rho) = 20$ years, such that the survival probability rate ρ is 0.95.

Table 2: Calibrated parameters.

Parameter	Interpretation	Value	Source/Target
		Parameters set outside the model:	
Demographics:			
J	Duration of working stage	41	
ρ	Duration of retirement is $1/(1-\rho)$	0.95	
Frictional access to pension and housing markets:			
$\pi_H(y)$	Mortgage approval rates by income	Figure 5	SCF
Financial contracts and returns:			
σ_P	Employer's share paid into account	0.8	SCF
χ_W	Contribution rate for workers	0.05	SCF
χ_E	Contribution rate for entrepreneurs	0.09	$(1+\sigma_P)\chi_W$
Ψ	Business wealth multiplier	1.5	SCF
r_M	Interest rate on mortgage	0.02	FRED
Income taxation and Social Security:			
(τ_1, τ_2)	Income tax level and progressivity	(0.08, 0.22)	Heathcote et al. (2020)
R	Replacement rate	0.39	Brendler (2023)
		Internally calibrated parameters:	
Labor market:			
$(\rho_W, \sigma_\epsilon^2)$	AR(1) process for labor income	$(0.95, 0.06)$	Variance of labor income by age
(ρ_E, σ_η^2)	AR(1) process for bus. income	$(0.89, 0.16)$	Bus. wealth/labor income by age
P_W	Probability to receive job offer	0.48	Average job tenure = 7 years
P_E	Probability to receive business idea	0.02	Extensive margin for business wealth = 13%
Retirement accounts:			
β	Discount factor	0.98	Extensive margin for pension wealth = 64%
r_P	Return on pension wealth	0.004	Intensive margin for pension wealth = 1.0
$\pi_P(l)$	Offer rates for employer-provided plans	Schedule	See text.
Housing:			
α	Weight on non-durable consumption	0.96	Extensive margin for home equity = 68%
κ	Amortization rate	0.06	Intensive margin for home equity = 1.6
(μ_h, σ_h^2)	Log-normal distrib. of housing offers $\tilde{h}/Y$	$(1.45, 0.84)$	Mean(h/Y) = 3.0, Var(h/Y) = 11.1
$\lambda(\tilde{h}=0)$	Probability to receive rental housing	10.0%	Mortgage application rate = 20%
$\underline{h}$	Housing for renters	0.1	Lowest positive point in the housing grid $\mathcal{H}$

Labor markets

The idiosyncratic job offers $\tilde{l}$ follow an AR(1) process given by eq. (5), which we approximate using a discrete-state Markov chain. We match the variance profile of log income for workers by age from the SCF. The slope of the profile informs the dispersion of offers σ_ϵ^2, and the autocorrelation ρ_W is informed by the curvature of the profile. Importantly, the calibration of these parameters cannot be done outside the model because agents have an acceptance-rejection decision for job offers, such that the age profile of income is endogenous.[22] In Appendix B (Figure 25(a)), we show the fit to the empirical variance profile for labor income by age. Finally, the probability of receiving a job offer, P_W, is calibrated to match the average job tenure in the U.S. labor market of 7 years (Jung and Kuhn, 2019).

The idiosyncratic process for business ideas of entrepreneurs, $\tilde{b}$, follows an AR(1) process equivalent to eq. (5) with corresponding parameters (ρ_E, σ_η^2). As for the process of labor mar-

[22] Low et al. (2010) demonstrate the importance of accounting for selection when estimating income risk due to the job change.

ket opportunities, we approximate it using a discrete-state Markov chain. As the number of observations for business wealth in some age groups is very small in the SCF, we target the average ratio of business wealth to income by age rather than the variance to calibrate (ρ_E, σ_η^2). As workers in the model decide whether to become entrepreneurs or stay with their current employer, the business wealth-to-labor income profile is informative about the dispersion of business opportunities. Figure 25(b) in the appendix shows the model fit of this margin.

Finally, we calibrate the arrival rate of business ideas P_E to match the average extensive margin for business wealth by income equal to 13% from the SCF sample. Note that the intensive margin of business wealth in the model is constant along the income distribution and equal to the business wealth multiplier Ψ in eq. (4). We set $\Psi = 1.5$ to match the intensive margin in the SCF data (Figure 3(b)).

Housing and pension markets

We parameterize the income-specific access rates to the mortgage market, $\pi_H(y)$, outside the model using the income-dependent approval rates for mortgage applications from the empirical analysis (Figure 5). To smooth out these probabilities, we fit a function of the form $\alpha_0+\alpha_1/y$. We proceed differently when calibrating the offer rates for employer-provided pension plans, $\pi_P(y)$. As already mentioned in the empirical section, the observed offer rate in the SCF data is a *stock* of households having received an offer rather than the share of households receiving an offer per period of time. For this reason, we assume that the (unobserved) flow offer rates, $\pi_P(y)$, follow the same functional form as the mortgage approval rates, and we calibrate its parameters inside the model matching the profile of stock offer rates observed in the data (Figure 6). Appendix B.2 displays the probability schedules $\pi_H(y)$ and $\pi_P(y)$ used in the model and provides additional discussion on their specification.[23]

We calibrate four remaining parameters of the housing and pension markets inside the model, targeting the average extensive and intensive margin for home equity and pension wealth (we do not target heterogeneity by income). We calibrate the relative weight on non-durable consumption, α, as specified in equation (2), by targeting the average extensive margin for home equity at 68%. Similarly, the amortization rate κ is calibrated to match the average intensive margin of home equity, equal to 1.6. Furthermore, the discount factor β is determined by targeting the average extensive margin for pension wealth at 60%.

Housing offers $\tilde{h} \in \mathcal{H}$ follow a (discretized) log-normal distribution with mean μ_h and variance σ_h^2. We normalize housing values in the model and data by the respective (economy-wide) average household income and calibrate the two parameters to match the mean and variance of housing (conditional on homeownership) as observed in the SCF. We add to the housing grid the lowest state equal to zero, corresponding to the case when the household receives a renting offer, and set the probability of receiving such an offer inside the model, targeting the average mortgage application rate of 20%. The subsistence level of housing, $\underline{h}$, is set to the

[23] We demonstrate that conditional on household income, there is little variation in offer rates for employer-sponsored pension plans across different age groups in the SCF, which is why we chose to specify offer rates in the model as a function of household income only. We provide similar evidence for mortgage approval rates.

lowest positive value in the grid for housing offers $\mathcal{H}$. Finally, we fix the rate of return on a mortgage, r_M, at 2%.[24]

In the SCF sample, the average employee's share of income paid into the retirement account is 5% and the employer adds an average contribution rate of 80%, so we set χ_W and σ_P in eq. (6) to 0.05 and 0.8, respectively. Entrepreneurs contribute the sum of the employer's and the employee's shares, i.e., $\chi_E = \chi_W(1 + \sigma_P)$ is equal to 9%.

Income taxation and Social Security

Following Heathcote et al. (2020), we specify the income tax function as:

$$T(\bar{y}) = \bar{y} - \mathcal{I}(1 - \tau_1)(\bar{y}/\mathcal{I})^{1-\tau_2},$$

where $\bar{y}$ is the agent's taxable income defined in eq. (9) and $\bar{y}/\mathcal{I}$ is the agent's taxable income normalized by the economy-wide average income $\mathcal{I}$. We set income tax progressivity τ_2 to 0.22 based on the estimates in Heathcote et al. (2020) and the income tax level τ_1 to 0.08, which implies that the mean-income agent pays an average income tax of 8%.

We assume retirees receive each period a constant fraction $R = 39\%$ of the economy-wide average income $\mathcal{I}$ as public pension benefit (Brendler, 2023). The present discounted value of all future pension benefits is then given by $g = R\mathcal{I}(1 + r_P)/(1 + r_P - \rho)$, which we use to compute the value of retirement in eq. (3).

5 Findings

In this section, we first demonstrate that a model with income-dependent asset market access closely fits the empirical facts on wealth accumulation, the joint distribution of income and wealth, and saving rate differences by income. In the second step, we conduct comparative statics to explain the model's wealth accumulation mechanism.

5.1 Model fit

Figure 7 shows the intensive and extensive margin of the three components of wealth, home equity, pension wealth, and business equity, along the income distribution from the calibrated model together with their empirical counterparts. Panels (a) and (b) show the model fit of the extensive margin and the intensive margin for home equity. The model tracks the empirical profiles of both margins closely. Importantly, it matches the strong increase at the extensive margin and the declining intensive margin in the lower half of the income distribution. Panels (c) and (d) display the model fit for the extensive and intensive margin of pension wealth.

[24] According to the Federal Reserve Economic Data (FRED), the 30-year fixed rate mortgage average in the U.S. during 2010–2019 was 4% (see https://fred.stlouisfed.org/series/MORTGAGE30US). Assuming an inflation rate of 2%, we obtain the parameterized value of r_M.

Again, the model accurately matches both margins across income groups.[25] Panels (e) and (f) show the fit for the extensive and intensive margin for business wealth by income. Despite being a stylized model of entrepreneurial activity, the model matches the profiles of the extensive margin and intensive margin of business wealth strikingly closely. Note that this close fit across asset classes does not follow as part of our calibration strategy since we only target the average extensive and intensive margins but not their heterogeneity by income. By contrast, the close fit corroborates the view of this paper that heterogeneous asset market access is a key driver of differences in wealth accumulation.[26]

Figure 8(a) presents wealth-to-income ratios by income as a measure of the joint distribution of income and wealth. It shows that the model matches closely the characteristic property of increasing wealth-to-income ratios by income.[27] The model misses only the non-linear increase of wealth-to-income ratios towards the top decile. Matching wealth holdings at the very top of the distribution arguably requires a model with additional drivers of wealth accumulation. As the focus of our analysis is on households outside the very top of the distribution, we abstain from incorporating such drivers in our model.

Notably, the endogenous income process in our model closely matches the overall Gini index for income, which is 0.52 in the model compared to 0.51 in the SCF (Table 3). Looking at the Gini index for wealth, we find that the model generates a Gini index of 0.68, which implies that wealth inequality is 31% (16pp) higher than income inequality. In the data, we find a similar amplification effect since the (core) wealth Gini index of 0.71 exceeds the income Gini by 39% (20pp). Decomposing the wealth Gini index in between-income-group and within-income-group inequality (Pyatt, 1976), we find that between-group inequality accounts for 89% of wealth inequality in the SCF compared to 98% in the model. The model's higher contribution of between-group inequality indicates that the lower overall wealth inequality must result from too low within-group wealth inequality, which is consistent with our modeling choice of abstracting from transitory income risk and other drivers of within-income-group wealth inequality such as return heterogeneity (see Appendix A.4.4). However, the model accounts for between-income-group inequality as the main driver of wealth inequality (Figure 8(a)).

Table 3: Gini index for income and net worth distributions in the model and data.

	Income	Wealth	Amplification, %
SCF	0.51	0.71	39.2
Model	0.52	0.68	30.8

Notes: *Amplification* is a percentage change in the Gini index for net worth relative to the Gini index for income.

Figure 8(b) shows the evolution of wealth-to-income ratios by age, which is an additional non-

[25] As we reported in Section 2.4, the share of workers with an active retirement account is 60% in the SCF. This number comes very close to the respective value in the model which is 56%.

[26] Appendix C.1 documents the model fit of the demand side for housing and retirement accounts. We show that the model is consistent with the empirically observed constant mortgage application rates by income and increasing pick-up rates for pension accounts by income.

[27] To provide a more granular decomposition, Figure 28 in Appendix C.1 shows how the model fits the wealth-to-income ratios for each wealth component. Table 12 in the same appendix further reports the shares of each wealth component by income quintile in the model and in the data.

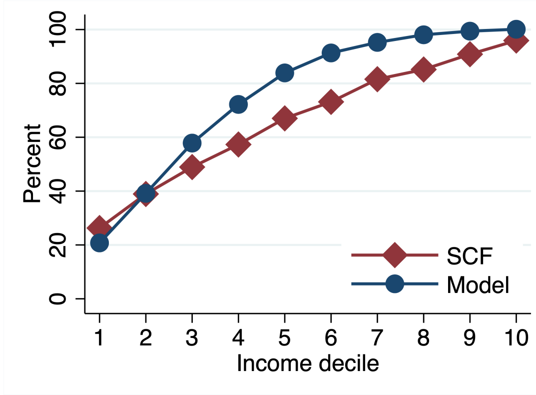

(a) Home equity: Extensive margin

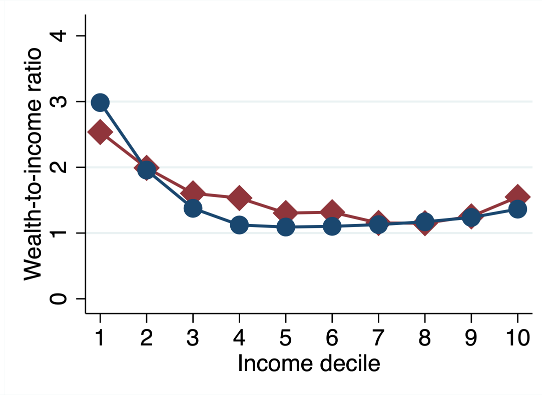

(b) Home equity: Intensive margin

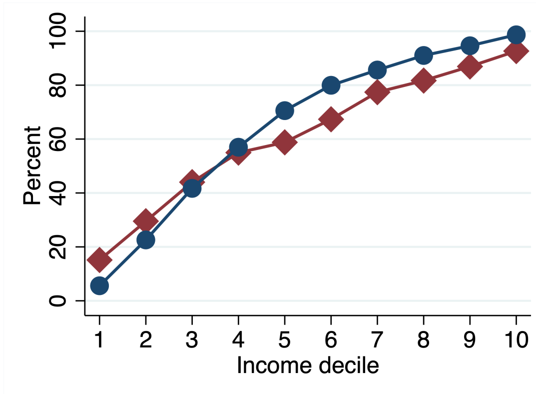

(c) Pension wealth: Extensive margin

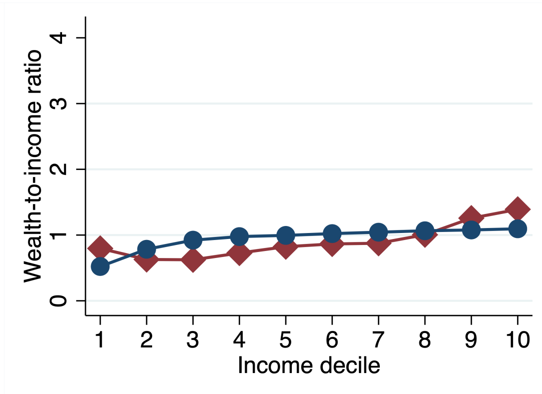

(d) Pension wealth: Intensive margin

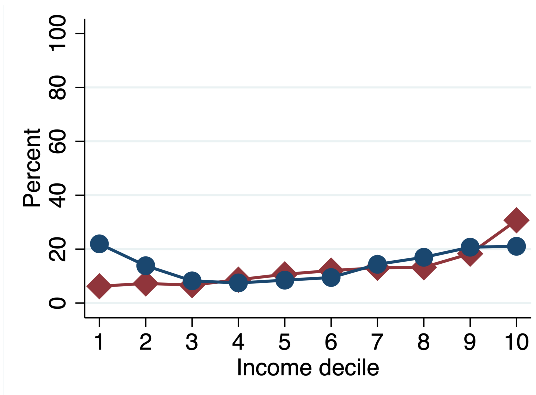

(e) Business wealth: Extensive margin

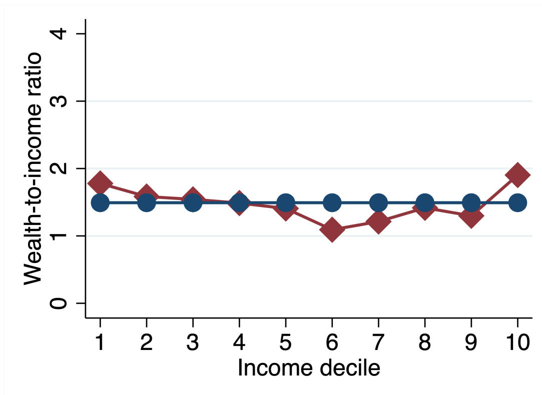

(f) Business wealth: Intensive margin

Figure 7: Extensive and intensive margins in the model and data.

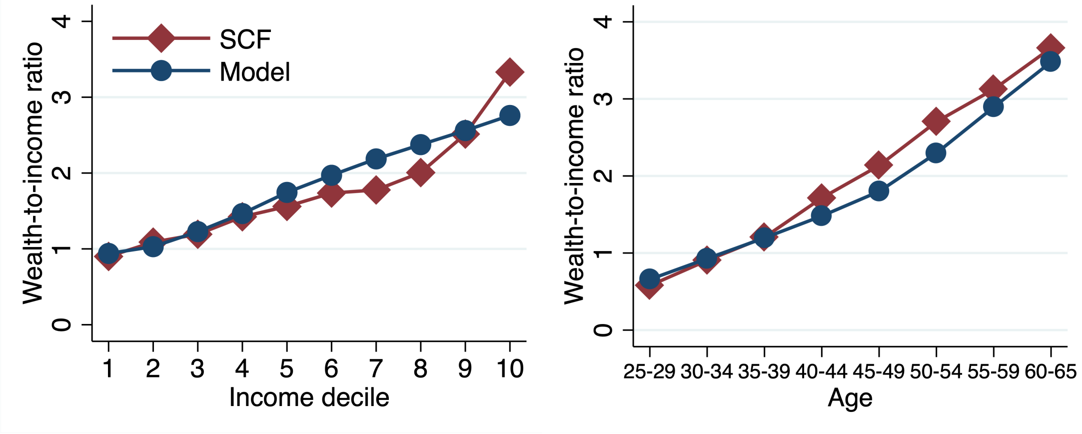

(a) Income profile

(b) Age profile

Figure 8: Wealth-to-income ratios for wealth in the model and data.

targeted dimension of wealth accumulation. The model achieves an excellent fit along this dimension. Models of life-cycle wealth accumulation typically struggle to fit the almost linear increase of the wealth-to-income ratio as they combine a deterministic life-cycle income profile and frictionless access to asset markets. Households know that income grows later in life and, therefore, start accumulating wealth only when income is above the average deterministic life-cycle trend. It is the combination of pre-specified financial contracts and frictional market access in our model that leads to a constant accumulation of savings starting early in the life cycle.

Finally, Table 4 compares saving rates by income in the model with the empirical evidence from Dynan et al. (2004), where we generate model-simulated data to conduct the counterpart to their empirical estimation. We see that the model is consistent with the empirical fact of very low saving rates at the bottom of the income distribution and strongly increasing saving rates with income. Although the model saving rates tend to be lower than the point estimates in each income quintile, they are not statistically significantly different from their empirical counterparts. The largest discrepancy is the increase in saving rates between the 4^{th} and 5^{th} quintiles. The model underestimates the 6.3pp difference between these two quintiles, generating a 2.4pp difference only. It should be noted, however, that Dynan et al. (2004) include non-core wealth in their estimation. As we have documented in the empirical section (Figure 1), there is a widening gap between total and core wealth at the top of the income distribution where most non-core wealth is held, such that the increase in saving rates between income groups is naturally higher at the top income quintiles.

Taking this into account, the model provides a close fit of the empirical differences in saving rates by income and, therefore, offers a new interpretation of these differences. Through the lens of the model, the variations in saving rates are not primarily due to differences in the proportion of income saved but rather due to differences in the share of households that save. These differences, in turn, arise from variations in market access rather than a differing demand for wealth accumulation. The idea that disparities in asset market access can significantly influence observed differences in saving rates has been proposed in the discussion by Dynan et al. (2004, p. 436). Our paper provides a conceptual framework that provides support for this conjecture.

Table 4: Saving rates by current income in the model and data.

	Q1	Q2	Q3	Q4	Q5
Data	0.014	0.090	0.111	0.173	0.236
	(0.019)	(0.029)	(0.032)	(0.027)	(0.040)
Model	0.001	0.098	0.105	0.137	0.161

Notes: Row *Data* reprints the estimation results of median regressions of saving rate on current income from Dynan et al. (2004) (their Table 3, second column). The authors define the saving rate variable as the change in real net worth in the SCF. We reproduce standard errors from their regression in brackets below each coefficient estimate. We display only the estimates for the income quintile dummies, omitting the results for other regressors. Row *Model* shows the estimation results from the same regression run inside our model.

5.2 Income-dependent market access and wealth accumulation

In the second step, we apply the calibrated framework to explain the model's wealth accumulation mechanism and to quantify the impact of income-dependent market access. We conduct two counterfactual experiments. In the first experiment, we assume that agents across the entire income distribution face the same market access in the mortgage and pension markets. Such an homogeneous market access is the prevailing assumption in the class of portfolio choice models with financial frictions discussed in the literature review. Our experiment allows us to quantify how much of the observed differences in wealth and saving rates arise due to the heterogeneity in market access. In the second experiment, we remove financial frictions in the mortgage market, while leaving frictions in the pension market unaffected. This experiment quantifies the extent of asset reallocation that occurs when one market becomes more accessible, for example, due to policy changes or innovations in financial markets.

5.2.1 Homogeneous access to mortgage and pension markets

In the first experiment, we assume that all agents have the same access to the mortgage and pension market independent of their income. Specifically, we set the mortgage approval rate, $\pi_H(y)$, and the offer rate for employer-sponsored accounts, $\pi_P(y)$, to their respective economy-wide averages, 62% and 64% (Figure 9). This change implies that market access to the mortgage market improves for agents with below-median incomes and worsens for those with above-median incomes. The change in access rates is the strongest for the lowest-income households, for whom mortgage approval rates triple from 20% to 60%. The change at the top is smaller with approval rates decreasing by one-third from 90% to 60%. Similarly, offer rates for employer-sponsored accounts increase for all workers with incomes below the 6^{th} decile and deteriorate for those above. The lowest-income workers see their probability of finding an employer offering a retirement plan increase by a factor of 2.3, from 30% to 70%, while chances for workers in the top income decile decline only slightly, from 80% to 70%.

Figure 10 shows the resulting changes at the extensive and intensive margins for home equity and pension wealth compared to the baseline economy. Panel (a) reveals that participation in the housing market rises for all households up to the 6^{th} income decile, with the most significant

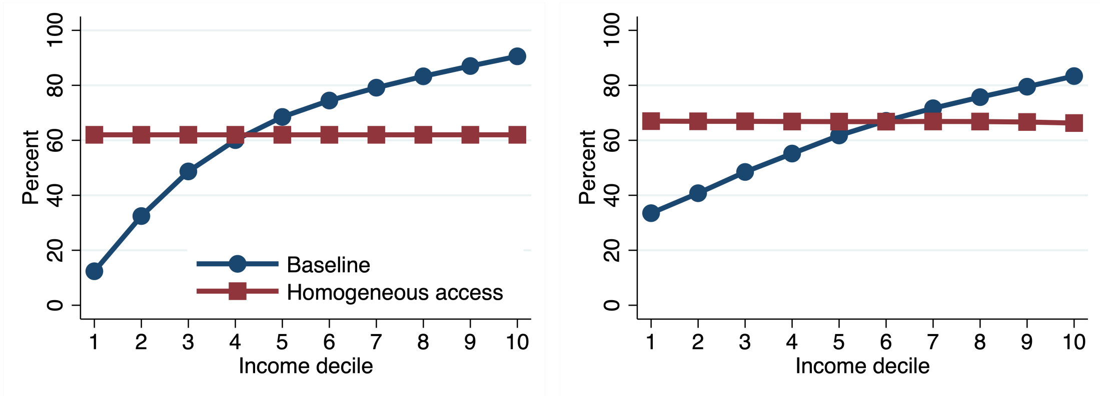

(a) Mortgage approval rates $\pi_H(y)$ (b) Offer rates $\pi_P(y)$

Figure 9: Financial frictions in the baseline model and homogeneous access experiment.

Notes: The figure shows the exogenous change in financial frictions in the experiment compared to the baseline model. The homogeneous access experiment sets both the mortgage approval rates, $\pi_H(y)$, and the offer rates for employer accounts, $\pi_P(y)$, to their respective economy-wide averages (62% and 64%) for all income groups.

increase occurring in the bottom decile, where participation jumps by a factor of 2.5 from 20% to 50%. In contrast, participation rates in the top three income deciles remain virtually unchanged. Although high-income households experience some deterioration in market access, this has only a negligible effect on their extensive margin. However, at the intensive margin in Panel (b), there is a downward shift in wealth-to-income ratios along the entire income distribution with the most substantial reduction of 23% (from 3.0 to 2.3) in the bottom income decile. In the top decile, the decline is smaller, but with 12% still substantial. The lower intensive margin at the bottom is a direct result of better market access: households that were largely excluded from the housing market in the baseline model purchase homes more frequently now, but on average, they buy smaller homes compared to the baseline. Conversely, the lower intensive margin at the top is due to worse market access, which makes it harder for incumbent homeowners to upgrade to larger homes when their incomes improve over the life cycle. As a result, the intensive margin decreases at higher incomes.

Looking at the extensive margin for pension wealth in Figure 10(c), we find that the most substantial changes happen in the lower half of the income distribution, where participation surges on average by 17pp. This strong increase in participation highlights that lower-income households are severely constrained in accessing the pension market in the baseline model. Note, however, that we do not find any increase in participation in the pension market at the very bottom of the income distribution. The reason is the availability of the social security system that provides households with wealth in retirement and therefore lowers the incentives for life-cycle savings, consistent with Hubbard et al. (1995). Furthermore, we also find a modest increase in participation among households with above-median incomes, specifically those in the 6^{th} and 7^{th} income deciles. This increase results from better access to the pension market earlier in life when the agents' incomes are lower on average. In contrast, agents in the top two income deciles show no change in participation. Even though these agents face a roughly 20pp lower probability of being offered a retirement account, they receive a sufficiently large number

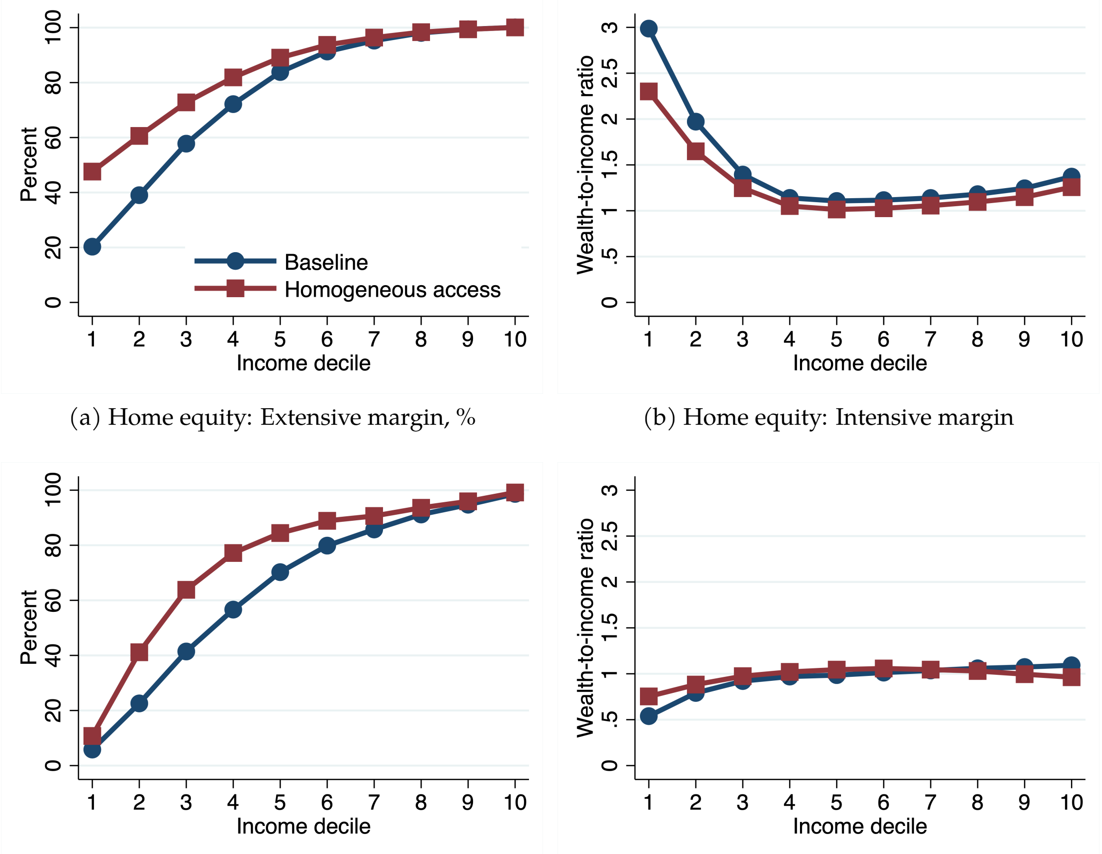

(a) Home equity: Extensive margin, %

(b) Home equity: Intensive margin

(c) Pension wealth: Extensive margin, %

(d) Pension wealth: Intensive margin

Figure 10: Extensive and intensive margins in the home equity and pension markets in the baseline model and homogeneous access experiment.

Notes: The homogeneous access experiment sets mortgage approval rates and offer rates for employer-sponsored retirement accounts to their economy-wide averages for all income levels (Figure 9).

of offers during the early and mid stages of their careers to get access to the pension market.

As for the intensive margin in the pension market, recall that employer-sponsored pension plans in our framework are assumed to be pre-specified financial products, with each participant required to contribute the same percentage of their income. This assumption implies that the intensive margin of pension wealth responds only insignificantly to changes in market access. Changes at the intensive margin primarily reflect changes in the timing when agents start participating in the pension market rather than how much more they contribute. In line with this intuition, Figure 10(d) documents only a minor change at the intensive margin of the pension market. In the policy experiment below, the intensive margin will adjust after we introduce an additional financial product with varying contribution rates.

Combining these changes, Figure 11 shows the percentage change in wealth-to-income ratios across income deciles. We report the percentage change in wealth-to-income ratios for wealth, home equity, and pension wealth in a given income decile in the experiment relative to the respective ratio in the benchmark economy.[28]

[28]Income deciles change little across experiments. Also, business wealth (not shown) rises, on average, by 8.5%.

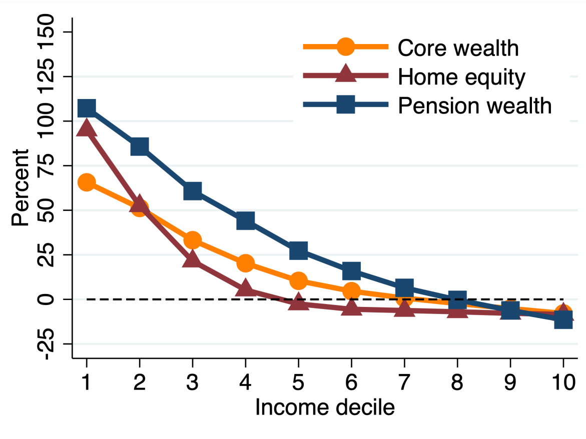

Figure 11: Change in wealth-to-income ratios in the homogeneous access experiment.

Notes: The figure shows the percentage change in wealth-to-income ratios for wealth, home equity, and pension wealth in the homogeneous access experiment relative to the respective ratio in the baseline model. The homogeneous access experiment sets mortgage approval rates and offer rates for employer-sponsored retirement accounts to their economy-wide averages for all income levels (Figure 9). Wealth defined in eq. (1) is the sum of home equity, business wealth, and retirement accounts.

Several important observations emerge from the figure. First, there is a pronounced surge of 34% in home equity within the bottom half of the income distribution. This increase reflects the higher participation rates, as now more lower-income households buy homes, albeit they buy smaller ones on average. The total stock of home equity for this group increases, suggesting that the extensive margin effect (higher participation) dominates quantitatively over the intensive margin (smaller house size). Conversely, there is a modest decrease of 7% in home equity for households in higher income brackets. Since participation rates remain virtually unchanged at the top, the observed decline results solely from these households purchasing smaller homes. These findings highlight the differential impact of homogeneous frictional access on home equity accumulation across the income distribution.

Second, pension wealth rises for all households except those in the top two income deciles. This increase is due to higher participation rates at an approximately constant intensive margin. Combining the effects on home equity and pension wealth, we conclude that wealth rises for households in the bottom half of the income distribution while it drops for households in the top two income deciles. For households in the 6^{th} to 8^{th} income deciles, the increase in pension wealth dominates, leading to an overall increase in wealth.

These results show that a model with homogeneous access to asset markets overstates wealth accumulation opportunities for low-income households while understating them for high-income households. More specifically, wealth-to-income ratios for wealth in the bottom half of the income distribution are 32% higher with homogeneous market access, while they are lower by 3% at the top half relative to the baseline model with heterogeneous access.

While Figure 11 displayed percentage changes in wealth and its components, Figure 12 shows levels of wealth-to-income ratios in the baseline model and the homogeneous access experiment. The model with homogeneous access to asset markets results in a flatter profile of wealth-to-income ratios along the income distribution and thus to a substantial reduction of

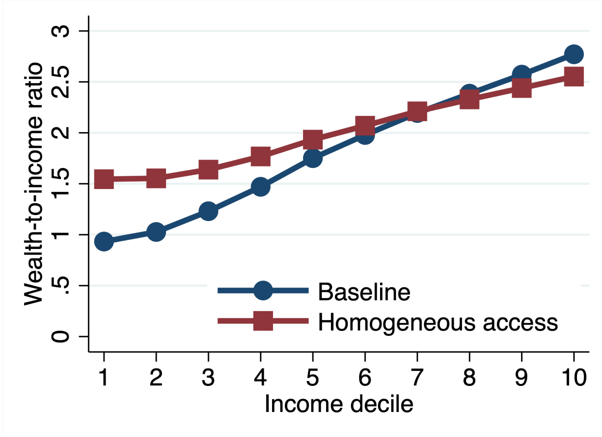

Figure 12: Wealth-to-income ratios for wealth in the baseline model and homogeneous access experiment.

Notes: The homogeneous access experiment sets mortgage approval rates and offer rates for employer-sponsored retirement accounts to their economy-wide averages for all income levels (Figure 9).

between-income-group wealth inequality.

Heterogeneity in saving rates by income is the important driver of the between-income-group wealth inequality. To explore this point, Table 5 reports percentage point changes in saving rates by income quintiles. As a reference, we repeat the saving rates from the baseline model from Table 4. Strikingly, we see that saving rates in the bottom quintile, which were virtually zero in the baseline model, increase by more than 6pp. By contrast, we find a large negative change in the top quintile where the saving rate drops by more than 3pp. The lower saving rate for these households results from worse trading opportunities to upgrade their homes, which reduces saving rates in home equity at the intensive margin. Finally, saving rates in the 2^{nd} and 3^{rd} quintiles hardly change, consistent with the small changes in wealth-to-income ratios.

Table 5: Saving rates by income in the homogeneous access experiment compared to the baseline model.

	Q1	Q2	Q3	Q4	Q5
Heterogeneous access (baseline)	0.001	0.098	0.105	0.137	0.161
Homogeneous access, pp.	+6.1	+0.1	-0.3	-3.7	-3.3

Notes: The table shows the estimation results of median regressions of saving rate on current income as in Dynan et al. (2004) (their Table 3, second column). The authors define the saving rate variable as the change in real net worth in the SCF. For brevity, we display only the estimates for the income quintile dummies, omitting the results for other regressors. The results are displayed in percentage point deviations from saving rates in the baseline model.

5.2.2 Frictionless access to the mortgage market

In the second experiment, we study the portfolio changes if we unilaterally improve market access in only one asset market. This experiment provides important insights to understand the results of the policy analysis below where only the frictions in the pension market are changed.

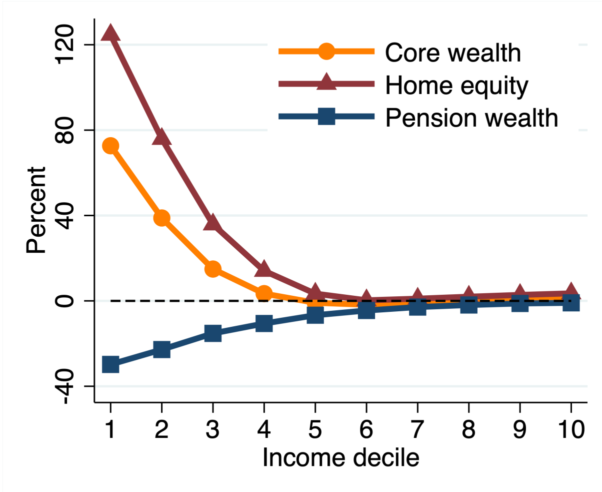

Figure 13: Change in wealth-to-income ratios in the experiment with frictionless access to the mortgage market.

Notes: The figure shows the percentage change in wealth-to-income ratios for wealth, home equity, and pension wealth in the homogeneous access experiment relative to the respective ratio in the baseline model. The experiment sets the mortgage approval rates to one for all income levels. Wealth defined in eq. (1) is the sum of home equity, business wealth, and retirement accounts.

Here, we consider the mortgage market where we set the probability of being approved for a mortgage, $\pi_H(y)$, to one for all income levels so that all households who apply for a mortgage will also be approved. We keep access to the pension market and all other parameters at their calibrated values. In this case, it becomes easier for all households to accumulate wealth in the housing market, and hence, in relative terms, the frictions in the pension market increase. From a macroeconomic perspective, such changes could, for example, result from financial innovations in mortgage markets.

Figure 13 shows the percentage change in wealth-to-income ratios across income deciles in this experiment. Comparing it to Figure 11, where we removed only the *heterogeneity* of market access in the mortgage and the pension market, we now find strong asset substitution from pension wealth to home equity, especially at the bottom of the income distribution, where access to asset markets is lowest. In the bottom decile, households reduce their pension wealth-to-income ratio by 30% while simultaneously increasing their home equity. On the net, they increase their wealth by almost 80%. In the upper half of the distribution, the gross changes are much smaller and asset substitution nets out, leading to roughly unchanged wealth, which shows that high-income households have sufficient saving opportunities. To summarize, wealth accumulation in the upper half of the distribution is, on average, not affected, but the shift in the asset market access leads to asset substitution towards the asset with lower frictions for income-poor households, thereby changing their composition of wealth.

6 Policy experiment

There has been a longstanding debate in the macroeconomic literature on designing fully-funded pension systems that allow households to accumulate sufficient wealth for retirement.

This literature has predominantly focused on policy measures targeted at the demand side of the private pension market, such as myopic behavior in financial planning, lack of financial literacy, liquidity constraints and uninsurable idiosyncratic risks (Hubbard et al., 1995; Conesa and Krueger, 1999; Dynan et al., 2004; Nishiyama and Smetters, 2007; Ábrahám et al., 2024). By contrast, limited access to the market for retirement saving as a supply-side factor has received no attention in the academic literature, despite ongoing public debates.[29] Our framework provides a novel quantitative laboratory to study the role of limited access to the retirement market for household wealth building.

We conduct a policy experiment that implements the key provisions of the recently proposed bipartisan federal *Retirement Savings for Americans Act*. The Act proposes that workers without access to an employer-sponsored retirement plan become eligible for a government pension plan and includes provisions for subsidizing pensions for low-income workers. To account for these provisions in the model, we assume that workers without access to an employer plan or those workers who have been offered such a plan but reject it now can enroll in a government plan.[30] Similar to employer-sponsored plans, we model the government plan as a predefined financial product. The worker's contribution rate is fixed at $\chi_G = 3\%$. The worker's labor income state y determines the copay rate σ_G and the base rate $\bar{\chi}_G$, both financed by the government. For low-income workers who earn below half of the median income y_{med}, the government pays a base contribution rate of $\bar{\chi}_G = 1\%$ and matches the worker's contributions at rate $\sigma_G = 100\%$. It is important to note that our reform requires workers to actively contribute to the account to qualify for the government contribution. For workers with above-median income ($y > y_{med}$), the copay rate and the base rate are both set to zero. For agents with intermediate incomes $y \in [y_{med}/2, y_{med}]$, the base contribution rate $\bar{\chi}_G$ and the copay rate σ_G phase out linearly to zero. As for employer plans, the workers' contributions to government plans are tax-deductible.

Figure 14 illustrates the profile of contribution rates by income in the policy experiment. The dashed line corresponds to the government plan. The maximum contribution rate including copay for low-income workers is $\bar{\chi}_G + \chi_G(1+\sigma_G) = 7\%$. It phases out linearly from 7% to 3% for intermediate incomes and then remains constant at a contribution rate of 3% that is only paid by the worker. The solid line corresponds to the employer-sponsored plan, whose conditions are the same as in the baseline model (the contribution rate is $\chi_W = 5\%$ and the employer copay rate is $\sigma_P = 80\%$), such that the cumulative contribution rate of $\chi_W(1 + \sigma_P)$ is constant at 9% along the income distribution.

The top row of Figure 15 shows the pick-up rates for employer-sponsored and government plans in the baseline model and the experiment. The pick-up rate for employer plans in Panel (a) is defined as the percentage of workers who have an active employer-sponsored plan among all workers eligible for such a plan. The pick-up rate for government plans in Panel (b) is the percentage of workers with an active government plan among all eligible workers, which includes all workers as government plans are freely accessible. It is by construction zero in the baseline economy where this option is not available.

[29] See, for example, a broad discussion of limited access for low-income households to the private pension market by the Economic Innovation Group (EIG) at `https://inclusivewealth.eig.org/`.

[30] Entrepreneurs are not eligible for government plans, and agents cannot enroll in both plans simultaneously.

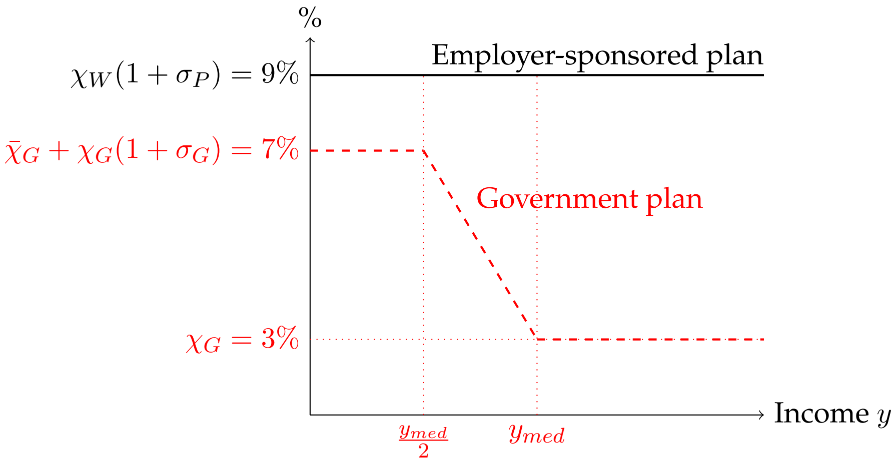

Figure 14: Cumulative contribution rates to the pension accounts in the policy experiment.

Notes: The figure shows the profile of contribution rates by income in the policy experiment. The solid line corresponds to the employer-sponsored plan, whose conditions are the same as in the baseline model (the worker's contribution is $\chi_W = 5.0\%$ and the employer copay is $\sigma_P = 80\%$). The dashed line corresponds to the government plan, where the worker's contribution is fixed at $\chi_G = 3\%$. The worker's labor income y determines the government copay rate σ_G and the base rate $\bar{\chi}_G$. For incomes below half of the median income $y_{med}/2$, $\bar{\chi}_G = 1\%$ and $\sigma_G = 1$. Both phase out linearly to 0% at the median income y_{med}.

Figure 15(a) shows that the policy reduces the demand for employer-sponsored plans for all households up to the 5^{th} income decile but leaves the demand of high-income households unaffected. At the same time, the policy boosts the demand for government plans along the entire income distribution (Panel (b)). To understand this result, note that the policy has two major effects on the pension market. The first effect is to offer access to the pension market for all workers, to which many workers did not have access in the baseline. Following the introduction of the policy, workers without prior access start demanding the now accessible government plan.

The second effect of the policy is to introduce a new saving product whose (pre-specified) conditions depend on the worker's income. Low-to-medium-income workers who were enrolled in employer plans in the baseline model now choose to shift to the newly available government plan due to the lower contribution rate. This crowding-out effect accounts for the reduction in pick-up rates for employer plans and an increase in demand for government plans in the lower half of the income distribution. By contrast, high-income workers have strong incentives to save, so they prefer to stay with employer plans as these plans allow for a 9% saving rate compared to the 3% for high-income workers in the government plan.

Consistent with the previous figure, Figure 15(c) confirms that the policy reform significantly enhances participation in the pension market (+17pp). The most substantial increase is observed in the lower half of the income distribution, with the exception of the bottom income decile, where participation remains virtually unchanged. The reason is that the lowest-income households prefer not to save for retirement in the presence of social security (Hubbard et al., 1995). Households in the top two income deciles also show no change in participation because they were little affected by missing access to start with.

Panel (d) of Figure 15 shows the intensive margin for pension wealth. We see a reduction in

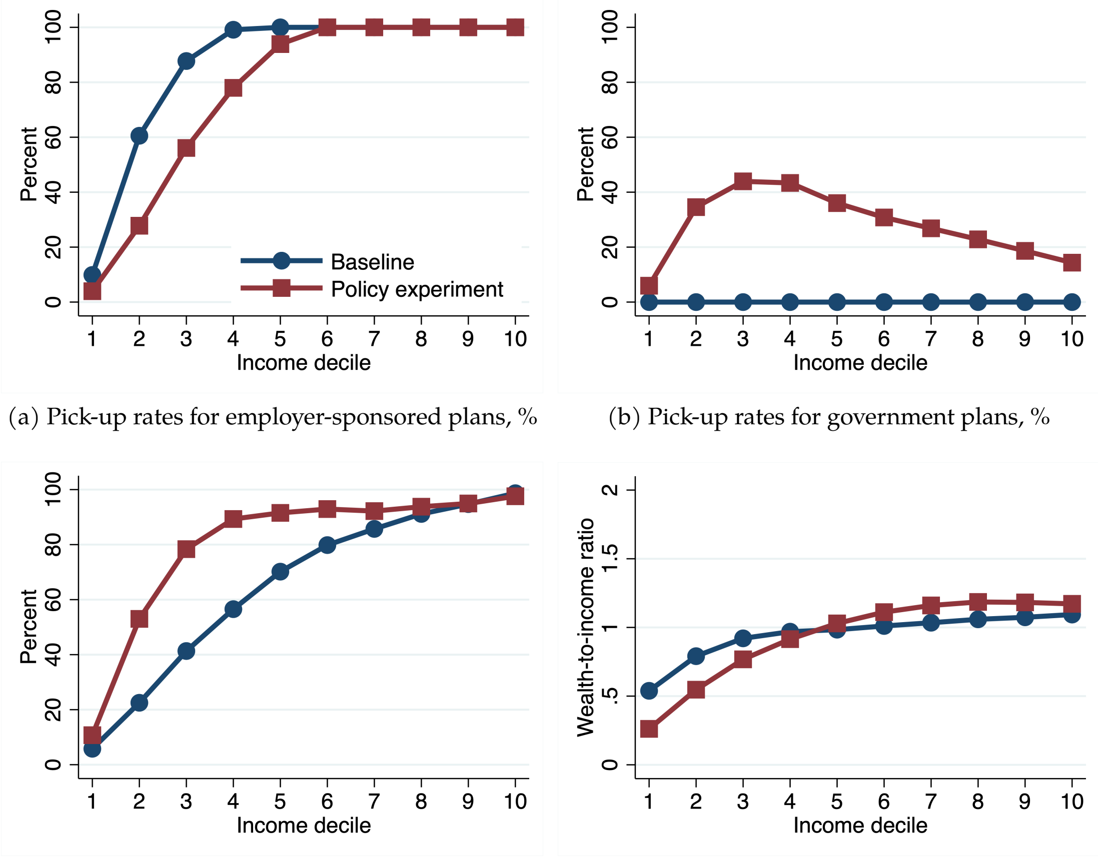

(a) Pick-up rates for employer-sponsored plans, %

(b) Pick-up rates for government plans, %

(c) Pension wealth: Extensive margin, %

(d) Pension wealth: Intensive margin

Figure 15: Pension market in the baseline model and policy experiment.

Notes: In Panel (a), the pick-up rate is defined as the percentage of workers who hold an active employer-sponsored account among all workers eligible for an employer-sponsored plan. In Panel (b), the pick-up rate is defined as the percentage of workers who hold an active government plan among all workers in the economy.

the intensive margin for households with below-median incomes and an increase for higher-income households. The reason is that many households with below-median incomes enroll in government plans and as these plans have lower contribution rates, pension wealth at the intensive margin declines on average. For high-income households, we find an increase in pension wealth at the intensive margin resulting from a life-cycle effect. When young, high-income households start in the lower parts of the income distribution and the newly available government-provided plan offers them earlier entry to the pension market so that when their income is high, they have already accumulated more pension wealth.

Figure 16 documents the effect of the policy on wealth accumulation by showing the percentage change in wealth-to-income ratios by income for total wealth, home equity, and pension wealth under the policy compared to the baseline model. Three important messages emerge from the figure. First, the policy has a pronounced impact on the pension wealth accumulation of households with below-median incomes. Looking at the entire bottom half of the income distribution, pension wealth increases, on average, by 69%. Pension wealth for the bottom income decile increases by 80%. The substantial increase in retirement wealth for low-income agents suggests that improving access to this market has a strong effect on their wealth accu-

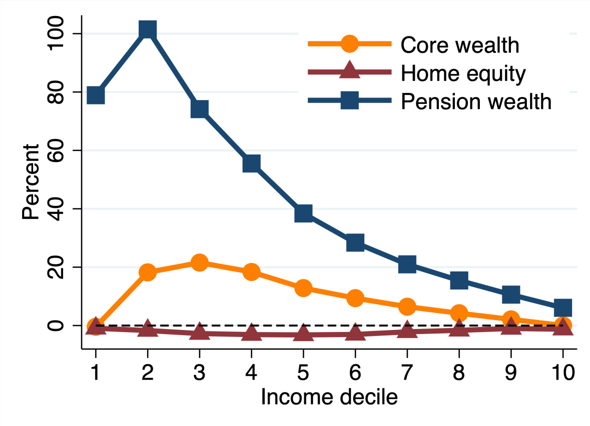

Figure 16: Change in wealth-to-income ratios in the policy experiment compared to the baseline model, %.

Notes: The figure shows the percentage change in wealth-to-income ratios for wealth, home equity, and pension wealth in the policy experiment relative to the respective ratio in the baseline model. Wealth is the sum of home equity, business wealth, and retirement accounts.

mulation. Yet, this percentage increase has to be interpreted with care as the *level* of retirement wealth for some low-income households has been very low before the reform.

Second, we observe a weak substitution effect from the housing market to the pension market after access to the pension market has improved. The moderate decrease in housing wealth (on average, 2%) indicates that there is asset substitution, mainly occurring in the lower part of the income distribution, consistent with the results from Section 5.2.2. In the upper part of the distribution, we do not observe any asset substitution. For most of these households, government-provided plans are dominated by employer-sponsored plans, so the policy has no effect. Overall, the policy leads to an increase in wealth by 9%, with the most pronounced increase concentrated in the lower half of the income distribution.

Looking at the costs of implementing the reform, we find them to be fairly low, amounting to less than 1% of the economy-wide average pre-tax income. The largest government expense should fall on households with incomes below $y_{med}/2$ who qualify for both a full base rate $\bar{\chi}_G = 1\%$ and a full copay rate of $\sigma_G = 100\%$, but Figure 15(b) shows that pickup rates for government plans in the bottom income decile are less than 10%. Recall that the reform requires workers to actively contribute to the account to qualify for the government contribution, and the mandatory 3% contribution rate deters low-income agents from participating.

How does the policy affect the saving rates of low-income households? The comparative statics analysis from Section 5.2.1 shows that saving rates of low-income households increase a lot when market access improves. Table 6 confirms this finding in the context of the policy experiment. Throughout the income distribution, saving rates increase after the reform but changes are very asymmetric. In the bottom income quintile, saving rates rise from roughly zero percent in the baseline model to almost 7pp, while in all other quintiles the increase is between 13bp and 64bp. The reason is that lower contribution rates in the government plan are particularly attractive for low-income households with weaker life-cycle savings motives.

In summary, the policy experiment demonstrates that implementing a policy designed to al-

Table 6: Saving rates by income in the policy experiment compared to the baseline model, pp.

	Q1	Q2	Q3	Q4	Q5
Policy Experiment	+6.88	+0.37	+0.24	+0.64	+0.13

Notes: The table shows the estimation results of median regressions of saving rate on current income as in Dynan et al. (2004) (their Table 3, second column). The authors define the saving rate variable as the change in real net worth in the SCF. For brevity, we display only the estimates for the income quintile dummies, omitting the results for other regressors. The results are displayed in percentage point deviations from saving rates in the baseline model.

leviate limited access to the pension market significantly enhances broad-based wealth accumulation. By providing an alternative savings product with subsidies and lower contribution rates, the policy enables low-income households to build retirement savings in a government-provided market. However, the bottom 10% of households remain largely unaffected by the policy, not because they do not have the opportunity to save but because they do not want to save given the existing government programs in place.

6.1 External validity

The model suggests a large effect in response to a policy that improves access to the market for retirement plans. To put these quantitative results into perspective, we compare them to the evolution of employer-sponsored retirement plans during the 1980s in the United States. This period marked a rapid growth in employer-sponsored retirement accounts, largely driven by the adoption of the 1981 *Economic Recovery Tax Act* (ERTA). Prior to the ERTA, retirement savings were limited to defined-benefit plans and Individual Retirement Accounts (IRAs). Moreover, strict eligibility requirements and limited flexibility in choosing plans restricted broader participation. ERTA was an instrumental policy that alleviated supply-side frictions in the pension market. First, it introduced 401(k) plans, which allowed employees to defer a portion of their income into retirement accounts on a tax-free basis, reducing their immediate tax burden. Second, the Act increased tax deductions available to employers for contributions, thereby making it financially more attractive for employers to offer such plans. Lastly, ERTA reduced administrative costs associated with setting up and managing pension plans. These provisions improved access to the pension market, very similar to the change in our policy experiment, so we take this historical policy as an opportunity to validate our model's predictions regarding the macroeconomic effects of such changes.

We rely on data from the SCF+ (Kuhn et al., 2020) and adhere to the same variable definitions and sample selection criteria as described in Section 2. In the SCF+ data, we verify that no wealth was reported for employer-sponsored plans before 1983.[31] Figure 17 depicts the dynamics of participation rates in the pension market in Panel (a) and the intensive margin for pension wealth in Panel (b) from 1983 to 1998 as blue squares across five survey waves. Following the ERTA adoption, participation in defined-contribution plans surged from 0% to 33%, which is a sizable increase compared to the 9pp rise in the average extensive margin in our policy ex-

[31] The SCF+ follows the SCF convention to report only marketable wealth, which excludes defined-benefit plans as they are not marketable. Note further that the SCF+ omits the 1986 survey.

periment. The participation rates continued to rise thereafter, almost doubling to 64% until 1998. Meanwhile, the intensive margin for pension wealth also showed rapid growth, from 0.3 in 1983 to 0.7 in 1998. While participation rates for home equity remained virtually unchanged at 70% during this period (Panel (c)), the intensive margin for home equity declined by 28% from 1.8 to 1.3 (Panel (d)). Both of these trends are consistent with the model-predicted asset substitution from home equity to pension wealth when access to the pension market improves.

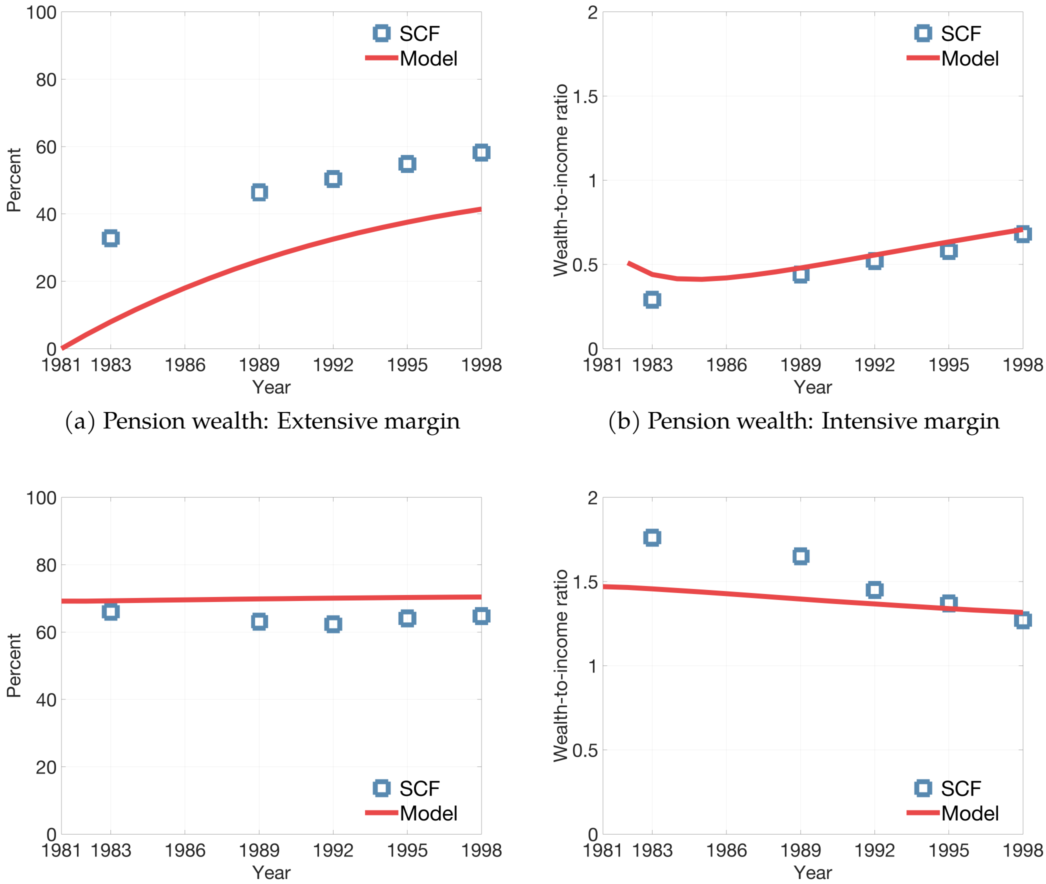

(a) Pension wealth: Extensive margin

(b) Pension wealth: Intensive margin

(c) Home equity: Extensive margin

(d) Home equity: Intensive margin

Figure 17: Average extensive and intensive margins for retirement wealth and home equity in the simulation and data.

Notes: The figure shows the evolution of pension wealth and home equity in the U.S. following the adoption of the 1981 Economic Recovery Tax Act. The empirical moments are taken from the SCF+ (Kuhn et al., 2020) and adhere to the same variable definitions and sample selection criteria as described in Section 2.

To assess whether our calibrated framework replicates the described trends in the historical data, we conduct the following experiment. First, we solve the baseline model assuming that access to the private pension market is entirely shut down by setting the probability of receiving an employer-sponsored retirement account, $\pi_P(y)$, to zero for all income levels (government plans are not offered in this experiment). We maintain all other parameters at their calibrated values. This situation represents the U.S. economy in 1981, prior to the adoption of the ERTA. Then, we introduce an unanticipated and permanent reform, which sets the probabilities of accessing employer plans to the levels of the calibrated model. We simulate a panel of households applying the decision rules obtained from the solution of the baseline model and analyze the behavioral response of agents to changes in the supply of pension plans.

Figure 17 (solid line) shows the simulated paths for the extensive and intensive margins for pension wealth (Panels (a) and (b)) and home equity (Panels (c) and (d)). The model successfully captures the increasing participation rates and wealth-to-income ratios for pension wealth in the data. For home equity, the model predicts a stable extensive margin and a decreasing wealth-to-income ratio, closely matching the empirical trends. Quantitatively, the model explains 58% of the total growth in the extensive margin for pension wealth during 1983–1998 and replicates remarkably closely the trajectory for the wealth-to-income ratios of pension wealth. Finally, the model accounts for 44% of the total drop in the intensive margin for home equity. We take this as strong additional evidence to support our framework and its mechanism.

7 Conclusions

This paper provides a new answer to the long-standing question of the causes of differences in saving rates by income. It thereby provides two new insights for the distribution of wealth. Using the SCF data and focusing on the three main asset classes in the portfolio of U.S. households (home equity, business equity, and retirement accounts), we find that differences in the extensive margin along the income distribution are key in accounting for the unequal distribution of wealth. Second, we provide extensive evidence that this heterogeneity is closely related to variation in asset market access.

We develop a new model of life-cycle wealth accumulation that focuses on asset market participation (extensive margin) rather than intensive margin variation of savings of asset market participants. Households in the model execute pre-specified financial contracts and have consumption dynamics like wealthy hand-to-mouth households who are typically not on their Euler equation. We demonstrate that this model accounts for the empirical evidence regarding the joint distribution of income and wealth. Increasing saving rates by income in the model do not result from higher shares of saved income of asset market participants but from differences in asset market participation, with higher-income households participating in more markets because of better market access. Our counterfactual analyses show that low-income households are typically constrained in their wealth accumulation due to missing asset market access. After improving market access, they start accumulating more wealth and wealth differences by income decline. The conclusion is that asset markets do not offer a level playing field to all households as opportunities to build wealth vary systematically with income.

Our policy experiment analyzes the consequences of the *Retirement Savings for Americans Act* on wealth accumulation. The Act aims to improve market access for low-income households in the pension market. We find that the targeted low-income households accumulate substantially more retirement wealth after their market access improves. We validate the model prediction from the reform to the similar historical episode of the proliferation of employer-sponsored retirement accounts in the 1980s. The policy experiment, therefore, highlights the important role of income-dependent market access for wealth building of American households.

References

Ábrahám, A., P. Brendler, and E. Carceles-Poveda (2024). Optimal income redistribution. *Working Paper*.

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics 109*(3), 659–684.

Alan, S., K. Atalay, and T. F. Crossley (2015). Do the rich save more? evidence from Canada. *Review of Income and Wealth 61*(4), 739–758.

Bayer, C., B. Born, and R. Luetticke (2024). Shocks, frictions, and inequality in U.S. business cycles. *American Economic Review 114*(5), 1211–1247.

Bayer, C., R. Lütticke, L. Pham-Dao, and V. Tjaden (2019). Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk. *Econometrica 87*(1), 255–290.

Benhabib, J., A. Bisin, and M. Luo (2019). Wealth distribution and social mobility in the U.S.: A quantitative approach. *The American Economic Review 109*(5), 1623–1647.

Bewley, T. (1986). Stationary monetary equilibrium with a continuum of independently fluctuating consumers. *Contributions to mathematical economics in honor of Gérard Debreu 79*.

Bhutta, N., J. Bricker, A. C. Chang, L. J. Dettling, S. Goodman, J. W. Hsu, K. B. Moore, S. Reber, A. Henriques Volz, and R. Windle (2020). Changes in US family finances from 2016 to 2019: Evidence from the Survey of Consumer Finances. *Federal Reserve Bulletin 106*(5).

Bilbiie, F. O. (2024). Monetary policy and heterogeneity: An analytical framework. *Review of Economic Studies*.

Bloom, N., C. I. Jones, J. Van Reenen, and M. Webb (2020). Are ideas getting harder to find? *American Economic Review 110*(4), 1104–44.

Brendler, P. (2023). Rising earnings inequality and optimal income tax and social security policies. *Journal of Monetary Economics 134*, 35–52.

Cagetti, M. (2003). Wealth accumulation over the life cycle and precautionary savings. *Journal of Business & Economic Statistics 21*(3), 339–353.

Cagetti, M. and M. De Nardi (2006). Entrepreneurship, frictions, and wealth. *Journal of Political Economy 114*(5), 835–870.

Cagetti, M. and M. De Nardi (2008). Wealth inequality: Data and models. *Macroeconomic Dynamics 12*(S2), 285–313.

Castañeda, A., J. Díaz-Giménez, and J. Ríos-Rull (2003). Accounting for the U.S. earnings and wealth inequality. *Journal of Political Economy 111*(4), 818–857.

Conesa, J. C. and D. Krueger (1999). Social Security Reform with Heterogeneous Agents. *Review of Economic Dynamics 2*(4), 757–795.

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

De Nardi, M. and G. Fella (2017). Saving and wealth inequality. *Review of Economic Dynamics 26*, 280–300.

De Nardi, M., E. French, and J. B. Jones (2010). Why do the elderly save? The role of medical expenses. *Journal of Political Economy 118*(1), 39–75.

Dynan, K. E., J. Skinner, and S. P. Zeldes (2004). Do the rich save more? *Journal of Political Economy 112*(2), 397–444.

Fagereng, A., M. B. Holm, B. Moll, and G. Natvik (2021). Saving behavior across the wealth distribution: The importance of capital gains. *Working Paper*.

Gomes, F., M. Haliassos, and T. Ramadorai (2021). Household finance. *Journal of Economic Literature 59*(3), 919–1000.

Gourinchas, P.-O. and J. A. Parker (2002). Consumption over the life cycle. *Econometrica 70*(1), 47–89.

Guiso, L., M. Haliassos, and T. Jappelli (2002). *Household portfolios*. MIT press.

Haliassos, M. and C. C. Bertaut (1995). Why do so few hold stocks? *Economic Journal 105*(432), 1110–1129.

Heathcote, J., K. Storesletten, and G. L. Violante (2020). Presidential address 2019: How should tax progressivity respond to rising income inequality? *Journal of the European Economic Association 18*(6), 2715–2754.

Hubbard, R. G., J. Skinner, and S. P. Zeldes (1995). Precautionary saving and social insurance. *Journal of Political Economy 103*(2), 360–399.

Hubmer, J., P. Krusell, and A. A. Smith. (2021). Sources of US wealth inequality: Past, present, and future. *NBER Macroeconomics Annual 35*, 391–455.

Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control 17*(5-6), 953–969.

Jung, P. and M. Kuhn (2019). Earnings losses and labor mobility over the life cycle. *Journal of the European Economic Association 17*(3), 678–724.

Kaplan, G., B. Moll, and G. L. Violante (2018). Monetary policy according to HANK. *American Economic Review 108*(3), 697–743.

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

Kaplan, G., G. L. Violante, and J. Weidner (2014). The wealthy hand-to-mouth. *V Brookings Papers on Economic Activity, Spring*.

Krueger, D., K. Mitman, and F. Perri (2016). Macroeconomics and household heterogeneity. In *Handbook of macroeconomics*, Volume 2, pp. 843–921. Elsevier.

Krusell, P. and A. A. Smith, Jr (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy 106*(5), 867–896.

Kuhn, M. and J.-V. Rıos-Rull (2016). 2013 update on the U.S. earnings, income, and wealth distributional facts: A view from macroeconomics. *Federal Reserve Bank of Minneapolis Quarterly Review 37*(1).

Kuhn, M., M. Schularick, and U. I. Steins (2020). Income and wealth inequality in America, 1949–2016. *Journal of Political Economy 128*(9), 3469–3519.

Low, H., C. Meghir, and L. Pistaferri (2010). Wage risk and employment risk over the life cycle. *American Economic Review 100*(4), 1432–1467.

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

Mankiw, N. G. and S. P. Zeldes (1991). The consumption of stockholders and nonstockholders. *Journal of Financial Economics 29*(1), 97–112.

Nishiyama, S. and K. Smetters (2007). Does Social Security privatization produce efficiency gains? *The Quarterly Journal of Economics 122*(4), 1677–1719.

Pyatt, G. (1976). On the interpretation and disaggregation of Gini coefficients. *The Economic Journal 86*(342), 243–255.

Quadrini, V. (2000). Entrepreneurship, saving, and social mobility. *Review of Economic Dynamics 3*(1), 1–40.

Saez, E. and G. Zucman (2014). Wealth inequality in the United States since 1913: Evidence from capitalized income tax data. NBER Working Papers 20625, National Bureau of Economic Research, Inc.

Straub, L. (2019). Consumption, savings, and the distribution of permanent income. *Unpublished manuscript, Harvard University 17*.

Online Appendix

A Empirical analysis

A.1 Variable definitions and sample selection criteria

We use the publicly available data from the Survey of Consumer Finances (SCF). Our main analysis spans the years 2010 to 2019, but we also conduct a robustness check pooling together all publicly available datasets since 1989 below. For the results on the relationship between income and wealth, we rely on the Summary Extract file. When analyzing access to the pension and home equity markets as well as the role of prices for asset-market participation, we resort to the full dataset with all variables.

A small set of sample selection restrictions is applied. First, we focus on household heads aged 25–65. Second, we restrict the sample to those households whose head works in dependent employment or is self-employed (variable *occat1*). Third, we remove households who report running a business (*actbus*) but not actively managing it (*nonactbus*), where *actbus* is the net value of businesses in which the household has an active interest and *nonactbus* is the net value of businesses in which the household does not have an active interest.

We define income as the sum of wage income (*wageinc*), income from sole proprietorship, a farm, or other businesses or investments (*bussefarminc*), income from non-taxable investments (*intdivinc*), income from government transfers (*transfothinc*), and income from Social Security, other pensions, annuities or other disability or retirement programs (*ssretinc*). We follow the SCF definitions of total wealth (*networth*), pension wealth (*retqliq*), and business wealth (*bus*). We define home equity as the value of both the primary residence (*houses*) as well as the value of other residential real estate (*oresre*), net of any related debts (*mrthel* and *resdbt*).

We winsorize the top 1% of the wealth-to-income ratios for total wealth and each core wealth component to exclude outliers. Then, we construct wealth-to-income ratios for core wealth from the winsorized data.

When analyzing access to the pension market and the price effects for participation in the pension market, we impose an additional sample selection restriction that the household head works in dependent employment (*x4106* and *x4100*). However, income deciles are defined on a full sample to maintain consistency with the preceding analysis.

A.2 Alternative sample period

Figure 18 confirms that the documented patterns for total wealth and core wealth, shown in Figure 1 in the main text for 2010–2019, also hold when we pool all publicly available SCF datasets since 1989. Similarly, Figure 19 shows that the documented patterns for the extensive and intensive margins (Figure 3) are also robust with respect to the sample period.

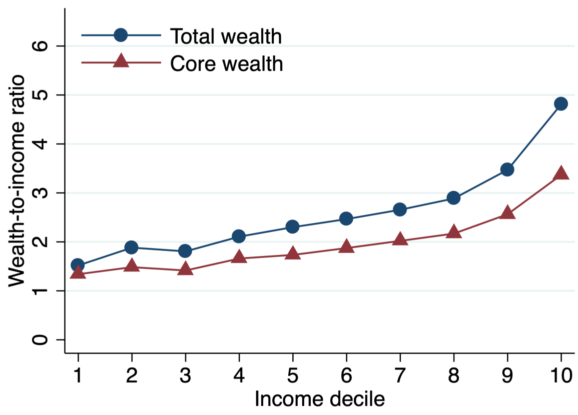

Figure 18: Wealth-to-income ratios for total wealth and core wealth in the SCF (1989–2019).

Notes: The figure is constructed using the full publicly available SCF dataset covering 1989–2010. *Core wealth* is the sum of business wealth, home equity, and retirement accounts. *Total wealth* comprises core wealth plus stocks, bonds, mutual funds, liquid assets, and remaining assets. The value of each asset component is computed net of debts. See Appendix A.1 for the sample selection criteria.

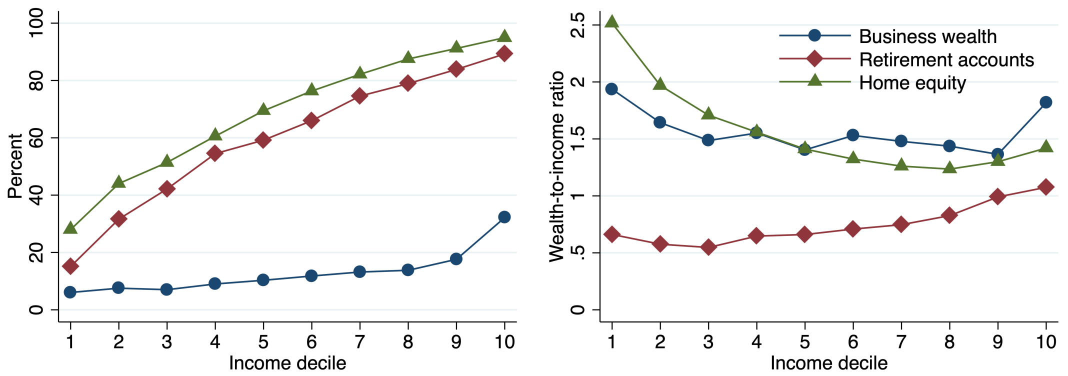

(a) Extensive margin

(b) Intensive margin

Figure 19: Extensive margin and intensive margin variations by income in the SCF (1989–2019).

Notes: The figure is constructed using the full publicly available SCF dataset covering 1989–2010. See Appendix A.1 for the sample selection criteria.

A.3 Non-core wealth

We demonstrate in the main text that core wealth accounts for 70% of total household wealth (Table 1. Below, we report wealth-to-income ratios for non-core wealth by income and over the life cycle.

The left panel of Figure 20 plots wealth-to-income ratios for non-core wealth by income and demonstrates that non-core wealth is substantially lower and hardly varies between the 2^{nd} and 9^{th} income deciles. Only the top decile has substantial holdings of non-core wealth consistent with stock holdings concentrated in that group. The right panel of Figure 20 further dissects the wealth-to-income ratios for non-core wealth presented in the left panel by categorizing households into two groups: those with positive amounts of core wealth and those without any core wealth. We can see that households lacking core wealth also hold virtually no non-core wealth. Put together, our findings indicate that there is no evidence to suggest a substitution effect from core to non-core wealth among households without core wealth.

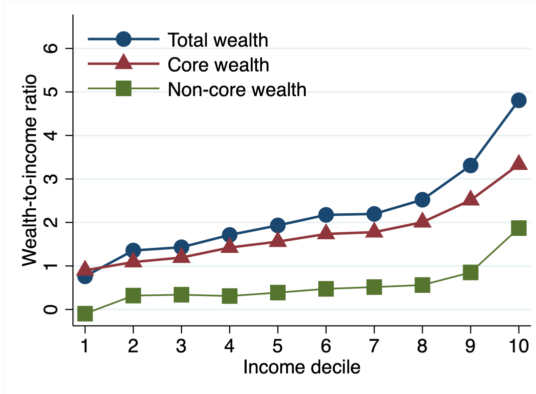

(a) Core and non-core wealth.

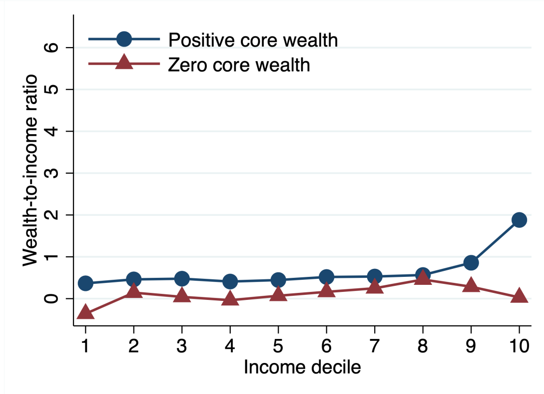

(b) Non-core wealth by core wealth holdings.

Figure 20: Joint distribution of non-core wealth and income in the SCF.

Notes: The left panel shows wealth-to-income ratios for total wealth, core wealth, and non-core wealth by income. The right panel shows wealth-to-income ratios for non-core wealth by income for households with positive holdings of core wealth and for households with zero holdings of core wealth. The figure is constructed using the full publicly available SCF dataset covering 1989–2010. *Core wealth* is the sum of business wealth, home equity, and retirement accounts. *Total wealth* comprises core wealth plus stocks, bonds, mutual funds, liquid assets, and remaining assets. *Non-core wealth* is Total wealth net of Core wealth. The value of each asset component is computed net of debts. See Appendix A.1 for the sample selection criteria.

Figure 21 plots wealth-to-income ratios for total wealth, core wealth, and non-core wealth over the life cycle. The figure shows that roughly 80% of accumulated wealth at the end of the life cycle is core wealth. Over the life cycle, core wealth-to-income ratios increase from about 0.5 to 4, whereas non-core wealth increases from zero to one. Hence, the largest share of life-cycle wealth accumulation is accounted for by core wealth. This finding is consistent with our model in Section 3 which focuses on life-cycle wealth accumulation.

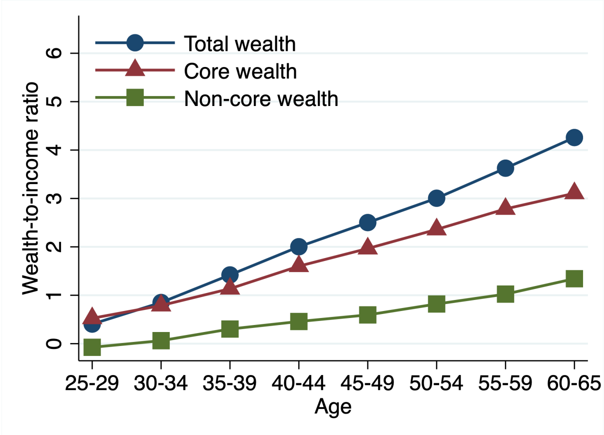

Figure 21: Non-core wealth over the life cycle in the SCF.

Notes: The figure is constructed using the full publicly available SCF dataset covering 1989–2010. *Core wealth* is the sum of business wealth, home equity, and retirement accounts. *Total wealth* comprises core wealth plus stocks, bonds, mutual funds, liquid assets, and remaining assets. *Non-core wealth* is Total wealth net of Core wealth. The value of each asset component is computed net of debts. See Appendix A.1 for the sample selection criteria.

A.4 Sensitivity analysis

In this subsection, we demonstrate that the relationship between income and wealth, documented in Figure 1, remains economically and statistically significant, even when accounting for a broad set of household characteristics. Additionally, we perform sensitivity analyses on the income dependence of mortgage approval rates and offer rates for employer-sponsored retirement plans to show that conditional on various household characteristics, income deciles exert a statistically significant and positive effect on both variables.

Section A.4.1 presents the regression results based on our primary SCF dataset. To corroborate our results on income dependence of market access, Section A.4.2 uses *Home Mortgage Disclosure Act* (HMDA) data to study the relationship between mortgage approval rates and income. Section A.4.3 relies on data from the *National Compensation Survey* (NCS) to corroborate the evidence for employer-provided pension plans. Finally, Section A.4.4 addresses some concerns regarding the role of prices for the extensive margin for pension wealth and home equity.

A.4.1 Survey of Consumer Finances

Variable definitions

Below we define the main variables used in our analysis (see the SCF definitions for other variables):

- *Offered* – An indicator variable that equals one if a household head has access to an employer-provided pension plan and zero otherwise. A household head is defined to have access if the following two conditions hold: their employer offers such a pension plan (*X4136*) and they are eligible to enroll in the plan (*X4137*).

- *Applied* – An indicator variable that equals one if at least one of the following three conditions holds: the household applied for a mortgage or a home-based loan (*X435*), or requested to refinance a mortgage (*X436*), or did not apply for any credit because it did not expect to get approved (*X441*). The indicator variable equals zero when the household responds *No* to all three questions.

- *Rejected* – An indicator variable that equals one if variable *Applied* equals one and at least one of the following two conditions is fulfilled: the request for credit is turned down or the household received less credit than what was applied for (*X407*), or the household did not apply for any credit because it did not expect to get approved (*X441*). The indicator variable equals zero when variable *Applied* equals one and no credit application was rejected in the last twelve months (*X407*).

- *Approved* – An indicator variable that equals one if the variable *Rejected* equals zero and zero if *Rejected* equals one.

- *Income Decile Dummies* – A set of indicator variables that indicate the household head's income decile. The first income decile constitutes the reference group for the probit regressions. In all figures that rely on the SCF data, the boundaries of the income deciles are kept the same throughout the paper to preserve consistency.

- *Age* – A variable between 25 and 65 that indicates the household head's age. When included as a control variable, we always use both age and age squared.

- *College* – An indicator variable that equals one if the household head holds at least some college or associate degree and zero otherwise. We use the categorical variable *edcl* to determine the household head's educational level.

Results

Table 7 presents the OLS regression results of wealth-to-income ratios on a broad set of household characteristics. The table confirms that the rising pattern of wealth-to-income ratios for total and core wealth documented in Figure 1 in the main text are not fully accounted for by the household life-cycle saving motives and that the household income has a statistically significant effect on the wealth-to-income ratios, conditional on other observable characteristics.

Next, we perform sensitivity analyses on the income dependence of the extensive margin and market access. Table 8 shows the results of the probit regressions of the extensive margin in each asset component of core wealth on income deciles and a broad set of control variables. Table 9 presents the results of the probit regressions of mortgage approval rates, while Table 10 displays the results of the probit regressions of access to employer-provided pension plans. The tables confirm that conditional on various household characteristics, income deciles exert a statistically significant and positive effect on the extensive margin and access rates.

Table 7: OLS regressions of wealth-to-income ratios (SCF).

	(1) Total	(2) Total	(3) Core	(4) Core
Income decile:				
-2nd	0.596	0.531	0.188	0.131
	(0.183)	(0.171)	(0.095)	(0.084)
-3rd	0.672	0.566	0.294	0.202
	(0.182)	(0.176)	(0.083)	(0.076)
-4th	0.947	0.696	0.517	0.316
	(0.161)	(0.154)	(0.075)	(0.070)
-5th	1.161	0.846	0.650	0.399
	(0.180)	(0.175)	(0.079)	(0.066)
-6th	1.405	0.961	0.827	0.483
	(0.147)	(0.146)	(0.065)	(0.061)
-7th	1.416	0.931	0.859	0.471
	(0.169)	(0.168)	(0.071)	(0.067)
-8th	1.753	1.192	1.097	0.648
	(0.156)	(0.159)	(0.071)	(0.077)
-9th	2.545	1.811	1.607	1.026
	(0.158)	(0.160)	(0.083)	(0.082)
-10th	4.011	2.947	2.394	1.566
	(0.217)	(0.208)	(0.110)	(0.104)
Constant	0.756	-1.434	0.897	-1.046
	(0.151)	(0.473)	(0.060)	(0.200)
Demographic controls		x		x
Year controls		x		x
Observations	62,290	62,290	62,859	62,859

Notes: The regressions are based on the SCF 2010–2019 datasets. Demographic controls include age, age squared, gender, race, education, and marital status. Standard errors are reported in brackets below each coefficient estimate. See Appendix A.1 for the sample selection criteria.

Table 8: Probit regression of participation rates in the pension, home equity, and business equity markets (SCF).

	(1) Pensions	(2) Home equity	(3) Business equity
Income decile:			
-2nd	0.507	0.344	0.070
	(0.048)	(0.061)	(0.060)
-3rd	0.908	0.587	-0.005
	(0.049)	(0.054)	(0.078)
-4th	1.192	0.746	0.081
	(0.058)	(0.056)	(0.061)
-5th	1.294	0.993	0.152
	(0.043)	(0.050)	(0.070)
-6th	1.509	1.126	0.198
	(0.056)	(0.049)	(0.065)
-7th	1.820	1.376	0.204
	(0.050)	(0.058)	(0.067)
-8th	1.963	1.474	0.193
	(0.054)	(0.051)	(0.067)
-9th	2.170	1.713	0.362
	(0.061)	(0.070)	(0.065)
-10th	2.465	1.997	0.694
	(0.061)	(0.063)	(0.061)
Constant	-0.972	-2.284	-3.094
	(0.193)	(0.144)	(0.258)
Demographic controls	x	x	x
Year controls	x	x	x
Observations	66,345	66,454	67,995

Notes: The regressions are based on the SCF 2010–2019 datasets. Demographic controls include age, age squared, gender, race, education, and marital status. Standard errors are reported in brackets below each coefficient estimate. See Appendix A.1 for the sample selection criteria.

Table 9: Probit regression of mortgage approval rates (SCF).

	(1) Approved	(2) Approved
Income decile:		
-2nd	0.756	0.685
	(0.171)	(0.164)
-3rd	0.825	0.787
	(0.179)	(0.179)
-4th	1.142	1.058
	(0.164)	(0.158)
-5th	1.393	1.245
	(0.153)	(0.166)
-6th	1.589	1.430
	(0.180)	(0.183)
-7th	1.996	1.797
	(0.201)	(0.206)
-8th	1.825	1.553
	(0.147)	(0.160)
-9th	2.052	1.791
	(0.177)	(0.193)
-10th	2.196	1.876
	(0.202)	(0.210)
Constant	-0.994	-0.071
	(0.127)	(0.564)
Demographic controls		x
Year controls		x
Observations	7,087	7,087

Notes: The regressions are based on the SCF 2010–2019 datasets. Standard errors are reported in brackets below each coefficient estimate. An applicant is a household who submitted a loan application in the past 12 months or who did not submit an application in the past 12 months because he/she thought it would be denied. An application is treated as approved if the household receives full funding. An application is treated as rejected if the request for credit is turned down, or the household receives less credit than what was applied for, or the household did not apply for any credit because it thought it would be denied. Applications for a mortgage, home-based loan, and refinancing are considered. Demographic controls include age, age squared, gender, race, education, and marital status. See the variable definitions above for more details. See the variable definitions in Appendix A.4.1. See Appendix A.1 for the sample selection criteria.

Table 10: Probit regression of offer rates for employer-sponsored retirement accounts (SCF).

	(1) Offered	(2) Offered
Income decile:		
-2nd	0.512	0.545
	(0.099)	(0.098)
-3rd	0.828	0.893
	(0.093)	(0.094)
-4th	1.041	1.141
	(0.093)	(0.094)
-5th	1.309	1.447
	(0.086)	(0.088)
-6th	1.436	1.598
	(0.089)	(0.093)
-7th	1.501	1.694
	(0.094)	(0.101)
-8th	1.666	1.879
	(0.086)	(0.092)
-9th	1.796	2.000
	(0.082)	(0.086)
-10th	1.871	2.057
	(0.093)	(0.100)
Constant	-0.846	-0.220
	(0.077)	(0.247)
Demographic controls		x
Year controls		x
Observations	37,715	37,715

Notes: The regressions are based on the SCF 2010–2019 datasets. Standard errors are reported in brackets below each coefficient estimate. A household is considered to be offered an employer-sponsored plan if the household head works for an employer who offers a retirement plan and is eligible to participate in the plan. Demographic controls include age, age squared, gender, race, education, and marital status. See the variable definitions in Appendix A.4.1. See Appendix A.1 for the sample selection criteria.

A.4.2 Home Mortgage Disclosure Act

To corroborate our findings regarding the mortgage approval rates by income obtained based on the SCF data (Section 2.4), we use the Home Mortgage Disclosure Act (HMDA) data, which provides a comprehensive data set on the US mortgage market. The HMDA data is published on an annual basis by the Consumer Financial Protection Bureau. Importantly, it includes the data on loan-to-value ratio and debt service-to-income ratio.

Sample selection

The available dataset spans the years 2018 to 2021, covering four survey waves. We restrict the sample to applicants aged 25–64 with positive income. To remove outliers, we exclude observations where the applicant is in the top 5% of the income distribution or the loan amount falls into the top 5% of the loan amount distribution. Additionally, observations with missing data for either sex or race are excluded. We classify a mortgage application as approved if the loan was originated or if the loan application was approved but not accepted. Conversely, we classify an application as rejected if the loan application or preapproval request was denied by the financial institution. Our sample includes loan applications for all purposes, but our results remain robust if we restrict the sample to mortgage applications or to both mortgage applications and mortgage refinancing.

Variable definitions

Some of the main variables exploited in our analysis are as follows:

- *Application Approved* – An indicator variable that equals one either if the loan was originated or if the loan application was approved but not accepted by the applicant. If, instead, the loan application was denied by the financial institution or the preapproval request was denied by the financial institution, the indicator variable equals zero. The indicator variable is constructed from the responses captured in the categorical variable *action_taken*.

- *Age Group* – A set of indicator variables that indicate the applicant's age. Age groups are 25–34, 35–44, 45–54, 55–64.

- *Debt-to-Income Ratio* – An indicator variable that equals one if the applicant's debt-to-income ratio is above 35% and zero otherwise. This indicator variable is constructed from the categorical variable *debt_to_income_ratio*, defined as the applicant's total monthly debt service to total monthly income.

- *Income Decile Dummies* – A set of indicator variables that indicate the applicant's income decile. The first income decile constitutes the reference group for the probit regressions.

- *Loan-to-Value (LTV) Ratio* – A continuous variable that captures the ratio of the total amount of debt secured by the property relative to the value of the property for which the loan is taken out.

- *State* – A set of indicator variables that indicate the applicant's state in which a given loan application was made.

Results

Table 11 presents the probit estimation results of mortgage approval rates, with the last column being our prefererd specification since it controls for loan-to-income ratios and debt-to-income ratios. We confirm that conditional on various household characteristics, income deciles exert a statistically significant and positive effect on mortgage approval rates, consistent with our findings from the SCF data.

A.4.3 National Compensation Survey

We corroborate the results regarding the heterogeneity of offer rates for employer-sponsored retirement plans by income, documented in Figure 6 in the main text, using data from the National Compensation Survey (NCS) on employer-provided pension plans conducted by the Bureau of Labor Statistics. This publicly accessible dataset includes offer and pick-up rates computed for selected percentiles of the wage distribution. For more granularity, we have acquired additional results categorized by income deciles from a special request to the BLS. Our analysis is confined to the period from 2010 to 2020, aligning with our examination of the SCF data. Additionally, we limit our focus to full-time workers employed in the private industry.[32]

Figure 22 presents the results. Two of the findings obtained in the SCF sample (Section 2.4) continue to hold: We observe a large variation in the offer and pick-up rates by income, with the variation in the offer rates playing a more dominant role.

[32]The series we use is "Percent of private industry workers with access to retirement benefits: all retirement plans; for average wage category is within X percent", where X refers to a given percentile, and "Take-up rate for private industry workers with access to and participating in benefit retirement benefits: all retirement plans; for average wage category is within X percent".

Table 11: Probit regression of approval rates for mortgages (HMDA).

	(1)	(2)	(3)	(4)
Income decile:				
-2nd	0.439	0.424	0.384	0.366
	(0.001)	(0.001)	(0.001)	(0.001)
-3rd	0.562	0.551	0.497	0.466
	(0.001)	(0.001)	(0.001)	(0.001)
-4th	0.631	0.622	0.559	0.515
	(0.001)	(0.001)	(0.001)	(0.001)
-5th	0.697	0.693	0.624	0.568
	(0.001)	(0.001)	(0.001)	(0.001)
-6th	0.736	0.741	0.666	0.596
	(0.001)	(0.001)	(0.001)	(0.001)
-7th	0.793	0.794	0.715	0.627
	(0.001)	(0.001)	(0.001)	(0.001)
-8th	0.836	0.846	0.763	0.658
	(0.001)	(0.001)	(0.001)	(0.001)
-9th	0.865	0.886	0.801	0.677
	(0.001)	(0.001)	(0.001)	(0.001)
-10th	0.872	0.904	0.819	0.668
	(0.001)	(0.001)	(0.001)	(0.001)
Loan-to-value ratio			-0.000	-0.000
			(0.000)	(0.000)
Debt-to-income ratio $> 35\%$				-0.310
				(0.001)
Demographic controls		x	x	x
State & Year controls		x	x	x
Observations	37,819,686	37,819,684	35,097,385	35,097,385

Notes: The regressions are based on the publicly available dataset spanning the survey years 2018, 2019, 2020, and 2021. The sample is restricted to applicants aged 25—64 with positive income. Standard errors are reported in brackets below each coefficient estimate. A loan application is approved if the loan was originated or if the loan application was approved but not accepted by the applicant. A loan application is rejected if the loan application is denied by the financial institution or the preapproval request is denied by the financial institution. Loan applications for all purposes are considered. All regressions include a constant. Demographic controls include age, gender, and race. See the sample selection criteria and the variable definitions in Appendix A.4.2.

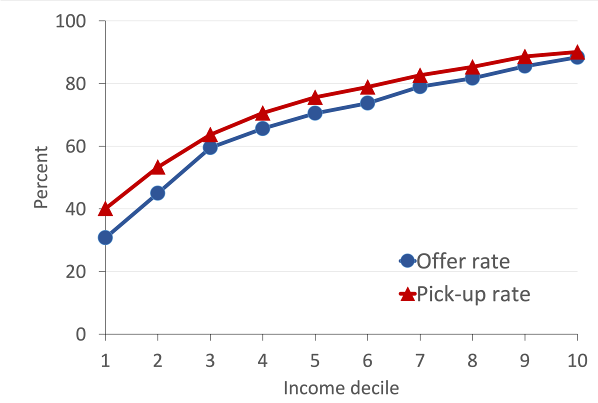

Figure 22: Offer rates and pick-up rates for employer-provided retirement accounts (NCS).

Notes: The figure is constructed using (partially) publicly available data from the National Compensation Survey (NCS) on employer-provided pension plans for 2010–2020. Tabulation by wage income deciles was provided upon a special request to the BLS. *Offer rate* denotes the fraction of employees whose employer offers a retirement plan and who are eligible to participate. *Pick-up rate* is defined as the proportion of employees who actually enroll in the employer-sponsored retirement plan out of the total pool of employees who have access to such a plan and are eligible to enroll.

A.4.4 Contribution rates, copay rates, and mortgage interest rates

One reason for heterogeneity in participation rates in employer-provided retirement plans, documented in Figure 3(a) in the main text, could be differences in copay rates offered by employers. We rely on data from the SCF to study copay rates by employers for active plans along the income distribution.

Figure 23a(a) shows the distribution of copay rates by income. We approximate this distribution by reporting the median, the 10^{th}, and the 90^{th} percentiles of the income-group-specific distribution. Looking first at the median, we find copay rates to vary little with income for most households. For the bottom 60%, median copay rates are at 100% and they decline linearly towards the upper part of the distribution. In the model, we set an average copay rate of 80% in line with the empirical average. Looking at the 10^{th} and 90^{th} percentiles, we find that the variation within income groups is substantially larger than the variation across income groups. We abstract from this heterogeneity in the model and assume a uniform copay rate. Such heterogeneity will be a driver of within-income-group inequality. Incorporating this additional heterogeneity would further increase wealth heterogeneity conditional on income. As we discuss in Section 5, this within-income-group inequality is currently too low in the model compared to the data.

Besides copay rates, the SCF data also provides information on contribution rates to employer-provided pension plans. Figure 23b(b) shows the distribution of contribution rates by income. Again, we summarize this distribution by reporting the median, 10^{th} and 90^{th} percentile.

Looking at the distribution by income groups, we observe, as for copay rates, that the within-income-group variation is much larger than the between-income-group variation, almost by an order of magnitude. However, contrary to copay rates, median contribution rates are slightly increasing with income, while we assume a uniform contribution rate by income in the model.

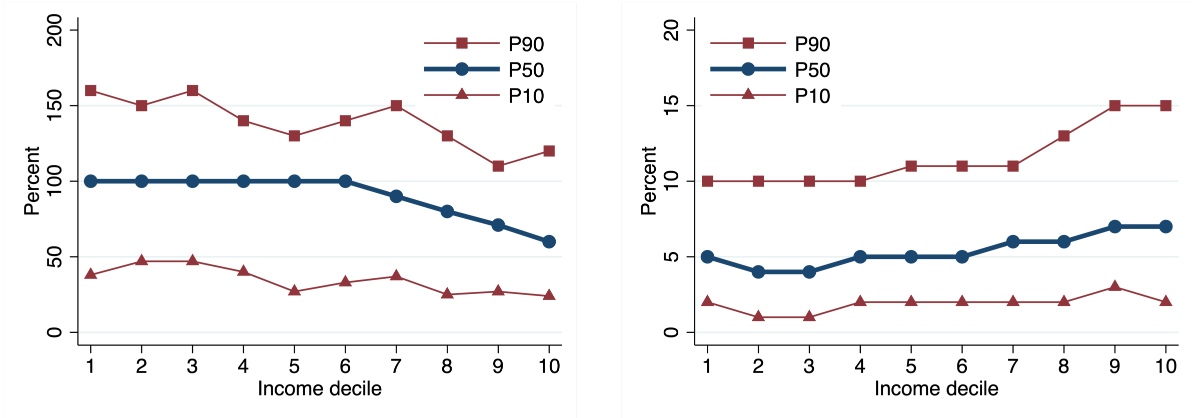

(a) Copay rates

(b) Contribution rates

Figure 23: Copay and contribution rates for active employer-provided retirement plans in the SCF.

Notes: Panel (a) shows the distribution of copay rates by income deciles. The distribution is approximated by the decile-specific median, 10^{th}, and 90^{th} percentiles. Panel (b) shows the equivalent distribution for contribution rates. The results are based on the SCF 2010–2019 datasets. See Appendix A.1 for the sample selection criteria.

Note that copay rates are declining with income whereas contribution rates are increasing with income. If we consider the total inflow of savings by income, the combined copay and contribution flows, we find the two trends by income almost cancel each other as the contribution rate is almost 50% higher in the top decile but the copay rate is 50% lower. Our model assumption of constant contribution and copay rates, therefore, matches the *total* inflow, which is important for the differences in saving rates. However, our model underestimates the *worker*'s share of the contributions at the top of the income distribution.

Similarly, one reason for heterogeneity in participation rates in home equity market, documented in Figure 3(a) in the main text, could be differences in mortgage interest rates. To study mortgage interest rates, we rely on HMDA data (see Section A.4.2 for details). In the HMDA data, we only observe interest rates for *approved* mortgages but these data are still informative about how mortgage rates differ with income and if there is selection on mortgage rates. If approval rates depend on the income-interest rate relationship in a way so that low-income households only get mortgages approved if the interest rate is sufficiently high, then we should see a distribution of mortgage interest rates shifted towards higher values for lower income households.

Figure 24 shows the distribution of mortgage interest rates by income. To describe the distribution, we show the median mortgage rate and the 10^{th} and 90^{th} percentiles of the income-group-specific mortgage interest rate distribution. Two observations are important for our analysis. First, there is very little variation in interest rates at the median along the income distribution. The median interest rate is roughly at 3.5% along the entire income distribution. This finding supports our assumption of a uniform mortgage interest rate in the model. Second, the within-income-group dispersion is orders of magnitude larger than the between-income-group variation but the support of the interest rate distribution varies little with income. The constant support of interest rates by income suggests that there is little selection on interest rates at the

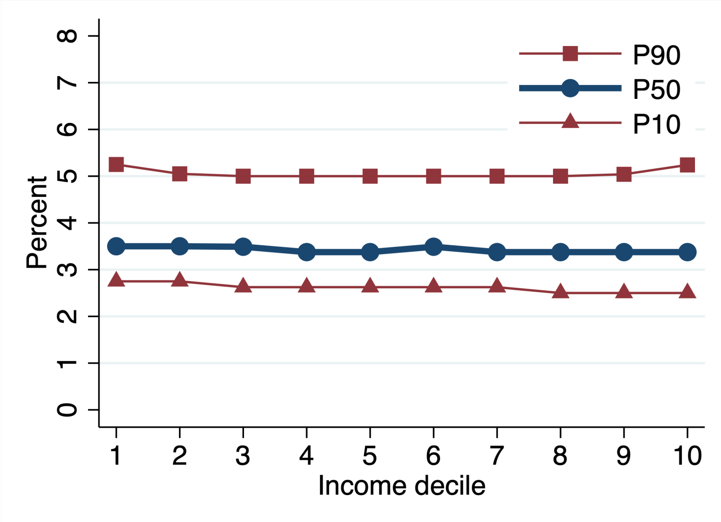

Figure 24: Mortgage interest rates in the HMDA data.

Notes: Distribution of interest rates by income deciles for approved mortgages. The distribution is approximated by the decile-specific median, the 10^{th}, and the 90^{th} percentiles. The results are based on the publicly available dataset spanning the survey years 2018, 2019, 2020, and 2021. The sample is restricted to applicants aged 25–64 with positive income.

approval margin. However, the large dispersion conditional on income supports the idea that the within-income-group wealth dispersion is driven by return differences which we abstract from in our analysis.

B Calibration

B.1 Stochastic process for labor income and business ideas

Figure 25(a) shows that the model matches the empirical variance profile of log income for workers by age in the SCF. Figure 25(b) shows that the model accurately fits the age profile of business wealth-to-income ratios.

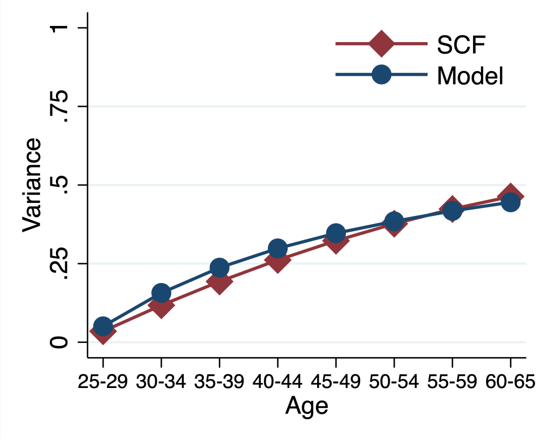

(a) Variance of log labor income by age.

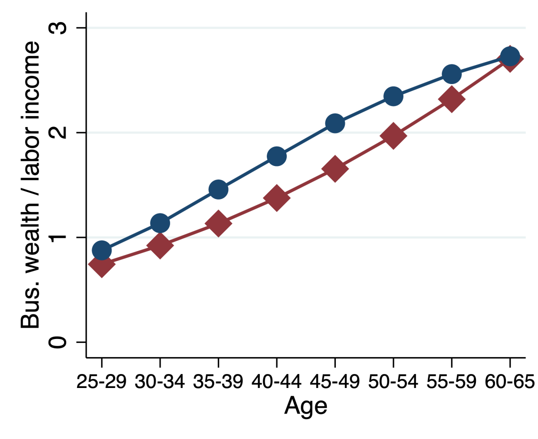

(b) Business wealth-to-labor income by age.

Figure 25: Calibration of the stochastic process for labor income and business ideas.

Notes: In Panel (a), the raw variance profile in the data is smoothed by fitting a quadratic polynomial in age. The intercept of the estimated variance profile in the data is set to zero. In Panel (b), we construct the profile in the model and data by dividing the average business wealth among business owners of a given age by the average labor income among workers of the same age.

B.2 Income-dependent financial frictions

Figure 26 displays the probability schedule of mortgage approval rates $\pi_H(y)$ (Panel (a)) and the schedule of (flow) offer rates for employer-sponsored retirement accounts $\pi_P(y)$ (Panel (b)) used as inputs in our model. For mortgage approval rates, we fit outside of the model a functional form $\pi_H(y) = \alpha_0 + \alpha_1/y$ to the approval rates in the SCF data. We specify the probability to be offered an employer plan to follow $\pi_P(y) = \beta_0 + \beta_1/y$ and estimate parameters (β_0, β_1) inside the model to match the stock of offer rates observed in the SCF dataset.

To justify specifying the probability of being offered an employer plan as a function of the agent's income only, we present Figure 27(a), which shows the marginal effects of income deciles on offer rates obtained from the probit regression analysis of Appendix A.4.1. We report the marginal effects based on the full sample as well as three subsamples depending on the household head's age. The figure demonstrates that conditional on household income, we do not observe large variation in offer rates for employer-sponsored pension plans by age in the SCF, which is why we chose to specify offer rates in the model, $\pi_P(y)$ as a function of household income only.

Figure 27(b) provides similar evidence for mortgage approval rates using the SCF data.

Finally, since we observe the applicant's debt-to-income ratio in HMDA data, we construct Panel (c), which reports the marginal effects of income deciles on mortgage approval rates obtained from the probit regression analysis of Appendix A.4.2. We plot the marginal effects from a full sample as well as two subsamples based on the applicant's debt-to-income ratio, where we distinguish between high debt-to-income ratio (above 35%) and low debt-to-income ratio (below 35%). The figure demonstrates that, conditional on income, debt-to-income ratios do not substantially affect the probability of being approved for a mortgage. This serves again as

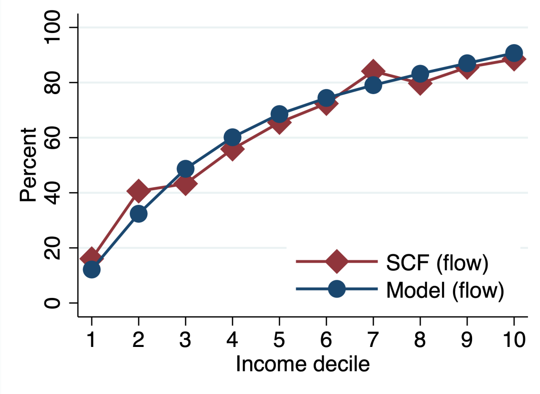

(a) Mortgage approval rates

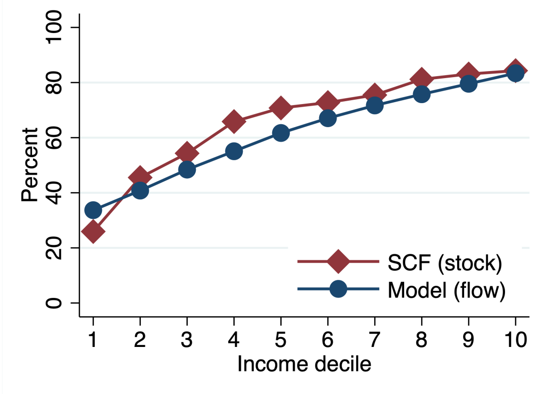

(b) Offer rates

Figure 26: Calibration of the financial frictions.

Notes: Panel (a) shows the mortgage approval rates in the SCF data and the model. The data line is taken from Figure 5 and represents the flow of approved mortgage applications in the past 12 months. The model line corresponds to the approval rates $\pi_H(y)$, which represent the flow of households whose mortgage application is approved in each period. We fit a functional form $\pi_H(y) = \alpha_0 + \alpha_1/y$ to the approval rates in the data. Panel (b) shows the offer rates for employer-sponsored retirement accounts in the data and the model. The data line is taken from Figure 6 and represents the stock of households having received an offer. The model line corresponds to the offer rates $\pi_P(y)$, which represent the flow of households receiving an offer in each period. We assume that $\pi_P(y) = \beta_0 + \beta_1/y$ and calibrate (β_0, β_1) inside the model to match the stock offer rates in the data.

a justification for why we chose the probability of being approved for a mortgage to depend on the agent's income only in the model.

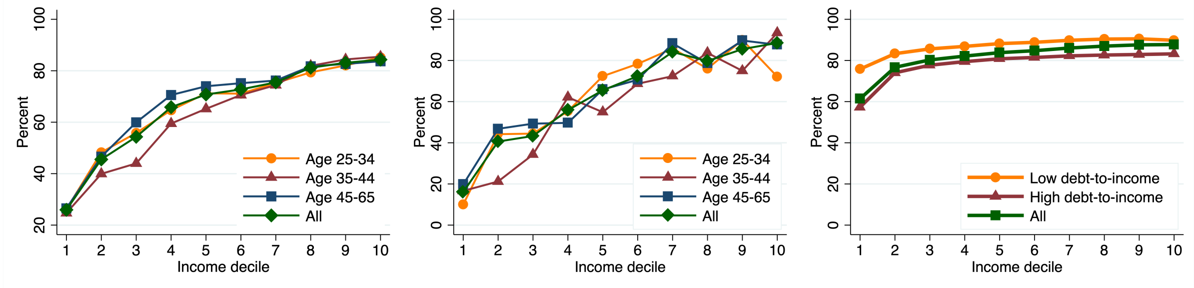

(a) Offer rates by age (SCF) (b) Approval rates by age (SCF) (c) Approval rates by debt-to-income ratio (HMDA)

Figure 27: Marginal effects of income on offer rates and mortgage approval rates in the SCF.

Notes: The figure illustrates the marginal effects of income decile dummies on two probabilities: being offered a retirement plan (Panel (a)) and being approved for a mortgage (Panels (b) and (c)). Panels (a) and (b) are derived from the SCF data detailed in Appendix A.4.1, with effects estimated from a full sample and three subsamples based on the household head's age. Panel (c), based on HMDA data (Appendix A.4.2), estimates effects from a full sample and two subsamples categorized by the applicant's debt-to-income ratio: *Low* (below 35%) and *High* (above 35%).

C Findings

C.1 Model fit

Figure 28 shows that the baseline model fits well the (unconditional) wealth-to-income ratios for each wealth component of core wealth in the SCF data.

Table 12 reports the shares of each core wealth component as well as the share of core wealth held by each income quintile in the model and data.

Figure 29(a) contrasts the mortgage application rates by income in the SCF data and the model. Consistent with the data, the model endogenously generates a flat profile of application rates across the income distribution.[33] Figure 29(b) shows the pickup rates for employer-sponsored retirement accounts. These rates are defined as the fraction of workers with active pension accounts relative to workers without an active account but whose employers offer them the opportunity to participate. Following the strategy in the empirical section, we restrict the sample in the model to workers but keep the income deciles the same as in the full sample for the sake of comparability. The model successfully matches a rising enrollment rate with income, generating a monotone demand for retirement wealth.

[33] See discussion to Figure 5 in the empirical section for details on the empirical measurement.

Figure 28: Wealth-to-income ratios for wealth components in the model and SCF data.

Notes: The SCF sample covers periods 2010–2019. The value of each asset component in the data is computed net of debts. See Appendix A.1 for the sample selection criteria.

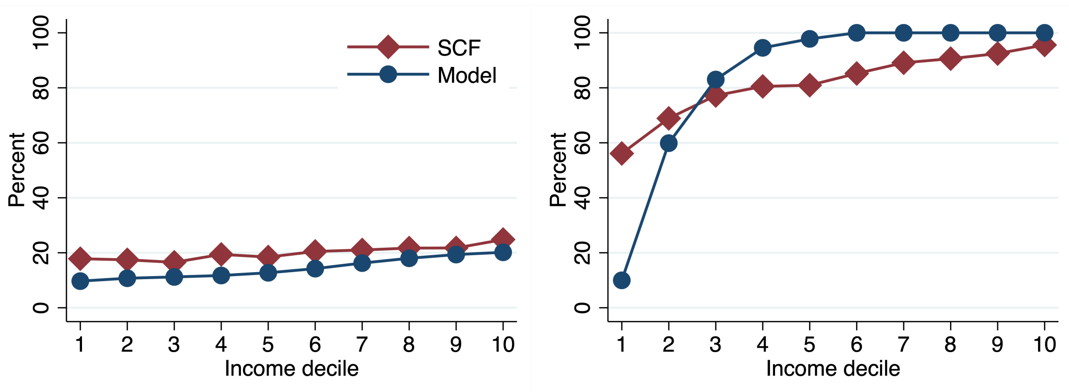

(a) Mortgage application rates

(b) Pick-up rates for retirement accounts

Figure 29: Demand for housing and retirement accounts in the model and data.

Notes: In Panel (a), an applicant in the SCF is a household who submitted a loan application in the past 12 months or who did not submit an application in the past 12 months because he/she thought it would be denied. Applications for a mortgage, home-based loan, and refinancing are considered. In Panel (b), *Pick-up rate* is defined as the proportion of employees who enroll in the employer-sponsored retirement plan out of the total pool of employees who have access to such a plan and are eligible to enroll.

Table 12: Joint distribution of wealth and income in the model and data.

	Income quintile				
	First	Second	Third	Fourth	Fifth
	A. Core wealth				
SCF	1.84	4.84	8.80	14.99	69.52
Model	1.37	4.06	9.24	21.86	63.47
	B. Pension wealth				
SCF	0.82	3.51	8.10	16.84	70.73
Model	0.74	4.09	10.12	22.97	62.07
	C. Housing wealth				
SCF	2.79	6.44	10.45	16.13	64.18
Model	1.76	4.64	9.96	21.89	61.74
	D. Business wealth				
SCF	1.11	2.31	4.04	7.51	85.04
Model	1.80	1.84	4.10	18.52	73.75

Notes: The table reports the shares of wealth components held by each quintile of the income distribution in the model and SCF. The reported values take non-participants in the respective asset market into account.