Essays in Income and Wealth Inequality

Inauguraldissertation
zur Erlangung des Grades eines Doktors
der Wirtschafts- und Gesellschaftswissenschaften

durch

die Rechts- und Staatswissenschaftliche Fakultät
der Rheinischen Friedrich-Wilhelms-Universität
Bonn

vorgelegt von
Ulrike Isabel Steins
aus Kleve

Bonn, 2020

Dekan:	Prof. Dr. Jürgen von Hagen
Erstreferent:	Prof. Dr. Moritz Kuhn
Zweitreferent:	Prof. Dr. Moritz Schularick
Tag der mündlichen Prüfung:	30. November 2020

Acknowledgements

I would like to thank the following people, without whom this dissertation would not have been possible and in any case not finished.

I am particularly grateful to my main supervisors and co-authors Moritz Kuhn and Moritz Schularick for continuous support, encouragement and fruitful discussions through the various projects over the last years. Especially I would like to thank Moritz Kuhn for the steady motivation both during the time at the BGSE and after I left it. I am also very grateful to Keith Kuester who supervised me at the beginning of the PhD and greatly supported me to find by field of economic interest.

In addition, I would like to express my gratitude to the Bonn professors as well as my BGSE colleagues for sharing their thoughts and insights. A special thanks goes to my friends Yao Chen, Nina Bobkova and Jana Willrodt. Without them I probably would not have made it through the first year. I am deeply grateful that I had them at my side during our time at the BGSE and that we are still close friends.

The members of the "Köln-Crew", first and formost, my best friends Antje Wessels and Sarah Kloppich as well as Inga Felde, Mikko Bayer and Sascha Bayer, have always made sure that I do not forget other important things in life besides the PhD. I am very thankful for this and that we master life together for so many years now.

A very special thanks goes to my family: my parents Günter and Inge, my sister Anne together with Stefan and Gabriel (in his very special, amusing way) Trummer as well as my brother David. They have always been tremendously supportive and their steady encouragement was a major factor for not giving up and finally being able to finish the PhD. This is especially true for my parents, who have created a perfect writing environment for me at home over the past few years.

Finally, I would like to acknowledge the financial support from the BGSE, the German Research Foundation as well as the Science Foundation of the Sparkassen-Finanzgruppe.

Contents

List of Figures

List of Tables

1

Introduction

The distribution of income and wealth across households in the U.S. has changed dramatically since the 1950s. How are rising income and wealth inequality linked? This thesis aims at providing new insights into the relationship between income, wealth accumulation and inequality. The first two chapters focus on trends of income and wealth portfolios over time. In the third chapter a new model is introduced that simulates the observed joint distribution of income and wealth.

Chapter 1 "Income and Wealth Inequality in America, 1949-2016" introduces a new long-run data set based on archival data from historical waves of the Survey of Consumer Finances. Studying the joint distribution of household income and wealth, we expose the central importance of portfolio composition and asset prices for wealth dynamics in postwar America. Asset prices shift the wealth distribution due to systematic differences in household portfolios along the wealth distribution. Middle-class portfolios are dominated by housing, while rich households predominantly own business equity. Differential changes in equity and house prices shaped wealth dynamics in postwar America and decoupled the income and wealth distribution over extended periods.

Chapter 2 "Modigliani Meets Minsky: Inequality, Debt and Financial Fragility in America, 1949-2016" studies the secular increase in household debt, and its relation to growing income inequality and financial fragility. We exploit a household-level dataset that covers the joint distributions of debt, income, and wealth in the U.S. over the past seven decades. The data show that increased borrowing by middle class families with low income growth played a central role for rising indebtedness. Debt-to-income ratios have risen most strongly for households between the 50th and 90th percentile of the income distribution. While income growth was low, middle class families borrowed against sizable housing wealth gains from rising home prices. Home-equity borrowing accounts for about half of the increase in U.S. household debt between the 1970s and 2007. The resulting debt increase made balance sheets more sensitive to income and house price fluctuations, and turned the American middle class into the epicenter of growing financial fragility.

In Chapter 3 "To Have or Not to Have: Understanding Wealth Inequality" the drivers of wealth accumulation and the wealth distribution are examined. The core of existing models of savings behavior is a mapping from the income process to wealth accumulation with saving motives arising from precautionary or life cycle reasons. Exploring a benchmark incomplete markets model calibrated to match the marginal distributions of income and wealth, we find that it struggles to account for the *joint* distribution of income and wealth. Using data from the Survey of Consumer Finances, we document two new facts about the American wealth distribution: First, three main asset classes, home equity, business equity, and retirement accounts, account for about 67 percent of household wealth and are as unequally distributed as total wealth. Second, for a given level of income, the largest part of the differences in holdings of these assets is accounted for by the extensive margin — a question of *to have or not to have*. We develop a model of wealth accumulation focusing on these three asset classes and access to them (the extensive margin). In the model, the labor market (income) situation of households affects their ability to invest in assets generating a further link between income and wealth accumulation. The calibrated model jointly generates large wealth inequality and is consistent with the joint distribution of income and wealth. The key innovation over existing models is that we emphasize financial frictions limiting access to asset investment (extensive margin) rather than affecting the intensive margin of saving decisions. In a policy counterfactual, we confirm the prominent policy conjecture that improving access to housing especially for poor households increases aggregate household wealth and reduces wealth inequality significantly.

2

Income and Wealth Inequality in America, 1949-2016

Joint work with Moritz Kuhn and Moritz Schularick

2.1 Introduction

We live in unequal times. The causes and consequences of widening disparities in income and wealth have become a defining debate of our age. Recent studies have made major inroads into documenting trends in either income or wealth inequality in the United States (Piketty and Saez (2003), Kopczuk et al. (2010), Saez and Zucman (2016)), but we still know little about how the joint distributions of income and wealth evolved over the long run. This paper fills this gap.

The backbone of this study is a newly compiled dataset that builds on household-level information and spans the entire U.S. population over seven decades of postwar American history. We unearthed historical waves of the Survey of Consumer Finances (SCF) that were conducted by the Economic Behavior Program of the Survey Research Center at the University of Michigan from 1947 to 1977. In extensive data work, we linked the historical survey data to the modern SCFs that the Federal Reserve redesigned in 1983.[1] We call this new resource for inequality research the SCF+.

The SCF+ complements existing datasets for long-run inequality research that are based on income tax and social security records, but also goes beyond them in a number of important ways. Importantly, the SCF+ is the first dataset that makes it possible to

[1] A few studies such as Malmendier and Nagel (2011) or Herkenhoff (2013) exploited parts of these data to address specific questions, but no study has attempted to harmonize modern and historical data in a consistent way. Note that we leave the post-1983 modern SCF unchanged. Its value for studying distributional trends has been demonstrated in recent contributions by Bricker et al. (2016) and Wolff (2017).

study the joint distributions of income and wealth over the long run. As a historical version of the SCF, it contains the same comprehensive income and balance sheet information as the modern SCFs. This means that we do not have to combine data from different sources or capitalize income tax data to generate wealth holdings. Moreover, the SCF+ contains granular demographic information that can be used to study dimensions of inequality —such as long-run trends in racial inequality— that so far have been out of reach for research.

Our analysis speaks to the quest to generate realistic wealth dynamics in dynamic quantitative models (Benhabib and Bisin (2018), Fella and De Nardi (2017), Gabaix et al. (2016), Hubmer et al. (2017)). A key finding of our paper is that a channel that has attracted little scrutiny so far has played a central role in the evolution of wealth inequality in postwar America: asset price changes induce large shifts in the wealth distribution. This is because the composition and leverage of household portfolios differ systematically along the wealth distribution. While the portfolios of rich households are dominated by corporate and non-corporate equity, the portfolio of a typical middle-class household is highly concentrated in residential real estate and, at the same time, highly leveraged. These portfolio differences are persistent over time. We document this stylized fact and expose its consequences for the dynamics of the wealth distribution.

An important upshot is that the top and the middle of the distribution are affected differentially by changes in equity and house prices. Housing booms lead to substantial wealth gains for leveraged middle-class households and tend to decrease wealth inequality, all else equal. Stock market booms primarily boost the wealth of households at the top of the wealth distribution as their portfolios are dominated by listed and unlisted business equity. Portfolio heterogeneity thus gives rise to a race between the housing market and the stock market in shaping the wealth distribution. We show that over extended periods in postwar American history, such portfolio valuation effects have been predominant drivers of shifts in the distribution of wealth.

A second consequence of portfolio heterogeneity is that asset price movements can introduce a wedge between the evolution of the income and wealth distribution. For instance, rising asset prices can mitigate the effects that low income growth and declining savings rates have on wealth accumulation. Looking at income and wealth growth of different parts of the wealth distribution, we find such a divergence played a prominent role in the four decades before the financial crisis. The middle class (50th-90th percentile) rapidly lost ground to the top 10% with respect to income but, by and large, maintained its wealth share thanks to substantial gains in housing wealth. The SCF+ data show that incomes of the top 10% grew 80% more than incomes of middle-class households (50th-90th percentile) and 120% more than incomes in the bottom 50% of households.

In line with previous research, the SCF+ data thus confirm a strong trend toward growing income concentration at the top (Piketty and Saez (2003); Kopczuk et al. (2010)). However, when it comes to wealth, the picture is different. For the bottom 50% of the wealth distribution, wealth grew 100% in excess of income between 1971 and 2007. A particularly pronounced difference using CPI inflation that leads to zero income growth and a doubling of wealth. For the middle class and for the top 10%, wealth grew at approximately the same rate despite diverging income paths. As a result, wealth-to-income ratios increased most strongly for the bottom 90% of the wealth distribution. That the SCF+ data reach back to the 1950s and 1960s, that is, before the income distribution started to widen substantially, makes it possible to expose these divergent trends.

Importantly, price effects account for a major part of the wealth gains of the middle class and the lower middle class. We estimate that between 1971 and 2007, wealth of the bottom 50% grew almost entirely because of price effects — essentially a doubling of wealth compared to household income without any (active) saving. Price-related wealth growth is high for the bottom 50% despite below-average homeownership rates because virtually all existing wealth of this group is invested in highly leveraged housing wealth. Even in the middle and at the top of the distribution, asset price induced gains accounted for close to half of total wealth growth over the 1971-2007 period, comparable to the contribution of savings flows. From a political economy perspective, it is conceivable that the strong wealth gains for the middle and lower middle class helped to dispel discontent about stagnant incomes. They may also help to explain the disconnect between trends in income and consumption inequality that have been the subject of some debate (Attanasio and Pistaferri, 2016). When house prices collapsed in the 2008 crisis, the same leveraged portfolio position of the middle class brought about substantial wealth losses, while the quick rebound in stock markets boosted wealth at the top. Relative price changes between houses and equities after 2007 have produced the largest spike in wealth inequality in postwar American history. Surging post-crisis wealth inequality might in turn have contributed to the perception of sharply rising inequality in recent years.

Thanks to its demographic detail, we can also exploit the SCF+ to shed new light on the long-run evolution of racial inequalities. The SCF+ covers the entire postwar history of racial inequality and spans the pre- and post-civil rights eras. With information on income and wealth at the household level, we do not only complement recent studies of the long-run evolution of racial wage inequality (Bayer and Charles, 2017), but we add new dimensions. Most importantly, the SCF+ data offer a window on long-run trends in racial *wealth inequality* that have so far remained uncharted territory. We expose persistent and, in some respects, growing inequalities between black and white Americans. Income disparities today are as big as they were in the pre-civil rights era. In 2016, black household income is still only half of the income of white households. The racial wealth gap is even wider and is still as large as it was in the 1950s and 1960s. The median black household persistently has less than 15% of the wealth of the median white household. We also find that the financial crisis has hit black households particularly hard and has undone the little progress that had been made in reducing the racial wealth gap during the 2000s (Wolff, 2017). The overall summary is bleak. The typical black household remains poorer than 80% of white households.

Related literature: Research on inequality has become a highly active field, and our paper speaks to a large literature. Analytically, the paper is most closely related to recent contributions emphasizing the importance of heterogeneity in returns on wealth for the wealth distribution. On the empirical side, this literature has mainly worked with European data, while our paper addresses the issues with long-run micro data for the United States. Bach et al. (2016) study administrative Swedish data. With regard to heterogeneity in returns along the wealth distribution, Fagereng et al. (2016) use administrative Norwegian tax data and document substantial heterogeneity in wealth returns and intergenerational persistence. For France, Garbinti et al. (2017) analyze the long-run distribution of wealth as well as the role of return and savings rate differentials. In the American context, Wolff (2016) demonstrates the sensitivity of middle-class wealth to the house price collapse in the Great Recession, and his earlier research (Wolff, 2002) is closely related as it discusses the sensitivity of the U.S. wealth distribution to asset price changes. In the policy debate, the role of asset prices for the wealth distribution has also been discussed, for example, by Yellen (2014). Moreover, Kuhn and Ríos-Rull (2016) argue that housing wealth plays an important role for the wealth distribution.

With respect to data production and the emphasis on long-run trends, our paper complements the pioneering work of Piketty and Saez (2003), Kopczuk and Saez (2004), and Saez and Zucman (2016), as well as the work of Kopczuk et al. (2010). Our paper also speaks to the more recent contribution of Piketty et al. (2017), who combined micro data from tax records and household survey data to derive the distribution of income reported in the national accounts. Saez and Zucman (2016) estimate the wealth distribution by capitalizing income flows from administrative data. This approach is advantageous for households at the top of the distribution that hold a significant part of their wealth in assets that generate taxable income flows. Yet many assets in middle-class portfolios do not generate taxable income flows — housing being a prime example. The SCF+ provides long-run data on all sources of income (including capital and non-taxable income) as well as the entire household balance sheet with all assets (including residential real estate) and liabilities (including mortgage debt). Playing to the strength of our data, our paper focuses on the bottom 90% of households, not on changes in inequality at the very top. We also connect our paper to the recent paper by Bricker et al. (2016) that demonstrates the potential of the modern SCFs to study distributional trends even at the top, and discuss the differences between the more advanced modern SCF and the historical SCF waves.[2]

Theoretical work modeling the dynamics of wealth inequality has grown quickly. A common thread is that models based on labor income risk alone typically produce too little wealth concentration and cannot account for substantial shifts in wealth inequality that occur over short time horizons. Our paper speaks to recent work by Benhabib

[2]Work in labor economics often relies on data from the CPS. Examples are Gottschalk and Danziger (2005) and Burkhauser et al. (2009). Most relevant for our work is Burkhauser et al. (2012), who show that trends in income inequality derived from the CPS are similar to the inequality series based on tax data in Piketty and Saez (2003). They also provide a detailed discussion of the conceptual differences in measuring income in the tax and CPS data.

and Bisin (2018), Benhabib et al. (2017), and Gabaix et al. (2016), who discuss the importance of heterogeneous returns for the wealth distribution and its changes over time. In another recent paper, Hubmer et al. (2017) use variants of incomplete markets models to quantify the contribution of different drivers for rising wealth inequality and point to return differences and portfolio differences as a neglected line of research. Our findings support the emphasis on asset returns.[3] Glover et al. (2017) quantify the welfare effects of wealth changes resulting from portfolio differences and asset price changes during the Great Recession. Fella and De Nardi (2017) survey the existing literature and discuss different models from the canonical incomplete market model to models with intergenerational transmission of financial and human capital, rate of return risk on financial investments, and more sophisticated earnings dynamics.

Outline: The paper is divided into three parts. The first part documents the extensive data work that we have undertaken over the past years to construct the SCF+ and what we did to align the historical and modern SCF data. The second part then exploits the new data and presents stylized facts for long-run trends in income and wealth inequality, including racial inequalities, that emerge from the SCF+. The third part studies the joint distributions of income and wealth and exposes the central importance of asset price changes for the dynamics of the wealth distribution in postwar America. The last section concludes.

2.2 Constructing the SCF+

The SCF is a key resource for research on household finances in the United States. It is a triennial survey, and the post-1983 data are available on the website of the Board of Governors of the Federal Reserve System[4]. Yet the first consumer finance surveys were conducted as far back as 1947. The early SCF waves were directed by the Economic Behavior Program of the Survey Research Center of the Institute for Social Research at the University of Michigan. The surveys were taken annually between 1947 and 1971, and then again in 1977. The raw data are kept at the Inter-University Consortium for Political and Social Research (ICPSR) at the Institute for Social Research in Ann Arbor, Michigan.

For this paper, we linked the archival survey data to the post-1983 SCF. To do this, we harmonized and re-weighted the historical data to make them as compatible as possible with the modern SCF. Note that we do not adjust the post-1983 SCF data. On the contrary, we take the advanced survey design of the modern SCF as the benchmark and adjust the historical surveys so that they come as close as possible to this benchmark. We discuss in detail below and in the Appendix 2.A.2 how we proceeded and how consistent the historical and modern data are, especially when it comes to the top of the

[3] See also Castaneda et al. (2003) for a benchmark model of cross-sectional income and wealth inequality and Kaymak and Poschke (2016) for another recent attempt to explain time trends.

[4] `https://www.federalreserve.gov/econres/scfindex.htm`. See Bricker et al. (2017) for results from the 2017 SCF data and for general information on the SCF data and sampling.

distribution. The combined dataset adds four decades of household-level micro data, effectively doubling the time span covered by the SCF. As a new resource for long-run research on household finances, we refer to this historically extended version of the SCF as the *SCF+*.

The SCF+ complements the data sets for long-run trends in the distribution of income and wealth in the U.S. that Piketty and Saez (2003), Kopczuk and Saez (2004), and Saez and Zucman (2016) have compiled using administrative tax data. Other researchers have used the 1962 Survey of Financial Characteristics of Consumers (SFCC) that provides a snapshot of the financial conditions of U.S. households in 1962 (Wolff, 2017).[5] But so far the tax data constitute the only data covering the entire post-war period on a continuous basis. The SCF+ provides an opportunity to corroborate and improve our understanding of postwar trends in the distribution of income and wealth.

For future researchers, it is important to have a good understanding of the relative strengths and weaknesses of the SCF+ for inequality research. A key advantage of the long-run tax data is their compulsory collection process resulting in near-universal coverage at the top of the distribution, whereas survey data have to cope with non-response of rich households. This being said, the tests carried out in a recent paper by Bricker et al. (2016) show that the modern SCF with its combined administrative and survey data methodology also captures households at the very top of the distribution.

The strengths of the administrative data in terms of accuracy and coverage at the top of the distribution also have to be weighed against the attractive properties of survey data in other respects. Most importantly, the survey data contain direct measurements of assets and debt plus the information to stratify households by demographic characteristics. The survey data also cover people who do not file taxes, and the unit of analysis is the household, not the tax unit. This structure is in line with economic models in which the household is the relevant unit for risk and resource sharing.[6]

Moreover, specific challenges arise when income tax data are used to construct wealth estimates. The capitalization method of Saez and Zucman (2016) relies on observable income tax flows that are capitalized to allocate aggregate wealth positions in the cross section. While ingenious as an approach, some gaps remain because a substantial part of wealth does not generate taxable income flows and has to be imputed (often on the basis of survey data). The key asset here is owner-occupied housing as well as its corresponding liability, mortgage debt. Pension assets also do not generate taxable income flows, and unrealized capital gains do not show up on tax returns until they are realized.

In the estimates of Saez and Zucman (2016), about 90% of the total wealth outside the top 10% has to be imputed. And even for the top 10%, the share of imputed wealth

[5]For the construction of the SCF+, we have set the distributional information from the 1962 SCF against the SFCC data and generally found the differences to be small. More details below.

[6]In 2012, there were about one-third more tax units (160.7 million) than households (121.1 million) in the United States. Bricker et al. (2016) argue that relying on tax units could lead to higher measured income concentration toward the top of the distribution.

stands at 40%. Saez and Zucman (2016) correctly stress that the exact distribution of these assets is of minor importance for the very top of the wealth distribution. Yet for researchers interested in long-run distributional changes outside the very top, these are binding constraints that the SCF+ overcomes. The capitalization method also has to derive returns for individual asset classes from a combination of capital income from tax data and aggregate estimates from the financial accounts. Kopczuk (2015) provides an illustration of how this method can lead to an upward bias of wealth concentration during low interest rate periods, and the recent paper by Bricker et al. (2018) quantifies this bias and discusses in detail other conceptual differences between survey estimates and estimates based on tax data.

2.2.1 Variables in the SCF+

The variables covered in the historical surveys of the SCF+ correspond to those in the contemporary SCF, but the exact wording of the questions can differ from survey to survey. Some variables are not continuously covered, so we have to impute values in some years. We explain the imputation procedure in the following section. Our analysis focuses on the four variables that are of particular importance for household finances: income, assets, debt, and wealth. In the analysis, we use all data and abstain from any sample selection. We adjust all data for inflation using the consumer price index (CPI) and report results in 2016 dollars.[7] Table 2.2.2 provides a general overview over variables and years when imputation is used. Online Appendix 2.A.1.1 contains additional information.

Income: We construct total income as the sum of wages and salaries, income from professional practice and self-employment, rental income, interest, dividends, transfer payments, as well as business and farm income. Note that we do not include imputed rental income of homeowners in the baseline, but we provide additional results in Appendix 2.A.4.2.

Assets: The historical SCF waves contain detailed information on household assets. We group assets into the following categories: liquid assets, housing, bonds, stocks and business equity, mutual funds, the cash value of life insurance, defined-contribution retirement plans[8], other real estate, and cars. Liquid assets comprise the sum of checking, savings, call/money market accounts, and certificates of deposits. By contrast, Social Security as a key asset for most families is not measured as part of household wealth.

[7] We use CPI data from the Macrohistory Database (Jordà et al., 2017). The series combines the CPI-U-RS series (1978-2016) from the Bureau of Labor Statistics, and the CPI-All Urban Consumers for 1948-1977. The CPI shows higher inflation rates relative to the personal consumption expenditure index (PCE) as discussed by Furth (2017). Comparisons of relative income and wealth trends between groups are unaffected by the choice of the deflator, but caution is warranted for absolute statements about income and wealth growth. We provide a sensitivity analysis using the PCE in Appendix 2.A.4.4.

[8] Data on defined-contribution retirement plans are only available from 1983 onward. However, according to the financial accounts of the United States, this variable makes up a small part of household wealth before the 1980s, so missing information before 1983 is unlikely to change the picture meaningfully. Up to 1970, defined-contribution plans correspond to less than 1% of average household wealth.

The wealth concept used here hence follows the literature by focusing on marketable household wealth. A more detailed discussion of the importance of Social Security for household wealth and its distribution can be found in Bricker et al. (2016).

Debt: Total debt consists of housing and non-housing debt. Housing debt is calculated as the sum of debt on owner-occupied homes, home equity loans, and home equity lines of credit. For 1977, only the origination value (instead of the current value) of mortgages is available. Using information on the year the mortgage was taken out, remaining maturity, and an estimated annual interest rate, we create a proxy for debt on homes for 1977. All debt other than housing debt refers to and includes car loans, education loans, and other consumer loans.

Wealth: We construct wealth as the consolidated value of the household balance sheet by subtracting debt from assets. Wealth constitutes households' net worth.

2.2.2 Weights and imputations

The SCF is designed to be representative of the U.S. population. As Bricker et al. (2016) discuss, the modern SCF applies a sophisticated dual-frame sampling scheme to oversample wealthy households, combining administrative and survey data. The historical surveys did not oversample households at the top. In this section and in Appendix 2.A.2, we outline how we dealt with the issue and discuss the implications for the representativeness of the SCF+. In addition to the adequate coverage of wealthy households, we also need to ensure representative coverage of demographic characteristics such as race, age, and education.

Oversampling of wealthy households: Since its redesign in 1983, the SCF consists of two samples. The first sample is drawn using area probability sampling of the entire U.S. population based on Census information. In addition, a second so-called *list sample* is drawn based on tax information.[9] For both samples, survey weights are constructed separately. In the list sample, survey weights have to be disproportionally adjusted for non-response. The weight of each household corresponds to the number of similar households in the population. In a final step, both samples are combined and survey weights are adjusted so that the combined sample is representative of the U.S. population (Kennickell and Woodburn, 1999).

Before 1983, the historical SCF data are not supplemented by a list sample. As a consequence, the challenge of adequately representing households at the very top is likely to be more pervasive (Sabelhaus et al., 2015). Missing households at the top could potentially lead to an under-representation and bias historical inequality measures downwards.

For the construction of the SCF+, we use information from the 1983 list sample to adjust

[9]The methodology has evolved over time and uses a combination of income capitalization with income from several tax years and regression evidence based on existing surveys. See Kennickell (2017) and Bricker et al. (2017) for details and further references.

the coverage of rich households in the pre-1983 data. In a first step, we determine the proportion of households in the 1983 list sample relative to all households. Their share corresponds to approximately 2%. In a second step, we determine where the households from the list sample are located in the income and wealth distribution in 1983. We find that most observations are among the top 5% of the income and wealth distribution. Using this information, we adjust survey weights in all survey years before 1983 in two steps. First, for each year we extract all observations that are simultaneously in the top 5% of the income and wealth distribution. Secondly, we increase the weighting of these households in such a way that we effectively add 2% of wealthy households to the sample and adjust the remaining weights accordingly. This approach is similar in spirit to Bricker et al. (2018), who adjust SCF weights inversely proportional to the overlap of the SCF sample with the Forbes list.

The list sample of the modern SCF is concentrated in the top 1% of the wealth distribution, and great effort goes into identifying these households as Bricker et al. (2016) discuss. A potential concern with our adjustment of the historical data is that we can only increase the weight of households that are sampled. This could be problematic if non-response rates have changed over time, or if the older surveys failed to identify and contact wealthy households in sufficient numbers. One way to get a better sense of how pervasive these issues are, is to compare the 1983 data with the 1962 Survey of Financial Characteristics of Consumers (SFCC). The SFCC was the only historical survey that also used a dual-frame sampling scheme similar to the 1983 list sample on the basis of income tax records.

Table 2.2.1: Share of respondents from list sample at the top of the distribution

	Income		Wealth	
	top 10%	top 5%	top 10%	top 5%
SFCC 1962	21 %	35 %	20 %	28 %
SCF 1983	17 %	34 %	17 %	32 %

Notes: Share of respondents in the 1962 SFCC and 1983 SCF data from list sample in different parts of the income and wealth distribution. The left part of the table shows the shares of households in the top 10% and top 5% of the income distribution in the 1962 SFCC and the 1983 SCF data from the list sample. The right part of the table shows the corresponding shares for the top 10% and top 5% of the wealth distribution.

Table 2.2.1 compares the non-response patterns at the top of the income and wealth distribution from these two surveys. Importantly, we find little evidence for a pronounced time trend in non-response of wealthy households. For the modern SCF surveys, the reported response rates for the list sample also do not point to any trends in non-response rates for the list sample over time (see Bricker et al. (2016), Bricker et al. (2017)). In Appendix 2.A.2 we apply a battery of tests to the historical data that were proposed in a recent paper by Bricker et al. (2016) to examine how well the modern SCFs perform in capturing the top of the distribution relative to the tax data. More precisely, we test how many households in the SCF+ are above the 99th percentile threshold for income and wealth from the tax data, and how mean income and wealth above this threshold compare. Although these tests do not signal systematic deviations, the strength of the SCF+ data clearly lies in their comprehensive coverage of the lower ranks of the distribution. A lot of research has already been devoted to small groups at the very top of the distribution, but less is known about long-run distributional trends among the bottom 99% of households. Consequently, in this paper we will not talk about the top 1% of households, but will focus on income and wealth trends of the remaining 99% of American families.

Demographic characteristics: We compare the demographic characteristics in the surveys before 1983 with data from the U.S. Current Population Survey (CPS). To obtain samples that match the CPS data, we subdivide both the CPS and the SCF+ data into demographic subgroups. Subgroups are determined by age of the household head, college education, and race. In addition to these demographic characteristics, we include homeownership as an additional dimension. We adjust SCF+ weights by minimizing the difference between the share of each subgroup in the SCF+ and the respective share in the CPS. As Census data are only available since 1962, we rely on data from the Decennial Census and linear inter- and extrapolation for the earlier years.

Figure 2.2.1: Population shares of age groups, college households, and black households

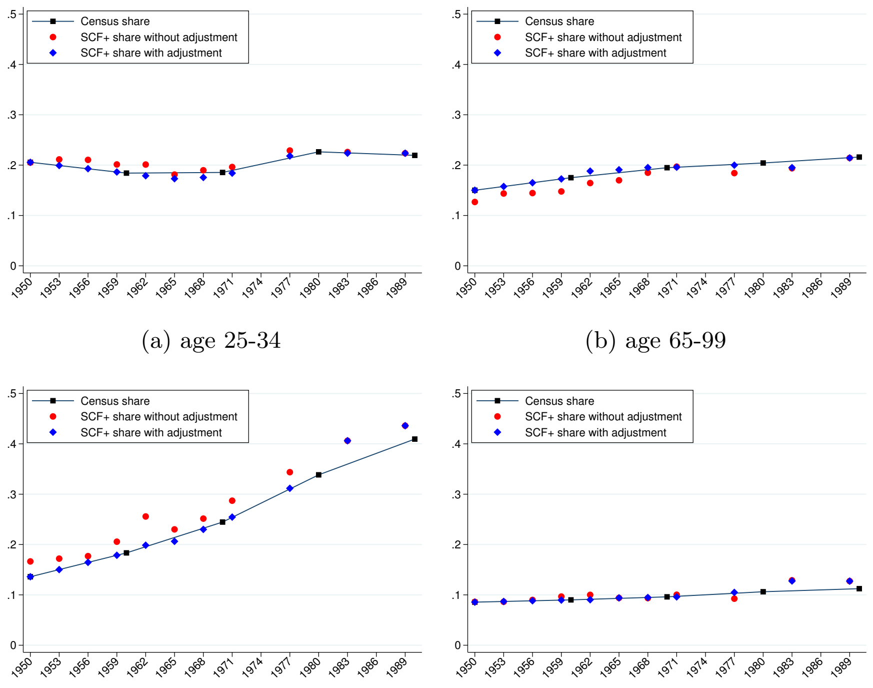

(a) age 25-34

(b) age 65-99

(c) college

(d) black

Notes: Lines with black squares show the population share of the respective demographic group in the CPS/Census data. The red dots show the population shares of the respective group using the original (unadjusted) survey data. The blue diamonds show the population shares using the adjusted survey data.

Figure 2.2.1 shows the shares of 10-year age groups, college households, and black households in the CPS/Census (black squares) and in the SCF+ with the adjustment of survey weights (red dots). Using adjusted weights (blue diamonds), the distributions of age, education, and race closely match the CPS/Census data. We match the homeownership rate equally well after the adjustment (see Figure 2.A.1).

Missing variables: The imputation of missing variables is done by predictive mean matching as described in Schenker and Taylor (1996). This multiple imputation method assigns variable values by finding observations that are closest to the respective missing observations. In line with the post-1983 data, we impute five values for each missing observation. We account for a potential undercoverage of business equity before 1983 and follow the method proposed by Saez and Zucman (2016) to adjust the observed holdings in the micro data with information from the financial accounts. A detailed description of these steps is provided in Appendix 2.A.1.1.

Table 2.2.2: Data availability

	income			financial assets			nonfinancial assets			debt			
Survey year	total	labor	labor + business	liquid assets	bonds	equity	housing	other real estate	business	total	housing	other real estate	nonhousing
1949	O	O	O	O	O	O	O	I	I	O	O	O	O
1950	O	O	O	O	O	O	O	O	O	O	O	O	O
1951	O	O	O	O	O	I	O	I	I	O	O	O	O
1953	O	O	O	O	O	O	O	O	O	O	O	O	O
1954	O	O	O	O	O	I	O	I	I	O	O	O	O
1955	O	O	O	O	O	O	O	I	I	O	O	O	O
1956	O	O	O	O	O	I	O	I	I	O	O	I	O
1957	O	O	O	O	O	I	O	I	I	O	O	O	O
1958	O	O	O	O	O	I	O	I	I	O	O	O	O
1959	O	O	O	O	O	I	O	I	I	O	O	O	O
1960	O	I	O	O	O	O	O	O	O	O	O	I	O
1962	O	I	O	O	O	O	O	O	O	O	O	I	O
1963	O	I	O	O	O	O	O	O	O	O	O	I	O
1965	O	I	O	O	O	I	O	I	I	O	O	I	O
1967	O	O	O	O	I	O	O	I	I	O	O	I	O
1968	O	O	O	O	I	O	O	O	I	O	O	O	O
1969	O	O	O	O	I	O	O	O	I	O	O	O	O
1970	O	O	O	O	O	O	O	O	O	O	O	O	O
1971	O	O	I	O	I	I	O	I	I	O	O	I	O
1977	O	O	I	O	O	O	O	O	I	O	O	O	O

Notes: Data availability across survey years. The first column shows the survey year. Other columns show variables. The letter *O* indicates that original observations from that survey year are used, the letter *I* indicates that the variable has been imputed using data from other survey years. In some years totals are available but components are not separately reported and had to be imputed. If the total is constructed as sum of components, then totals are marked as imputed if any component is imputed. Equity includes stocks and mutual funds.

Table 2.2.2 shows the variables and their coverage, as well as the years in which we imputed data.[10] In Appendix 2.A.1.1, we also explain how we impute the value of cars for selected years based on model and purchasing year.

The final SCF+ data set comprises 35 survey years with cross-sectional data, totaling 102,304 household observations. The number of observations varies from a minimum of 1,327 in 1971 to a maximum of 6,482 in 2010. Table 2.A.1 in the appendix reports the number of observations for all survey years.

2.2.3 Aggregate trends

Before looking in detail at the evolution of the income and wealth distributions since World War II, the first step is to benchmark aggregate trends from the SCF+ to the national income and product accounts (NIPA) and the financial accounts of the United States (FA). To do so, we have to take into account that even high-quality micro data do not always correspond one-to-one to aggregate data as measurement concepts differ. We follow Henriques and Hsu (2014) and Dettling et al. (2015) to account for the conceptual differences when constructing income and wealth series. We relegate the details to Appendix 2.A.1.5. For the modern SCF data, Henriques and Hsu (2014) and Dettling et al. (2015) conclude that after accounting for the conceptual differences between micro and macro data, the data align well. They also provide detailed discussions for observed differences. Figure 2.2.2 compares indexed time series for average household income and wealth from the SCF+ with the corresponding series constructed from NIPA and FA.

Figure 2.2.2: Comparison of income and wealth from SCF+, NIPA, and FA data

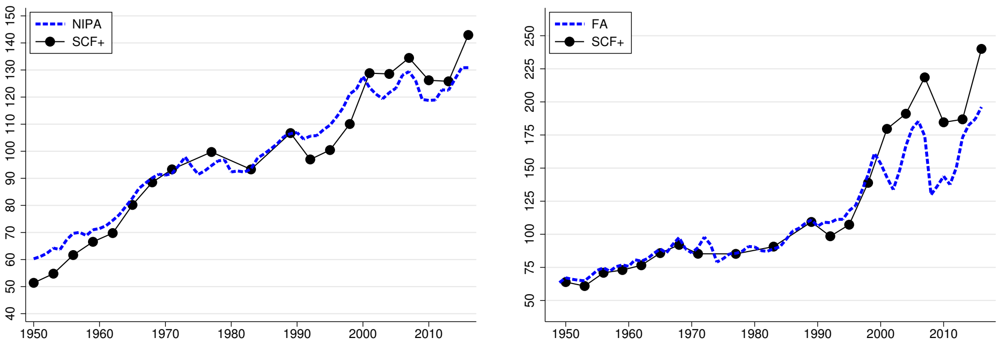

(a) Income (b) Wealth

Notes: Income and wealth data from SCF+ in comparison to income data from NIPA and wealth data from FA. All data have been indexed to the 1983-1989 period (= 100). SCF+ data are shown as black lines with circles, NIPA and FA data as a blue dashed lines. For the indexing period, SCF+ data correspond to 87% of NIPA income and 90% of FA wealth.

Figure 3.2.1a shows that the trend in income is very similar for SCF+ and NIPA data throughout the 1949-2016 time period. For the base period of 1983-1989, the SCF+

[10] We exclude the survey years 1947, 1948, 1952, 1961, 1964 and 1966 because we lack information on housing, mortgages, and liquid assets. These three wealth components are held by a large fraction of households but can only be poorly inferred from information on other variables. For 1977, we impute income using original data for income intervals.

matches 87% of income from the NIPA. Looking at wealth, the trends differ only slightly. Before 1995, wealth trends from the SCF+ and FA hardly differ. There appears to be a persistent level shift in the late 1990s that Henriques and Hsu (2014) trace back to differences in business wealth and owner-occupied houses.

Figure 2.2.3: Comparison of asset and debt components from SCF+ and FA data

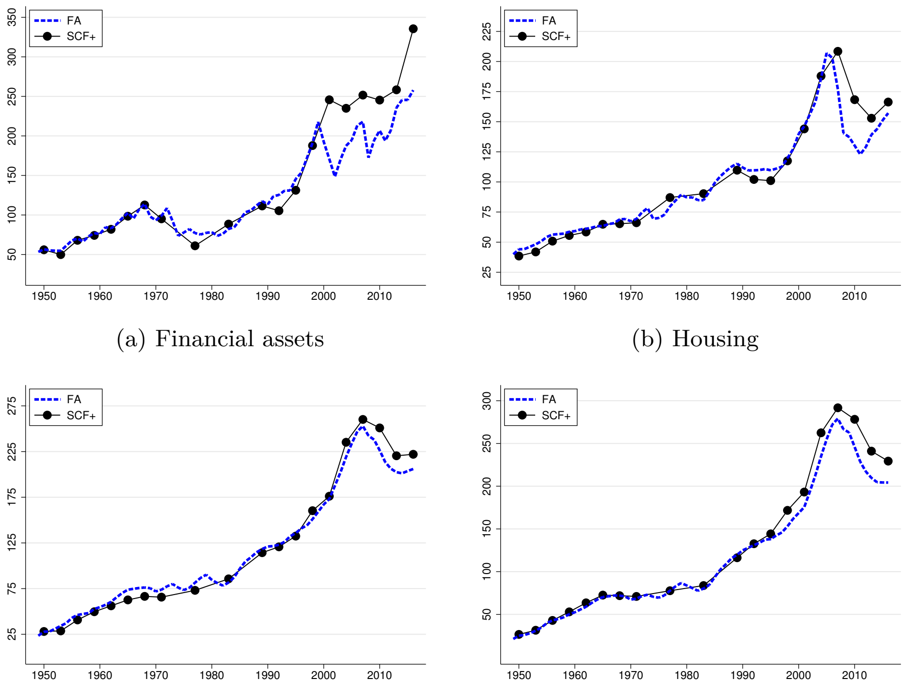

(a) Financial assets (b) Housing

(c) Total debt (d) Housing debt

Notes: Asset and debt components of household balance sheets from SCF+ and FA data. All data have been indexed to the 1983-1989 period (= 100). SCF+ data are shown as black lines with circles, FA data as a blue dashed lines. For the indexed period, SCF+ data correspond to 68% of financial assets, 98% of housing, 73% of total debt, and 75% of housing debt from the FA data.

Looking at different wealth components, we find that financial assets in the SCF+ (Figure 2.2.3a) increase more strongly in the early 2000s than the corresponding FA values. Henriques and Hsu (2014) attribute most of the difference to the coverage of retirement accounts in the SCF data. Figure 2.2.3b shows that housing as the most important non-financial asset is covered well in the survey data. Debt is the household balance sheet component for which the SCF+ matches the aggregate best, as shown in Figure 3.2.1b. Summing up, the SCF+ matches aggregate trends of NIPA data and FA asset and debt positions. In particular, the SCF+ data and the FA show similar trends for the important categories of housing wealth and mortgage debt.

2.3 Income and wealth inequality in the SCF+

This section presents stylized facts for long-run trends in income and wealth inequality that the SCF+ data bring to light. We begin by documenting the evolution of Gini coefficients for income and wealth as a comprehensive measure of inequality. We go on and look at the income and wealth inequality trends in different parts of the distribution. For this step, we will rely on the strength of the SCF+ data in covering the bottom 90% of the distribution. We also look at the long-run trends in the share of hand-to-mouth households and use the demographic information in the SCF+ data to analyze the importance of demographic factors in distributional change. Importantly, we present evidence on long-run trends in inequalities in income and wealth between black and white households.

2.3.1 Gini coefficients

The Gini coefficient is a comprehensive summary measure of inequality along the entire distribution. Table 2.3.1 reports Gini coefficients for income and wealth at selected points in time. The first row reports the Gini coefficient for all households; the other rows focus on the bottom 99% and the bottom 90%, respectively.[11]

Table 2.3.1: Gini coefficient (×100) for income and wealth

		1950	1971	1989	2007	2016
income	all	44	43	53	55	58
	bottom 99%	39	39	46	47	49
	bottom 90%	32	33	39	38	39
wealth	all	83	79	79	82	86
	bottom 99%	75	74	72	74	79
	bottom 90%	64	62	62	63	70

Notes: Gini coefficients for income and wealth for different years. All Gini coefficients are multiplied by 100. Survey years shown in columns. Upper part of the table shows Gini coefficients for income. First row shows Gini coefficients for all households, second row shows Gini coefficients for all households in the bottom 99% of the income distribution, and third row shows Gini coefficients for all households in the bottom 90% of the income distribution. Bottom part of the table shows Gini coefficients for wealth. First row shows Gini coefficients for all households, second row shows Gini coefficients for all households in the bottom 99% of the wealth distribution, and third row shows Gini coefficients for all households in the bottom 90% of the wealth distribution.

The Gini coefficients show that income and wealth inequality has increased not only across the entire population (across all households) but also among the bottom 99% and bottom 90% of households. The overall income Gini has risen from its postwar low of 0.43 in 1971 to 0.58 in 2016 (Figure 2.3.1a). Unsurprisingly, there is a substantial drop in inequality once we look at the bottom 99% of the distribution, but the increase in the Gini coefficient is still substantial. Also, within the bottom 90% income inequality has

[11] We report the full time series in Table 2.A.5 in the Appendix. We include negative-wealth households in the calculation. Figure 2.A.9 of the Appendix shows that time trends are very similar when we restrict the analysis to positive income and wealth households.

widened, yet this has mainly occurred between 1971 and 1989. In Section 2.3.3 below, we explore in detail changes within the bottom 90% over time.[12]

Turning to wealth, it is well known that wealth is considerably more unequally distributed than income. The wealth Gini has fluctuated around 0.8 for most of the postwar period and did not change much, if at all, between 1950 and 2007 (Figure 2.3.1b). By 2007, it stood at 0.82 and was only marginally higher than in both 1950 and 1971. The marked decline in the wealth Gini between 1971 and 1977 stands out. We will trace this decline back to asset price shifts in Section 2.4.3 below. A substantial increase in the Gini coefficient occurred between 2007 and 2016, and the wealth Gini reached its postwar peak in 2016.

The confidence bands in Figure 2.3.1 also show that Gini coefficients for both income and wealth are tightly estimated, although the confidence bands are somewhat wider in the historical data.[13] The observed long-run trends are clearly statistically significant. America is considerably more unequal today than it was in the 1970s, with respect to both income and wealth.

Figure 2.3.1: Gini coefficients for income and wealth with confidence bands

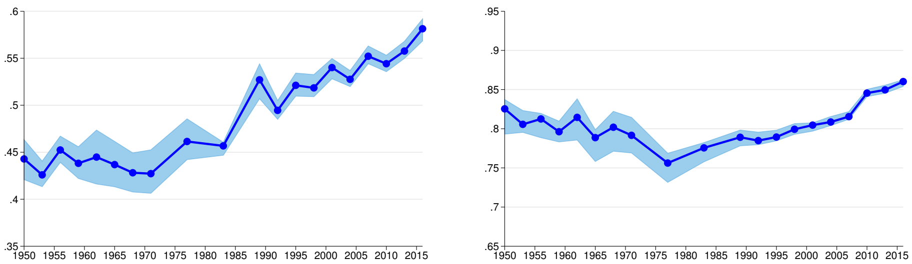

(a) Income (b) Wealth

Notes: Gini coefficients of income (panel (a)) and wealth (panel (b)) with 90% confidence bands. Confidence bands are shown as light blue areas. Confidence bands are bootstrapped using 999 different replicate weights constructed from a geographically stratified sample of the final dataset.

2.3.2 Income and wealth shares

Decomposing inequality trends, we start with an exploration of changes in income and wealth shares at the top, following the recent literature.[14] Broadly speaking, the SCF+ data corroborate the trajectories of the U.S. income and wealth distribution that emerged

[12] Our baseline income does not include rental income of owner-occupiers. As a sensitivity check, we imputed this rental income using historical rental yields from Jordà et al. (forthcoming) in Appendix 2.A.4.2. We find the Gini coefficient for income after imputing rents to be slightly lower.

[13] All confidence bands are computed using 999 replicate sample weights. Replicate weights are provided for the modern SCF surveys after 1983. For the historical surveys, we construct comparable 999 replicate weights (see Appendix 2.A.1.2).

[14] We follow the recent literature in considering synthetic income and wealth groups. Households from a group need not be the same across surveys due to mobility. Exploring wealth mobility using data from the Panel Study of Income Dynamics (PSID), we find high persistence of households within wealth groups. More than 70% of households typically remain within wealth groups between survey dates. We report detailed results in Online Appendix 2.A.3.

from the well-known studies by Piketty and Saez (2003) and Saez and Zucman (2016).

Figure 2.3.2a compares the income shares of the top 10% and 5% of the income distribution in the SCF+ to those calculated by Piketty and Saez (2003) using IRS data.[15] On the right-hand side, Figure 2.3.2b compares top wealth shares from the SCF+ with those from Saez and Zucman (2016). Figure 2.3.2 also shows estimated 90% confidence bands resulting from sampling error in the SCF+ data for the top income and wealth shares. The confidence bands underscore that the reported increases in income and wealth inequality are statistically significant. Despite some minor discrepancies, the SCF+ and tax data align in levels and trends of inequality so that they tell a similar story about the long-run trajectory of wealth and income inequality in postwar America.

The 1962 SFCC constitutes an alternative data point to our SCF+ data for that year. The survey shows a nearly identical top 10% wealth share and a somewhat lower top 10% income share. This implies that the increase in income concentration at the top since the 1960s is even stronger when using the 1962 SFCC datapoint. However, the income tax data are actually much closer to the SCF+ data and also show a higher top 10% share than the SFCC. For consistency reasons, we use the SCF+ data throughout.

Figure 2.3.2: Top 5% and top 10% income and wealth shares

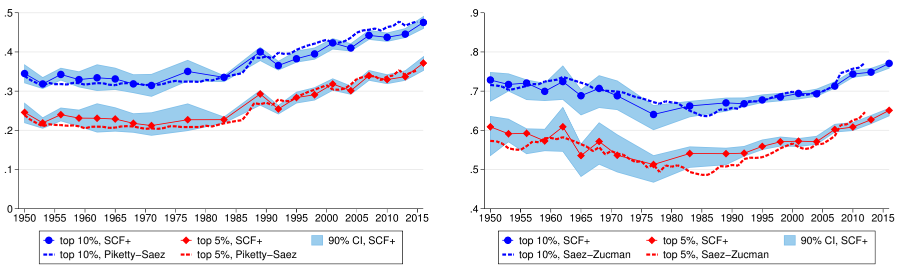

(a) Income (b) Wealth

Notes: Top 5% and top 10% income and wealth shares from SCF+ data, Piketty and Saez (2003), and Saez and Zucman (2016). Left panel shows income shares. Blue dots show top 10% income shares from SCF+ data, blue dashed line top 10% income shares from Piketty and Saez (2003) using IRS tax data. Red diamonds show top 5% income shares from SCF+ data, red dashed line top 5% income shares from Piketty and Saez (2003). Light blue areas show 90% confidence bands for SCF+ estimates. Right panel shows wealth shares. Blue dots show top 10% wealth shares from SCF+ data, blue dashed line top 10% wealth shares from Saez and Zucman (2016) using IRS data and capitalization method. Red diamonds show top 5% wealth shares from SCF+ data, red dashed line top 5% wealth shares from Saez and Zucman (2016). Light blue areas show 90% confidence bands for SCF+ estimates.

In the next step of our decomposition, we move down the distribution and turn to the evolution of income and wealth among the bottom 90% (Table 2.3.2). The mirror image of increasing concentration of income in the hands of the top 10% must, by definition, be (relative) income losses among the bottom 90%. But which strata of the bottom 90% were hit particularly hard by the growing income share of the top 10%?

Table 2.3.2 reports the income shares of different groups of the income distribution and

[15] Piketty and Saez (2003) include salaries and wages, small business and farm income, partnership and fiduciary income, dividends, interest, rents, royalties, and other small items reported as other income.

wealth shares of different strata of the wealth distribution.[16] Starting with income on the left, the SCF+ shows that the top 10% have grown their income share by close to 15 percentage points from 34.5% to 47.2% between 1950 and 2016. The income share of the bottom 50% of Americans has fallen by roughly a third from 21.6% to 14.6%, and middle-class households (50th to 90th percentiles) have lost about 6 percentage points in income shares. In other words, we do observe a hollowing out of middle-class America, with households around the median having witnessed the largest relative income losses.

Table 2.3.2: Shares in aggregate income and wealth

	Income					Wealth				
	1950	1971	1989	2007	2016	1950	1971	1989	2007	2016
bottom 50%	21.6	21.6	16.3	15.5	14.6	2.3	3.3	3.0	2.5	1.1
0%- 25%	6.0	6.2	5.0	4.5	4.5	-0.4	0.0	-0.1	-0.1	-0.5
25%-50%	15.6	15.4	11.3	11.0	10.1	2.7	3.4	3.1	2.6	1.6
50%-90%	43.9	47.0	43.7	40.3	37.9	24.8	27.9	30.0	26.2	21.8
50%-75%	23.5	25.0	22.4	20.3	18.4	9.8	10.7	11.9	10.3	7.4
75%-90%	20.4	22.0	21.3	20.0	19.5	15.0	17.2	18.1	15.9	14.4
top 10%	34.5	31.4	40.0	44.2	47.5	72.8	68.8	67.0	71.3	77.1

Notes: Shares in aggregate income and wealth in different years. First column shows household groups of the income and wealth distribution. Left part of the table shows the share in total income for household group from the first column of the table. Households are sorted by income. Shares of bottom 50%, 50%-90%, and top 10% add to 100%. Shares for bottom 50% and 50%-90% are further split into subgroups. Columns show shares for different years. Right part of the table shows corresponding wealth shares across household groups. Households for this part of the table are sorted by wealth.

The right side of Table 2.3.2 studies the change in wealth shares (households are now stratified by wealth). The main insight here is that until the 2008 financial crisis, changes in wealth shares were modest. If anything, the bottom 90% wealth share was slightly higher in 2007 than it was in 1950, and very close to its 1971 level. In contrast to the observed changes in the income distribution, middle-class households managed to maintain their wealth shares until the mid-2000s. The 50%-90% wealth share was higher in 2007 than in 1950, and only slightly lower than in 1989. It is equally clear that the financial crisis had a substantial effect on the wealth distribution. Middle-class wealth shares collapsed across the board, while the wealth share of the top 10% surged by 6 percentage points within less than a decade. The decade since the financial crises witnessed the largest spike in wealth concentration in postwar America.

The overall outcome was a more pronounced shift in the income distribution than in the wealth distribution since the 1970s. We return to this important fact in Section 2.4. In

[16] Appendix 2.A.5.3 reports the full time series.

the next step, we zoom in on the bottom 90% and study long-run distributional trends in the lower parts of the distribution, as well as low and negative-wealth households.

2.3.3 The bottom 90%

Much of recent research on seminal trends in inequality has focused on developments at the very top of the distribution. This emphasis on the top 1% (and beyond) plays to the strength of the tax data that were, at least so far, the only source spanning the postwar decades on a continuous basis. However, the tax data can only provide a relatively coarse picture of developments in the lower parts of the distribution. The SCF+ fills this gap.

We start the analysis with income and wealth trends for percentiles across the bottom 90%. Figure 2.3.3a documents that income grew at a similar rate across the 25th, 50th, and 75th percentiles in the first two postwar decades. From the 1970s to the 1990s, the 25th and 50th percentile experienced real income losses while incomes at the 75th percentile stagnated. All groups saw a return to real income growth from the mid-1990s to mid-2000s, but only incomes at the 75th percentile have recovered from the 2008 crisis.

Looking at percentile ratios in Figure 2.3.3b, we see that since 1980 income at the median evolved similarly to income at the 25th percentile so that the 50-25 ratio did not change much over the last four decades. By contrast, since the 1970s the 75th and 90th percentile left the median behind leading to a pronounced widening of the 75-50 and 90-50 percentile ratios.

Figure 2.3.3c presents the same analysis for wealth. The picture is markedly different. First, wealth between groups started to persistently diverge only in the 2000s, not in 1970s as in the case of income. Second, households at all three percentiles saw major wealth drops after 2007, but there was considerable variation. The outcome is a substantial polarization of wealth and pronounced widening of the 90-50 and 50-25 ratios. The Figure also shows that nearly all wealth gains that households at the 25th percentile had made since 1971 have been wiped out by the crisis.

2.3.3.1 Low- and negative-wealth households

Low- and negative-wealth households (net debtors) are key groups when it comes to the consequences of wealth inequality for macroeconomic dynamics (Krusell and Smith, 1998). Using the SCF+ data, we show in Figure 2.3.4 how the shares of low- and negative-wealth households evolved over the last seven decades. The share of net debtors has doubled from its low of the 1980s, but remains within its postwar range that fluctuated between 5% and 12% (see Figure 2.3.4a). Starting in the 1980s, the average debt of net debtors increased from slightly less than 60% of average annual income in the period from 1950 to 1977 to over 140% in 2010. In 2016, the average debt balance of

Figure 2.3.3: Percentile growth and percentile ratios for income and wealth

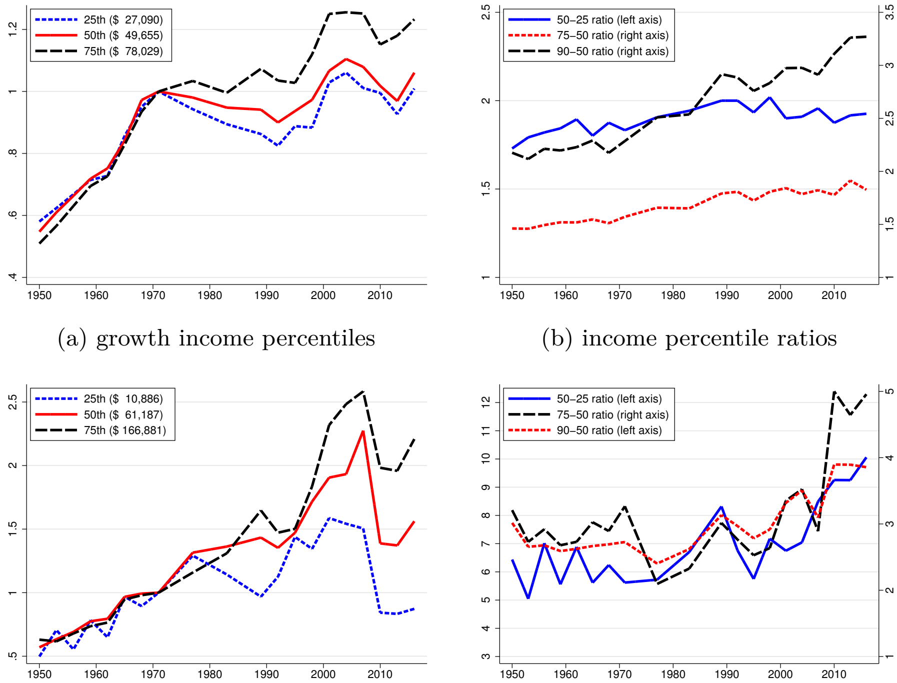

(a) growth income percentiles

(b) income percentile ratios

(c) growth wealth percentiles

(d) wealth percentile ratios

Notes: Top left panel shows growth of the 25th, 50th, and 75th percentile of income relative to 1971 (= 1). Level of percentiles in 1971 are shown in legend (2016 dollars). Top right panel shows 90-50, 75-50, and 50-25 percentile ratios for income. Bottom left panel shows growth of the 25th, 50th, and 75th percentile of wealth relative to 1971 (= 1). Level of percentiles in 1971 are shown in legend (2016 dollars). Bottom right panel shows 90-50, 75-50, and 50-25 percentile ratios for wealth.

net debtors corresponds to 113% of their average income.[17] A broader measure of low-wealth households includes all households that have positive wealth but whose wealth is low relative to their income. We use a threshold of three months of income, implying a wealth-to-income ratio of 0.25 or below. This group can self-insure only to a limited extent by accessing savings, for instance in the case of a job loss. The share of this group is large: close to one quarter of American households are low-wealth households according to this definition. The share of these households has risen since the crisis, but remains within its postwar range. One reason why households have negative wealth is negative home equity, and Figure 2.3.4b reports the share of homeowners among negative-wealth households. The ratio reached its all-time high in 2007 when house prices collapsed and highly leveraged households ended up *under water*.

2.3.3.2 Wealthy hand-to-mouth households

Kaplan and Violante (2014) argue that the group of households that behave like hand-to-mouth consumers, i.e., as if they had no wealth for consumption smoothing, is much larger as many households hold wealth in illiquid assets that cannot be easily accessed.

[17] Results on average debt of net debtors are available from the authors upon request.

Figure 2.3.4: Low-wealth households

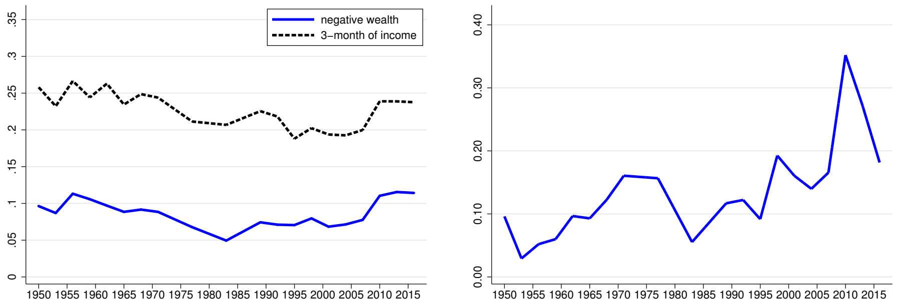

(a) low-wealth households

(b) share of homeowners among negative-wealth households

Notes: Share of low-wealth households and share of homeowners among negative-wealth households (*net debtors*). Left panel shows shares for two measures of low-wealth households: Black dashed line shows share of households with wealth less than three months of income ($\frac{3}{12}$ of annual income). Blue solid line shows share of negative-wealth households. Right panel shows the share of homeowners among households with negative wealth.

Kaplan and Violante (2014) coined the term *wealthy hand-to-mouth households* and documented that from 1989 to 2010 about one in three American households can be classified as "hand-to-mouth" and that about two-thirds of these households are wealthy hand-to-mouth consumers.

Using the SCF+, we can provide estimates for the share of wealthy hand-to-mouth consumers for the entire postwar period. We follow Kaplan et al. (2014) in identifying hand-to-mouth households in the data and relegate details to Appendix 2.A.1.7. We also provide estimates for cash holdings of households going back until 1973 using data from the National Crime Victimization Survey (NCVS).[18]

Figure 2.3.5 provides two different estimates of the share of hand-to-mouth and wealthy hand-to-mouth households in the United States over the post-WW2 period.[19] Figure 2.3.5a shows the baseline estimate following the approach for cash holdings in Kaplan et al. (2014). It shows a slight downward trend for hand-to-mouth households over the time period of the historical data and a rising proportion of wealthy-hand-to-mouth households. Figure 2.3.5b provides estimates for hand-to-mouth households using our estimates for cash holdings from the NCVS data. Both Figures show a relatively stable ratio of wealthy hand-to-mouth households since 1950, albeit with some variation over shorter horizons.

[18] Although the survey is designed to collect data on victims of crime, it also records details of the incidence, including theft. Online Appendix 2.A.1.7 provides details on the construction of estimates and Figure 2.A.2 shows cash estimates as fraction of median SCF+ income.

[19] We do not provide estimates for 1977 because income in 1977 is reported in intervals so that the share of hand-to-mouth is estimated imprecisely.

Figure 2.3.5: Shares of poor and wealthy hand-to-mouth households

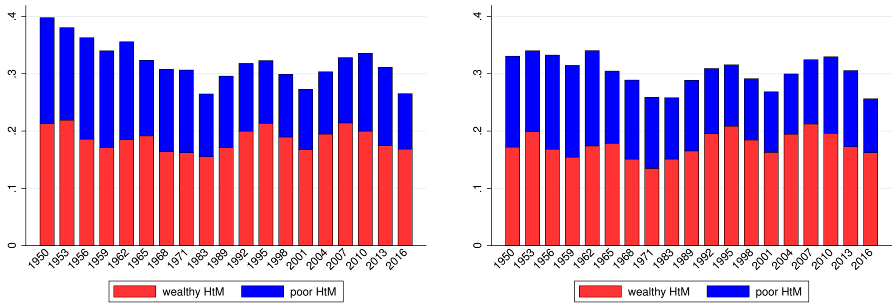

(a) baseline cash estimate

(b) alternative cash estimate

Notes: Shares of poor and wealthy hand-to-mouth households (HtM) for two alternative cash estimates. Bars show total share of hand-to-mouth households in the population. Red bars indicate share of wealthy hand-to-mouth households, blue bars the share of poor hand-to-mouth households. Panel (a) shows estimated shares using the cash estimates following Kaplan et al. (2014). Panel (b) shows estimated shares using cash estimates based on NCVS data. See text for details. Estimates for 1977 are omitted due to data limitations (see footnote 19).

2.3.4 Demographic change

What were the effects of secular changes in terms of educational attainment, age structure, and household size of the U.S. population on income and wealth inequality? Using the demographic information in the SCF+, we provide answers to these questions. In a first step, we implement an approach proposed by Fortin et al. (2011) to remove changes in the age structure and educational attainment over time.[20] In a second step, we account for changes in household size by adjusting income and wealth at the household level to per-adult equivalents using the OECD equivalent scale.[21]

Figure 2.3.6a shows Gini coefficients for the original data and the two counterfactual cases where we add the marginal effects from fixing educational attainment and the age structure at the 1971 distributions. The effect from shifts towards more highly educated household heads on income appear rather small, but the effects coming from an older population are more sizable. Note that this finding is in line with a rising college wage premium as we only consider the effect from changes in quantities (number of households) not prices (wages). In the case of wealth (Figure 2.3.6b), the effect of changing educational attainment and aging are small.[22] All in all, demographic changes have some effects, but do not change the overall pattern of income and wealth inequality in the United States since World War II.

A second secular trend in the United States has been the decrease in average house-

[20] We use 1971 as our base year for which we fix the distribution of demographic characteristics. We then estimate a probit model including age, educational attainment, the number of adults and children in a household, and the race of the household head as controls to derive adjustment weights. We relegate a detailed description to Appendix 2.A.1.6.

[21] The OECD equivalence scale assigns a value of 1 to the first household member, 0.7 to each additional adult, and 0.5 to each child (see OECD `http://www.oecd.org/eco/growth/OECD-Note-EquivalenceScales.pdf`).

[22] Bartscher et al. (2018) provide a detailed analysis on the trends in the financial situation of college and non-college households in the United States based on the SCF+ data.

Figure 2.3.6: Gini coefficients accounting for change in demographic composition

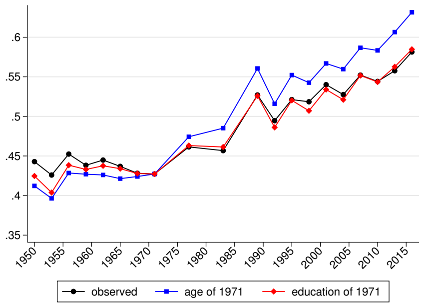

(a) income Gini coefficients

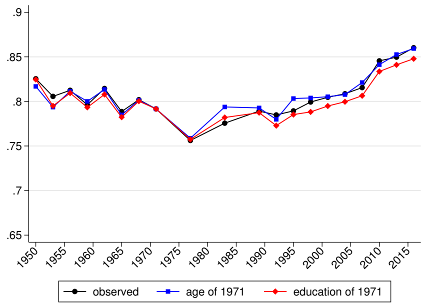

(b) wealth Gini coefficients

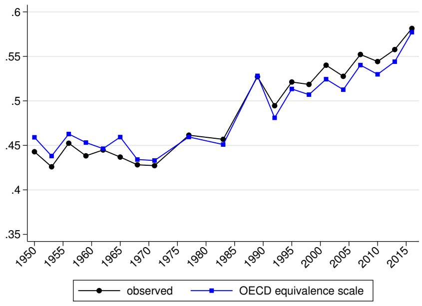

(c) income Gini coefficients

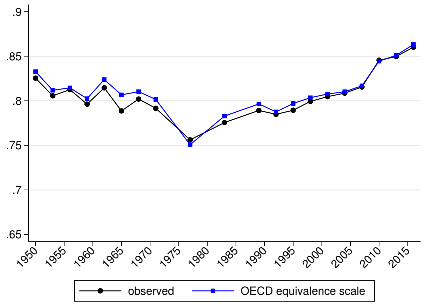

(d) wealth Gini coefficients

Notes: Gini coefficients for income and wealth accounting for changes in the age, education, and household size composition over time. Top left panel shows three time series for Gini coefficients of income. Black dots show the observed time series of Gini coefficients. Blue squares show counterfactual Gini coefficients for a constant 1971 age composition of households. Red diamonds show counterfactual Gini coefficients for a constant 1971 education composition of households. Age and education refer to the head of household. Top right panel shows the corresponding three time series for Gini coefficients of wealth. Bottom left panel shows two time series for Gini coefficient of income. Black dots show the observed time series of Gini coefficients. Blue squares show Gini coefficients after dividing income by household size using the OECD equivalence scale (see footnote 21). Bottom right panel shows the corresponding two time series for Gini coefficients of wealth.

hold size from an average of 3.4 household members in 1949 to an average of 2.5 in 2013 according to Census data. Given that the SCF+ is a household survey, changes in household size can potentially affect measures of household-level inequality. We adjust income and wealth to per-adult-equivalent member of the household with the OECD equivalence scale. Figure 2.3.6c reports that income concentration at the top falls somewhat when we look at adult-equivalent income. This trend is consistent with stronger assortative mating and increasing female labor force participation. For wealth (Figure 2.3.6d), we do not observe big effects.

2.3.4.1 The persistence of racial disparities in income and wealth

Race is an important stratifying dimension of the U.S. population. In a recent paper, Bayer and Charles (2017) provide long-run evidence on the black-white earnings gap using data from the U.S. Census Bureau and the American Community Survey. They document persistent earnings differences for working-age men. The SCF+ data complement recent work on the long-run evolution of racial inequality along three dimensions.[23]

First, in addition to earnings, we study household income from all sources. Second, our unit of observation is the household, not working-age male individuals. We thus capture the effects of changing marriage patterns, higher labor force participation of women, as well as changes in transfers, education, and retirement decisions of households. Third, the SCF+ data also allow us to analyze the long-run evolution of *wealth* differentials between black and white households. So far, the racial wealth gap has remained uncharted territory as long-run data were simply not available. With data reaching back to the pre-civil rights era, our analysis extends recent work by Wolff (2017), who studied wealth differences between black, white, and Hispanic households in the modern SCF data starting in 1983. For the analysis, we group households into black and white households according to the race of the household head.[24]

Figure 2.3.7 shows the trends in income and wealth of the median household and of the household at the 90th percentile for both white and black households. The racial divide will fall if black households' income or wealth increases more strongly over time. A lockstep evolution of the series for black and white households (equal growth rates) signals persistence of existing racial disparities.

Three facts stand out. First, income has grown at a comparable rate for black and white households. This means that pre-civil rights era disparities have largely persisted as black income growth did not accelerate relative to white households. Second, as the numbers indicate, the size of the racial income divide remains substantial. The median black household has about half of the income of the median white household. Third, the wealth gap is much larger than the income gap and equally persistent. The median black household disposes of 12% of the wealth of a median white household. In the 1980s, the wealth of the median black household stood at about $13,000 in 2016 prices — equivalent to the value of a car. The median white household had about $115,000 — corresponding to the value of a small house.

Looking at the time trends in more detail, we note two periods when the racial disparities narrowed temporarily. In the 1970s, the income of the median black household grew about 20% faster than the income of the median white household. However, the trend reversed in the 1980s when the share of black households headed by women increased strongly.[25] The 2000s are the second period in which the racial income gap narrowed somewhat for the median household.

Figure 2.3.7b exposes an equally persistent racial wealth gap. The difference in wealth

[23] Thompson and Suarez (2017) and Dettling et al. (2017) analyze racial inequality using SCF data.

[24] The number of interracial marriages is growing but remains small. Fryer (2007) reports that for whites about 1% of marriages were interracial and about 5% for black Americans. We drop all other racial categories. The survey questions for race in the SCF changed little over time. An important change happened in 1989, when the information was obtained as part of the interview rather than coded directly by the interviewer.

[25] When adjusting incomes for household size, the decline in relative incomes for black households during the 1980s becomes less pronounced.

Figure 2.3.7: Income and wealth trends for black and white households

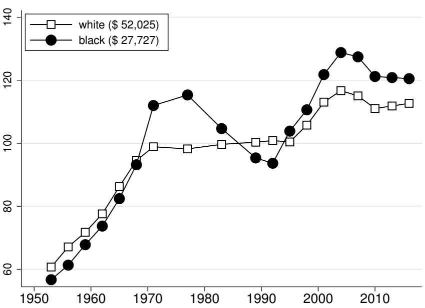

(a) Median income

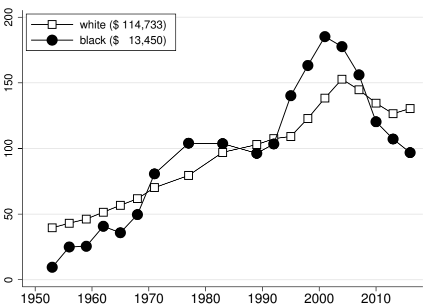

(b) Median wealth

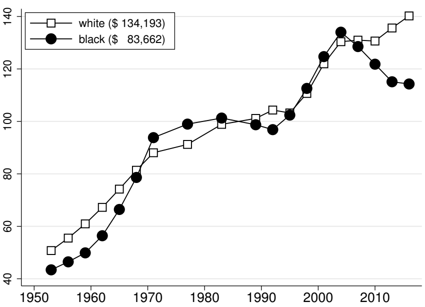

(c) Income 90th percentile

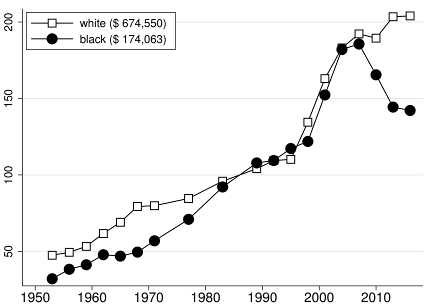

(d) Wealth 90th percentile

Notes: Trends of medians and 90th percentiles of income and wealth for black and white households. Top left panel shows trends of median income for black and white households indexed to the period 1983-1989 (=100). Average income levels at the median over the indexing period are shown in the legend (2016 dollars). Top right panel shows the corresponding time series for median wealth. Bottom left panel shows trends at the 90th percentile of income for black and white households. Time series are indexed to the period 1983-1989 (=100). Average income levels at the 90th percentile over the indexing period are shown in the legend (2016 dollars). All time series show moving averages over three neighboring observations. Medians and 90th percentile always refer to the respective income and wealth distribution of black and white households.

narrowed temporarily in the housing boom of the 1990s and early 2000s, but widened again after the financial crisis. After 2007, the wealth levels of households at the 90th percentile of the black wealth distribution collapsed, while the 90th percentile of the white wealth distribution remained largely unaffected.

As an alternative to study the evolution of earnings differences over time, Bayer and Charles (2017) apply the concept of a racial “rank gap”. Adapted for wealth, the rank gap is the percentage point difference between the rank of a given percentile in the black and white wealth distribution. For instance, a number of -30 for the median of the black wealth distribution means that the place of that household would be 30 percentage points lower in the wealth distribution of white households, that is, only at the 20th percentile.

Figure 2.3.8a shows the wealth rank gap at the median and the 90th percentile. For the median, the long-run average is close to -30, implying that the median black household is at the 20th percentile of the wealth distribution of white households. Put differently, the typical black household is poorer than 80% of white households. The rank gap fluctuates,

Figure 2.3.8: Racial rank gaps for wealth

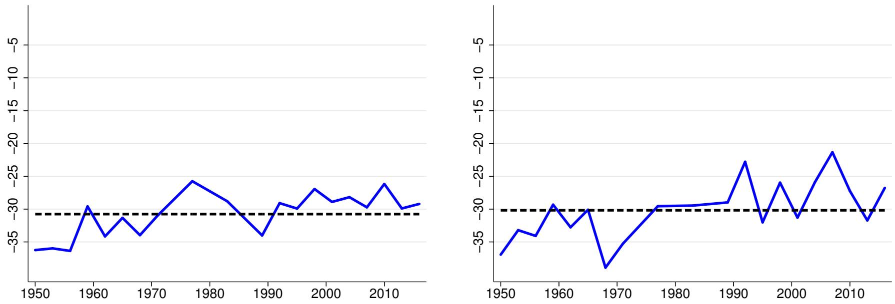

(a) median (b) 90th percentile

Notes: Racial rank gaps for wealth at the median and 90th percentile. Left panel shows the racial rank gap at the median. The racial rank gap is the difference in percentage points between the rank that the wealth level of the median black household takes in the wealth distribution of white households and the rank of the median white household. Dashed line shows the long-run average of the racial wealth rank gap. Right panel shows the corresponding racial rank gap at the 90th percentile. Dashed line shows again the long-run average of the racial wealth rank gap.

tracking what we have seen for levels in Figure 2.3.7b, but is highly persistent over time. We find an equally large and persistent rank gap at the 90th percentile of the wealth distribution. Our main conclusion is that virtually no progress has been made over the past 70 years in reducing wealth inequality between black and white households.

2.4 Asset prices and the wealth distribution

In the previous section, we discussed changes in the income and wealth distributions separately, as in the existing literature. Yet it is precisely the link between the income and wealth distributions that plays a central role in theoretical models of wealth inequality. A central advantage of the SCF+ is that it allows us to study the long-run evolution of the *joint* distribution of income and wealth. This topic is what we turn to now.

In the simplest model of the dynamics of the wealth distribution, changes in the income and wealth distributions are closely linked. With saving rates that are constant over time and uniform across wealth classes, and uniform returns on wealth along the wealth distribution, changes in the wealth distribution would be solely driven by changes in the income distribution. Or, put differently, the differential growth rates of wealth would be a function of the differential growth rates of income. Recent studies have questioned this assumption, as models based on labor income risk typically produce too little wealth concentration at the top and cannot account for substantial shifts in wealth inequality that occur over short time horizons (Benhabib and Bisin (2018), Gabaix et al. (2016), Hubmer et al. (2017)).

As a first check, in Figure 2.4.1 we compare the time path of income and wealth growth in the United States since 1971. Note that we stratify all households by wealth and index income and wealth levels to 1 in 1971. Figure 2.4.1a highlights a substantial divergence in income growth for different groups of the wealth distribution. Income growth was low for the bottom 90% and particularly meager for households in the lower half. For the bottom 50%, real incomes have stagnated since the 1970s. For households between the 50th and 90th percentiles of the wealth distribution, real incomes rose modestly by about a third over nearly 40 years, implying annual growth rates of much less than 1% per year. By contrast, income growth at the top was strong. The incomes of households within the top 10% of the wealth distribution have doubled between 1971 and 2007.[26]

Figure 2.4.1: Income and wealth growth along the wealth distribution

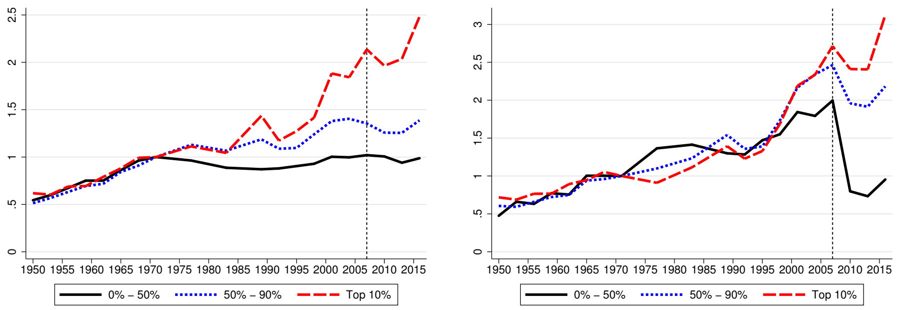

(a) Income growth

(b) Wealth growth

Notes: Income and wealth growth for different groups along the wealth distribution. Left panel shows income growth for three groups in the wealth distribution: the bottom 50% (black solid line), the middle class 50%-90% (blue dotted line), and the top 10% (red dashed line). All income time series are indexed to 1 in 1971. Right panel shows wealth growth for the same three groups along the wealth distribution. The vertical lines in both panels indicate the 2007 survey.

Yet when we turn to wealth growth for the same groups in Figure 2.4.1b, the contrast is stark. From 1971 to 2007 (the last pre-crisis survey), wealth growth has been, by and large, identical for the top 10% and the bottom 90% of the wealth distribution. More precisely, middle-class (50%-90%) wealth increased by 140% the same rate as top 10% wealth. And even the bottom 50% did not do too badly when it comes to wealth growth, as their wealth still doubled between 1971 and 2007. Wealth and income growth rates have decoupled over an extended period, in marked contrast to the simple model sketched above. We will return to this point below.

Figure 2.4.1b also shows how devastating the 2007-2008 financial crash was to lower middle-class wealth, while the impact of the crisis on wealth at the top was rather minor. By 2013, the absolute level of real wealth below the median household was 20% below its 1971 level. Within a few years, the crisis wiped out all gains in household wealth that the bottom 50% of the distribution had made over the preceding four decades. As of 2016, still close to half of the American population dispose of less wealth in real dollar amounts than in 1971.[27]

One upshot is that before the financial crisis, wealth-to-income ratios increased most

[26] Online Appendix 2.A.5.3 reports income shares for households along the wealth distribution comparable to the income shares along the income distribution in Table 2.3.2.

[27] We provide a sensitivity analysis to these growth trends when using the PCE instead of the CPI for inflation adjustment in Appendix 2.A.4.4. See footnote 7 for further discussion.

strongly in the middle and at the bottom of the wealth distribution and then also fell most strongly for those groups during the crash. Figure 2.4.2 illustrates this phenomenon. Figure 2.4.2a shows the strong increase in wealth-to-income ratios for the bottom 50% (left axis) and the middle class (right axis) until 2007, followed by a substantial decline. Figure 2.4.2b shows wealth-to-income ratios for the top 10% that are much higher on average. Wealth-to-income ratios increased only slightly between 1971 and 2007 at the top and hardly changed after 2007.

Figure 2.4.2: Wealth-to-income ratios by wealth group, 1950-2016

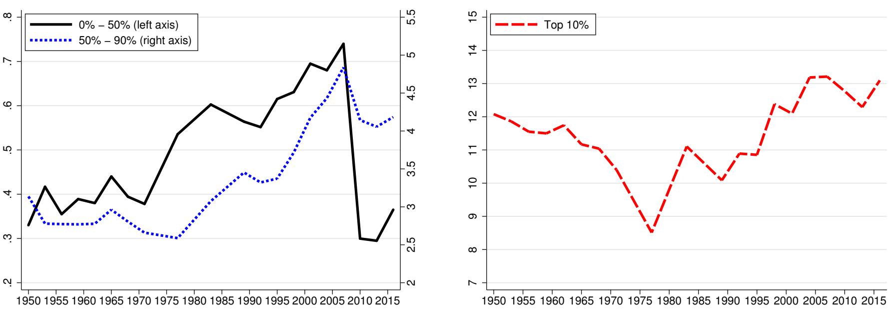

(a) Wealth-to-income ratios bottom 90% (b) Wealth-to-income ratios top 10%

Notes: Wealth-to-income ratios by wealth groups. Left panel shows wealth-to-income ratios for the bottom 50% of the wealth distribution (black solid line) and for the 50%-90% of the wealth distribution (blue dotted line). Right panel shows wealth-to-income ratio for the top 10% of the wealth distribution. Wealth-to-income ratios are constructed as ratio of averages within wealth groups.

2.4.1 The dynamics of the wealth distribution

If we are to understand the dynamics of the wealth distribution in America over the past seven decades, we must look beyond income growth. In the following, we demonstrate that asset price changes played an important role in the observed dynamics of wealth inequality in postwar America.[28]

Asset prices affect the dynamics of the wealth distribution through two channels. First, asset prices lead to differential capital gains if portfolios differ across the distribution. We document this important stylized fact for the U.S. economy below. As changes in asset prices revalue existing wealth, they can induce shifts in wealth shares that are unrelated to income changes. Moreover, they can do so over short horizons as they immediately affect the value of accumulated assets. We will document that in America, persistent differences in portfolio composition between middle-class households and rich households essentially give rise to a race between the stock market and the housing market in shaping the dynamics of the wealth distribution.

The second channel through which asset prices matter for the dynamics of wealth inequal-

[28]For the French case, Garbinti et al. (2017) show that price effects played an important role in shaping the French wealth distribution over the past 200 years. In the American context, Saez and Zucman (2016) discuss that price effects can change inequality trends relative to those implied by income and saving rate differences but focus on saving rate differences in their discussion.

ity is through their effect on wealth-to-income ratios. The level of the wealth-to-income ratio determines the relative importance of savings flows for wealth dynamics. When wealth-to-income ratios are high, income growth and savings flows become relatively less important for the wealth distribution, simply because the stock of wealth is high relative to income flows. This second channel is particularly relevant for the time period from 1970 to 2007, when the aggregate wealth-to-income ratio increased from 4 to more than 7, driven by rising house prices and a booming stock market.

To see how asset prices affect wealth inequality when portfolios are heterogeneous, consider a household i in period t with a portfolio of assets $\{A^i_{j,t}\}^J_{j=1}$, for instance, houses, stocks, and saving accounts. For household i, the capital gain Π^i_t from asset price changes between period t and $t+1$ is the asset-weighted average of price changes

$$\Pi^i_t = \sum_{j=1}^{J} \left(\frac{p_{j,t+1}}{p_{j,t}} - 1 \right) A^i_{j,t},$$

where $p_{j,t}$ denotes a (real) price index for asset j in period t. Denote the household's wealth in t by W^i_t and divide both sides of the equation by wealth to get

$$\begin{aligned} \frac{\Pi^i_t}{W^i_t} &= \sum_{j=1}^{J} \left(\frac{p_{j,t+1}}{p_{j,t}} - 1 \right) \frac{A^i_{j,t}}{W^i_t} \\ q^i_t &= \sum_{j=1}^{J} \left(\frac{p_{j,t+1}}{p_{j,t}} - 1 \right) \alpha^i_{j,t}, \end{aligned} \tag{2.4.1}$$

where $\alpha^i_{j,t}$ denotes the portfolio share $\frac{A^i_{j,t}}{W^i_t}$ of asset j for household i in period t and q^i_t is the growth rate of household wealth from capital gains. Equation (2.4.1) shows that portfolio differences (i.e., differences in $\alpha^i_{j,t}$ across households) lead to differences in capital gains q^i_t.

To fix ideas and structure the discussion about how this affects the wealth distribution, we rely on an illuminating accounting framework adapted from Saez and Zucman (2016).[29] Consider a simplified law of motion for wealth of household i:

$$W^i_{t+1} = W^i_t(1 + r^i_t + q^i_t) + Y^i_t - C^i_t$$

where r^i_t are returns on wealth other than capital gains (e.g., dividends), Y^i_t denotes income from all other sources, and C^i_t denotes consumption.[30] The savings flow S^i_t of household i in period t corresponds to total income net of consumption $S^i_t = r^i_t W^i_t + Y^i_t - C^i_t$. Define further the saving rate s^i_t as $s^i_t = \frac{S^i_t}{Y^i_t}$, so that the law of motion for wealth becomes

$$W^i_{t+1} = W^i_t(1 + q^i_t) + S^i_t = W^i_t(1 + q^i_t) + s^i_t Y^i_t = (1 + q^i_t + \sigma^i_t) W^i_t \tag{2.4.2}$$

[29] A micro-founded analysis of household saving behavior and portfolio choice requires a more complex approach. Hubmer et al. (2017) discuss why such a framework remains beyond reach for the time being but must become a topic for future research.

[30] Income denotes income from all sources excluding capital. This approach simplifies wealth dynamics by abstracting from bequests and death, divorce, marriage or other life-cycle events that affect households' wealth accumulation. We adjust the framework by Saez and Zucman (2016) slightly with respect to the timing convention by assuming that capital gains accrue together with savings flows, but the underlying mechanism remains the same. Saez and Zucman (2016) focus on the heterogeneity in savings behavior and assume homogeneity of capital gains. Our discussion focuses on the comparison of the relative importance of savings and capital gains as drivers of wealth accumulation.

with $\sigma_t^i = \frac{s_t^i Y_t^i}{W_t^i}$ capturing the contribution of savings to wealth growth and the term q_t^i captures the effect of capital gains to wealth growth.

In the next step, we move from the law of motion for wealth *levels* to a law of motion for wealth *shares* of different wealth strata. We construct "synthetic" saving flows and capital gains for specific wealth groups, again taking the lead from Saez and Zucman (2016). Savings flows and capital gains for wealth groups are "synthetic" in the sense that they assume that households stay in their wealth group from one period to the next. Using PSID data, we show in Online Appendix 2.A.3 that while there is some mobility in practice, the synthetic method yields a good approximation of wealth dynamics.

Note that we will now use i to refer to the group of households in a specific wealth stratum. The wealth share of group i in period t is $\omega_t^i = \frac{W_t^i}{W_t}$ where W_t is aggregate wealth in period t. All aggregate variables are defined according to group-level variables; for example, the aggregate savings rate is $s_t = \frac{S_t}{Y_t}$. Applying some straightforward transformations to equation (2.4.2) yields the law of motion for the wealth share ω_t^i:

$$\omega_{t+1}^i = \frac{1 + q_t^i + \sigma_t^i}{1 + q_t + \sigma_t} \omega_t^i \qquad \Longleftrightarrow \qquad \frac{\omega_{t+1}^i}{\omega_t^i} = \frac{1 + q_t^i + \sigma_t^i}{1 + q_t + \sigma_t}. \tag{2.4.3}$$

The law of motion has an intuitive interpretation: the wealth share of any group i increases if the group's wealth growth rate exceeds the average wealth growth rate in the economy. Differences in growth rates result from the two components of wealth growth in equation (2.4.2). First, group i's capital gains q_t^i can be higher (or lower) than the average capital gain q_t in the economy. Second, different rates of wealth growth can result from the difference between group i's savings component σ_t^i relative to the average savings σ_t. This is the channel through which differences in income growth translate into wealth inequality: higher incomes of group i will, all else equal, increase the group's saving flows and its wealth level relative to other groups in the economy.

The savings term σ_t^i comprises the inverse of the wealth-to-income ratio. With higher wealth-to-income ratios, the importance of savings flows declines and tends to zero. This implies that the relative importance of differences in savings flows diminishes with higher wealth-to-income ratios. This effect is independent of portfolio composition and is operative even when households have identical wealth portfolios (so that capital gains are identical). Also note that the two components q_t^i and σ_t^i are independent, so that a low savings component relative to the average can go hand-in-hand with a large capital gains component relative to the average for group i, and vice versa. This can decouple the evolution of the income and wealth distributions.

2.4.2 Portfolio heterogeneity and asset price exposures

If portfolios differ systematically along the wealth distribution, asset price changes will lead to differential capital gains along the wealth distribution. These in turn can induce changes in the wealth distribution that are unrelated to changes in the income distribution. The necessary condition for such effects is that portfolios are heterogeneous. For the first time, the SCF+ provides us with long-run balance sheet information to study the composition of household portfolios over a long time horizon. The evidence points to systematic and highly persistent differences in wealth portfolios across groups and hence a potentially important role for asset prices in shifting the wealth distribution, as we will now demonstrate.

Figure 2.4.3 displays the heterogeneity of household portfolios. It tracks the portfolio composition of the bottom 50%, the 50%-90%, and the top 10% of the wealth distribution since 1949. As a benchmark, we also track the average portfolio of the macroeconomy. In the figures, assets enter with positive values and debt as negative values. Household wealth corresponds to the consolidated value of all portfolio positions and is indicated by a dashed line in each of the figures. The degree of leverage in household portfolios can be inferred by looking at the sum of assets in excess of wealth. We provide time series in the Appendix 2.A.5.5.

Figure 2.4.3: Heterogeneity in household portfolios for different wealth groups

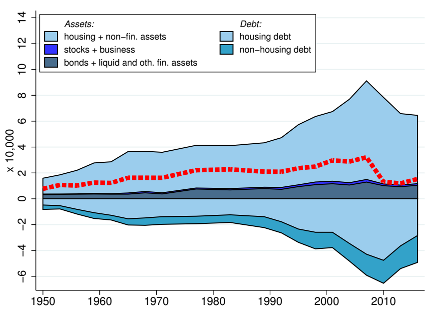

(a) bottom 50%

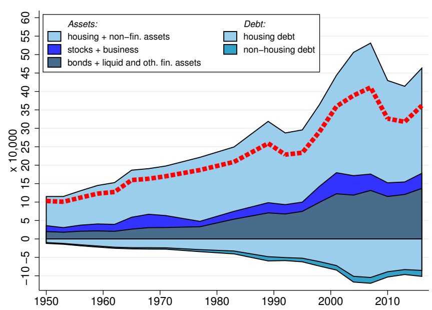

(b) 50-90%

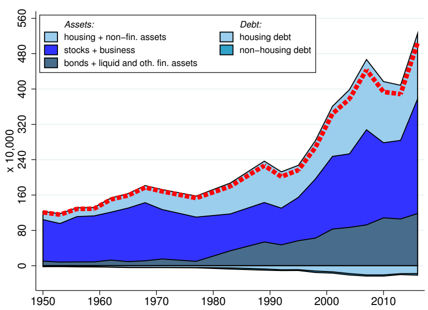

(c) top 10%

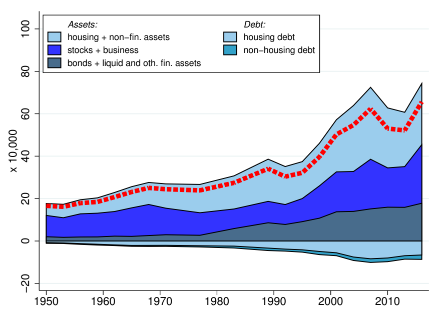

(d) all households

Notes: Heterogeneity in household portfolios for different groups along the wealth distribution. Top left panel shows portfolio composition of the bottom 50% of the wealth distribution. Assets are shown as positive portfolio components, debt as negative portfolio components. The red dashed line indicates mean wealth of this group. All portfolio components and wealth levels are shown in 10,000 dollars (2016 dollars). Top right panel shows the equivalent portfolio of the middle class (50%-90%) together with mean wealth of this group (red dashed line). Bottom left panel shows the equivalent portfolio of the top 10% of the wealth distribution and average wealth for this group (red dashed line). Bottom right panel shows the equivalent portfolio for all households and mean wealth as red dashed line. Wealth groups are separately defined for each survey year.

Bottom 50%: The bottom 50% have little wealth, yet their small wealth position masks substantial gross positions. Houses and other nonfinancial assets, mainly cars, make up more than 80% of the asset side of the balance sheet. Financial assets play a minor role in bottom 50% portfolios. On the liability side, housing debt is the dominant form of debt, but compared to other wealth groups, the bottom 50% also have a high share of non-housing debt. In recent years, education loans make up a growing share of this debt (Bartscher et al. (2019)). Assets exceed wealth by a large margin, indicating a high degree of leverage.

Middle class (50%-90%): The middle-class portfolio is dominated by nonfinancial assets. About two-thirds of the middle-class portfolio consists of houses and other non-financial assets. Direct stock holdings are typically below 5%. The large growth of other financial assets in the portfolio comes mainly from defined-contribution pension plans. The middle class is also leveraged, with housing debt being the dominant debt component and assets exceeding wealth by 10% to 30%.

Top 10%: The top 10% are different when it comes to portfolio composition. The bulk of wealth is held in stocks and business equity. Houses as an asset class gained in importance for the top 10% but constitute a comparatively small fraction of assets. Other financial assets have grown strongly, mainly because of the proliferation of defined-contribution pension plans. Leverage is low, so that for the top 10%, the sum of assets corresponds approximately to wealth.

Figure 2.4.4: Shares in total asset holdings by wealth group for selected years

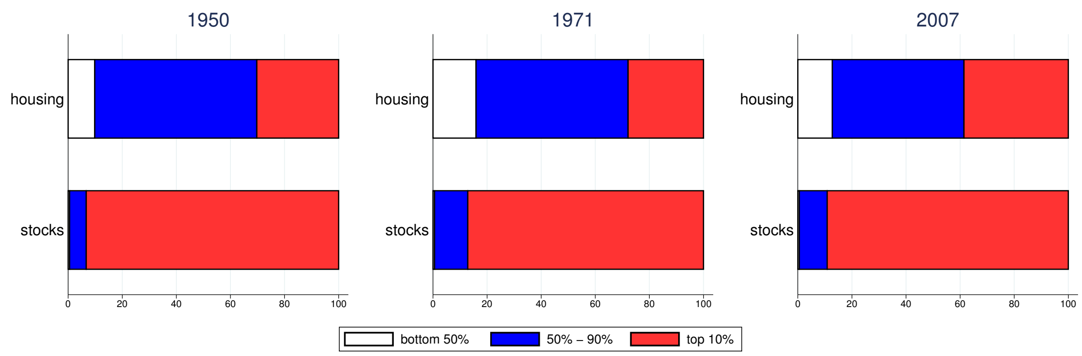

Notes: Shares in total asset holdings by wealth group for 1950, 1971, and 2007. Each bar shows the share of total assets held by different wealth groups. Wealth groups are the bottom 50%, the 50%-90%, and the top 10% of households. Top bars show housing assets, bottom bars show stocks including mutual funds. The left panel shows shares for 1950, the middle panel for 1971, and the right panel for 2007. The white part of each bar shows the share of the asset in that year held by the bottom 50%, the blue part the share held by the 50%-90%, and the red part the share held by the top 10%. Shares across wealth groups for each asset and year add up to 100%.

Portfolio composition thus varies substantially along the wealth distribution. These differences are also highly persistent. The portfolios of the bottom 90% are non-diversified and highly leveraged. Houses are *the* asset of the bottom 90%, making residential real estate the most egalitarian asset. Figure 2.4.4 highlights this point by showing the ownership structure of housing and stocks at different points in time.[31] The bottom 90%

[31] We include mutual funds in the stock holdings. Results change little if we only consider direct stock holdings.

hold about half of all housing wealth, but only a tiny fraction of stocks. Stocks are *the* asset of the wealthy in the sense that the top 10% consistently hold more than 90% of stocks. Looking at the Gini coefficient for individual asset classes confirms that housing is the most equally distributed asset and that the distribution of housing wealth has not changed substantially over time. We report Gini coefficients for all asset classes in the Appendix 2.A.5.1.

An important consequence of non-diversified and leveraged portfolio positions is that the wealth of middle-class households is highly exposed to changes in house prices. We quantify this exposure as the elasticity of wealth with respect to house prices, which is equal to $\frac{\text{Housing}}{\text{Wealth}}$, the ratio of the asset value of housing to wealth. Figure 2.4.5 shows the resulting exposure to house prices for middle-class households and households in the top 10% over time. The figure confirms that the elasticity of middle-class wealth to house prices is three to four times higher than at the top. A 10% increase in house prices increases middle-class wealth by 6%-7%.[32] For changes in stock prices, the exposures are reversed. The top 10% are highly exposed, the rest very little.

Figure 2.4.5: House price exposure $\left(\frac{\text{housing}}{\text{wealth}} \times 100\right)$ by wealth group

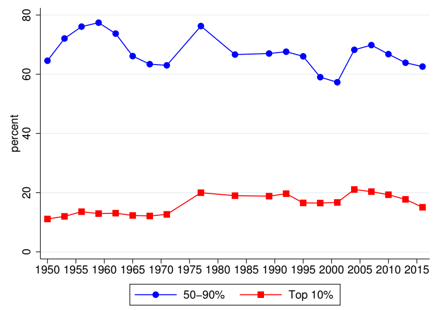

Notes: House price exposure for different wealth groups. House price exposure is measured by the elasticity of household wealth with respect to house price changes $\left(\frac{\text{housing}}{\text{wealth}} \times 100\right)$. Blue dots show house price exposure for the 50%-90% of households in the wealth distribution, red squares show house price exposure for the top 10% of households in the wealth distribution.

2.4.3 The race between the stock market and the housing market

Such pronounced portfolio differences between households along the wealth distribution give rise to what we call the race between the stock market and the housing market: owing to their larger exposure, the middle class gains relatively more than top-wealth households when house prices rise. All else equal, rising house prices make the wealth distribution more equal, while stock market booms have the opposite effect: they primarily boost wealth at the top and lead to a more unequal distribution of wealth.

To explore how important this race between the stock market and the housing market has

[32] We do not show the bottom 50% in this graph because of their large exposure. Appendix 2.A.5.2 provides further results on house price exposure along the wealth distribution and changes over time and decomposes the house price elasticity of wealth further into a *diversification component* and a *leverage component*.

Table 2.4.1: The race between the stock and the housing market

	(1)	(2)	(3)	(4)
β_h	-0.062	-0.071	-0.084	-0.073
β_s	0.024	0.026	0.024	0.026
θ^{top10}	no	yes	no	yes
$\frac{Y}{W}$	no	no	yes	yes
N	19	19	19	19
R^2	0.06	0.23	0.21	0.23

Notes: Regression of changes in the top 10% wealth share on asset price growth and controls. Growth rates computed using log differences. θ^{top10} denotes the income share of the top 10% of households in the wealth distribution. $\frac{Y}{W}$ denotes controls for the inverse of the wealth-to-income ratio of the top 10% of households in the wealth distribution and for the aggregate economy. All observations from the surveys from 1950 to 2016 are used for the regressions.

been for the wealth distribution in postwar America, we estimate the following regression relating changes in the top 10% wealth share over the three-year survey intervals to asset price movements:

$$\Delta \log(\omega_{t+1}^{top10}) = \beta_0 + \beta_h \Delta \log(p_{t+1}^h) + \beta_s \Delta \log(p_{t+1}^s) + \varepsilon_t,$$

where Δ is the first-difference operator, $\Delta x_{t+1} = x_{t+1} - x_t$, the superscript h denotes house prices and the superscript s stock prices. We use the S&P 500 stock market index and the Case-Shiller House Price Index obtained from the latest version of the Macrohistory Database (Jordà et al., 2017).

Table 2.4.1 reports the estimated coefficients for the baseline regression in the first column. The signs of the estimated coefficients demonstrate how the race between the housing market and the stock market shaped wealth dynamics: rising house prices are associated with a falling top 10% wealth share. Rising stock prices boost the top 10% wealth share. Note that in the baseline specification, the error term comprises all other effects related to differences in savings or wealth-to-income ratios. Wolff (2002) estimated similar regressions relating the top 1% and top 5% wealth share to fluctuations in stock and house prices.

We control for these factors in additional regressions in columns (2)-(4) by adding the income share of the top 10% and changes in the ratio of income to wealth as regressors. The coefficients β_h and β_s become larger and are economically significant, but the sample is so small that they are not statistical significant.

The estimated regression coefficients have an intuitive interpretation as they correspond to the average elasticity of the top 10% wealth share with respect to asset prices.[33] From the law of motion for the wealth share, we can derive the elasticity of the wealth share of group i with respect to the price of asset j:

[33] The portfolio share $\alpha_h = \frac{H}{W}$ in Figure 2.4.5 corresponds to the elasticity of the wealth level.

$$\frac{\partial\left(\frac{\omega^i_{t+1}}{\omega^i_t}\right)}{\partial\left(\frac{p^j_{t+1}}{p^j_t}\right)} = (1 + q_t + \sigma_t)^{-1}\Big(\alpha^i_{t,j} - \alpha_{t,j}\frac{\omega^i_{t+1}}{\omega^i_t}\Big).$$

Figure 2.4.6 shows the time series for the elasticity constructed from the portfolio shares in the SCF+ data. The house price elasticity of the top 10% fluctuates around a mean of -0.17, reasonably close and within the confidence bounds of the point estimates in the wealth share regressions above. All else equal, an elasticity of -0.17 implies that a 10% increase in house prices will lower the top 10% wealth share by 1.7%. Assuming a top 10% wealth share of 75%, this corresponds to a decrease in the top 10% wealth share of 1.3 percentage points.

The magnitudes are large. A hypothetical 40% increase in real house prices would reduce the wealth share of the top 10% from 75% to 70%, bringing it back to its 1971 level. For stock prices, the long-run average elasticity stands at 0.036 and is slightly larger than the point estimates from the regression above. A 130% real increase in the stock market — comparable to the period between 1998 to 2007 — increases the wealth share of the top 10% by about 6 percentage points. By the same token, the fall in the stock market in the 1970s contributed to the decline in wealth inequality in the 1970s that we discussed above. Because of the substantial portfolio heterogeneity along the wealth distribution, asset prices have first-order effects on the evolution of wealth inequality.

Figure 2.4.6: Asset price elasticity of top 10% wealth share

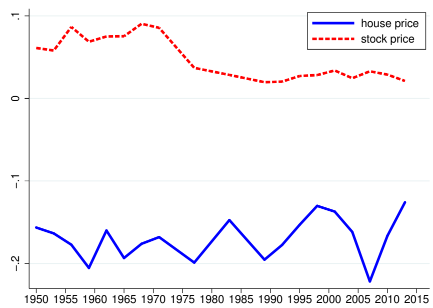

Notes: Elasticity of the top 10% wealth share with respect to asset prices. Red dashed line shows elasticity of the top 10% wealth share with respect to the stock price. Blue solid line shows elasticity of top 10% wealth share with respect to the house price.

2.4.4 Wealth gains from asset prices

The results of the previous section demonstrate that over short horizons, asset price fluctuations are closely associated with changes in wealth shares. We now quantify the contribution of asset price changes to wealth accumulation of different groups over the past four decades.

We concentrate on wealth growth over two distinct periods. The first period comprises the nearly four decades from 1971 to the 2007-2008 financial crisis (the last pre-crisis survey was carried out in 2007). This was a period in which the income distribution widened substantially, but measures of wealth inequality changed very little. We will see that house-price-induced wealth gains for the bottom 90% of the population played a central role in the observed stability of the wealth distribution.

The second period that we study covers the decade after the financial crisis. As discussed above, this decade has witnessed the largest increase in wealth inequality in postwar history. The income distribution, by contrast, changed only modestly over this period. We will see that asset prices are again central in accounting for the large shift in the distribution of wealth over a relatively short period.

Figure 2.4.7: Wealth growth from asset price changes, 1971-2007

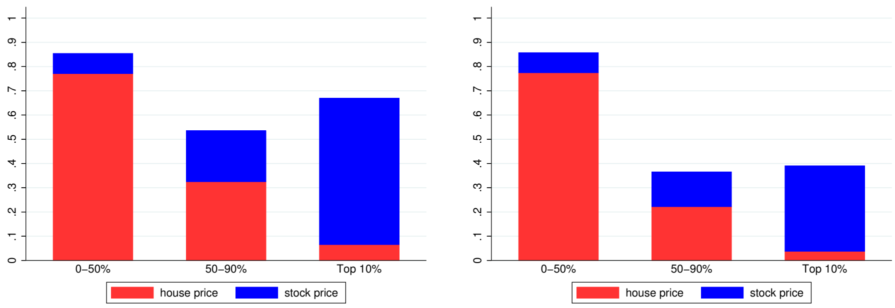

(a) Wealth growth from asset prices (b) Share in total wealth growth

Notes: Wealth growth from changes in house prices and stock prices in levels and as share of total wealth growth for different wealth groups. Left panel shows wealth growth from asset price changes between 1971 and 2007 for the bottom 50%, 50%-90%, and top 10% of the wealth distribution. Wealth growth from asset prices is computed as asset price change multiplied with 1971 stock of asset holdings of wealth group divided by the 1971 wealth level. Wealth groups are the bottom 50%, 50%-90%, and top 10% of the wealth distribution. Left panel shows wealth growth relative to 1971 wealth levels. Right panel shows share of wealth growth from asset prices from left panel as share in wealth growth of the respective group. Only asset price changes in housing and stock prices are considered. Red parts of bars show contribution from house price changes, blue part of bar show contribution from stock price changes.

Figure 2.4.7a shows how much wealth grew between 1971 and 2007 because of price changes in the stock market and the housing market. As before, we track wealth growth across three different wealth groups: the bottom 50%, the 50%-90%, and the top 10%. While Figure 2.4.7a displays the wealth growth from asset price changes, Figure 2.4.7b highlights the share of such asset-price-induced wealth growth in total wealth growth, including changes in wealth coming from savings flows.

Two observations stand out. First, over the entire period, wealth gains from rising asset prices were substantial across the distribution. As Figure 2.4.7a shows, the wealth of the bottom 50% grew by 90% only because of price effects. Also for the 50%-90% group and the top 10%, asset price changes induced sizable wealth growth of about 60%. It is easy to spot the race between the stock market and the housing market in the data: wealth growth from house price effects dominate for the bottom 90% of the wealth distribution, and stock price gains account for the bulk of the wealth gains at the top.

How much of the total wealth increase in the different groups is explained by price effects? Figure 2.4.7b shows that for the bottom 50% virtually all wealth growth over the 1971-2007 period came from higher asset prices. But even in the middle and at the top, asset prices still account for 40% of total wealth growth, with the rest accounted for by savings. Note that these estimates for the role of asset prices for wealth growth are likely conservative as households increased their exposure to the housing market over this period. The variation in the contribution to wealth growth along the distribution helps us to understand why wealth and income growth decoupled in the decades before the crisis: relative to income, wealth grew most strongly for the bottom 90%.

Figure 2.4.8: Wealth growth from asset price changes, 2007-2016

Notes: Wealth growth from changes in house prices and stock prices for different wealth groups between 2007 and 2016. Wealth groups are the bottom 50%, 50%-90%, and top 10% of the wealth distribution. Wealth growth from asset prices is computed as asset price change multiplied with 2007 stock of asset holdings of wealth group divided by the 2007 wealth level. Wealth groups are the bottom 50%, 50%-90%, and top 10% of the wealth distribution. Only asset price changes in housing and stock prices are considered. Red parts of bars show contribution from house price changes, blue part of bar show contribution from stock price changes.

Asset price movements also explain why wealth concentration spiked after the financial crisis. House prices plummeted after 2007 and recovered only slowly in recent years. By 2016, they were still 10% below their 2007 peak level. By contrast, the stock market recovered more quickly and the main stock market indices were about 30% above their 2007 levels in real terms in 2016. Figure 2.4.8 shows the race between the housing market and the stock market between 2007 and 2016. The bottom 50% lost 15% of wealth relative to 2007 levels, mainly because of lower house prices. By contrast, the top 10% were the main beneficiary from the stock market boom and were relatively less affected by the drop in residential real estate prices. The consequence of substantial wealth losses at the bottom and in the middle of the distribution, coupled with wealth gains at the top, produced a large spike in wealth inequality.

What would the distribution of wealth in America look like today without asset price effects? To construct a counterfactual, we use the law of motion for wealth levels from equation (2.4.2), keeping all parameters constant but adjusting the asset return term, q_t^i, so that nominal house prices (or stock prices) only increased with the rate of CPI inflation since 1971 (i.e., we keep prices constant in real terms). Note that this is an accounting exercise, not a general equilibrium analysis. Our aim here is to illustrate the potential of asset prices to shift the wealth distribution.

Table 2.4.2 shows the measured change in the wealth shares of the three groups relative to 1971, as well as the counterfactual change under two scenarios. In the first scenario, we keep real house prices constant. In the second scenario, we fix real stock prices at their 1971 level. The table shows counterfactual changes in wealth shares under these two corner assumptions.

Table 2.4.2: Changes in wealth shares relative to 1971

		1989	2007	2016
bottom 50 %	observed change	-0.3	-0.8	-2.2
	constant house prices	-0.5	-1.7	-2.8
	constant stock prices	-0.2	-0.4	-2.0
50% - 90%	observed change	2.1	-1.7	-6.1
	constant house prices	1.7	-3.6	-7.6
	constant stock prices	2.6	1.4	-2.9
Top 10%	observed change	-1.8	2.4	8.2
	constant house prices	-1.2	5.3	10.4
	constant stock prices	-2.4	-1.1	4.9

Notes: Changes in wealth shares for different wealth groups relative to 1971. First part of the table shows the bottom 50% of the wealth distribution, middle part shows the 50%-90% of the wealth distribution, and bottom part shows the top 10% of the wealth distribution. For each wealth group, the first row shows the observed change in wealth shares between 1971 and the indicated year at the top of the column. The second row ("constant house prices") shows the counterfactual change in wealth shares for constant real house prices at the 1971 level. The third row ("constant stock prices") shows the counterfactual change in wealth shares with constant real stock prices at the 1971 level. Changes in wealth shares are shown in percentage points relative to the 1971 level of the wealth share for each group.

The key message from Table 2.4.2 is that the asset price effects were potentially large. From 1971 to 2007, rising house prices slowed down wealth concentration in the hands of the top 10% by 2.9 percentage points. Without higher house prices, the increase in wealth concentration at the top would have been four times higher than what we observe in the data. The house price crash after 2007 largely reversed these effects, but even in 2016, the observed increase in the top 10% wealth share of 8.1 percentage points was still about one-quarter lower than the counterfactual increase of 10.4 percentage points. Again, it is important to realize the magnitude of these shifts in wealth shares. A difference of 2 percentage points corresponds to over 14% of total annual household income. The last row for each wealth group in Table 2.4.2 also reports the corresponding counterfactual for constant stock prices. Without the stock market boom, the top 10% wealth share would have been 1.1 percentage points lower in 2007 than in 1971, and even over the whole period, the middle class (50%-90%) would not have lost substantially less ground relative to the top. These counterfactual simulations are suggestive only, but they highlight to what extent the wealth distribution is sensitive to asset price dynamics.

2.5 Conclusions

This paper makes three contributions to the literature on income and wealth inequality. First, we introduced the SCF+, an extended version of the Survey of Consumer Finances. The SCF+ covers as household-level dataset the financial situation of U.S. households since World War II. The SCF+ complements existing datasets for long-run inequality research that are based on tax and social security records. The SCF+ makes it possible to study the joint distribution of income and wealth over time as it contains both income and balance sheet information, coupled with extensive demographic information. We expect that the SCF+ data will become a valuable resource for future empirical and theoretical research on inequality, household finance, political economy, and beyond.

Second, we exploited the new data to study the trends in income and wealth inequality. Previous research documented a trend toward increasing polarization of income and wealth since the 1970s. The data confirm this finding and underscore that the American middle class was the main loser of increasing income concentration at the top. We also track the racial wealth gap between black and white households over the long run. The picture that the SCF+ paints is one of a persistent income and wealth divide between white and black households. Importantly, we expose divergent trends in the income and wealth distribution for extended periods. Before the crisis in 2008, income growth of households in the lower half of the distribution fell substantially short of wealth growth for these households. However, a large part of these wealth gains have been undone in the crisis. We link these findings to the importance of asset price changes for wealth dynamics.

The third main contribution of the paper is to expose systematic and highly persistent differences in portfolio composition and leverage of households along the wealth distribution. An important consequence of these differences is that asset price changes have first-order effects on the wealth distribution. They lead to capital gains and losses that induce shifts in the wealth distribution that are unrelated to changes in the income distribution. The magnitude of changes in the wealth distribution induced by this asset price channel can be large. By highlighting the crucial role that portfolio composition, leverage, and asset prices play for the wealth distribution, our paper opens up new avenues for future empirical and theoretical research on the determinants of wealth inequality.

2.A Appendix

This appendix accompanies the paper *Income and Wealth Inequality in America, 1949-2016*. Section 2.A.1 provides further details on the SCF+ data and its construction. We also explain how we implement the approaches for demographic adjustment and how we identify hand-to-mouth households in the data. Section 2.A.2 discusses the coverage of the SCF+ at the top of the income and wealth distribution. Section 2.A.3 discusses evidence on mobility between wealth groups based on Panel Study of Income Dynamics (PSID) data. Section 2.A.4 provides sensitivity tests for Gini coefficients excluding negative wealth or income observations and for imputing rental income from owner-occupied housing. Section 2.A.5 provides additional results on Gini coefficients for different asset classes, a decomposition of house price exposure for different wealth groups over time, and presents time series for income and wealth shares, Gini coefficients for income and wealth, and for the portfolio composition along the wealth distribution.

2.A.1 Data

This section provides details on the imputation of missing variables, construction of replicate weights, number of household observations across years, home ownership rates after weight adjustment, and construction of consistent income and wealth concepts for comparison with NIPA and FA data.

2.A.1.1 Imputation of missing variables

This section provides further details on the imputation of missing variables by predictive mean matching as described in Schenker and Taylor (1996) and the imputation of car values for selected years. Following the modern SCF, we use multiple imputation and produce five imputed values for each missing variable. The imputation involves several steps. First, a linear regression model of the variable of interest is estimated on a sample with non-missing observations. For each of the multiple imputations, a random realization of the regression coefficients is drawn using the estimated variance-covariance matrix. Using this coefficient vector and the linear regression model, a prediction for the variable of interest is generated. The predicted values on missing and non-missing observations are compared to find non-missing observations that produce the closest prediction. For each missing observation, we choose the three observations among the non-missing observations that have predicted values most similar to the respective missing observation. Out of these three, we choose one observation randomly and assign the value of the variable of interest to the corresponding missing observation. Hence, the linear regression model is only used to define the distance between missing and non-missing observations. The imputed values for the variables are all observed values. We refer to Schenker and Taylor (1996) for an in-depth discussion of the topic.

For each missing variable, several adjacent surveys could in principle be used as nonmissing samples for the imputation. In order to determine which adjacent survey years are most suitable for imputing missing values, we implement the following optimization before imputation. First, we determine all income, asset, debt, and demographic variables that are available in the year for which the variable is missing. For each combination of adjacent years, we then construct a subset of variables that are available both in the year with missing values and in the adjacent years. As the predictive accuracy decreased with the number of explanatory variables, we select those variables with the highest predictive power by using the lasso method. This method sets regression coefficients to zero for variables with small predictive power. For each combination of survey years, we then regress the variable of interest on those variables selected by the lasso method. Finally, we calculate the R^2 for each regression. We use the R^2 as a measure of how well the combination of adjacent years is able to predict the missing variable. The combination with the highest R^2 is chosen for the imputation.

We account for a potential undercoverage of business equity before 1983 and follow the method proposed by Saez and Zucman (2016) to adjust the observed holdings in the micro data with information from the FA. We rely on data from the 1983 and 1989 surveys and adjust business wealth and stock holdings including other managed assets in the earlier surveys so that the ratio of business wealth and stocks matches the 1983 and 1989 values. Let X_{it} be business wealth or stocks of observation i in period t. The variable $\bar{\mathrm{X}}_t$ is the respective mean in period t, and X_t^{FA} is the corresponding FA position per household in t. The adjusted values of business wealth and stocks are then calculated as follows

$$\mathrm{X}_{it}^{adj} = \mathrm{X}_{it} \frac{\mathrm{X}_t^{FA}}{\bar{\mathrm{X}}_t} \frac{\bar{\mathrm{X}}_{1983,1989}}{\mathrm{X}_{1983,1989}^{FA}}$$

For cars, the current value is available in the historical files for 1955, 1956, 1960, and 1967. We impute the value in other years using information on age, model, and size of the car. Surveys up to 1971 include information on age, model, and size of the car a households owns. If a household bought a car during the previous year, the purchasing price of this car is also available. We impute the car value using the average purchasing price of cars bought in the previous year that are of the same age, size, and model. In 1977, only information on the original purchasing price and the age of the car is given. For this year, we construct the car value assuming a 10% annual depreciation rate.

2.A.1.2 Confidence intervals for income and wealth shares

Bricker et al. (2018) provide a discussion of *modeling error* in the capitalization approach to the wealth distribution. Survey data are immune to modeling error as income and wealth are directly observed from the answers of survey participants. The survey data, however, contain measurement and sampling error. To provide estimates of the variability that results from these errors, replicate weights for the modern SCF data have been constructed and provided to researchers. We construct replicated weights for the historical surveys following as closely as possible the practice for the modern samples. We draw 999 samples from the data stratified by geographical units and adjust weights to get a nationally representative sample. We use these replicate weights to construct all confidence bounds of our estimates shown in the main paper.

2.A.1.3 Sample size

Table 2.A.1 reports the number of household observations for the different sample years. As described in Section 2.A.1.1, there are five implicates for each household observation in the final data. For the historical data, we pool sample years when surveys were conducted annually to increase the accuracy of our estimates. We show the number of observations for single survey years and after pooling sample years in Table 2.A.1. The years highlighted in bold are used for all results in the main part of the paper using the pooled samples.

Table 2.A.1: Number of household observations

year	observations	pooled	year	observations	pooled	year	observations
1949	2,988		1960	2,708		**1983**	4,103
1950	2,940	5,928	**1962**	1,922	4,630	**1989**	3,143
1951	2,938		1963	1,819		**1992**	3,906
1953	2,663	5,601	**1965**	1,349	3,168	**1995**	4,299
1954	2,599		1967	3,165		**1998**	4,305
1955	2,766		**1968**	2,677	5,842	**2001**	4,442
1956	2,660	8,025	1969	2,485		**2004**	4,519
1957	2,726		1970	2,576		**2007**	4,417
1958	2,764		**1971**	1,327	6,388	**2010**	6,482
1959	2,790	8,280	**1977**	2,563		**2013**	6,015
Total number observations: 102,304						**2016**	6,248

Notes: Number of household observations in SCF+ data. The first column shows the survey year. Survey years in bold are used for time series in the main paper. The second column shows the number of household observations for different survey years. The third column shows the number of observations after pooling survey years. For results in the main paper, pooled survey years are always used. Horizontal lines indicate the pooled annual survey years.

2.A.1.4 Homeownership rates

In Section 2.2.2, we discuss the adjustment of survey weights to match population shares for age of the household head, college education, and race to be consistent with Census data. In the same step, we also target homeownership rates. Figure 2.A.1 shows the homeownership rate in the CPS/Census (black squares) and in the SCF+ with the adjustment of survey weights (red dots) and without adjustment (blue diamonds).

Figure 2.A.1: Homeownership rates

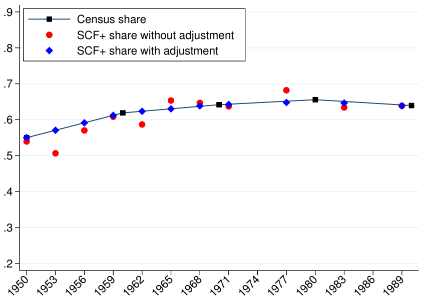

Notes: The black squares show homeownership rates in the CPS/Census data. The red dots show the homeownership rate using the original (unadjusted) survey data. The blue diamonds show the homeownership rate using the adjusted survey data.

2.A.1.5 Comparison to NIPA and FA

It is well known that even high-quality micro data do not correspond one-to-one to aggregate data due to differences in measurement. For instance, Heathcote et al. (2010) discuss that data from the NIPA and CPS differ substantially. Indirect capital income from pension plans, nonprofit organizations, and fiduciaries, as well as employer contributions for employee and health insurance funds, are measured in the NIPA but not in household surveys such as the CPS or the SCF. Henriques and Hsu (2014) and Dettling et al. (2015) provide detailed discussions of the differences between SCF micro data and FA and NIPA data and explain how to construct equivalent income and wealth measures.

For the construction of incomes for the comparison in Section 2.2.3, we follow Dettling et al. (2015). We start with personal income from NIPA table 2.1 and subtract/add components as described in Table A1 of Dettling et al. (2015). Our SCF+ income measure does not include capital gains. We also abstain from subtracting any components from *other income* as we do not have the detailed breakdown by source that is used in the adjustment by Dettling et al. (2015) for this position. We adjust NIPA income by CPI and divide totals by the number of households from the Census.

To construct wealth from the FA data, we add the following FA positions using annual data:

NW (wealth)	=	FIN (financial assets) + NFIN (nonfinancial assets) – TDEBT (debt)
FIN	=	DEPOS (deposits) + BONDS (bonds)
		+ CORPEQUITY (corporate equity) + MFUND (mutual funds)
		+ DCPEN (defined contribution pension wealth)
NFIN	=	BUS (noncorporate businesses) + HOUSE (houses)
TDEBT	=	HDEBT (housing debt) + PDEBT (personal debt)
DEPOS	=	LM153091003(A) + FL153020005(A) + FL153030005(A) + FL153034005(A)
BONDS	=	FL153061105(A) + FL153061705(A) + FL153062005(A) + FL153063005(A)
CORPEQUITY	=	LM153064105(A)
MFUND	=	LM153064205(A)
DCPEN	=	FL574090055(A) + FL224090055(A) + FL344090025(A)
BUS	=	LM152090205(A)
HOUSE	=	LM155035015(A)
HDEBT	=	FL153165105(A)
PDEBT	=	FL153166000(A)

This construction of wealth excludes the identifiable components of wealth of nonprofit organizations that comprise real estate, equipment, intellectual property products, open market paper, other loans and advances, municipal securities, commercial mortgages, and trade payables. We exclude consumer durables and assets related to life insurance, security credit, and miscellaneous assets and loans. These positions are excluded either because they belong to nonprofit organizations that are not part of the SCF household survey, or because of non-comparability. Dettling et al. (2015) explain the reasons for this selection in detail in Appendix B of their paper. We do not subtract the wealth of the Forbes 400 from the FA total, because these estimates only exist since 1982. We adjust FA values by CPI and divide totals by the number of households from Census data.

2.A.1.6 Accounting for demographic change

We implement the approach proposed by Fortin et al. (2011) to construct counterfactual wealth and income distributions when holding individual demographic characteristics constant over time. We choose 1971 as the base year. The counterfactual relies on changing the conditional distribution of a certain demographic variable X_d. This is achieved by a re-weighting approach. We follow the steps in Fortin et al. (2011) for the case of general covariates. In the first step, we pool the data from the base year and the year for which we want to compute the counterfactual. In the second step, we estimate a re-weighting factor

$$\hat{\Psi}_X(X) = \frac{\hat{P}(D_Y = 1|X)/\hat{P}(D_Y = 1)}{\hat{P}(D_Y = 0|X)/\hat{P}(D_Y = 0)}.$$

D_Y is a dummy which is equal to 1 in the base year and else zero. The estimate is obtained by combining the predicted probabilities from a probit or logit regression of D_Y on the covariates X, and the pooled sample proportions of the two groups, $\hat{P}(D_Y = j)$, $j = 0, 1$. In the third step, we estimate a similar re-weighting factor using all covariates except X_d, $\hat{\Psi}_{X_{-d}}(X_{-d})$. Finally, we adjust survey weights by the factor $\frac{\hat{\Psi}_X(X)}{\hat{\Psi}_{X_{-d}}(X_{-d})}$. In order to obtain the marginal effect of covariate X_d, we subtract from the statistic of interest derived using the adjusted weights the statistic derived using survey weights adjusted by $\hat{\Psi}_X(X)$. We add the marginal effect to the baseline estimate without adjustment. We used a probit model, including age, educational attainment, the number of adults and children, as well as race as explanatory variables in X. We first fix the age distribution to its 1971 level, and then fix the share of households whose head has at least attained some college to the base year.

2.A.1.7 Identifying hand-to-mouth households

We follow Kaplan et al. (2014) and restrict the sample to households with a household head between 22 and 79 years of age. We also exclude households without other income than self-employment income and households with negative income. We refer to this sample as the restricted sample in the empirical results. Following Kaplan et al. (2014), we classify a household as *hand-to-mouth* if one of the following two conditions applies:

$$0 \leq a \leq \frac{1}{2}\frac{y}{\theta}$$

$$a < 0 \qquad \text{and} \qquad a \leq \frac{1}{2}\frac{y}{\theta} - \underline{a}$$

where a denotes liquid assets of the household, y denotes monthly household income net of capital income, θ is the payment frequency for income, and $\underline{a}$ denotes the borrowing limit. The factor $\frac{1}{2}$ is due to the assumption that resources are consumed at a constant rate over the month.

We follow the baseline parametrization in Kaplan et al. (2014) and set $\theta = 2$ and the borrowing limit to one times monthly income ($\underline{a} = y$). According to the first definition, a household is *hand-to-mouth* if liquid assets are positive but less than 25% of monthly labor income. Alternatively, a household is *hand-to-mouth* according to the second definition if it has negative liquid assets and the remaining credit line $\underline{a} + a$ is less than 25% of monthly labor income.

Following Kaplan et al. (2014), we distinguish between *wealthy* and *poor* hand-to-mouth households. We classify a household as a *wealthy* hand-to-mouth household if the household is hand-to-mouth according to one of the definitions from above and if the household has positive illiquid assets. Illiquid assets are the sum of housing, other real estate, certificates of deposits, retirement accounts, the cash value of life insurances, and savings bonds. From this asset position, we subtract housing debt associated with the primary residence and other real estate. The net position is the estimate for illiquid assets. Liquid assets comprise all assets held in checking, savings and call accounts, money market deposit accounts, money market mutual funds and (since 2016) prepaid debit cards.

Cash holdings are not recorded in the SCF data and have to be estimated using other data sources. We rely on different estimates for cash holdings for the historical data. The first estimate relies on Kaplan et al. (2014) who use evidence from the 2010 Survey of Consumer Payment Choice (SCPC). Based on this evidence, they adjust the assets in transaction accounts by a factor $1 + \frac{138}{2{,}500}$ to account for cash holdings, i.e. they increase the money held in transaction accounts by roughly 5.5%. This estimate is the ratio of average individual cash holdings from the SCPC (excluding large-value holdings), divided by the median of assets in transaction accounts from the restricted 2010 SCF sample. It seems plausible that households historically relied more on cash during the era when credit or debit cards and ATMs were not yet widely available.

Figure 2.A.2: Cash holdings estimated from NCVS relative to median income

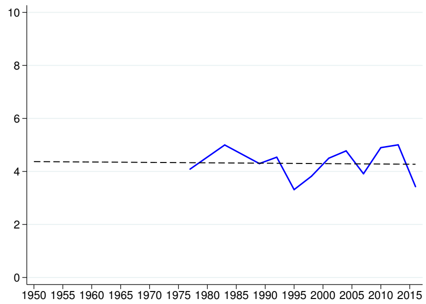

Notes: Estimated cash holdings relative using National Crime Victimization Survey (NCVS) data. The figure shows the ratio of the average amount of cash stolen in theft incidences from NCVS data relative to median total household income in SCF+ data. The series was extended backwards before 1973 by linear extrapolation. The NCVS data are at the individual level.

Providing alternative estimates on cash holdings is difficult, as data on cash holdings are notoriously hard to find. We provide new estimates on U.S. individual cash holdings by relying on data from the National Crime Victimization Survey (NCVS). The survey explores if individuals have been victims of crime and asks about details of the incidence, including theft. An advantage of using data on the theft of cash is that victims are likely to accurately remember the amount. Like Kaplan et al. (2014), we exclude particularly large cash amounts stolen as they might be more systematically selected and less exogenous. In particular, we exclude values above the 99th percentile of the stolen cash distribution. The data from the NCVS exist since 1973. Figure 2.A.2 shows the cash estimates as a share of median income from the SCF. For the period prior to 1973, we extrapolate this series using a linear fit, which is plotted as a dashed line in the figure.

2.A.2 Coverage at the top of the distribution

The main part of the paper analyzes inequality trends using three wealth groups: the bottom 50%, 50%-90%, and the top 10% of households of the wealth distribution. A large part of the literature on trends in wealth and income concentration focuses on smaller groups in the right tail of the income and wealth distribution such as the top 1% (Piketty and Saez (2003), Saez and Zucman (2016), Bricker et al. (2016)). Capturing the far right tail is challenging for survey data, because non-response becomes prevalent at the top of the distribution (Sabelhaus et al., 2015). Starting in 1983 the modern SCF introduced a dual-frame sampling scheme that heavily oversamples households based on information from tax data. The historical SCF data do not feature a similarly sophisticated sampling scheme. As explained in the paper, we use a re-weighting approach to account for more prevalent non-response at the top of the distribution. This re-weighting approach is calibrated to the 1983 data and described in detail in the paper. A core assumption is that non-response patterns do not change systematically over time.

Bricker et al. (2016) provide a comprehensive discussion of the modern SCF data and the oversampling approach. They also propose a series of tests to examine the coverage of the (modern) SCF data at the top of the distribution in comparison to the tax data. In this section, we apply the tests suggested by Bricker et al. (2016) to the historical data. The aim is to better understand how the historical SCFs perform in capturing the right tail of the distribution relative to the tax data.

2.A.2.1 The top of the distribution: SCF+ vs. tax data

Income
For the modern SCFs, Bricker et al. (2016) use the income thresholds from tax data and compute the share of households in the survey data that is above this threshold. The data collected by Piketty and Saez (2003) allow us to implement the same test for the historical survey data. If the share of households is lower in the historical data, it could be an indication that top-income households are under-represented in the historical surveys.

We use the 90th, 95th, and 99th income percentiles reported by Piketty and Saez (2003). We adjust all data so that incomes are expressed in 2016 dollars, and calculate income thresholds excluding capital gains (Table A4 from Piketty and Saez (2003)) to align the tax data with income from the SCF+ data.

Figure 2.A.3 shows the share of households in the SCF+ data who are above the 99th percentile (Figure 2.A.3a), the 90th, and the 95th percentile (Figure 2.A.3b). Looking first at the share of households above the 99th percentile in Figure 2.A.3a, we observe that the share of households is in all years larger than 1%. Bricker et al. (2016) discuss this finding and the reasons behind it for the post-1989 data. If the historical data systematically miss households above the 99th percentile, we either expect a positive time trend in the share of households in case the coverage of the top improves over time, or a pronounced level shift from the 1970s to the 1980s, when we transition from the historical data to the modern SCF data. This is not the case. The data show neither a time trend nor any indication that there has been a pronounced level shift between the 1970s and 1980s in the share of households above the 99th percentile. The share fluctuates around its long-run mean (dashed line). We repeat the exercise for households above the 90th and 95th percentiles of the income tax distribution in Figure 2.A.3b. We confirm the same pattern: the shares fluctuate around their long-run means with no indication of a secular trend or level shift from the 1970s to the 1980s.

As a second test, Bricker et al. (2016) propose to compare the incomes of households above the 99th percentile threshold. Even if there is no trend in the share of households

Figure 2.A.3: Shares of SCF+ households above income thresholds from tax data

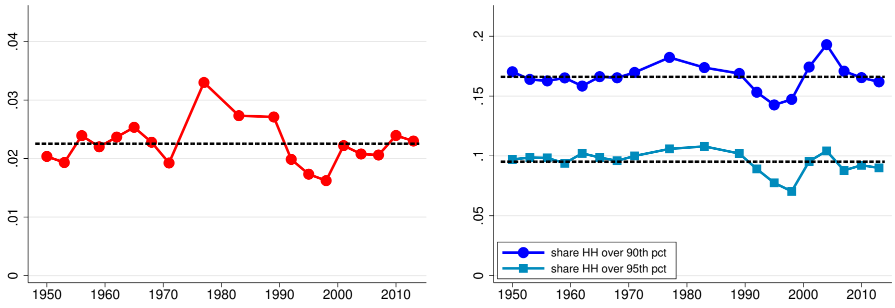

(a) 99th percentile (b) 90th/95th percentile

Notes: Share of households in the SCF+ data above 90th, 95th, and 99th percentile of the income distribution from the tax data. Left panel: Share of households in the SCF+ data above the 99th percentile from the tax data. Dashed line shows sample average. Right panel: Share of households in the SCF+ data above the 90th and 95th percentile of the income distribution from the tax data. Dashed line shows sample average for both time series. All income levels for tax data are taken from Piketty and Saez (2003).

above the 99th percentile, the historical surveys might have too few households in the very right tail of the income distribution. If this is the case, average income for households above the 99th percentile will be too low and we should see mean incomes jump (compared to the tax data) when we move from the historical to the modern SCFs.

Figure 2.A.4 compares mean income of households who are above the 99th percentile of the tax-data income distribution. The two series co-move closely over time. Importantly, there is no indication of a level shift in the average income series from the 1970s to the 1980s, when the historical and modern SCF data are combined.

Figure 2.A.4: Income level of households above the 99th percentile of tax data

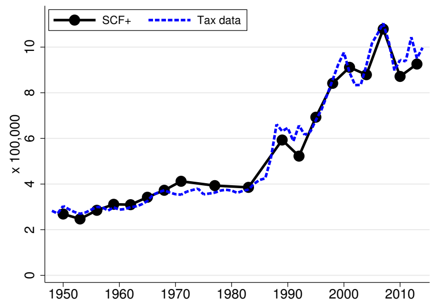

Notes: Mean income of households in the SCF+ and tax data above the 99th percentile of the income distribution from the tax data. All data have been transformed to 2016 dollars using the CPI and levels are shown in 100,000 dollars. All income levels for tax data are taken from Piketty and Saez (2003).

Wealth For wealth the comparison is less comprehensive because we have to rely on indirect estimates for wealth thresholds from Saez and Zucman (2016) that are based on the capitalization method. These estimates start in 1962 only. As in the case of income, we report in the first step households above the 90th, 95th, and 99th percentiles from the estimated wealth distribution in Saez and Zucman (2016). In the second step, we again compare the conditional mean wealth level above the 99th percentile using the tax-data estimated and the estimate from the SCF+ data.

Figure 2.A.5: Shares of SCF+ households above estimated wealth thresholds from tax data

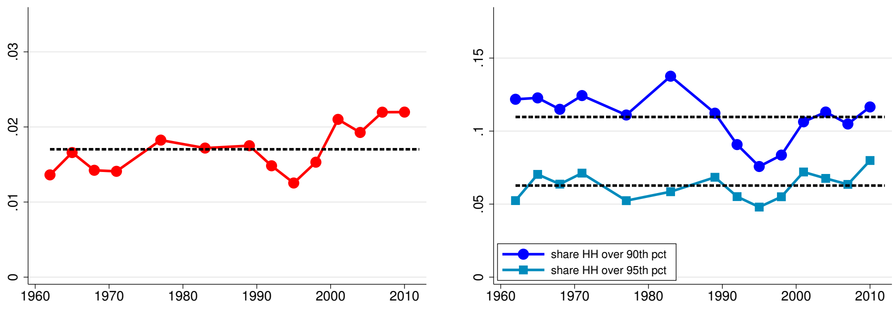

(a) 99th percentile (b) 90th/95th percentile

Notes: Share of households in the SCF+ data above 90th, 95th, and 99th percentile of the estimated wealth distribution based on tax data and the capitalization method. Left panel: Share of households in the SCF+ data above the 99th percentile of the estimated wealth distribution from tax data. Dashed line shows sample average. Right panel: Share of households in the SCF+ data above the 90th and 95th percentile of the estimated wealth distribution from tax data. Dashed line shows sample average. All estimated wealth levels from tax data are taken from Saez and Zucman (2016).

Looking at the share of households above the 99th percentile of the wealth distribution in Figure 2.A.5a, we see a similar picture as for the income distribution. The share is greater than 1% and fluctuates around its long-run mean over the entire sample period. Again, there is no evidence of a structural break between the 1970s to the 1980s. Figure 2.A.5b shows results for the 95th and 90th percentile of the wealth distribution.

How does the conditional mean wealth level for these households above the 99th percentile compare between the two sources? To get comparable estimates in terms of covered asset classes, we remove retirement accounts from wealth in the SCF+ and tax data.[34] As in the case of income, we find a close alignment for wealth levels. Comparing the pre-1983 to the post-1983 period, we find no evidence of a break in the time series when the historical and modern SCF data are merged.

It is important to add a cautionary note. The evidence presented here does not support the conclusion that the SCF+ data match the distribution *within* the top 1% over time. The evidence presented only offers support for the notion that the re-weighted historical data capture the representative top 1% household reasonably well (measured against the benchmark of the tax data). The distribution within the top 1% is beyond reach without the heavy oversampling of the modern SCFs.

[34] The tax-based estimates by Saez and Zucman (2016) include defined-benefit retirement accounts while retirement accounts in the SCF+ data only comprise defined-contribution retirement accounts.

Figure 2.A.6: Wealth level of households above the 99th percentile of capitalized tax data

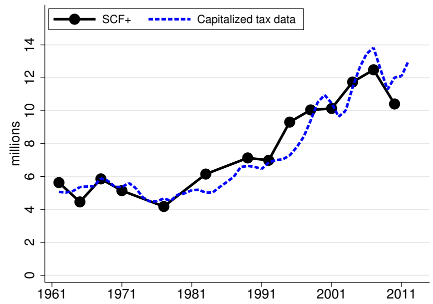

Notes: Mean wealth of households in the SCF+ data and tax data above the 99th percentile of the estimated wealth distribution from tax data. All data have been transformed to 2016 dollars using the CPI and levels are shown in million dollars. All estimated wealth levels from tax data are taken from Saez and Zucman (2016).

2.A.2.2 Percentiles

Another way to gauge the coverage at the top is to follow the levels of income and wealth of the top percentiles over time. Figure 2.A.7 shows the 90th, 95th, and 99th percentiles of the income and wealth distribution over time, as well as the growth of these percentiles relative to 1971.

Starting with income levels, all three percentiles evolve smoothly over time, and there is no indication of a shift in the early 1980s when we move from the historical to the modern surveys. The levels of the wealth percentiles in Figure 2.A.7b also do not display level shifts. Figures 2.A.7c and 2.A.7d show the growth relative to the base 1971. Once more the transition from the adjusted historical surveys to the modern SCF does not lead to breaks in the time series.

2.A.2.3 Top income and wealth shares

Figure 2.A.8 shows top 1% income and wealth shares from the SCF+ and compares them to the estimates by Piketty and Saez (2003) and Saez and Zucman (2016). Looking at the income share of the top 1% in Figure 2.A.8a, the SCF+ and the tax data align closely in levels and trends for the historical period. The tax data and the combined SCF and SCF+ data both show a large increase of the top 1% income share between 1971 and 2016. Importantly, the increase in the top 1% income share is concentrated in the period after 1980 where we rely on modern SCF data that capture the very right tail of the income distribution. For wealth in Figure 2.A.8b there is some divergence in trends over the last two decades. Bricker et al. (2016) and Kopczuk (2015) argue that trends from the capitalized tax data overstate the increase in wealth concentration over this time period.

Figure 2.A.7: Top percentiles for income and wealth over time

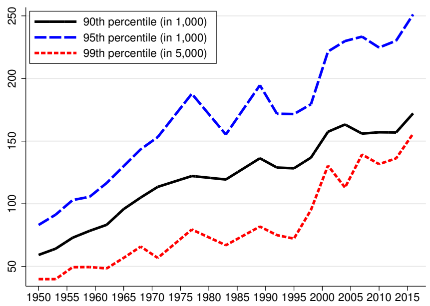

(a) income

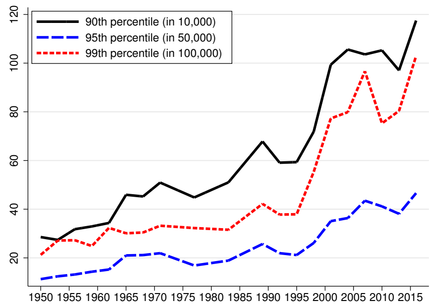

(b) wealth

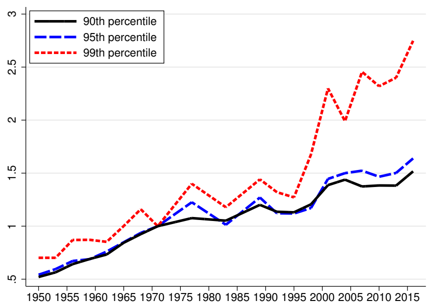

(c) income growth

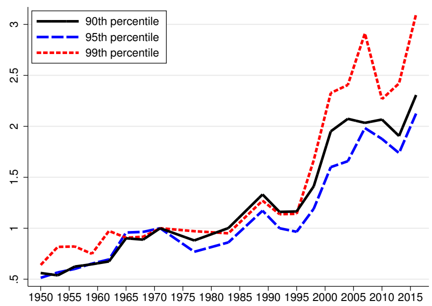

(d) wealth growth

Notes: Percentile levels for income and wealth and growth of these percentiles relative to 1971. Top left panel: Levels of the 90th, 95th, and 99th percentile of the income distribution (2016 dollars). Top right panel: Levels of the 90th, 95th, and 99th percentile of the wealth distribution (2016 dollars). Bottom left panel: Growth of 90th, 95th, and 99th percentile of the income distribution relative to 1971 (= 1). Bottom right panel: Growth of 90th, 95th, and 99th percentile of the wealth distribution relative to 1971 (= 1).

Figure 2.A.8: Top 1% income and wealth shares

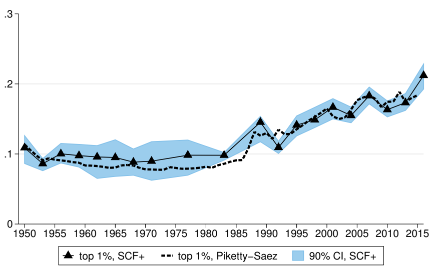

(a) Income

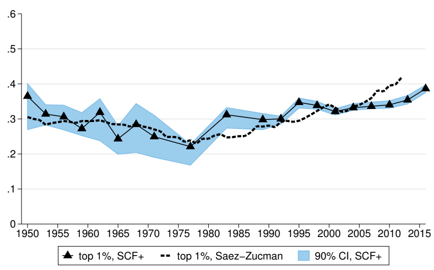

(b) Wealth

Notes: Top 1% income and wealth shares from SCF+ data and Piketty and Saez (2003) and Saez and Zucman (2016). The triangles show income and wealth shares from SCF+ data, the dashed lines income shares from Piketty and Saez (2003) using IRS tax data or wealth shares from Saez and Zucman (2016) using IRS data and the capitalization method. Light blue areas around the SCF+ estimates show 90% confidence bands. Confidence bands are bootstrapped using 999 different replicate weights constructed from geographically-stratified sample of the final data set.

2.A.3 Wealth mobility

The SCF+ data provide detailed information on the financial situation of U.S. households over the last seven decades. The survey is a repeated cross section and does not track households over time.[35] For our analysis in the main part of the paper, we follow the synthetic panel approach of Piketty and Saez (2003) and Saez and Zucman (2016). Households are grouped by wealth or income and then group-level income and wealth are traced over time.

In this section, we use additional data from the Panel Study of Income Dynamics (PSID) to quantify the degree of income and wealth mobility between survey dates. Díaz-Giménez et al. (2011) provide a similar analysis for the period from 2001 to 2007. We extend their analysis for the time period from 1983 to 2010. The main advantage of the PSID data is that it tracks the same households over time in a panel, but there are also limitations because the PSID is not designed as a wealth survey. When the PSID started in 1968, there was no systematic coverage of households' balance sheets. This changed in 1983, but the collected financial data are generally seen as lower quality compared to the SCF data as the PSID relies heavily on imputations. While this is less of a concern for cross-sectional analysis, imputation in the panel dimension tends to increase mobility due to imputation errors. Furthermore, it is important to note that for the period from 1983 to 1993, the time intervals between wealth surveys are five years and therefore longer than after 1998, when the time intervals reduce to two years. Hence, the persistence of households in wealth groups should be expected to be lower during the 1983-1993 period compared to the post-1998 period with shorter intervals between survey dates.

With these caveats in mind, we use the PSID data to explore mobility across the three wealth groups from the analysis in the main part of the paper: the bottom 50%, 50%-90%, and top 10%. For this analysis, we follow Kaplan et al. (2014) by restricting the sample to household heads from the SRC sample but abstain from any further sample selection.

Table 2.A.2 shows the share of households that remain within their respective wealth group between survey dates. The share is generally high.[36] For two-year periods, the data show that more than 83% of households from the bottom 50% of the wealth distribution stay within their wealth group across surveys. We find equally high stability for the 50%-90% where around 80% of households remain in this wealth group between survey dates, and also the persistence at the top of the wealth distribution recently increased to close to 80%. Moreover, most households that move out of the top 10% between survey dates, remain close to the top. For instance, the share of households from the top 10% that is still within the top 20% at the next survey date is around 90% for the two-year periods and 86% for the five-year period.

[35] The 1983 and the 2007 SCF provide information on a subset of households from the initial survey at a second interview three years later.

[36] We always require households to be present at both survey dates when computing flows.

Table 2.A.2: Wealth mobility

	bottom 50%	50%-90%	top 10%
1983	0.774	0.748	0.647
1988	0.750	0.740	0.672
1993	0.748	0.758	0.679
1998	0.833	0.793	0.681
2000	0.830	0.781	0.687
2002	0.838	0.796	0.729
2004	0.840	0.803	0.742
2006	0.846	0.802	0.735
2008	0.829	0.790	0.739
2010	0.848	0.811	0.774

Notes: Wealth mobility between survey dates based on PSID data. Columns show the wealth group and rows the initial survey year. Mobility is shown as the share of households who remain in the wealth group between survey dates. Difference between survey dates is five years for the first three surveys and two years starting in 1998.

2.A.4 Sensitivity analysis

This section provides a sensitivity analysis of Gini coefficients for income and wealth and explores the diverging income and wealth inequality trends using quantile ratios instead of group averages. For Gini coefficients, we explore the effect of excluding negative-wealth and income observations and the effect of imputing rents from owner-occupied housing to income. The sensitivity analysis complements the discussion from Section 2.3. For quantile ratios, we explore changes of the 95-50 and 90-50 ratios for income and wealth relative to 1971 to demonstrate that the pattern of diverging income and wealth inequality trends from Section 2.4 do not depend on trends in the very right tail of the distribution.

2.A.4.1 Excluding negative income and wealth observations

Figure 2.A.9 shows the effect of excluding negative wealth (income) observations when computing Gini coefficients. We find that the Gini coefficients are lower when excluding negative wealth observations, but the time trends remain unaffected. For income, the effect from excluding negative income observations is negligible. Negative income observations can result, for example, from business losses.

2.A.4.2 Imputed rental income of owner occupiers

The main part of the paper does not include rental income of owner-occupied housing in total income. This section uses rental yields from Jordà et al. (forthcoming) to impute

Figure 2.A.9: Gini coefficients for income and wealth

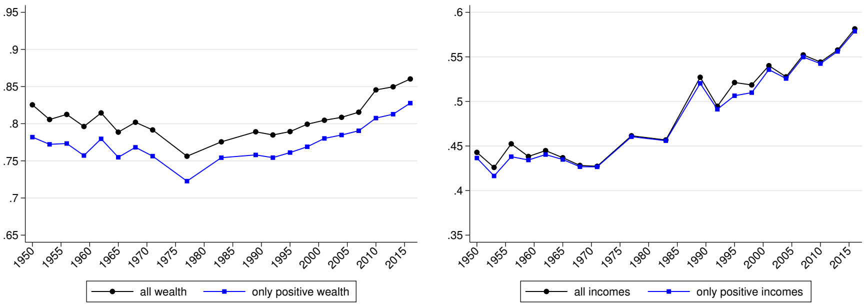

(a) wealth (b) income

Notes: Gini coefficients for income and wealth over time. Black lines with circles show Gini coeffients if all observations are included, blue lines with squares show Gini coefficients if only positive values for income or wealth are included.

rental income for homeowners. Rental yields are average rental yields for the U.S. On average, imputed rental income accounts for 8% to 12% for households in the 50% - 90% of the wealth distribution and 6% to 10% for the top 10%. For the bottom 50%, the share accounts for only 2% to 6% due to lower homeownership rates in this part of the distribution. The share of imputed income is rising over time for all wealth groups. Including imputed rental income slightly decreases measured income inequality, but the overall effects are small (Figure 2.A.10).

Figure 2.A.10: Gini coefficient for income with imputed rents

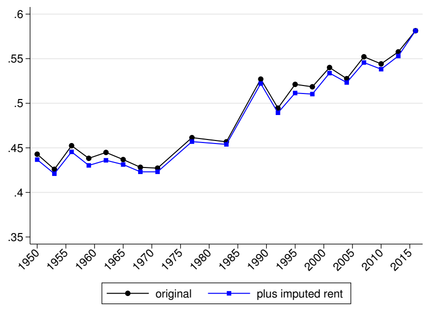

Notes: Gini coefficient for income over time. Black line with dots shows income Gini for baseline definition of income. Blue line with squares shows income Gini with imputed rents for owner-occupied housing using rental yields from Jordà et al. (forthcoming).

2.A.4.3 Trends in quantile ratios for income and wealth

In Section 2.4, we document diverging income and wealth inequality trends for the period between 1971 and 2007 based on group averages. We also document a strong rise in wealth inequality after 2007. For the top 10%, the trends in average income or wealth could be driven by changes in the very right tail of the distribution. Figure 2.A.11 shows changes in the 95-50 and 90-50 ratios for income and wealth relative to 1971. These quantile ratios are robust to particularly high income or wealth growth in the very right tail of the distribution. We find that the stylized fact of hardly changing wealth inequality between 1971 and 2007 is a robust finding also when considering quantile ratios.[37] For income inequality, we also find the documented trends of rising income inequality since 1971 to be robust. The increase in wealth inequality after 2007 shows up even more strongly for the quantile ratios. Both the 95-50 and the 90-50 ratio jump up after 2007 and remain higher compared to historical levels.

Figure 2.A.11: Changes in quantile ratios over time

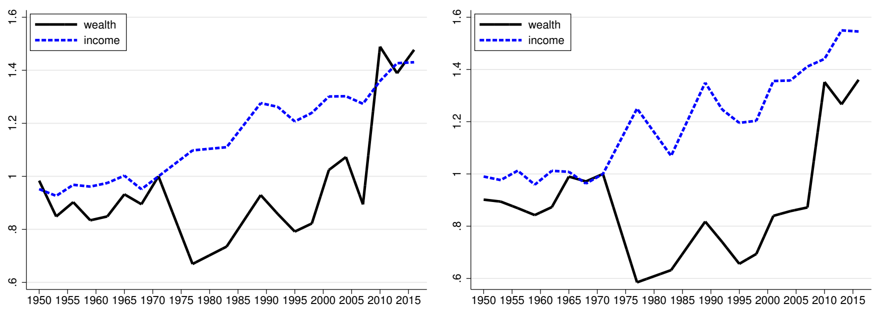

(a) 90-50 ratios

(b) 95-50 ratios

Notes: Change in quantile ratios for income and wealth relative to 1971 (= 1). The left panel shows the growth of the 90-50 quantile ratios of income and wealth relative to 1971. The right panel shows the growth of the 95-50 quantile ratios of income and wealth relative to 1971.

2.A.4.4 Effect of price deflator on growth trends

In the main part of the paper, we adjust all income and wealth levels for inflation using the consumer price index (CPI). We discuss concerns that the CPI may overstate inflation and bias income and wealth growth downward over time in Section 2.2.1. In Figure 2.A.12, we replicate the results on income and wealth growth from Figure 2.4.1 using PCE inflation rates instead of CPI inflation rates. Importantly, different price indices will affect the level of growth rates but will not change the relative growth trends between groups that are important for changes in inequality. Comparing Figure 2.A.12 to Figure 2.4.1 from the main text, we find that income growth between 1971 and 2007 is higher. While the differences in growth rates between 1971 and 2007 are of minor importance for the 50%-90% and the top 10% of the wealth distribution, there is an important qualitative change in growth rates for the bottom 50% of the wealth distribution. For this group, we now obtain positive income growth between 1971 and 2007. Wealth growth has been very high across all wealth groups between 1971 and 2007 so that changes in growth rates from using PCE instead of CPI are of minor importance.

[37] We also explored trends in the 99-50 ratios and found very similar results.

Figure 2.A.12: Income and wealth growth along the wealth distribution (PCE adjusted)

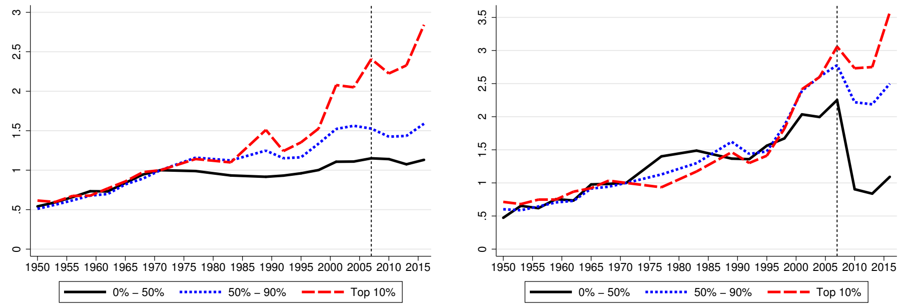

(a) Income growth (b) Wealth growth

Notes: Income and wealth growth for different groups along the wealth distribution using PCE inflation adjustment. Left panel shows income growth for three groups in the wealth distribution: the bottom 50% (black solid line), the middle class 50%-90% (blue dotted line), and the top 10% (red dashed line). All income time series are indexed to 1 in 1971. Right panel shows wealth growth for the same three groups along the wealth distribution. The vertical lines in both panels indicate the 2007 survey.

2.A.5 Additional results

This section provides additional results that complement the analysis on portfolio composition from Section 2.4. First, we show Gini coefficients for different asset classes. Second, we provide a decomposition of house price exposure for different wealth groups. Finally, we report the estimated time series for Gini coefficients, wealth and income shares, and portfolio shares.

2.A.5.1 Gini coefficients for different asset classes

Section 2.4.2 in the main paper documents systematic differences in the household portfolios along the wealth distribution. We document that the distribution of stock holdings is highly skewed, while houses are in comparison relatively equally distributed. Here we offer an alternative view on the distribution of assets in the population by looking at Gini coefficients within asset classes. Kuhn and Ríos-Rull (2016) report large differences in inequality of asset holdings within asset classes. The SCF+ data allow us to extend such an analysis over the long run and document that such inequalities in the asset distributions have been a long-run phenomenon. Figure 2.A.13 presents Gini coefficients for different asset classes. The time series are reported in Table 2.A.3.

Corroborating the pattern from Figure 2.4.4, we find that housing is the most equally distributed asset, with a Gini coefficient fluctuating around 0.6 and only recently exceeding 0.7. We observe a slight upward trend since 1960. By contrast, business equity and stocks show a very high degree of inequality, with high and stable Gini coefficients in excess of 0.95.

Table 2.A.3: Gini coefficients for different asset classes over time

year	housing + nonfin. assets	stocks + business	bonds + liq. and oth. fin. assets
1950	0.65	0.96	0.81
1953	0.63	0.96	0.80
1956	0.61	0.95	0.81
1959	0.59	0.95	0.81
1962	0.62	0.96	0.83
1965	0.61	0.94	0.81
1968	0.63	0.94	0.78
1971	0.63	0.94	0.83
1977	0.61	0.96	0.79
1983	0.66	0.96	0.81
1989	0.69	0.96	0.82
1992	0.67	0.96	0.81
1995	0.64	0.96	0.82
1998	0.65	0.95	0.79
2001	0.67	0.95	0.81
2004	0.67	0.96	0.82
2007	0.67	0.96	0.82
2010	0.68	0.96	0.84
2013	0.69	0.97	0.84
2016	0.71	0.97	0.84

Notes: Gini coefficients for different asset classes over time. Survey year reported in first column. Second column reports Gini coefficients for the sum of housing and non-financial assets, Gini coefficients for sum of stocks and business equity reported in third column, fourth column reports Gini coefficients for the sum of bonds, liquid assets, and other financial assets. See text for details on definitions of asset components.

Figure 2.A.13: Gini coefficients for different asset classes over time

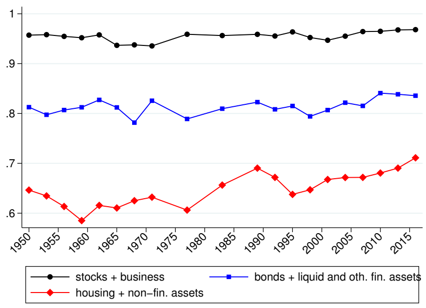

Notes: Gini coefficients for different asset classes over time. Black circles show Gini coefficients for sum of stocks and business equity. Blue squares show Gini coefficients for the sum of bonds, liquid assets, and other financial assets, and red diamonds show Gini coefficients for the sum of housing and non-financial assets. See text for details on definitions of asset components.

2.A.5.2 Decomposition of house price exposure

The house price elasticity of wealth from Figure 2.4.5 in the main part of the paper can be further broken down into a *diversification component* that is determined by the share of housing in assets and a *leverage component* measured by the debt-to-wealth ratio

$$\frac{\text{Housing}}{\text{Wealth}} = \underbrace{\frac{\text{Housing}}{\text{Assets}}}_{\text{diversification}} \times \left(1 + \underbrace{\frac{\text{Debt}}{\text{Wealth}}}_{\text{leverage}}\right).$$

Figure 2.A.14 shows these two components of house price exposure for the middle class and the top 10% over time. Panels (a) and (c) show the diversification component, panels (b) and (d) the leverage component.

We find that the bottom 90% are always less diversified in their asset holdings and that the share of housing in total assets remained more or less constant at 60% over the past seven decades. The top 10% are substantially more diversified with housing accounting for at most 20% of their assets (Figures 2.A.14a and 2.A.14c). The middle class (50%-90%) is substantially much more leveraged than the top 10% and their leverage increased over time. In 1950, debt was equivalent to only about 10% of wealth. By 2016, this number has almost tripled (Figure 2.A.14b). For the top 10%, as share of wealth also increased over time but accounts for only 5% of their total wealth by 2016. Leverage for the bottom 50% is substantially higher (Figure 2.A.14d). Until 2007, debt was on average roughly as much as their total debt. This changed dramatically during the financial crisis and its aftermath. In 2016, the ratio of debt to wealth for the bottom 50% is roughly 3. Strong exposure from low diversification and high leverage is not itself the result of rising house prices. Even in the 30 years between 1950 and 1980 — when real house prices were relatively stable (Knoll et al., 2017) — the bottom 90% were highly exposed to the housing market.

Figure 2.A.14: Components of house price exposure by wealth group

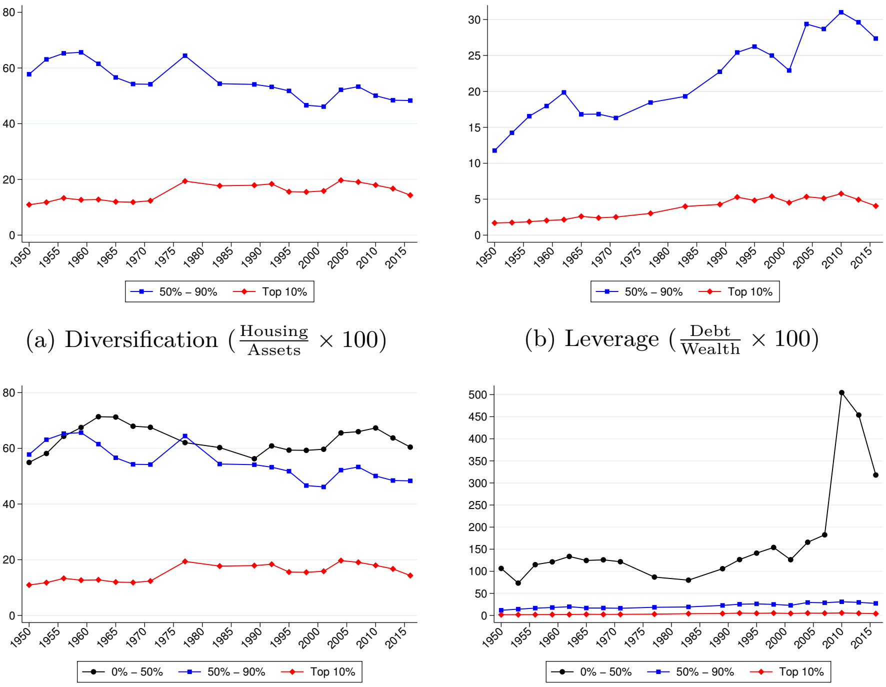

(a) Diversification ($\frac{\text{Housing}}{\text{Assets}} \times 100$)

(b) Leverage ($\frac{\text{Debt}}{\text{Wealth}} \times 100$)

(c) Diversification ($\frac{\text{Housing}}{\text{Assets}} \times 100$)

(d) Leverage ($\frac{\text{Debt}}{\text{Wealth}} \times 100$)

Notes: Decomposition of house price exposure for households in the bottom 50%, the 50%-90%, and the top 10% of the wealth distribution. Panels (a) and (c) show the *diversification* component, panels (b) and (d) the *leverage* component. See text for further details on decomposition.

2.A.5.3 Time series of income and wealth shares

Table 2.A.4 shows income shares for three income and wealth groups over time. The groups are the bottom 50%, the 50%-90%, and the top 10%. Table 2.3.2 in the main paper shows the data for selected years. The last three columns show the income shares by wealth groups corresponding to the discussion in Section 2.4 in the paper.

Table 2.A.4: Shares in aggregate income and wealth

	income shares by income groups			wealth shares by wealth groups			income shares by wealth groups		
year	0%-50%	50%-90%	top 10%	0%-50%	50%-90%	top 10%	0%-50%	50%-90%	top 10%
1950	21.6	43.9	34.5	2.3	24.8	72.8	33.5	37.7	28.7
1953	22.2	45.9	31.8	3.3	25.0	71.7	34.6	39.1	26.3
1956	21.2	44.6	34.3	2.9	25.1	72.1	34.6	38.7	26.7
1959	21.3	45.8	32.9	3.4	26.7	69.9	35.7	39.4	24.9
1962	20.9	45.7	33.4	3.0	24.6	72.4	34.1	38.8	27.0
1965	21.5	45.4	33.1	3.5	27.6	68.8	34.1	39.7	26.2
1968	21.6	46.6	31.9	3.2	26.1	70.7	34.4	38.8	26.8
1971	21.6	47.0	31.4	3.3	27.9	68.8	33.9	40.5	25.6
1977	19.9	45.1	35.0	4.6	31.3	64.1	30.6	42.8	26.6
1983	19.6	46.9	33.5	4.1	29.7	66.2	30.1	43.2	26.7
1989	16.3	43.7	40.0	3.0	30.0	67.0	25.9	42.0	32.1
1992	17.7	45.7	36.5	3.4	29.8	66.8	28.7	42.4	28.9
1995	16.4	45.3	38.3	3.6	28.6	67.8	28.5	41.2	30.3
1998	16.6	43.9	39.5	3.0	28.5	68.5	26.7	42.5	30.7
2001	15.9	41.7	42.3	2.8	27.8	69.4	24.7	40.4	34.9
2004	16.6	42.4	41.0	2.6	28.1	69.3	24.6	41.3	34.2
2007	15.5	40.3	44.2	2.5	26.2	71.3	24.0	38.0	37.9
2010	16.0	40.2	43.7	1.2	24.4	74.4	25.3	37.6	37.1
2013	15.3	40.1	44.6	1.1	24.1	74.8	23.7	37.7	38.7
2016	14.6	37.9	47.5	1.1	21.8	77.1	21.9	36.7	41.4

Notes: Income and wealth shares in total income and wealth. First column reports survey year. Columns 2-4 report shares in total income for bottom 50%, 50%-90%, and top 10% of the income distribution. Columns 5-7 report shares in total wealth for bottom 50%, 50%-90%, and top 10% of the wealth distribution. Columns 8-10 report shares in total income for bottom 50%, 50%-90%, and top 10% of the wealth distribution.

2.A.5.4 Time series of Gini coefficients

Table 2.A.5 shows the time series of Gini coefficients over time. We discuss the observed time trends in Section 2.3.1 of the main paper.

Table 2.A.5: Gini coefficients for income and wealth

	Income			Wealth		
year	all	bottom 99%	bottom 90%	all	bottom 99%	bottom 90%
1950	44	39	32	83	75	64
1953	43	39	32	81	74	59
1956	45	41	34	81	75	62
1959	44	39	33	80	74	60
1962	44	40	33	81	75	61
1965	44	39	32	79	74	61
1968	43	39	33	80	74	61
1971	43	39	33	79	74	62
1977	46	42	34	76	71	58
1983	46	41	35	78	69	59
1989	53	46	39	79	72	62
1992	49	45	38	78	71	61
1995	52	46	39	79	70	60
1998	52	45	38	80	72	62
2001	54	46	38	80	73	63
2004	53	46	38	81	73	64
2007	55	47	38	82	74	63
2010	54	47	37	85	79	71
2013	56	48	38	85	79	71
2016	58	49	39	86	79	70

Notes: Gini coefficients for income and wealth for all households, bottom 99%, and 90% of the income and wealth distribution. Gini coefficients for income shown in the left part of the table (columns 2-4). Bottom 99% and bottom 90% refer to the income distribution. Gini coefficients for wealth shown in the right part of the table (columns 5-7). Bottom 99% and bottom 90% refer to the wealth distribution.

2.A.5.5 Time series of portfolio composition

Tables 2.A.6, 2.A.7, and 2.A.8 show the portfolio composition of households for the three wealth groups considered in the main paper. These groups are the bottom 50%, the 50%-90%, and the top 10%. The first six columns show shares in assets, the next two columns show shares in debt, and the last column shows the debt-to-asset ratio.

Table 2.A.6: Shares of wealth components in wealth portfolios of bottom 50% (in%)

year	other nonfin. assets	real estate	bus. wealth	equity	liquid assets, bonds	other fin. assets	non housing debt	housing debt	debt-to-asset ratio
1950	22.1	54.9	1.0	2.2	19.8	0.0	41.5	58.5	51.6
1953	21.9	58.1	0.8	0.6	18.6	0.0	32.2	67.8	42.3
1956	18.4	64.3	0.5	0.7	16.1	0.0	30.8	69.2	53.5
1959	17.2	67.5	0.6	1.4	13.4	0.0	29.3	70.7	54.8
1962	14.8	71.3	0.6	0.7	12.5	0.0	22.2	77.8	57.2
1965	16.4	71.2	0.5	1.9	9.9	0.0	23.6	76.4	55.4
1968	16.7	67.9	0.7	2.1	12.6	0.0	27.9	72.1	55.7
1971	19.6	67.5	0.5	2.0	10.4	0.0	29.9	70.1	54.9
1977	17.8	62.0	0.5	1.4	18.3	0.0	30.0	70.0	46.5
1983	19.1	61.7	1.0	1.4	10.2	6.6	32.5	67.5	44.6
1989	22.4	56.9	1.1	1.4	9.5	8.7	37.9	62.1	51.5
1992	19.4	62.0	2.0	1.1	7.6	7.9	32.6	67.4	55.7
1995	20.9	60.2	1.3	1.1	6.1	10.4	30.5	69.5	58.8
1998	19.0	60.6	1.4	1.8	6.4	10.8	33.2	66.8	60.8
2001	19.6	60.2	1.1	1.8	6.4	10.8	31.4	68.6	56.0
2004	17.6	66.2	1.0	1.3	5.2	8.6	28.3	71.7	62.6
2007	16.5	67.1	1.2	1.0	4.8	9.4	27.5	72.5	64.8
2010	16.5	68.8	1.3	0.4	4.3	8.7	27.2	72.8	83.3
2013	18.9	65.0	1.2	0.7	5.4	8.8	32.6	67.4	81.9
2016	20.3	61.5	1.2	0.8	6.1	10.0	42.3	57.7	76.2

Notes: Portfolio composition for the bottom 50% of households in the wealth distribution. Asset components in columns 2-7 reported as share in total assets. Debt components in columns 8-9 reported as share in total debt. Column 10 reports debt-to-asset ratio for total debt to total assets. See text for definitions of asset and debt components.

Table 2.A.7: Shares of wealth components in wealth portfolios of 50%-90% (in%)

year	other nonfin. assets	real estate	bus. wealth	equity	liquid assets, bonds	other fin. assets	non housing debt	housing debt	debt-to-asset ratio
1950	11.0	57.8	9.5	4.4	17.3	0.0	19.1	80.9	10.5
1953	10.2	63.1	8.8	2.1	15.8	0.0	14.9	85.1	12.5
1956	6.1	65.3	8.4	4.1	16.1	0.0	14.5	85.5	14.2
1959	6.3	65.6	6.2	6.9	15.0	0.0	13.6	86.4	15.2
1962	12.7	61.5	5.0	7.3	13.5	0.0	11.6	88.4	16.6
1965	11.9	56.6	6.9	10.3	14.3	0.0	11.9	88.1	14.4
1968	10.7	54.3	6.7	12.3	16.1	0.0	13.6	86.4	14.4
1971	13.8	54.2	6.9	9.5	15.7	0.0	13.0	87.0	14.0
1977	14.0	64.4	2.7	4.0	14.9	0.0	16.0	84.0	15.6
1983	13.9	56.1	5.9	2.7	13.3	8.2	19.2	80.8	16.2
1989	14.1	55.1	5.4	3.3	10.7	11.4	19.2	80.8	18.7
1992	13.3	54.3	5.1	3.7	10.4	13.2	14.0	86.0	20.4
1995	13.5	52.8	4.4	4.1	8.4	16.9	15.7	84.3	21.0
1998	13.1	47.7	4.7	7.2	9.1	18.2	16.5	83.5	20.2
2001	12.0	47.6	5.7	7.2	7.9	19.7	14.4	85.6	19.0
2004	12.4	53.7	5.1	5.3	7.2	16.3	13.3	86.7	23.1
2007	11.8	55.1	4.2	4.2	6.9	17.9	13.1	86.9	22.6
2010	12.8	51.8	5.1	3.6	7.3	19.5	13.8	86.2	24.0
2013	12.4	50.2	4.0	4.2	7.1	22.0	13.8	86.2	23.3
2016	11.4	50.2	4.4	4.4	7.3	22.2	15.9	84.1	21.9

Notes: Portfolio composition for the middle class (50%-90% of households) in the wealth distribution. Asset components in columns 2-7 reported as share in total assets. Debt components in columns 8-9 reported as share in total debt. Column 10 reports debt-to-asset ratio for total debt to total assets. See text for definitions of asset and debt components.

Table 2.A.8: Shares of wealth components in wealth portfolios of top 10% (in%)

year	other nonfin. assets	real estate	bus. wealth	equity	liquid assets, bonds	other fin. assets	non housing debt	housing debt	debt-to-asset ratio
1950	4.1	10.9	51.2	25.2	8.6	0.0	35.4	64.6	1.7
1953	7.0	11.8	49.3	24.8	7.0	0.0	26.1	73.9	1.7
1956	1.9	13.3	44.9	33.1	6.7	0.0	20.2	79.8	1.8
1959	1.6	12.7	43.2	35.8	6.6	0.0	15.0	85.0	2.0
1962	8.2	12.8	37.8	32.9	8.4	0.0	8.7	91.3	2.1
1965	8.1	12.0	35.1	39.1	5.7	0.0	11.7	88.3	2.5
1968	9.5	11.8	32.9	39.5	6.2	0.0	11.1	88.9	2.3
1971	13.9	12.3	34.1	30.8	8.9	0.0	8.7	91.3	2.5
1977	10.8	19.4	46.1	17.3	6.3	0.0	13.0	87.0	2.9
1983	18.3	19.0	30.7	13.9	12.1	6.0	35.9	64.1	4.0
1989	20.9	18.6	28.7	9.0	12.2	10.5	27.6	72.4	4.2
1992	19.5	19.1	28.1	11.0	10.6	11.6	17.1	82.9	5.1
1995	15.6	16.3	27.5	15.6	10.4	14.6	17.9	82.1	4.8
1998	14.1	16.5	26.6	20.7	7.4	14.6	22.3	77.7	5.4
2001	14.4	17.0	25.1	20.4	7.6	15.4	17.6	82.4	4.6
2004	15.7	20.8	25.3	16.5	8.2	13.6	15.7	84.3	5.3
2007	14.2	19.9	29.6	16.5	6.2	13.6	11.6	88.4	5.0
2010	14.4	18.8	25.2	15.5	9.1	16.9	11.7	88.3	5.6
2013	12.9	17.6	26.1	17.5	7.7	18.2	10.4	89.6	4.9
2016	13.1	15.1	27.6	21.8	6.9	15.6	17.5	82.5	4.1

Notes: Portfolio composition for the top 10% of households in the wealth distribution. Asset components in columns 2-7 reported as share in total assets. Debt components in columns 8-9 reported as share in total debt. Column 10 reports debt-to-asset ratio for total debt to total assets. See text for definitions of asset and debt components.

3

Modigliani Meets Minsky: Inequality, Debt and Financial Fragility in America, 1949-2016

Joint work with Alina Bartscher, Moritz Kuhn and Moritz Schularick

3.1 Introduction

Rising indebtedness of U.S. households is a much-debated phenomenon. The numbers are eye-catching. Between 1950 and the 2008 financial crisis, American household debt has grown fourfold relative to income. In 2010, the household debt-to-income ratio peaked at close to 120%, up from 30% on the eve of WW2. Figure 3.1.1 below shows the trajectory of this secular increase over the past seven decades. The underlying drivers of the process remain, however, controversial.

Rising income inequality is frequently invoked as an important factor. The second line in Figure 3.1.1 shows that the share of the richest 10% of households in total household income has increased from below 35% to almost 50% between 1950 and 2016. Rajan's (2011) influential book *Fault Lines* has popularized the view that growing income inequality and indebtedness are two sides of the same coin. The idea is that households with stagnant incomes have increasingly relied on debt to finance consumption – be it out of sheer necessity to "get by", or to "keep up with the Joneses" at the top of the income distribution whose incomes were growing nicely (cf. Fligstein et al. 2017). A recent paper by Mian et al. (2019) discusses how rising income concentration at the top brought about a "savings glut of the rich" that supplied the funds for increased borrowing by non-rich households.

But we still know surprisingly little about the borrowers and their financial situation.

From the borrower perspective, the financial history of the growth of U.S. household debt and its distribution remains largely unwritten. This paper closes this gap. We study the dynamics of household debt over the entire postwar period. We ask which households borrowed so much more, and why. Without long-run household-level data for the joint distributions of income, debt, and assets this would be a daunting task. However, we can rely on a new dataset that combines historical waves of the Survey of Consumer Finances (SCF), going back to 1949, with the modern SCF that the Federal Reserve Board administers since 1983 (see Kuhn et al. forthcoming). This long-run "SCF+" makes it possible to follow the evolution of household borrowing across the entire income distribution over seven decades. Where needed, we also combine information from the cross-sectional SCF+ data with data from the Panel Study of Income Dynamics (PSID), which provide panel data on housing wealth and mortgage since 1968.

The data support the much talked-about association between rising income inequality and increased borrowing. Debt growth was concentrated among households with low income growth. Debt-to-income ratios have risen most strongly for households whose share in aggregate income has fallen. Middle-class households, defined here as households between the 50th and 90th percentile of the income distribution, went most into debt. Higher borrowing by middle-class households accounts for 55% of the total increase in household debt since 1950. By contrast, households in the bottom 50% of the income distribution account for a relatively small share of the total debt increase (15%). While their debt-to-income ratio has risen, too, their share in total debt has fallen. The American household debt boom of the past decades is first and foremost a middle-class affair.

The transformation of middle-class balance sheets in the past four decades was compre-

Figure 3.1.1: Debt-to-income ratio and top-10 income share, 1950-2016

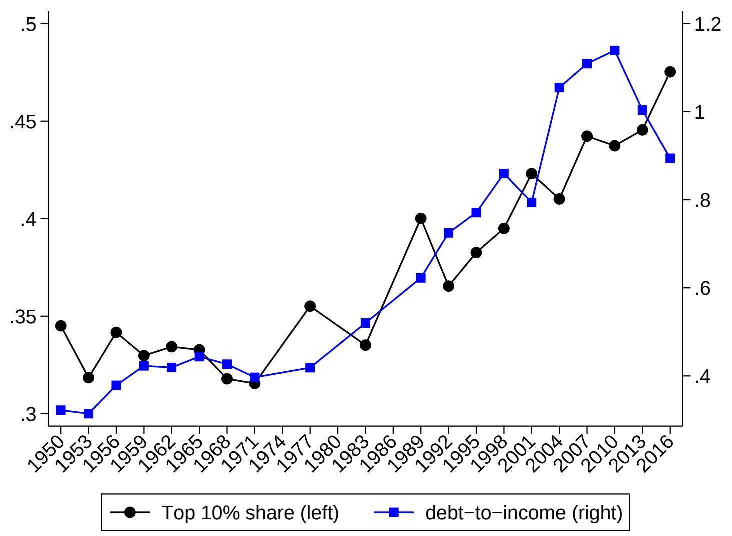

Notes: The graph shows the share of the top 10% of the income distribution (left axis) and the household debt-to-income ratio (right axis) over time.

hensive. Adjusting by the consumer price index (CPI), the average incomes of households in the 50th to 90th percentile of the income distribution have grown by about 25% since the 1970s, or less than half a percent per year. Over the same period, the amount of debt taken out by these households grew by 250% until the crisis, about ten times faster. A similar picture emerges for households right at the median of the income distribution. Here, income growth was barely positive in CPI-adjusted terms between 1971 and 2007, but debt grew by a factor of almost ten at the median. This association between low income growth and high borrowing is puzzling. In standard economic logic, households are typically expected to borrow against the expectation of higher, not lower or stagnant future income.

How can one rationalize this behavior? Here the strength of the SCF+ data with respect to its comprehensive coverage of the entire household balance sheet comes into play and leads to an important insight. A plausible suspicion would be that with rising debt, the net wealth of middle-class households decreased. After all, the liability side of the typical middle-class balance sheet grew substantially. Yet this is not the case. The net wealth position of middle class households actually improved. Households borrowed more, but at the same time became (wealth-) richer. Simple balance sheet accounting dictates that this is possible only if the value of household assets increased even faster than their debt. In the absence of a substantial increase of savings out of stagnant incomes, this can happen only if the value of existing assets rises. The explanation for the U.S. household debt boom that we put forward in this paper builds on this disconnect between income and asset growth that is evident in the SCF+.

The housing market played the central role in this process. We will show that owing to their high exposure to house prices, middle-class American families made sizable wealth gains when their main asset, residential real estate, appreciated in price. In inflation-adjusted terms, quality-adjusted house prices in the U.S. increased by 75% between the mid-1970s and the mid-2000s. Housing wealth-to-income ratios of middle-class households more than doubled from 140% of income to 300% in 2007, with price effects alone accounting for close to 50% of this increase. In other words, income growth of middle-class households was low, but at the same time their housing wealth grew strongly. Wealth-to-income ratios increased even more for these households than at the top.

From here our analysis essentially follows the logic of the canonical Modigliani life cycle model (Modigliani et al. 1954). When middle-class households racked up sizeable gains in housing wealth, they used debt to turn higher life-time wealth into additional expenditures. We show that the combined effects of home equity extraction through refinancing, HELOCs, and second mortgages were quantitatively large and explain a substantial share of the increase in household debt since the 1970s. Debt is key for the response to the wealth shock because housing is a peculiar asset. A key characteristic is that it is indivisible, meaning it cannot be sold in small increments, unlike, for instance, equities. When the stock market rises, households can sell some shares and use the proceeds for consumption. Turning housing wealth gains into additional expenditures (while continuing to live in the same house) is possible only by taking on debt.

The Panel Study of Income Dynamics (PSID) contains data on housing wealth and mortgages that allow us to identify home equity-extracting households and quantify the aggregate effects of home equity-based borrowing since the 1980s. Using the PSID, we decompose the debt increase into additional debt incurred by extractors, new homeowners, and upgraders moving to a larger home. We find that home equity-based borrowing against existing owner-occupied real estate accounts for around 50% of the increase in housing debt since the 1980s. From the early 1980s to the 2008 crisis, equity extraction only pushed total the household debt-to-income ratio up by more than 30 percentage points.

Without equity extraction the housing debt-to-income ratio would have stayed at around 50% of income until 2008. Home equity extraction averaged around 1.5% of annual income until the mid-1980s, and rose to around 4.5% thereafter. Over a twenty-year period, the cumulative effects of equity extraction were substantial. Importantly, we find that home-equity-based borrowing was responsible for a significant fraction of the rise in U.S. household debt already before the extraction boom of the 2000s that has been studied by Greenspan and Kennedy (2008), Klyuev and Mills (2007), and Mian and Sufi (2011), among others. This is consistent with the findings of Guren et al. (2018), who report substantial housing wealth effects already since the 1980s.

Stratifying equity extraction by income groups, we show that about half of total home-equity-based borrowing is accounted for by middle-class households (50-90%). Local projections at the state level not only confirm a close association between house prices and equity extraction. They also corroborate a much higher elasticity of equity extraction for growth middle-class households whose portfolios are most concentrated in housing and more strongly leveraged.

A large share of the increase in household debt can be rationalized as a Modigliani-style response of middle-class households to capital gains they made in housing markets. We will show that the observed equity extraction is qualitatively and quantitatively in line with the predictions of recent models such as Berger et al. (2017). In their model, a consumption response to housing wealth gains arises as soon as the strict assumptions that underlie the model in Sinai and Souleles (2005) are relaxed. Sinai and Souleles (2005) argue that if houses are handed from generation to generation, and there is no mobility and adjustment in housing size, then housing tenure becomes infinite and house price changes will not affect household consumption. Yet in the presence of life cycle variation in housing size, contemporaneous ownership of housing of parent and children generations, or imperfectly correlated local housing markets and household mobility, rising housing wealth triggers consumption responses of homeowners also in their model.[1]

[1] The positive net response in Berger et al. (2017) also results from an additional substitution effect that Sinai and Souleles (2005) rule out by construction. Berger et al. (2017) interpret the net effect as an endowment effect with income, substitution, and collateral effects canceling out. Campbell and Cocco (2007) also discuss the result from Sinai and Souleles (2005) and argue that changing life cycle housing

The intuition is straightforward. When homeowners make capital gains in the housing market, they are richer than they expected when originally making their financial planning. As housing is indivisible, households need to liquidize some of their home equity if they want to smooth consumption over time. In principle, households could also sell their house and move to a smaller home. However, this would involve substantial transaction, search and potentially also emotional costs (see Aladangady 2017), and few households do this in practice, as the PSID shows. The remaining option is to engage in negative savings (equity extraction) after the deviation from the life-cycle wealth profile. The reason for the house price increase is irrelevant, as long as it was unexpected when financial plans were made, and is assumed to persist.

Empirical evidence for recent years supports the theoretical argument that housing wealth effects are substantial. Based on matched microdata, Aladangady (2017) estimates a causal effect of house prices on consumption of around 5 cents per dollar increase of home value. Mian and Sufi (2014) explicitly consider the response of household debt to house price shocks. They exploit regional heterogeneity in the U.S. and also find substantial effects that can be rationalized in the context of recent models with liquidity-constrained consumers, such as (Kaplan and Violante, 2014).

Taken together, these findings lead us to a more nuanced interpretation of the postwar household debt boom. It is true that middle-class families with low average income growth were chiefly responsible for increased borrowing. It is also true that these households relied on debt to finance consumption in the face of stagnant incomes. But they could do so because they had become richer, at least for the time being. It is obviously possible that households, in particular in the later years of the boom of the 2000s, mistakenly treated house price increases as persistent when they were not.

Note that this history of household debt in America is compatible with the idea of a savings glut, arising either from global factors (Bernanke 2005) or growing income concentration at the top (Mian et al. 2019), that lowered interest rates, made borrowing cheaper and increased housing values. Our analysis does not speak to the initial trigger of this process. Rising income inequality might well have played a role. The argument we make is that once the house price increase was underway, home-owning middle-class households made large wealth gains from rising house prices and increasingly turned those wealth gains into spending via home-equity-based borrowing. Clearly, the fact that interest rates kept on falling despite rising borrowing volumes meshes nicely with the idea of a credit-supply-driven household debt boom. We discuss the importance of enabling factors such as financial deregulation and the 1986 tax reform that maintained interest deductibility for mortgage lending and thereby created incentives to switch to home-equity-based products. Story (2008) describes how banks advertised these new products heavily with catchphrases such as: "Now, when the value of your home goes up, you can take credit for it."

In the last part of the paper we discuss how this rational response of Modigliani house-

demand leads to an age-varying endowment effect from house price shocks.

holds leads to a more fragile, high-debt economy. Home-equity-based borrowing may be optimal from an individual household's point of view, but in the process balance sheets are extended and become more sensitive to shocks. We document this "Minsky" aspect of the debt build-up by conducting a quantitative assessment of household balance sheets akin to stress test for banks, similar to Fuster et al. (2018). We trace the results of this stress test over seen decades of postwar history and show the increased vulnerability of households. This connects our paper to a lively research agenda concerned with the effects of shocks to household balance sheets on macroeconomic activity (see e.g. Mian and Sufi 2009, Mian and Sufi 2017, Jordà et al. 2013), as well as the interactions between housing and credit markets (Guerrieri and Uhlig 2016).

In any given year, we "shock" households with an exogenous income decline based on estimates for earnings losses in recessions Davis and von Wachter (2011). We then construct a measure for the total value of mortgage debt that is owed by "at risk" households whose liquidity is severely weakened after the shock. Following the literature, we define households as being "at risk" if their debt service ratio crosses 40% of income.

Across the stress scenarios, the increase in financial fragility, measured by the value of loans at risk, turns out to be sizable, especially for middle-class households. From the 1950s to the 1970s, the value of outstanding middle-class mortgage debt that was at risk following an income shock was in the range of 1-2% of total income. By 2007, owing to larger balance sheets and higher debt, the value of mortgage debt at risk had risen to 15% of middle-class income, or about 150% of total bank equity in the system. Clearly, not all of these "at risk" mortgages end up in default, but the strain on household balance sheets and on the financial system is considerable.

Literature: The analysis of household balance sheets and their importance for the business cycles and financial fragility has become an active research field for macroeconomists (Mian and Sufi 2014, 2017, Zinman 2015, Jordà et al. 2013, Adelino et al. (2018), Albanesi et al. 2017). A large empirical and theoretical literature has examined wealth effects due to house price increases and their consequences for household borrowing and consumption.[2] Empirical trends in household indebtedness have been discussed in Dynan and Kohn (2007) and Wolff (2010). Dynan and Kohn (2007) provide an early analysis of the 1990s debt boom and discuss potential sources for the rise in indebtedness of U.S. households. They point already to an important role of mortgage debt and document its co-movement with house prices. Wolff (2010) provides a broader perspective on the change in household finances that emphasizes the rise in middle-class debt since 1983.

Regarding house prices and credit conditions, several important papers have traced house price increases to regulatory changes since the 1980s (e.g. Hoffmann and Stewen 2019, Favara and Imbs 2015, Di Maggio and Kermani 2017). Recent research has also

[2]Iacoviello (2005), Hurst and Stafford (2004), Calomiris et al. (2013), Aladangady (2017), Cloyne et al. (2017), Guren et al. (2018), Andersen and Leth-Petersen (2019), Campbell and Cocco (2007), Berger et al. (2017), and Kaplan et al. (2017) among others.

emphasized link between rising inequality and household borrowing (De Stefani 2018, Mian et al. 2019). In their influential work, Mian and Sufi (2009, 2011) argued that household borrowing in low-income regions of the U.S. grew particularly strongly before the 2008 crisis, and was then followed by severe output and employment losses. In a theoretical model, Kumhof et al. (2015) show that higher savings of the rich may lead to a decline in interest rates, which leads to higher borrowing by low- and middle-income households and higher financial fragility. However, Coibion et al. (2020) find that low-income households face higher borrowing costs and reduced access to credit as inequality increases. Adelino et al. (2016) and Albanesi et al. 2017 provide complementary evidence on the debt boom during the 2000s and highlight the important role of the middle class for the debt boom during these years. Adelino et al. (2016) conclude that growth of middle class debt played an important role. Similarly, Foote et al. (2016) study debt growth in the early 2000s across the income distribution and discuss the implications for theoretical models of the debt boom. Our study is also linked to work that discuss policy option to limit the accumulation of excessive leverage when there are externalities on the macro level (Korinek and Simsek 2016, Schmitt-Grohé and Uribe 2016).

Researchers from related fields have also contributed to the discussion. For instance, political sociologists like Krippner (2012) have linked the debt build-up to growing socio-economic pressures. Economic historians have pointed out that the U.S. economy experienced two major financial and economic crises in the 20th century – the Great Depression and the Great Recession – and both were preceded by a sharp rise of income inequality and growing household indebtedness (Piketty and Saez 2006, Schularick and Taylor 2012). In his history of household borrowing in America, Wright (2012) tied the growth of household debt in America to widening income disparities. However, so far a long-run empirical analysis of the dynamics and changing distribution of household debt in the U.S. has been hampered by the absence of high-quality long-run micro-data for the joint distribution of debt, income, and wealth that the SCF+ now provides.

The structure of the paper is as follows. We first introduce and discuss the historical SCF data and show that the micro data closely match aggregate trends. Second, we show that mortgage borrowing of households between the 50th and 90th percentiles of the income distribution accounts for the lion's share of the debt increase. Third, using PSID data, we show that equity extraction in response to higher housing wealth played a central role for the aggregate debt increase. Fourth, we rationalize our empirical findings in the context of a Modigliani life-cycle model. Finally, we turn to the Minsky-side of the debt increase and show that in particular the financial fragility of middle-class households has risen substantially over time.

3.2 Data

Our paper relies on a new data source that allows us to track the financial history of debt in the United States since World War II along the income distribution. The "SCF+"

combines historical waves of the Survey of Consumer Finances (SCF) going back to 1949 with the modern waves available since 1983. The historical files are kept at the Inter-University Consortium for Political and Social Research (ICPSR).

Kuhn et al. (forthcoming) give a detailed description of the construction of SCF+, including demographic details, the coverage of rich households, and its strength in providing the joint distributions of income, debt, and wealth. The early surveys were carried out annually between 1947 and 1971, and then again in 1977. We follow Kuhn et al. (forthcoming) and use data since 1949, which is the first year in which all relevant variables are available, and pool the early waves into three-year bins.

In the following, we will briefly introduce the data set and discuss how the data match trends from the National Income and Product Accounts (NIPA) and the Financial Accounts (FA). We will also briefly introduce our second main data source, the Panel Study of Income Dynamics (PSID), that we rely on the corroborate the cross-sectional information from the SCF+ with data that provide a panel dimension.

We complement the microdata with data from the *Macrohistory Database* (Jordà et al. 2017), in particular house prices and the consumer price index (CPI). The house price index in the *Macrohistory Database* is based on the index of Shiller (2009) until 1974, and the repeat sales index of the Federal Housing Finance Agency (FHFA, former OFHEO) since 1975. These indices are designed to filter out changes in the average quality and size of homes (cf. Rappaport 2007). If not explicitly stated otherwise, all presented results are in real terms, converted to 2016 dollars using the CPI.

3.2.1 Household debt in the SCF+

The SCF is a key resource for research on household finances. Data for the modern survey waves after 1983 are readily available from the website of the U.S. Federal Reserve. The surveys are conducted every three years by the Federal Reserve Board (see Bricker et al. 2017). The comprehensiveness and quality of the SCF data explain its popularity among researchers (see Kuhn and Ríos-Rull 2016 and the references therein).

Adding data from the historical surveys results in a dataset that contains household-level information over the entire postwar period and provides detailed demographic information in addition to financial variables. Important for the current analysis, the SCF+ data contain all variables needed to construct long-run series for the evolution of household debt including its sub-components. The SCF+ data are weighted with post-stratified cross-sectional weights which assure representativeness along several socio-economic characteristics, in particular race, education, age and homeownership.

Total debt consists of housing and non-housing debt. Several recent papers have stressed the importance of real estate investors for the debt boom prior to 2007 (Haughwout et al. 2011, Bhutta 2015, Mian and Sufi 2018, Albanesi et al. 2017, DeFusco et al. 2017). Real estate investors are defined as borrowers with multiple first-lien mortgages.

They accounted for a disproportionately large share of mortgage growth before 2007 compared to their relatively small population share. However, mortgage debt on the principal residence is on average 8 times larger than mortgage debt on other real estate (see Figure 3.A.8). When it comes to housing debt, in this paper we focus only on debt incurred for owner-occupied housing. We treat investment in non-owner-occupied housing like business investment, and only use the net position when calculating wealth.

Non-housing debt includes car loans, education loans, and loans for the purchase of other consumer durables. Data on credit card balances become available after 1970 with the introduction and proliferation of credit cards. Note that the appearance of new financial products like credit cards does not impair the construction of consistent data over time. Implicitly, these products are counted as zero for years before their appearance.

The core of our analysis studies the dynamics of debt along the income distribution. For this, we calculate total income as the sum of wages and salaries plus income from professional practice and self-employment, rental income, interest, dividends, transfer payments as well as business and farm income.

We abstain from any sample selection for most of our analysis. One exception is the decomposition of changes in debt-to-income ratios in section 3.3.3. Here we use households-level ratios, and drop observations with extreme debt-to-income ratios larger than 50 in absolute value. Moreover, we use household-level loan-to-value ratios and debt-service-to-income ratios in Section 3.6, after trimming the largest percent. Our analysis in this part explicitly relies on individual ratios. Otherwise, we use ratios of averages instead of averages of ratios due to their greater robustness to outliers.

3.2.2 Panel data from the PSID

The key strength of the SCF+ is that it allows to study the joint distribution of income and wealth over seven decades. However, the data are in the form of repeated cross-sections, and thus do not allow to track individual households over time. As the analysis in Section 3.4.2 requires a panel dimension, we use data from the Panel Study of Income Dynamics (PSID). While the SCF+ is at the household level, the PSID is at the family level. Therefore, PSID families living together were aggregated into one household for better comparability (cf. Pfeffer et al. 2016). Additional details are given in Appendix 3.A.1.

Following Kaplan et al. (2014), we only use data from the PSID's "SRC sample". Post-stratified cross-sectional survey weights are provided on the PSID web page only for the waves between 1997 and 2003. Therefore, we use the longitudinal family weights provided on the PSID homepage, and post-stratify them to match the same Census variables which we targeted in the post-stratification of the historical SCF waves. We verified that all reported results are similar when using the unweighted PSID data or the original longitudinal PSID weights without post-stratification. Figure 3.A.1 in the appendix compares the PSID data to the SCF+. Overall, the two data sets align very well.

3.2.3 Aggregate trends in SCF+ and NIPA

Aggregated household surveys are not always easy to match to data for the macroeconomy. Measurement concepts can differ, such that even high-quality microdata may not match aggregate data one-to-one. To judge the reliability of the SCF+ data, we start by comparing the aggregate trends in income and household debt in the SCF+ to data from the National Income and Product Accounts (NIPA) and the Financial Accounts (FA).

We index the series to 100 in 1983-1989 to abstract from level differences that can be attributed to different measurement concepts, and focus on comparing growth trends over time. During the base period 1983-1989, the SCF+ data correspond to 89% of NIPA income and 78% of FA debt in levels.[3]

[3]Income components of the NIPA tables that are included are wages and salaries, proprietors' income, rental income, personal income receipts, social security, unemployment insurance, veteran benefits, other transfers, and the net value of other current transfer receipts from business. Mortgages and consumer

Figure 3.2.1: Income and debt in the SCF+ versus NIPA and FA

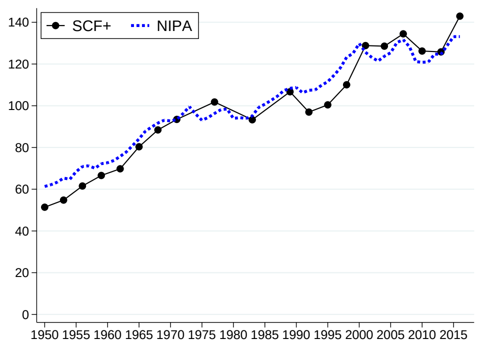

(a) Income

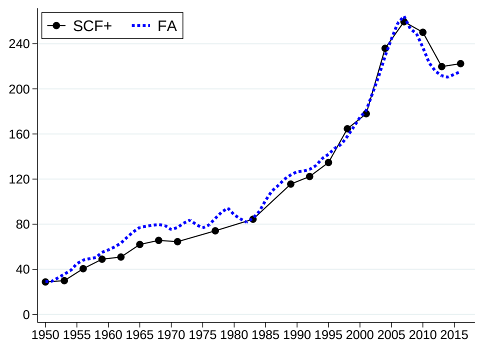

(b) Total debt

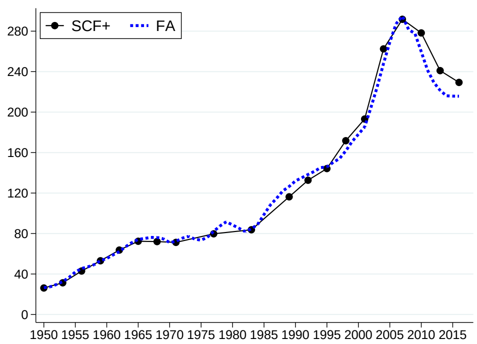

(c) Housing debt

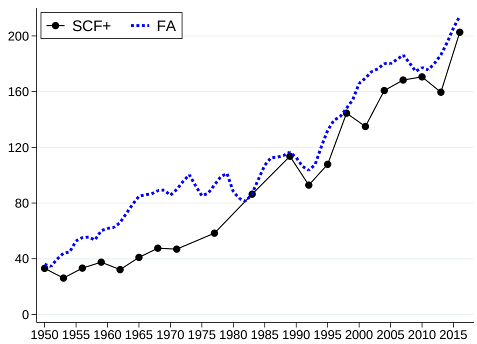

(d) Non-housing debt

Notes: The figure shows income and total debt from the SCF+ in comparison to income data from the NIPA and total debt data from the FA. All series have been indexed to the period 1983 - 1989 (= 100). The SCF+ data are shown as black lines with circles, NIPA and FA data as a blue dashed line. Over the index period, the SCF+ values correspond to 89% for income, 78% of total debt, 80% of housing debt, and 73% for non-housing debt.

Figure 3.2.2: Total and housing debt-to-income ratios

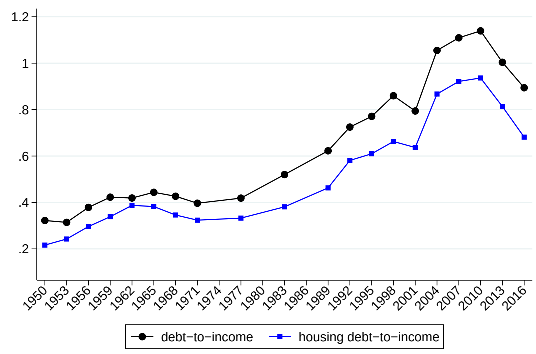

Notes: The graph shows the debt-to-income ratio for total debt and housing debt from the SCF+ over time.

Figure 3.2.1 shows the comparison of growth trends between the SCF+ and aggregate data for 1950 to 2016. Overall, the aggregate data and the aggregated microdata show very similar trends. With respect to housing debt, the SCF+ data and the FA match almost perfectly. Non-housing debt also aligns well with the FA data, albeit there is a certain discrepancy before the 1980s. All in all, the close alignment in growth trends effectively alleviates concerns that the microdata systematically miss parts of the distributional changes underlying the observed macroeconomic growth trends.

Figure 3.2.2 shows the evolution of debt-to-income ratios over the last seven decades. Debt-to-income ratios effectively quadrupled between 1950 and the 2007 crisis. They have fallen by about 20 percentage points since then. Housing debt accounts for 78% of the increase in the debt-to-income ratio from 30% to 92% between 1950 and 2016.

This long-run increase in household indebtedness is well documented on the macro level in the FA statistics. However, with the SCF+ data, we are in a position to track the historical evolution of the distribution of household debt and study its drivers.

3.3 The American household debt boom, 1950-2016

In this section, we will use the SCF+ to track the growth and distribution of household debt and its relation to income dynamics over the past seven decades. Which households borrowed so much more, and for what purposes?

The analysis will proceed in three steps. We will first look at the distribution of debt among income groups over time, and establish that the middle class accounts for the largest part of both outstanding debt and new borrowing. In a second step, we will decompose the overall debt increase into changes at the intensive and extensive margins of different debt components. In a last step, we will shift our focus from the income to the age distribution, and study the changing life-cycle patterns of household debt.

credit are included as FA debt components.

3.3.1 The distribution of household debt

How is household debt distributed among rich and poor households, and how has this distribution changed over time? To address these questions, we stratify households by income. Following standard practices in the literature, we divide the population into three groups according to their position in the income distribution (see Piketty and Saez 2003, Saez and Zucman 2016, Alvaredo et al. 2018).

The first group is made up by households in the bottom 50% of the income distribution, and the second covers households between the 50th and 90th percentile. We refer to this group as the "middle class" throughout the paper. The third group consist of the top 10% of the income distribution. We will only occasionally talk about the top 1% to illustrate dynamics at the very top. Even very rich households owe considerable amounts of debt despite high net wealth (with tax considerations likely playing an important role). Yet as borrowers they are not central for trends in aggregate debt. This being said, very top incomes might have played an important role for the supply of funds (see Mian et al. 2019). Before we study the evolution of debt shares and debt-to-income ratios of these different groups over time, it is important to recognize that the SCF+ is a repeated cross-section. This means that households can move between income groups over time. Our groups are reasonably large so that inter-group mobility can be expected to be low, but we will use PSID panel data to test this assumption, along the lines of Díaz-Giménez et al. (2011). The PSID reveals that around 84% of households in the bottom 50% already were in this group two years ago. The numbers for the 50-90% and top 10% are 75% and 66%, respectively. When we extend the intervals to six years, the share of households who are in the same group six years later is still 77% for the bottom half, 68% for the middle class, and 53% for the top 10%. Moreover, households that change income groups tend to remain close to the "border" with the previous group. For instance, among households who changed into the middle-class between, 64% were no more than two deciles away from this group two years earlier. On average, households remain in the same income group for 77% of the periods in which we observe their income.[4]

Figure 3.3.1 shows the share of total debt owed by the three different income groups. Debt shares have been rather stable over time. Over the entire postwar period, middle-class households have always accounted for the largest share of total debt, on average about 50 to 60% of total outstanding debt. Low-income households in the bottom half make up another 20%. The debt share of the top 10% fluctuated around 20% before the 1980s, and then increased to around 30%. It is clear from Figure 3.3.1 that the upper half of the income distribution has always accounted for about 80% of total household debt outstanding.

[4] As a further robustness check, Figure 3.A.9 in the Appendix presents additional evidence for income group stability. It shows income and housing debt, two key variables for our analyses, for households aged 30 to 55. We examine if the trends in debt look different depending on whether we sort households using their contemporaneous income, or the initial income at the beginning of a decade. The trends look very similar.

Figure 3.3.1: Debt shares by income group

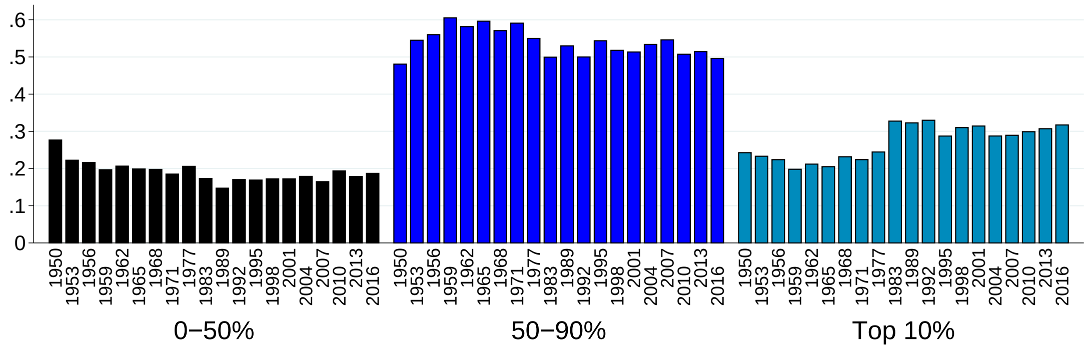

Notes: The figure shows shares in total debt for the different income groups over time.

Figure 3.3.2: Share of increase in debt, 1950-2007

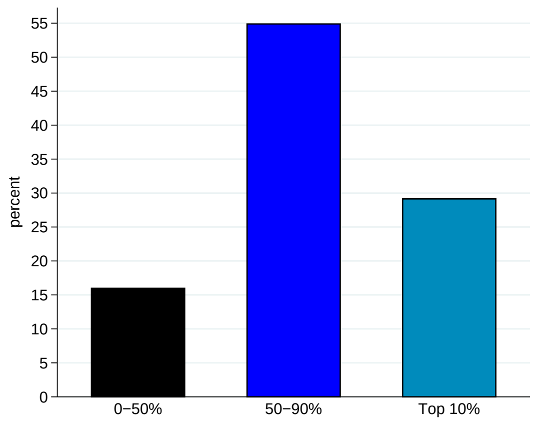

Notes: The graph shows the share of each income group in the total increase of household debt from 1950 to 2007.

It follows from the relative stability of the debt shares over the past seven decades that the middle class also played the dominant role for the growth of debt. Figure 3.3.2 confirms this visually. From 1950 to 2007, middle-class households accounted for 55% of the total debt increase, whereas households from the bottom 50% of the income distribution only contributed 15%, even less than the top 10% with almost 30%. This is an important insight in itself. 85% of the increase of U.S. household debt occurred within the upper 50% of the income distribution. The explanation for soaring household debt in the U.S. lies in the borrowing behavior of these incomes groups, and in particular of middle-class households (see also Adelino et al. 2018).

We next turn to debt-to-income ratios. There have been substantial changes in the distribution of income in the U.S. over the past 70 years. On a CPI-adjusted basis, the average income of households in the top 10% increased by a factor of 2.5 between 1971 and 2016, while the average income of the middle class grew by only 25%, and that of the bottom 50% stagnated in real terms. Figure 3.3.3 displays the diverging income growth trajectories of the different parts of the American income distribution.

Figure 3.3.3: Income growth

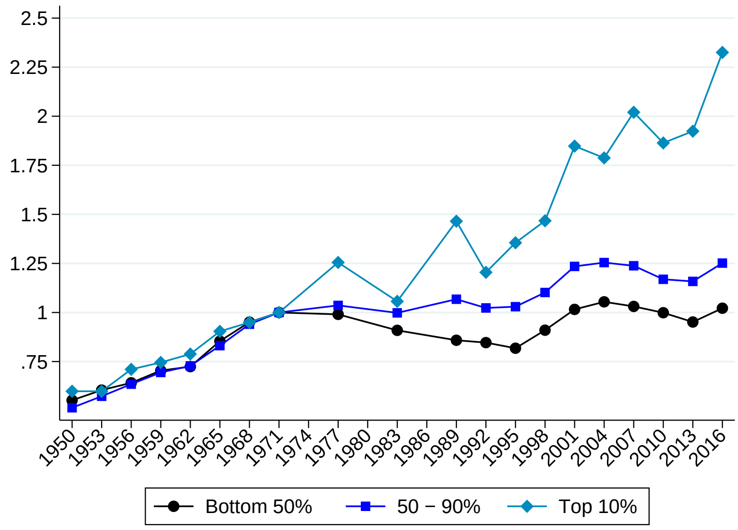

Notes: The graph shows average income of the three income groups from the SCF+. All series are normalized to one in 1971.

These differential trends in income growth across the groups have important consequences for resulting trends in debt-to-income ratios that are shown in Figure 3.3.4. Figure 3.3.4a shows surging debt-to-income ratios for middle class and low income households. For both income groups, debt-to-income ratios rose from around 40% in the early 1950s to close to 140% in 2007. For the top 10%, the increase is much more muted, despite the fact that the group accounts for a higher share in total debt compared to the 1950s. This is because their incomes have risen almost proportionally. Figure 3.A.11 in the Appendix shows that from the 1950s to the 1970s, debt and income have grown at almost identical rates for all three groups, resulting in the observed stability of debt-to-income ratios over this period.

Figure 3.3.4b shows debt-to-income ratios of the top 1%, compared to the bottom 90%. The chart underscores the divergent debt trajectories at the top and in the rest of the economy. For the very top, debt ratios have remained relatively constant. The bottom 90% witnessed a sharp rise in debt-to-income over the past decades. The chart nicely captures that debt-to-income ratios at the top and the bottom evolved in tandem until the late 1970s, and then sharply diverged as income concentration at the top increased. In the past four decades, debt ratios have increased most for parts of the population whose income growth was low.[5]

A more comprehensive picture of the distributional dimension of the American household debt boom emerges from Figure 3.3.5. For different survey waves, the Figure shows the evolution of debt-to-income ratios across the entire distribution. The left-hand side shows total household debt relative to income, and the right-hand side shows housing

[5]Figure 3.A.12 in the Appendix shows that the ratio of aggregate debt to aggregate assets has equally stayed largely flat for high-income households. Both debt-to-income and debt-to-asset ratios have increased most strongly for the middle class.

Figure 3.3.4: Debt-to-income ratios

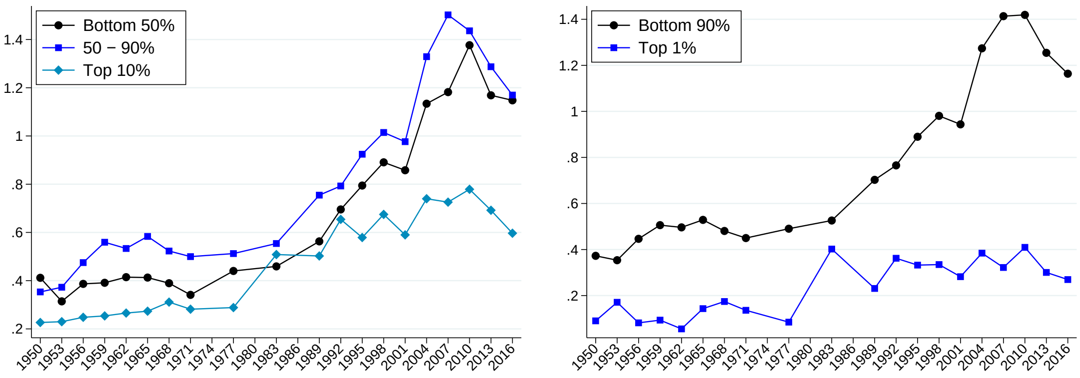

(a) Debt-to-income ratio

(b) Debt-to-income ratio

Notes: The left panel shows housing debt-to-income ratios for the bottom 50%, 50-90% and top 10% of the income distribution. The right panel compares debt-to-income ratios of the bottom 90% and top 1%.

debt ratios only. Debt-to-income ratios were relatively constant in 1950, with debt ratios being less than 50% across the entire income spectrum. By 1983, debt-to-income ratios had increased somewhat, but were not far off their levels in the 1950s. Since then, indebtedness has risen strongly across all income groups, but soaring debt ratios of middle-class households stand out. For households between the 50th and 90th percentile, debt-to-income ratios have approximately tripled within 25 years.[6]

3.3.2 The composition of household debt

As Figure 3.3.5 illustrates, housing debt plays an important role for debt trends of households in the upper half of the income distribution. Adding information on the number of households with outstanding debt and the type of debt, we can decompose the debt increase into its extensive and intensive margin. In other words, we can answer to what extent the total number of indebted households has increased, and to what extent indebted households have taken on larger amounts of debt. Additionally, we can calculate the ex- and intensive margin effects separately for different types of debt, i.e. housing and non-housing debt, including car loans, credit cards and student debt.

Let $d_{i,t}$ stand for the mean debt-to-income ratio of income group i in period t. $s_{i,t}^{H+}$ is the share of households with positive housing debt, i.e., the extensive margin, and $d_{i,t}^{H+}$ is the average housing debt-to-income ratio of households with positive housing debt, i.e. the intensive margin. $s_{i,t}^{N+}$ and $d_{i,t}^{N+}$ are the respective values for non-housing debt. The mean debt-to-income ratio, $d_{i,t}$ can be written as follows: $d_{i,t} = s_{i,t}^{H+} d_{i,t}^{H+} + s_{i,t}^{N+} d_{i,t}^{N+}$. The percentage-point change in debt-to-income ratios between period t and $t-1$ is then

[6] In Appendix Figure 3.A.13, we show that leverage has also increased most strongly for households from the middle of the income distribution.

Figure 3.3.5: Debt along the income distribution

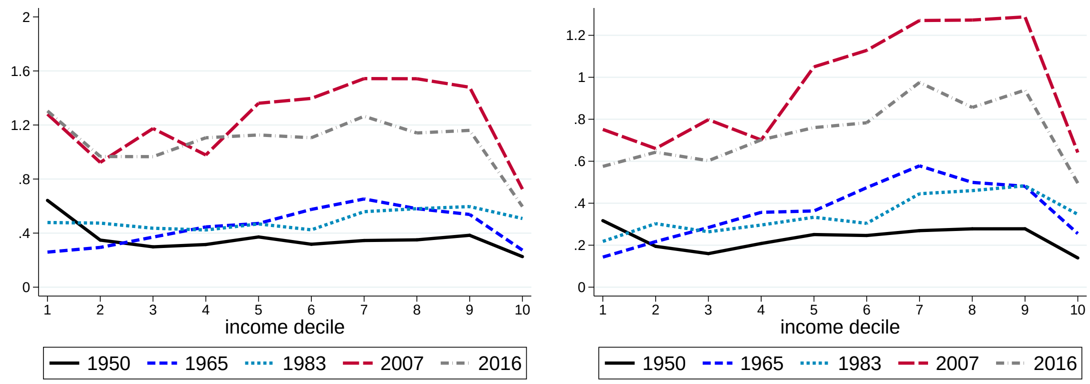

(a) Total debt-to-income

(b) Housing debt-to-income

Notes: The chart shows the evolution of average total (left) and housing (right) debt-to-income ratios by deciles of the aggregate income distribution for the SCF+ waves 1950, 1965, 1983, 2007 and 2016. We exclude households with total income below 10% of the annual wage of a household with a single earner receiving the contemporaneous minimum wage.

calculated as

$$d_{i,t} - d_{i,t-1} = \underbrace{\left(s_{i,t}^{H+} - s_{i,t-1}^{H+}\right) d_{i,t-1}^{H+}}_{\Delta \text{ extensive housing}} + \underbrace{s_{i,t}^{H+} \left(d_{i,t}^{H+} - d_{i,t-1}^{H+}\right)}_{\Delta \text{ intensive housing}} + \underbrace{\left(s_{i,t}^{N+} - s_{i,t-1}^{N+}\right) d_{i,t-1}^{N+}}_{\Delta \text{ extensive non-housing}} + \underbrace{s_{i,t}^{N+} \left(d_{i,t}^{N+} - d_{i,t-1}^{N+}\right)}_{\Delta \text{ intensive non-housing}} . \tag{3.3.1}$$

The first part of this expression is the change in household indebtedness due to a

Table 3.3.1: Decomposition of the increase in debt-to-income ratios between 1950 and 2016

housing debt	intensive margin	32.9
	extensive margin	19.7
non-housing debt	intensive margin	14.5
	extensive margin	7.5
total		74.5

Notes: The table shows the percentage point change in the average debt-to-income ratio between 1950 and 2016, decomposed into extensive and intensive margin effects for housing and non-housing debt according to (3.3.1).

change in the extensive margin of housing debt. In other words, it captures by how much household indebtedness would have risen if only $s_{i,t}^{H}$ had changed, everything else being at the level of period $t-1$. The second part is the effect due to variations in the intensive margin, i.e. changes in household indebtedness due to an increase in d_t^H if the extensive margin of housing debt, $s_{i,t}^{H}$, had been constant at the level of period t, and all non-housing debt components had remained at the level of period $t-1$. The third and fourth parts are the respective effects for non-housing debt.

Table 3.3.1 shows the extensive- and intensive-margin effects of the increase in the average debt-to-income ratio between 1950 and 2016. Overall, we find that the intensive margin of housing debt accounts for 31.5 percentage points of the 75-percentage-point increase in average household debt-to-income. Another 20 percentage points are due to the extensive margin of housing debt. The remaining 23.5 percentage points are due to non-housing debt. This confirms that mortgage lending has played the dominant role for the debt boom.

Figure 3.3.6: Extensive and intensive margin of debt-to-income

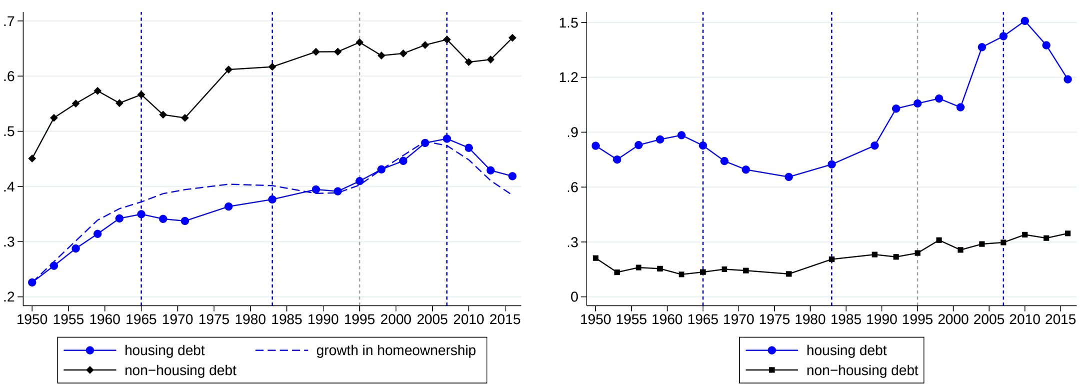

(a) Extensive

(b) Intensive

Notes: The left panel shows the share of households with positive housing debt (blue line with dots) and positive non-housing debt (black line with squares). Moreover, it shows the growth rate of the homeownership rate since 1950, normalized to extensive-margin housing debt in 1950 for comparison. The right panel shows the (non-)housing debt-to-income ratio of households with positive (non-)housing debt. Black vertical lines indicate pivotal dates related to the debt boom. The gray dashed line marks 1995, when house price growth accelerated and homeownership started to increase.

Figure 3.3.6 shows both margins of indebtedness over time for both types of debt. The extensive margin in the left panel captures the share of households with positive (non-) housing debt balances. A closer look at Figure 3.3.6 reveals that the extensive margin of housing debt tracks changes in the homeownership rate closely (dashed line). The intensive margin in the right panel is represented by the debt-to-income ratio for households with positive levels of (non-)housing debt. Overall, more households have personal debt than housing debt. In particular, the roll-out of credit cards in the 1970s lead to a substantial increase in the share of households with personal debt (cf. Figure 3.A.10). Yet the amount that households owe is small compared to the average amount owed on housing debt, as the right-hand side shows.

3.3.3 Four phases of the postwar debt boom

From Figures 3.3.6, we identify four different phases of the postwar debt increase that we explore in detail in Figure 3.3.7a. We decompose the change in debt-to-income ratios into the extensive and intensive margin, and stratify by income. The Figure shows two boom phases (1950-1965 and 1983-2007), followed by two periods of deleveraging. Figure 3.3.7b shows a similar picture for loan-to-value ratios. There are substantial differences between the four periods.

The postwar homeownership boom: The first period is characterized by the rise of homeownership after World War II until the mid-1960s, aided by public policies to increase homeownership (Fetter 2013, 2014). The debt-to-income ratios approximately doubled in this period (see Figure 3.2.2), mainly driven by the extensive margin of housing debt and by the upper half of the income distribution. Likewise, average loan-to-value ratios increased, driven predominantly by the extensive margin and some increase in LTVs in the lower half of the distribution.

Stability, 1965-1983: The second period spans the years from roughly 1965 to 1983. It is characterized by almost stable debt-to-income ratios and a slight decline of intensive-

Figure 3.3.7: Decomposition of changes in debt-to-income and loan-to-value by income group

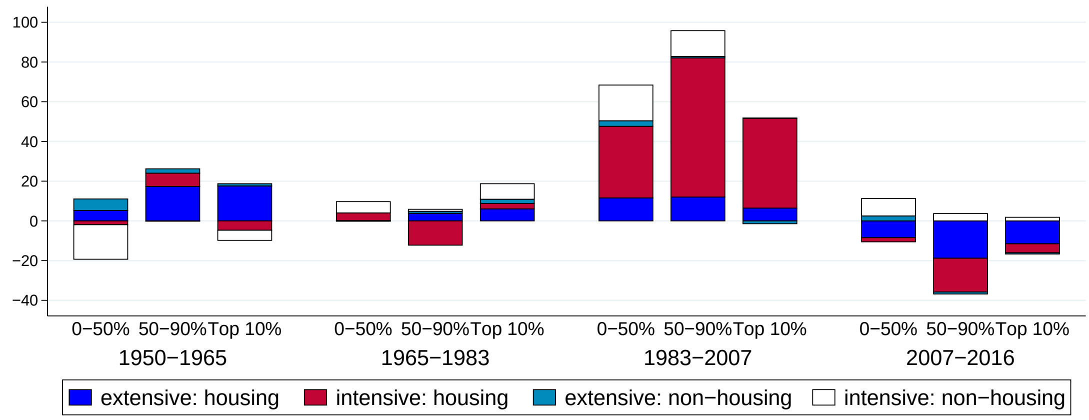

(a) Debt-to-income

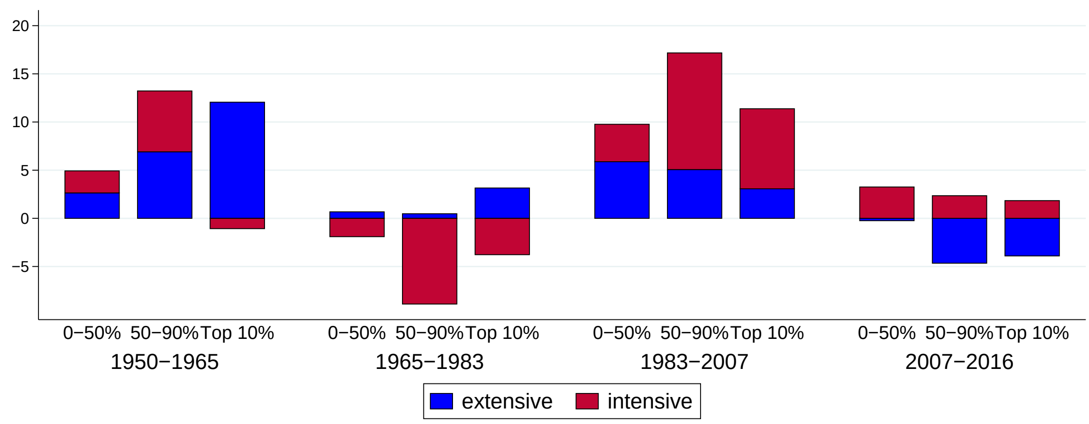

(b) Loan-to-value

Notes: The upper panel shows the decomposition into extensive and intensive margin effects from 3.3.1 over the four phases of the debt boom, stratified by income. The lower panel shows an analogous decomposition of the loan-to-value ratio. Observations with debt-to-income ratios above 50 in absolute value were excluded.

margin housing debt of the middle class, with marginal increases at the extensive margin. Both at the top and in the bottom 50%, non-housing debt (car loans and credit cards) make a small but positive contribution to debt ratios. Loan-to-value ratios decrease across income groups.

The second debt boom, 1983-2007: Starting in the 1980s, the United States entered a second debt boom, which came to an end with the crisis. Debt-to-income ratios more than doubled within the 25 years between 1983 and 2007, from roughly 60% of income to above 130%. This time, the increase was mainly driven by higher intensive margins of housing debt, as Figure 3.3.7a shows. Overall, the extensive margin made a relatively small contribution, but the effect is larger in the 2000s, as we will see below. The boom was fueled by households from all parts of the income distribution. In particular, the intensive-margin effect of the middle class (50-90%) stands out, both for debt-to-income and loan-to-value ratios.

Crisis and deleveraging, 2007-2016: The final period covers the decade after the crisis. It is marked by deleveraging. Overall, the debt-to-income ratio fell by about 30 percentage points. For the bottom 50%, non-housing debt, mainly education loans, showed positive growth. The middle class and the top 10% de-leveraged at both margins, but chiefly at the extensive margin. Homeownership rates have fallen across all income groups. The decline in LTVs was mainly driven by a decline in the extensive margin.

Recently, the consequences of strongly rising student debt have received increasing attention (Looney and Yannelis, 2015; Mueller and Yannelis, 2019, see, for example,). Rising student debt shows up in Figure 10a as part of the intensive margin of non-housing debt. Since 1983, we find a significant contribution from the intensive margin of non-housing debt especially in the lower half of the income distribution. These increasing debt levels might shape financial decision making of young generations of American households in the future but what Figure 10a also shows is that from a macroeconomic perspective, the contribution of student debt is still much smaller than the increase in housing debt over the same period of time.

Figure 3.3.8 zooms in on the post-1980 debt boom. In its first phase, from 1983 to 1995, the debt increase was similar for all income groups, and intensive-margin housing debt played the central role. In the second phase, from 1995 to 2007, the quality of the debt boom changed considerably. Middle-class debt-to-income grew twice as much as that of the other income groups. The significant increase of the debt ratio in the top 10% is also noteworthy, as it effectively outpaced the increase of debt ratios in the bottom half of the income distribution. In the middle and the lower half of the distribution, the extensive margin also made a substantial contribution to rising debt levels after 1995. This reflects the homeownership boom of the 2000s, partly driven by lending to households from the lower half of the distribution. Over the entire boom from 1983 to 2007, middle-class debt-to-income increased by 82 percentage points, predominantly because of higher intensive-margin indebtedness.

Figure 3.3.8: Two stages of the second debt boom

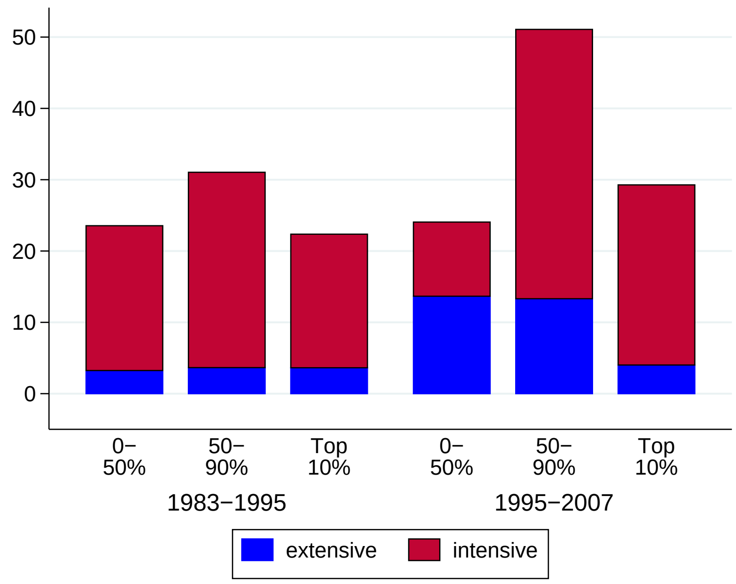

Notes: The graph repeats the analysis from Figure 3.3.7a, zooming in on the second debt boom. Observations with debt-to-income ratios above 50 in absolute value were excluded.

3.3.4 Life-cycle profiles of household debt

So far, we have shown that the middle class and the intensive margin of housing debt were the main drivers of the debt boom in the past decades. In this section, we will ask how the debt increase has affected households at different stages of their life cycle. We will encounter substantial changes in the life-cycle of debt in the U.S. Importantly, we will see that the slope of debt-to-income profiles flattened substantially over time.

Instead of stratifying the data by income group, we trace different generations of American households. The long time span of the SCF+ data gives us the possibility, for the first time, to follow individual birth cohorts and their indebtedness over several decades. Since the SCF+ is not a panel, we construct synthetic birth cohorts. Households with heads born between 1915 and 1924 are our oldest cohort, and households with heads born between 1965 and 1974 are our youngest cohort. Correspondingly, our oldest cohort is on average 30 in 1950, and our youngest cohort is on average 46 in 2016. We estimate life-cycle profiles of total and housing debt-to-income for each synthetic cohort by regressing individual ratios on six age group dummies. We focus on households between 25 and 85 years of age. The groups comprise households with a head of 25-34, 35-44, 45-54, 55-64, 65-74, and 75-85 years, respectively.[7]

The resulting life-cycle profiles are shown in Figure 3.5.1. Debt-to-income increased from one cohort to the next, leading to an upward shift of life-cycle profiles across cohorts. For instance, the generations born before WW II started with an average debt-

[7]We exclude households with extreme debt-to-income or housing-to-income ratios of larger than 50 in absolute value. Very small incomes of less than 10 in absolute value and house values of less than 500 dollars (in real terms) are treated as zero.

Figure 3.3.9: Debt over the life cycle

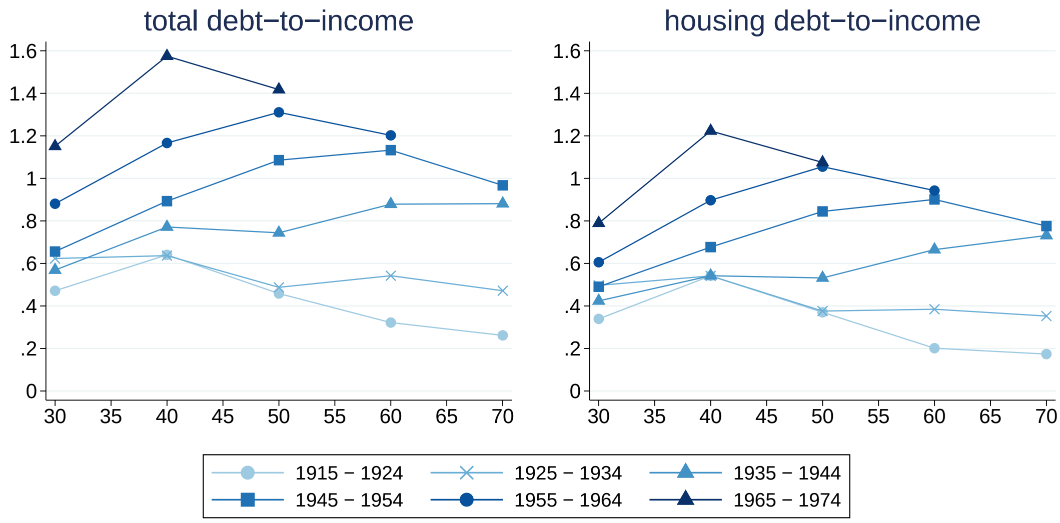

Notes: The panel shows life-cycle profiles of total and housing debt-to-income for our synthetic cohorts.

to-income ratio of around 0.5. Debt ratios of the two baby boomer cohorts, born in the two decades after WW II, were slightly higher at the beginning of the life cycle. At age 30, they started with debt ratios between 0.5 and 0.6, possibly reflecting the effects of the postwar credit policies that encouraged homeownership and sustained markedly higher LTVs (Fetter 2013).

Apart from the level shift, we also observe a turning of the life-cycle profiles. This upward rotation occurs when the average household from the 1915-1924 cohort is 60, the average household from the 1925-1934 cohort is 50, and the average household from the 1935-1944 cohort is 40, i.e. around 1980 at the onset of the second debt boom. These households reach retirement age with substantially elevated debt levels compared to previous cohorts (see also Lusardi et al. 2018).

At age 70, the visual contrast is stark. The prewar generations typically entered retirement with much reduced debt ratios of around 30% to 50% of income. Yet households in the first baby boomer cohort (1945-1954) had debt ratios of almost 120% on average at the same age, i.e., more than twice as high. Younger cohorts reach retirement age with considerably higher debt levels than before. The shift in the slope of the life-cycle profiles is considerably more pronounced than the upward shift of the profiles at the beginning of the life cycle.[8]

An explanation for the increase in American debt will therefore have to be able to account for the fact that households are no longer reducing, or even increasing housing debt over the life cycle. The drivers of this change in debt profiles over the life cycle is what we turn to next.

[8]Appendix 3.A.7 shows that the same patterns are visible in the PSID data, which allow to follow actual instead of synthetic cohorts.

3.4 House prices, wealth growth, and the debt boom

We have established that the intensive margin of middle-class housing debt was the key driver for the increase in household debt. At the same time, income growth of middle-class households was low, at best. Is this evidence supportive of the popular view that those parts of the population that were cut off from income growth increasingly had to rely on debt to finance consumption? How can we rationalize this substantial middle-class debt accumulation in the presence of stagnant incomes? To address these question in this section, we exploit a key strength of the SCF+ data. They provide a comprehensive picture of the entire household balance sheet, including the asset side. We also complement the analysis with data from the PSID, which has a panel structure allowing to study debt accumulation of individual households over time.

We start the discussion by pointing to an important fact, displayed in Figure 3.4.1. The graph shows the long-run trend in debt-to-income ratios for the bottom 90% next to the trajectory of their (net) wealth-to-income ratios. The chart demonstrates that the increase in debt is dwarfed by the rise in net wealth. The figure tells us that the average value of assets grew by a larger absolute amount than the average value of debt.[9]. Put differently, despite the pronounced rise in debt-to-income ratios since the 1980s, middle-class households became considerably richer. Middle-class wealth and income growth diverged substantially.

Figure 3.4.1: Debt-to-income vs. wealth-to-income

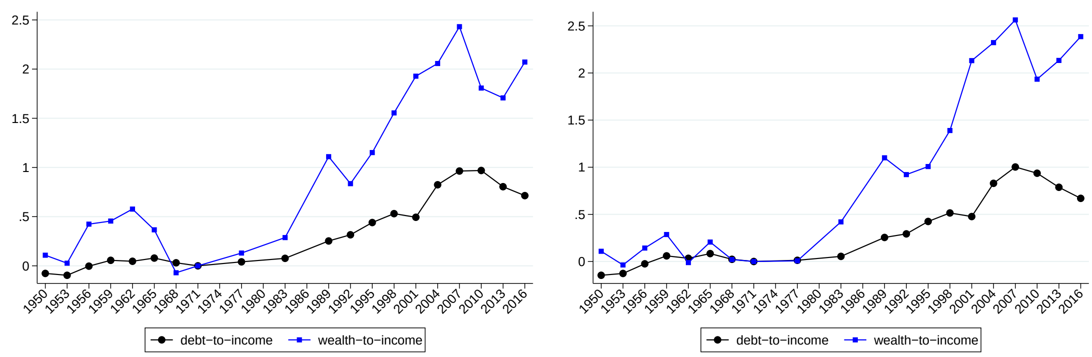

(a) Bottom 90%

(b) 50 - 90%

Notes: The right panel shows average debt-to-income and wealth-to-income for the 50th to 90th percentile of the income distribution, normalized to zero in 1971. The left panel shows the same series for the bottom 90% of the income distribution.

There are two potential sources for an increase in asset holdings. First, higher savings may lead to a more rapid accumulation of assets. Second, there may have been valuation gains on existing assets. For the first channel to be quantitatively important at a time of low income growth for low- and middle-class households, we would have to see a substantial rise in savings rates. However, the data show that savings rates actually

[9]Given relatively low initial debt-to-asset ratios, which only increased moderately over time (see Figure 3.A.12), this outcome is not surprising.

decreased for these households over time (Mian et al. 2019, Saez and Zucman 2016, Zandi 2019) so that we are left with the second channel: capital gains on existing assets. We will argue that such valuation gains, predominantly on residential real estate, played the dominant role for rising middle-class wealth in the face of stagnant incomes. Rising house prices, against the background of the high exposure of the typical middle-class household portfolio to the housing market, led to substantial equity gains that pushed up middle class net worth (Kuhn et al., forthcoming).

Figure 3.4.2a shows that between the early 1980s and 2007 real house prices, adjusted for quality changes, increased by almost 70%. Figure 3.4.2b shows the increase of housing assets relative to income across the income distribution. Housing-to-income rose most strongly for middle- and low-income households, considerably more than at the top. Between the late 1970s and the 2008 crisis, the average housing-to-income ratio of the middle class increased by more than 160 percentage points (Figure 3.4.2b) and thereby more than doubled from a level of 145% to 300%. Price increases can account for about two thirds of this increase, according to our data.

Figure 3.4.2: House prices and housing wealth-to-income ratios

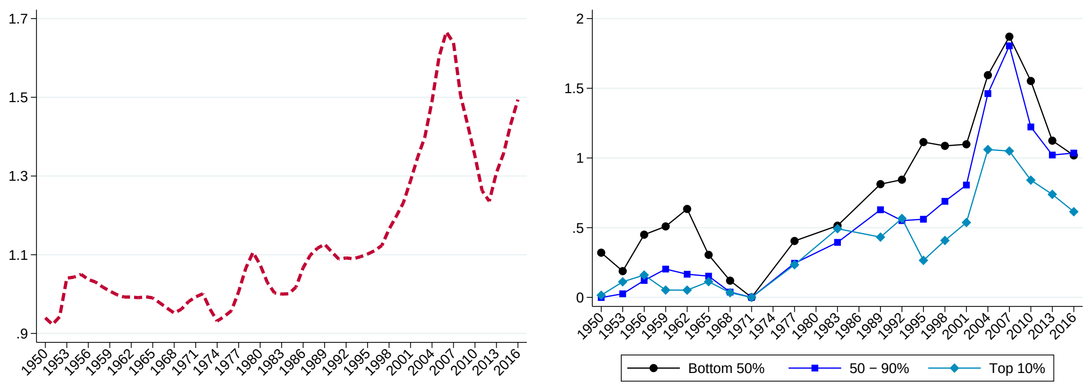

(a) House prices (b) Housing-to-income

Notes: The left panel shows the house price index from the *Macrohistory Database*, deflated by the CPI. The right panel shows average housing wealth relative to average income from the SCF+, normalized to zero in 1971.

We will argue that these housing wealth gains hold the key to understanding the middle-class borrowing surge of the past decades. This is because a substantial share of the debt increase was a *reaction* to such house-price-induced wealth gains. As the value of their real estate increased, middle-class households became wealthier and turned part of this new wealth into additional spending through home-equity-based borrowing. We will show that a significant share of the debt build-up was a Modigliani-style consumption smoothing response of (mainly) middle-class households to large wealth gains resulting from concentrated housing portfolios.

When putting the empirical facts together, we still find low-income grówth middle class households at the center of the debt boom, yet in a way that challenges existing hypotheses. While most of the borrowing was done by households from groups with stagnant incomes, it turns out that until 2007 the same groups also experienced high wealth growth. Rapid debt growth can, to a large extent, be rationalized as a consumption smoothing response to this prie-induced growth of middle-class wealth. Clearly, this "rational" explanation of the debt growth does not preclude that behavioral factors also played a role at some point in the process. For instance, households might have mistakenly assumed housing wealth gains to persist when they did not. But the data suggest that households acted as if these wealth gains were assumed to be persistent.

To make the argument, we will proceed in three steps: First, we will substantiate the idea that the net wealth position of households in the bottom 90% of the income distribution is particularly exposed to house prices, and that rising real estate prices led to substantial capital gains for middle-class households. In a second step, we will show that households reacted to these capital gains by extracting home equity in a way that is quantitatively important for the overall trajectory of household debt. For this step, we complement the SCF+ data with housing and mortgage panel data from the PSID that allow us to decompose debt dynamics and quantify the contributions of equity extraction, new ownership, and upgrading to the debt increase.

In the last step, we will contend that the observed home-equity-based borrowing is consistent with optimizing household behavior in state-of-the-art life-cycle models (Berger et al. 2017). The discussion will also deal with the question whether households are "right" to treat wealth gains from house prices in a similar way to, say, gains in the stock market, and what the financial stability implications are.

3.4.1 House prices and middle-class wealth

To quantify the exposure of middle-class households to the housing market, Figure 3.4.3a presents the elasticities of household wealth with respect to house price changes for our three income groups. The elasticity of around 0.5 that we observe on average for the bottom 50% and the middle class (50%-90%) implies that a 1% increase in house prices increases the wealth of these households by 0.5%. Clearly, also the top 10% own houses, and the average amounts of their housing wealth is high. Yet as a share of total wealth, houses constitute a smaller share for this group, and leverage is lower. Consequently, we find a substantially smaller elasticity for the top 10%, varying around 0.2. House price exposure of the bottom 90% is, hence, on average more than twice as large. Figure 3.4.3a shows little variation in house price exposure between the bottom 50% and the middle class (50%-90%). Yet the average level of housing assets is much smaller for the bottom 50%, which implies that this group matters less for aggregate household debt.[10]

Figure 3.4.3b combines the information from Figures 3.4.2a and 3.4.3a for a first approximation of housing capital gains along the income distribution. We multiply housing

[10]For the bottom 50%, housing is, with 55,800 dollars across survey years, substantially smaller compared to the middle class (50-90) with an average of 135,000 dollars across survey years (see also Adelino et al. 2018).

Figure 3.4.3: House price exposure and capital gains

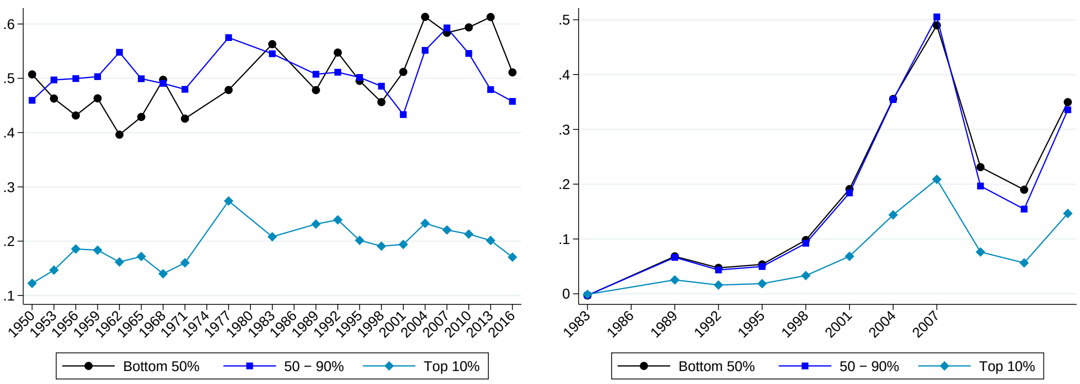

(a) House price exposure

(b) Capital gains over assets

Notes: The left panel shows house price exposure, computed as $\frac{house}{wealth}$. Right panel shows capital gains, computed by multiplying average housing assets of each group in 1983 with the increase of the house price index from Figure 3.4.2a since then.

assets of the income group in period t with the observed rate of constant-quality house price growth from t to $t+1$, and sum these capital gains over time. We normalize the resulting series by the average wealth of each group in 1983. Without saving any income, the average household from the bottom 90% experienced capital gains equivalent to 50% of its 1983 wealth until the peak of the housing boom in the 2000s, in contrast to only 20% for the average top 10% household.

3.4.2 Quantifying home-equity-based borrowing

How did households react to these gains in housing wealth, and what role did the reaction play for the increase in household debt? To quantify the contribution of home equity-based borrowing for the debt increase, we complement the SCF+ data with panel data from the PSID. As discussed in Section 3.2.2, we use the SRC sample, which tracks the original households from the first PSID wave in 1968 over time, as well as the new households formed by former members of these households, e.g., adult children moving out. We will focus the analysis on housing debt as the largest component of debt that has driven the overall increase in debt, as discussed above in Section 3.3. Information on net wealth is available in the PSID since 1984. However, information on housing is available since 1968, and on mortgage balance since 1969 (with the exceptions of 1973-1975 and 1982). The initial sample size was about 2,930 households in 1968, and increased to 5,601 by 2017. The PSID was conducted at an annual frequency until 1997, and every two years thereafter. To ensure consistency over time, we discard all even years from the sample.[11]

To isolate the contribution of home equity withdrawal (HEW), we need to separate

[11] The only information we use from the even years is whether a household has moved over the last year. We use this information to construct a measure of whether the household has moved during the last two years, as in the modern waves.

it from other channels that affect debt levels over time: transitions from renting to ownership and vice versa, upgrading to bigger or better homes, and downgrading. We employ the following definitions: **New owners** are defined as households who (1) bought a house and (2) were not homeowners in the previous survey.

Upgraders are households who (1) were homeowners before, (2) bought a new house and (3) either explicitly stated upgrading as a reason to move, or moved to a home with a larger number of rooms.

Downgraders are the mirror image of upgraders.[12]

Extracters are defined following a similar approach to Bhutta and Keys (2016) and Duca and Kumar (2014). In particular, these are households who (1) did not purchase a new home and (2) increased their nominal mortgage balance from one survey to the next. [13] The debt change is computed in real terms.

The sum of first and second mortgages is our outcome variable. Since 1996, the PSID provides detailed information on mortgage types. These reveal that on average, 92% of first mortgages are conventional mortgages, and 5% are home equity loans. Before 1994, the PSID only reports the remaining balance on first and second mortgages in one variable. However, the largest part of extraction happens via first mortgages, as the overall quantity of second mortgages is small (see Figures 3.A.18 to 3.A.20). Even at the peak of the boom in 2007, only 9% of households had a second mortgage according to the PSID, with an average balance of 4,200 dollars. By contrast, 46% had a first mortgage, with an average balance of about 70,000 dollars.

Figure 3.4.4 shows the extensive and intensive margin of the different groups over time. At each point in time, we report the share of households who extracted equity, upgraded, or bought a new home (extensive margin).[14] We see a pronounced increase in the share of extracters since the mid-1980s, whereas the shares of upgraders and new owners remained relatively constant over time.

The right-hand side of Figure 3.4.4 documents a surge in the amount by which households change their debt conditional on extracting, upgrading, or changing from renting to owning (intensive margin). In the PSID, the average extraction amount is approximately 35,000 dollars between 1999 and 2010. This number is close to the estimate by Bhutta and Keys (2016) of 40,000 dollars for this period. In the SCF, there is a question on equity extraction related to first mortgages since 2004. Despite some differences in mortgage classifications between the SCF and the PSID, the SCF also shows an average

[12]The number of rooms was averaged across all years a household is living in a given house to avoid spurious classifications due to one-time misreporting. Households who increased (decreased) both the size and value of their house by more than 1.5 (0.5) were defined as upgraders (downgraders) even if they did not explicitly indicate to have moved.

[13]We also include a relatively small number of households who increased their nominal mortgage balance but moved to a less expensive, smaller, or same-sized home.

[14]We focus on these groups, because they will be most important for our following analysis. A full version with downgraders and households who sell their homes to become renters can be found in Figure 3.A.17 in the Appendix.

Figure 3.4.4: Intensive and extensive margin by type

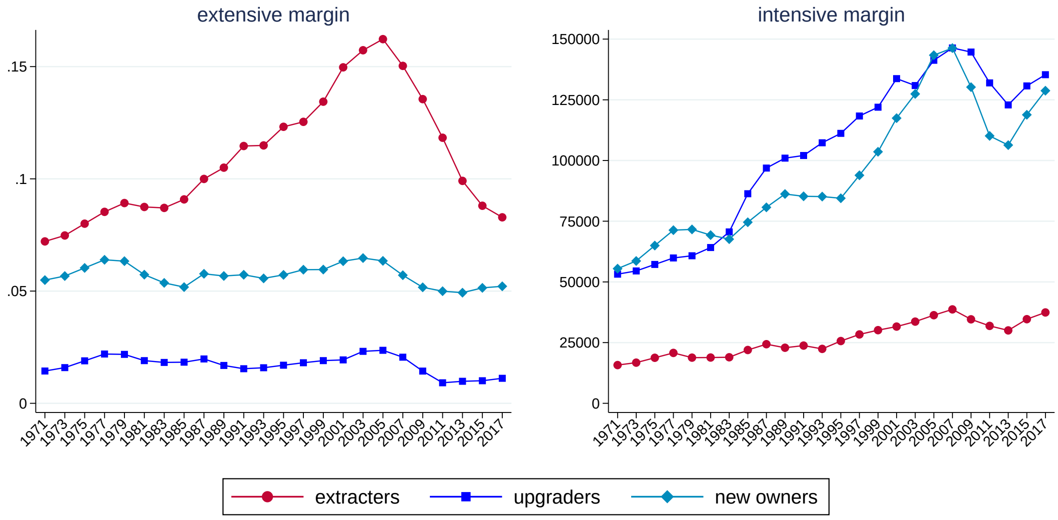

Notes: The left panel shows the share of households who extracted equity, upgraded or bought a new home over time. The right panel shows the average debt increase of these households. The series were smoothed by taking a moving average across three neighboring waves.

extraction amount of 39,000 dollars between 2004 and 2010. Table 3.A.1 compares equity extraction estimates between the PSID and SCF.

To quantify relative importance of extractors, new owners, and upgraders for the growth of household debt, we use an accounting approach. Let D_t denote the stock of housing debt in period t. D_t^+ is new debt taken out by extracters, upgraders or new owners. D_t^- is debt paid back by households who downgrade or switch to renting. A_t is regular amortization of households who do not move or refinance. Then the law of motion for aggregate housing debt is

$$D_t = D_{t-1} + D_{t-1}^+ - D_{t-1}^- - A_{t-1}. \tag{3.4.1}$$

Between the mid-1960s and early 1980s, the aggregate debt stock was relatively constant (see Figure 3.2.1c). In other words, we had a situation in which $D_{t+1} - D_t \approx 0$, and therefore $D_t^+ \approx D_t^- + A_t$. For D_{t+1} to increase beyond D_t, we need to observe increases in D_t^+, or decreases in D_t^- or A_t.

For a specific example consider equity extraction. There are two reasons for additional debt due to equity extraction: First, there may be *more households* extracting equity (extensive margin). Second, conditional on extracting equity, households may extract *larger amounts* (intensive margin). Let b denote the base year, and let $\Delta_t D$ denote the average debt change of households who extracted equity in period t, i.e., the intensive margin. Further let s_t denote the sample share of extracters in period t, i.e. the extensive margin. The additional debt due to increases in the share of extracters since the base year is $\Delta D_t^{ext} = \Delta D_t \times (s_t - s_b)$. The additional debt due to changes in the average amount by which households increase their debt at the time of extracting is $\Delta D_t^{int} = s_b \times (\Delta D_t - \Delta D_b)$.

Adding these two numbers yields our estimate for the amount by which average housing debt would have been lower each period if the share and amount of extracters had stayed at their base-year levels. We can then cumulate this series to compute the amount by which the stock of housing debt would have been lowered over time in the absence of additional equity extraction. Analogous calculations are done for upgraders, downgraders, and new homeowners.

Figure 3.4.5 reports the results and plots the contribution of the different household types to the increase in housing debt relative to the base year. We consider data between 1981 and 2007 to cover the whole debt boom period since the 1980s. The dashed line in the figure shows the observed increase in housing debt since 1981.

Figure 3.4.5: Decomposition of the housing debt boom

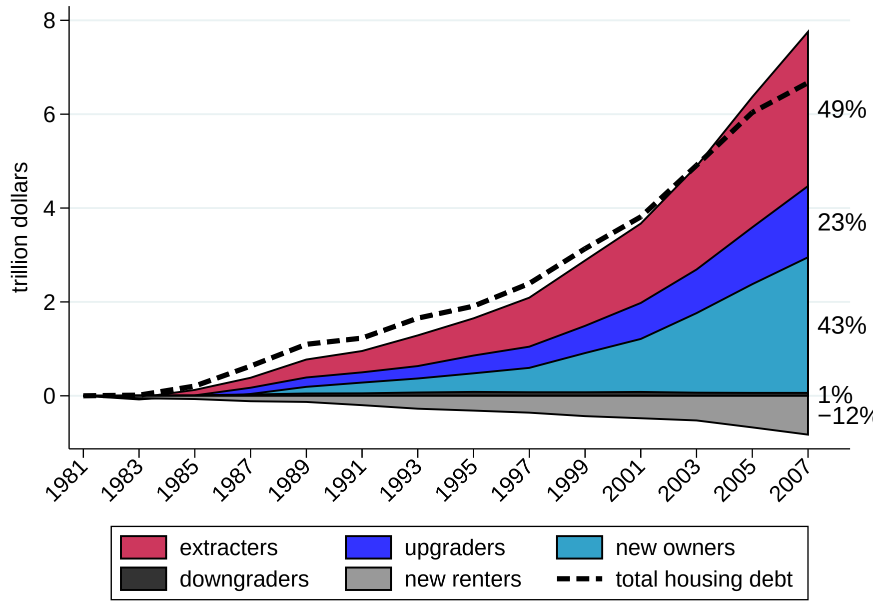

Notes: The graphs shows the change in total housing debt since 1981 as a black dashed line, together with estimates of the change in the stock of housing debt due to HEW, upgrading, downgrading, new home ownership and giving up home ownership. Please refer to the text for details on the construction of these estimates. The percentages on the right side are the shares of each shaded area relative to the actual increase (indicated by the dashed line) in 2007.

The first important observation is that the accounting decomposition matches the total housing debt increase between 1981 and 2007 closely. The combined growth of debt for the individual groups account for almost all of the overall debt increase. Equation (3.4.1) implies then that there were no major changes in amortization behavior.

The second key result is that home equity extraction has played a quantitatively large role in driving the debt boom. It accounts for about 49% of the total increase in housing debt. In other words, about half of the increase in housing debt is driven by incumbent owners borrowing against their home equity. New owners account for a slightly smaller share, with around 43%. Upgraders account for about 23%, while new renters contribute negatively to the total increase. The net contribution of downgraders was negligible over the considered period.

Together, upgrading and home equity extraction account for more than 70% of additional housing debt since 1981. This corroborates our previous finding that the intensive margin of housing debt is the key driver of the debt boom. Note that both extracters and upgraders tap into home equity for additional spending. Upgraders increase housing consumption by buying a larger house, while extractors may use the funds for home improvements or other consumption purposes.[15]

The relative contribution of new ownership rose in the mid-1990s, reflecting the increase in home ownership rates prior to the 2008 crisis. While rising house prices bring capital gains to existing homeowners, they imply less purchasing power for prospective homeowners who have saved for the down payment. With falling purchasing power, prospective homeowners have to accumulate more savings out of income, or rely on additional debt to finance their home purchase. As most households who change from renting to owning are young, this drove young households deeper into debt than in previous generations. Appendix Figure 3.A.21 shows that loan-to-value ratios of young homeowners increased from around 40% in 1950 to almost 80% by 2007. Yet the overall picture is dominated by incumbent homeowners and variations in the intensive margin of debt.

3.4.3 Regulatory and tax changes

Home equity-based borrowing started to surge in the mid-1980s. The timing is not coincidental, as changes in taxation in regulation prepared the ground. The most important change came with the Tax Reform Act of 1986 that limited the deductibility of interest on debt to interest on debt secured by first and second homes. This meant that homeowners could retain the tax deductibility of interest payments by shifting other debt to housing debt, e.g., home equity lines (HELs) (Kowalewski 1987). In addition, interest rates charged on such HELs were considerably lower than credit card debt (Canner et al. 1988). Maki (1996, 2001) shows how households took advantage of this and changed their debt portfolios from consumer towards housing debt after the abolition of the consumer interest rate deductibility.

Financial institutions started to market new home equity borrowing products strongly in the 1980s. In the mid-1980s, nearly half of the country's largest financial institutions spent more advertising dollars on these products than on anything else (Canner et al. 1988). For instance, Citibank advertised its new "Equity Source Account" by linking house prices to individual achievement: "Now, when the value of your home goes up, you

[15]In the SCF, households are asked for the purpose for which they extracted home equity since 1995. Among the households who extracted equity, around one third used the money for home improvements and repairs. Another thirty to forty percent spend the money on consumption and the repayment of other debts. Other important purposes are the purchase of vehicles, vacation properties, and investments in other assets, with average response rates of around 5-10% each.

can take credit for it" (Story 2008). Banks were successful in overcoming the negative connotation of second mortgage products, which were traditionally seen as a last resort for households in financial trouble. HELs were now branded as a cheap and convenient way to tap into home equity (Kowalewski 1987).

Within a few years in the 1980s, the HEL market grew from close to zero to 100 billion dollars in volume (Story 2008). Regulatory change played a role for kicking off the 1980s equity extraction boom, too. Until its amendment in 1982, the Truth in Lending Act gave consumers the right to rescind credit transactions secured by home equity within three days. This made second mortgage credit burdensome and expensive for the banks. Other Depression-era regulations on the mortgage market were also abolished, allowing mainstream banks to sell secondary mortgage products (Story 2008).

A second withdrawal boom got underway in the 1990s. Conforming real mortgage interest rates fell from around 6% in the mid-1990s to 3% in the 2000s (see Figure 3.A.14 in the Appendix), and house price growth accelerated. This provided strong incentives for households to refinance, and many of extracted home equity on the way via a cash-out refinancing. Bhutta and Keys (2016) show that cash-outs accounted for the largest share of equity extraction between the early 2000s and the crisis in 2008, followed by HELOCs and second mortgages.

In Appendix Figure 3.A.23, we show how mentions of the term "home equity loan" in American books have evolved over time. The data come from the Google Books Ngram Viewer, an online search engine which displays the frequency of search strings (*n-grams*) in sources printed until 2008 (see also Michel et al. 2011). The graph clearly mirrors the historical evidence: Until 1982, the term "home equity loan" was hardly mentioned at all. In 1983, the share of mentions starts to go up, and then rises steeply in 1986. After reaching a plateau in the late 1980s, the share surges rapidly again in 1995, consistent with the timing of the second withdrawal boom.

3.4.4 Middle-class equity extraction

How was the equity extraction boom distributed across the different income groups? Is there evidence that middle-class households, whose portfolios are most exposed to rising house prices, played an active role in the process? Based on the PSID data, we can show that households between the 50th and 90th percentile accounted for the dominant share of equity extraction. Middle-class households also exhibit particularly high extraction elasticities with respect to house price changes in their states.

Figure 3.4.6 shows total home equity extraction as a share of total annual household income for the bottom 50%, the 50%-90%, and the top 10% of the income distribution.[16] We smoothed the data by taking a moving average across three neighboring waves. Before 1986, the ratio of extraction to income was similar for all three groups, at around

[16] Note that our measure refers to total extraction over the previous two years. The results of Bhutta and Keys (2016) suggest that between 10 and 20% of households extract in two consecutive years.

2-3% for the bottom 90%, and 2% for the top 10%. In the mid-1980s, we see an increase in extraction relative to income, which is particularly pronounced for the top 10%. This again points to the Tax Reform Act of 1986, which arguably had larger effect on households with higher incomes. Figure 3.A.24 in the Appendix illustrates how rich households began to expand their housing debt and invest in owner-occupied real estate around this time.[17]

Over the 1990s, extraction rose from around 3% to more than 6% of annual income for households from the bottom 90%. After the crisis in 2008, it dropped to a level of around 3.5%, where it has remained since 2013. By contrast, extraction was falling over the 1990s for the top 10%, and only rose again in the early 2000s. Even at the peak of the debt boom, it did not exceed 5% of income.

For households with low income growth, additional extraction will translate almost one-to-one into higher debt-to-income ratios. To see this, let us reconsider equation (3.4.1), and divide by income Y_{t-1} on both sides:

$$\frac{Y_t}{Y_{t-1}}\frac{D_t}{Y_t} = \frac{D_{t-1}}{Y_{t-1}} + \frac{D^+_{t-1}}{Y_{t-1}} - \frac{D^-_{t-1}}{Y_{t-1}} - \frac{A_{t-1}}{Y_{t-1}}$$

To ease notation, we will express ratios relative to income in small letters, and denote the income growth rate by g:

$$d_t = (1+g)^{-1}\left[d_{t-1} + d^+_{t-1} - d^-_{t-1} - a_{t-1}\right].$$

For households with low income growth, we have $g \approx 0$. Iterating backwards, we obtain

$$d_t - d_0 = \sum_{i=0}^{t-1}\left[d^+_i - d^-_i - a_i\right] \qquad (3.4.2)$$

Until 1985, middle-class households on average extracted 2.6 percent of their annual income over a two-year period. For the period between 1986 and 2007, this figure increased by 2.3 percentage points. Over 20 years, this translates into a 23 percentage points higher housing debt-to-income ratio due to increased extraction alone.

3.4.4.1 State-level evidence

To estimate the association between house price growth and equity extraction relative to income over time, we estimate local projections (Jordà 2005) on state-level data. Previous research has stressed that house price exposure can vary considerably across geographies due to heterogeneity in house price developments (Bhutta and Keys 2016,

[17] The top 1% already increased their housing investment after 1983. In that year, a revision of the alternative minimum tax became effective, which limited deductions, with owner-occupied real estate being a rare exception (Bettner 1982).

Figure 3.4.6: Extraction relative to income, by income group

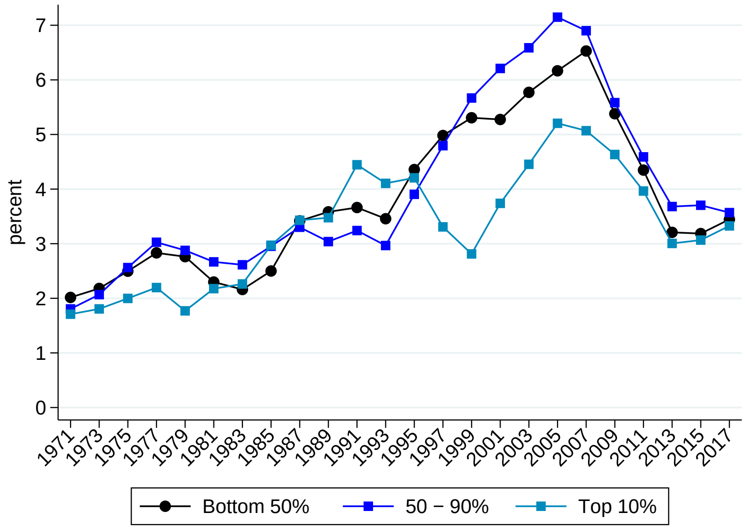

Notes: The graph shows total extraction relative to total income by income group. The series were linearly interpolated for 1973-1977, as mortgage information is not available for the years 1973 and 1975. Data series have been smoothed by taking a moving average over three neighboring waves.

Aladangady 2017, Fuster et al. 2018).[18] We use the state-level version of the FHFA house price index and estimate the following equation for different horizons h:

$$Y_{is,t+h} = \beta_0 + \beta_1 \, g_{st}^P + \Gamma' X_{ist} + \Psi' \delta_t + \Phi' \gamma_i + \epsilon_{it} \tag{3.4.3}$$

where $Y_{is,t}$ denotes extraction relative to income for a household i living in state s in year t. We focus exclusively on households who do not move. $Y_{is,t+h}$ denotes the cumulative extraction relative to income between period t and period $t+h$; g_{st}^P is the growth rate of the state-level FHFA house price index between two survey waves, and X_{ist} is a set of household-level demographic controls that are plausibly related to equity extraction.[19] The regressions also include time and household fixed effects δ_t and γ_i to capture aggregate conditions and time-invariant household characteristics. As mentioned before, the PSID changed its frequency from annual to biennial in 1997. In order to get consistent results over time, we discard the even survey waves before 1997, and re-compute equity extraction based on the remaining information. Hence, one period corresponds to two years.

Figure 3.4.7 plots the estimated coefficients $\hat{\beta}_1$ for $h = 1, ..., 5$ from equation (3.4.3). The results imply that after a 10-percent increase in house prices, which corresponds to one standard deviation of house price growth, the average homeowner extracts equity equal to about 0.9% of annual income over the following 6 years. Importantly, for the middle

[18] Figure 3.A.4 in the appendix combines regional information from the SCF+, where we observe the state of residence until 1971, with the PSID data. It illustrates the co-movement of housing and housing debt across geographies.

[19] We include age group dummies to capture the life cycle, as well as dummies for the total number of children, the birth of an additional child and business ownership.

Figure 3.4.7: Effect of house prices over time

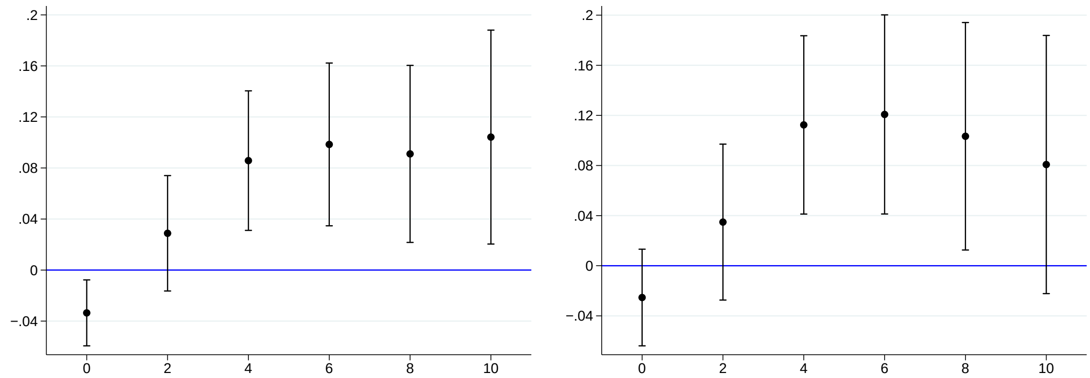

(a) All income groups

(b) 50-90% only

Notes: The left panel shows the coefficient on house price growth at various horizons from equation 3.4.3. The right panel shows the corresponding coefficients after restricting the sample to middle-class households. Observations with extraction amounts larger than twice the annual income or with negative income were excluded. Two-year periods were considered throughout. Controls include dummies for age, children and business ownership, as well as time fixed effects. Standard errors are clustered at the household level.

class, the effect is about one third larger, with a cumulative response of around 1.2% of annual income over six years. The association between house price growth and equity extraction is most pronounced in the middle of the income distribution.

We also estimated event-study regressions around the extraction date, using the reported value of a household's home as the outcome variable. The results show that the house values of extracters increased substantially more than those of non-extracters in the six years prior to extraction, consistent with the evidence from the local projections (see Figure 3.A.22 in the appendix).

3.4.4.2 Aggregate importance of middle-class equity extraction

In a final step, we aggregate the group-specific house price responses back to the level of the macroeconomy to calculate the importance of equity extraction by the different income groups. To do this, we compute the average amount of additional debt due to extraction as in Figure 3.4.5, and multiply it with the total number of households to obtain the aggregate effect. We then again add up the resulting series to find the amount by which the aggregate stock of housing debt would have been lowered each period. Finally, we subtract this estimate from total aggregate housing debt to provide an estimate of how much debt would have increased absent the equity extraction.[20] The black line in the figure shows the actual housing debt-to-income ratio from the PSID data.[21] The blue line shows the counterfactual housing debt-to-income ratio after subtracting our estimate of additional debt due to extraction.

[20] This simple estimate rules out behavioral and general equilibrium responses.

[21] Note that housing debt-to-income has increased somewhat less in the PSID than in the SCF+, reaching 0.84 in 2007, compared to 0.92 in the SCF+ (see Figure 3.A.2a).

Figure 3.4.8: Quantitative importance of middle-class extraction

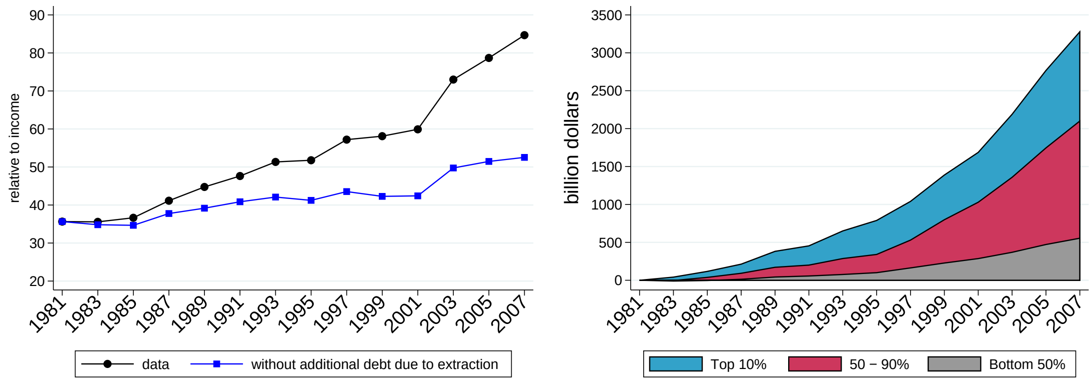

(a) Counterfactual housing debt-to-income

(b) Role of the middle class

Notes: The left panel shows housing debt-to-income from the PSID. The blue line with squares shows actual housing debt minus additional debt due to extraction relative to income. The right panel shows additional debt due to extraction by income group.

Without equity extraction, housing debt would have increased by half as much over the 1981 to 2007 period. Debt-to-income ratios would have stayed at around 40% until 2001. The uptick during the boom of the 2000s, when new homeowners increased aggregate housing debt (see also Figure 3.4.5), would have been much more modest. By 2007, we estimate that the housing debt-to-income ratio would hardly have exceeded 50% of income.

We can also approximate the effect on total household debt based on SCF+ data, which include non-housing debt. If we assume that housing debt had increased by 50% less from 1983 to 2007, and that non-housing debt had not been affected by the slower increase in housing debt, total household debt would have peaked a third lower in 2007 at around 74% of income (see Figure 3.2.2).[22] Figure 3.4.8b highlights the role of the middle class in this development. Equity extraction of the middle class accounts for the lion's share of total equity extraction, and the largest part of the increase in household debt.

3.5 Debt over the life cycle

How have the life-cycle profiles of debt changed? The previous section demonstrated that the basic Modigliani model correctly predicts that the consumption response to rising house prices will be the strongest for households that have the highest wealth gains. Owing to portfolio structure and leverage, in the case of the U.S. that's the middle class.

However, the basis Modigliani model makes an additional prediction. The consumption response will be stronger, the later in the life cycle the unexpected but permanent wealth change occurs. In other words, we are looking for changes in the life-cycle profile of debt

[22]In the PSID, information on non-housing debt is only available since 1984, and the quality and detail of the data is lower than in the SCF+. However, comparing the debt increase in the PSID since 1984 and the SCF since 1983 yields similar results.

Figure 3.5.1: Debt over the life cycle

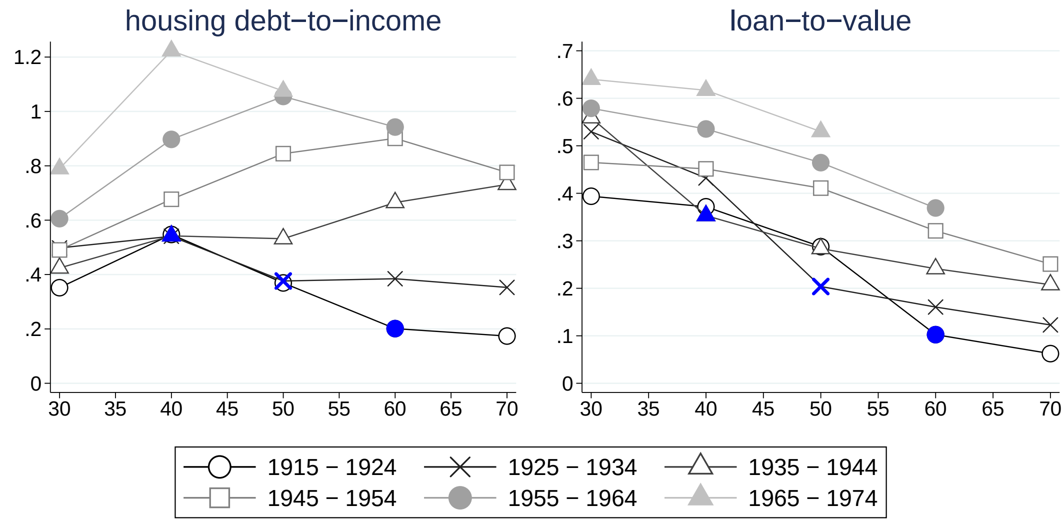

Notes: The left panel shows life-cycle housing debt-to-income profiles for our synthetic cohorts. The right panel shows life-cycle loan-to-value profiles. 1977 was excluded from the regressions, because housing debt is imputed in this year. Blue dots mark the onset of the second debt boom. Please refer to the text for details.

that should be particularly pronounced for older households.

Think about the following case. A household took out a mortgage in 1950, when the head was thirty years old. This household then paid back its debt over the next 30 years, and entered retirement without any debt. Now imagine a household who took out a mortgage in 1970, also at age thirty. Some years later, the head decides to move to a bigger house with a larger mortgage. Still ten years later, she decides to extract equity. When she retires, she still has not paid back a considerable part of her housing debt. Consistent with this reasoning, Lusardi et al. (2018) show that the fraction of people nearing retirement with non-zero debt has increased substantially over time.

The historical SCF+ data give us the possibility to examine this question by comparing many cohorts over a long period. Since the SCF+ is not a panel, we construct synthetic birth cohorts. We use households with heads born between 1915 and 1924 as our oldest cohort and households with heads born between 1965 and 1974 as our youngest cohort. Correspondingly, our oldest cohort is on average 30 in 1950 and our youngest cohort is on average 46 in 2016. The long time span covered by the SCF+ allows us to follow many cohorts over several decades. We estimate life-cycle profiles of housing debt-to-income and loan-to-value ratios for each synthetic cohort by regressing individual ratios on six age group dummies. We focus on households between 25 and 85 years of age. The groups comprise households with a head of 25-34, 35-44, 45-54, 55-64, 65-74, and 75-85 years, respectively.[23]

[23] We exclude households with extreme debt-to-income or housing-to-income ratios of larger than 50 in absolute value. Very small incomes of less than 10 in absolute value and house values of less than 500 dollars (in real terms) are treated as zero.

Figure 3.5.2: Homeowners' housing debt by age

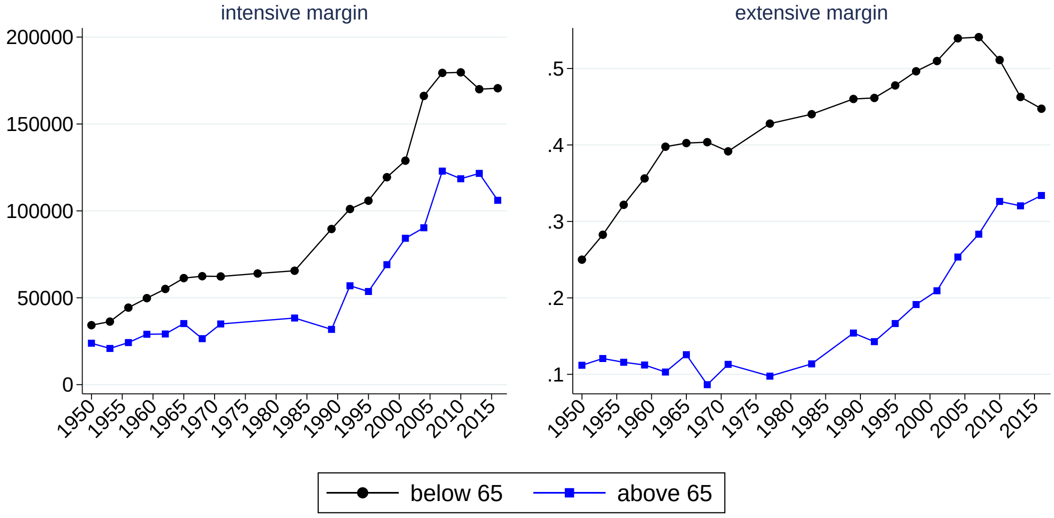

Notes: The graph shows the intensive margin (left panel) and extensive margin (right panel) of housing debt in the sub-sample of homeowners for households below and above 65 years of age over time. The intensive margin of housing debt is imputed in 1977. As the extensive margin of old-age housing debt is small in this year, the imputation does not capture the intensive margin within the subgroups of old households. Therefore, the corresponding data point was omitted.

The estimated life-cycle profiles are shown in Figure 3.5.1. Housing debt-to-income has increased massively from one cohort to the next, leading to an upward shift in loan-to-value ratios across cohorts. Apart from an upward shift of the life-cycle schedules, we can also observe a change in slopes for the debt-to-income and loan-to-value ratios. This change occurs when the average household from the 1915-1924 cohort is 60, the average household from the 1925-1934 cohort is 50, and the average household from the 1935-1944 cohort is 40, i.e. around 1980 at the onset of the second debt boom (see blue markers in Figures 3.5.1). Appendix 3.A.7 shows that the same patterns are visible in the PSID data, which allow to follow actual instead of synthetic cohorts.

The life-cycle debt-to-income and loan-to-value schedules illustrate that, consistent with our conjecture, households from more recent cohorts enter retirement with more debt than households from previous cohorts. Note that the share of households who upgrade (extract) *when they are already above 60* at the point of upgrading (extracting) has only increased *after* the peak of the second debt boom in 2007 (see Figure 3.A.28).

Yet if households buy houses and extract equity later in life, this still implies that there are more households with a non-zero mortgage balance at the age of retirement, and that they have larger debt balances on average, as many of them have taken out a mortgage only recently. Figure 3.5.2 shows the intensive margin (left panel) and extensive margin (right panel) of housing debt among homeowners below and above age 65. Both the extensive and intensive margin of old-age housing debt were constant at relatively low levels before 1980, and have increased markedly since then. While the intensive margin of homeowners' housing debt has evolved in tandem for households below and above 65, the extensive margin has increased more for old households.

Figure 3.5.3a compares the unweighted average age of households over time to the debt-weighted average. To compare the trends in both series, we index the weighted series to 1971 by removing level differences. Since the late 1980s, the debt-weighted average age has risen more quickly than the average age. While the average household age increased by four years between 1971 and 2016, the debt-weighted average age has risen by eight years, i.e. twice as much. This highlights the shift of debt towards older households.

Figure 3.5.3b shows that housing debt has grown particularly fast for households in retirement age during the house price boom. Before, debt had largely grown in parallel for households with heads above and below 65. However, the fact that the lines for all households and households below 65 lie almost on top of each other already indicates that old-age debt is of limited importance for the aggregate debt boom.

This is confirmed by Figure 3.5.4, which shows average housing debt, together with a counterfactual in which the intensive and extensive margin of housing debt in the group of households above age 65 were fixed to their 1950 levels. The two lines lie virtually on top of each other until the 1980s, and stay very close to each other until 2004. At the peak of the boom in 2007, aggregate housing debt would have been roughly 790 billion dollars lower if the intensive and extensive margin of old households had stayed at their 1950 levels. This corresponds to 8.1% of total housing debt, accounted for by 21% of the population.

By contrast, if the intensive and extensive margin of middle-class debt had remained at their 1950 levels, aggregate housing debt in 2007 would have been 4.9 trillion dollars lower, which corresponds to 50.1%. In other words, the aggregate contribution of old-age debt is non-negligible, but rather small in relative terms. There are nevertheless good reasons why we should care about increasing old-age debt. In particular, their

Figure 3.5.3: Aging of debt

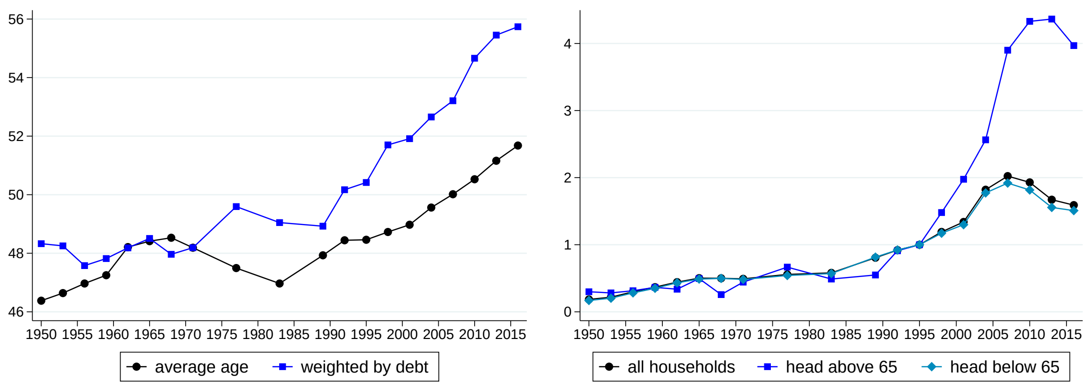

(a) debt-weighted age

(b) housing debt growth

Notes: The left panel shows average age in the SCF+ as a black line, and debt-weighted average age as a gray line. The weighted series is indexed to 1971 by removing level differences. The right panel shows average housing debt of households below and above 65, as well as of all households, normalized to 1995.

Figure 3.5.4: Aggregate role of old-age debt

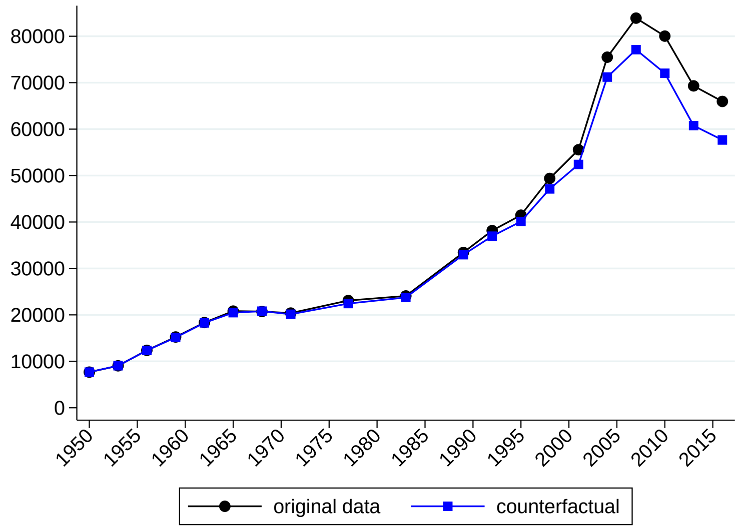

Notes: The graph shows average housing debt, together with a counterfactual in which the intensive and extensive margin of housing debt in the group of households above age 65 were fixed to their 1950 levels.

heightened debt levels make retirees financially fragile (see Lusardi et al. 2018).

Actually, the importance of old-age debt has increased much more *after* the peak of the debt boom than before. The share of old-age housing debt in total housing debt increased from around 5.9% in 2004 to 9% in 2007, and climbed up to 13.2% in 2016. This is consistent with a shift in the market towards buyers who are considered as less risky subsequent to the crisis (cf. Albanesi et al. 2017). By consequence, old households have remained particularly financially fragile even after the end of the crisis. In the final section of our paper, we will take a closer look at the implications of the house price and debt boom for financial fragility.

3.6 Modigliani meets Minsky

In the last section, we return to the macroeconomy to explore the consequences of the surge in debt-financed home equity extraction for financial stability. The gist of the argument will be that home equity-based borrowing, while potentially optimal at an individual level from a *Modigliani perspective*, has made the economy and especially the balance sheets of middle-class families more fragile. We will show that the sensitivity of households to income shocks has risen substantially as debt ratios have risen. The surge in home equity borrowing since the 1980s played an important role in this process. We call this the *Minsky aspect* of the equity extraction boom.

We will see that rising financial fragility of middle-class households was long in the making. The data document a steady increased in share of households that had to be considered "at risk" following an income shock. With the benefit of hindsight, it appears plain evident that financial fragility of U.S. households was rising materially in the background and also that a shock to middle-class income could lead to loan losses that would quickly overwhelm the banking sector's capital base.

Financial fragility is a complex and multidimensional issue. For the analysis here, we focus on household liquidity as one important dimension of financial risk. Liquidity has been emphasized in recent research as an important driver of household consumption decisions (Kaplan and Violante 2014) and mortgage defaults (Ganong and Noel 2018). We quantify growing vulnerability using a *stress testing* approach, not dissimilar to the stress tests used for financial institutions. Our analysis of the macroeconomic consequences of home-equity-based borrowing builds on the work of Mian and Sufi (2011) that explores the link between equity extraction and default rates in the crisis. In a similar spirit, Fuster et al. (2018) conduct a stress test for households based on Equifax CRISM data, shocking home equity positions. The latter paper focuses on a relatively short time period from 2005 to 2017. With our long-run SCF+ data, we will track the trends in financial fragility of the U.S. household sector over the entire post-WWII period.

Figure 3.6.1: Debt-service-to-income and loan-to-value

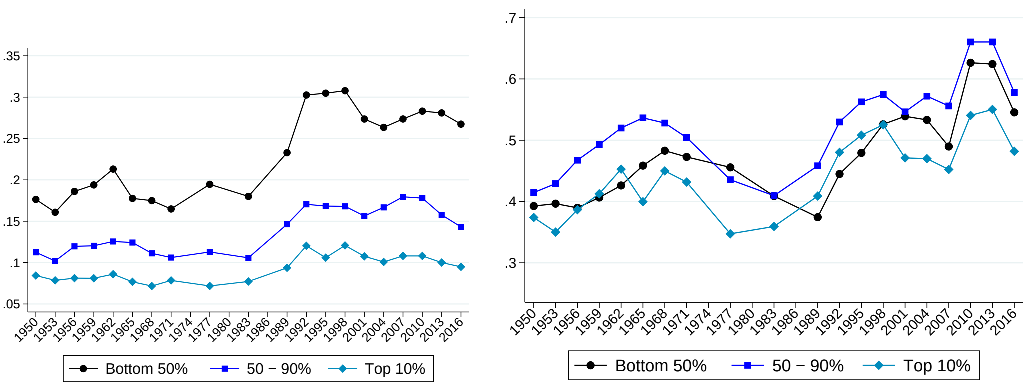

(a) by income group (b) Loan-to-value

Notes: The left panel shows average debt service relative to income among households with positive housing debt, stratified by income. The right panel shows debt-service-to-income for recent extracters, upgraders, new owners and all other households. Recent extracters are households who extracted during the past two years, and the other groups are defined analogously. The graph is based on PSID data.

We start the discussion by looking at the evolution of debt service ratios. The previous sections had shown a substantial rise of debt-to-income ratios among the bottom 90% of the income distribution. Figure 3.6.1 shows that the rise in debt-to-income ratios also led to rising debt service ratios. Falling interest rates cushioned the effect of rising debt on debt service ratios, but the share of income devoted to debt service went up by 50% since the early 1980s, with particularly pronounced increases for the bottom 50%, and the middle class. Overall, the bottom-50% devote 1.5 times more of their income to debt service than the middle class. A similar gap exists between the middle class and the top 10%. Loan-to-value ratios evolved similarly for the three groups, as shown in Figure 3.6.1b.

For our main stress test scenario, we construct shocks that constrain the debt servicing ability of households. Drops in income reduce liquidity and put households under *financial stress*. This in turn potentially requires them to cut their consumption or forces them into loan delinquency or even default. We employ estimates on earnings losses following job displacement by Davis and von Wachter (2011). The paper documents that earnings losses are particularly pronounced in recessions, amounting to losses of 39% in the first year after displacement (Figure 4 in Davis and von Wachter (2011)). We use this number and let the income of the main wage earner in the household drop by 39%.[24] Following the literature, we consider households to be under financial stress if the debt-service-to-income ratio exceeds 40% after the income shock.[25] We then report the *loan value at risk*, computed as the value of outstanding mortgages balances of households under financial stress.

In a second scenario, we consider a joint income and house price shock. The house price shock reduces home equity and pushes some households into negative equity territory. A strand of the literature has emphasized the importance of such "double trigger" events for fragility and default (see e.g. Adelino et al. 2018, Fuster et al. 2018). We present the results in Appendix 3.A.6. The "double trigger" scenario combines our baseline income shock scenario with an 8% drop in real house prices (calibrated to match the average decline from 2007 to 2008), and assumes households to be at risk if their home equity turns negative on top of having a debt-service-to-income ratio above 0.4 after the shock.

Figure 3.6.2a shows our estimate for the loan value at risk as a share of aggregate household income from 1950 to today. The first observation is that there is a pronounced increase of the value at risk over time, peaking in 2007. Second, the increase in the value at risk is quantitatively large, from below 5% to over 30% by 2007. It should however be noted that our shock is imposed on top of all actual shocks that hit the households in our data. Moreover, our stress test scenario hits *all* households by a large income shock while empirical estimates by Davis and von Wachter (2011) suggest that only between 1% and 3% of households are displaced in a recession. If displacement in a recession was not correlated with the amount of the outstanding mortgage, this would imply that loans of up to 9% of aggregate household income had to be serviced by households under financial stress.

Figure 3.6.2b zooms in on the income distribution. It shows that the increase in fragility was mainly attributable to households from the bottom 90% of the income distribution. While the "baseline level" of fragility was always between one and two percent for the bottom 50%, it was close to zero for the middle class. Yet in the 1980s, fragility of both groupsincreased substantially, and in particular for middle-income households. While

[24]We exclude households with negative income. Before 1956, we do not have separate information on the labor income of head and spouse. We therefore impute the earnings share of the principal earner based on data from 1956 to 1959. The average share of the main earner in total household labor income was between 88% and 93% in these years.

[25]The value of 0.4 lies between the two common threshold of 0.36 and 0.45 in the "eligibility matrix" used in the Financial Stability Reports of the Bank of England. We are grateful to Anil Kashyap for suggesting this matrix.

Figure 3.6.2: Value at risk as share of income

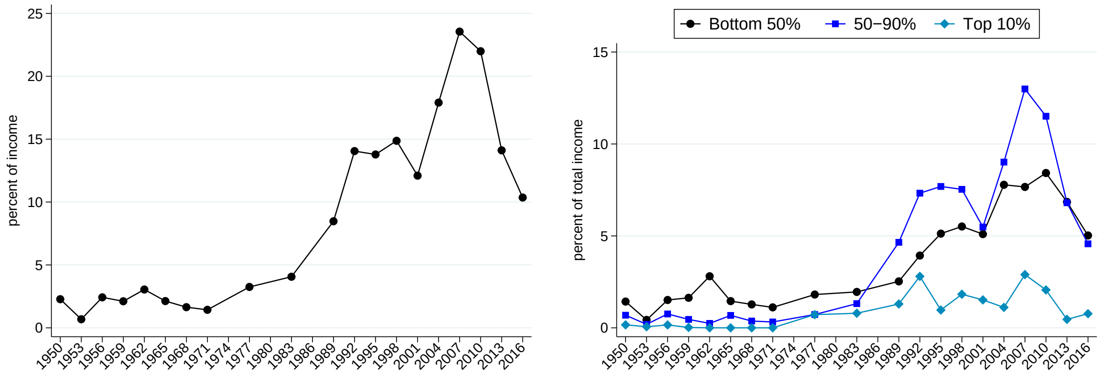

(a) Aggregate

(b) By income group

Notes: The figure shows the value at risk relative to total income after a 39% drop in the main earner's income. The graph is based on SCF+ data. Households are assumed to be at risk if they have a debt-service-to-income ratio > 40%. Right panel shows the value at risk relative to total income after a 20% drop in income and a 20% drop in house prices, stratified by income. Households are assumed to be at risk if they have negative home equity and a debt-service-to-income ratio > 40%. The graph is based on SCF+ data.

income groups evolved in lockstep until 1980, middle-class vulnerability surged over the past four decades.[26].

In Figure 3.6.3, we scale the amount of credit at risk by a different denominator. Instead of income, we use the total amount of equity capital in the U.S. banking system from Jordà et al. (2017). Figure 3.6.3 illustrates how rising debt and debt service ratios brought about a situation where an income shock that could previously have been absorbed overwhelms the ability of the financial system to deal with the potential losses in the post-1990 world. Note also that equity capital in the banks has been reasonable stable relative to income (Jordà et al. 2017) so that the household side was the main driver of this increase in underlying financial risks.

Figure 3.6.3b highlights how the middle class turned from being an anchor of financial stability over almost three decades to the epicenter of financial risks on the eve of the financial crisis. By the end of the 1970s, the value of risk from the middle class corresponded to less than 10% of banks' equity, while this would mean a meaningful shock to the banking sector, it still seems vanishingly small compared to the almost 150% of bank's equity that middle class' value at risk accounted for by 2007. These dynamics underscore the considerable macroeconomic consequences of microeconomic behavior of a middle class that responded to large wealth gains in the housing market by increasing its debt levels. Financial and macroeconomic fragility was rising in the background.

This also implies that a tension for regulation of financial markets exists when *Modigliani meets Minsky*. Through the lens of economic theory, equity extraction can be as welfare-enhancing life cycle consumption smoothing — a key gain of functioning financial mar-

[26] Figure 3.A.3 in the Appendix shows that qualitatively similar patterns emerge when using PSID data.

Figure 3.6.3: Value at risk as share of bank equity

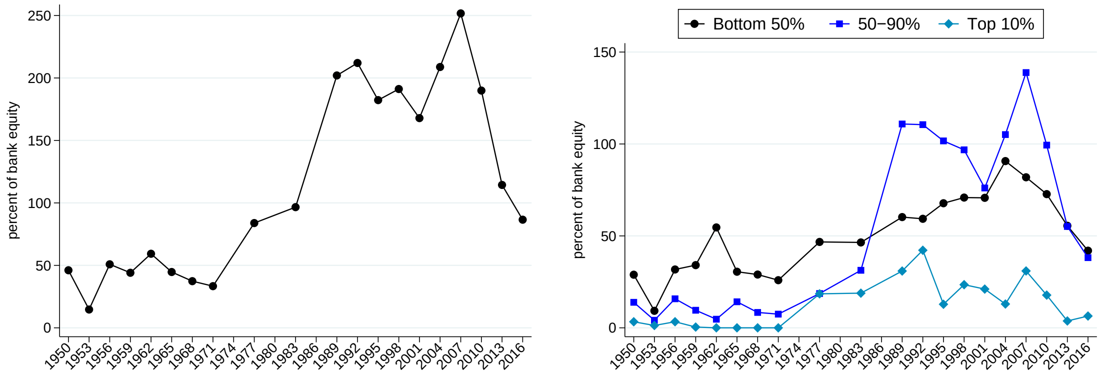

(a) Aggregate

(b) By income group

Notes: The figure shows the value at risk relative to total bank equity from Jordà et al. (2017) after a 39% drop in the main earner's income. The graph is based on SCF+ data. Households are assumed to be at risk if they have a debt-service-to-income ratio > 40%. Right panel shows the value at risk relative to total income after a 20% drop in income and a 20% drop in house prices, stratified by income. Households are assumed to be at risk if they have negative home equity and a debt-service-to-income ratio > 40%. The graph is based on SCF+ data.

kets. Yet, when turning to the macroeconomy, that such life cycle behavior came at the cost of elevated levels of financial fragility. This connects our paper to recent work on excessive leverage and aggregate demand externalities (Korinek and Simsek 2016, Schmitt-Grohé and Uribe 2016), as well as research that discusses the high sensitivity of high-leverage economies to business cycle shocks (Jordà, Schularick and Taylor 2017). According to our stress tests, financial fragility has been substantially reduced since the 2008 crisis. We are back to the fragility levels of the 1980s. Yet the trade-off between the Modigliani and Minsky aspects of equity extraction still remains. As house prices in the United States have been growing again and passed their 2007 peak in 2016, the financial stability implication of the equity extraction boom may soon become highly relevant again.

3.7 Conclusion

This paper studied the increase in household debt in the U.S. since World War II. Relative to income, household debt has risen by a factor of four. Yet the financial history of the United States' postwar surge in household debt remained unwritten. On the basis of long-run household-level data from the SCF+, this paper helps to close this gap. We document the growth of U.S. household debt, its composition and distribution, as well as the link to developments on the asset side of the household balance sheet. We emphasize the nexus between house prices, housing wealth, and equity extraction. House price increases lead to a substantial increase in household wealth, to which optimizing middle-class households responded by extracting home equity via debt. Such home-equity-based borrowing accounts for about half of the increase in U.S. household indebtedness in the past four decades. At the same time, our study documents the increase in financial stability risks that arise when households treat asset-price-induced wealth gains as permanent and borrow against them. This interaction between asset prices and home-equity-based borrowing is central to the surge in household debt since World War II. Our findings provide new and potentially important insights for future research on household portfolio choices and their implications for financial stability.

3.A Appendix

3.A.1 Comparison of PSID Housing Data with the SCF+

In this section, we compare the data on the two main variables of interest, housing and housing debt, from the PSID and the SCF+. Note that the SCF+ is at the household level, whereas the PSID is at the family level. Therefore, PSID families living together were aggregated into one household for better comparability (cf. Pfeffer et al. 2016).[27] The variables were taken from the two surveys as they are, without further harmonization of income, asset and debt concepts (cf. Pfeffer et al. 2016 for a comparison of the survey instruments with respect to wealth). Nominal variables were converted to 2016 dollars using the CPI from the *Macrohistory Database* (Jordà et al. 2017).

Figure 3.A.1 shows that the two data sets yield very similar results at the intensive margin, but there are some differences at the extensive margin. This seems consistent with the results of Pfeffer et al. (2016), who report several differences in asset ownership rates between the SCF and PSID. Conditional on homeownership, average housing debt is similar in both surveys.

Figure 3.A.3 compares the amount of value at risk from our stress testing exercise in Section 3.6. The income distribution differs somewhat between the SCF+ and PSID. While income and debt service are similar on average, the share of households with debt-service-to-income ratio above 0.4 is somewhat lower in the PSID. The qualitative pattern is however consistent with the SCF+ data.

[27] To identify the person among families sharing a household who would most likely have been identified as the head in the SCF+, we create scores based on (a) being male, (b) being the oldest person in the household below retirement age (set to 65), (c) having the highest income within the household, and (d) owning the house. Within each household, the person with the highest score is defined to be the head, and his or her demographics are kept. If there is a tie, we choose the homeowner as the head. If there is still a tie, we choose the senior person, and if there is still a tie, we choose the person with the higher income. Income and wealth variables are summed across all families in the household.

Figure 3.A.1: Comparison of average house value and housing debt: PSID vs. SCF+

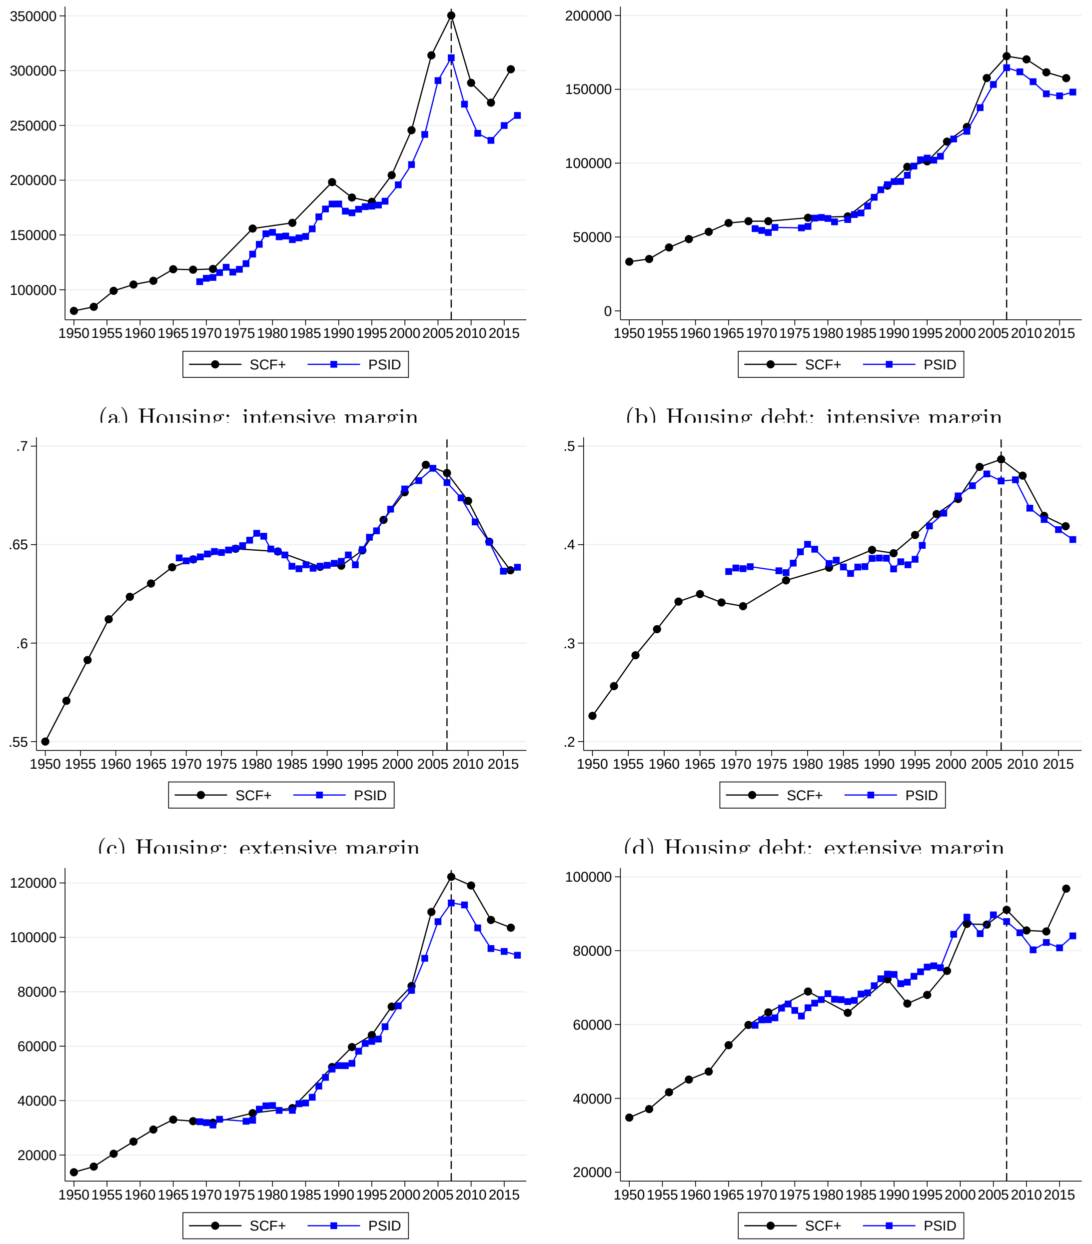

(a) Housing: intensive margin

(b) Housing debt: intensive margin

(c) Housing: extensive margin

(d) Housing debt: extensive margin

(e) Housing debt: homeowners

(f) Total income

Notes: Panel (a) shows the average value of a house, conditional on being a homeowner. Panel (b) shows the average value of housing debt, conditional on having any housing debt. Panel (c) shows the homeownership rate. Panel (d) shows the share of households with positive housing debt. Panel (e) shows average housing debt in the subsample of homeowners. Panel (f) shows total household income. Black lines with dots show SCF+ data, gray lines with squares show PSID data.

Figure 3.A.2: Housing debt-to-income in the SCF+ and PSID

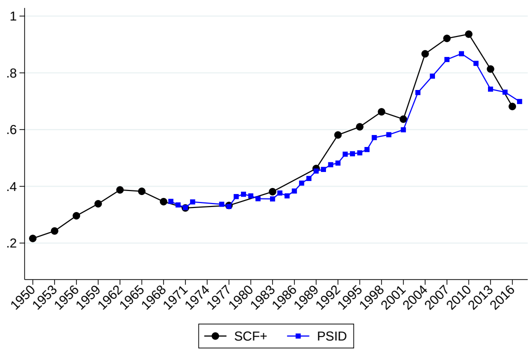

(a) Average

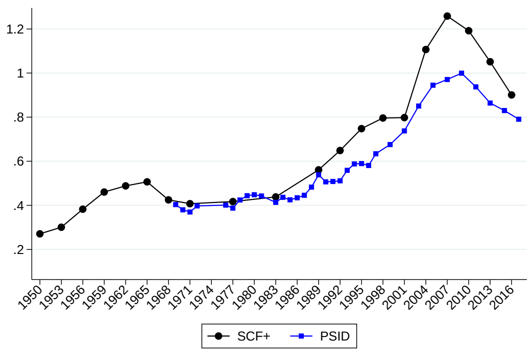

(b) 50-90%

Notes: The graph shows the housing debt-to-income ratio in the SCF+ and PSID over time. The right panel shows results for households from the 50th to 90th percentile of the income distribution only.

Figure 3.A.3: Value at risk relative to income by group

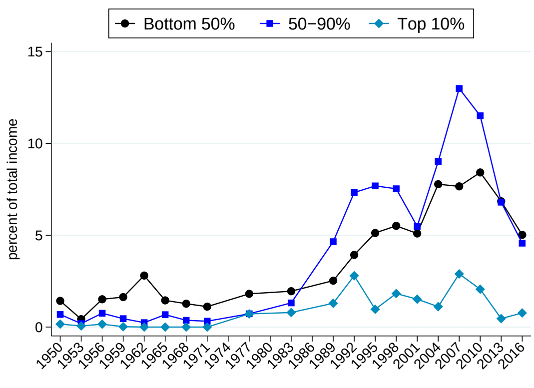

(a) SCF+

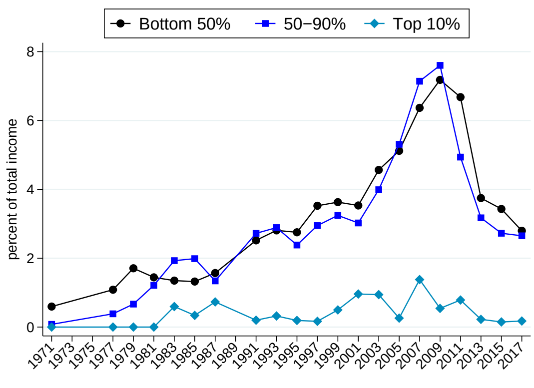

(b) PSID

Notes: The figure shows the value at risk relative to total income in the SCF+ (left panel) and PSID (right panel) after a 20% drop in income and a 20% drop in house prices. The data are stratified by income. Households are assumed to be at risk if they have negative home equity and a debt-service-to-income ratio in excess of 50%.

3.A.2 Geographic variation

This section provides additional information on geographic variation in house values and housing debt from the PSID. While Figure 3.A.4 shows levels. Figure 3.A.5 shows the growth of the average amount of extracted home equity since 1977, and compares this to the increase in the constant-quality house price index from the Federal Housing Finance Agency (FHFA, former OFHEO), which was converted to real terms using the CPI.

Figure 3.A.4: Housing and housing debt by Census region (intensive margin)

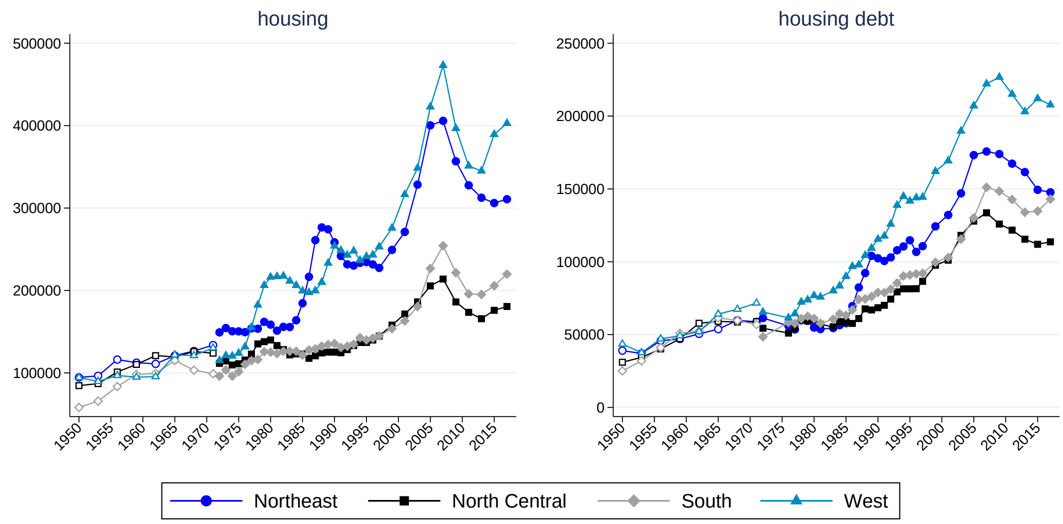

Notes: The graph shows the intensive margin of housing and housing debt by Census region. Filled markers show PSID data, and hollow markers show SFC+ data.

Figure 3.A.5: House prices and equity extraction by Census region

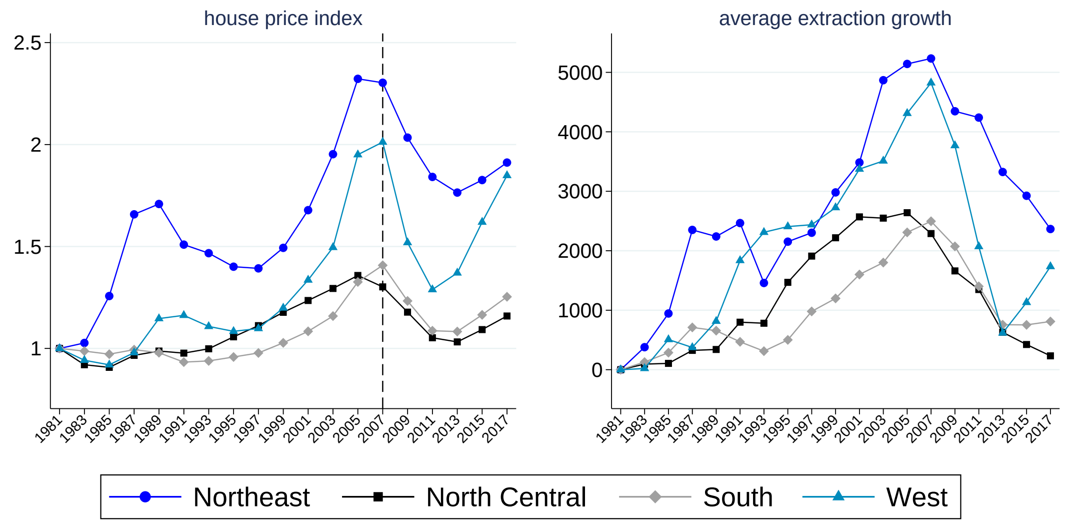

Notes: The left panel shows the growth of state-level FHFA house price indices since 1981, averaged by year and region. The right panel shows the average amount extracted by region, smoothed by taking a moving average across three neighboring waves and normalized by subtracting 1981 levels.

Figure 3.A.6: Growth of housing and extraction by state (intensive margin)

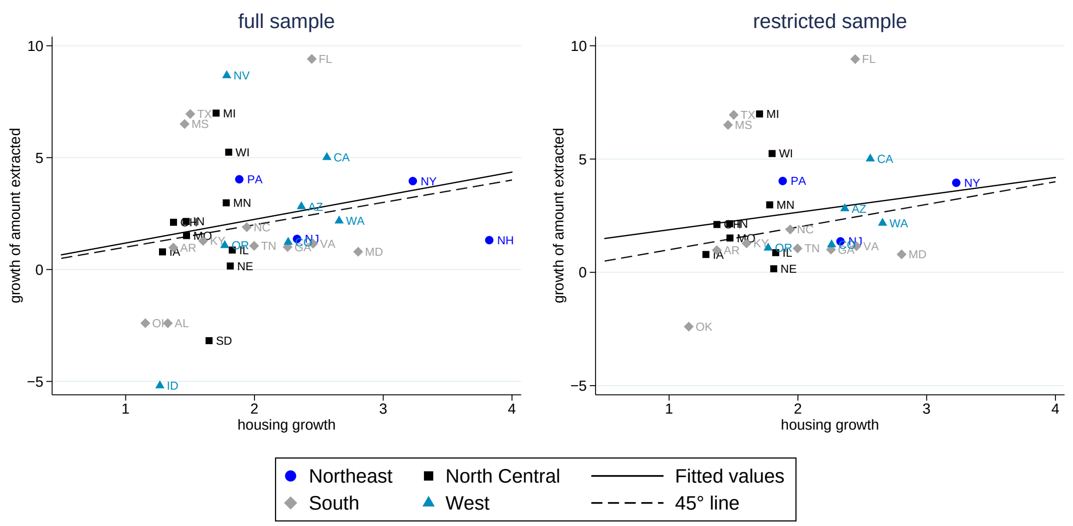

Notes: The graph plots the growth in average house values against the growth in average extraction between 1981 and 2007. Averages are computed at the intensive margin, i.e. conditional on having a house and a mortgage, respectively. The right panel excludes states with less than 50 observations. Massachusetts, Connecticut and South Carolina were excluded as outliers from both panels.

3.A.3 Literature on HEW

Several approaches have been made to quantify the importance of home equity extraction. Bhutta and Keys (2016) estimate that nearly one trillion dollars of equity were extracted between 2002 and 2005 via home equity loans, HELOCS, second mortgages and cash-out refinancings. They exclude the use of funds to move into a more expensive home or buy a second house. According to their calculations, households on average extracted 40,000 dollars between 1999 and 2010. The share of extractors among households with positive mortgage debt holdings varied over time, from 8.5% in 1999 to 18.4% at the peak in 2003. Canner et al. (2002) estimate that around 132 billion dollar were extracted via cash-out refinancings from 2001 to early 2002. They estimate that 16-23% of households with mortgage debt were refinancing, out of which 45% extracted equity.

In the modern SCF, questions on equity extraction via cash-out refinancings and home equity loans exist since 1995, and the amount is elicited since 2004. Out of the households surveyed in 2004, 6.4% had extracted equity between 2002 and 2004, which amounts to 13.4% of all households with positive housing debt. Among those households who extracted between the last and the current SCF wave, the average extracted amount across all available years was 41,200 dollars. Extraction information in the SCF only refers to the first mortgage according to the SCF classification. While the PSID counts mortgages consecutively irrespective of their type, the SCF reports HELOCs in a separate variable. The year of origination is only reported for non-HELOC mortgages. Moreover, there are households who are reported to have a third mortgage without having a first or second mortgage. Therefore, a comparison of the extensive margin of extraction with the PSID is not straightforward. However, the extracted amount conditional on extracting is of a broadly similar magnitude in both surveys.

Table 3.A.1: Average amount extracted

year	PSID	SCF+
1999	32724.29	.
2001	29245.52	.
2003	32835.51	.
2004	.	35185.82
2005	38884.87	.
2007	37185.00	47736.85
2009	39974.38	.
2010	.	34786.23
2011	26629.27	.
2013	29090.20	41825.63
2015	34378.86	.
2016	.	46413.52
2017	40473.45	.

Notes: The table reports the average amount extracted, conditional upon extracting, from the SCF and PSID. The SCF measure is based on first mortgages only, and refers to households who extracted over the current and previous two years.

Greenspan and Kennedy (2008) take a broader perspective, taking into account existing home sales as well. They estimate that on average, HEW generated around 590 billion dollars of free cash per year between 1991 and 2006, out of which two thirds were accounted for by existing home sales. However, their estimates are based on a so-called mortgage system, which was discontinued after 2008, as it did not adequately capture features of the housing market as experienced in the Financial Crisis of 2007 and 2008. Klyuev and Mills (2007) obtain slightly lower but similar estimates with a more simple method. They use the difference between all borrowing secured by dwellings (T_H) and the net acquisition of residential assets (T_{DH}) from the FA as a proxy. The FA mortgage transaction series T_{DH} includes all kinds of mortgages, except construction loans. The housing transaction series T_H includes gross fixed investment in residential structures, net of depreciation, as well as land sales from other sectors to the household sector. However, this "broad" HEW proxy is a somewhat coarse measure of equity extraction. For instance, if a household buys a new home for 100 dollars, and takes out a mortgage for 80 dollars, this measure would count it as *negative* equity extraction (equity injection) of 20 dollars. We compare this measure to our PSID-based equity extraction measure in Figure 3.A.7.

Figure 3.A.7: Comparison to FA measure of Klyuev and Mills (2007)

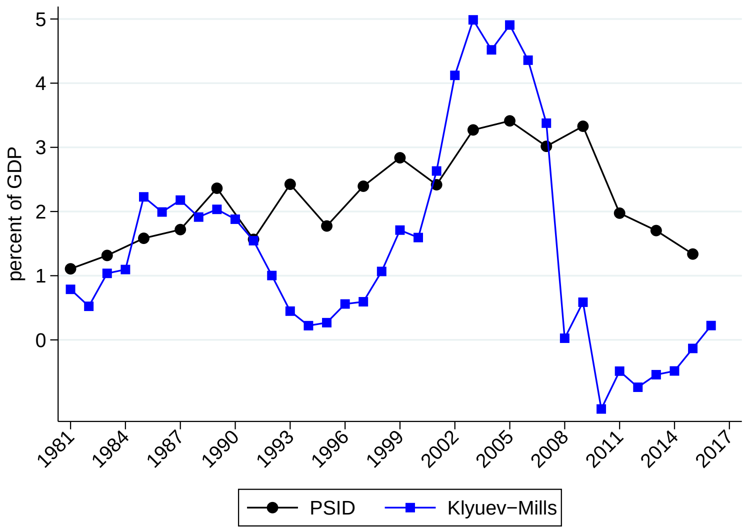

Notes: The figure shows the HEW measure proposed by Klyuev and Mills (2007) and the total amount extracted based on our computations with the PSID, both normalized by NIPA GDP.

3.A.4 Supplemental Figures and Tables

Figure 3.A.13a shows loan-to-value ratios along the income distribution for the same years as in Figure 3.3.5. A strong increase in loan-to-value ratios has occurred since 1983. In 2007, LTVs along the whole income distribution exceeded those from the peak of the first debt boom in 1965. Like debt-to-income, leverage has risen most strongly in the middle of the distribution. While middle-class debt-to-income had decreased again

Figure 3.A.8: Other real estate debt

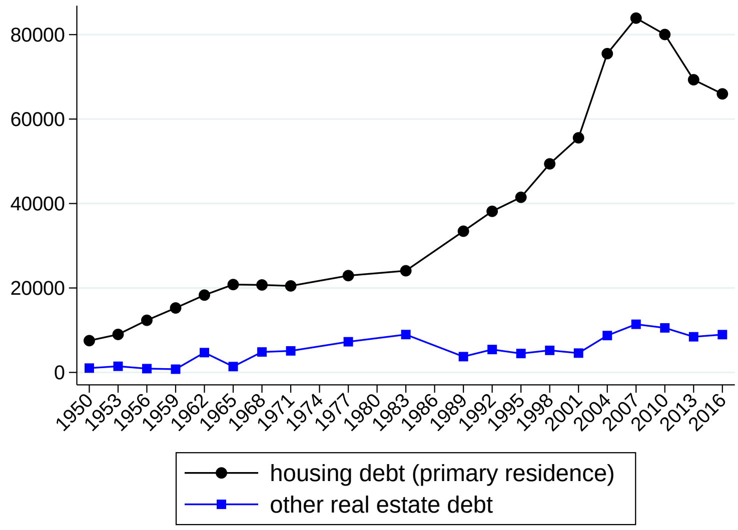

Notes: The graph shows housing debt on owner-occupied real estate in comparison to other real estate debt in the SCF+.

in 2016, LTVs were still similar to 2007 due to the simultaneous decline in house values.

Figure 3.A.15 shows the pairwise correlation of our indicator for equity extraction and the PSID indicator for refinancing of first mortgages. LaCour-Little et al. (2010) and Bhutta and Keys (2016) report extraction booms in 1998 and 2003. This is mirrored in a particularly high correlation around these years.

Figure 3.A.18 shows average first and second mortgages from the SCF since 1983, as well as the share of households having first and second mortgages, respectively, which we observe since 1955. As mentioned above, the SCF counts HELOCs separately, whereas the PSID counts them among the second (or if no other mortgage is held, even the first) mortgages. Therefore, we re-classify HELOCs, which are available in the modern SCFs since 1989, as first mortgages if no other mortgage is available, and as second mortgages if only a first mortgage is recorded. HELOCs were only introduced on a relevant scale in the mid-1980s (see Maki 2001).

Figure 3.A.9: Sensitivity: housing debt and income by income group

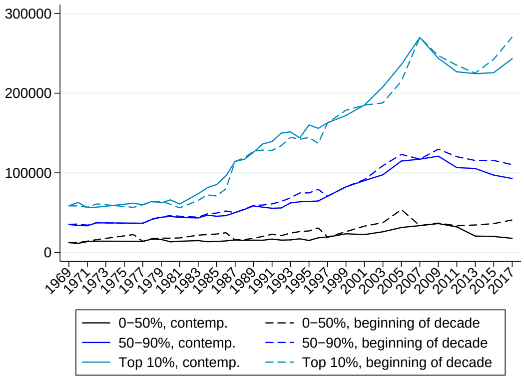

(a) Housing debt

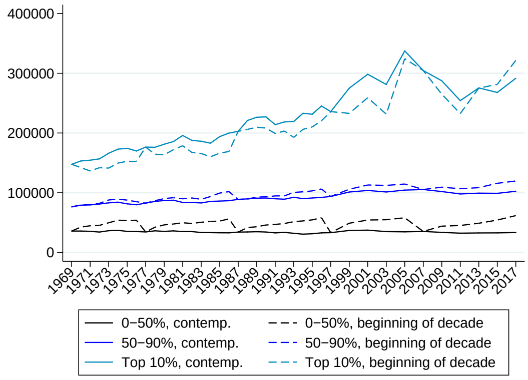

(b) Income

Notes: The graph shows average housing debt (left panel) and income (right panel) by income group for households between ages 30 and 55. We first sort households by their contemporaneous income, and show the results as solid lines. For comparison, we sort households by their income at the beginning of each decade (1969, 1977, 1987, 1997, 2007). These results are shown as dashed lines.

Figure 3.A.10: Personal debt, extensive margin

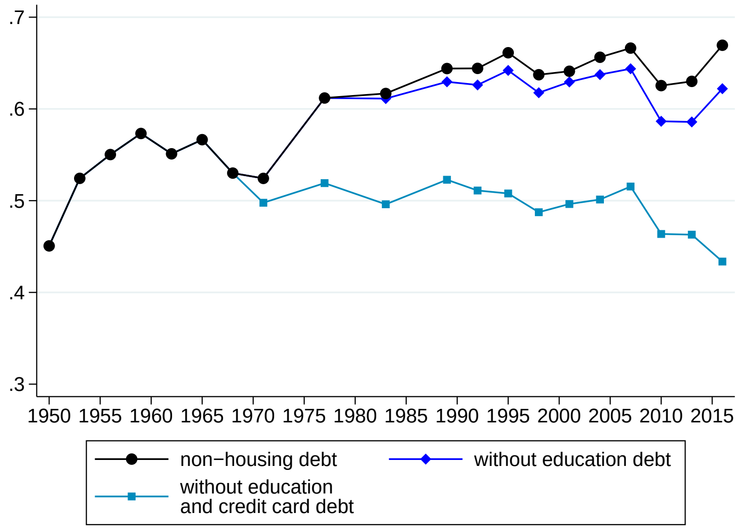

Notes: The graph shows the extensive margin of personal debt from Figure 3.3.6, together with counterfactuals in which credit card and education debt were set to zero.

Table 3.A.2: Income group stability

year	Bottom 50%	50 - 90%	Top 10%
1970	0.85	0.73	0.66
1971	0.85	0.74	0.69
1972	0.86	0.74	0.67
1973	0.86	0.74	0.64
1974	0.85	0.75	0.66
1975	0.85	0.75	0.67
1976	0.84	0.75	0.65
1977	0.85	0.75	0.62
1978	0.86	0.75	0.66
1979	0.86	0.74	0.64
1980	0.86	0.76	0.67
1981	0.86	0.77	0.65
1982	0.85	0.75	0.65
1983	0.83	0.75	0.69
1984	0.85	0.77	0.70
1985	0.86	0.75	0.65
1986	0.86	0.74	0.64
1987	0.83	0.74	0.63
1988	0.83	0.75	0.68
1989	0.85	0.74	0.71
1990	0.86	0.77	0.73
1991	0.86	0.77	0.70
1992	0.84	0.75	0.68
1993	0.83	0.75	0.64
1994	0.83	0.72	0.61
1995	0.83	0.74	0.60
1996	0.83	0.74	0.62
1997	0.83	0.72	0.63
1999	0.83	0.74	0.61
2001	0.81	0.73	0.64
2003	0.82	0.74	0.65
2005	0.84	0.76	0.67
2007	0.85	0.78	0.69
2009	0.85	0.76	0.64
2011	0.85	0.76	0.69
2013	0.86	0.77	0.70
2015	0.86	0.76	0.70
2017	0.84	0.77	0.74

Notes: The table reports the share of households who stayed in their respective income group since two years ago.

Figure 3.A.11: Debt and income growth

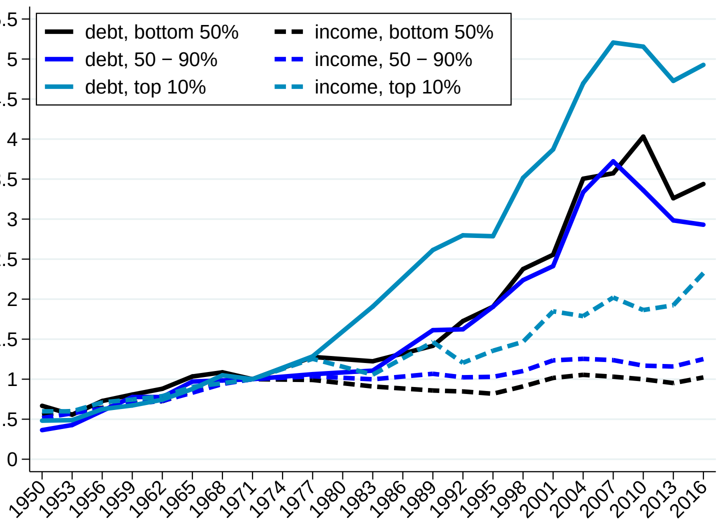

Notes: The graoh shows the growth of average total housing debt and income by income group, relative to 1971.

Figure 3.A.12: Debt-to-asset ratios

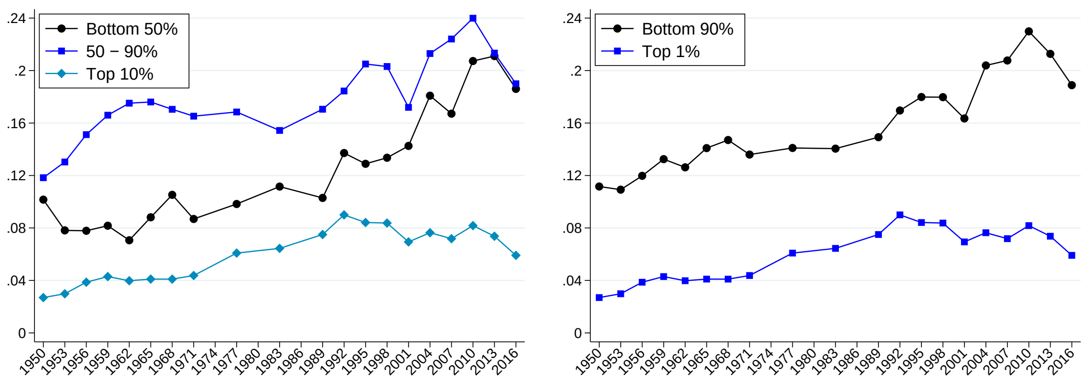

(a) Debt-to-asset ratio

(b) Debt-to-asset ratio

Notes: The left panel shows housing debt-to-asset ratios for the bottom 50%, 50-90% and top 10% of the income distribution. The right panel compares debt-to-asset ratios of the bottom 90% and top 1%.

Figure 3.A.13: LTV and debt-to-assets along the income distribution

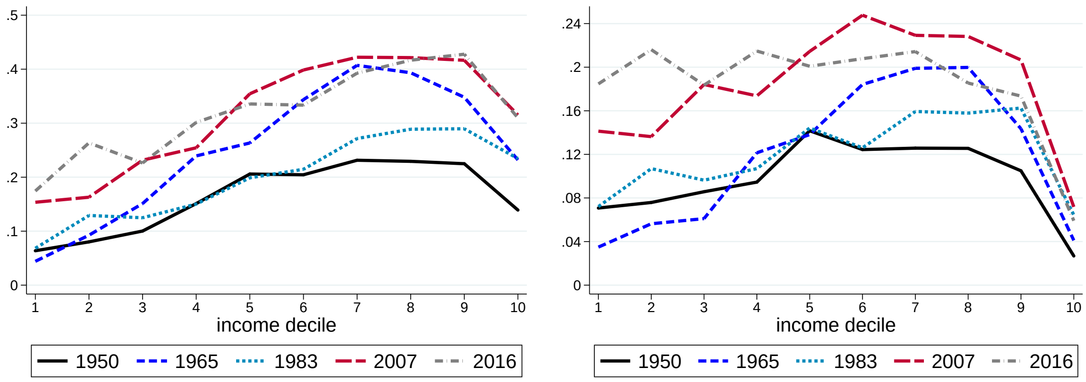

(a) Loan-to-value (b) Debt-to-asset

Notes: The left panel shows the evolution of average loan-to-value ratios by deciles of the aggregate income distribution for the SCF+ waves 1950, 1965, 1983, 2007 and 2016. The right panel shows the evolution of total debt to total assets. We exclude households with total income below 10% of the annual wage of a household with a single earner receiving the contemporaneous minimum wage.

Table 3.A.3: Portfolio shares

year	bonds + liquid and oth. fin. assets	stocks + business	non-fin. assets	housing	housing debt	non-housing debt
			Bottom 50%			
1950	15.05	26.65	12.73	45.57	55.54	44.46
1965	9.60	35.55	15.75	39.10	73.63	26.37
1983	18.15	15.21	16.44	50.21	66.17	33.83
2007	17.56	16.51	15.75	50.18	71.31	28.69
2016	18.57	20.43	17.46	43.53	62.44	37.56
			50 - 90%			
1950	13.70	36.21	9.58	40.51	77.80	22.20
1965	10.01	38.64	10.22	41.13	84.57	15.43
1983	20.17	16.20	17.15	46.48	79.08	20.92
2007	22.74	17.29	13.32	46.65	83.78	16.22
2016	29.15	19.54	13.34	37.97	77.01	22.99
			Top 10%			
1950	9.23	75.15	3.72	11.90	67.48	32.52
1965	6.92	69.21	7.39	16.48	89.21	10.79
1983	18.71	44.05	16.98	20.25	68.18	31.82
2007	20.60	45.22	13.26	20.91	88.41	11.59
2016	22.44	48.97	11.98	16.61	83.06	16.94

Figure 3.A.14: Mortgage interest rates (positive housing debt)

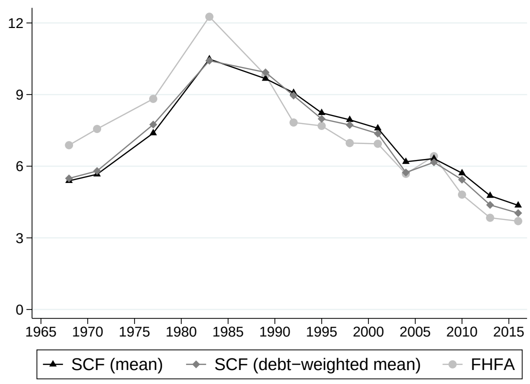

(a) Nominal

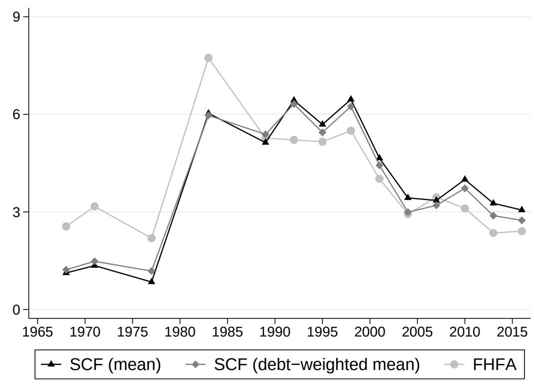

(b) Real

Notes: The graph shows average interest rates on first mortgages in the SCF+ among households with positive housing debt. The left panel presents nominal interest rates i_t^m. Real interest rates in the right panel were calculated as $r_t^m = \left[(1+i_t^m)/(1+\pi_t) - 1\right] \cdot 100$, where π_t denotes year-on-year CPI inflation. The black lines with triangles present the simple average, whereas the medium gray lines with diamonds present the housing-debt-weighted average. As a comparison, the light gray lines with dots show the average interest rate on conventional non-farm single-family mortgages on new and previously occupied homes from the Monthly Interest Rate Survey of the FHFA. The survey excludes FHA-insured and VA-guaranteed loans, loans on multifamily buildings and mobile homes, as well as refinancing loans. Note that the SCF+ data shown in this figure have not yet been subject to imputation. The data in 1967 were top-coded at 9.9%, 9.7% in 1968-1970, and 20% in 1977.

Figure 3.A.15: Extraction and refinancing

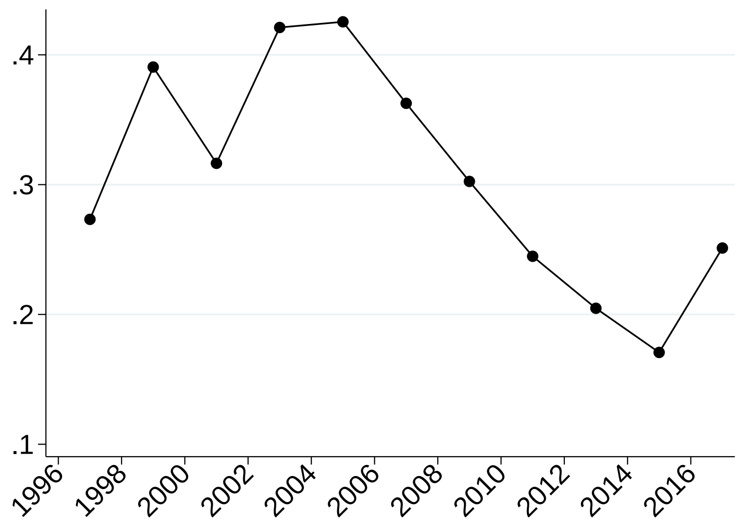

Notes: The graph shows the pairwise correlation of our indicator for equity extraction and the PSID indicator for refinancing of first mortgages over time.

Figure 3.A.16: Home buyers by type

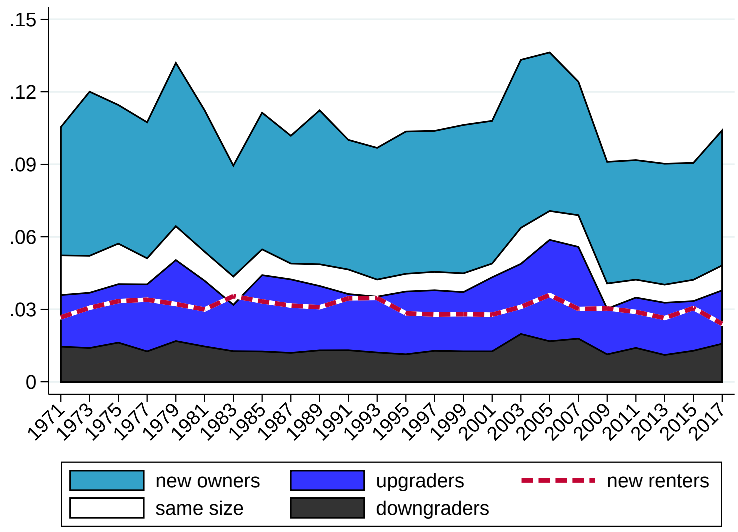

Notes: The figure decomposes home buyers into new owners, upgraders, downgraders and households moving to a similarly-sized home. Please refer to the text for the exact definitions of these groups. The dashed line shows the share of households who sell their home to become a renter for at least two waves.

Figure 3.A.17: Intensive and extensive margin by type

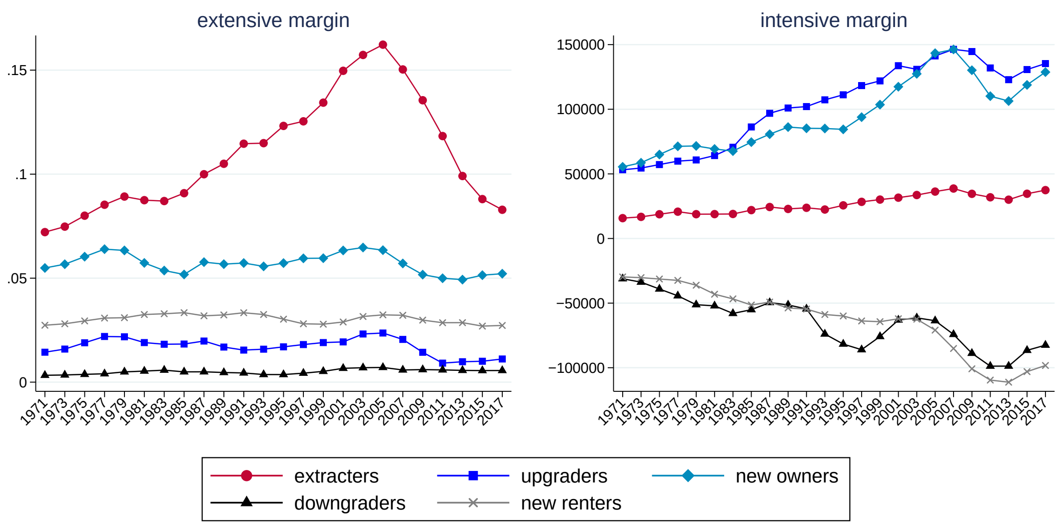

Notes: The left panel shows the share of households who extracted equity, upgraded, downgraded, bought a new home, or sold their home to become a renter. The right panel shows the average debt increase of these households. The series were smoothed by taking a moving average across three neighboring waves.

Figure 3.A.18: First and second mortgages, SCF+

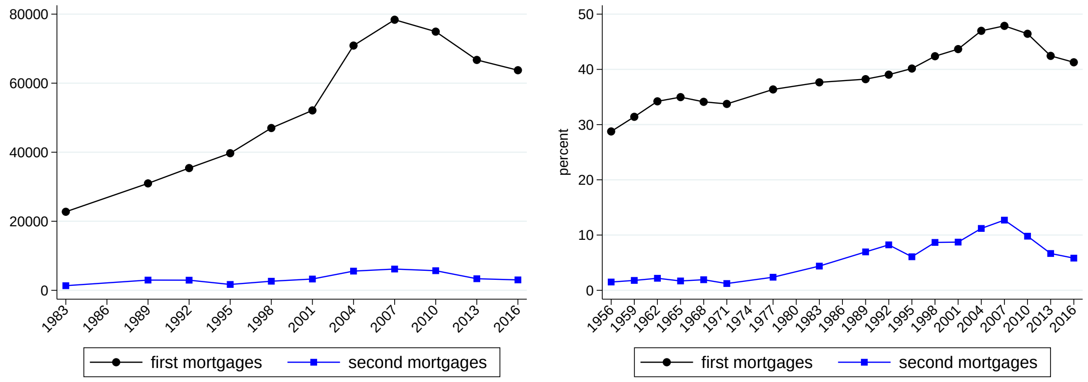

(a) Average (b) Extensive margin

Notes: The left panel shows average first and second mortgages from the SCF. The right graph shows the share of households who have first or second mortgages. HELOCs are included (see text for details).

Figure 3.A.19: First mortgages, PSID

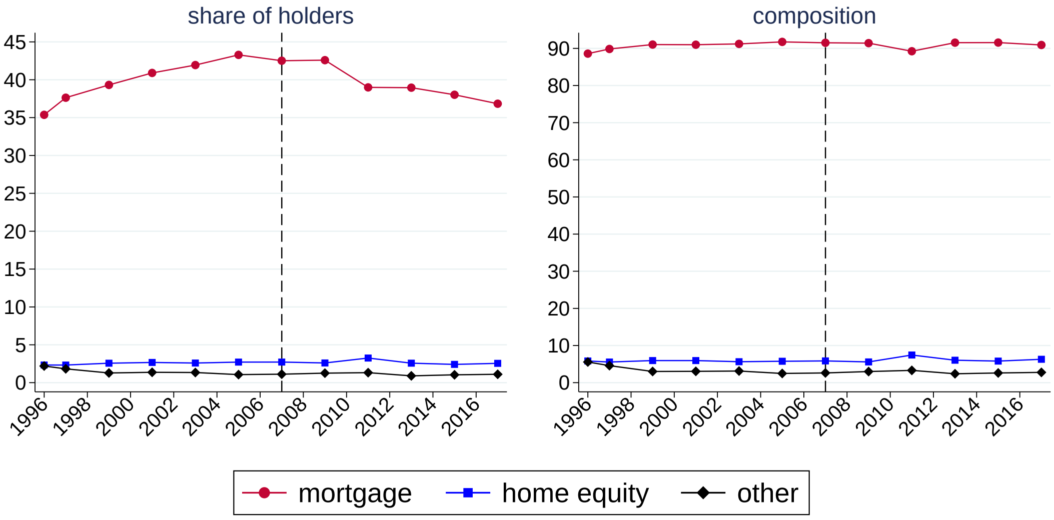

Notes: The left panel shows the share of households in the PSID who hold the respective type of mortgage. The right panel shows the share conditional upon having a first mortgage.

Figure 3.A.20: Second mortgages, PSID

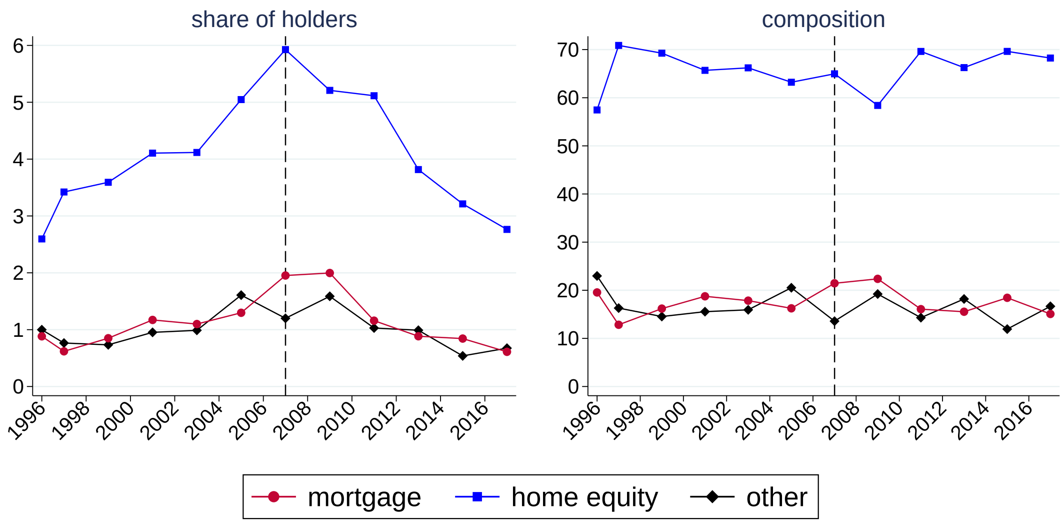

Notes: The left panel shows the share of households in the PSID who hold the respective type of mortgage. The right panel shows the share conditional upon having a second mortgage.

Figure 3.A.21: Rising house prices and debt of young homeowners

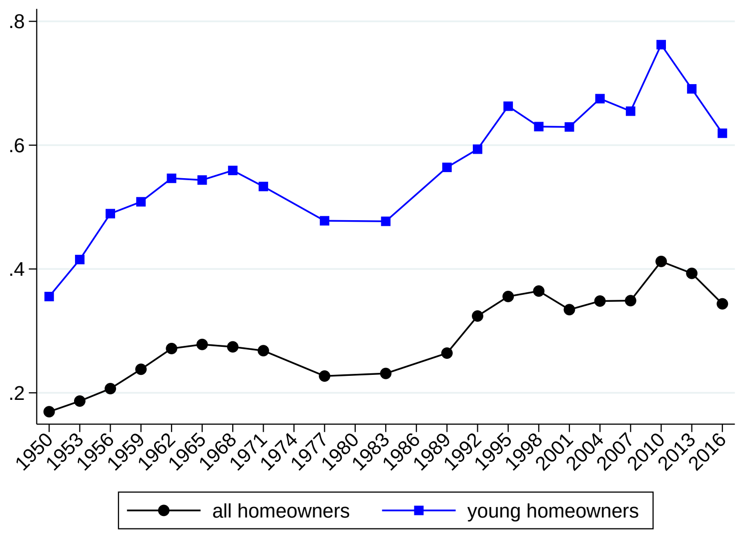

Notes: The graph shows average housing debt relative to average housing for all homeowners and homeowners with a head below age 35.

Figure 3.A.22: Event study: extraction

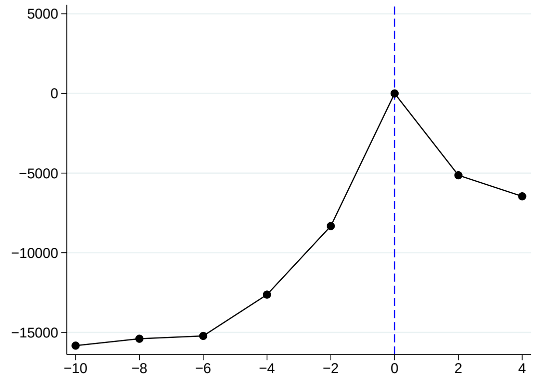

(a) All income groups

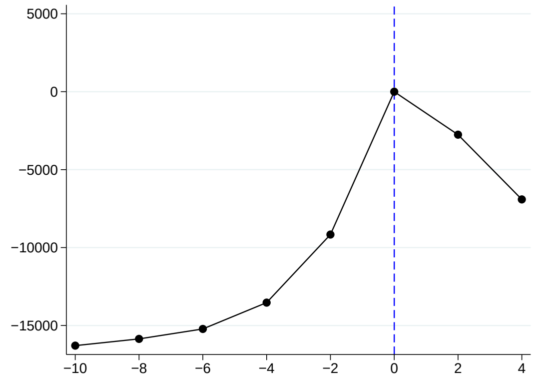

(b) 50-90% only

Notes: The graph shows regressions of the house value on leads and lags of the extraction dummy. Zero is the period of extraction. Even years were discarded from the data set to avoid a change in frequency, just as for the local projections in Section 3.4.4. We focus on households stay in their home upon extraction. The regressions include year and household fixed effects.

Figure 3.A.23: Google Books Ngram Viewer

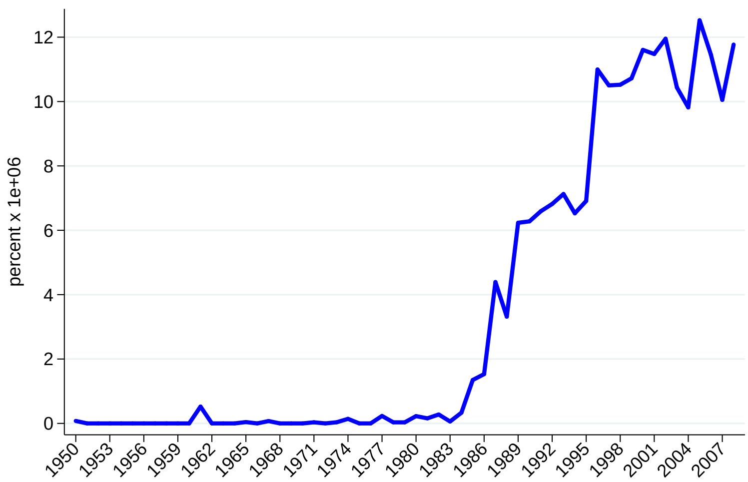

Notes: The graph shows how mentions on the 3-gram “home equity loan” have evolved over time. The figure is based on data from the Google Books Ngram Viewer. The y-axis shows the share of this 3-gram among all 3-grams contained in the Google sample of books written in English and published in the United States. The Google data are normalized with the total number of books published in each year.

Figure 3.A.24: Growth of housing and housing debt by income group, PSID

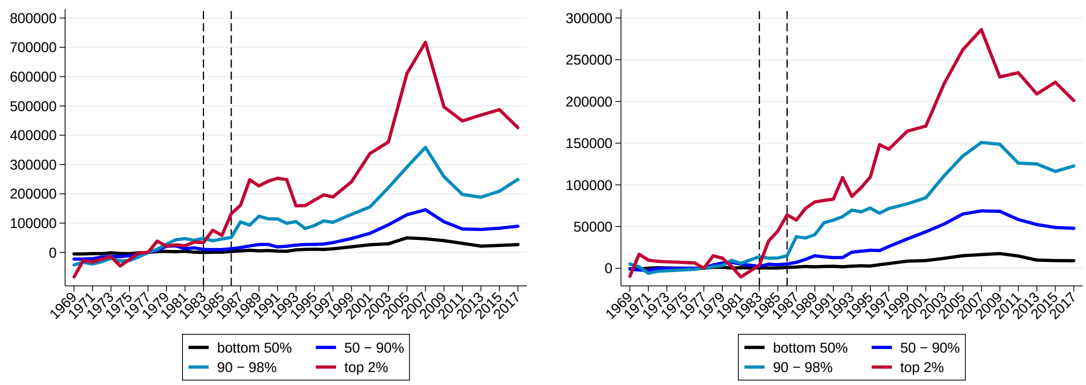

(a) Share of households

(b) Value at risk

Notes: The left panel shows the growth of housing by income group. The right panel shows the growth of housing debt. The graph is based on PSID data.

3.A.5 Details for life cycle model

3.A.5.1 Derivations

Here, we derive the optimal policies and law of motion for the Modigliani life cycle model from Section 4.4 in the main part of the paper. The agent's problem reads

$$\begin{aligned} \max_{\{c_j,h_j,d_{j+1}\}_{j=0}^{J}} \quad & \sum_{j=0}^{J} \beta^j \Big(\rho \log(c_j) + (1-\rho)\log(h_j) \Big) \\ s.t. \quad & c_j + p_h h_j - d_{j+1} = y_j - (1+r)d_j + (1-\delta)h_{j-1}p_h \\ & h_{-1},\ d_0 \qquad \text{given} \end{aligned} \tag{3.A.1}$$

First-order conditions deliver

$$\frac{1}{c_j}\rho p_h = (1-\rho)\frac{1}{h_j} + \beta\rho(1-\delta)p_h\frac{1}{c_{j+1}} \tag{3.A.2}$$

$$\frac{1}{c_j} = \beta(1+r)\frac{1}{c_{j+1}} \tag{3.A.3}$$

From equation (3.A.3), we get the optimal path of consumption growth

$$c_j = (\beta(1+r))^j c_0 \tag{3.A.4}$$

Using the Euler equation (3.A.3) in equation (3.A.2) delivers

$$\begin{aligned} \rho p_h &= (1-\rho)\frac{c_j}{h_j} + \beta\rho(1-\delta)p_h\frac{c_j}{c_{j+1}} \\ 1 &= \frac{1-\rho}{\rho}\frac{c_j}{p_h h_j} + \beta(1-\delta)(\beta(1+r))^{-1} \\ p_h h_j &= \frac{1-\rho}{\rho}c_j + \frac{1-\delta}{1+r}p_h h_j \\ p_h h_j &= \frac{1+r}{r+\delta}\frac{1-\rho}{\rho}c_j \end{aligned} \tag{3.A.5}$$

with the standard constant expenditure share result. Note that expenditures for housing are the user costs $\frac{r+\delta}{1+r}p_h h_j$. Combining equation (3.A.5) with the Euler equation delivers

$$p_h h_j = \frac{1+r}{r+\delta}\frac{1-\rho}{\rho}(\beta(1+r))^j c_0 \tag{3.A.6}$$

The law of motion for the debt level is

$$d_{j+1} = c_j - y_j + p_h h_j + (1+r)d_j - (1-\delta)h_{j-1}p_h \tag{3.A.7}$$

Using this law of motion and plugging in recursively delivers

$$d_{j+1} = \sum_{s=0}^{j}(c_s - y_s)(1+r)^{j-s} + p_h h_j + \sum_{s=0}^{j-1} p_h h_s (r+\delta)(1+r)^{j-1-s} - (1+r)^j((1-\delta)h_{-1}p_h - (1+r)d_0) \tag{3.A.8}$$

For $j = J$, we get

$$d_{J+1} = \sum_{s=0}^{J}(c_s - y_s)(1+r)^{J-s} + p_h h_J + \sum_{s=0}^{J-1} p_h h_s (r+\delta)(1+r)^{J-1-s} - (1+r)^J((1-\delta)h_{-1}p_h - (1+r)d_0) \tag{3.A.9}$$

Now we multiply both sides by $(1+r)$ and subtract $(1-\delta)p_h h_J$

$$d_{J+1}(1+r) - (1-\delta)p_h h_J = (1+r)^{J+1}\Bigg(\sum_{s=0}^{J} \frac{c_j - y_j}{(1+r)^s} + \frac{(1+r)p_h h_J - (1-\delta)p_h h_J}{(1+r)^{J+1}}$$

$$+ \frac{1}{1+r}\sum_{s=0}^{J-1} \frac{p_h h_s}{(1+r)^s}(r+\delta) - \Big((1-\delta)h_{-1}p_h - (1+r)d_0\Big)\Bigg)$$

$$\frac{d_{J+1}(1+r) - (1-\delta)p_h h_J}{(1+r)^{J+1}} = \sum_{s=0}^{J} \frac{c_j - y_j}{(1+r)^s} + \frac{(r+\delta)p_h h_J}{(1+r)^{J+1}}$$

$$+ \frac{r+\delta}{1+r}\sum_{s=0}^{J-1} \frac{p_h h_s}{(1+r)^s} - \Big((1-\delta)h_{-1}p_h - (1+r)d_0\Big)$$

$$\frac{d_{J+1}(1+r) - (1-\delta)p_h h_J}{(1+r)^{J+1}} = \sum_{s=0}^{J} \frac{c_j}{(1+r)^s} - \overbrace{\sum_{s=0}^{J} \frac{y_j}{(1+r)^s}}^{=Y}$$

$$+ \frac{r+\delta}{1+r}\sum_{s=0}^{J} \frac{p_h h_s}{(1+r)^s} - \underbrace{\Big((1-\delta)h_{-1}p_h - (1+r)d_0\Big)}_{=E}$$

$$\frac{d_{J+1}(1+r) - (1-\delta)p_h h_J}{(1+r)^{J+1}} = \sum_{s=0}^{J} \frac{c_j}{(1+r)^s} + \frac{r+\delta}{1+r}\sum_{s=0}^{J} \frac{p_h h_s}{(1+r)^s} - (E+Y) \quad (3.A.10)$$

Under the optimal policy it is always optimal that all resources are consumed in the last period, so that equity at the end of the life cycle is zero $E' = (1-\delta)p_h h_J - d_{J+1}(1+r) = 0$. This implies that the left-hand side of equation (3.A.10) must be zero for the solution to be optimal and we obtain

$$E + Y = \sum_{s=0}^{J} \frac{c_j}{(1+r)^s} + \frac{r+\delta}{1+r}\sum_{s=0}^{J} \frac{p_h h_s}{(1+r)^s} \quad (3.A.11)$$

Now we plug in equations (3.A.4) and (3.A.6) and obtain

$$\begin{aligned}
\underbrace{E+Y}_{=W} &= \sum_{s=0}^{J} \frac{c_0(\beta(1+r))^s}{(1+r)^s} + \frac{r+\delta}{1+r}\sum_{s=0}^{J} \frac{\frac{1+r}{r+\delta}\frac{1-\rho}{\rho}(\beta(1+r))^s c_0}{(1+r)^s} \\
W &= c_0\sum_{s=0}^{J}\beta^s + \frac{1-\rho}{\rho}c_0\sum_{s=0}^{J}\beta^s \\
W &= c_0\frac{1-\beta^{J+1}}{1-\beta} + \frac{1-\rho}{\rho}c_0\frac{1-\beta^{J+1}}{1-\beta} \\
\underbrace{\frac{1-\beta}{1-\beta^{J+1}}}_{=\alpha}W &= \frac{1}{\rho}c_0 \\
\rho\alpha W &= c_0^* \qquad (3.A.12)
\end{aligned}$$

The law of motion from equation (3.A.13) follows directly from iterating equation (3.A.7)

$$d_{j+1} = \sum_{s=0}^{j}(c_s - y_s)(1+r)^{j-s} + \sum_{s=0}^{j}(p_h h_s - (1-\delta)p_h h_{s-1})(1+r)^{j-s} + (1+r)^{j+1}d_0 \quad (3.A.13)$$

Rearranging terms, we get the expression from equation (3.A.8) and plug in the result for the constant expenditure shares to obtain

$$\begin{aligned}
d_{j+1} &= \underbrace{\sum_{s=0}^{j} c_s(1+r)^{j-s}}_{\text{consumption costs}} - \underbrace{\sum_{s=0}^{j} y_s(1+r)^{j-s}}_{\text{income}} + \underbrace{p_h h_j}_{\text{current housing}} \\
&\quad + \underbrace{\sum_{s=0}^{j-1} p_h h_s \frac{r+\delta}{1+r}(1+r)^{j-s}}_{\text{user costs}} - (1+r)^j \underbrace{\Big((1-\delta)h_{-1}p_h - (1+r)d_0\Big)}_{\text{initial endowment}} \\
d_{j+1} &= \sum_{s=0}^{j} c_s(1+r)^{j-s} - \sum_{s=0}^{j} y_s(1+r)^{j-s} + p_h h_j - (1+r)^j(1-\delta)h_{-1}p_h \\
&\quad + \sum_{s=0}^{j-1} \frac{1-\rho}{\rho} c_s(1+r)^{j-s} + (1+r)^{j+1}d_0 \\
\underbrace{\frac{d_{j+1}}{(1+r)^j}}_{\substack{\text{present value}\\ \text{of debt}}} &= \overbrace{\sum_{s=0}^{j} \frac{c_s}{(1+r)^s} + \sum_{s=0}^{j-1} \frac{1-\rho}{\rho}\frac{c_s}{(1+r)^s}}^{\substack{\text{present value of}\\ \text{total expenditures}}} - \overbrace{\sum_{s=0}^{j} \frac{y_s}{(1+r)^s}}^{\substack{\text{present value}\\ \text{of income}}} \\
&\quad + \underbrace{\left(\frac{p_h h_j}{(1+r)^j} - (1-\delta)h_{-1}p_h\right)}_{\substack{\text{present value of}\\ \text{housing adjustments}}} + \underbrace{(1+r)d_0}_{\substack{\text{(present value)}\\ \text{initial debt}}} \qquad (3.A.14)
\end{aligned}$$

3.A.5.2 Discussion

In the model, households will reduce housing consumption after a positive house price shock, but housing wealth $(1-\delta)p_h h$ will increase nonetheless.[28] This implies that our stylized model predicts that households will not upgrade to larger/better houses after a positive house price shock. A key reason for that is that the stylized model abstracts from borrowing constraints and adjustment costs.[29]

In turn, the model predicts too much *downgrading*: households buy less/worse housing after a positive house price shock. Introducing trading and adjustment costs would allow to match the empirically observed patterns more closely. Moreover, the model abstracts from renters. Current renters constitute the pool of potential new homeowners who are affected by rising house prices.

When house prices rise, households who switch from renting to owning have to pay more for a home of a given size. Hence, new homeowners will have to rely on additional debt to finance their home, buy a smaller house, or postpone home ownership. The data

[28] The elasticity of housing with respect to prices is $\frac{\partial h}{\partial p_h}\frac{p_h}{h} = \theta_h - 1$, so $\frac{\partial(p_h h)}{\partial p_h}\frac{p_h}{p_h h} = \theta_h$.

[29] Without borrowing constraints and adjustment costs, households react immediately to a positive shock to house prices and substitute away from housing. If, however, households are borrowing constrained, then a shock that increases home equity slackens the constraint and allows them to upgrade. The idea that upgrading households use (part of) their equity gain for the down payment of a new home has been discussed, for example, in Genesove and Mayer (1997).

suggest that during the housing boom, many new homeowners relied on additional debt to finance their new home (see Appendix Figure 3.A.21).

In our stylized environment, we do not consider other ways how extracted equity could be used other than for non-durable consumption. Empirical studies have found that home equity is also used for home improvements, the repayment of personal debt, or the foundation of a business (see Mian and Sufi 2011, Cloyne et al. 2017, Greenspan and Kennedy 2008).

Finally, it should be noted that we abstract from other factors beyond house prices which have likely contributed to an increase of debt financing since the 1980s, such as lower mortgage interest rates and higher inflation which raised the attractiveness of debt financing, falling mortgage transaction costs, the disappearing of mortgage prepayment penalties, or the rising costs of financing childrens' education (see e.g. Bhutta and Keys 2016, Canner et al. 2002, Greenspan and Kennedy 2008, Cooper 2010).

3.A.6 Fragility with double trigger

Figure 3.A.25: Value at risk as share of income, double trigger

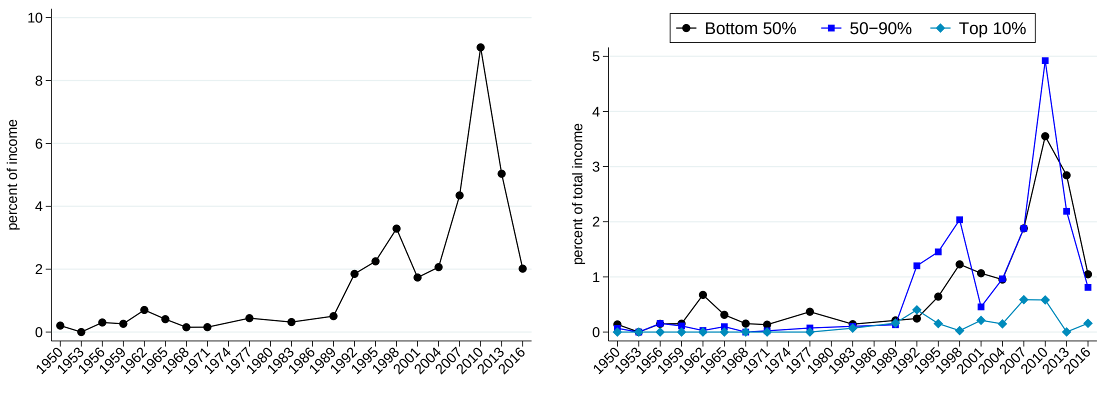

(a) Aggregate (b) By income group

Notes: The figure shows the value at risk relative to total income after a 39% drop in the main earner's income and an 8% drop in house prices. The graph is based on SCF+ data. Households are assumed to be at risk if they have negative home equity and a debt-service-to-income ratio > 40%. Right panel shows the value at risk relative to total income after a 20% drop in income and a 20% drop in house prices, stratified by income. Households are assumed to be at risk if they have negative home equity and a debt-service-to-income ratio > 40%. The graph is based on SCF+ data.

3.A.7 Life-cycle patterns and old-age debt in the PSID

Figure 3.A.26 shows life-cycle loan-to-value profiles obtained by regressing individual loan-to-value ratios on six age group dummies (25-34, 35-44, 45-54, 55-64, 65-74, and 75-85 years). The left panel repeats the results from the SCF+ data from section 3.5 for comparison. The middle panel shows PSID data treated analogously to the SCF+ data, and the right panel shows results which exploit the panel dimension of the PSID by

Figure 3.A.26: Comparison of Life-Cycle Loan-to-Value

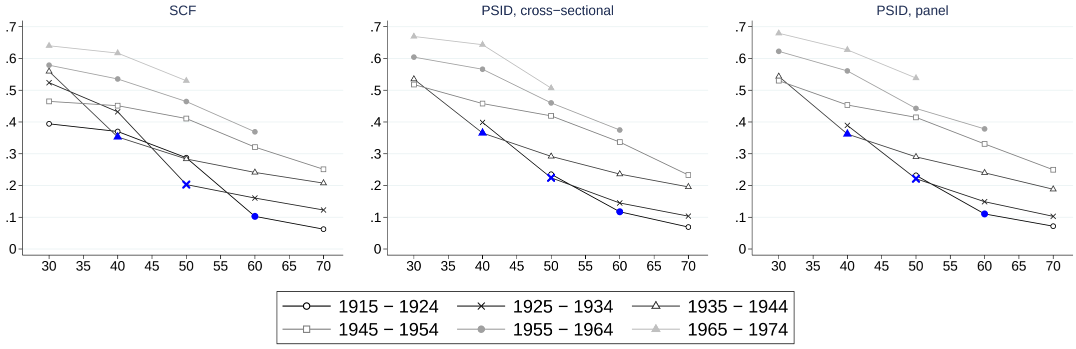

Notes: The graph shows life-cycle loan-to-value profiles for different cohorts. The left panel shows the SCF+ data, the middle panel shows PSID data when treating the data as cross-sectional, and the right panel shows PSID data when exploiting the panel dimension by including household fixed effects.

including household fixed effects. Note that the SCF+ data start in 1950, whereas the PSID data only begin in 1969.[30] Figure 3.A.27 shows analogous results for the housing debt-to-income ratio. It also includes a fourth panel, in which we exploited the PSID's panel dimension to replace income by its three-year moving average (MA) within each household. This helps to avoid extreme values due to temporary income fluctuations.

The result are quantitatively and qualitatively similar across both data sets and all specifications: Housing debt-to-income ratios and leverage (loan-to-value) have both shifted and turned upwards conspicuously. There is a visible shift in slopes around 1980 for all cohorts, no matter whether they were 40, 50 or 60 years at this point (see blue markers). The shift is most pronounced for households around age 40 in 1980. The results are very similar when controlling for household fixed effects in the PSID, which confirms that the results obtained with the SCF+ are not artifacts of working with synthetic cohorts.

Figure 3.A.28 looks at the share of different age groups among all home buyers (left panel) and equity extracters (right panel). We can see that the share of buyers and extracters between 40 and 49 years of age has increased substantially between the mid 1980s and late 1990s. Likewise, the share of buyers and extracters in their fifties has increased markedly since the mid 1990s. This implies that cohorts born after 1940 began to buy more houses and extract more home equity after 1980. In accordance with this evidence, the median age of buyers has increased from around 32 in the early 1980s to 39 in the early 2000s, and the median age of extracters has increased from 40 to 49.

[30]The first PSID wave from 1968 was excluded, as many important variables are still missing in this year.

Figure 3.A.27: Comparison of Life-Cycle Housing Debt-to-Income

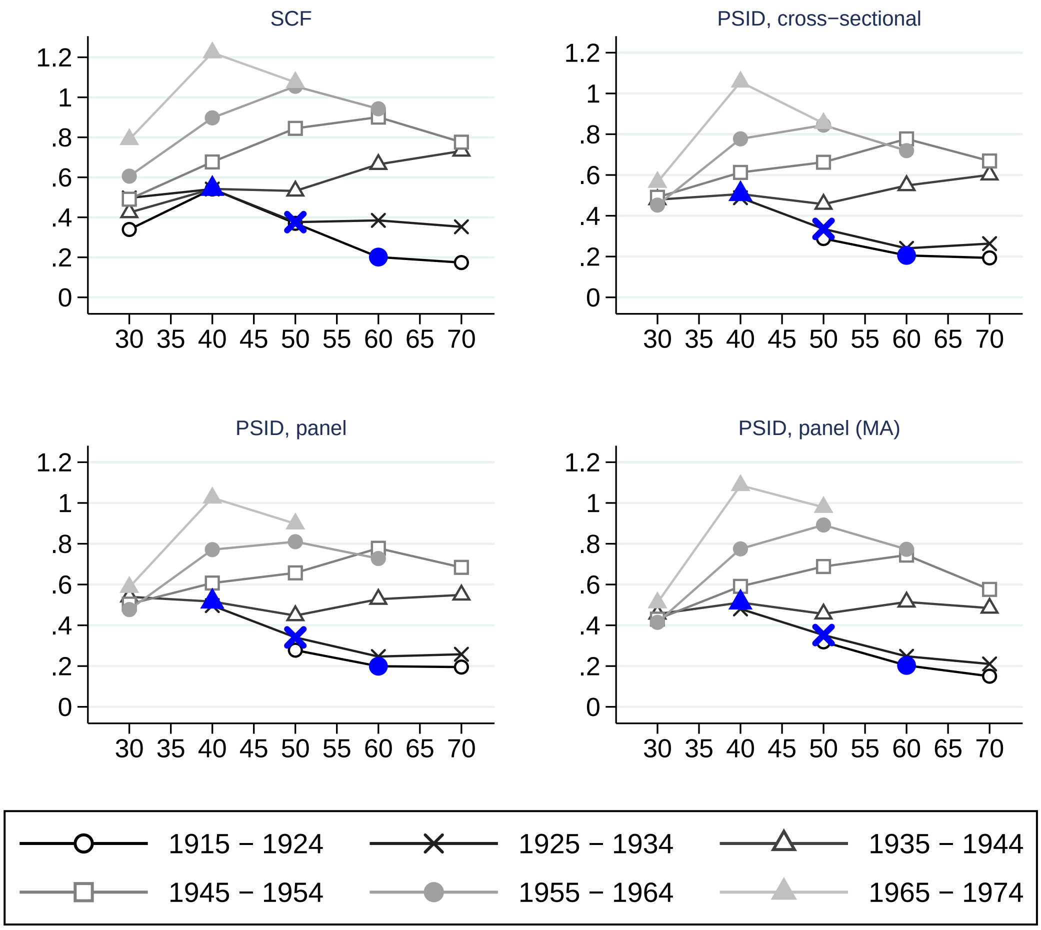

Notes: The graph shows life-cycle housing debt-to-income profiles for different cohorts. The upper left panel shows the SCF+ data, the upper right panel shows PSID data when treating the data as cross-sectional, and the lower left panel shows PSID data when exploiting the panel dimension by including household fixed effects. The lower right panel uses a three-year moving average of total household income in the denominator.

Figure 3.A.28: Buyers and extracters by age

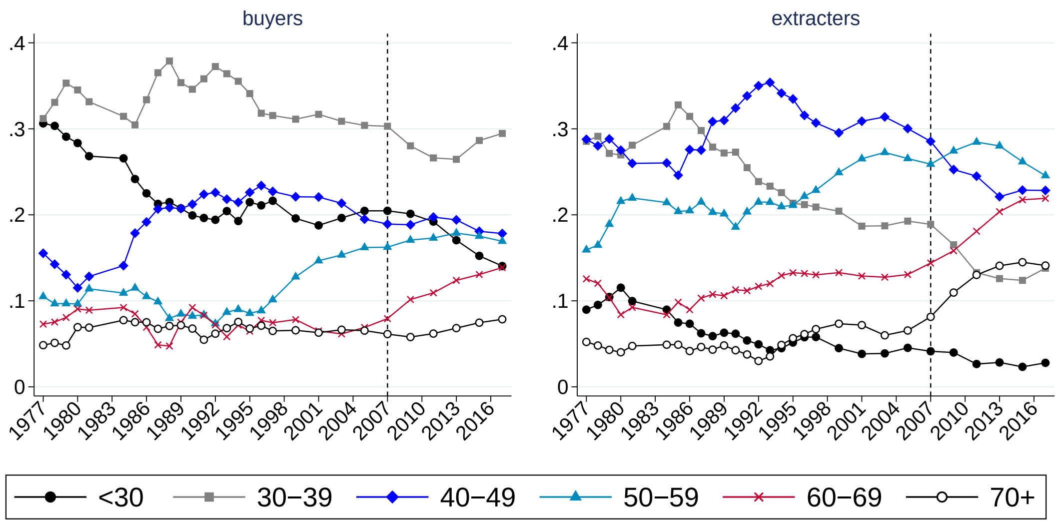

Notes: The graph shows the share of households who buy a house (left panel) and extract equity (right panel) by age over time.

4

To Have or Not to Have: Understanding Wealth Inequality

Joint work with Moritz Kuhn

4.1 Introduction

American wealth inequality is at historical highs (Saez and Zucman (2016), Kuhn et al. (forthcoming)). This fact has further spurred the interest in the old question of economic research about the drivers of wealth accumulation and the wealth distribution. A key challenge for economic theory is to account for the observed distribution of wealth (see De Nardi and Fella (2017) for a survey). The focus of existing research on the drivers of wealth accumulation is on the distribution of the consolidated value of all positions on the household balance sheet — a household's net worth. We propose a different angle to look at the wealth distribution by decomposing wealth into its components and explore the importance of the variation in access to asset classes rather than variation in the quantity of saving flows alone. The motivation for taking this new perspective is guided by the empirical evidence in this paper.

We deviate from the existing literature in two ways. First, instead of focusing on household's consolidated balance sheet, total net worth, we consider its three major components and their distribution. Using Survey of Consumer Finances (SCF) data, we document that home equity, retirement accounts, and business equity are the three main components of American household wealth accounting for 60% of total wealth and that this *core wealth* is as unequally distributed as total wealth. We explore how many households have positive asset holdings in each of these wealth component (the extensive margin) and how much these households have invested in each of these asset components (the intensive margin). We find that the extensive margin is a key determinant of wealth inequality. Second, we shift the focus from the wealth and income distribution in isolation to the joint distribution of income and wealth. This is motivated by the observation that existing theories of wealth accumulation provide at their core a mapping from the income process to wealth accumulation so that the joint distribution of income and wealth is essential to learn about the drivers of wealth accumulation. Motivated by our empirical evidence, we propose a model with frictional financial markets focusing on the three main asset classes and extensive margin variation across households. We demonstrate that this model accounts for key qualitative and quantitative characteristics of the distribution of wealth in the United States.

The main challenge in modeling the joint distribution of income and wealth is that wealth is much more unequally distributed than income (Kuhn and Ríos-Rull, 2016). Existing models of wealth inequality have introduced various mechanisms to account for the observed wealth concentration at the top such as bequest motives, entrepreneurial activity, or features of high earnings risk. The model by Castaneda et al. (2003) serves as a benchmark model in the macroeconomic literature to study wealth inequality. Castaneda et al. calibrate the income process to match the marginal distribution of wealth what results in high negative skewness of shocks for high incomes. This leads to a very strong precautionary savings motive at the top of the income distribution. Their mechanism has been adopted throughout the macroeconomic literature (see, for example (Kindermann and Krueger, 2014) or (Kaymak and Poschke, 2016)). We replicate the model and calibration to show that the model while being consistent with the empirically observed degree of income and wealth inequality generates a joint distribution that is not supported by the data. The failure to match the joint distribution scrutinizes the precautionary wealth mechanism as driver behind observed wealth inequality.

Motivated by this observation we explore the 2016 SCF data and document new facts on the joint distribution of income and wealth. We document that most variation in wealth, conditioning on household income, is explained by extensive margin variation across asset classes rather than by how much savings households have in each asset class (intensive margin). We argue that this extensive margin variation is tightly linked to households labor market situation, e.g. workers in more highly paid jobs are more likely to get access to a firm-sponsored retirement plan, workers in highly paid jobs are more likely to get access to mortgages to buy a home. We then build a model of asset accumulation that emphasizes frictions in access to assets. This also leads us to depart from the widely adopted approach to consider the consolidated household portfolio (networth) but instead we model the three major asset components home equity, retirement accounts, and business equity explicitly and the financial frictions shape access to these assets. We abstract from modeling an intensive margin choice by households in their consumption-saving decision but assume that the available financial contracts limit intensive margin variation in savings. Like wealthy hand-to-mouth consumers as in Kaplan et al. (2014), consumption dynamics in our model are typically not governed by the Euler equation. As a result, limited access to assets depending (partly) on household's labor market situation (income) is the key determinant of wealth inequality. Our model generates agents which are *savings but not borrowing* constrained. In other words, these agents would like to save part of their income by investing in retirement funds or home equity. However, they are not able to invest as they do not have access to these assets. This contrasts with standard incomplete markets models, in which agents can freely decide how much of their income they want to save but have a specific (net) borrowing limit which often depends on an agent's income. Through this mechanism our model is able to account for the non-continuous distribution of asset holdings. Furthermore, as the access to assets is intertwined with the income process, the model is also able to match heterogeneity along the extensive margin and the joint distribution of income and wealth.

We calibrate our model to the 2016 SCF data and perform a stylized policy experiment in order to analyze the impact of financial frictions on the joint distribution of wealth and income. This allows us to verify a prominent policy conjecture that access to housing is an important source of high wealth inequality. Specifically, we look at how variations in access to housing impact wealth inequality. The main findings are the following: tightening the access to home equity strongly increases the share of savings constraint households along the whole income distribution. Both aggregate wealth and aggregate welfare significantly decline. Losses are highest for the middle class, as the share of savings constraint households increases the most for this income group. Looser access leads to positive growth in aggregate wealth and welfare. Poor and middle class households gain the most from better access to home equity so that wealth inequality declines. This result supports the idea that improving the access to homeownership will reduce wealth inequality. The fact that many poor Americans cannot buy houses increases wealth inequality (oftentimes this argument is made of African American households to explain the large black-white wealth gap, see, for example, Hamilton and Famighetti (2019)).

The rest of the paper is structured as follows. Section 4.2 contains the empirical analysis of wealth inequality with the joint distribution of income and wealth. Section 4.3 provides a literature review on existing models of the joint distribution of income and wealth. In Section 4.4, we develop a stylized model of asset accumulation. The calibration of the model is described in Section 4.5. Section 4.6 presents results on the match both along targeted and non-targeted dimensions. In Section 4.7 we perform policy experiments by varying the access to mortgage financing and retirement funds. Finally, section 4.8 provides concluding remarks.

4.2 Empirical analysis

In this section we look at the empirical facts on the joint distribution of wealth and income. We use data from the 2016 Survey of Consumer Finances (SCF) for our empirical analysis. The 2016 SCF is the most recent wave of the triennial SCF. The SCF is a detailed survey on the financial situation of U.S. households. Its particular sampling scheme with an oversampling of rich households yields a representative data set even covering the financial situation of the richest households. This fact explains its popularity among researchers interested in facts on U.S. income and wealth inequality (see for example Kuhn and Ríos-Rull (2016)). The SCF provides detailed information on the portfolio composition of U.S. households. An overview on the various wealth components provided by the SCF can be found in Table 4.A.1 in the appendix. The sum of all asset classes listed in Table 4.A.1 in the appendix is defined as total wealth. Unless otherwise stated, we use the value of total annual household income provided by the SCF as income variable. Throughout the paper, we focus our analysis on working-age households, i.e. households whose head is aged between 25 and 65.

Figure 4.2.1: Wealth-to-income ratio

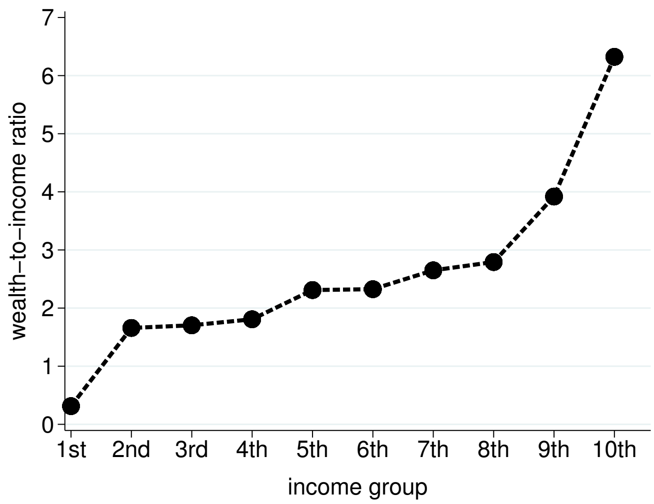

Notes: Data of the SCF 2016. Only working-age households aged between 25 and 65 are included. In order not to distort the results by outliers, the top 1% ratios in the population are excluded. These ratios are mainly due to income values near 0.

Figure 4.2.1 shows wealth-to-income ratios by income deciles.[1] If wealth and income were equally distributed across income groups, this graph would be a flat line. The fact that wealth-to-income ratios is increasing with income

However, we see in Figure 4.2.1 that wealth-to-income ratios are increasing with income: up to the first income decile, wealth corresponds on average to 24 % of annual income. In contrast, wealth of a top 10 % income household is on average more than 6 times its annual income. This tells us that wealth is significantly more concentrated at the top than income. Put another way, the higher the income a household earns, the more wealth it has been accumulated relative to its income.

In order to get a deeper understanding about the relation between wealth and income we decompose wealth in its core components and look how these are distributed across the income distribution. Core components are the three asset categories with the highest share in the portfolio of household wealth. We define the sum of these three components

[1] We pool households in income deciles and report ratios within each income decile. The first income group are households with income up to the 1st decile. The second group includes households whose income is between the 1st and 2nd decile and so forth. Finally, the 10th income group are households with income in the top 10 %.

as core wealth. To determine the core wealth components we use the SCF's fine-grained information on the household portfolio composition and calculate for each asset class the average share in total household wealth.[2] The three components with the biggest shares are business wealth (25.3 %), home equity (23.0 %), the difference between the asset value of a house and its housing debt, and retirement accounts (18.5 %).[3] In total, these three components correspond to 66.8 % of average household wealth. We define the sum of business wealth, home equity and retirement founds as core wealth. The major wealth components left out are mutual funds (10.5 %), stocks (5.8%) and liquid assets (5.3 %). The shares of each asset class differ across income groups. However, business wealth, home equity and retirement funds are the wealth components with the highest share in total household wealth across all income groups.[4]

Figure 4.2.2: Lorenz curve of total and core wealth

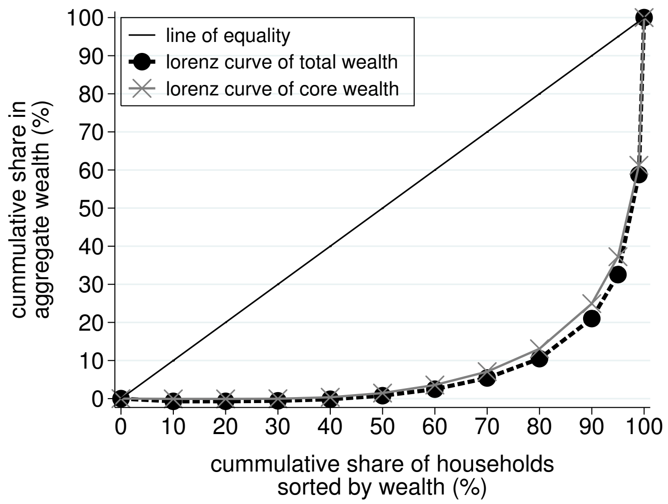

Notes: Data of the 2016 SCF. Only working age households aged between 25 and 65 are included.

Figure 4.2.2 and Table 4.2.1 compare the distribution of total and core wealth. In Figure 4.2.2 the Lorenz curves of total and core wealth are presented.[5] We see that the distribution of total and core wealth is very similar. Up to the bottom 60 % the shares in aggregate wealth are about the same. From then onwards, the Lorenz curve of core wealth is slightly above the curve of total wealth meaning that the distribution of core

[2] A list of these shares can be found in Table 4.A.1 in the appendix.

[3] Home equity includes both the primary residence and other residential real estate. Business equity is the sum of businesses where the household as an active or non-active interest. Retirement accounts include IRAs, thrift accounts and the value of future pensions.

[4] See the second to fourth column in Table 4.A.1 which show the shares of each asset class for the bottom 60 %, the 60-90% and to top 10 % of the income distribution

[5] Lorenz curves are graphs showing the degree of inequality of a distribution. To construct a Lorenz curves of wealth a population is sorted by the amount of wealth owned. Using this sorted population, the shares in aggregate wealth are calculated for various cumulative population shares. In other words, we see from the Lorenz curve how much wealth is held by the bottom x% of a wealth distribution. The line of equality you see in Figure 4.2.2 represents a Lorenz curve with aggregate wealth shares being completely equally, i.e. uniformly, distributed. In particular, the share in aggregate wealth of the bottom 10 % would be 10 % that of the bottom 20 % would be 20% and so on. The greater the area between the line of equality and a Lorenz curve, the more unequal the distribution of wealth. Dividing the area between the line of equality and a Lorenz curve by the total area under the line of equality gives us the Gini coefficients listed in Table 4.2.1.

wealth is slightly more equal than that of total wealth. The resulting Gini coefficients of total and core wealth listed in Table 4.2.1 are 0.84 and 0.79, respectively. We also report other measures of wealth inequality and the skewness of the distributions (see Kuhn and Ríos-Rull (2016)), Table 4.2.1 compares the variance of logs, the mean-to-median ratio as well as the ratio of the 9th to 5th decile of total and core wealth. The mean-to-median and the 90-50 ratio describe the skewness of the wealth distribution. In both cases, we find the distribution to be highly skewed. The mean of total wealth is 7.6 times as large as its median. Regarding core wealth, the mean is 6.2. times as large as the median. This is informative about the long right tail of the wealth distribution that increases the mean but leaves the median as the wealth level of the typical household unaffected. The long right tail of the wealth distribution also shows up in the high level of the 90-50 ratio. To enter the top 10 % of wealthiest American households, one has to have 13 times as much total and 11 time as much core wealth of the typical American households. While inequality of total wealth is slightly higher and the distribution is also more skewed, we find that overall the three major wealth components capture already very well the overall dispersion of wealth among American households. Understanding their distribution will therefore be instrumental to also understand better the distribution of total wealth.

Table 4.2.1: Inequality measures of total and core wealth

	total wealth	core wealth
Gini coefficient	0.84	0.79
variance of logs	4.59	3.47
mean-to-median ratio	7.61	6.24
90-50 ratio	12.72	11.08

Notes: Data of the 2016 SCF. Only working age households aged between 25 and 65 are included.

In Figure 4.2.3 we look at wealth-to-income ratios of total and core wealth along the income distribution. As in Figure 4.2.1, households are grouped into income deciles and wealth-to-income ratios are reported for each income decile. As the level of core wealth has to be below that of total wealth by construction we normalize the wealth-to-income ratios of total and core wealth so that there is no level-difference. In particular, wealth-to-income ratios of both total and core wealth of top 10 % income households are normalized to 100 %. The respective ratios of the other income groups are presented relative to wealth-to-income of the top 10 %. For example, looking at the 9th income group, normalized total wealth-to-income is about 70 % meaning that it is 30% below the ratio of top 10% income households. We see that the distribution of core wealth-to-income is very similar to that of total wealth-to-income: The higher the income group, the more wealth households have accumulated relative to their income. Only the ratios of the bottom 10% significantly differ.[6] In other words, not only the distributions of total

[6]The deviation of total and core wealth in the first decile is due to the fact that core wealth only includes mortgages and other residential debt. Other debt, such as installment debt, is left out. As in particular income poor households have high other debts relative to their income, total and core wealth-to-income ratios deviate in the first income decile.

and core wealth are about the same, but also the joint distribution of core wealth and income follows the same patterns to that of total wealth and income. For the remainder of the paper we focus our analysis on the joint distribution of core wealth and income. Unless otherwise stated, the words wealth and core wealth will from now on be used synonymously.

Figure 4.2.3: Wealth-to-income ratios: total and core wealth

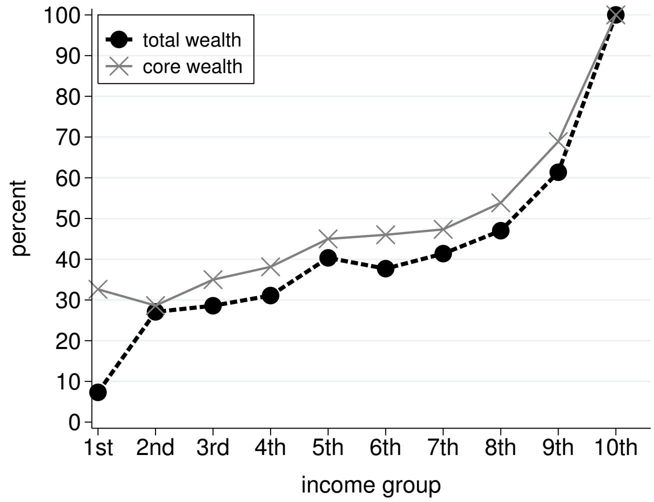

Notes: Data of the SCF 2016. Only working-age households are included. In order not to distort the results by outliers, the top 1 % ratios in the population are excluded. These ratios are mainly due to income values near 0. The ratios are normalized to the respective values of the 10th income group. The total and core wealth-to-income ratios of the top 10 % in the left graph are 6.11 and 4.15, respectively.

In order to identify the driving forces of wealth concentration at the top of the income distribution, we are now going to look at the distribution of each core wealth component separately. In Figure 4.2.4 the ratios of each core wealth component relative to income are shown. As in Figure 4.2.3 households are grouped by income deciles and the ratios are normalized to the respective values of top 10 % income households. Looking first at business wealth, up to the 7th decile business wealth is only slightly increasing. Between the 7th and 9th decile the slope becomes steeper and the greatest difference is between the 9th income group and the top 10 %. The slope of retirement funds to income ratios is approximately linear across the whole income distribution. Finally, regarding home equity, the ratio is only slightly rising up to the 9th decile. From the 9th decile to the top 10 % is a steep increase in home equity to income. While we have seen in Figure 4.2.3 that wealth-to-income ratios are increasing with income, we find here that this property is inherited by the core wealth components. We can conclude that all core components contribute to rising wealth-to-income ratios along the income distribution. The observed rising wealth-to-income ratios can either result from differences in the share of households holding assets (extensive margin variation) or from households higher up in the income distribution holding more assets (intensive margin variation). In the following, we explore how important extensive and intensive margin differences are to explain variations in wealth-to-income ratios along the income distribution.

Figure 4.2.4: Normalized ratios of core wealth components to income

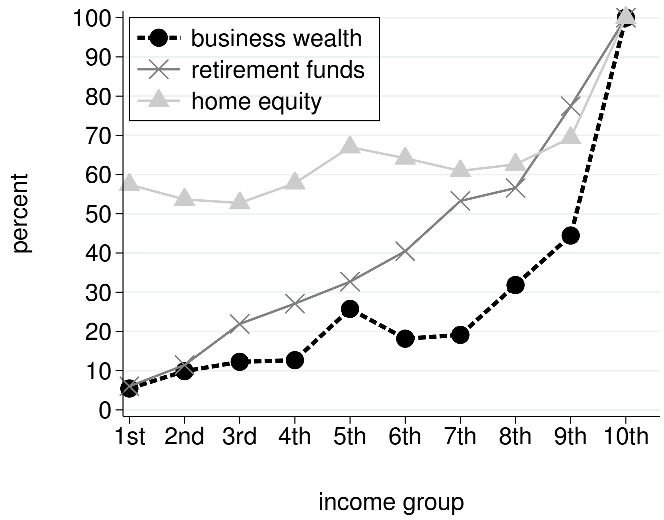

Notes: Data of the SCF 2016.Only working-age households are included. In order not to distort the results by outliers, the top 1 % ratios are excluded. The ratios are normalized to the respective values of the 10th income group. The business-to-income ratio of top 10 % income households is 0.71, their retirement-funds-to-income ratio is 1.47 and their home-equity-to-income ratio is 1.61.

We look at the distribution of ex- and intensive margins separately in order to analyze, which margin is the driving force of rising wealth-to-income ratios. In Figure 4.2.5a we see the extensive margins of each core wealth component. The margin of the top 10% income households is normalized to 100 %. If the probability of having invested in an asset was equal along the income distribution, i.e. extensive margins were constant, the lines in Figure 4.2.5a would be flat. However, we see in Figure 4.2.5a that the extensive margins of wealth components are increasing with income: the higher the income group, the more households within a group own each of the wealth components. The increase for retirement funds and home equity is approximately linear, whereas the slope of extensive margins of business wealth increases with income. This tells us that the more income a households earns the more likely it is to have invested into home equity, retirement funds or business wealth. The intensive margins of wealth components relative to income are shown in Figure 4.2.5b. Again the values of top 10 % income households are normalize to 100 %. We see that the intensive margins of business and retirement funds are slightly increasing. In contrast, the intensive margin of home equity-to-income is strongly decreasing with income up to the 8th decile. Between the 8th income group and the top 10 % the increase in the intensive margin is similar to that of retirement funds. This implies that once a household has invested in an asset, the amount relative to its income is about the same along the income distribution. Regarding home equity, this ratios is even decreasing with income. The distributions of ex- and intensive margins presented in Figure 4.2.5 highlight much more variation in the extensive than in the intensive margin along the income distribution. This finding points already to an important role of extensive margins to account for the variation in wealth-to-income ratios. Put differently, in order to explain wealth concentration at the top of the income distribution these results suggests that we have to account for how many many households have invested in wealth components across the income distribution. The impact of these variations in the extensive margin is quantified in the next paragraph.

Figure 4.2.5: Normalized ex- and intensive margins of core wealth components

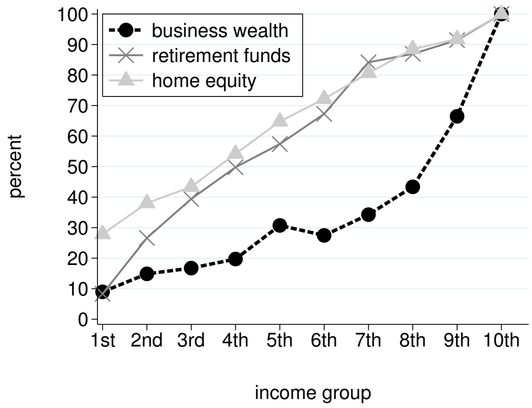

(a) Extensive margins

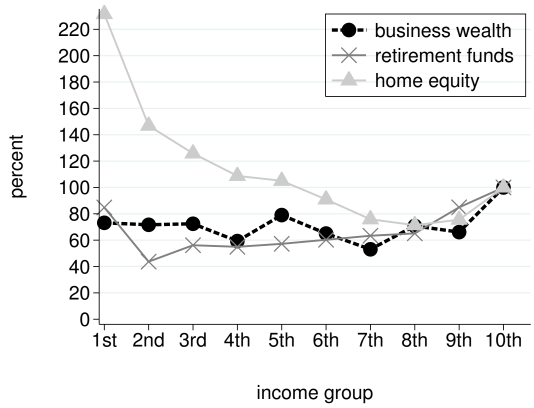

(b) Intensive margins

Notes: Data of the SCF 2016.Only working-age households are included. In order not to distort the results by outliers, the top 1% ratios are excluded. The margins are normalized to the respective values of the 10th income group. The extensive margin of business of top 10% income households is 0.35, that of retirement funds is 0.92 and that of home equity is 0.96. The intensive margin of business-to-income of top 10% income households is 2.18, that of retirement funds is 1.60 and that of home equity is 1.68.

We construct wealth-to-income ratios under two different scenarios. In the first scenario we assume the extensive margins of the three wealth components to be constant across income groups. In particular, extensive margins are set equal to that of top 10% income households. In other words, this scenario shows wealth-to-income ratios if the share of households with home equity, retirement funds or business wealth was equal to that of the top 10% across all income groups. In other words, this scenario assumes the probability of having invested in an asset to be equal along the income distribution. We refer to this as the *extensive margin effect*. The second scenario assumes intensive margins to be constant along the income distribution. Again, we fix it at the value of the 10th income decile. This is the ratio if all households having invested in an asset would hold the same amount relative to their income. This scenario is called the *intensive margin effect*.

In Figure 4.2.6 we see wealth-to-income ratios resulting from the ex- and intensive margin effect, respectively. The black dotted line shows the ratio using the original SCF data. These are identical to those shown in Figure 4.2.3. Results are striking, if we remove the variation in the extensive margin, we get that wealth-to-income ratios were decreasing up to the 7th decile while in the data ratios are increasing. In contrast, the effect of removing intensive margin variation is modest. The wealth-to-income profile becomes slightly steeper and remains strongly increasing, only slightly less convex. In summary, if extensive margins were constant and income groups would only differ with respect to

Figure 4.2.6: Ex- and intensive margin effects on wealth-to-income ratios

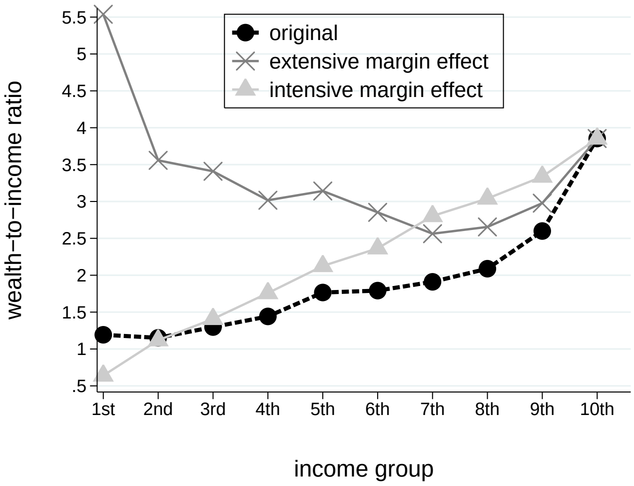

Data of the SCF 2016. Only working age households are included. Ex- and intensive margins of home equity, retirement funds and business equity are set to those of the top 10% of the income distribution.

their intensive margins, there would be a downward trend in wealth-to-income ratios. This means that the upward trend of wealth-to-income ratios in the original data would be reversed. In contrast, with constant intensive margins, i.e. if income groups would only differ with respect to their extensive margins, there would be a clear linear upward trend in wealth-to-income ratios. This is a striking result as this tells us that the increase in wealth-to-income ratios along the income distribution is mainly due to variations in extensive margins. In other words, differences in the probability to invest in an asset determine the distribution of wealth along the income distribution.

Analogously to Figure 4.2.6, Figure 4.A.1 in the appendix shows wealth-to-income ratios resulting from the ex- and intensive margin effect for different 10-year age groups. The trend of both effects is the same across all age groups. However, the older the age group, the higher are the deviations to the original data. In other words, the stronger are both the ex- and intensive margin effects.

Figure 4.2.7 decomposes the absolute change in wealth-income ratios into the ex- and intensive margin effects of each wealth component. Figure 4.2.7a shows the absolute change in wealth-to-income ratios due extensive margin effects along the income distribution. This corresponds to the difference between the dark gray line with cross and the black line in Figure 4.2.6.[7] The black bars in Figure 4.2.7a are the change in wealth-to-income if only the extensive margins of business wealth were constant across income groups. This is the extensive margin effect of business wealth. The dark gray bars represent the extensive margin effect of retirement funds and the light gray bars are the extensive margin effects of home equity. We see that the extensive margin effect is strongly decreasing with income: the wealth-to-income ratio of the bottom 10 % is more

[7]Figures 4.A.2 and 4.A.3 in the appendix show results the ex- and intensive margin effects of each wealth component by 10-year age groups.

Figure 4.2.7: Decomposition of ex- and intensive margin effects by wealth components

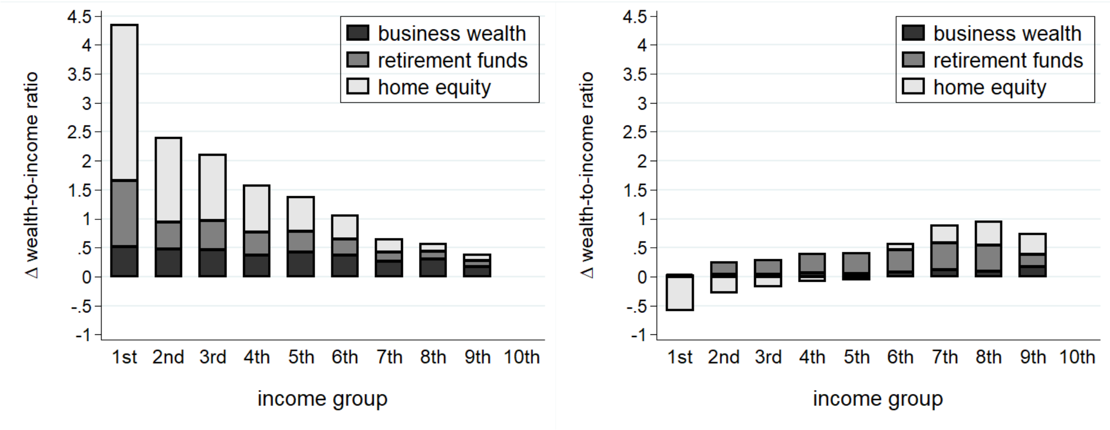

(a) Extensive margin effect (b) Intensive margin effect

Data of the SCF 2016. Only working age households are included. Ex- and intensive margins of home equity, retirement funds and business equity are set to those of the top 10% of the income distribution.

than four times that of the original data. In contrast, the ratio of households in the 9th income group only increases by about 0.4. The driving force is the extensive margin effect of home equity. The wealth-to-income ratio of the 1st income group roughly triples if extensive margin were equal to that of the top 10 % income households. The ratio of the 2nd and 3rd income groups increases by about 1.5. For the following income groups the effect decreases, being less than 0.1 for the 9th income group. Regarding retirement funds, the extensive margin has a strong effect on the 1st income groups. Their wealth-to-income ratio would rise by about 1.2. The 2nd to 5th income group have an extensive margin effect of about 0.5. As for home equity, the effect for the following income groups declines being less than 0.1 for the 9th income group. The effect of business wealth is roughly the same across income groups. It range from an increase of about 0.5 in the bottom 10 % to about 0.2 for the 9th income group. Looking now at Figure 4.2.7b, we see that the intensive margins affect wealth-to-income ratios significantly less compared to the extensive margins. This decrease of about 0.5 is caused by the intensive margin effect of home equity. The 2nd and 3rd income groups also have a slightly negative effect of home equity of about 0.2 and 0.1, respectively. Constant intensive margins of retirement funds increase wealth-to-income ratios for all income groups between the 2nd and the 9th decile. The effect ranges from 0.2 to about 0.5. Finally, the intensive margin effect of business wealth is negligibly low. The 9th income group has a slight increase of about 0.1. To sum up, the key takeaways are the following: first, setting the extensive margins to that of the top 10 % along the whole income distribution has a much stronger effect than constant intensive margins. Second, the lower the income group, the higher the positive extensive margin effect. Third, the driving force of this is home equity followed by retirement funds.

So far we have analyzed how much wealth relative to income is accumulated and how this is affected by the ex-and intensive margins of wealth components. Next, we look

Figure 4.2.8: Joint Lorenz curve of wealth by income: ex- and intensive margin effects

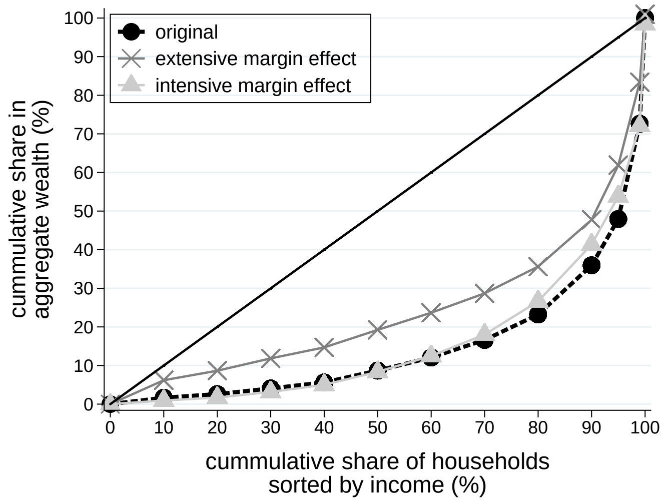

Notes: Data of the SCF 2016. Only working-age households are included (head aged 25-65). The scenarios referring to constant ex- and intensive margins assume margins being equal to the respective value of the 10th income decile across income groups.

at the ex- and intensive margins effects on the distribution of aggregate wealth across income groups. In order to do so, we construct a Lorenz curve of wealth where we sort households by income but look at their share in wealth. Combining the income ranking with the wealth share provides a measure of the joint distribution of income and wealth. The Gini coefficient resulting from this Lorenz curve is defined as the joint Gini coefficient of wealth by income.

In Figure 4.2.8 we see the joint Lorenz curve of wealth by income using the original 2016 SCF data (black dotted line) as well as the curves with the ex- and intensive margin effects (dark and light gray lines). The line of the intensive margin lies almost exactly on the black line. Only in the upper part it lies slightly above it. In contrast, the joint Lorenz curve of the extensive margin is significantly more closely to the line of equality. This underpins the fact, that the extensive margins are key in explaining the joint distribution of income and wealth.

The joint Gini coefficients resulting from the joint Lorenz curves of Figure 4.2.8 are presented in the right column of Table 4.2.2. Note that if wealth were distributed as unequally as incomes, then the Gini coefficient would be identical to the Gini coefficient of income. The difference between the Gini coefficient for the joint Lorenz curve and the Gini coefficient for income is therefore informative how much more unequally wealth is distributed compared to income. In the original SCF data, the Gini of income is 0.59 and that of wealth is 0.68 meaning that the Gini of income is 13 % below the Gini of wealth. With constant intensive margins, the joint Gini of wealth by income slightly decreases from 0.68 to 0.66. In contrast, the extensive margin effect reduces the joint Gini by about 20 % to 0.54. Thus, if extensive margins were constant along the income distribution, wealth would be more equally distributed than income.

Table 4.2.2: Gini coefficient with constant ex- and intensive margins

		Share in aggregate wealth (in %)			Gini coeff. of wealth by income
Income group		< 50%	50-90%	> 10%	
original data		8.7	27.3	64.0	0.68
extensive margin effect	home equity	15.2	26.6	58.2	0.60
	retirement funds	11.9	27.5	60.6	0.64
	business equity	10.8	28.8	60.5	0.65
	all	19.4	28.5	53.2	0.54
intensive margin effect	home equity	7.2	30.1	62.7	0.69
	retirement funds	9.8	30.2	60.1	0.65
	business equity	8.6	31.7	59.7	0.66
	all	8.5	33.6	57.9	0.66

Notes: Data of the SCF 2016. Only working-age households are included (head aged 25-65). The scenarios referring to constant ex- and intensive margins assume margins being equal to the respective value of the 10th income decile across income groups.

The other columns in Table 4.2.2 are shares in aggregate wealth for the bottom 50 %, the 50-90 % and the top 10 % of the income distribution respectively. From now on, we define these three income groups as poor, middle class and rich households, respectively. We see that with extensive margins being constant, poor households, the bottom 50 %, significantly gain whereas the rich, the top 10%, loose shares: the share in aggregate wealth of poor households increases by more than 10 pp whereas that of the rich decreases by about the same amount. The share of the middle class, the 50-90 %, remains about the same. In contrast, constant intensive margins only marginally increase wealth shares of both the poor and middle class households.

Hence, the stylized fact that wealth is much more unequally distributed than income hinges less on the fact that the income poor have fewer assets but rather that they do not have these assets at all. This suggests that understanding why some households do not invest in certain asset classes is key to understand the overall distribution of wealth. After discussing existing models of wealth accumulation, we develop in Section 4.4 a model of wealth accumulation focusing on extensive margin variation of asset investment.

4.3 Existing Models of Wealth Inequality

As described in the previous section, wealth is distributed much more unequally than income. Thus, models which aim at matching wealth inequality to the data have to incorporate a mechanism generating a wealth variance significantly higher compared to that of income. In other words, income rich households must have a motive to accumulate disproportionate high amounts of wealth.

Cagetti and De Nardi (2008) as well as De Nardi and Fella (2017) provide a review on models aiming at explaining wealth inequality. By introducing various saving mechanisms these models generate wealth concentration at the top. The main mechanisms described by Cagetti and De Nardi (2008) are the following: first, including entrepreneurship into a model (see for example Quadrini (2000)), in which agents can either become an entrepreneur or remain a worker. The saving behavior of entrepreneurs differs as they do not receive labor income but revenues from their business wealth. A second approach is to incorporate different preference types (see for example Krusell and Smith (1998)). Agents with higher discount factors are more patient so that they have a higher savings rate compared to impatient agents. Another way how of obtaining a saving motive is to model wealth as a luxury good form which agents receive utility (see for example Diaz et al. (2003)). The models described so far are so called dynasty models which assume infinitely lived agents and do not incorporate life cycle patterns. In contrast, a second class of models uses overlapping generations models as a basis. De Nardi (2004) generate wealth inequality by introducing bequest motives so that some new born agents start their life-cycle with wealth accumulated by their ancestors. Other mechanisms are the presence of medical expense risk and heterogeneity in life expectancy as in the model of De Nardi et al. (2010). Finally, there are models combining elements of both dynasty and life-cycle models. The work of Castaneda et al. (2003) belongs to this kind of category. Wealth inequality is obtained by modeling earnings risk which induces a high precautionary savings motive for rich agents. This model serves as a benchmark model in the macroeconomic literature to study wealth inequality. In the following we describe this model and its mechanisms to generate a wealth distribution which matches the data in more detail.

Castaneda et al. (2003) use a standard incomplete markets model with stochastic aging and stylized features of the social security and tax system. Their key innovation is to back out the earnings process from the observed distribution of earnings and wealth. They calibrate the earnings process to match a set of points on the Lorenz curves for earnings and wealth. This provides the model a very close fit to the observed *marginal* distribution of earnings and wealth. To generate the high concentration of wealth at the top of the distribution, the calibration requires one very high but transitory income state in which households accumulate large precautionary savings because this high income is only transitory. The model of Castaneda et al. (2003) matches the observed earnings and wealth inequality very well. While the model matches the *marginal* distribution well, its economic mechanism of a strong precautionary savings motive makes also predictions for the *joint* distribution of income and wealth. These predictions provide a way to assess if the data are consistent with the underlying economic mechanism to generate the wealth distribution. Large discrepancies in the *joint* distribution would indicate that the respective saving mechanism is likely to be not correctly specified.

In order to explore the model's prediction for the *joint* distribution of earnings and

Figure 4.3.1: Lorenz curves of earnings and complete wealth

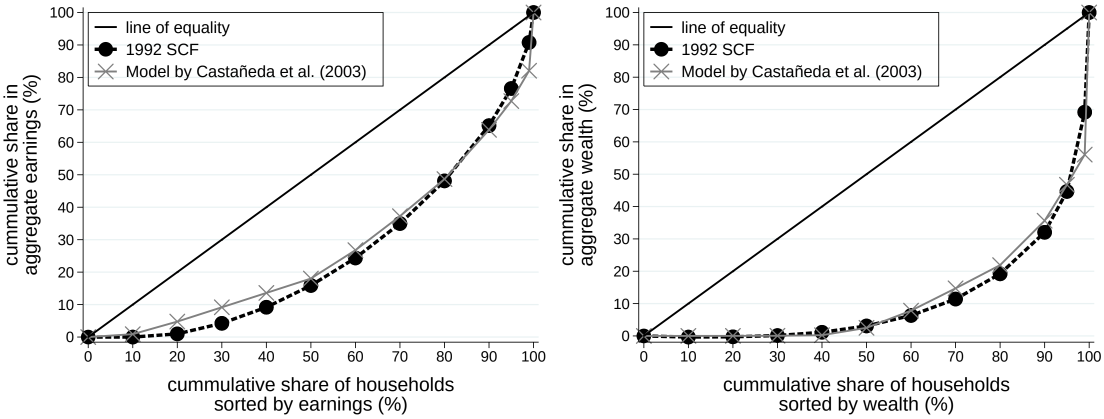

(a) Lorenz curve of earnings

(b) Lorenz curve of complete wealth

Notes: Data of the 1992 SCF. Only working-age households are included.

Table 4.3.1: Distribution of earnings and complete wealth

	Share in aggregate earnings (in %)						Gini coeff. of earnings
Earnings group	< 20%	20-40%	40-60%	60-80%	80-90%	> 10%	
SCF 1992	1.0	8.2	15.2	23.8	17.0	34.9	0.44
Castañeda et al. (2003)	4.8	8.8	13.2	21.9	15.2	36.2	0.45

(a)

	Share in aggregate wealth (in %)						Gini coeff. of wealth
Wealth group	< 20%	20-40%	40-60%	60-80%	80-90%	> 10%	
1992 SCF	-0.3	1.4	5.2	12.9	12.8	68.0	0.76
Castañeda et al. (2003)	0.0	0.2	7.6	14.1	13.7	64.4	0.78

(b)

Notes: Data of the 1992 SCF. Only working-age households are included.

wealth we replicate the benchmark model by Castaneda et al. (2003). Castaneda et al. (2003) use data of the 1992 SCF and target their model to the distributions of earnings and complete wealth. We therefore rely on 1992 SCF data for this part of the paper. Figure 4.3.1 shows the Lorenz curves of earnings and wealth. In the left graph of Figure 4.3.1 we see the Lorenz curve of earnings (black dotted line) and the curve generated by the model of Castaneda et al. (2003) (gray line with crosses). The right graph is the corresponding Lorenz curve of total wealth. The Lorenz curves generated by the model and those of the SCF are roughly on top of each other so that also the Gini coefficients listed in Table 4.3.1 only slightly differ. This shows us that the model of Castaneda et al. (2003) matches the distribution of both earnings and wealth extremely well. In particular, their model is capable of matching the high concentration of wealth at the top. This is a feature other existing models struggled with.

Figure 4.3.2: Joint Lorenz curves of earnings and complete wealth

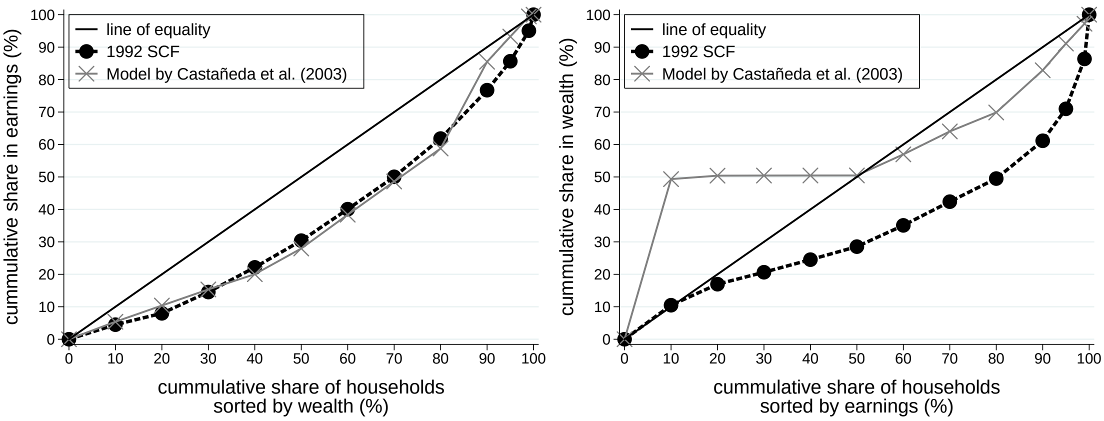

(a) Joint Lorenz curve of earnings by wealth (b) Joint Lorenz curve of wealth by earnings

Notes: Data of the 1992 SCF. Only working-age households are included.

Figure 4.3.2 and Table 4.3.2 are analogous to Figure 4.3.1 and Table 4.3.1, but now focus on the untargeted *joint* distribution of earnings and wealth. In the left graph we see the joint Lorenz curve of earnings with households being sorted by complete wealth. The corresponding Gini coefficients of the observed data and the model are listed in Table 4.3.2. The model fits to the observed data still quite well. The model now struggles to account for the income share of wealth-rich households at the upper end of the wealth distribution. Deviations to the data are bigger than for the *marginal* earnings and wealth distributions in Figure 4.3.1 and Table 4.3.1. The share in aggregate earnings of households with wealth between the 8th and 9th decile is about 10 percentage points higher compared to the data. The shares of top 10 % wealth households are about 10 percentage point too low. The right graph in Figure 4.3.2 shows the joint Lorenz curve of wealth by earnings.

The deviation now is stark. The strong divergence between model and data provides strong evidence against the underlying economic mechanism of a strong precautionary savings motive. In particular the share in aggregate wealth of households in the bottom of the earnings distribution is much too high. In the 1992 SCF about 10 % of aggregate wealth is owned by households in the bottom 10 % of the earnings distribution. In contrast, the model generates a share being about 40 percentage points higher.[8] The high wealth share at the bottom leads to a joint Lorenz curve lying above the line of

[8] In addition to targeting the model of Castaneda et al. (2003) to the marginal distributions of earnings and wealth we also performed calibrations using the joint distribution of earnigns and wealth as targets. However, as mentioned in Section 4.3, the model mechanism of having very wealth-rich households falling back into the lowest earnings state always results in much too high aggregate wealth shares in the bottom of the earnings distribution.

Table 4.3.2: Joint Distribution of earnings and complete wealth

	Share in aggregate earnings (in %)						Joint Gini of earnings by wealth
Wealth group	< 20%	20-40%	40-60%	60-80%	80-90%	> 10%	
SCF 1992	8.0	14.2	17.9	21.8	14.9	23.3	0.29
Castañeda et al. (2003)	10.4	9.7	18.3	20.4	26.7	14.5	0.31

(a)

	Share in aggregate wealth (in %)						Joint Gini of wealth by earnings
Earnings group	< 20%	20-40%	40-60%	60-80%	80-90%	> 10%	
1992 SCF	17.0	7.6	10.6	14.4	11.6	38.9	0.42
Castañeda et al. (2003)	50.4	0.1	6.5	12.9	13.0	17.1	0.65

(b)

Notes: Data of the SCF 1992. Only working-age households are included.

equality so that the resulting Gini coefficient cannot be reasonably interpreted. The poor model fit is a direct consequence of the earnings process in combination with the strong precautionary savings motive in the model. More precisely, the estimated earnings process results in earnings of the rich to be more than three thousand times higher than those of the poorest households. In addition, the probability of the rich to become a poor household is relatively high, about 10% every period, so that they save a large fraction of their earnings and accumulate high levels of wealth over time. Due to these components the model generates a higher variance of wealth compared to earnings and is able to match wealth and earnings inequality to the data. However, if a earnings rich household drops down to the lowest earnings state it has accumulated so much wealth, that it is not worth for it to work anymore as it can just use up its wealth. Hence, such a wealth rich household would end up at the bottom of the earnings distribution. Due to this mechanism wealth shares of earnings poor households are highly overestimated by the model compared to the data.

To conclude, the model of Castaneda et al. (2003) is able to match the high wealth concentration at the top by incorporating productivity shocks, which generates a strong precautionary savings motive for earnings rich households. However, the model does not match to the joint distribution of earnings and wealth. This provides strong evidence against this economic mechanism. The findings of Section 4.2 indicate that extensive margins of asset components are essential in explaining the joint distribution of income and wealth. In the next section we propose a stylized model with an extensive margin in wealth accumulation generated by differences in labor market experience.

4.4 Financial Frictions Model

This section proposes a partial equilibrium model of wealth accumulation. Financial frictions in the model govern the consumption-saving decision. Access to financial products is frictional and financial products are pre-determined so that savings flows cannot be adjusted easily. This implies that households are typically not on their Euler equation resembling wealthy hand-to-mouth agents. Wealth inequality does not arise from a savings motive, but from limited access to assets depending (partly) on the income of a household. This is a key deviation from existing models. We refer to the model framework in the following as the Financial Frictions Model.

Agents in the model can invest in housing, retirement accounts, and businesses. The key assumption is that investment opportunities in these assets arise randomly. In addition, they depend in the case of houses and retirement accounts on the labor market situation of households. The income dependence stems from banks only offering credit based on an agent's labor market situation and employers do not offer retirement plans to all of their employees. We do not model this process in detail but in a reduced form as random investment opportunities. In particular, the higher the income of an agent, the more likely the access to the frictional asset market. This allows us to account for the empirical fact, that the share of households holding assets such as housing or retirement funds is increasing with income. Agents can only invest in illiquid assets in our models, which can only be consumed at the end of the life-cycle. This effectively imposes hand-to-mouth behavior for all agents (see Kaplan et al. (2014)). Some of the agents will be poor hand-to-mouth agents but the majority of agents in the model will be like wealthy hand-to-mouth agents. What differs in our model is that wealthy hand-to-mouth agents will have positive savings rates in the financial products they hold. For example, homeowners will always repay their mortgage but have marginal propensities to consume like wealthy hand-to-mouth agents. This ingredient of our model induces households to accumulate wealth over their cycle.

An agents' investment decision involves two steps: first, as a prerequisite for an investment, an agent must get access to the frictional asset market. Second, in case of access the agent must decide whether to invest at all or, if so, how much she wants to invest. For this it is evaluated whether its worth to reduce consumption during working life in order to be able to consume the investment during retirement. In a nutshell, the key element of our model is the question of "to have or not to have" invested in an asset. In the following we describe our model in detail.

There is a continuum of finitely-lived agents. Each agent lives for J periods when she is active in the labor market and retires at age $J + 1$. The retirement period is of stochastic length $\frac{1}{\rho}$. At the end, the agent dies, receives a continuation utility of zero, and will be replaced by a newborn agent. We denote the age of the agent by j. Agents receive either labor income or income from being a business owner. During their time in the labor market, agents receive earnings l which is idiosyncratic and drawn from a distribution with a support on a finite number of elements $L = \{l_1, l_2, \ldots, l_s, \ldots, l_S\}$. The unconditional probability of a draw l_s is denoted by $\pi^l(l_s)$. Instead of working as an employee an agent can decide to become self-employed and to run a business. We assume that a prerequisite to run a business is what we call a *business idea*. Business ideas can be of different quality, and therefore, result in different business wealth. Business wealth is, in line with the empirical analysis, the equity of the firm. Our way of modeling corresponds to a situation where business equity is intangible capital. The fact that agents do not have to put in capital to start a firm can be considered as the outcome of well-functioning venture capital markets where agents with an idea receive financing to start the idea. In line with the data, we only consider business equity and consolidate all assets and debt of the business. We model a business opportunity as the seed of a Lucas tree. If the agent decides to take the business opportunity, she plants the tree and it provides dividends ("fruits") in every period. Note that such a tree will immediately have value as it can be sold at the present value of dividends. We denote the value of the business by b. The business wealth is the present value of the sequence of dividends:

$$b = \sum_{t=0}^{\infty} \frac{d_t}{(1+r_b)^t} = \frac{d}{r_b} \tag{4.4.1}$$

It is assumed that there is finite number of business ideas $B = \{b_1, b_2, \ldots, b_t, \ldots, b_T\}$. We denote the unconditional probability of receiving a business idea that creates value, b_t, by $\pi^b(b_t)$.

In each period, each agent receives either a job offer or a business opportunity. Thus, a worker has never business wealth and business owners do not receive labor income. Put another way, being in state $l \in L$ implies $b = 0$ and having a business idea $b \in B$ results in labor income $l = 0$. In addition, the probability of receiving any job offer plus the probability of receiving any business idea has to be equal to one:

$$\sum_{s=1}^{S} \pi^l(l_s) + \sum_{t=1}^{T} \pi^b(b_t) = 1 \tag{4.4.2}$$

Workers and self-employed have the opportunity to set up retirement funds that have a tax advantage. We denote the balance of assets in the retirement fund by p and use ψ as an indicator if an agent has a retirement fund ($\psi = 1$) or not ($\psi = 0$). We assume that access to an employer-sponsored retirement fund depends on the income state of the agent. When an agent receives a new job opportunity l_s there is a probability that this job comes with a retirement plan. We denote this probability by $\pi_p^l(\psi|l_s)$. We also allow self-employed to run retirement plans for themselves. When starting a new business, the agent must also decide if she sets up a retirement plan. The probability that she gets access to set up a retirement plan is $\pi_p^b(\psi|b_t)$ and depends on the current business opportunity b_t. We assume that all working agents take retirement plans if offered to them and for self-employed, we assume that they only have the opportunity to set up a retirement plan when they start the business. Retirement plans are pre-specified financial products. We denote by χ^l the share of labor income l a worker pays into the plan each period. The employer matches employee contributions with a share σ of labor income l. For self-employed, χ^b denotes the share of dividend payments d that the business owner has to pay into the retirement fund each period. The payments into the retirement funds of both the workers and the business owners, $\chi^l l$ and $\chi^b d$ respectively, are assumed to be tax-exempt. In addition, balances of retirement plans are perfectly portable in the market. The law of motion for the retirement plan is:

$$p' = p(1 + r_p) + \psi'((1 + \sigma)\chi^l l' + \chi^b d') \tag{4.4.3}$$

Each period an agent gets a return r_p on the current retirement fund balance. Tomorrow's balance of an agent's retirement fund depends on whether or not she has a retirement plan tomorrow.[9] If an agent does not have a retirement plan tomorrow ($\psi' = 0$) tomorrow's balance of an agent's retirement plan is equal to the current balance p plus the return $p \cdot r_p$ on it. Agents deciding to stay or become a worker with a retirement plan tomorrow ($\psi' = 1$ and $l' \neq 0$) pay $\chi^l l'$ into their fund and their employer further adds $\sigma\chi^l l'$ into it. Agents who are going to stay or become a business owner with a retirement plan ($\psi' = 1$ and $b' \neq 0$) pay $\chi^b d'$ into the fund.

In addition to the accumulation of assets through business wealth and retirement funds, agents encounter the opportunity to buy a house. We denote the value of the house by h and assume a finite number of house types $H = \{h_1, h_2, \ldots, h_e, \ldots, h_E\}$. A house with value h provides utility φh (housing services). To buy a house, an agent must take out a mortgage. We abstract from any financing frictions and assume that all houses are bought without equity and finances by a mortgage m with loan-to-value ratio $\frac{m}{h} = 1$. As retirement plans, mortgages come as pre-specified credit contracts. The contracts have the interest r_m and a required fixed repayment schedule κh each period. We allow for refinancing to buy new housing h' so the law of motion for the mortgage is:

$$m' = \begin{cases} h' - h & \text{if } m = 0 \\ m(1 + r_m) - \kappa h + h' - h & \text{if } m > 0 \end{cases} \tag{4.4.4}$$

The first row of Equation 4.4.4 is tomorrow's mortgage balance of agents without any debt today. In case an agent owns a house of value h which is fully paid off, i.e. $m = 0$, and buys a new house h' tomorrow, the mortgage balance is equal to the difference between the two housing values $h' - h$.[10] If no new house is bought, i.e. $h' = h$, the agent remains out of mortgage debt. Agents buying a house for the first time ($h = m = 0$) which is of value h' have mortgage debt of the same amount ($\frac{m'}{h'} = 1$). In the second row of equation 4.4.4 we see the law of motion regarding home owners with mortgage debt m. Agents have to pay interest r_m on this debt so that tomorrows mortgage balance increases by mr_m. Debt is reduced by κh, the amount that is repaid by agents. As mentioned above, agents are allowed to refinance their current debt in order to buy a new house h'. The resulting new mortgage debt is the current balance m plus interest mr_m minus repayments κh as well as the difference of h' to today's housing value, $h' - h$.

[9] For example, an agent might get a job offer without a retirement plan but with a much higher labor income than today. In this case it might be optimal to switch the job despite loosing the retirement plan

[10] It is assumed that agent remain homeowners once they bought a house and they do not trade their house for one worth less.

Hence, by repaying their mortgage agents build up home equity over time.

We document empirically that the extensive margin of housing is tightly linked to income. The higher the income of a household, the more likely it is to own a house. We assume that search for housing is frictional and that an agent finds a house of size h_e with probability $\lambda(h_e)$. We interpret this evidence as the outcome of a financing friction that we model in reduced form. We assume that the probability to obtain financing for a house depends on the current job l or business idea b, respectively. The higher the agent's income, the more likely financing for a house is available. We denote the probabilities of being offered a finance plan by $\theta^l(l_s)$ for workers with job l_s and by $\theta^b(b_t)$ for business owners with business b_t. These probabilities are assumed to be increasing with respect to labor and business income, respectively. In order to pin down these probabilities we are going use the extensive margins of housing along the income distribution (see Section 4.5).

The model components described above enable agents to accumulate wealth by investing in three different asset classes: business wealth b, retirement funds p and home equity $h - m$. Thus, total wealth w of the agent is given by:

$$w = h + b + p - m \tag{4.4.5}$$

This wealth definition is equivalent to core wealth described in the empirical part in Section 4.2.

We assume that the agent's period felicity function has the following separable additive form of utility from consumption and housing:

$$u(c, h) = \alpha log(c) + (1 - \alpha)log(\underline{h} + \varphi h) \tag{4.4.6}$$

where $\underline{h}$ denotes some minimum level of housing that each agent receives and that we introduce to abstract from a separate renting market for housing.

Agents have to pay taxes. We assume that mortgage payments and contributions to retirement plans are tax deductible so that taxable income $\tilde{y}$ is given by

$$\tilde{y} = l + d - \chi^l l - \chi^b d - \kappa h \tag{4.4.7}$$

Agents are hand-to-mouth consumers of after tax disposable income:

$$c = \bar{y} + (1 - \tau) \cdot \tilde{y} - \chi^l l - \chi^b d - \kappa h \tag{4.4.8}$$

with $\bar{y}$ being non-taxable income. In other words, agents consumption is equal to their after tax disposable income minus the sum of payments into retirement funds and repayments on their mortgage. As mentioned above, this is a key element of our model. In contrast to a standard consumption-saving decision agents consumption in a given period does not depend on their total wealth, but only on their current income. Wealth that is accumulated by investing in retirement funds and housing can only be consumed after working life, i.e. in period $J + 1$. Through this mechanism agents can be what has been called wealthy hand-to-mouth consumers (see Kaplan et al. (2014)). These agents have accumulated high amounts of illiquid wealth they cannot use for consumption during working life but only when they retire.

We present the dynamic decision problem directly in recursive form. We split each period in two stages. The first stage is the labor market stage when job offers and business opportunities are received and decided upon. In the second stage, agents look for housing and get opportunities to buy housing if they receive financing. The state variables of an agent are (j, l, b, ψ, p, h, m) where j captures age, l the current job state (0 if business owner), b denotes business equity (0 if no business), ψ is an indicator whether the agent currently pays into a retirement account, h the current housing stock (0 if no housing) and m is the mortgage balance.

At the first stage, labor market decisions have to be made. Households receive job offers and get business opportunities. They face a discrete choice problem of choosing the alternative that yields the highest continuation value. At the end of the period at stage 2, households get the opportunity to find and buy housing. The housing decision is taken conditional on the labor market decision from stage 1 of the period.

The value function of an agent with state vector (j, l, b, ψ, p, h, m) at the first stage is

$$V(j, l, b, \psi, p, h, m) = u(c, s) + \sum_{l' \in L} \pi^l(l') \sum_{\psi_l' \in 0,1} \pi_p^l(\psi_l'|l') \sum_{b' \in B} \pi^b(b') \sum_{\psi_b' \in 0,1} \pi_p^b(\psi_b'|b') \tilde{V}_{max} \tag{4.4.9}$$

$\tilde{V}_{max}(j, l', b', \psi', p', h, m)$ is the value that maximizes the continuation value at stage two of the agent's decision problem.

Given $\tilde{V}_{max}$, the expected value at the beginning of stage one depends on the probability distributions of labor and business states, $\pi^l(l')$ and $\pi^b(b')$, as well as on the probability to receive a pension plan conditional on the agent's labor market situation, $\pi_p^l(\psi_l'|l')$ and $\pi_p^b(\psi_b'|b')$. Having received a job l' or business offer b' at stage one, the labor market decisions is made by the agent. In case the offer includes a pension plan, i.e. ψ_l' or ψ_b' is equal to one, the agent also has to decide whether to invest in this plan. At the end of the first stage the state variables l', b', ψ' are set. The state of pension wealth p' is endogenously determined according to equation 4.4.3.

At the beginning of stage two the state vector of the household is $(j, l', b', \psi', p', h, m)$ where primes denote that these variables might have changed depending on the household's decision at stage one. The value function of an agent at the beginning of the second stage is

$$\tilde{V}(j, l', b', \psi', p', h, m) = \beta \sum_{h' \in H} \lambda(h') \theta(l', b') V(j+1, l', b', \psi', p', h', m') \tag{4.4.10}$$

Given the labor market and pension plan decision of stage 1, the expected value at the beginning of the second stage depends on the probability distribution of finding a house, $\lambda(h')$, and the probability of being offered a finance plan, $\theta(l', b')$. The latter is conditional on the agent's labor market decision of stage 1.

At stage two the agent now may find a house and receives a finance plan. In case the agent is offered a house, h', in conjunction with a finance plan, she has to decide whether or not to invest in this house. If now house if found or the agent does not get a finance plan, there is no decision to be made so that $h' = h$. At the end of stage 2 the state h' is set. The state of mortgages m' is endogenously determined according to equation 4.4.4.

At the end of working life, households enter into the retirement phase. We assume that their continuation utility in that period only depends on their total accumulated wealth so that $V(J,w) = \delta log(w)$.

4.5 Calibration

In this section we define the specific functional forms of the model as well the empirical target statistics used for the calibration.

Labor income and business wealth: As described in Section 4.4, there is a finite number of job offers $L = \{l_1, l_2, \ldots, l_s, \ldots, l_S\}$ and business ideas $B = \{b_1, b_2, \ldots, b_t, \ldots, b_T\}$ with unconditional probabilities

$$\Pi^l = \{\pi^l(l_1), \pi^l(l_2), \ldots, \pi(l_s), \ldots, \pi(l_S)\} \text{ and}$$

$$\Pi^b = \{\pi^b(b_1), \pi^b(b_2), \ldots, \pi(b_t), \ldots, \pi(b_T)\}, \text{ respectively.}$$

We neither impose any specific functional form on the distribution of job and business states nor on their probability distributions. We only impose the states l_s and b_t to be increasing with s and t, respectively. Besides, we impose the probabilities $\pi^l(l_s)$ and $\pi^b(b_t)$ to be decreasing with s and t, respectively. In other words, the better the job offer or business idea, the lower the probability of receiving it. The older an agents, the more likely she is to have received a high paid job offer or a high quality business idea during her course of life. Hence, both labor income as well as the extensive margin of business owners is expected to be increasing with age which fits to the data. This mechanism is equivalent to assuming that agents get promotion over time or increase their quality of their business idea. We normalize the lowest job state l_1 to 1 and its probability is $\pi^l(l_1) = 1 - \sum\limits_{s=2}^{S} \pi^l(l_s) + \sum\limits_{t=1}^{T} \pi^b(b_t)$.[11]

According to equation (4.4.1) in Section 4.4 business income depends on the interest rate r_b. As our model is a partial equilibrium model, we treat this interest as an additional parameter. In total, there are $2(S-1+T)+1$ parameters to estimate the distribution of labor income and business wealth.

In order to solve the model numerically we set the number of job and business opportunities to $S = 5$ and $T = 4$, respectively.[12] Hence, the distributions of jo offers and business ideas are determined by the parameters $\{l_2, l_3, l_4, l_5\}$ and $\{b_1, b_2, b_3, b_4\}$, respectively. The probability distribution of receiving either a job offer or a business opportunity requires 8 additional parameters, $\{\pi^l(l_2), \pi^l(l_3), \pi^l(l_4), \pi^l(l_5)\}$ and $\{\pi^b(b_1), \pi^b(b_2), \pi^l(b_3), \pi^l(b_4)\}$. For the estimation of these 16 parameters we use empirical targets that account for the

[11] As mentioned in Section 4.4 $\sum\limits_{s=1}^{S} \pi^l(l_s) + \sum\limits_{t=1}^{T} \pi^b(b_t) = 1$

[12] As the number of possible states in the model increases overproportionally high with the number of job, business, pension or housing states we choose relatively small numbers for these states.

life cycle patterns of workers and business owners as well as for the distribution of income and business wealth.

Figure 4.5.1 shows life cycle profiles of earnings and business owners using 2016 SCF data and grouping households into 10-year age groups. In Figure 4.5.1a we see mean earnings by 10-year age groups relative to mean earnings of households aged 55-65. Earnings are increasing with age: average earnings of households aged 25-35 correspond to 60 % of those of the 55-65 year olds. Households aged between 35-45 earn about 20 % less than households aged 55-65 and the 45-55 year olds about 10 % less. In order to match the life cycle patterns of earnings we target the Financial Frictions Model to the values presented in Figure 4.5.1a. The share of business owners is shown in Figure 4.5.1b. We see that these shares are increasing up to the age group of the 45-55 year olds: 6 % of households aged between 25 and 35 are business owners and 18 % of the 45-55 year olds. The share of business owners among households aged 55-65 is slightly less, about 17 %. We target our model to these values to match the life cycle pattern of business owners.

The distribution of income is estimated by targeting different parts on the Lorenz curve of income. Regarding business wealth, we both match the extensive margin of business owners as well as the share in aggregate business along the income distribution.

Figure 4.5.1: Life cycle profiles of earnings and the share of business owners

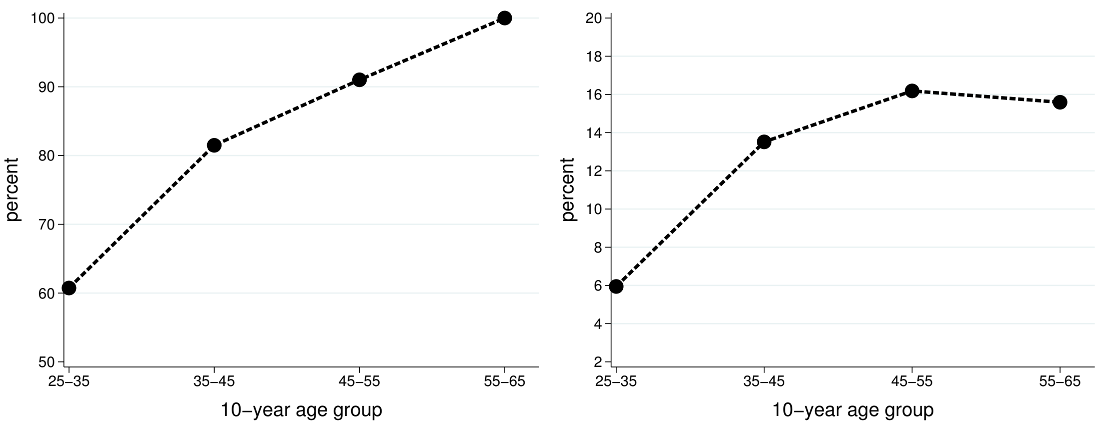

(a) earnings

(b) share of business owners

Notes: Data of the 2016 SCF. Only households with earnings greater than 0 are considered to calculate values of the left figure. Mean earnings in the right figure are normalized to the respective value of households aged between 55 and 65. Mean earnings of these households are equal to \$112868.

Retirement funds: Agents get with probability $\pi_p^l(\psi|l)$ a job offer in conjunction with a pension plan and with probability $\pi_p^b(\psi|b)$ a business idea with a pension plan. Again, we do not assume any specific functional form for the probability distributions

$$\Pi_p^l(\psi) = \{\pi_p^l(\psi|l_1), \pi_p^l(\psi|l_2), \ldots, \pi_p^l(\psi|l_s), \ldots, \pi_p^l(\psi|l_S)\} \text{ and}$$

$$\Pi_p^b(\psi) = \{\pi_p^b(\psi|b_1), \pi_p^b(\psi|b_2), \ldots, \pi_p^b(\psi|b_t), \ldots, \pi_p^b(\psi|b_S)\}, \text{ respectively.}$$

Since $\psi \in 0, 1$, $\pi_p^l(0|l_s) = 1 - \pi_p^l(1|l_s)$ and $\pi_p^b(0|b_t) = 1 - \pi_p^b(1|b_t)$. In other words, the probability of not being offered a pension plan is residually determined by 1 minus the probability of receiving a pension plan offer. Hence, it is sufficient to estimate the respective probabilities of receiving a pension plan offer. We impose the probability to get a plan to be increasing with the value of the job offer and business idea. In other words, the better the job or business opportunity, the more likely the agent is also offered a pension plan.[13] It is assumed that the lowest job opportunity comes always without a pensions plan, i.e. $\pi_p^l(1|l_1) = 0$. Thus, we have $S - 1 + T$ parameters to determine the probability distributions of pension fund offers.

As stated in Equation (4.4.3), the value of an agent's pension plan is determined by the parameters r_p, χ_l and χ_b, as well as σ.[14] In other words, the accumulation of retirement funds depends on how much interest is received, how large the share of labor and business income an agent pays into the fund and how large the employer's share that is paid into the fund. In order to solve the model numerically, we assume a finite number of pension states $P = \{p_1, \dots, p_k, \dots, p_K\}$. We find a new pension state $p' \in P$ in the following way: first, we calculate p'_{cont} according to Equation 4.4.3. Second, we the determine the states

$$p'_{\text{lower}} = \min_{p_k \in P \,\wedge\, p_k \leq p'_{\text{cont}}} |p'_{\text{cont}} - p_k|$$

and

$$p'_{\text{upper}} = \min_{p_k \in P \,\wedge\, p_k \geq p'_{\text{cont}}} |p'_{\text{cont}} - p_k|$$

Finally, we randomly draw[15] $\eta^p \in \{0, 1\}$ and the new pension state is then given by

$$p' = \eta^p \, p'_{\text{lower}} + (1 - \eta^p) p'_{\text{upper}}$$

Put another way, we first determine the two pension states, that are closest to the endogenously determined pension state p'_{cont}. The closest state from below is p'_{lower} and the closest from above is p'_{upper}. Hence, p'_{cont} always lies between these two states. Next, it is randomly determined whether the final pension state p' deviates from the endogenously determined state p'_{cont} upwards or downwards. By this we avoid a bias of pension states p' in the upper or lower direction. We interpret the realization of η^p as a random returns component on the retirement portfolio.

In order to determine the states $P = \{p_1, \dots, p_k, \dots, p_K\}$, we assume the distribution of pension states to be a piecewise linear function with $I < K$ pieces:

$$p_k = \begin{cases} \rho_1 & \text{if } \; k = 1 \\ \dots & \\ \dfrac{\rho_i - \rho_{i-1}}{k_i - k_{i-1}}(k - k_{i-1}) + \rho_i, & \text{if } \; i - 1 < k \leq i \; \text{ for } 1 < i < I \\ \dots & \\ \rho_I, & \text{if } \; k = K \end{cases}$$

[13] $\pi_p^l(1|l_{s-1}) < \pi_p^l(1|l_s)$ and $\pi_p^b(1|b_{t-1}) < \pi_p^b(1|b_t)$

[14] Due to the partial equilibrium model we treat the the interest on pension funds r_p as a parameter just like the interest on business wealth.

[15] We randomly draw from a uniform distribution with support $(0, 1)$ and round this draw to its nearest integer.

with $k_i = 1 + \frac{K-1}{I-1}(i-1)$ so that $P = \{p_1, \dots, p_k, \dots, p_K\}$ is determined by the parameters $\rho_1, \dots, \rho_I$. In other words, setting specific values for the states $\rho_1, \dots, \rho_I$ endogenously determines the remaining $K - I$ pension states by assuming these to lie on one of the piecewise linear functions between $\rho_1, \dots, \rho_I$. As a result, the state and probability distribution of pension funds requires the estimation of $4 + S - 1 + T + I$ parameters in total.

We set the number of retirement fund states to $K = 11$ and determine these by using $I = 5$ parameters.[16] Thus, there are 5 parameters which determine the distribution of possible pension plan states, $\{\rho_1, \rho_2, \rho_3, \rho_4, \rho_5\}$. The probabilities to receive a job offer or business opportunity in conjunction with a pension plan, $\{\pi_p^l(1|l_2), \pi_p^l(1|l_3), \pi_p^l(1|l_4), \pi_p^l(1|l_5)\}$ and $\{\pi_p^b(1|b_1), \pi_p^b(1|b_2), \pi_p^b(1|b_3), \pi_p^b(1|b_4)\}$, involve the estimation of 8 more parameters. To match the empirical distribution of retirement funds we target the extensive margin of retirement funds as well as the shares in aggregate retirement funds at different parts of the income distribution. In addition, the share of retirement funds in the wealth portfolio of a household is targeted for various income groups.

Housing: There are E different housing types $H = \{h_1, h_2, \dots, h_e, \dots, h_E\}$ with the unconditional probability distributions

$$\Lambda = \{\lambda(h_1), \dots, \lambda(h_e), \dots, \lambda(h_E)\} \text{ of finding a house of type } h_e \in H.$$

Besides, the probability distributions of being able to finance a house conditional on being in job state $l_s \in L$ or business state $b_t \in B$, respectively are given by

$$\Theta^l = \{\theta^l(l_1), \dots, \theta^l(l_s), \dots, \theta^l(l_S)\} \text{ and } \Theta^b = \{\theta^b(b_1), \dots, \theta^b(b_t), \dots, \theta^b(b_T)\}.$$

The probability of being able to buy a house of size h given that the agent is in state l or b is then given by $\lambda(h)\theta^l(l)$ for workers and $\lambda(h)\theta^b(b)$ for business owners with probability matrices $\Lambda'\Theta^l$ and $\Lambda'\Theta^b$, respectively. We neither impose any specific functional form on the distribution of housing states nor on the probability distributions. We only impose the housing value h_e to be increasing with e. Regarding the probabilities, we assume $\lambda(h_e)$ to be decreasing with e and $\theta^l(l_s)$ and $\theta^b(b_t)$ are increasing with s and t, respectively. In other words, the more a house is worth, the less likely it is found. The higher an agent's labor or business income, the more likely she can finance the house. We normalize the probability to find a house with the lowest value h_1 to 1, i.e. $\lambda(h_1) = 1$. In addition we assume, that an agent in the highest job state l_S is always able to finance a house so that $\theta(l_S) = 1$. Hence, there are $E - 1 + S - 1 + T$ parameters to determine the distribution of housing states and its probability distributions.

As stated in Equation (4.4.4) in Section 4.4, the mortgage value is determined by the parameters r_m, the interest on mortgages, and by κ, the share of debt that is repaid each period.[17] As for pensions, we assume a finite number of mortgage states $M = \{m_1, \dots, m_n, \dots, m_N\}$ with $N > E$ to solve the model numerically. The procedure

[16] As we have to choose a relatively small number for I due to an overproportionally high increase in computing time with the number of pension states, we use a cubic interpolation instead of linear interpolation for the current estimation of pension states.

[17] Due to the partial equilibrium model we treat the the interest on mortgages r_m as a parameter just like the interest on pensions funds and business wealth.

to find a new mortgage state m' is equivalent to that of pensions. It is given by[18] $m' = \eta^m \, m'_{\text{lower}} + (1-\eta^m)m'_{\text{upper}}$. We first determine the two mortgage values in $M = \{m_1, \ldots, m_n, \ldots, m_N\}$ which are closest to the endogenously determined value m'_{cont}. Second, to avoid any bias we choose randomly which of these two states is used as the new mortgage state m'. We interpret η^m to be an inflation shock. The interest rate in our model is real whereas mortgage contracts in the data are typically nominal. A change in inflation rates thus affects the remaining real balance of a mortgage. Similar to pension states we assume the distribution of mortgage states $M = \{m_1, \ldots, m_n, \ldots, m_N\}$ to be a piecewise linear function with E pieces:

$$m_n = \begin{cases} 0 & \text{if } \; n = 1 \\ \ldots & \\ \dfrac{h_e - h_{e-1}}{n_e - n_{e-1}}(n - n_{e-1}) + h_e, & \text{if } \; e-1 < n \leq e \; \text{ for } 1 < e < E \\ \ldots & \\ h_E, & \text{if } \; e = E \end{cases}$$

with $n_e = 1 + \frac{N-1}{E-1}(e-1)$. Simply put, we assume the mortgage states to lie on the piecewise linear functions between the housing states $H = \{h_1, h_2, \ldots, h_e, \ldots, h_E\}$. Thus, by choosing housing states, the distribution of mortgage states $M = \{m_1, \ldots, m_n, \ldots, m_N\}$ is determined. In total, we need $2 + E + E - 1 + S - 1 + T$ parameters to determine the state and probability distribution of housing.

For the estimation we use $E = 4$ different housing states and $N = E * 4 + 1$ mortgage states.[19] The distribution of housing states is thus determined by 4 parameters, $\{h_1, h_2, h_3, h_4\}$. The probabilities to find a house of a specific size and be able to finance it include 11 parameters in total, $\{\lambda(h_2), \lambda(h_3), \lambda(h_4)\}$, $\{\theta(l_1), \theta(l_2), \theta(l_3), \theta(l_4)\}$ and $\{\theta(b_1), \theta(b_2), \theta(b_3), \theta(b_4)\}$. Similar to pensions, we use as targets the extensive margin of home owners as well as the shares in aggregate home equity along the income distribution. Besides, we match the share of home equity in the wealth portfolio of a household for various income groups.

Other parameters:
J is the number of periods of an agent's working life. β is the time discount factor. An agent's within period utility is determined by τ, the income tax rate, α, which indicates how much an agent values consumption relative to housing, as well as by $\underline{h}$ and φ, which determine the value of housing service. Finally, the parameter δ determines how much an agent values the retirement period.

One model period is set to 5 years. The working live of an agent is assumed to last for $J = 9$ periods so that agents work for 45 years. We set the time discount factor β to 0.95 which corresponds to an annual discount factor of 0.99. and the income tax rate τ

[18] $m'_{\text{lower}} = \min_{m_n \in M \,\wedge\, m_n \leq m'_{\text{cont}}} |m'_{\text{cont}} - m_n|$, $m'_{\text{upper}} = \min_{m_n \in M \,\wedge\, m_n \geq m'_{\text{cont}}} |m'_{\text{cont}} - m_n|$ and $\eta^m \in \{0, 1\}$.

[19] As for pensions we have to choose a relatively small number for E due to an overproportionally high increase in computing time with the number of housing states. Hence, we use a cubic interpolation instead of linear interpolation for the current estimation of mortgage states.

to 0.1. We assume that the employer pays the same share χ_l into the retirement fund as an agent so that $\sigma = 1$. In addition χ_l is fixed to be 0.2 and χ_b is assumed to be $\chi_b = 2\chi_l$ as the employer's share is not paid for business owners. α, which indicates how much an agent values consumption relative to housing, is set to 0.5. φ which determines housing service is set to 0.5. It is assumed that $\underline{h}$ is endogenously determined and equal to $0.05\varphi h_1$, i.e. 5% of the lowest value of housing service. Housing service of new born households, who have not yet been able to buy a house, is always equal to this value. δ which determines the value assigned to the retirement period is set to 20. κ, the share of mortgage debt repaid is set to 0.2. Hence, on average mortgage debt is repaid after 5 periods which corresponds to 25 years. As descried in Section 4.4 we assume interest rates of business wealth, pension funds and mortgages to differ. The interest rate r_b is set to 0.34 which corresponds to an annual interest rate of 6% , r_p and r_m are set to 0.10 implying an annual interest of 2%.

To sum up, the calibration of the model requires the estimation of 44 parameters in total. We choose model parameters so that our model economy matches closely the targets described above. However, due to the nonlinearity and discontinuity of the model, there might be multiple solutions. In addition, as mentioned above, the space of endogenously determined states, such as retirement funds and mortgages, are approximated by a finite and relatively small number of states due to limited computing power. This adds discontinuous elements to the model when solving it numerically. To account for this, different sets of starting values were used. The results of the calibration are presented in the next section.

4.6 Results

This section discusses the estimated parameters of the Financial Frictions Model as well as on the resulting model fit. We target the model to data of the 2016 SCF.[20] Total wealth in our model is equivalent to core wealth, the sum of business wealth, retirement funds and home equity. As in Section 4.2, empirical wealth statistics using data of the SCF thus only include core wealth components. In Section 4.2 we have seen that the extensive margin is a central determinant of wealth inequality. More precisely, to understand wealth inequality, the key question is *how many* households have invested in wealth along the income distribution rather than how much. To account for this the targets used for the calibration include to a large extent extensive margins of core wealth components for various income and age groups. In contrast to Castaneda et al. (2003), we do not target our model directly to the distribution of wealth. In the following paragraphs we first analyze the results of the model with respect to the estimated income

[20] As in Section 4.2 only working-age households, i.e. households aged between 25 and 65, are considered. Both the targets using data of the SCF and the corresponding model statistics are calculated using only this group of households. Regarding the SCF, this implies that households with a head being younger than 25 or older than 65 are excluded. With regard to model, agents in the first and last period of their life cycle are excluded.

process of workers and business owners. In this context we also look at the distribution of business wealth. Afterwards, results regarding the other two wealth components, retirement funds and home equity, are presented. Finally, we evaluate how well our model matches both the marginal and in particular the joint distributions of income and wealth. We compare the results of the Financial Frictions Model to those of Castaneda et al. (2003)'s model.

Table 4.6.1: Parameters determining the distribution of labor states l

Job offer states	l_s	1	7	9	18	50
Probability to get job offer (in %)	$\pi^l(l_s)$	34.5	34.0	22.8	3.3	1.7

Notes: Estimation results of Financial Frictions Model targeted to 2016 SCF. One model period corresponds to 5 years.

In Table 4.6.1 we present estimated parameters related to the distribution of labor states. The different job offer states l_s are presented in the first row of Table 4.6.1. The corresponding probabilities of receiving an offer, $\pi^l(l_s)$, are displayed in the second row. The job offer states range from 1 to 50 meaning that workers with the highest earnings receive 50 times as much as the poorest worker.[21] Summing up the probabilities in the second row of Table 4.6.1 gives use the probability of receiving any job offer or, in other words, the probability of receiving no business idea. This value is equal to 96.3 % which implies that the probability of getting a business idea is 3.7 %. An agent is thus much more likely to stay a worker than becoming a business owner. Looking at the probability distribution of job offers, we see that by far most probability mass is concentrated in the three lowest labor states: each period the probability of receiving a job offer in state 1, 7 or 9 is 34.5 % , 34.0 % and 22.8 % respectively. In contrast, the second best offer with an income of 18 has a probability of only 3.3 %, the best offer with an income of 50 a probability of 1.7 %. As mentioned in Section 4.4, receiving a better job offer can also be understood as agents being internally promoted in their current jobs without changing their employer.

Analogously to Table 4.6.1 we see in Table 4.6.2 estimation results for parameters related to business states. The first row shows estimated business opportunity states b_t. In order to relate this to job offer states, the resulting business income, $r_b b_t$, is presented in the second row. The third row shows the probabilities of receiving a business opportunity, $\pi^b(b_t)$. In the fourth row the probability of getting a retirement plan with a business opportunity, $\pi_p^b(1|b_t)$, is presented. In the fifth row we see the probability of being offered a housing finance plan, $\theta^b(b_t)$, for business owners in state b_t. The business income received with the worst business idea is equal to 7 while agents receive 253 with the best business opportunity. The latter is about 5 times the highest labor income a

[21] As mentioned in Section 4.4, the lowest labor states is normalized to one.

Table 4.6.2: Parameters determining the distribution of business states b

Business opportunity states	b_t	19	36	126	744
Business income	$r_b b_t$	7	12	43	253
Probability to get business opportunity (in %)	$\pi^b(b_t)$	1.4	0.9	0.7	0.6

Notes: Estimation results of Financial Frictions Model targeted to 2016 SCF. One model period corresponds to 5 years. $r_b = 0.34$.

worker can earn. The other two business states generate a business income of 12 and 43, respectively. This implies that only agents with the best business idea end up with income higher than that for the best paid workers. Up to the third business state it could be worth for agents to become a worker again as a job offer might exceed their current business income. This is in particular true for agents in the lowest business state with an income of 7. These poor self employed can be seen as mom-and-pop store owners, for example.

As mentioned above, the probability of receiving any of these business opportunities is only 3.7 % each period. The best business idea is offered with a probability of 0.6 %. The other states with business income of 43, 12 and 7 are obtained with probabilities of 0.7 %, 0.9 % and 1.4 %. Similar to job offers, a change in business ideas does not necessarily mean that agents switch their business but can also be interpreted as an increase in the quality of a business idea.

Table 4.6.3: Realized distribution of labor states l and business states b

Job offer states l_s	1	7	9	18	50	$\sum$
Share of agents in job offer l_s (in %)	8.9	26.9	37.0	9.5	6.7	89.0

(a) labor states l

Business opportunity states b_t	19	36	126	744	$\sum$
Share of agents in business state b_t (in %)	2.0	3.3	3.2	2.5	11.0

(b) business states b

Notes: Estimation results of Financial Frictions Model targeted to 2016 SCF. One model period corresponds to 5 years.

Table 4.6.3 shows the realized distribution of labor and business states. Looking first at the distribution of labor states, we see in Table 4.6.3a that almost two thirds of the agents, 63.9 %, are in the second and third labor states. 8.9 % of the population are in the lowest state and 6.7 get the maximum amount of earnings. In total, 89 % of the agents are workers. The distribution of business owners, which correspond to 11 % of the population, is presented in Table 4.6.3b. Agents are relatively equally distributed across business states. 2 % are in the lowest business state which generate income similar to that in the second labor state. The highest business state in which agents receive income about 4 times the highest possible earnings is reached by 2.5 % of the population. The distribution of labor and business states significantly changes with agents' age as agents get a new job or business offer each period. Among households aged 25 to 35, 23.9 % are in the lowest labor state and only 2.6 % have already reached the highest job offer state. In addition, only 4.9 of the 25-35 year olds have become business owners. In contrast, 10.1 % of agents aged 55-65 receive the highest possible earnings and 16.9 % are business owners.

Figure 4.6.1: Financial Frictions Model targeted to 2016 SCF: life cylce profiles of earnings and the share of business owners

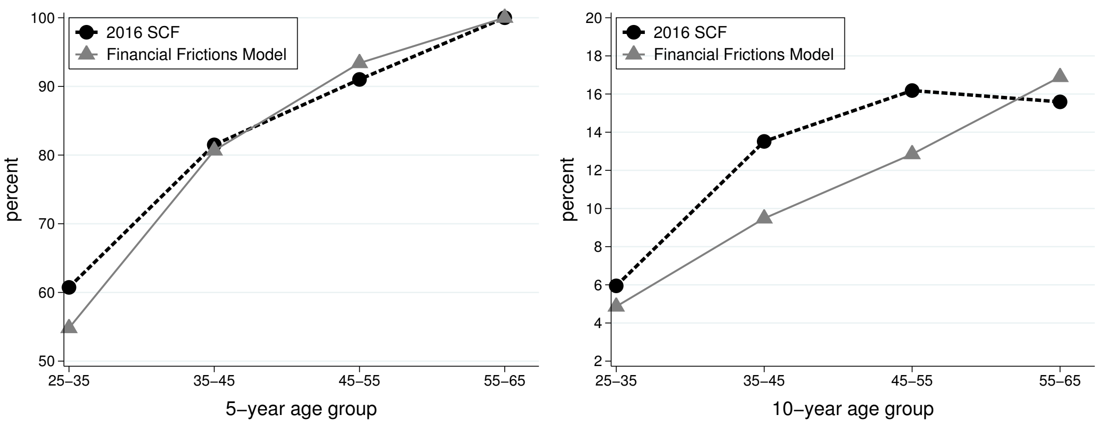

(a) earnings

(b) share of business owners

Notes: Estimation results of Financial Frictions Model targeted to 2016 SCF. Only working-age households, i.e. aged between 25 and 65, are included. Wealth is defined as the sum of nethousing, business equity and pension funds.

As mentioned in Section 4.5, the parameters determining the distribution of the labor and business states are targeted by using life cycle profiles of average earnings and the share of business owners. Figure 4.6.1 compares the life cycle profiles of earnings and the share of business owners generated by the Financial Frictions Model to the data. Figure 4.6.1a shows average earnings of workers by 10-year age groups. The black dotted line represents values using data the 2016 SCF. These are identical to those shown in Figure 4.5.1 in Section 4.5. The gray solid line are the respective values generated by the Financial Frictions Model. We see that the life cycle profile of average earnings is matched quite well. Average earnings of the 25-35 year olds are little bit too low and those of the 45-55 age group are a little bit too high. However, the model is consistent with an increasing life cycle profile. Regarding business owners, the model replicates this upward trend.In the model the share of business owners is increasing with age across all groups, whereas the share is slightly decreasing for the 55-65 year olds in the data.[22]

Table 4.6.4: Financial Frictions Model targeted to 2016 SCF: business wealth targets

	2016 SCF	Financial Frictions Model
extensive margin of business wealth of top 10% income households	0.35	0.33
extensive margin of business wealth of 80-90% income households	0.22	0.23
extensive margin of business wealth of 60-80% income households	0.14	0.17
share in aggregate business wealth of top 10% income households	0.82	0.81
share in aggregate business wealth of 80-90% income households	0.08	0.12
share in aggregate business wealth of 60-80% income households	0.05	0.05

Notes: Estimation results of Financial Frictions Model targeted to 2016 SCF. Only working-age households, i.e. aged between 25 and 65, are included. Wealth is defined as the sum of nethousing, business equity and pension funds.

In addition to life cycle profiles of earnings and business owners, the model is targeted to the empirical distribution of business wealth across income groups in order to estimate the parameters of the labor and business process. The results are presented in Table 4.6.4. The first three rows compare the extensive margin of business wealth for various income groups, the remaining rows are the shares in aggregate business wealth of each income group. As business wealth is concentrated in the top of the income distribution, we only use targets including households with income at the 6th decile and above. Looking at the data we see that the extensive margin of business owners is increasing with income: 14 % of households with income between the 6th and and 8th decile are business owners. For the top 10% income households, this proportion is more than twice as high at 35 %. The respective extensive margins generated in the model match the data closely. Typical differences range between 1 to 3 percentage points. The empirical distribution of aggregate business wealth to that in the model is compared in the lower three lines in Table 4.6.4. Business wealth is highly concentrated in the top of the distribution: top 10 % income households own 82 % of aggregate business assets. The Financial Frictions model replicates this highly unequal distribution very well. The share of the top 10 % differs by only one percentage point. The business share of households with income between the 8th and 9th decile is 4 percentage point higher compared to the data. Finally, the 6th to 8th decile income group has a share that perfectly

[22] Currently, our model does not include parameters such as life cycle components of labor or business income. We plan to introduce such components into the model in order to improve the fit to the life cycle profiles.

matches the data. To conclude, the empirical targets are matched by the model very well. In particular, as in the data, business owners are mainly found at the top of the income distribution and their business wealth is highly concentrated at the top. As mentioned in Section 4.5, we also use different point of the Lorenz curve of income in order to estimate labor and business parameters. These results are described further below where we contrast the marginal distributions of income and wealth generated by the Financial Frictions Model to those in the Model of Castaneda et al. (2003).

Table 4.6.5: Probability distribution of getting a pension plan or housing finance plan

Job offer states	l_s	1	7	9	18	50
Probability to get retirement fund (in %)	$\pi_p^l(1\|l_s)$	0.0	0.2	54.0	71.0	79.2
Probability to be offered a housing finance plan (in %)	$\theta^l(l_s)$	0.0	11.4	68.3	76.4	86.4

(a) probabilities depending on labor states l

Business opportunity states	b_t	19	36	126	744
Business income	$r_b b_t$	7	12	43	253
Probability to get retirement fund (in %)	$\pi_p^b(1\|b_t)$	43.8	97.8	98.0	98.6
Probability to be offered a housing finance plan (in %)	$\theta^b(b_t)$	41.7	48.0	85.3	99.8

(b) probabilities depending on business states b

Notes: Estimation results of Financial Frictions Model targeted to 2016 SCF. One model period corresponds to 5 years.

Next, we look at the estimated parameters related to retirement funds and housing, and evaluate how well the Financial Friction Model matches the empirical targets used for the estimation. Table 4.6.5 shows the probability of receiving a retirement fund or housing finance plan conditional on getting a job offer or business idea. The probabilities related to job offers l_s are presented in Table 4.6.5a. By assumption, agents being in the lowest job state neither get a retirement plan nor a housing finance plan. A job offer of 7 is also very unlikely to include a retirement plan: this probability is only 0.2 %. To get a housing finance plan in this state is with a probability of 11.4 % more likely. More than half of offers with earnings of 9, 54.0 %, contain a retirement plan and more than two-thirds, 68.3 %, include the offer of a housing finance plan. Finally, about three-quarters of job offer states 18 and 50 come with a pension plan: these probabilities are 71.0 % and 79.2 %, respectively. A finance plan is received with 76.4 % in state 18 and 86.4 % in state 50. In Table 4.6.5b we see the probabilities of getting a retirement plan or house financing in conjunction with a business idea. The probabilities are significantly higher compared to those of job offers. This is in particular true for the probability of getting a retirement plan. 43.8 % of worst business ideas involve a retirement plan and 41.7 % a housing finance plan.Regarding the second business state, nearly all ideas, 97.8 %, come with a retirement plan but only about half of them, 48 %, include an offer of a finance plan. The third business opportunity has a pension plan with a probability of 98 % and a finance plan with 85.3 %. Roughly all agents receiving the best idea also get a retirement plan as well as a housing finance plan. These probabilities are 98,6 % and 99.8 %, respectively. To sum up, we see that the estimated probabilities of getting the opportunity to invest either in retirement funds or home equity are strongly increasing with income. This indicates that our model generates agents which are *savings but not borrowing* constrained. In other words, these agents would like to save part of their income by investing in retirement funds or home equity. However, they are not able to invest as they do not have access to these assets. This contrasts with standard incomplete markets models, in which agents can freely decide how much of their income they want to save but have a specific (net) borrowing limit which often depends on an agent's income.

Table 4.6.6: Retirement fund states and housing states with probabilities of finding a house

Retirement fund states	ρ_i	2	17	51	123	350

(a) Retirement funds

Housing states	h_e	27	68	198	755
Probability to find house of size h_e (in %)	$\lambda(h_e)$	55.6	15.1	15.0	14.3

(b) Housing states and probabilities of finding a house

Notes: Estimation results of Financial Frictions Model targeted to 2016 SCF.

In Table 4.6.6 the distributions of retirement funds as well as housing states are presented. Levels of pension wealth range from 2 to 350 meaning that the maximal level of pensions agents can accumulate is about twice as much as an agent's average lifetime income. Regarding housing, the lowest state is 27 and the biggest house an agent can buy has a value of 755. This is about 4 times average lifetime income. A house with the lowest value is found with by far the highest probability of 55.6 %. The other houses are found with relatively similar probabilities that range from 14.3 to 15.1 % (see second row of Table 4.6.6 in the appendix).

Table 4.6.7 compares the empirical targets used to estimate the retirement funds and housing parameters to the respective values generated by the Financial Frictions Model. The first six rows compare the extensive margins of home equity and retirement funds, respectively, along the income distribution. The model fit with respect to home equity

Table 4.6.7: Financial Frictions Model targeted to 2016 SCF: retirement funds and home equity

	home equity		retirement funds	
	2016 SCF	Financial Frictions Model	2016 SCF	Financial Frictions Model
extensive margin of top 10% income households	0.95	0.95	0.93	0.90
extensive margin of 80-90% income households	0.89	0.90	0.86	0.84
extensive margin of 60-80% income households	0.81	0.81	0.79	0.74
extensive margin of 40-60% income households	0.65	0.62	0.60	0.60
extensive margin of 20-40% income households	0.46	0.49	0.39	0.43
extensive margin of bottom 20% income households	0.32	0.21	0.15	0.16
share in aggregate of top 10% income households	0.52	0.48	0.55	0.52
share in aggregate of 80-90% income households	0.13	0.12	0.19	0.15
share in aggregate of 60-80% income households	0.15	0.17	0.16	0.14
share in aggregate of 40-60% income households	0.11	0.08	0.07	0.04
share in wealth of top 10% income households	0.28	0.24	0.24	0.29
share in wealth of 80-90% income households	0.36	0.29	0.40	0.38
share in wealth of 60-80% income households	0.46	0.44	0.39	0.41
share in wealth of 40-60% income households	0.57	0.67	0.30	0.33

Notes: Estimation results of Financial Frictions Model targeted to 2016 SCF. Only working-age households, i.e. aged between 25 and 65, are included. Wealth is defined as the sum of nethousing, business equity and pension funds.

is quite good. The model deviates a little bit more from the data, in particular in the bottom of the income distribution: the share of households in the bottom 20 % of the income distribution is 32 % in the model but only 21 % in the data. The share of households having a retirement plan is also matched to the data. The maximum deviation is 5 percentage points for households with income between the 6th and 8th decile. The second block in Table 4.6.7, the 7th to 10th row, compares the shares in aggregate home equity and retirement funds, respectively, for income groups from the 4th decile upwards. Regarding home equity, the greatest deviation is in the top 10 %: their share is 52 % in the model but only 48 %. Most shares differ upwards meaning that the share in home equity of the bottom 40 % is slightly underestimated. The shares in aggregate retirement funds have a similar model fit as home equity. They deviate at most by 3 percentage points from the data. All of these differences are upward deviations so that the bottom 40 % receive too little of aggregate retirement funds.

Figure 4.6.2: Lorenz curves of income and wealth

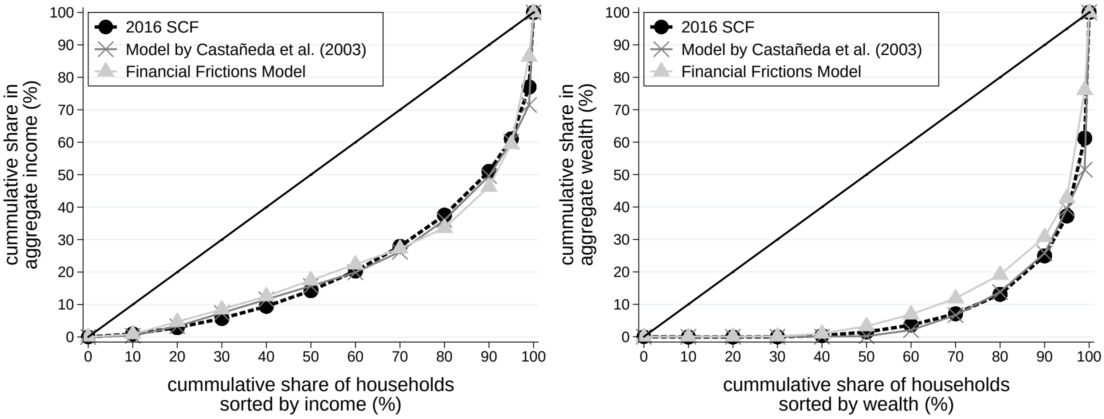

(a) Lorenz curve of income

(b) Lorenz curve of wealth

Notes: Data of the 2016 SCF. Only working-age households, i.e. aged between 25 and 65, are included.

Having described the estimated parameters of the Financial Frictions Model and the model fit to the empirical targets of the 2016 SCF, we now compare the marginal and joint distributions of income and wealth generated by the Financial Frictions Model to the respective distribution in the data as well as to the respective outcome using the model of Castaneda et al. (2003) we discussed in Section 4.3. For comparability we recalibrated the model of Castaneda et al. (2003) to the 2016 SCF data.[23] The marginal distributions of income and wealth are presented in Figure 4.6.2 and Table 4.6.8. As mentioned above, among others, the income process of the Financial Frictions Model is estimated using points on the Lorenz curve of income. Thus, the marginal distribution of income has to match that in the 2016 SCF by construction. For the calibration of the model of Castaneda et al. (2003) points of the Lorenz curve of income and wealth are used as targets. This implies that both the marginal distribution of income and wealth have to fit the data. In Figure 4.6.2 we see the Lorenz curves of income and wealth, respectively. The resulting Gini coefficients as well as share in aggregate income and wealth are shown Table 4.6.8. The presentation of results is analogously to Figure 4.3.1 and Table 4.3.1 in Section 4.3. The black dotted lines are Lorenz curves using data of the 2016 SCF. The dark gray lines with crosses are the curves resulting from the model by Castaneda et al. (2003). The results of the Financial Frictions Model are the light gray lines with triangles. By construction, the model by Castaneda et al. (2003) as well as the

[23]The recalibration is performed using targets analogously to those in Section 4.3.

Table 4.6.8: Distribution of income and wealth: SCF 2016 and modelCastaneda et al. (2003)

	Share in aggregate income (in %)						Gini coeff. of income
Income group	< 20%	20-40%	40-60%	60-80%	80-90%	> 10%	
2016 SCF	2.9	6.6	10.9	17.2	13.5	49.0	0.59
Castañeda et al. (2003)	3.3	8.3	8.3	16.1	13.5	50.4	0.57
Financial Frictions Model	4.7	7.9	9.8	11.2	12.7	53.7	0.58

(a) shares in income by income groups

	Share in aggregate wealth (in %)						Gini coeff. of wealth
Wealth group	< 20%	20-40%	40-60%	60-80%	80-90%	> 10%	
2016 SCF	-0.1	0.4	3.2	9.5	11.9	75.1	0.79
Castañeda et al. (2003)	0.0	0.0	2.0	11.8	11.9	74.3	0.80
Financial Frictions Model	0.0	1.0	5.9	12.3	11.6	69.3	0.79

(b) shares in wealth by wealth groups

Notes: Data of the SCF 2016. Only working-age households, i.e. aged between 25 and 65, are included.

Financial Frictions Model replicate almost exactly the empirical Lorenz curve of both income. Both the dark and the light gray line cover almost exactly the dotted black line. The corresponding Gini coefficients of income in Table 4.6.8a also fit to the empirical counterpart quite well: the Gini coefficient of 0.59 in the 2016 SCF is slightly lower both in the model of Castaneda et al. (2003), which generates a Gini of 0.57, and in the Financial Frictions Model with a Gini of 0.58. Looking at the shares in aggregate income for specific income groups we see that the two models slightly deviate from the data. The model of Castaneda et al. (2003) overstates income of households in the bottom 40 %, while the Financial Frictions Model overstates income of the top 10 %. Looking next at the distribution of wealth, the Lorenz curve generated by the Castaneda et al. (2003) model in Figure 4.6.2b matches the empirical curve almost exactly. As mentioned above, points of this curve are used as targets for the calibration of the model. The curve of the Financial Frictions Model also matches the data closely yet overstates middle class wealth. As can be seen in Table 4.6.8b, this deviation is mainly due to households with wealth between the 4th and 8th getting too much of aggregate wealth. However, the empirical Gini coefficient of 0.79 is matched by the Financial Frictions Model. The Gini coefficient of the Castaneda et al. (2003) model is also only marginally higher with a value of 0.80. In contrast to the model of Castaneda et al. (2003), we did not target the marginal distribution of wealth but only statistics of the wealth components which have been described above.

Figure 4.6.3: Joint Lorenz curves of income and wealth

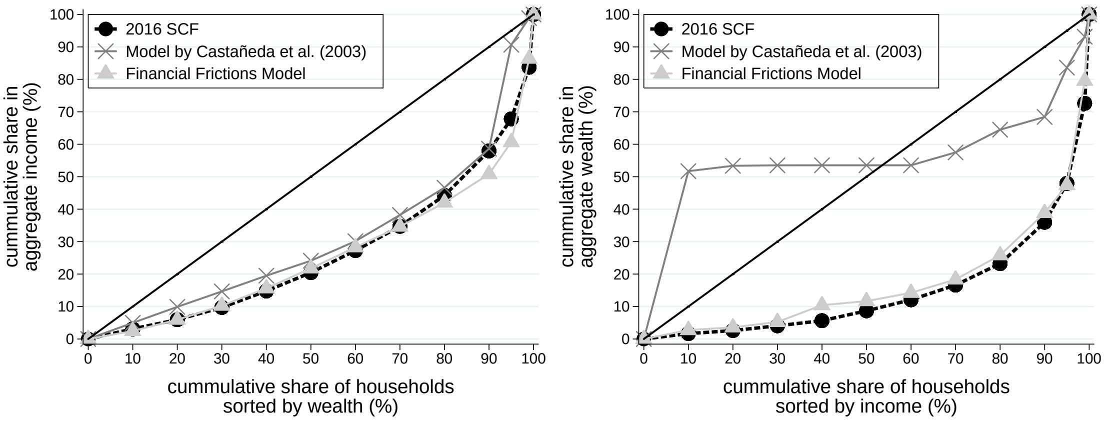

(a) Joint Lorenz curve of income by wealth (b) Joint Lorenz curve of wealth by income

Notes: Data of the SCF 2016. Only working-age households, i.e. aged between 25 and 65, are included.

Table 4.6.9: Joint distribution of income and wealth: SCF 2016 and modelCastaneda et al. (2003)

	Share in aggregate income (in %)						Joint Gini of income by wealth
Wealth group	< 20%	20-40%	40-60%	60-80%	80-90%	> 10%	
2016 SCF	5.9	8.8	12.5	17.0	13.7	42.1	0.48
Castañeda et al. (2003)	9.9	9.6	10.6	16.5	12.0	41.4	0.42
Financial Frictions Model	5.9	9.7	12.7	13.8	8.8	49.2	0.49

(a) shares in income by wealth groups

	Share in aggregate wealth (in %)						Joint Gini of wealth by income
Income group	< 20%	20-40%	40-60%	60-80%	80-90%	> 10%	
2016 SCF	2.6	3.0	6.4	11.1	12.7	64.0	0.71
Castañeda et al. (2003)	53.4	0.2	0.0	11.0	3.9	31.6	0.72
Financial Frictions Model	3.6	6.8	3.8	11.6	13.0	61.2	0.69

(b) shares in wealth by income groups

Notes: Data of the SCF 2016. Only working-age households, i.e. aged between 25 and 65, are included.

Next, we analyze how well the joint distributions of income and wealth are replicated by the two models. In contrast to the marginal distributions, these are not used as targets in order to solve either of either the model of Castaneda et al. (2003) or the Financial Frictions Model. In Figure 4.6.3 we see the joint Lorenz curve of income by wealth as well as the joint Lorenz curve of wealth by income. In Table 4.6.9 we see the respective joint Gini coefficients as well as the shares in income by wealth groups and the share in wealth by income groups. The presentation of results is analogously to Figure 4.3.2 and Table 4.3.2 in Section 4.3. As in Figure 4.6.2, the black dotted lines are Lorenz curves using data of the 2016 SCF. The dark gray lines with crosses are the curves resulting from the model by Castaneda et al. (2003). The results of the Financial Frictions Model are the light gray lines with triangles. Looking first at the the joint distribution of income by wealth displayed in Figure 4.6.3a, we see that the joint Lorenz curve generated by the Castaneda et al. (2003) model lies slightly above the empirical curve represented by black dotted line. The deviation is greatest at the top of the distribution. This implies that the joint distribution of income by wealth is more equally distributed compared to the data. Comparing the joint Gini coefficients of income by wealth in Table 4.6.9a underpins this result. We see that the model of Castaneda et al. (2003) generates a coefficient being too low compared to the data: the joint Gini of the 2016 SCF is 0.48 whereas that of the model is only 0.42. The light gray line of the Financial Frictions Model matches the black dotted line for the most part. Only in the upper part of the distribution it is slightly below the empirical curve. The joint distribution of income by wealth is thus a little bit more unequally distributed compared to the data. The resulting joint Gini coefficient of 0.49 slightly exceeds the empirical value of 0.48.

Looking now at the joint distribution of wealth by income in Figure 4.6.3b we see that the model of Castaneda et al. (2003) strongly deviates from the data. The shape of the joint Lorenz curve is similar to that of Figure 4.3.2b in Section 4.3 for the 1992 data.[24] We discussed that the fit is due to the model mechanism of having very wealth-rich households falling back into the lowest income state. This always results in much too high aggregate wealth shares in the bottom of the income distribution even when the model is directly targeted to the joint distribution of income and wealth. In Table 4.6.9b we see that the aggregate share in wealth of households in the bottom 20 % of the income distribution is more than 50 percentage points higher compared to the data. As the Lorenz curve lies party above the line of equality, the corresponding joint Gini coefficient cannot reasonably interpreted. As mentioned in Section 4.3, the strong deviation of the model of Castaneda et al. (2003) with regard to the joint distribution of income and wealth is due to the following mechanism: very rich households face a relatively high risk of becoming poor so that they have a strong precautionary savings motive and accumulate high amounts of wealth. Due to the high risk, the model is able to generate a high wealth concentration at the top. However, a significant pro-

[24] As for the result in 4.3, in addition to targeting the model of Castaneda et al. (2003) to the marginal distributions of income and wealth we also performed calibrations using the joint distribution of income and wealth as targets. However, as mentioned in Section 4.3, the model mechanism of having very wealth-rich households falling back into the lowest income state always results in much too high aggregate wealth shares in the bottom of the income distribution.

portion of the wealth rich does indeed slip into the lowest income group. As a result of these "crashed" households, the wealth share of the income poor becomes much too high compared to the data. Thus, the model matches the high wealth concentration at the top of the wealth distribution, but strongly overestimates wealth shares in the bottom of the income distribution so that wealth concentration at the top of the income distribution are significantly underestimated. In contrast to Castaneda et al. (2003), the Financial Frictions model matches the empirical curve almost exactly. Only in the middle the distribution the light gray line of the model is a little above the black dotted line representing the data. This means that the model distribution is slightly more equally distributed than that in the data. The resulting joint Gini coefficient of wealth by income is 0.69 being only a little bit below the empirical value of 0.71. As mentioned in Section 4.5, for the calibration of the Financial Frictions Model, we do not target it to the joint distribution of income and wealth. The concentration of wealth at the top is not generated through a strong precautionary savings motive as in Castaneda et al. (2003) but by introducing financial frictions: The more income a household earns the more likely it is to get access to wealth components which enables the household to accumulate wealth. A key feature of this mechanism is that high wealth concentration is generated both at the top of the wealth and income distribution.

To conclude, linking the access to assets to an agent's working state enables the replication of both the separate distributions of income and wealth as well as the joint distribution. Most importantly, in contrast to Castaneda et al. (2003) the accumulation of wealth along the income distribution matches the data very well. Both the marginal as well as the joint distributions of income and wealth are closely accounted for by the Financial Friction Model. The good fit of both the marginal wealth and the joint distribution of wealth and income is achieved without targeting our model directly to these distributions. Instead, we use statistics of wealth components as targets focusing on the extensive margins. The close fit supports the model's mechanism of limited asset access to account for the distribution of wealth in the data.

4.7 Policy Experiment

As illustrated in Section 4.2, the extensive margins of wealth components are essential in explaining how aggregate wealth is distributed along the income distribution. The model presented in Section 4.4 accounts for this by modeling explicitly the extensive margin to asset access. In particular, the probabilities of receiving a business opportunity, a pension plan or being able to finance a house depend on a household's labor market situation and its income as a reduced form to capture, for example, credit ratings. The income richer an agent, the more likely she is getting access to the housing market or employer sponsored retirement accounts. Despite being very restrictive regarding the intensive margin choice, we have seen in Section 4.6 that this mechanism generates both marginal and joint distributions of income and wealth that match to the data very well. This underpins that differences in the access to assets are a key component in understanding the accumulation of wealth as a function of income.

A key difference of the mechanism to standard incomplete markets models is that agents might be *savings but not borrowing constrained*. They would like to invest part of their income but do not have access to an investment opportunity. In particular income poor households are likely to be denied access to assets so that they cannot build up wealth and move up in the wealth distribution. There is an ongoing debate that missing access to housing for the poor is a key barrier to wealth accumulation and a driver of wealth inequality (see, for example, Hamilton and Famighetti (2019), Arundel and Hochstenbach (2019) or Iacoviello (2011)). In this section we are going to perform a simple policy experiment that aims at examining the effects of access to housing on wealth inequality. To do so we look at two dimensions of housing access: The first dimension examines the impact of variations in the probability of getting the opportunity to buy a house in the first place. We define this as the direct access effect. The second dimension evaluates how changes in financial conditions, such as the maturity of mortgage debt, affect wealth accumulation. This is defined as the indirect access effect. We are going to analyze the direct and indirect effects of asset access both separately and jointly.

The direct access to home equity is measured by the parameters θ^b and θ^l, the probabilities to be offered a housing finance plan conditional on being in a specific job and business state, respectively.[25] We simulate the Financial Frictions Model under tighter and looser direct access to home equity. The scenario of tighter access is generated with θ^b and θ^l being reduced by 50 percentage points compared to the model targeted to the 2016 SCF. In this scenarios, access to a housing financing plan has become harder across all agents. In the second scenario, we set θ^b and θ^l to 1 for all agents. This means, whenever an agent finds a house, she is always offered a financing plan. This scenario represents in simplified form a situation similar to the house price boom in the early 2000s in the U.S. when also households with bad credit scores could get financing to become home owners. As a measure for the indirect access to home equity we use the parameter κ. This parameter is the share of debt being repaid each period. The higher κ the smaller is the maturity of the agent's debt so that a household has to repay a higher share of its income each period. In the original model κ is set to 20 % so that average mortgage debt is repaid after 5 periods which corresponds to 25 years. The scenario of tighter indirect access to home equity assumes κ to be equal to 50 %. In this scenario the the maturity of mortgage debt is reduced to 10 years. Hence, in this simulation, agents have higher repayment rates. In the scenario of looser indirect access κ is decreased to 12.5 %. An average, agents repay their debt after 40 years. As the maturity of mortgages increases, agents have to repay less each period. We start our policy experiment with an analysis of the direct and indirect access effects on aggregate wealth. In a second step, we examine the aggregate welfare effects by calculating the consumption equivalent for each of the scenarios. Finally, we look at how access to home

[25] The parameter λ, the probability of finding a house, also influences the direct access to home equity. However, we do not apply this to our analysis as λ is independent from an agent's income.

equity affects the distribution of wealth across income groups and thus wealth inequality.

Figure 4.7.1: Access effects of home equity on aggregate household wealth (in %)

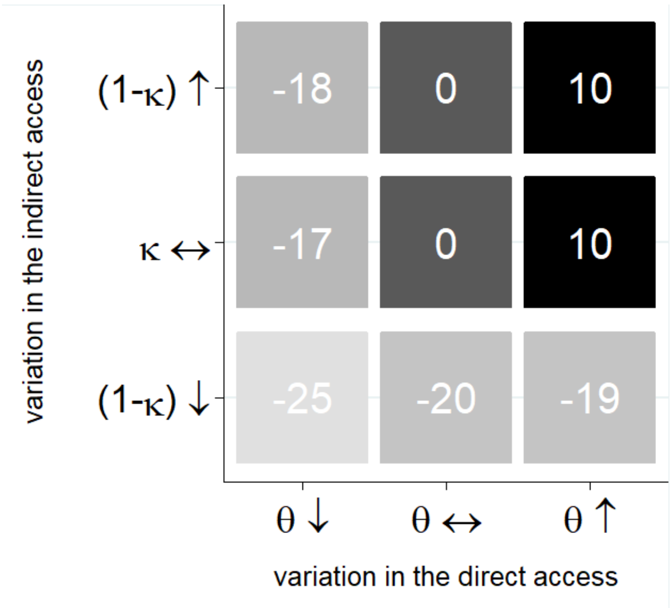

Notes: Simulation results of the Financial Frictions Model. Values in the row $(1-\kappa)\leftrightarrow$ are generated with κ being equal to the estimated parameter of the Financial Frictions Model targeted to data of the 2016 SCF. Values in the row $(1-\kappa)\downarrow$ are generated by increasing κ to 50 %. Values in the row $(1-\kappa)\uparrow$ are generated by decreasing κ t o12.5 %. Values in the row $\theta \leftrightarrow$ are generated with θ^b and θ^l being equal to the respective estimated parameter of the Financial Frictions Model targeted to data of the 2016 SCF. Values in the column $\theta\downarrow$ are generated by decreasing the estimation results for θ^b and θ^l of the targeted model by 50 percentage points. Values in the column $\theta\uparrow$ are generated by setting θ^b and θ^l to 1.

We see in Figure 4.7.1 how the level of aggregate household wealth changes with variations in the access to home equity. In particular, each tile shows the percentage change of aggregate wealth relative to the original model targeted to the 2016 SCF. On the x-axis the direct access to home equity, θ, is varied. The indirect access to housing, κ, is varied on the y-axis. The further right and the higher up a tile is, the better is the direct and indirect access, respectively. The middle tile represents the model targeted to the 2016 SCF so that its value has to be 0 by construction. The brighter the tile color of a particular scenario, the lower is the level of aggregate wealth generated by the respective scenario compared to the targeted model. In contrast, a dark tile color implies that aggregate wealth is higher compared to the targeted model. We see that both tighter direct and indirect access to home equity decreases aggregate wealth, whereas looser direct and indirect access generates an increase. Looking first at the direct access effect in isolation, we see that a tightening reduced aggregate wealth in the economy decreases by 17 %. In this scenario the probability of getting a financing plan is reduced: the extensive margin of housing, i.e. the share of homeowners in the population, declines from 61 % in the original model to 26.1 % (see Table 4.A.5 in the appendix). Put another way, the share of savings constrained households increases by 34.9 %. These households cannot accumulate wealth through an investment in housing anymore. Thus, aggregate wealth declines. This mechanism works exactly the other way round if the probability of getting a finance plan is increases, i.e. the direct access is loosened: aggregate wealth is 17 % higher in this scenario. There are less households which are savings constrained so that more households are able to invest in a home, thereby accumulating wealth. We do not observe this kind of symmetric mechanism for variation in the indirect access to housing. If indirect access to home equity is tightened, aggregate wealth decreases by 20 %. Households have to repay there debt much faster so that payment rates increase. In contrast to tighter direct access, households do not become savings constrained. However, for some households repayment rates are too high now so that they decide to not enter into the housing market and save less in total: the extensive margin of housing declines even more than in the scenario with tighter direct access to 19.3 % (see Table 4.A.5 in the appendix). Besides the effect on the home-ownership rate, households may decide to buy smaller houses so that repayment rates are still manageable, which further induces aggregate wealth to decline. In contrast to the direct access effect, loosening indirect access does not affect aggregate wealth in our simulation. In the original calibration the majority of households without a home are savings constrained households which would like to use part of their income to invest in housing, but are denied access to financing. If only the indirect access is improved, these households stay savings constrained: we see in Table 4.A.5 that the share of homeowners only slightly increases by 2.6 percentage points. Hence, without decreasing the number of savings constraint households by loosening the direct access, there is no significant effect on aggregate wealth. Analyzing now the joint effects of direct and indirect access to home equity, we see that they reinforce each other regarding tighter access. Once the direct access to home equity has become more difficult, making the indirect access also more burdensome further reduces aggregate wealth: in this case aggregate wealth is reduced by 25 %. In this scenario the share of savings constraint households is increases and those who get a finance plan now have to pay higher rates. Loosening both the direct and indirect access results in the same increase in aggregate wealth of 10 % as in the scenario with only looser direct access. As mentioned above, the original calibration most household without a house are savings constrained. Hence, an improving the access to financing, i.e. the direct access, is the main driver to generate more investments in housing and thus an increase in aggregate wealth. Through the lens of our model, the results so far suggest that financial frictions reducing access to credit prevents households from building up wealth. If more households would get the opportunity to get financing and become homeowners, their core wealth rates and total wealth would increase. As mentioned above, while many models of household consumption-saving behavior focus on credit constrained households, our model features a large share of savings constrained households, i.e. households who, if they get the opportunity to save more, will do so.

Next, we evaluate the welfare effects of variations in the access to home equity. To do so, we calculate the consumption equivalent of each scenario relative to the original model targeted to 2016 SCF data. The consumption equivalent tells us how much more or less agents have to consume in the original model so that their resulting lifetime utility is

Figure 4.7.2: Consumption equivalent for Access effects of home equity (in %)

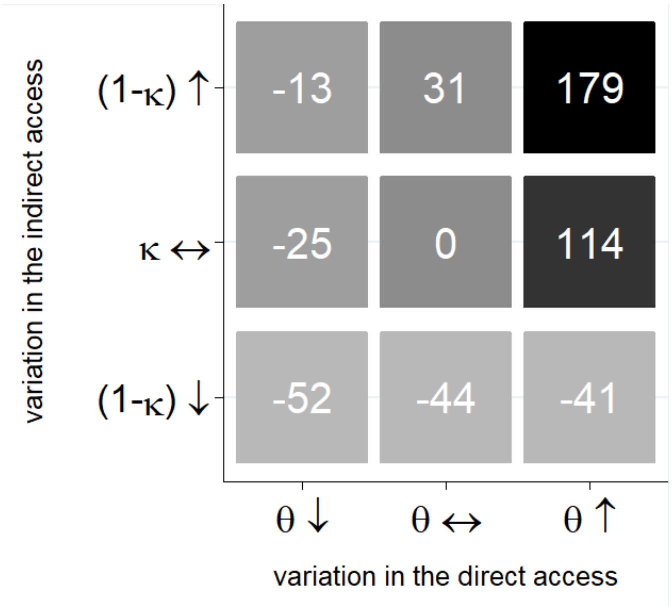

Notes: Simulation results of the Financial Frictions Model. Values in the row $(1-\kappa) \leftrightarrow$ are generated with κ being equal to the estimated parameter of the Financial Frictions Model targeted to data of the 2016 SCF. Values in the row $(1-\kappa) \downarrow$ are generated by increasing κ to 50 %. Values in the row $(1-\kappa) \uparrow$ are generated by decreasing κ t o12.5 %. Values in the row $\theta \leftrightarrow$ are generated with θ^b and θ^l being equal to the respective estimated parameter of the Financial Frictions Model targeted to data of the 2016 SCF. Values in the column $\theta \downarrow$ are generated by decreasing the estimation results for θ^b and θ^l of the targeted model by 50 percentage points. Values in the column $\theta \uparrow$ are generated by setting θ^b and θ^l to 1.

equal to that in the respective scenario.[26] Housing affects an agent's utility in two ways: on the one hand, homeowners have to reduce consumption in order to repay their mortgages during working life. On the other hand, households benefit from housing services by living in their home and from higher consumption during retirement as they use up their accumulated housing wealth. In Figure 4.7.2 we see the consumption equivalent for each of the scenarios. Similar to aggregate wealth, both tighter direct and indirect access reduce consumption: regarding direct access, it is 25 % below the consumption in the original calibration. Tighter indirect access reduces consumption by 44 %. If both direct and indirect access is reduced, consumption is more than halved, it decreases by 52 %. The mechanisms are the same as described above. The extensive margins of home equity are reduced in both scenarios. A reduction in the direct access generates more households which are savings constrained. Tighter indirect access, i.e. higher repayment rates, makes the payments for some households no longer affordable so that the homeownership rate declines. The effects of looser access are partly different to that of

[26] Let $V^m = \sum_{j=1}^{J-1} \beta^{j-1} u(c_j^m, h_j^m) = \sum_{j=1}^{J-1} \beta^{j-1}(\alpha log(c_j^m) + (1-\alpha) log(\underline{h} + \varphi h_j^m))$ be the lifetime value of an agent in the original Financial Frictions Model targeted to the 2016 SCF. V^s is the respective lifetime value of an agent in one of the scenarios with the access to housing being varied. The consumption equivalent is obtained by solving the equation $V^s = \sum_{j=1}^{J-1} \beta^{j-1} u(c_j^m(1+\Delta c^s), h_j^m)$ for Δc^s. Hence, the consumption equivalent of a scenario s is given by $\Delta c^s = e^{(V_s - V_m)\frac{1-\beta}{1-\beta^{J-1}}\frac{1}{\alpha}}$

aggregate wealth. A loosening of both the direct and indirect effect increases consumption and the two effects reinforce each other. With an increase in the direct access agents more than double their lifetime consumption: the consumption equivalent is 114 %. In this scenario, the share of savings constrained households is strongly reduced as the access to financing is improved. According to the policy function of savings constraint households it would be optimal to invest in a house but they are denied access to housing. In other words, the increase in utility from housing service and consumption would outweigh the decrease in consumption due to mortgage payments. Hence, households which are savings constrained in the original calibration and now become homeowners can increase their lifetime utility. The effect of the indirect access is significantly lower. In this scenario agents lifetime consumption is increased by 31 %. However, as looser indirect access only marginally affects aggregate wealth this increase in the consumption equivalent is remarkable. As mentioned above, the extensive margin of housing is only slightly increased in this scenario as the share of savings constrained households remains constant. However, looser indirect access reduces the payment rates during working life. Thus, agents can increase their consumption and benefit from housing service simultaneously, which increases lifetime utility. With both looser direct and indirect access we see a strong reinforcement of both effects: consumption increases by 179 %. As mentioned above, looser direct access reduces the share of savings constraint households. Households now get access to housing so that they can increase their lifetime utility through housing service and higher consumption when they are retired. In addition, due to looser indirect access each homeowner can increase consumption during working life as repayment rates are reduced. This suggests that for wealth accumulation it is important to have the opportunity to save but at low saving rates.

So far we have analyzed the effects of varying home equity access on the economy as a whole. We are now going to look at how the distribution of wealth changes across income groups and and how this in turn affects wealth inequality. Figures 4.7.3 and 4.7.4 show the access effects of home equity on the joint Gini coefficient and the joint Lorenz curve, respectively.[27] The brighter the tile color in Figure 4.7.3, the more equal is wealth distributed across the income distribution compared to the original calibration. In contrast, a dark tile color implies that wealth is more unequal distributed. The tile in the middle is the joint Gini coefficient of the original model. This value is identical to that in Table 4.6.9b in Section 4.6. Similar to the consumption equivalent, loosening the direct and indirect access to home equity affects the distribution of wealth across income groups more strongly than a tightening. Tighter direct an indirect access results in a Gini of 0.72 which is only slightly above the original value. If both the direct as well as the indirect access to home equity is made more accessible, the joint Gini coefficient decreases from 0.69 in the targeted model to 0.61 which corresponds to a reduction of about 12 %. As for aggregate wealth, loosening the direct access effect seem to be the key driver for the reduction in inequality: improving the only direct access results in

[27] The corresponding results for the standard Gini coefficient of wealth are shown in Figure 4.A.4 in the appendix.

Figure 4.7.3: Access effects of home equity on the joint Gini of wealth by income

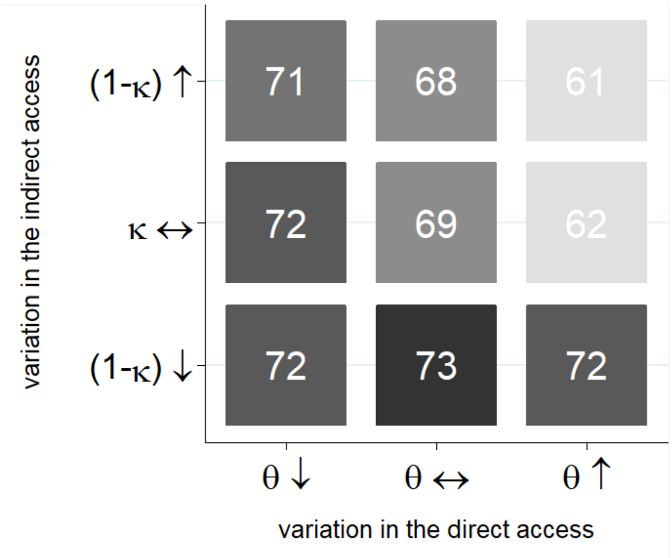

Notes: Simulation results of the Financial Frictions Model. Values in the row $(1-\kappa) \leftrightarrow$ are generated with κ being equal to the estimated parameter of the Financial Frictions Model targeted to data of the 2016 SCF. Values in the row $(1-\kappa) \downarrow$ are generated by increasing κ to 50 %. Values in the row $(1-\kappa) \uparrow$ are generated by decreasing κ to 12.5 %. Values in the row $\theta \leftrightarrow$ are generated with θ^b and θ^l being equal to the respective estimated parameter of the Financial Frictions Model targeted to data of the 2016 SCF. Values in the column $\theta \downarrow$ are generated by decreasing the estimation results for θ^b and θ^l of the targeted model by 50 percentage points. Values in the column $\theta \uparrow$ are generated by setting θ^b and θ^l to 1.

a joint Gini coefficient of 0.62 whereas the respective value of the indirect access effect is 0.68. The Gini coefficient does not tell us, in which way the distribution of wealth changes across income groups. It is possible that only the poor benefit while wealth shares of the middle class stay constant or vice versa.

The effect of access variations on the whole distribution of wealth is presented in Figure 4.7.4 which shows the joint Lorenz curve of wealth by income. The black line is the Lorenz curve of the original calibration. The dark gray line are results with tighter direct and indirect access. Looser direct and indirect access is represented by the light gray line. We first look at the effect of tighter access. We have seen in Figure 4.7.3 that this scenario the Gini coefficient slightly increases from 0.69 to 0.72. The corresponding Lorenz curve in Figure 4.7.4, the dark gray line, lies slightly above the black line up to the 5th income decile and slightly below it afterwards. Income poor households, the bottom 50 %, as well as the income rich, top top 10 %, gain wealth shares: their shares increase by 3.4 and 6.6 &, respectively (see Table 4.A.6 in the appendix). As a result, the middle class, households with income between the 5th and 9th decile, looses shares in aggregate wealth of 10.1 percentage points. Although the Gini coefficient only slightly increases in this scenario, the losses of middle class households are significant. We have seen above in Figure 4.7.1 that aggregate wealth decreases by 25 % in the scenario with both tighter direct and indirect access. The respective Lorenz curve of Figure 4.7.4 tells us that the losses in wealth shares of middle class households are disproportionately high compared to the income poor and rich. Tighter access decreases the extensive margin

Figure 4.7.4: Access effects of home equity on the joint Lorenz curve of wealth by income

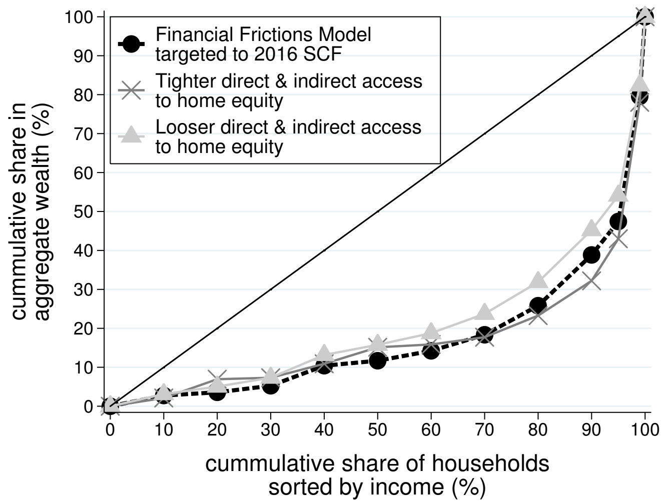

Notes: Simulation results of the Financial Frictions Model. The black dotted line is calculated using parameters of the Financial Frictions Model targeted to data of the 2016 SCF. The dark gray line with crosses is calculated by increasing κ to 50 % and by decreasing θ^b and θ^l of the targeted model by 50 percentage points. The light gray line with triangles is calculated by decreasing κ to 12.5 % and setting θ^b and θ^l to 1.

of the middle class by 73.8 %, whereas the homeownership rates of the poor and rich only drops by 40.6 % and 38.9 %, respectively. This implies that the number of savings constraint households which cannot accumulate wealth by investing in housing anymore, increases by far the most among middle class households. As the decline in savings is highest for this income group, wealth declines relatively more than for the poor and rich. Next, we analyze changes in the distribution of wealth with looser direct and indirect access to housing. This scenario generates a Gini coefficient of about 12 % below that of the original model. The respective Lorenz curve lies slightly above the black line which represents results of the original model. The poor and the middle class gain wealth shares of 4.1 and 2.3 percentage points, respectively, so that rich households loose (see Table 4.A.6 in the appendix). Aggregate wealth in this scenario is 10 % higher than in the original calibration (see Figure 4.7.1). The poor and the middle class receive a disproportionately large share of this increase so that wealth inequality decreases. With looser access, every household gets a financing offer so that there are no savings constraint households anymore. In addition, due to lower repayment rates the share of households which can afford to buy a home increases. The increase in the homeownership rate in this scenario is decreasing with income: among poor households the rate rises by 29.3 percentage points, among the middle class by 12.2 percentage points and the income rich observe only an increase by 2.7 percentage points (see Table 4.A.6 in the appendix). Thus, the growth in agents, which now get the opportunity to invest in home equity, is highest among poor and middle class households.

To conclude, tightening the access strongly increases the share of savings constraint households along the whole income distribution. Both aggregate wealth and aggregate welfare significantly decline. As the number of savings constraint households rises by far the most among middle class households, wealth declines disproportionately high for this income group. In the oppposite scenario with looser access there is a positive growth both in aggregate wealth and welfare. The induced increase in the extensive margin is highest for poor households, followed by the middle class. Thus, wealth shares of these income groups increase which leads to a decline wealth inequality. This result supports the idea that improving the access to homeownership will reduce wealth inequality. The fact that many poor Americans cannot buy houses increases wealth inequality (oftentimes this argument is made of African American households to explain the large black-white wealth gap, see, for example, Hamilton and Famighetti (2019)).

4.8 Conclusions

This paper provided two key new insights on the sources behind the distribution of wealth using data of the 2016 SCF. First, the empirical evidence presented in this paper documented that the extensive margin, meaning the question who has it and who does not, is an important driver of wealth inequality. We split the core wealth components, i.e. home equity, retirement funds and business equity, in their ex- and intensive margins. We show that differences in the extensive margins across households are key to account for the distribution of wealth. In contrast, the effect of variations in the intensive margins were of minor importance. In addition, as opposed to existing papers on wealth inequality, we looked at the joint distribution of income and wealth in detail. We show that variations in the extensive margins of wealth components along the income distribution are critical for explaining wealth inequality. In particular, if extensive margins were equal across income groups, the joint gini coefficient of wealth by income would be would be greatly reduced, so that wealth would be much more evenly distributed.

Second, to account for the empirical evidence on both the joint distribution of income and wealth as well as on the importance of the extensive margins of wealth components we build a new stylized model of asset accumulation that focuses on variation in the extensive margin in asset access rather than intensive margin variation. The latter are the focus of models where saving dynamics are determined by the Euler equation. Households in our model are like wealthy hand-to-mouth agents who are typically not on on their Euler equation. Indeed, we find that many households are *savings constraint*. After improving access to financial assets, they start to accumulate more wealth than before. This model is defined as the Financial Frictions Model. We show that the benchmark model on wealth inequality by Castaneda et al. (2003) which obtained wealth concentration at the top by introducing a strong precautionary savings motive, generated a joint distribution of income and wealth that was inconsistent with the data. The Financial Frictions Model departed from the widely adopted approach to consolidate portfolio positions from the household balance and to consider only net worth. Instead, we modeled the three major components home equity, retirement accounts, and business equity explicitly and introduced financial frictions regarding the access to these assets.

We abstracted from a consumption-savings decision of households: wealth inequality does not arise from a savings motive, but from limited access to assets depending (partly) on the income of a household. This is a key deviation from existing models. Our model generates agents which are *savings but not borrowing* constrained. In other words, these agents would like to save part of their income by investing in retirement funds or home equity. However, they are not able to invest as they do not have access to these assets. This contrasts with standard incomplete markets models, in which agents can freely decide how much of their income they want to save but have a specific (net) borrowing limit which often depends on an agent's income. Through this mechanism our model is able to account for the non-continuous distribution of asset holdings. More precisely, we show that once we have access to assets intertwined with the income process, the model is able to match heterogeneity along the extensive margin and the joint distribution of income and wealth.

The impact of financial frictions on the joint distribution of wealth and income was analyzed by a simple policy experiment in which the access to housing is varied. Tightening the access to home equity strongly increases the share of savings constraint households along the whole income distribution. Both aggregate wealth and aggregate welfare significantly decline. Losses are highest for the middle class, as the share of savings constraint households increases the most for this income group. Looser access leads to positive growth in aggregate wealth and welfare. Poor and middle class households gain the most from better access to home equity so that wealth inequality declines. This result supports the idea that improving the access to homeownership will reduce wealth inequality. The fact that many poor Americans cannot buy houses increases wealth inequality (oftentimes this argument is made of African American households to explain the large black-white wealth gap, see, for example, (Hamilton and Famighetti, 2019)).

In this paper we present new empirical facts on the joint distribution of income and wealth as well as on the importance of extensive margins, the question of who has it and who does not. The empirical evidence serves to inform the model building. The results of the simple stylized model, the Financial Fictions Model, show that incorporating financial frictions with respect to asset access into a model generates extensive margins as well as a joint distribution of income and wealth that matches the data.

4.A Appendix

4.A.1 Empirical analysis

Table 4.A.1: Composition of core and total wealth (2016 SCF)

	Wealth component	Share of component in complete wealth (in %)			
		all households	income groups		
			<50%	50-90%	>10%
	business wealth	25.3	14.7	14.5	30.1
+	home equity	23.0	47.0	32.9	17.1
+	retirement funds	18.5	18.2	30.0	14.6
	core wealth	**66.8**	**79.9**	**77.4**	**61.8**
+	mutual funds	10.5	2.5	4.1	13.4
+	stocks	5.8	1.7	2.7	7.3
+	liquid assets	5.7	6.8	6.9	5.3
+	other managed assets	4.3	4.0	3.4	4.6
+	value of vehicles	3.8	13.1	8.1	1.4
+	equity in non-resid. real est.	3.8	5.1	3.0	3.9
+	bonds	1.2	0.5	0.3	1.7
+	cash value of life insurance	1.0	1.48	1.6	0.8
+	other financial assets	0.7	1.3	0.5	0.7
+	other non-financial assets	0.5	0.9	0.5	0.4
-	credit card debt	-0.5	-1.7	-1.2	-0.1
-	other debt	-3.4	-15.4	-7.2	-0.9
	total wealth	**100.0**	**100.0**	**100.0**	**100.0**

Notes: Data of the 2016 SCF. Only working-age households (aged 25-65) are included. Business wealth is the value both active and non-active business interests. Home equity is the value of the household's primary residence minus debt on this residence plus the value of other residential real estate minus debt on this property. Retirement fund include IRAs, thrift accounts and the value of future pensions. Liquid assets are the sum of checking accounts, savings accounts, money market accounts, call accounts, prepaid cars and certificates of deposit. Other managed assets include trusts, annuities and managed investment accounts. Bonds include state and local bonds, government and government agency bonds and bills, corporate and foreign bonds, mortgage-backed bonds and savings bonds. Other financial assets include, among others, future proceeds, royalties and deferred compensation. Other non-financial assets are, among others, gold, jewelry and antiques. Other debt contains, among others, other lines of credit, installment debt and margin loans. Mean total wealth of all households is equal to \$607,161. Mean total wealth of the bottom 50% of the income distribution is \$88,137, that of households with income between the 5th and 9th decile is \$357,266 and that of the top 10% income households is \$4,201,865

Figure 4.A.1: Constant ex- and intensive margins

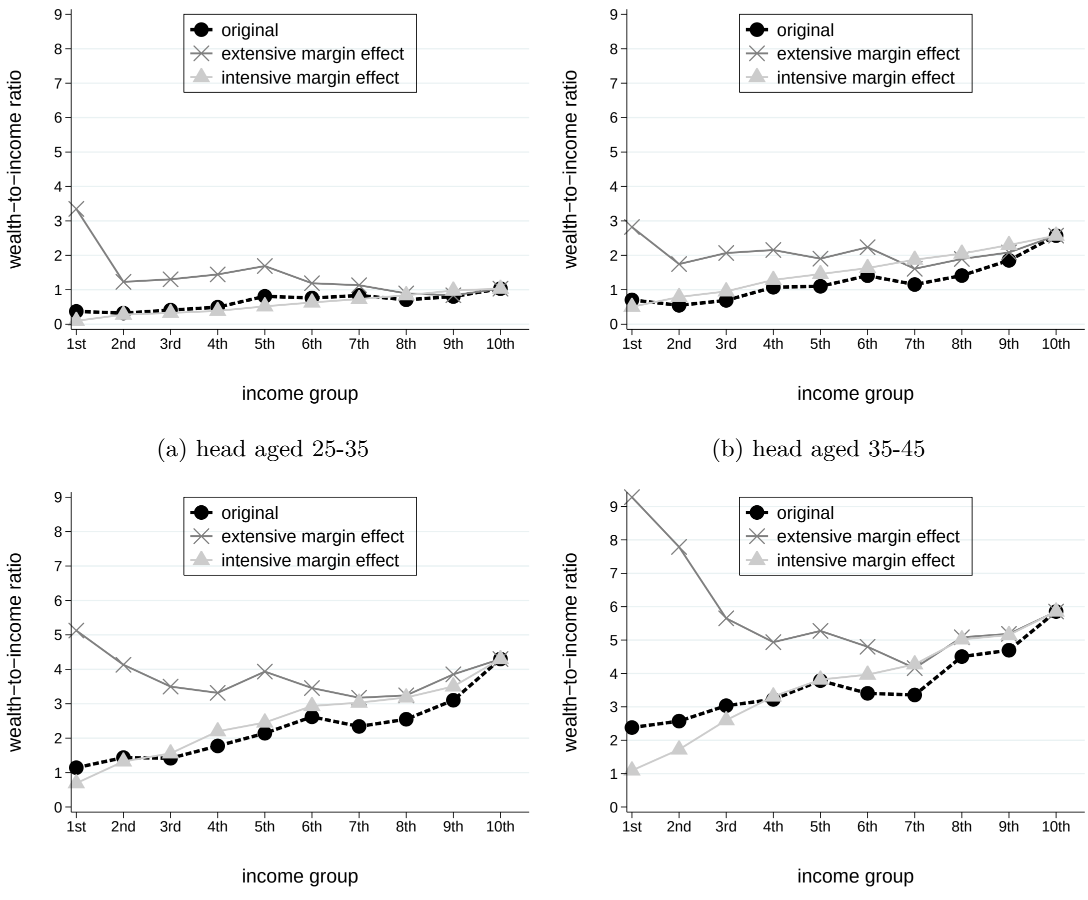

(a) head aged 25-35

(b) head aged 35-45

(c) head aged 45-55

(d) head aged 55-65

Data of the SCF 2016. Only working age households are included. Extensive margins of home equity, retirement funds and business equity are set to those of the top 10% of the income distribution of the respective 10-year age group.

Figure 4.A.2: Change in wealth-to-income ratios: Constant extensive margins

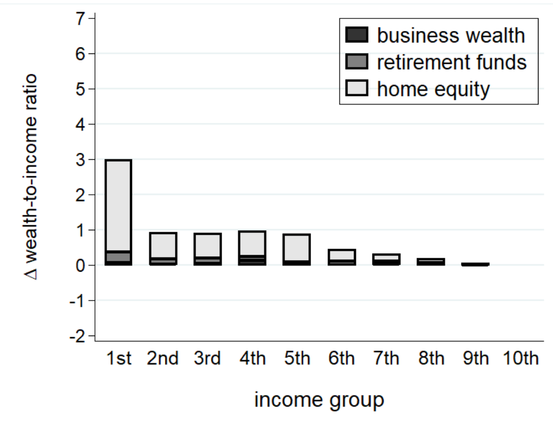

(a) households aged 25-35

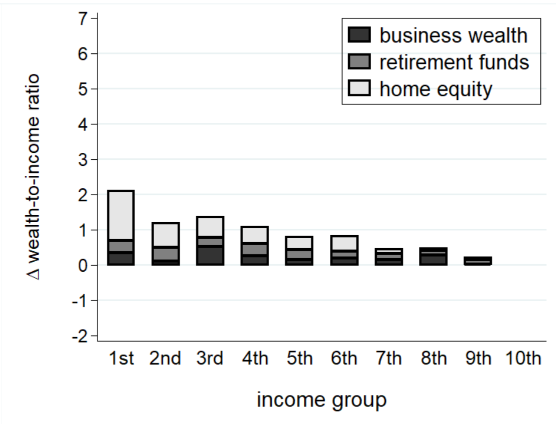

(b) households aged 35-45

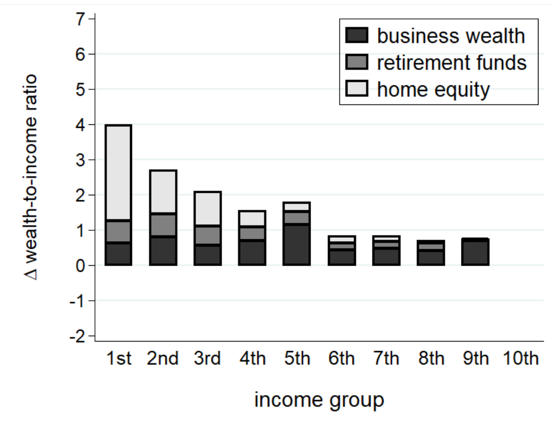

(c) households aged 45-55

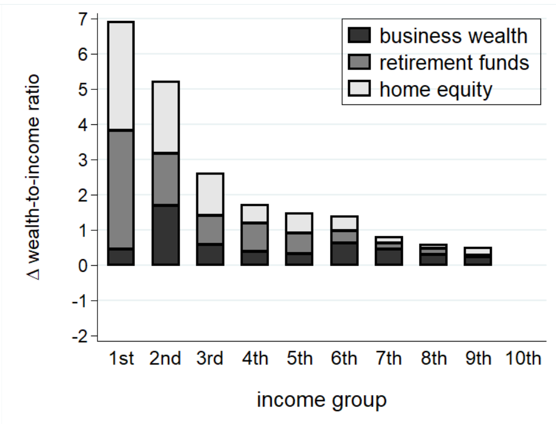

(d) households aged 55-65

Data of the SCF 2016. Only working age households are included. Extensive margins of home equity, retirement funds and business equity are set to those of the top 10% of the income distribution of the respective 10-year age group.

Figure 4.A.3: Change in wealth-to-income ratios: Constant intensive margins

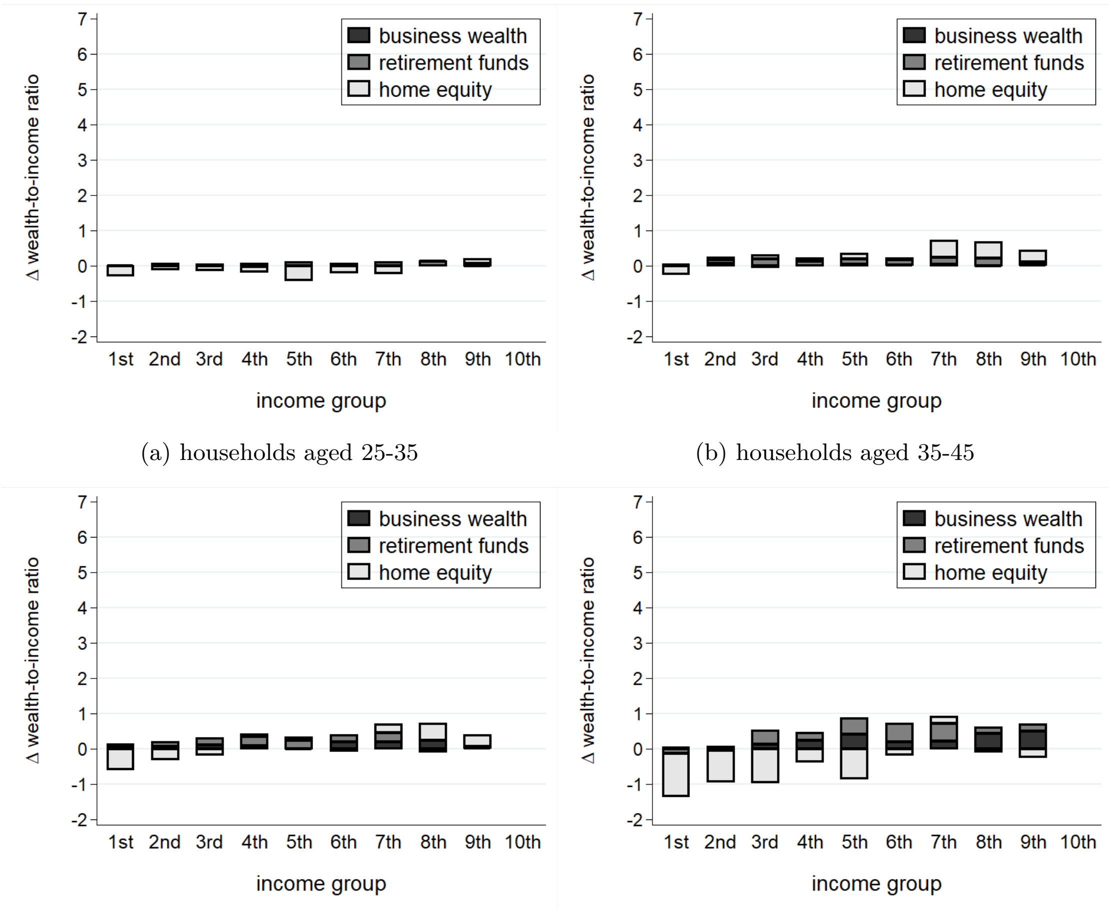

(a) households aged 25-35

(b) households aged 35-45

(c) households aged 45-55

(d) households aged 55-65

Data of the SCF 2016. Only working age households are included. Extensive margins of home equity, retirement funds and business equity are set to those of the top 10% of the income distribution of the respective 10-year age group.

4.A.2 Results

Table 4.A.2: Financial Frictions Model targeted to 2016 SCF: income targets

	2016 SCF	Financial Frictions Model
gini coefficient of income	0.59	0.58
share in aggregate income of top 10% income households	0.49	0.54
share in aggregate income of 80-90% income households	0.14	0.13
share in aggregate income of 60-80% income households	0.17	0.11

Notes: Estimation results of Financial Frictions Model targeted to 2016 SCF. Only working-age households, i.e. aged between 25 and 65, are included. Wealth is defined as the sum of nethousing, business equity and pension funds.

Table 4.A.3: Financial Frictions Model targeted to 2016 SCF: wealth targets

	2016 SCF	Financial Frictions Model
gini coefficient of wealth	0.16	0.17
share in aggregate wealth of top 10% wealth households	0.16	0.09
share in aggregate wealth of 80-90% wealth households	0.06	0.05
share in aggregate wealth of 60-80% wealth households	0.06	0.05

Notes: Estimation results of Financial Frictions Model targeted to 2016 SCF. Only working-age households, i.e. aged between 25 and 65, are included. Wealth is defined as the sum of nethousing, business equity and pension funds.

4.A.3 Policy Experiments

Figure 4.A.4: Access effects of home equity on the gini of wealth

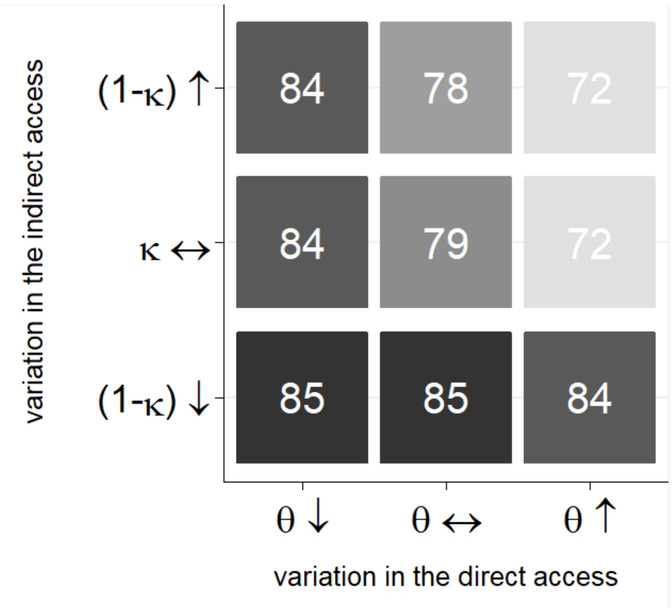

Notes: Simulation results of the Financial Frictions Model. Values in the row $(1-\kappa) \leftrightarrow$ are generated with κ being equal to the estimated parameter of the Financial Frictions Model targeted to data of the 2016 SCF. Values in the row $(1-\kappa) \downarrow$ are generated by increasing κ to 50 %. Values in the row $(1-\kappa) \uparrow$ are generated by decreasing κ to 12.5 %. Values in the row $\theta \leftrightarrow$ are generated with θ^b and θ^l being equal to the respective estimated parameter of the Financial Frictions Model targeted to data of the 2016 SCF. Values in the column $\theta \downarrow$ are generated by decreasing the estimation results for θ^b and θ^l of the targeted model by 50 percentage points. Values in the column $\theta \uparrow$ are generated by setting θ^b and θ^l to 1.

Figure 4.A.5: Access effects of home equity on the extensive margin of home equity by income groups

Income group		$< 50\%$	50-90%	$> 10\%$	‰
		Extensive margin of home equity (in %)			
Model targeted to 2016 SCF		44.0	83.7	95.3	61.0
Direct access effect of home equity	tighter ($\theta \downarrow$)	10.7	34.8	68.6	26.1
	looser ($\theta \uparrow$)	71.9	94.3	99.0	83.6
Indirect access effect of home equity	tighter ($(1-\kappa) \downarrow$)	4.5	21.2	85.8	19.3
	looser ($(1-\kappa) \uparrow$)	41.0	84.0	95.3	63.6
Access effect of home equity	tighter ($\theta \downarrow$, $(1-\kappa) \downarrow$)	3.4	9.9	56.4	11.3
	looser ($\theta \downarrow$, $(1-\kappa) \downarrow$)	73.3	95.9	99.0	84.9

Notes: Simulation results of the Financial Frictions Model. Values in the row $(1-\kappa) \leftrightarrow$ are generated with κ being equal to the estimated parameter of the Financial Frictions Model targeted to data of the 2016 SCF. Values in the row $(1-\kappa) \downarrow$ are generated by increasing κ to 50 %. Values in the row $(1-\kappa) \uparrow$ are generated by decreasing κ t o12.5 %. Values in the row $\theta \leftrightarrow$ are generated with θ^b and θ^l being equal to the respective estimated parameter of the Financial Frictions Model targeted to data of the 2016 SCF. Values in the column $\theta \downarrow$ are generated by decreasing the estimation results for θ^b and θ^l of the targeted model by 50 percentage points. Values in the column $\theta \uparrow$ are generated by setting θ^b and θ^l to 1.

Figure 4.A.6: Shares in aggregate wealth by income groups

Income group		Share in aggregate wealth (in %) < 50%	50-90%	> 10%
Model targeted to 2016 SCF		11.7	27.2	61.2
Direct access effect of home equity	tighter ($\theta \downarrow$)	12.5	21.3	66.2
	looser ($\theta \uparrow$)	16.4	27.2	56.4
Indirect access effect of home equity	tighter ($(1-\kappa) \downarrow$)	12.2	20.8	67.0
	looser ($(1-\kappa) \uparrow$)	11.9	28.8	59.3
Access effect of home equity	tighter ($\theta \downarrow$, $(1-\kappa) \downarrow$)	15.1	17.1	67.8
	looser ($\theta \downarrow$, $(1-\kappa) \downarrow$)	15.8	29.5	54.8

Notes: Simulation results of the Financial Frictions Model. Values in the row $(1-\kappa) \leftrightarrow$ are generated with κ being equal to the estimated parameter of the Financial Frictions Model targeted to data of the 2016 SCF. Values in the row $(1-\kappa) \downarrow$ are generated by increasing κ to 50 %. Values in the row $(1-\kappa) \uparrow$ are generated by decreasing κ t o12.5 %. Values in the row $\theta \leftrightarrow$ are generated with θ^b and θ^l being equal to the respective estimated parameter of the Financial Frictions Model targeted to data of the 2016 SCF. Values in the column $\theta \downarrow$ are generated by decreasing the estimation results for θ^b and θ^l of the targeted model by 50 percentage points. Values in the column $\theta \uparrow$ are generated by setting θ^b and θ^l to 1.

Bibliography

Adelino, M., A. Schoar, and F. Severino (2016): "Loan originations and defaults in the mortgage crisis: The role of the middle class," *The Review of Financial Studies*, 29, 1635–1670.

——— (2018): "The Role of Housing and Mortgage Markets in the Financial Crisis," *Annual Review of Financial Economics*, 10, 25–41.

Aladangady, A. (2017): "Housing Wealth and Consumption: Evidence from Geographically-Linked Microdata," *American Economic Review*, 107, 3415–46.

Albanesi, S., G. D. Giorgi, and J. Nosal (2017): "Credit Growth and the Financial Crisis: A New Narrative," *Working Paper*.

Alvaredo, F., L. Chancel, T. Piketty, E. Saez, and G. Zucman (2018): *World Inequality Report 2018*, Harvard University Press.

Andersen, H. Y. and S. Leth-Petersen (2019): "Housing Wealth Effects and Mortgage Borrowing: The Effect of Subjective Unanticipated Changes in Home Values on Home Equity Extraction in Denmark," CEPR Discussion Papers 13926, C.E.P.R. Discussion Papers.

Arundel, R. and C. Hochstenbach (2019): "Divided access and the spatial polarization of housing wealth," *Urban Geography*, 0, 1–27.

Attanasio, O. P. and L. Pistaferri (2016): "Consumption Inequality," *Journal of Economic Perspectives*, 30, 3–28.

Bach, L., L. E. Calvet, and P. Sodini (2016): "Rich Pickings? Risk, Return, and Skill in the Portfolios of the Wealthy," Tech. Rep. 11734, Centre for Economic Policy Research.

Bartscher, A., M. Kuhn, M. Schularick, and U. I. Steins (2019): "The Great American Debt Boom, 1949-2016," Tech. rep., University of Bonn.

Bartscher, A. K., M. Kuhn, and M. Schularick (2018): "The College Wealth Divide: Education and Inequality in America, 1956-2016," Tech. rep., University of Bonn.

Bayer, P. and K. K. Charles (2017): "Divergent Paths: A New Perspective on Earnings Differences Between Black and White Men Since 1940," *The Quarterly Journal of Economics*.

Benhabib, J. and A. Bisin (2018): "Skewed wealth distributions: Theory and empirics," *Journal of Economic Literature*, 56, 1261–91.

Benhabib, J., A. Bisin, and M. Luo (2017): "Earnings Inequality and Other Determinants of Wealth Inequality," *American Economic Review*, 107, 593–97.

Berger, D., V. Guerrieri, G. Lorenzoni, and J. Vavra (2017): "House Prices and Consumer Spending," *The Review of Economic Studies*, 85, 1502–1542.

Bernanke, B. S. (2005): "The Global Saving Glut and the U.S. Current Account Deficit," Speech 77, Board of Governors of the Federal Reserve System.

Bettner, J. (1982): "Your Money Matters: Investors May Find Sheltering Income Costlier Thanks to New, Alternative Minimum Tax," The Wall Street Journal.

Bhutta, N. (2015): "The Ins and Outs of Mortgage Debt during the Housing Boom and Bust," *Journal of Monetary Economics*, 76, 284 – 298.

Bhutta, N. and B. J. Keys (2016): "Interest Rates and Equity Extraction during the Housing Boom," *American Economic Review*, 106, 1742–74.

Bricker, J., L. J. Dettling, A. Henriques, J. W. Hsu, L. Jacobs, K. B. Moore, S. Pack, J. Sabelhaus, J. Thompson, and R. A. Windle (2017): "Changes in U.S. Family Finances from 2013 to 2016: Evidence from the Survey of Consumer Finances," *Federal Reserve Bulletin*, 103.

Bricker, J., A. Henriques, and P. Hansen (2018): "How much has wealth concentration grown in the United States? A re-examination of data from 2001-2013," Tech. rep., Discussion Paper.

Bricker, J., A. Henriques, J. Krimmel, and J. Sabelhaus (2016): "Measuring income and wealth at the top using administrative and survey data," *Brookings Papers on Economic Activity*, 2016, 261–331.

Burkhauser, R. V., S. Feng, and S. P. Jenkins (2009): "Using the P90/P10 index to measure US inequality trends with Current Population Survey data: A view from inside the Census Bureau vaults," *Review of Income and Wealth*, 55, 166–185.

Burkhauser, R. V., S. Feng, S. P. Jenkins, and J. Larrimore (2012): "Recent trends in top income shares in the United States: reconciling estimates from March CPS and IRS tax return data," *Review of Economics and Statistics*, 94, 371–388.

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

CALOMIRIS, C. W., S. D. LONGHOFER, AND W. MILES (2013): "The Housing Wealth Effect: The Crucial Roles of Demographics, Wealth Distribution and Wealth Shares," *Critical Finance Review*, 2, 49–99.

CAMPBELL, J. Y. AND J. F. COCCO (2007): "How do house prices affect consumption? Evidence from micro data," *Journal of monetary Economics*, 54, 591–621.

CANNER, G. B., K. E. DYNAN, AND W. PASSMORE (2002): "Mortgage refinancing in 2001 and early 2002," *Federal Reserve Bulletin*, 469–481.

CANNER, G. B., J. T. FERGUS, AND C. A. LUCKETT (1988): "Home Equity Lines of Credit," *Federal Reserve Bulletin*, 74, 361.

CASTANEDA, A., J. DÍAZ-GIMÉNEZ, AND J.-V. RÍOS-RULL (2003): "Accounting for the U.S. Earnings and Wealth Inequality," *Journal of Political Economy*, 111, 818–857.

CLOYNE, J., K. HUBER, E. ILZETZKI, AND H. KLEVEN (2017): "The Effect of House Prices on Household Borrowing: A New Approach," *Working Paper*.

COIBION, O., Y. GORODNICHENKO, M. KUDLYAK, AND J. MONDRAGON (2020): "Greater Inequality and Household Borrowing: New Evidence from Household Data," *Journal of the European Economic Association*.

COOPER, D. (2010): "Did Easy Credit Lead to Overspending? Home Equity Borrowing and Household Behavior in the Early 2000s," *Working Paper*.

DAVIS, S. J. AND T. VON WACHTER (2011): "Recessions and the Costs of Job Loss," *Brookings Papers on Economic Activity*.

DE NARDI, M. (2004): "Wealth inequality and intergenerational links," *The Review of Economic Studies*, 71, 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, 39–75.

DE STEFANI, A. (2018): "Debt, Inequality and House Prices: Explaining the Dynamics of Household Borrowing Prior to the Great Recession," *Journal of Housing Economics*.

DEFUSCO, A. A., C. G. NATHANSON, AND E. ZWICK (2017): "Speculative Dynamics of Prices and Volume," Working Paper 23449, National Bureau of Economic Research.

DETTLING, L. J., S. DEVLIN-FOLTZ, J. KRIMMEL, S. PACK, AND J. P. THOMPSON (2015): "Comparing micro and macro sources for household accounts in the United States: evidence from the Survey of Consumer Finances," Tech. rep., Federal Reserve Board.

Dettling, L. J., J. W. Hsu, L. Jacobs, K. B. Moore, and J. P. Thompson (2017): "Recent trends in wealth-holding by race and ethnicity: evidence from the survey of consumer finances," Tech. rep., Board of Governors of the Federal Reserve System (US).

Di Maggio, M. and A. Kermani (2017): "Credit-Induced Boom and Bust," *The Review of Financial Studies*, 30, 3711–3758.

Diaz, A., J. Pijoan-Mas, and J.-V. Rios-Rull (2003): "Habit formation: Implications for the wealth distribution," *Journal of Monetary Economics*, 50, 1257–1291.

Díaz-Giménez, J., A. Glover, and J.-V. Ríos-Rull (2011): "Facts on the Distributions of Earnings, Income, and Wealth in the United States: 2007 Update," *Federal Reserve Bank of Minneapolis Quarterly Review*, 34.

Duca, J. V. and A. Kumar (2014): "Financial Literacy and Mortgage Equity Withdrawals," *Journal of Urban Economics*, 80, 62 – 75.

Dynan, K. E. and D. L. Kohn (2007): "The rise in US household indebtedness: Causes and consequences," .

Fagereng, A., L. Guiso, D. Malacrino, and L. Pistaferri (2016): "Heterogeneity and Persistence in Returns to Wealth," Tech. rep., Discussion Paper.

Favara, G. and J. Imbs (2015): "Credit Supply and the Price of Housing," *American Economic Review*, 105, 958–92.

Fella, G. and M. De Nardi (2017): "Saving and Wealth Inequality," *Review of Economic Dynamics*.

Fetter, D. K. (2013): "How Do Mortgage Subsidies Affect Home Ownership? Evidence from the Mid-century GI Bills," *American Economic Journal: Economic Policy*, 5, 111–147.

——— (2014): *The Twentieth-Century Increase in U.S. Home Ownership: Facts and Hypotheses*, University of Chicago Press, chap. 10, 329–350.

Fligstein, N., O. P. Hastings, and A. Goldstein (2017): "Keeping up with the Joneses: How Households Fared in the Era of High Income Inequality and the Housing Price Bubble, 1999–2007," *Socius*, 3.

Foote, C. L., L. Loewenstein, and P. S. Willen (2016): "Cross-sectional patterns of mortgage debt during the housing boom: evidence and implications," Tech. rep., National Bureau of Economic Research.

Fortin, N., T. Lemieux, and S. Firpo (2011): "Chapter 1 - Decomposition Methods in Economics," in *Handbook of Labor Economics*, ed. by O. Ashenfelter and D. Card, Elsevier, vol. 4a, 1 – 102.

Fryer, Roland G., J. (2007): "Guess Who's Been Coming to Dinner? Trends in Interracial Marriage over the 20th Century," *Journal of Economic Perspectives*, 21, 71–90.

Furth, S. (2017): "Measuring Inflation Accurately," *Backgrounder*, June.

Fuster, A., B. Guttman-Kenney, and A. F. Haughwout (2018): "Tracking and Stress-testing U.S. Household Leverage," *Economic Policy Review*, 35–63.

Gabaix, X., J.-M. Lasry, P.-L. Lions, and B. Moll (2016): "The dynamics of inequality," *Econometrica*, 84, 2071–2111.

Ganong, P. and P. Noel (2018): "Liquidity vs. Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession," Working Paper 24964, National Bureau of Economic Research.

Garbinti, B., J. Goupille-Lebret, and T. Piketty (2017): "Accounting for wealth inequality dynamics: Methods, estimates and simulations for France (1800-2014)," Tech. rep., Discussion Paper.

Genesove, D. and C. J. Mayer (1997): "Equity and Time to Sale in the Real Estate Market," *The American Economic Review*, 87, 255–269.

Glover, A., J. Heathcote, D. Krueger, and J.-V. Rios-Rull (2017): "Intergenerational Redistribution in the Great Recession," *Working Paper.*

Gottschalk, P. and S. Danziger (2005): "Inequality of wage rates, earnings and family income in the United States, 1975–2002," *Review of Income and Wealth*, 51, 231–254.

Greenspan, A. and J. Kennedy (2008): "Sources and Uses of Equity Extracted from Homes," *Oxford Review of Economic Policy*, 24, 120–144.

Guerrieri, V. and H. Uhlig (2016): "Chapter 17 - Housing and Credit Markets: Booms and Busts," in *Handbook of Macroeconomics*, ed. by J. B. Taylor and H. Uhlig, Elsevier, vol. 2, 1427 – 1496.

Guren, A. M., A. McKay, E. Nakamura, and J. Steinsson (2018): "Housing Wealth Effects: The Long View," *Working Paper.*

Hamilton, D. and C. Famighetti (2019): "State of the Union 2019: Housing," Article, Stanford Center in Poverty and Inequality.

Haughwout, A., D. Lee, J. S. Tracy, and W. Van der Klaauw (2011): "Real Estate Investors, the Leverage Cycle, and the Housing Market Crisis," *FRB of New York Staff Report.*

Heathcote, J., F. Perri, and G. L. Violante (2010): "Unequal we stand: An empirical analysis of economic inequality in the United States, 1967-2006," *Review of Economic Dynamics*, 13, 15–51.

Henriques, A. M. and J. W. Hsu (2014): "Analysis of wealth using micro-and macro-data: A comparison of the survey of consumer finances and flow of funds accounts," in *Measuring Economic Sustainability and Progress*, University of Chicago Press, 245–274.

Herkenhoff, K. F. (2013): "The impact of consumer credit access on unemployment," *University of Califomia Los Angeles mimeo.*

Hoffmann, M. and I. Stewen (2019): "Holes in the Dike: The Global Savings Glut, U.S. House Prices, and the Long Shadow of Banking Deregulation," *Journal of the European Economic Association.*

Hubmer, J., P. Krusell, and A. A. Smith (2017): "The Historical Evolution of the Wealth Distribution: A Quantitative-Theoretic Investigation," *Working Paper.*

Hurst, E. and F. Stafford (2004): "Home is where the equity is: Mortgage refinancing and household consumption," *Journal of Money, credit and Banking*, 985–1014.

Iacoviello, M. (2005): "House prices, borrowing constraints, and monetary policy in the business cycle," *American economic review*, 95, 739–764.

Iacoviello, M. M. (2011): "Housing wealth and consumption," *FRB International Finance Discussion Paper.*

Jordà, O. (2005): "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 95, 161–182.

Jordà, Ò., K. Knoll, D. Kuvshinov, M. Schularick, and A. M. Taylor (forthcoming): "The rate of return on everything, 1870–2015," *Quarterly Journal of Economics.*

Jordà, O., B. Richter, M. Schularick, and A. M. Taylor (2017): "Bank Capital Redux: Solvency, Liquidity, and Crisis," *Working Paper.*

Jordà, O., M. Schularick, and A. M. Taylor (2013): "When Credit Bites Back," *Journal of Money, Credit and Banking*, 45, 3–28.

Jordà, Ò., M. Schularick, and A. M. Taylor (2017): "Macrofinancial History and the New Business Cycle Facts," *NBER Macroeconomics Annual*, 31, 213–263.

Kaplan, G., K. Mitman, and G. L. Violante (2017): "The housing boom and bust: Model meets evidence," Tech. rep., National Bureau of Economic Research.

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

Kaplan, G., G. L. Violante, and J. Weidner (2014): "The Wealthy Hand-to-Mouth," *V Brookings Papers on Economic Activity, Spring.*

KAYMAK, B. AND M. POSCHKE (2016): "The evolution of wealth inequality over half a century: The role of taxes, transfers and technology," *Journal of Monetary Economics*, 77, 1–25.

KENNICKELL, A. B. (2017): "Modeling wealth with multiple observations of income: Redesign of the sample for the 2001 Survey of Consumer Finances," *Statistical Journal of the IAOS*, 33, 51–58.

KENNICKELL, A. B. AND R. L. WOODBURN (1999): "Consistent Weight Design for the 1989, 1992 and 1995 SCFs, and the Distribution of Wealth," *Review of Income and Wealth*, 45, 193–215.

KINDERMANN, F. AND D. KRUEGER (2014): "High marginal tax rates on the top 1%? Lessons from a life cycle model with idiosyncratic income risk," Tech. rep., National Bureau of Economic Research.

KLYUEV, V. AND P. MILLS (2007): "Is Housing Wealth an "ATM"? The Relationship between Household Wealth, Home Equity Withdrawal, and Saving Rates," *IMF Staff Papers*, 54, 539–561.

KNOLL, K., M. SCHULARICK, AND T. STEGER (2017): "No Price Like Home: Global House Prices, 1870-2012," *American Economic Review*, 107, 331–53.

KOPCZUK, W. (2015): "What Do We Know about the Evolution of Top Wealth Shares in the United States?" *Journal of Economic Perspectives*, 29, 47–66.

KOPCZUK, W. AND E. SAEZ (2004): "Top Wealth Shares in the United States, 1916-2000: Evidence From Estate Tax Returns," *National Tax Journal*, 57, 445–487.

KOPCZUK, W., E. SAEZ, AND J. SONG (2010): "Earnings inequality and mobility in the United States: evidence from social security data since 1937," *The Quarterly Journal of Economics*, 125, 91–128.

KORINEK, A. AND A. SIMSEK (2016): "Liquidity Trap and Excessive Leverage," *American Economic Review*, 106, 699–738.

KOWALEWSKI, K. J. (1987): "Home Equity Lines: Characteristics and Consequences," *Economic Commentary*.

KRIPPNER, G. R. (2012): *Capitalizing on Crisis Capitalizing on Crisis - The Political Origins of the Rise of Finance*, Harvard University Press.

KRUSELL, P. AND A. A. SMITH, JR (1998): "Income and wealth heterogeneity in the macroeconomy," *Journal of political Economy*, 106, 867–896.

KUHN, M. AND J.-V. RÍOS-RULL (2016): "2013 Update on the US Earnings, Income, and Wealth Distributional Facts: A View from Macroeconomics," *Federal Reserve Bank of Minneapolis Quarterly Review*, 37.

KUHN, M., M. SCHULARICK, AND U. I. STEINS (forthcoming): "Income and Wealth Inequality in America, 1949-2016," *Journal of Political Economy*.

KUMHOF, M., R. RANCIÈRE, AND P. WINANT (2015): "Inequality, Leverage, and Crises," *The American Economic Review*, 105, 1217–1245.

LACOUR-LITTLE, M., E. ROSENBLATT, AND V. YAO (2010): "Home Equity Extraction by Homeowners: 2000–2006," *Journal of Real Estate Research*, 32, 23–46.

LOONEY, A. AND C. YANNELIS (2015): "A crisis in student loans?: How changes in the characteristics of borrowers and in the institutions they attended contributed to rising loan defaults," *Brookings Papers on Economic Activity*, 2015, 1–89.

LUSARDI, A., O. S. MITCHELL, AND N. OGGERO (2018): "The Changing Face of Debt and Financial Fragility at Older Ages," *AEA Papers and Proceedings*, 108, 407–11.

MAKI, D. M. (1996): "Portfolio Shuffling and Tax Reform," *National Tax Journal*, 317–329.

——— (2001): "Household Debt and the Tax Reform Act of 1986," *American Economic Review*, 91, 305–319.

MALMENDIER, U. AND S. NAGEL (2011): "Depression Babies: Do Macroeconomic Experiences affect Risk Taking?" *The Quarterly Journal of Economics*, 126, 373–416.

MIAN, A., L. STRAUB, AND A. SUFI (2019): "The Saving Glut of the Rich and the Rise in Household Debt," *Working Paper*.

MIAN, A. AND A. SUFI (2009): "The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis," *The Quarterly Journal of Economics*, 124, 1449–1496.

——— (2011): "House Prices, Home Equity-Based Borrowing, and the U.S. Household Leverage Crisis," *American Economic Review*, 101, 2132–56.

——— (2014): "House Price Gains and Us Household Spending from 2002 to 2006," Tech. rep., National Bureau of Economic Research.

——— (2017): "Household Debt and Defaults from 2000 to 2010: The Credit Supply View," in *Evidence and Innovation in Housing Law and Policy*, ed. by L. A. Fennell and B. J. Keys, Cambridge University Press, 257–288.

——— (2018): "Credit Supply and Housing Speculation," Working Paper 24823, National Bureau of Economic Research.

MICHEL, J.-B., Y. K. SHEN, A. P. AIDEN, A. VERES, M. K. GRAY, J. P. PICKETT, D. HOIBERG, D. CLANCY, P. NORVIG, J. ORWANT, S. PINKER, M. A. NOWAK, AND E. L. AIDEN (2011): "Quantitative Analysis of Culture Using Millions of Digitized Books," *Science*, 331, 176–182.

Modigliani, F., R. Brumberg, et al. (1954): "Utility Analysis and the Consumption Function: An Interpretation of Cross-section Data," *Post-keynesian economics*, 1, 338–436.

Mueller, H. M. and C. Yannelis (2019): "The rise in student loan defaults," *Journal of Financial Economics*, 131, 1–19.

Pfeffer, F. T., R. F. Schoeni, A. Kennickell, and P. Andreski (2016): "Measuring Wealth and Wealth Inequality: Comparing Two U.S. Surveys," *Journal of Economic and Social Measurement*, 41, 103–120.

Piketty, T. and E. Saez (2003): "Income Inequality in the United States, 1913-1998," *The Quarterly Journal of Economics*, 118, 1–39.

——— (2006): "The Evolution of Top Incomes: A Historical and International Perspective," *American Economic Review*, 96, 200–205.

Piketty, T., E. Saez, and G. Zucman (2017): "Distributional national accounts: methods and estimates for the United States," *The Quarterly Journal of Economics*, 133, 553–609.

Quadrini, V. (2000): "Entrepreneurship, saving, and social mobility," *Review of economic dynamics*, 3, 1–40.

Rajan, R. G. (2011): *Fault Lines: How Hidden Fractures Still Threaten the World Economy*, Princeton University press.

Rappaport, J. (2007): "A Guide to Aggregate House Price Measures," *Federal Reserve Bank of Kansas City Economic Review*, 41–71.

Sabelhaus, J., D. Johnson, S. Ash, D. Swanson, T. I. Garner, J. Greenlees, S. Henderson, J. Sabelhaus, S. Ash, and G. Garner (2015): "Is the Consumer Expenditure Survey Representative by Income?" *Improving the Measurement of Consumer Expenditures*, 74, 241.

Saez, E. and G. Zucman (2016): "Wealth inequality in the United States since 1913: Evidence from capitalized income tax data," *The Quarterly Journal of Economics*, 131, 519–578.

Schenker, N. and J. M. Taylor (1996): "Partially parametric techniques for multiple imputation," *Computational Statistics & Data Analysis*, 22, 425–446.

Schmitt-Grohé, S. and M. Uribe (2016): "Downward Nominal Wage Rigidity, Currency Pegs, and Involuntary Unemployment," *Journal of Political Economy*, 124, 1466–1514.

Schularick, M. and A. M. Taylor (2012): "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870-2008," *American Economic Review*, 102, 1029–61.

Shiller, R. J. (2009): *Irrational Exuberance*, New York: Broadway Books, 2 ed.

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

Story, L. (2008): "Home Equity Frenzy Was a Bank Ad Come True," *The New York Times*.

Thompson, J. P. and G. Suarez (2017): "Updating the Racial Wealth Gap," Tech. rep., Federal Reserve Board.

Wolff, E. N. (2002): "Top Heavy: The Increasing Inequality of Wealth in America and What Can Be Done about It," .

——— (2010): "Recent trends in household wealth in the United States: Rising debt and the middle-class squeeze-an update to 2007," .

——— (2016): "Household Wealth Trends in the United States, 1962 to 2013: What Happened over the Great Recession?" *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 2, 24–43.

——— (2017): "Household Wealth Trends in the United States, 1962 to 2016: Has Middle Class Wealth Recovered?" Working Paper 24085, National Bureau of Economic Research.

Wright, R. (2012): "Louis Hyman. Debtor Nation: The History of America in Red Ink. Princeton, NJ: Princeton University Press, 2011. 378 pp. ISBN 978-0-691-14068-1, $35.00 (cloth); 978-1-400-83840-0$,35.00 (e-book)." *Enterprise and Society*, 13, 195–197.

Yellen, J. (2014): "Perspectives on Inequality and Opportunity from the Survey of Consumer Finances," Remarks by Janet L. Yellen, Chair Board of Governors of the Federal Reserve System at the Conference on Economic Opportunity and Inequality, Federal Reserve Bank of Boston, Boston, Massachusetts, October 17, 2014 [Accessed : March 22, 2019].

Zandi, M. (2019): "How Is the Middle-Class Family Faring in Today's Economy?" *Written Testimony before U.S. House Committee on Ways and Means*.

Zinman, J. (2015): "Household Debt: Facts, Puzzles, Theories, and Policies," *Annual Review of Economics*, 7, 251–276.