

MODIGLIANI MEETS MINSKY: AMERICAN HOUSEHOLD DEBT, 1949-2016*

Alina K. Bartscher[†] Moritz Kuhn[‡] Moritz Schularick[§]
Ulrike I. Steins[¶]

This version: October 2018
PRELIMINARY AND INCOMPLETE

Abstract

Using seven decades of household level micro data, this paper makes the first systematic attempt to explain the household debt boom in postwar America. We document major changes in life-cycle profiles of household debt and loan-to-value ratios, driven mainly by housing debt. At the same time, home equity profiles remained constant across cohorts. We show that standard life-cycle savings behavior can explain a substantial part of the debt increase in a macroeconomic environment where house prices increase. Older households increase debt to remain on their targeted life-cycle wealth trajectories, while younger households have to take on more debt to purchase more expensive homes. Although net housing wealth remains constant, the extension of balance sheets endogenously leads to rising fragility and makes the economy highly sensitive to house price fluctuations.

*We thank organizers and participants of the “Tipping Points” Conference of the Pricate Debt Initiative, as well as Lukas Gehring for outstanding research assistance, and Stefania Albanesi, Christian Bayer, Karen Dynan, Olivier Godechot, Ethan Ilzetzki, Dirk Krüger, Felix Kubler, Costas Meghir, Atif Mian, Thomas Piketty, Morten Ravn, José-Víctor Ríos-Rull, and Gianluca Violante.

[†]University of Bonn, Adenauerallee 24-42, 53113 Bonn, Germany, alina.bartscher@uni-bonn.de

[‡]University of Bonn, CEPR, and IZA, Adenauerallee 24-42, 53113 Bonn, Germany, mokuhn@uni-bonn.de

[§]University of Bonn and CEPR, Adenauerallee 24-42, 53113 Bonn, Germany, schularick@uni-bonn.de

[¶]University of Bonn, Adenauerallee 24-42, 53113 Bonn, Germany, ulrike.steins@uni-bonn.de

1 Introduction

Why did American household debt rise from 15% in 1946 to more than 100% of income on the eve of the 2008 financial crisis? The dramatic surge of household debt in the postwar era has become a major topic in the American public debate and an active research agenda for macroeconomists. Understanding the interaction between household borrowing and the macroeconomic outcomes has become a central goal for the profession, but puzzles and debates still permeate the field (Zinman 2014). In particular, the absence of long-run micro data has hampered a more detailed analysis of the long-run evolution and the driving forces of American households' rising indebtedness.

This paper lets, for the first time, long-run micro-level data speak about the rise of household debt in the U.S. over nearly seven decades. We rely on a newly compiled data set that combines historical waves of the Survey of Consumer Finances (SCF) with its "modern" counterparts. Kuhn, Schularick, and Steins (2018) explain how the *Historical Survey of Consumer Finances* (HSCF) is compiled. Using this new dataset, we study the financial position of the cross-section of U.S. households. We organize the discussion around questions that have featured prominently in recent debates in economics and in the other social sciences.

Using the SCF data, we first present a number of new stylized facts that describe the patterns of the increase of U.S. household debt over the past seven decades:

- Sharply higher debt-to-income ratios, driven by mortgage borrowing and increasingly concentrated among households in 50th to 80th percentiles of the income distribution;
- Substantial increases in aggregate loan-to-value ratios for all income groups, with slightly higher increases in the middle and lower parts of the income distribution;
- Stable home equity positions over time and across the distribution.

The core hypothesis of this paper is that a combination of rising house prices and individually rational life-cycle savings behavior is able to produce the observed patterns and quantitatively matches the U.S. debt increase. Moreover, we argue that the same life-cycle optimizing behavior (a la "Modigliani") endogenously creates financial fragility that can erupt when debt levels are high (a la "Minsky").

In essence, when house prices rise prospective homeowners have to borrow larger amounts to buy more expensive houses. In other words, they start their savings path with expanded balance sheets for a given amount of housing equity. They will follow their life-cycle savings plan even if they enter the housing market at elevated price levels.

Existing homeowners make capital gains when house prices rise that lead to an increase in home equity. They are richer than planned when originally making their financial planning for old age. To smooth their consumption over time, households therefore need to liquidize some of the equity in the house to finance consumption. In other words, they engage in negative savings (“equity extraction”) after the deviation from their life-cycle wealth profile. Also note that the reason for the house price increase is irrelevant, as long as it was unexpected when making financial plans and is assumed to be permanent. Clearly, temporary increases would not affect the life-cycle wealth accumulation if they do not last until retirement.

However, unlike stocks or other assets, housing cannot be sold off in parts to reduce equity on the balance sheet. It is indivisible. However, what households can do is to raise additional debt, thereby lowering home equity and rebalancing the household portfolio. This mortgage equity withdrawal appears completely consistent with a optimizing life-cycle saver from the Modigliani life-cycle perspective. We will provide evidence that this simple mechanism is able to explain a large part of the debt increase in America post-WW2.

While rational on the individual level, this life-cycle smoothing has substantial implications for financial stability as larger balance sheets increase the sensitivity of the economy to house price fluctuations. The 2008 financial crisis has opened up a lively research agenda concerned with the effects of the composition of household balance sheets on macroeconomic activity (Mian and Sufi 2009, 2014; Jorda, Schularick and Taylor 2013), as well as the interactions between housing and credit markets (Guerrieri and Uhlig 2016). We use the micro data and quantify how rising leverage has increased financial vulnerabilities of individual strata of the income distribution.

We propose a quantitative assessment of household balance sheets akin to a stress test for banks. We “shock” households’ balance sheet with exogenous declines of house prices and construct a measure for the value of households’ home equity and the total value of mortgage debt at risk for an exogenous 10, 20 and 30% decline in nominal house prices.

We then track how these risk measures have evolved over time as the leverage of American

households has risen. We demonstrate that the vulnerability of the American economy to asset price shocks has increased substantially. In 1970, a 20% decline in house prices created negative home equity equivalent to about 1% of aggregate income. Today, the same drop in house prices leads to about 5% of household income or 600 billion Dollars of negative home equity in the system. Assuming all households with negative home equity defaulted on their mortgage, the total value of loans in default would exceed 20% of aggregate income, up from 5% in the 1980s.

The structure of the paper follows from the discussion above. We first introduce and discuss the historical SCF data and show that the micro-data match the aggregate trends closely. In the next step, we look at debt trends. In the next section, we show that the combination of rising house prices and life-cycle savings behavior is able to produce to quantitatively explain the surge in American household debt and match the debt profiles of birth cohorts over time. In the following section, we discuss the effects of rising household debt on financial fragility. The last section concludes.

2 Data

We use data from the Historical Survey of Consumer Finances (HSCF) — newly compiled microdata covering seven decades of U.S. financial history at the household level. Kuhn, Schularick, and Steins (2017) (KSS henceforth) describe how to compile the HSCF data from *historical* and *modern* waves of the Survey of Consumer Finances (SCF) data. We will briefly introduce the data and discuss how the survey data match the aggregate debt trends from the Flow of Funds. We complement the microdata with data, for example for house prices, from the *Macroeconomic History Database* (Jordà, Schularick, and Taylor (2017)).

2.1 Historical SCF Data

The SCF is a key resource for research on household finances. Data for the modern survey waves after 1983 are readily available for download from the website of the Federal Reserve. The comprehensiveness and quality of the SCF data explain its popularity among researchers (see Kuhn and Rios-Rull (2016) and references therein). However, the first consumer finances surveys were conducted much earlier, namely as far back as 1948. These historical SCF waves were conducted annually between 1948 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. KSS describe the details to extract, combine, and harmonize the historical and modern survey waves to construct the HSCF data. We rely on this work in the current paper.

A key advantage of the HSCF data is that it covers household-level information over the entire post-war period and provides in addition to financial variables detailed demographic information of households. Important for the current analysis, the HSCF data contain all variables needed to construct long-run series for the evolution of household debt including its subcomponents. For most of our analysis, we use total income, total debt, and the distinction in housing and non-housing debt from the HSCF data. Total income is 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. Total debt consists of housing and non-housing debt. Housing debt is debt on owner-occupied

housing.¹ 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 after their introduction and proliferation.² We abstain from any sample selection for our analysis to provide an analysis of a representative sample of U.S. households over time. We follow KSS and pool data over three surveys for the first part of the sample when annual surveys exist.

2.2 FFA and HSCF

Household surveys often struggle to match aggregate data when the microdata is aggregated to the level of the macroeconomy. We compare the aggregate trends in income and household debt in the HSCF to trends in aggregate data from the National Income and Product Accounts (NIPA) and the Flow of Funds (FoF). In many cases, measurement concepts differ between household surveys and macrodata, explaining at least in part why even high quality microdata do not correspond to aggregate data. We remove therefore level difference that we attribute to different measurement concepts and compare growth trends over time (see KSS for further discussion and references). We choose the period from 1983 to 1989 as our base period and remove level differences based on this period. During this period, the HSCF data correspond to 84 percent of NIPA income and 86 percent of FoF debt.³

Figure 1 shows the comparison of growth trends between the HSCF and aggregate data throughout the 1950-2016 time period. Overall, the aggregate data and the aggregated microdata show similar trends. Comparing the data series for income, housing, and non-housing debt highlights the strong increase in household debt. With respect to housing debt, the HSCF data and the FFA match almost perfectly. Non-housing debt also aligns well with the FFA data, albeit the increase in the HSCF data appears slightly stronger than in the FFA.

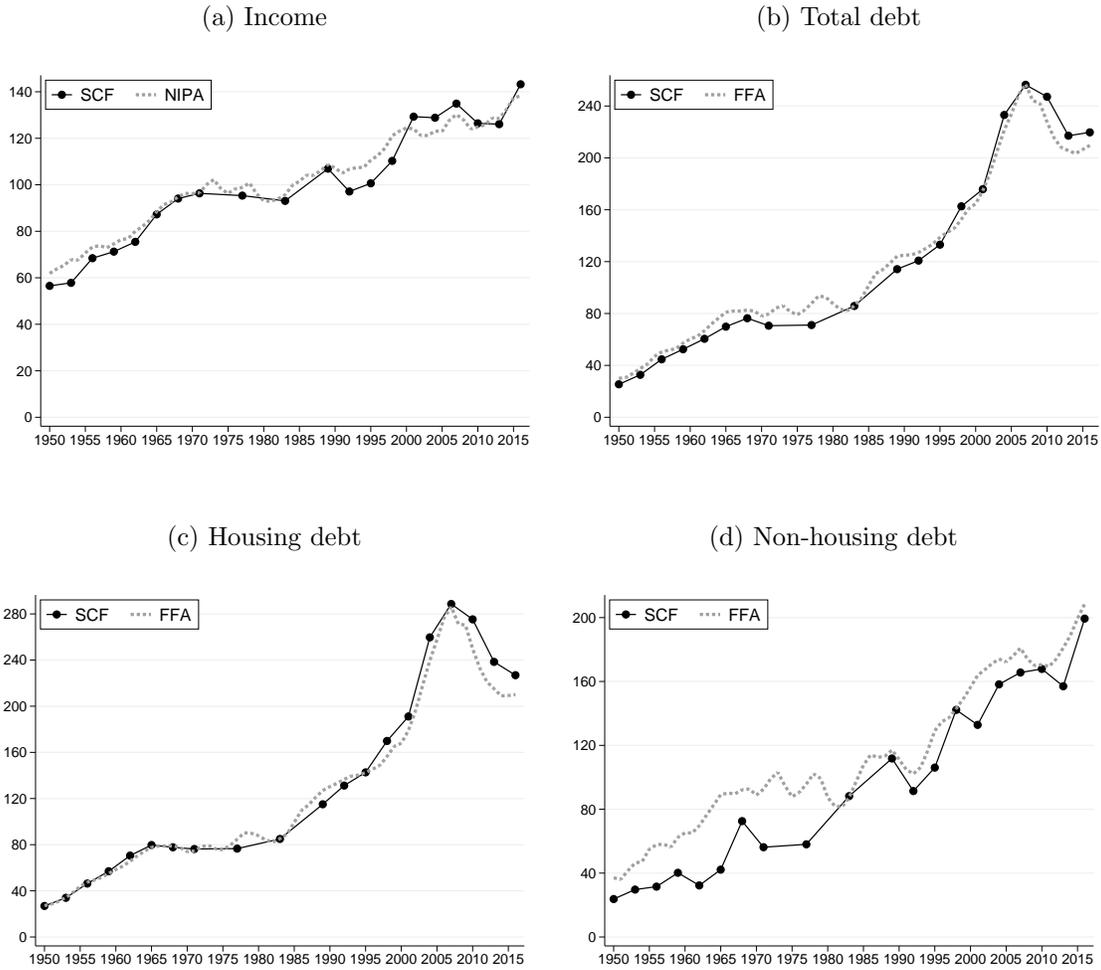
¹We treat investment in non-owner-occupied housing like business investment and only use the net position when calculating wealth. Other real-estate debt typically accounts for about 10% of total household debt. See Figure XXX in Appendix YYY.

²Note 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.

³Income components of the NIPA tables that are included are wages and salaries, proprietors income, rental income, personal income receipts, social security, unemployment insurance, veterans benefits, other transfers, and the net value of other current transfer receipts from business. Mortgages and consumer credit are included as FoF debt components.

All in all, the close alignment in growth trends effectively alleviates concerns that the microdata systematically misses parts of the distributional changes underlying the observed macroeconomic growth trends. KSS demonstrate further that the HSCF data also closely match aggregate trends for other major asset series. Relying on the HSCF microdata, the next section explores for the first time the distributional changes in household debt in the United States over the long run.

Figure 1: HSCF, NIPA, and FFA: Income and Total Debt



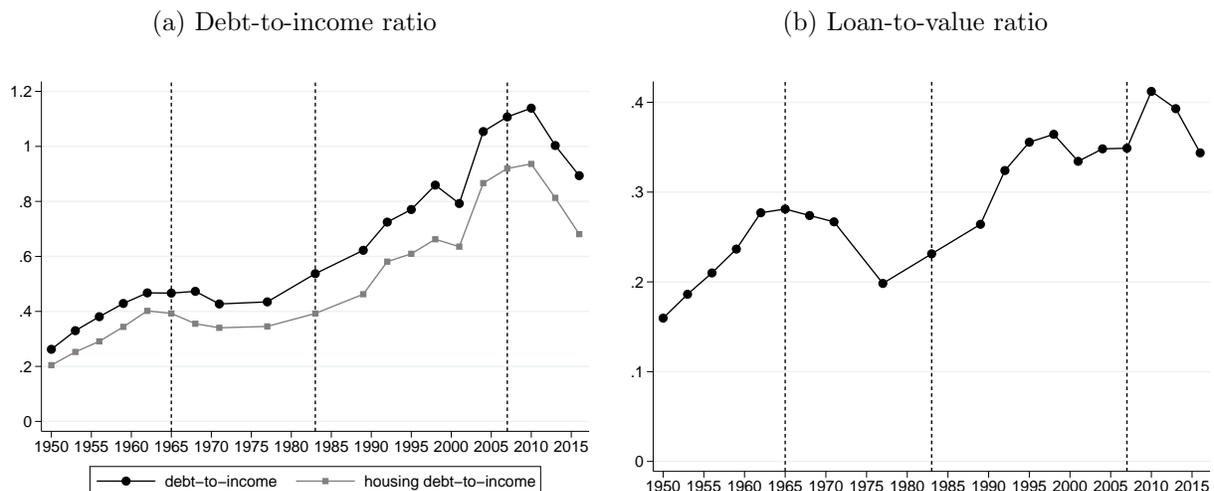
Notes: Income and total debt data from SCF in comparison to income data from NIPA and total debt data from FoF. All data has been indexed to the 1983 - 1989 period (= 100). SCF data is shown as red lines with circles, NIPA and FoF data as a blue dashed line. Over the index period, SCF values correspond to 80% for income, 86% of total debt, 93% of housing debt, and 70% of non-housing debt.

2.3 Aggregate trends

Figure 2a shows the evolution of debt-to-income and debt-to-asset ratios over the last seven decades. Debt-to-income ratios more than quadrupled between 1950 and the 2007 crisis. They have fallen by about 20pp. since the crisis. Housing debt accounts for 76% of the increase in the debt-to-income ratio from 26% to 89% between 1950 and 2016.

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

Figure 2: Aggregate debt-to-income and loan-to-value ratios



Notes: Time series of debt-to-income ratio for total debt and housing debt over time. Horizontal axis shows calendar time and vertical axis level of the debt-to-income ratio. Vertical lines indicate end and starting points of debt booms used in analysis.

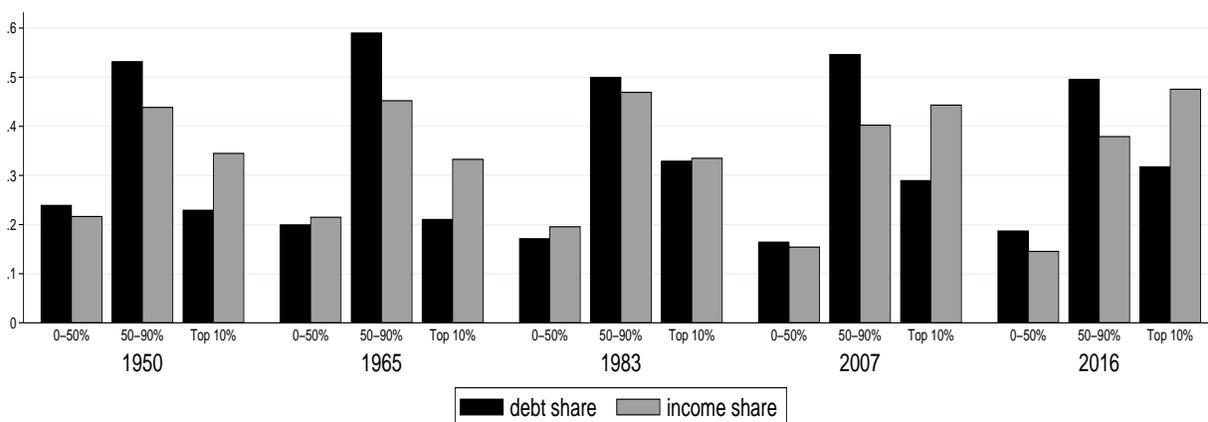
2.4 Distribution

How debt is household debt distributed among rich and poor households and how has this distribution evolved over time? For distributional questions, we will typically sort households by income and divide households into three groups whenever it is impractical to show

the entire distribution. The first group are the bottom 50% of households in the income distribution, the second group covers households between 50th and 90th percentile, and the third group refers to the top 10% of households. We have also studied income quintiles, but found that the key trends can be well represented across these three groups.

Figure 3 reports the share of total debt owed by the different income groups. To put these numbers into perspective, we also report the income share of each of the groups in total income.

Figure 3: Shares in debt and income along the income distribution



Notes: Shares in total debt and income for different income groups along the income distribution for different years. The income groups are the bottom 50% of households in the income distribution, the middle 50-90%, and the top 10%. The different data years are shown below the labels for income groups. The black bars show for each income group the share of aggregate household debt that is owed by this group. The gray bars show the share of aggregate income received by each of the income groups.

Two facts are striking. First, the richest households in the top 10% of the income distribution owe a large share of total debt. Like their income share, the debt share increased over time. While the top 10% owed one out of five dollars of aggregate debt in 1950, this increased to one out of three dollars in 2016. At the same time, their income share increased from 35% to 48%. Second, the share of debt owed by the middle class (50% - 90%) remained quite stable over the past seven decades.

The middle class owes typically between 50% and 60% of all household debt in the United States. Taken together the middle class and the top 10% of the income distribution owe with the exception of 1950 always more than 80% of all household debt. This high concentration of debt in the upper half of the income distribution and the stability of this debt share over

time speaks strongly against soaring household debt at the bottom of the income distribution as the driver of the aggregate debt increase.

Table 1 shows the shares of different income groups in aggregate housing and non-housing debt, respectively. The Table suggest that, if anything, a greater share of aggregate debt is owed by the upper echelons of the income distribution today. The share of housing debt of the top 10% income households increased from 23% in 1950 to 34% in 2013.

Table 1: Shares in aggregate housing and non-housing debt

(a) Shares in aggregate housing debt

year	0-50%	50-90%	Top 10%	90-95%	95-99%	Top 1%
1950	19.9	57.1	23.0	11.4	8.5	3.2
1965	17.8	59.7	22.5	10.4	8.3	3.8
1983	15.3	54.0	30.7	12.4	13.2	5.0
1998	14.5	52.7	32.8	13.2	14.2	5.4
2007	14.1	55.1	30.8	12.4	12.8	5.6
2016	15.3	50.1	34.6	14.4	13.7	6.5

(b) Shares in aggregate non-housing debt

year	0-50%	50-90%	Top 10%	90-95%	95-99%	Top 1%
1950	37.0	37.2	25.7	10.4	12.1	3.2
1965	31.4	55.2	13.3	5.9	4.6	2.8
1983	22.0	39.0	39.0	8.9	15.2	14.9
1998	26.3	48.6	25.1	10.2	7.7	7.2
2007	27.8	52.4	19.8	8.7	7.5	3.6
2016	29.5	47.9	22.6	7.4	9.2	6.0

3 The American household debt boom, 1949-2016

This section begins by providing stylized facts for the growth and distribution of household debt and its two major components, housing debt and non-housing debt, over the past 70 years. After that, we will take closer look at the time line of the debt increase and identify four different phases of the debt boom.

3.1 Stylized facts

This section shows that the following core facts characterize the growth of household debt in the U.S. since World War 2:

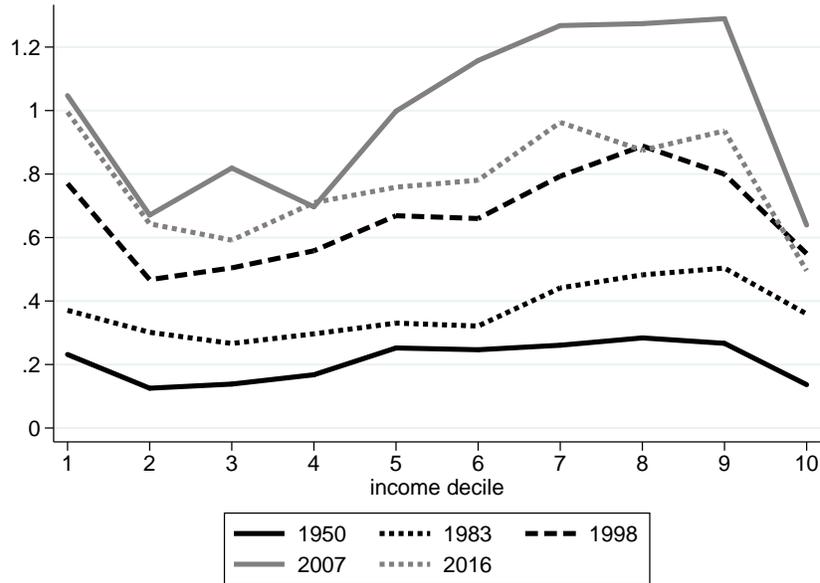
- Sharply higher debt-to-income ratios, driven by rising housing debt and increasingly concentrated among households in 50 to 80th percentile of the income distribution;
- Substantial increases in aggregate loan-to-value ratios, with somewhat faster increases for households in the middle and lower part of the income distribution;
- Stable home equity positions over time and across the distribution.

3.1.1 Debt and income

How have debt-to-income ratios evolved over the distribution? Figure 4 shows debt-to-income ratios along the income distribution for the different SCF waves. The indebtedness of households was relatively evenly distributed in 1950 with mean debt-to-income being less than 40% across all income groups. By 1983, debt-to-income ratios had increased somewhat both at the lower and upper ends of the distribution, with relatively little change in the middle.

This change markedly since the 1980s. Indebtedness has not only risen strongly across all income groups, but the strongest increase took place in the middle of the distribution between the 50th to 70th percentile. For households in these income brackets, housing-debt-to-income ratios quadrupled from 30% to 120% within a 25 year period.

Figure 4: Debt-to-Income Ratios

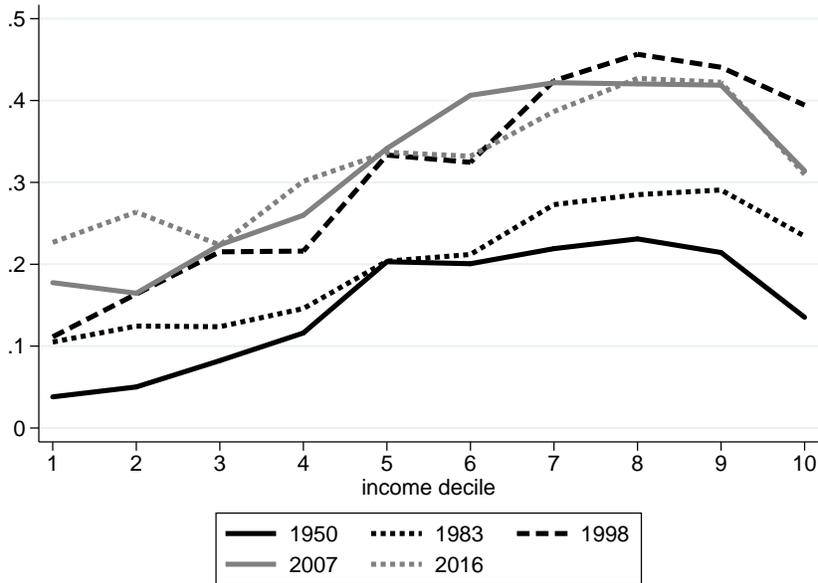


Notes: The graph shows the evolution of average housing-debt-to-income ratios by deciles of the aggregate income distribution for the HSCF waves 1950, 1965, 1983, 2007 and 2016.

3.1.2 Debt and assets

Housing debt was the major driver of the household debt increase. Figure 5 studies loan-to-value ratios along the income distribution at different points in time. The strongest increase in loan-to-value ratios occurred since 1983. Once more, leverage rose most dramatically in the middle of the distribution.

Figure 5: loan-to-value ratios



Notes: The graph shows the evolution of average loan-to-value ratios by deciles of the aggregate income distribution for the HSCF waves 1950, 1965, 1983, 2007 and 2016.

3.1.3 Home equity

Does higher indebtedness and higher loan-to-value ratios mean that American households are poorer today than they were before?⁴

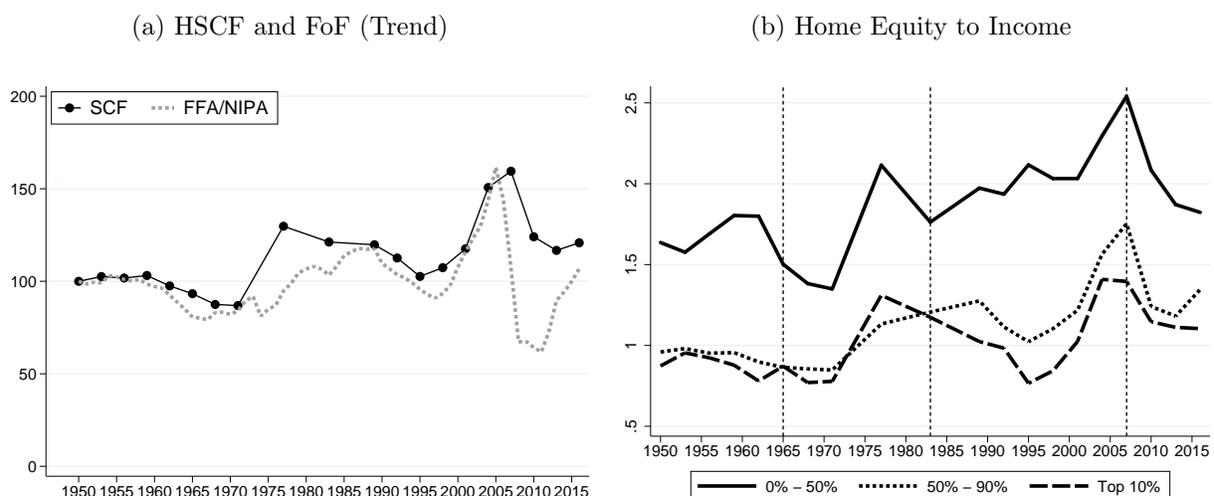
Figure 6a shows the evolution of the home equity-to-income ratio from the HSCF data and from aggregate FoF and NIPA data. It reveals a striking pattern. Seventy years of skyrocketing debt levels have left the value of home equity relative to income on the household balance sheet largely unaffected. Moreover, as Figure 6b shows, this stability holds across the income distribution.

In other words, attempts to understand the increase in American household debt will have to account for the fact that debt-to-income and loan-to-value ratios have increased, but home-equity-to-income ratios have remained stable. The American debt boom led to a substantial

⁴Note that loan-to-value ratios are not sufficient to answer this question. For example, a household with a loan-to-value ratio of 70% for a house worth 100,000 Dollars has home equity of 30,000 Dollars whereas a household with a lower loan-to-value ratio of 50% on a house worth only 50,000 Dollar has home equity of 25,000 Dollars. The household with the higher loan-to-value ratio is wealthier.

extension of household balance sheets, but the (net housing) wealth of Americans barely changed.

Figure 6: Home equity to income over time



Notes: The left panel presents the ratio of home equity (housing net of housing debt) from the Flow of Funds (FFA) relative to income from the National Income and Product Accounts (NIPA). The right panel compares this aggregate ratio of home equity to income with SCF data, indexing both ratios to the 1980s.

3.2 A chronology of the American debt boom

The previous section has identified the core stylized facts of the household debt boom. In this section, we provide a more detailed chronology of the debt build-up, including a decomposition into the intensive and extensive margin of household debt.

3.2.1 Decomposition

Using information on the number of household with outstanding debts and their debt levels, we can decompose the increase of debt in its extensive and intensive margin and ask: to what extent has the total number of indebted households increased, or have already indebted households taken on more debt? We calculate the ex- and intensive margin effects separately for housing and non-housing debt.

More precisely, $d_{i,t}$ stands for the mean debt-to-income ratio of income group i . $s_{i,t}^{H+}$ is the share of households having positive housing debt, i.e., the extensive margin, and $d_{i,t}^{H+}$ is the mean 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 of 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 calculated by:

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

The first part of this expression is the change in household indebtedness due to a change in the extensive margin of housing debt. In other words, by how much household indebtedness would have risen if only the extensive margin of housing debt, $s_{i,t}^H$, 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. a change in household indebtedness due to an increase in d_t^H if the extensive margin of housing debt had been constantly at the level of period t and all non-housing debt components being at the level of period $t - 1$. The third and fourth parts are the respective effects of non-housing debt.

Table 2 shows the ex-an intensive margin effects of the increase in aggregate debt-to-income ratios between 1950 and 2016. Overall, we find that 33pp of the 82pp. increase in household debt-to-income ratios was driven by the intensive margin of housing debt, and 22pp. due to the extensive margin of housing debt. The remaining 25pp. are due to non-housing debt. Mortgage lending hence played the dominant role for the increase in household debt.

Table 2: Decomposition of the increase in aggregate debt-to-income ratios between 1950 and 2013

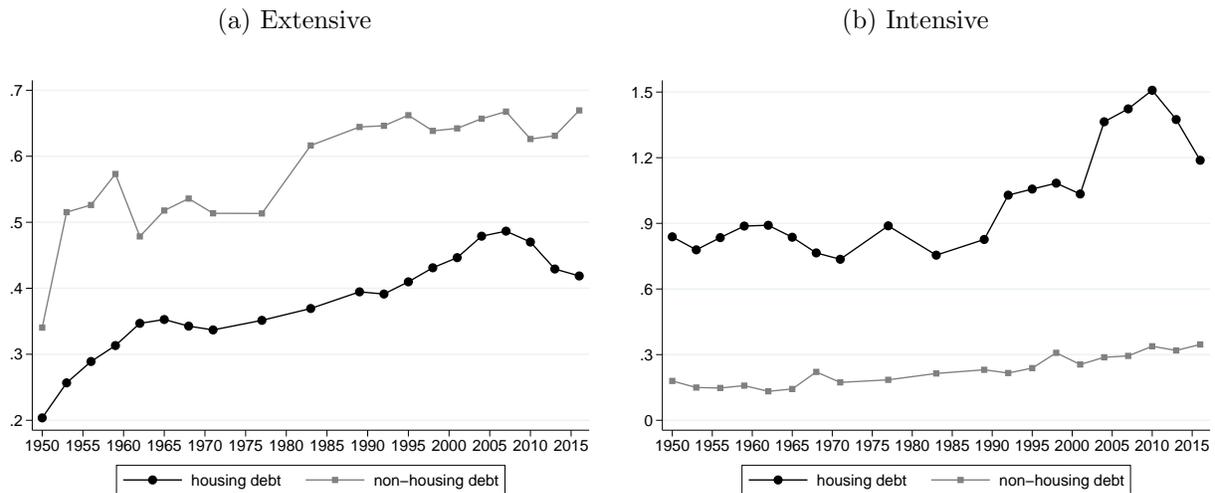
Extensive margin	housing debt	21.7
	non-housing debt	9.6
Intensive margin	housing debt	33.3
	non-housing debt	17.8
total increase		82.4

Notes: Percentage point change in aggregate debt-to-income between 1950 and 2016.

Figure 7 shows both margins of indebtedness of American households over time and for

both types of debt, i.e., housing and non-housing, and across the distribution. The extensive margin captures the share of households with positive debt balances and the intensive margin reports the debt-to-income ratio for households owing debt. The extensive margin expands from 1950 to 1965, while the intensive margin of housing debt remains flat. This is in stark contrast to the debt boom after 1983 when the intensive margin increases strongly, while the extensive margin follows the increase in homeownership rates in the early 2000s but returns almost to its 1983 level by 2016.

Figure 7: Extensive and intensive margin of debt-to-income



Notes: The left panel shows the share of households with positive housing debt (black line with dots) and positive non-housing debt (gray line with squares). The right panel shows the (non-)housing debt-to-income ratio of households with positive (non-)housing debt.

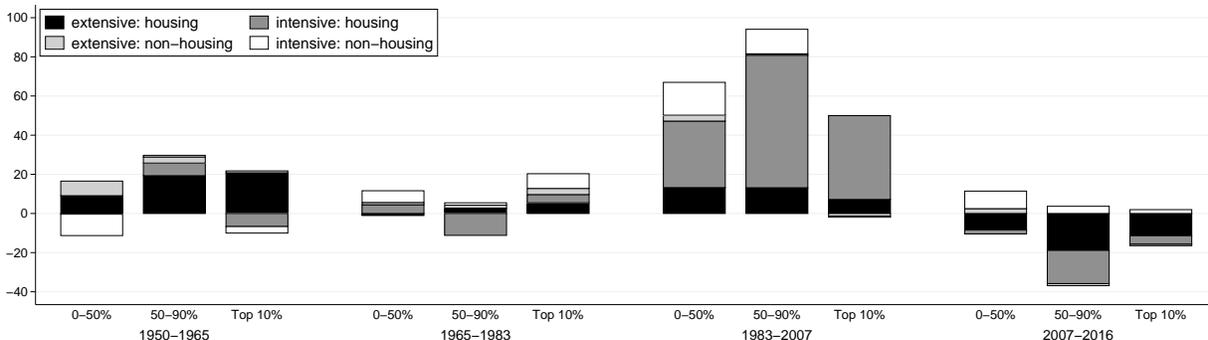
3.2.2 The four phases

We next split the cross-section into three groups according to their income and look at changes in debt-to-income ratios between 1950, 1965, 1983, 1998, 2007 and 2013. The black and dark gray bars are the ex- and intensive margin effects of housing debt between two time periods. The light gray and white bars are the corresponding effects of non-housing debt. The sum of all four bars is the total percentage point change of mean debt-to-income ratios between two time periods.

This decomposition in ex- and intensive margin effects reveals that different forces caused

variations in debt-to-income ratios over time. We identify four distinct periods in the American debt build-up.

Figure 8: Decomposition of changes in debt-to-income ratios over time



Notes: Observations with debt-to-income ratios and debt-to-asset ratios above 50 in absolute value were excluded.

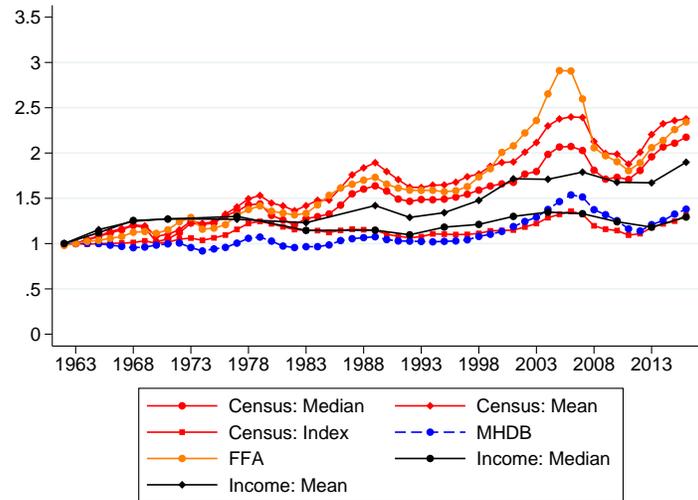
- The postwar home ownership boom, 1950-1965:** The first period is characterized by the rise of homeownership rates after the World War II between 1950 and 1965. During this period, debt-to-income ratios increased strongly and almost doubled within 15 years. Looking at loan-to-value ratios for housing debt, we also find a strong increase in the same ballpark within 15 years. The increase between 1950 and 1965 was mainly driven by higher extensive margins of housing debt, i.e., more people took on housing debt.
- Stability, 1965-1983:** The second period covers the years from 1965 to 1983 with almost stable debt-to-income ratios and even slightly falling loan-to-value ratios.
- The 25-year boom, 1983-2007:** The year 1983 marks the beginning of the second large debt boom in the United States from 1983 to 2007. Debt-to-income ratios now more than doubled within 25 years and loan-to-value ratios increased by about 60%. In contrast, both the rise from 1983 to 1998 and 1998 to 2007 was mainly driven by higher intensive margins of housing debt. An exception are poor households – their increase in debt-to-income ratios between 1983 and 1998 was mainly due to higher intensive margins of non-housing debt.
- Crisis and deleveraging, 2007-2016:** The final period covers the financial crisis and its aftermath from 2007 to 2016 with falling debt-to-income ratios and a hump

in loan-to-value ratios. In 2016, debt-to-income and loan-to-value ratios still are at historical highs and substantially above their 1983 level.

4 House prices and housing debt

House prices have risen in recent decades both in the U.S. and globally (Knoll, schularick and Steger 2017). Figure 9 shows the trajectory of house prices in the U.S. since 1949. For the U.S., the house price increase amounts to approximately 50%, with most of the increase occurring in the post-1983 period.

Figure 9: Comparison of House Price and Income Growth (CPIAUCNS)



The mechanism how rising house prices lead to increasing home equity on the household balance sheet is the following. An unanticipated and permanent increase in house prices boosts the housing wealth of home-owning households.

Unless households decide to pass on more wealth to the next generation, the rise in house prices and the associated increase in home equity will affect households life-cycle savings and consumption decisions. A typical households can be expected to adjust life-cycle wealth accumulation by extracting home equity using additional housing debt. At the same time, younger cohorts have to increase debt-to-income to buy more expensive homes from older cohorts.

Consider a simple example of a household with income of 1 Dollar, a house of 5 Dollar, and housing debt of 3 Dollar. The leverage ratio is 60% and the home equity-to-income ratio is 2. Now, consider the simplest case where house prices increase by 20% so that the household now owns a house worth 6 Dollars. To balance home equity to income the household has

to raise 1 Dollar of additional debt, i.e., to extract all the additional equity. The leverage ratio on the additional equity is 100% and the resulting leverage ratio for housing increases to 67% ($= 4/6$).

Yet is a 50% increase in house prices quantitatively sufficient to explain the observed quadrupling of household debt-to-income ratios? A simple accounting example helps illustrate the empirical relevance of this mechanism. Household income is given by Y_t , and all variables are expressed ratios to income. Constant quantity of houses $H = 1$ at price of house P_t . Households use debt D_t and equity E_t to finance a house, and housing demand grows in line with income:

$$HP_t = D_t + E_t$$

Period 0: $P_0 = 1$, $E_0 = 0.8$, $D_0 = 0.2$

Period 1: $P_1 = 1.5$, $E_1 = 0.8$, $D_1 = 0.7$

In other words, as a result of a 50% increase in house prices, the debt-to-income ratio increases from 0.2 to 0.7 (3.5-fold), and the loan-to-value ratio increases from 0.2 to 0.5 (2.5-fold). Clearly, we have not picked these numbers at random. the chosen example matches the increase in U.S. housing debt observed between 1950 and 2016. In our example, loan to value ratios increased from 0.2 to 0.5 over the same period. This matches again the long-run increase as observed in the HSCF data. We will show below that the cohort profiles over time also correspond closely to this mechanism.

Note that in our example income increases do not affect these conclusions as long as housing demand H and equity perfectly scale with income. In other words, people will increase the quantity of housing in line with income so that housing expenditure to income ratios remain constant in the long-run. This, too, matches the U.S. postwar evidence.

4.1 Evidence on equity extraction

Standard life-cycle behavior suggests that when house price rise unexpectedly and permanently, home-owning households adjust their home equity – be it through HELOCs or cash-out refinancing. Yet does the available evidence confirm that the amounts of equity withdrawals observed in the data are quantitatively large enough to account for the substantial increase in housing debt in postwar America?

Direct estimates of home equity extraction are difficult for the 1980s as data availability is poor. We address the question by relying on two quantitatively plausible estimates.

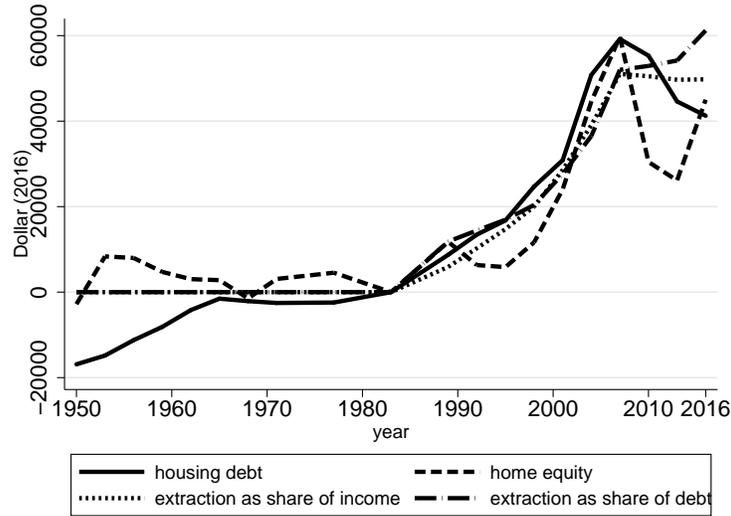
Figure 10 shows the increase in average housing debt since 1983 in real Dollars. Between 1983 and 2016 average housing debt per household increased by 40,000 Dollars. We start by estimating the potential for home equity extraction for existing home-owners by multiplying the 1983 stock of housing assets by the change in the house price level. The change in asset values relative to 1983 represents the additional home equity that is a result of the price increase. The average increase per household is also 40,000 Dollars, and the two series correlate closely except for a short time period in the late 1990s.

Yet this only shows the total amount of additional home equity due to the price increase. What is a realistic estimate of home equity extraction and what share of the debt increase does it explain? We consider two scenarios. In the first scenario, we let households extract home equity equivalent to 1% of their annual income each year after 1983 and 2% of their annual income starting in 1992, corresponding to the estimates of Greenspan and Kennedy (2005). For the period after the financial crisis, we let households accumulate home equity (deleveraging) at 1% of their annual income per year.

In a second scenario, we construct an estimate for equity extraction using data from the *Cash Out Refinancing Report* by Freddie Mac (Q2 2013). The data report the share of refinancing households who increase their unpaid principal balance by at least 5%, the median age of the refinanced loan, and the median appreciation of the refinanced property. The data are annual, starting in 1985. We combine this data with information on annual inflation rates and the share in aggregate debt of households 35 years and older. We take this age threshold as a proxy for the group of “older” households who are homeowners and have capital gains triggered by house price appreciation.⁵ These are our estimates for extraction rates of

⁵More precisely, we proceed as follows. We first derive the inflation and duration adjusted median

Figure 10: Housing debt, home equity, and home equity extraction counterfactuals



Notes: Mean housing debt increase, capital gain in home equity from house price increase, and two counterfactual scenarios for home equity extraction relative to 1983. The estimate for capital gains in home equity is derived by multiplying the stock of housing in 1983 by the house price index. The two counterfactual simulations for home equity extraction show the accumulated equity extraction if households extract a share of annual total income per year or a part of the outstanding debt balance per year. All values are expressed as changes relative to the level in 1983 and equity extraction is assumed to be zero in 1983 and before. All values are expressed in 2016 dollars using the CPI.

existing homeowners. To account for the fact that younger households have to raise more debt to buy houses in the presence of rising house prices, we add to our estimate for equity extraction a component from new homeowners.⁶ We add these two effects to the estimates

appreciation. For this, we multiply the median duration of the loan D_t by the inflation rate π_t and subtract this from the median appreciation of the refinance property Q_t . This provides us with an estimate of the real capital gain on the property R_t

$$R_t = Q_t - D_t\pi_t.$$

We then use the median duration of the loan to determine the share of loans that is refinanced $s_t = \frac{1}{D_t}$ and multiply this share by the share of loans that increase their loan balance e_t and the share of housing debt held by households 35 years and older h_t . The extraction rate as share of the outstanding loans is then

$$x_t = h_t \times s_t \times e_t \times R_t$$

The HSCF data is not annual and the estimate for x_t in 1989 is particularly high. We replace therefore the extraction rate in 1989 by the average estimate for x_t between 1985 and 1989, that cuts the estimated extraction rate in 1989 in half. We assume that extraction rates are constant between survey dates and the high estimate for 1989 would otherwise apply for the entire time period between 1983 and 1989 which is not in line with the data.

⁶Note that in this case, the young pay the house price to the old who realize the capital gain on their property which is in the aggregate equivalent to equity extraction of old households. We therefore include this effect in our equity extraction estimate. We assume that demographic turnover leads to transactions

for the equity extraction of older households x_t described above. We get the annual extracted home equity amount by multiplying this extraction rate with the outstanding housing debt.

Figure 10 shows the Dollar estimates for equity extraction under the two scenarios. We find that both ways of constructing the data account for the increase in housing debt over time. They show some divergence in the aftermath of the financial crisis when households delevered and houses went into foreclosure. A refined estimate for equity extraction/delveraging quantifying these effects more closely would likely match the decline after 2007 much better. Figure 11 shows Figure 10 from above but normalizes all Dollar values by annual income, so that they correspond to percentage point changes in debt-to-income ratios over time. Finally, Figure 12 shows the estimated total effect for equity extraction and the component that comes from demographic turnover and the increase in homeownership rates during the 2000s. The figures shows that about 40% of the debt increase due to this broader definition of equity extraction results from demographic turnover in the presence of rising house prices and rising homeownership during the 2000s. Before 2000, this effect is negligible.

for one out of 45 households in the economy (buying at age 30 and selling at age 75 on average). These households have to put up the addition debt to account for the price increase. We derive the estimate as above for home equity by multiplying the stock of housing assets in 1983 by the house price change since 1983. This provides the home equity series from above. Multiplied by $\frac{1}{45}$, we take this as our estimate for debt increase due to demographic turnover. Note that with constant house prices, this estimate is always zero. Finally, we also account for the increase in homeownership after 2000. We assume that each year starting in 2001 3.3% of the existing stock of houses is taken out as additional debt for homeownership. This corresponds to 1% annual growth in the housing stock together with a 33% loan-to-value ratio. A loan-to-value ratio of 33% seems low but it is relative to the average housing value in the economy and the increase in homeownership happened for less expansive houses.

Figure 11: Housing debt, home equity, and home equity extraction counterfactuals

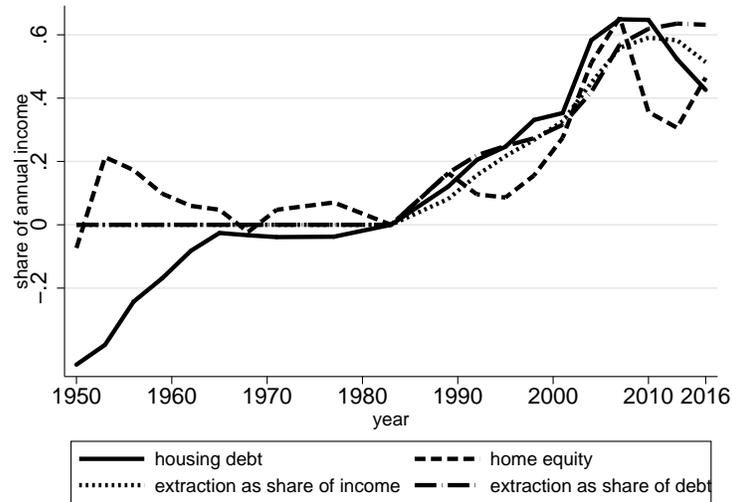
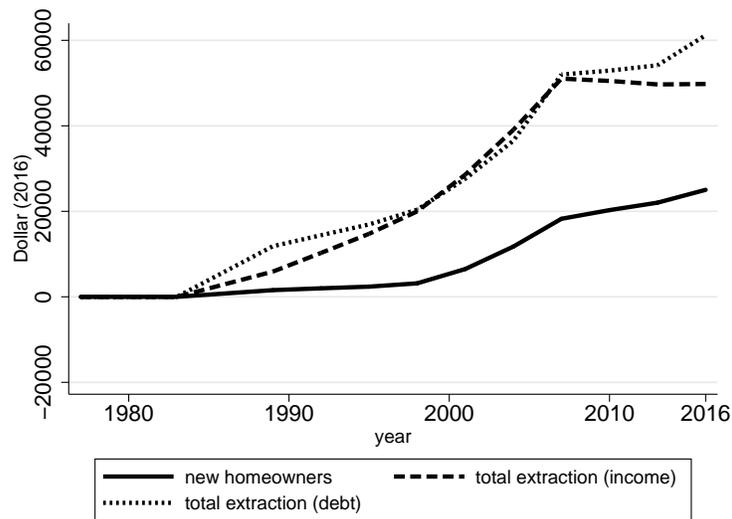


Figure 12: Equity extraction from new homeowners



4.2 A stylized model of household debt

This section develops a stylized model of the life-cycle of household debt to explore if the observed changes in the life-cycle of household debt can be rationalized from economic theory. The model builds on the life-cycle hypothesis focusing on the intertemporal smoothing of consumption in the presence changing income over the life-cycle. The model remains stylized and abstracts from savings motives due to idiosyncratic income risk by assuming perfect foresight.

Our stylized model is designed to focus on the link between the asset value of houses and housing debt. A key challenge is to generate debt dynamics with decreasing debt-to-income and loan-to-value ratios over the life-cycle. Our model uses a combination of market frictions, refinancing costs for long-term mortgage contracts, and a bequest motive to generate the stylized features of life-cycle debt profiles from the data.

Households live for J periods and work for $J - J_R$ periods before they retire. While working, a household of age j receives income y_j . Households face no income risk. Each household maximizes discounted life-time utility using time-discount factor β and period utility function $u(c, h) = \log(c) + \phi \log(h)$ where c denotes non-durable consumption and the service flow from housing is proportional to the stock of housing h . We abstract from a choice on housing size and assume each household is endowed with a house of size h . Housing does not affect household decisions except for its role as collateral for housing debt and we drop it as an argument from the utility function below. The household portfolio at age j is composed of (liquid) financial assets a_j , housing debt d_j , and the asset value of housing ph where p is the price of housing. Household wealth is the consolidated value of the household balance sheet $w_j = ph - d_j + a_j$. The period budget constraint reads

$$a_{j+1} + c_j + d_j = d_{j+1} - rd_j + y_j + a_j - \varphi(d_j, d_{j+1})$$

where r is the mortgage interest rate and there is no interest on financial assets.⁷ The function $\varphi(d_j, d_{j+1})$ captures mortgage refinancing costs.

In each period, households decide about their financial asset holdings and consumption. Mortgage debt follows a fixed amortization schedule $d_{j+1} = d_j - \alpha$. Opportunities to adjust

⁷Financial assets represent transaction accounts. This assumption for the return difference follows the recent literature on liquid-illiquid portfolio choice where liquid assets have a lower return.

mortgage debt arise stochastically with probability ρ . If the household has the opportunity to refinance, the household must decide to either extract a share θ of home equity at costs $\varphi(d_j, d_{j+1})$ or to follow the fixed amortization schedule with payment α . If households extract equity, debt next period is $d_{j+1} = d_j + \theta(ph - d_j)$.

We do not allow for uncollateralized borrowing ($a_j \geq 0$) and also abstract from collateral constraints. Instead, we assume that households have a bequest motive so that all households want to hold nonnegative wealth in the last period of life. We use bequest utility $v(w) = w$ if $w \geq 0$ and $-\infty$ otherwise. We formulate the household problem in recursive form (for any period $j < J$), we get

$$\begin{aligned}
V_{no}(a, d, j) &= \max_{a'_{no} \geq 0} u(c_{no}) + \beta \rho V_{adj}(a'_{no}, d - \alpha, j + 1) + (1 - \rho) V_{no}(a'_{no}, d - \alpha, j + 1) \\
c_{no} &= y - \alpha - rd + a - a' \\
\tilde{V}_{adj}(a, d, j) &= \max_{a'_{adj} \geq 0} u(c_{adj}) + \beta \rho V_{adj}(a'_{adj}, d', j + 1) + (1 - \rho) V_{no}(a'_{adj}, d', j + 1) \\
c_{adj} &= \theta(hp - d) - rd + y + a - \varphi(d, d') - a'_{adj} \quad d' = (1 - \theta)d + \theta hp \\
V_{adj}(a, d, j) &= \max\{\tilde{V}_{adj}(a, d, j), V_{no}(a, d, j)\}
\end{aligned}$$

The goal of the model is to illustrate the link between the asset value of housing ph and household debt. We abstain from calibrating the model but parametrize the model to highlight the economic mechanism. We set parameters so that the model is roughly in line with the empirical evidence. We set $J = 60$ and $J_R = 40$. For refinancing costs, we assume a constant function $\varphi(d, d') = \phi$ and set $\phi = 0.01$. We set the mortgage interest rate r at 6%, the time discount factor β to 0.95, the extraction share for equity $\theta = 0.5$, and the refinancing probability to 20%. Given our assumption on the initial portfolio, we set $\alpha = 0.065$ to approximate a 30-year fixed rate mortgage with constant amortization payments. For the initial household portfolio, we use $p = 1$ and $h = 1$. We set initial debt holdings to match a housing debt-to-income ratio of 1. We approximate life-cycle income pattern using a linearly increasing life-cycle profile starting at 0.9 ending at 1.1 during working life and set retirement income to 0.9. This parametrization constitutes our baseline model.

Figure 13 shows the baseline model and two counterfactuals with a 10% and 20% higher house price. When changing house prices, we also lower mortgage interest rates by 1pp and 2pp respectively to roughly compensate for higher interest rate payments. The results remain qualitatively unaffected without this adjustment but become harder to compare given

differences in average income.⁸

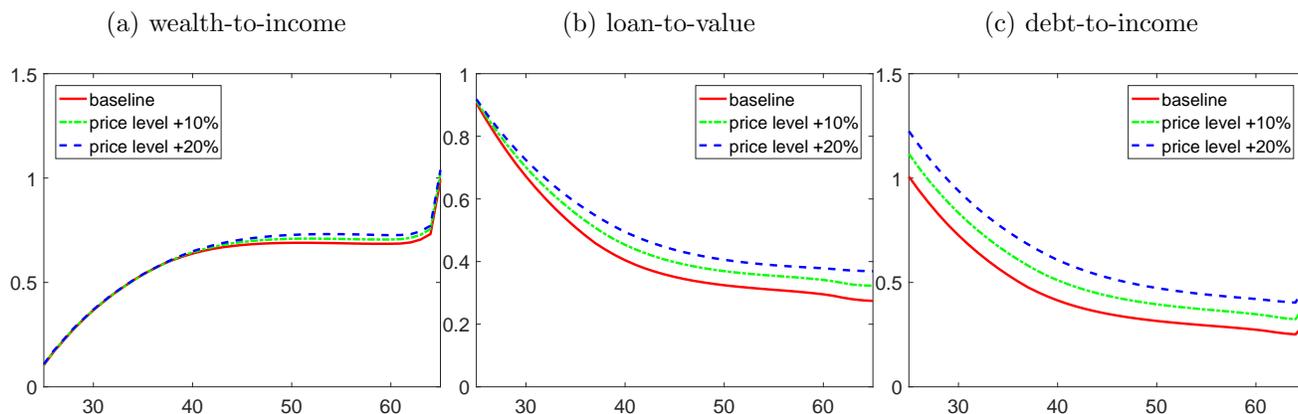


Figure 13: Life-cycle profiles for wealth, loan-to-value ratio, and debt-to-income ratio

Looking at figure 13, we see that households react to higher house prices by adjusting their debt position to balance life-cycle wealth. While debt-to-income and loan-to-value profiles both shift up, wealth accumulation remains virtually unaffected. The change in the slope of the loan-to-value ratio results from the fact that we do not adjust amortization rate α leaving amortization-to-income ratios unaffected.

The previous experiment compared three households who live in economies with different values of houses over their entire lifecycle. Figure 14 considers the case when a household is surprised by a house price increase in the middle of the lifecycle. For this experiment, we simulate the first part of the life-cycle using policy functions assuming a constant price level over the entire lifetime. After the shock, households reoptimize given the current state and use the policy corresponding to the new price level for their remaining lifetime. Note that in this case, housing debt d_j and financial assets a_j remain unaffected on impact. We set the refinancing probability in the period of the house price shock to 1 so that all households can refinance in that period. In line with the data, we assume 10-year-birth cohorts so that households between ages 35 and 44 are affected and we average over their life-cycle profiles.

House price shocks in the middle of the lifecycle result in profile changes that correspond the observed turning of debt-to-income and loan-to-value profiles observed across cohorts in

⁸The interest rate change only roughly accounts for the effect of higher interest rate payments. Falling mortgage interest rates are in line with the evidence of 30-year fixed rate conforming mortgages over the past decades.

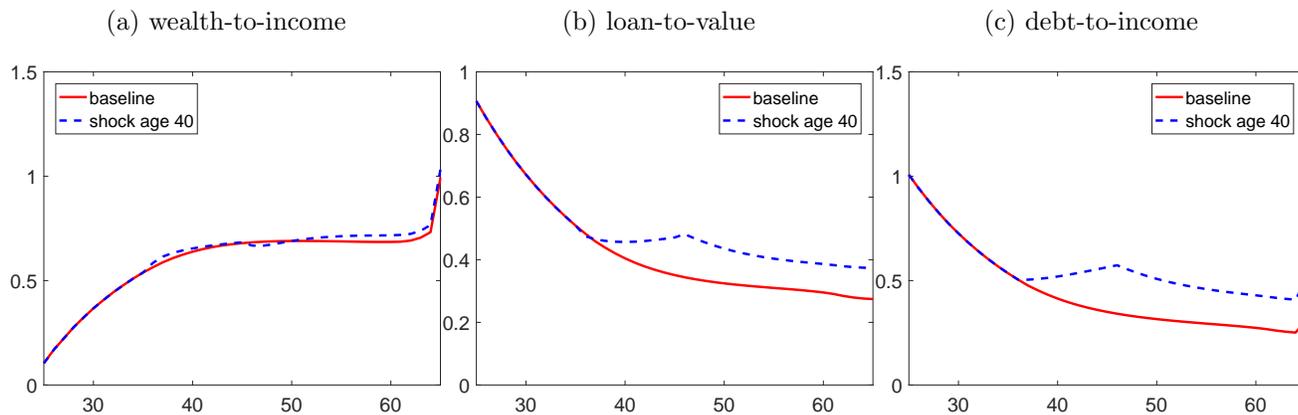


Figure 14: Life-cycle profiles for wealth, loan-to-value ratio, and debt-to-income ratio

the data. Again, the wealth profile remains virtually unaffected.

Taken together, figures 13 and 14 rationalize the changing life-cycle profiles of household debt in the presence of constant life-cycle wealth profiles. Changing house prices and the property that housing can serve as a collateral can at least qualitatively rationalize the observed changes in the life-cycle of household debt without resorting to behavioral explanations or households who enter on unsustainable debt paths to explain the changing pattern of household indebtedness over time. While the model can explain the changing debt pattern based on optimal individual behavior, the macroeconomic consequences from the changing debt pattern are not part of the model. We explore the macroeconomic consequences for financial stability in the next section.

4.3 Life cycle profiles of debt

The price effect might provide a good approximation in the aggregate but the question is whether it also holds across cohorts? Do cohorts that have been more strongly affected by increasing house prices show more pronounced changes in their life-cycle profiles? This section will show that the answer is likely affirmative.

4.3.1 Birth cohorts

Comparing the life-cycle debt profiles of the different cohorts, we uncover a secular change in the life cycle of debt of U.S. households. Debt profiles both shifted upwards and turned across cohorts leading to marked changes in the life-cycle debt pattern over time. We argue that house price increase account for a substantial fraction of the observed life cycle profile change of cohorts.

In a second step, we trace these changes in life-cycle debt profiles back to differential trends in house price and income growth over recent decades. We provide supporting evidence to argue that rising macroeconomic indebtedness of U.S. households is the reaction to fast growing house prices. Such a pattern is consistent with life-cycle savings dynamics targeting a fixed wealth-to-income profile over the lifecycle.

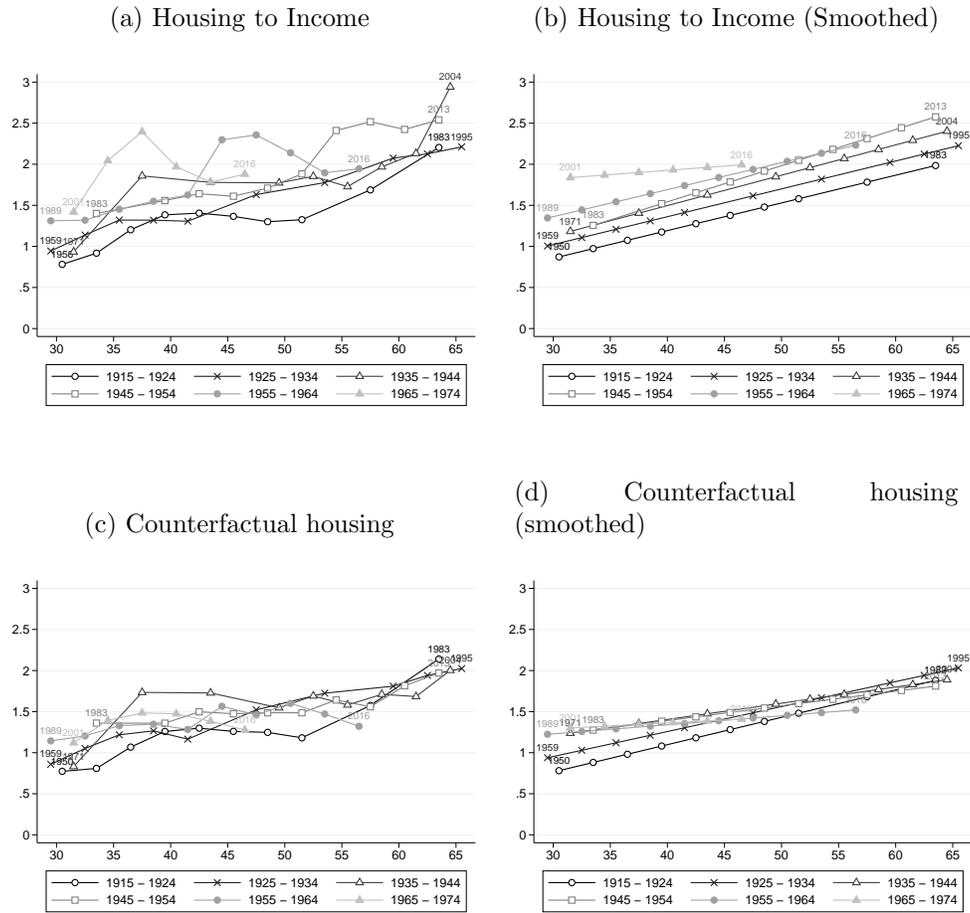
The HSCF data is not a panel and to follow cohorts of American households over time we rely on the construction of synthetic birth cohorts from the different survey waves. We group households in 10-year birth cohorts. The long time span covered by the HSCF data is important for such an analysis as it allows us to track several of these birth cohorts over time.

We will use households with heads born between 1915 and 1924 as our oldest cohort and households with heads born between 1955 and 1964 as our youngest cohort. Our oldest cohort is on average 30 in 1950 and our youngest cohort is on average 53 in 2013.⁹

For these cohorts, we trace the life-cycle profiles of housing-to-income ratios, housing debt-to-income ratios, and loan-to-value ratios from age 30 to age 65, except for the youngest cohort that we only observe until households are in their mid-50s.

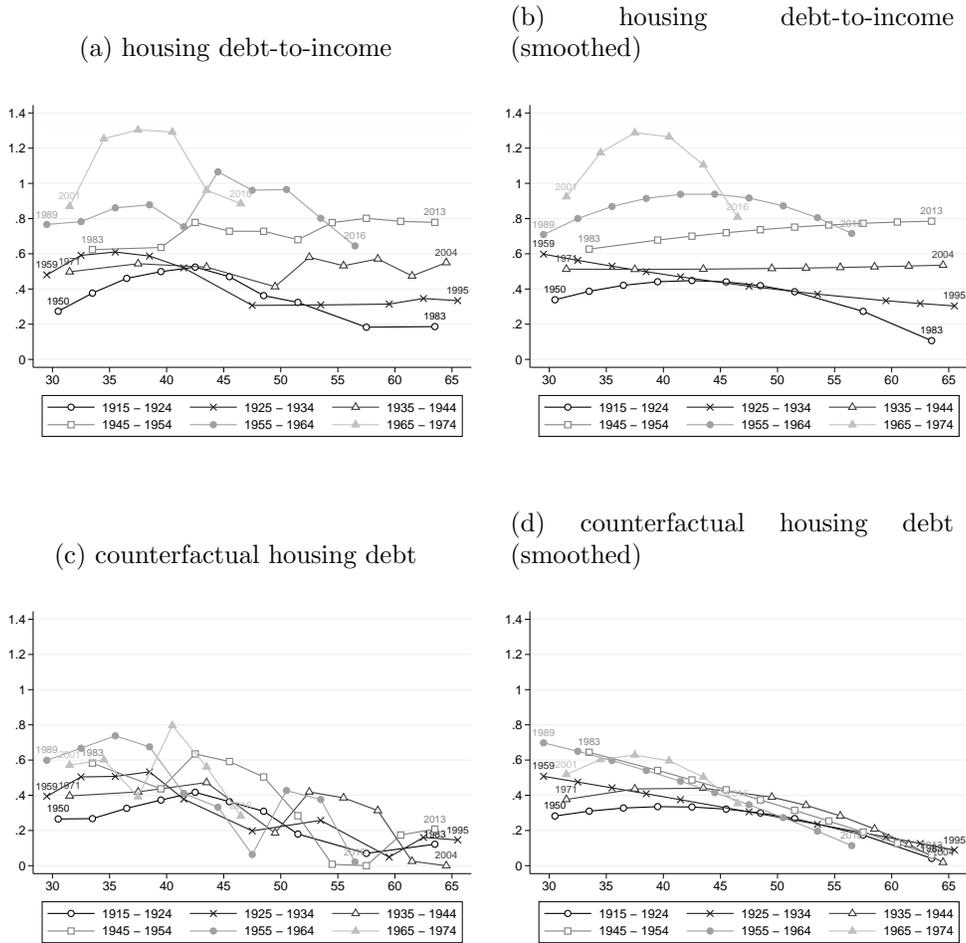
⁹We show additional cohorts in the appendix.

Figure 15: Life-cycle housing-to-income ratios



Notes: The graph shows average housing-to-income ratios for the synthetic cohorts at different stages of their life-cycle. Cohort 1 was born between 1915 and 1924, cohort 2 between 1925 and 1934, cohort 3 between 1935 and 1944, cohort 4 between 1945 and 1954, cohort 5 between 1955 and 1964, and cohort 6 between 1965 and 1974. The areas shaded in black represent the counterfactual average housing-to-income ratios if households from all cohorts had experienced the house prices of 1950 at all ages.

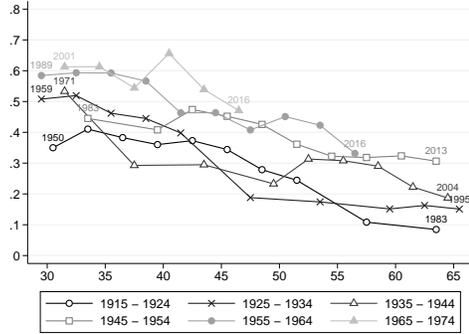
Figure 16: Ratio of housing and non-housing debt-to-income ratios



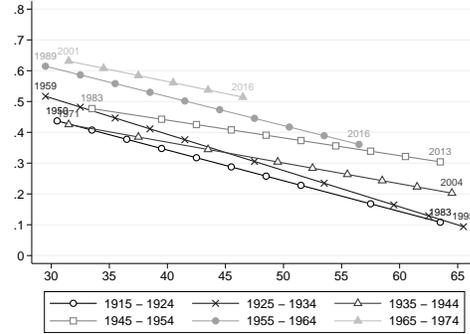
Notes: The graph shows average housing-debt-to-income ratios for the synthetic cohorts at different stages of their life-cycle. Cohort 1 was born between 1915 and 1924, cohort 2 between 1925 and 1935, cohort 3 between 1935 and 1944, cohort 4 between 1945 and 1954, cohort 5 between 1955 and 1964, and cohort 6 between 1965 and 1974. The areas shaded in black represent the counterfactual average housing-debt-to-income ratios if households from all cohorts had experienced the house prices of 1950 at all ages.

Figure 17: Loan-to-value ratios

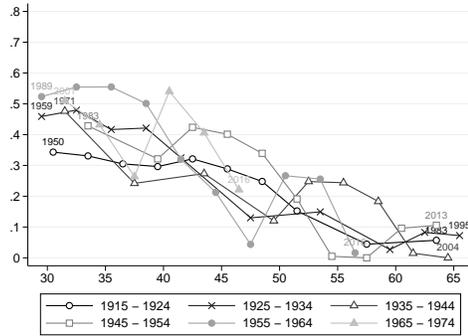
(a) Loan-to-value



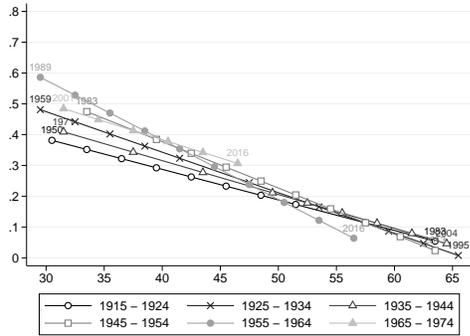
(b) Loan-to-value (smoothed)



(c) Counterfactual Loan to Value



(d) counterfactual Loan to Value (Smoothed)



Notes: The graph shows average loan-to-value ratios for the synthetic cohorts at different stages of their life-cycle. Cohort 1 was born between 1915 and 1924, cohort 2 between 1925 and 1934, cohort 3 between 1935 and 1944, cohort 4 between 1945 and 1954, cohort 5 between 1955 and 1964, and cohort 6 between 1965 and 1974. The areas shaded in black represent the counterfactual average housing-debt-to-income ratios if households from all cohorts had experienced the house prices of 1950 at all ages.

4.3.2 The aging of debt

Older households have profited most from rising house prices. In line with the hypothesis that life-cycle savings behaviour explains a substantial part of the debt increase, we expect to see that older households still hold substantial debt at the end their working lives. One way to show this is to look at the *average age of a dollar of debt*, over time.¹⁰

Figure 18a compares the average age of a dollar of debt to the average age of households over time. We see that since the late 1980s the average age of a dollar of debt has been rising more quickly than average household age. This shows that debt holdings of older households increased more quickly than debt of younger households. We index age to 1971 and remove level differences. We find that the debt shifting towards older households is substantial so that the average age of debt increases twice as fast after 1990. While the average household age increased by four years between 1992 and 2013, the average dollar of debt aged by eight years highlighting the shift of debt towards older households.

We have seen that across cohorts loan-to-value ratios at the end of working life increased (Figure 17). Figure 18b shows that much of this shift happened along the extensive margin of housing debt. While for more than 30 years after 1950 about 10% of households 65 and older had housing debt, this number tripled over the past three decades so that today a third of U.S. households in retirement still owe money on their home. For working-age households, the share of households with mortgage debt follows trends in home-ownership rates. It strongly rises between 1950 and 1965, it stays roughly constant afterwards up until the mid 1990s when it increased further. Over the entire post-war period, the share of U.S. households of working age with housing debt doubled from 25% to 50%.

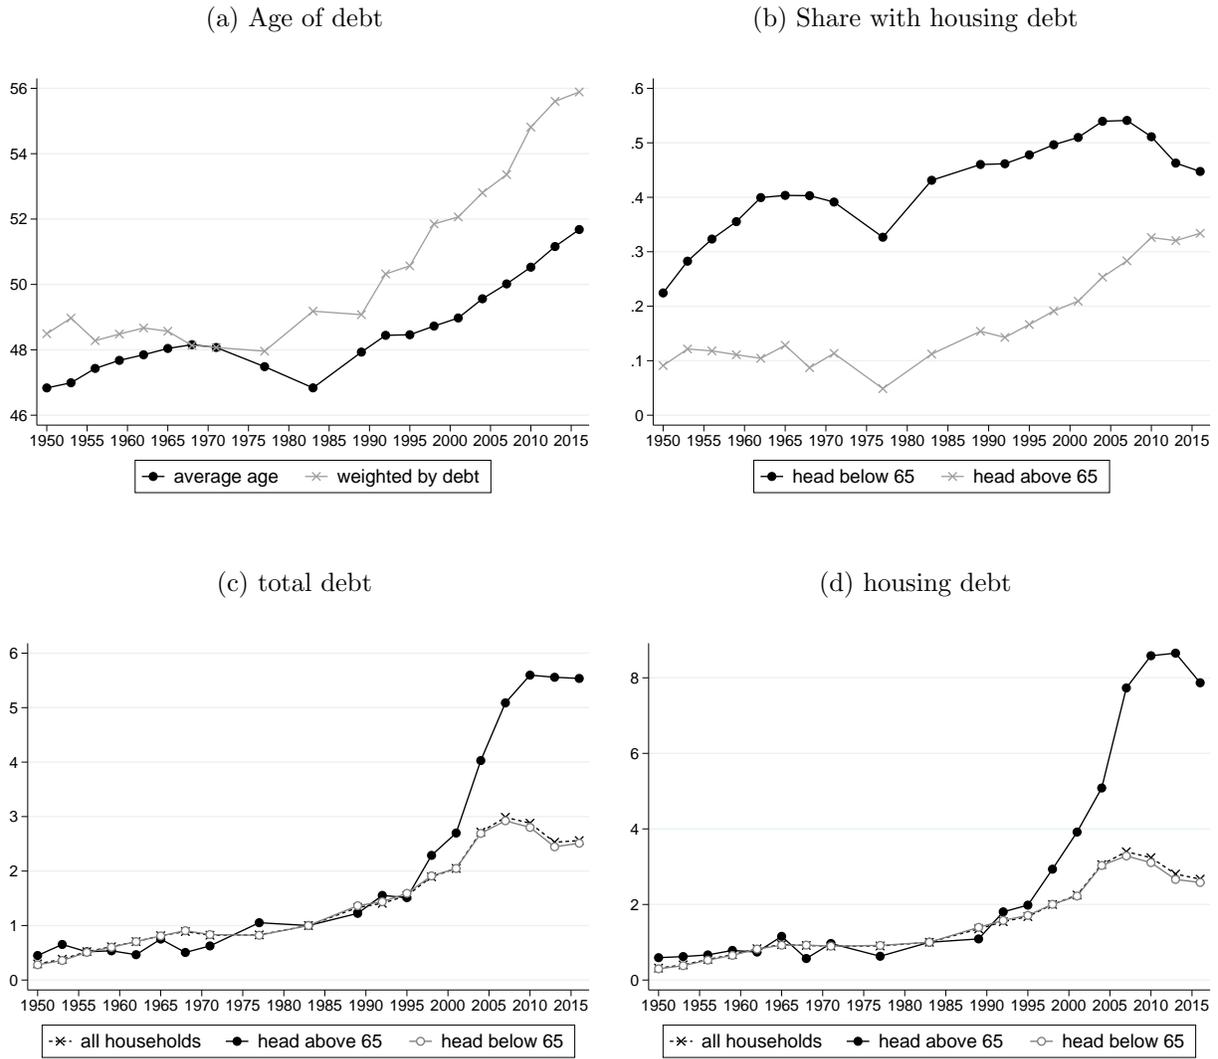
The old-age debt boom in the housing market since the 1980s can also be seen when looking at their debt share over time. Their share almost increased fourfold from less than 4% to almost 14% between 1983 and 2013. The same development can be seen when looking at newly originated mortgages. The past two decades have witnessed a boom in America's old-age debt, driven by home equity extraction.

Figure 18a also shows that debt largely grew in parallel for all four groups until the mid 1990s, but debt growth of households with heads older than 65 increased much more strongly

¹⁰The average age of a dollar of debt is the debt-weighted average household age.

since then. While average debt increased roughly 2.5 times from 1983 to 2007, debt of old households more than quadrupled. This divergence becomes even more extreme when looking at housing debt. Old households in 2007 held more than 8 times as much debt as in 1983 while the average household increase debt by a factor of slightly more than 3 and young households' debt not even tripled. In 2013, debt of young households just slightly more than doubled relative to 1983 while total debt tripled due to skyrocketing debt of the old.

Figure 18: Aging of debt



Notes: Mortgages include first and second mortgages, as well as other non-HELOC home equity loans. New mortgages are those taken out within the last 3 years.

5 Household Debt and Financial Fragility

Household debt increased substantially over the past decades. We demonstrated that this debt increase can be traced back to house prices rising in excess of incomes and the desire of American households to balance their home equity position. We have shown that the resulting housing net position did not change across cohorts but the required additional leverage of households might have substantially decreased the household sector’s resilience towards financial shocks, and thereby increased the financial fragility of the macroeconomy.

The 2008 financial crisis clearly demonstrated how a drop in house prices can lead to sizable amounts of negative equity in the system and trigger widespread defaults on mortgage contracts. The role of leveraged asset price fluctuations on household balance sheets and their knock-on effects on consumer spending have been studied intensively in recent years (Mian and Sufi (2009, 2011)). Empirical evidence shows that households have used mortgage equity withdrawals to fund consumption during the housing boom and then aggressively cut back spending when their financial position deteriorated during the housing bust.

In this last part of the paper, we aim to understand the effects of the increase of household debt on macroeconomic and financial fragility. We adopt an approach common for the banking sector and do *stress testing* of household balance sheets. We will exogenously “shock” households with lower house prices and then track the amount of negative equity and the share of negative net wealth households over time. When “shocking” both house prices and income, we follow the empirical evidence on foreclosure and household default and assume a *double trigger*: households have to hold negative home equity and a debt-service-to-income ratio in excess of 50%.

We consider two different scenarios. In the first, we consider a drop in house prices, in the second we consider a correlated drop in house prices and incomes of households effectively mimicking a large recession with drops in asset prices and rising unemployment.

We start the discussion in Figure 19 by showing the amount of home equity at risk if house prices dropped by 10, 20 or 30% respectively. More precisely, we calculate the absolute value of home equity that becomes negative following the house price drop. Repeating the exercise for every survey year yields a time series of home equity “under water” that is conditional on the asset price decline.

Put differently, home equity at risk measures the amount by which the difference between house value and outstanding mortgage debt changes after the fall in house prices. We also see the time line for home equity risk if house prices dropped by 20% for different age and income groups, allowing us to identify the effects on specific parts of the population.

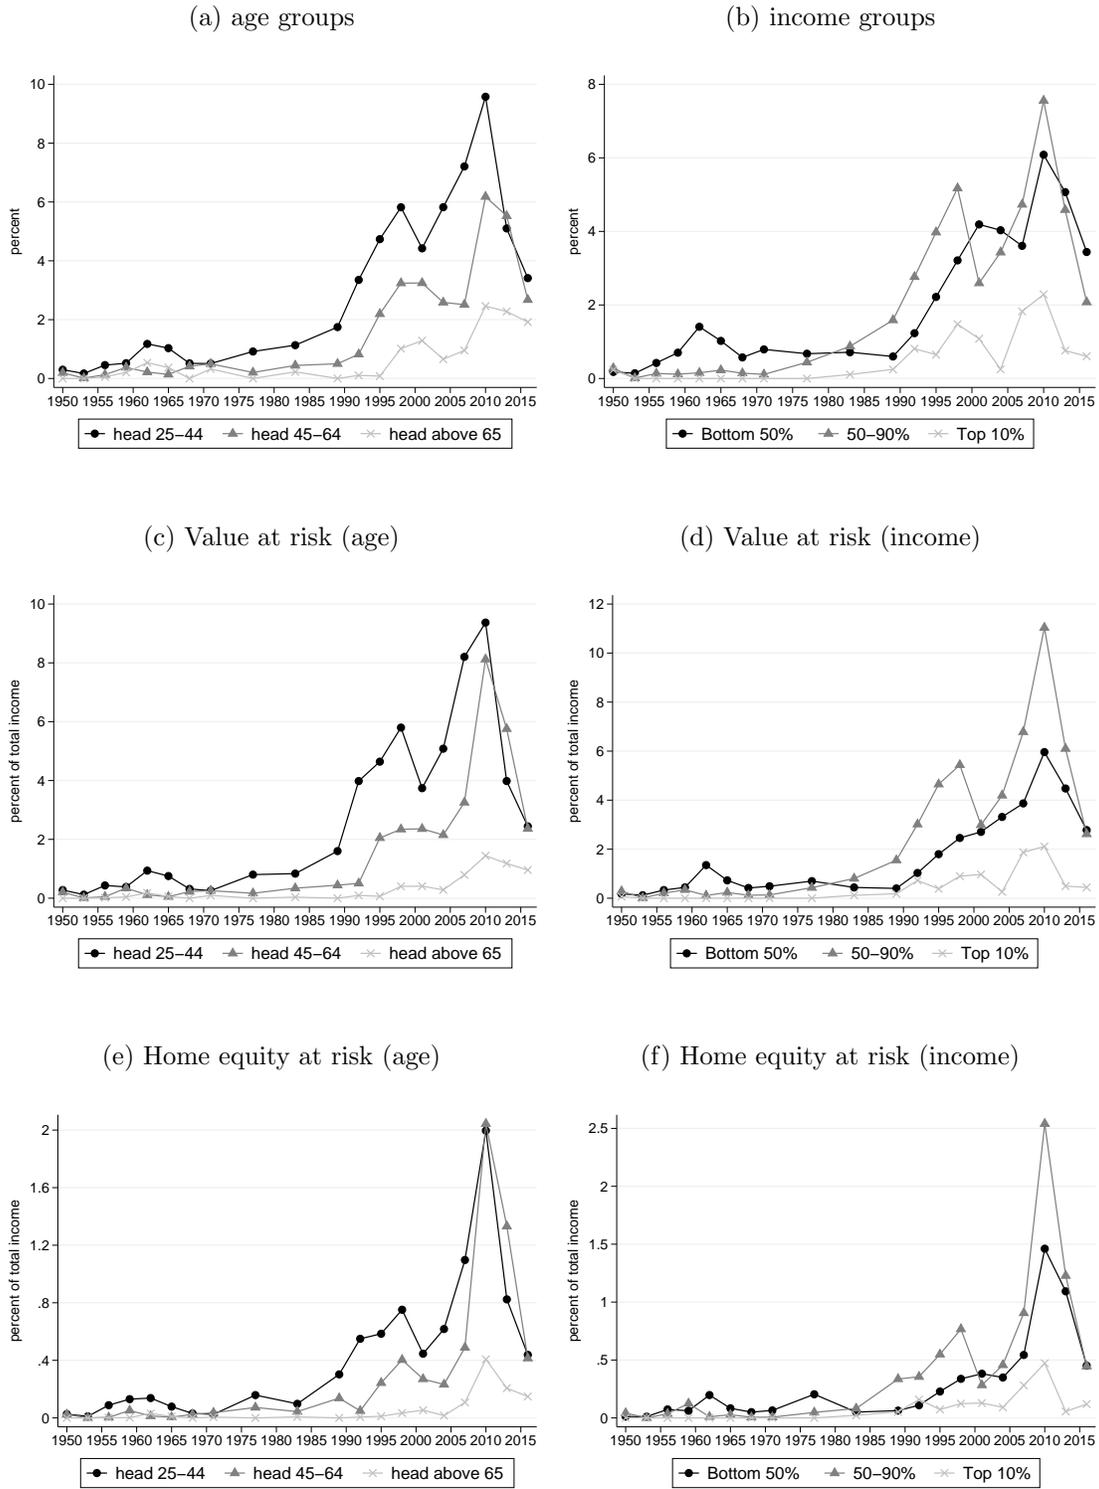
The charts demonstrate how much more sensitive U.S. households have become to house price fluctuations. While a 20% drop in house prices was associated with a drop in home equity equivalent to about 1.5-2% of aggregate income until the 1990s, the sensitivity is now more than three times as high. In 2013, a 20% drop in house prices would have led to a negative home equity of about 6% of aggregate income.

Owing to higher leverage and more expensive houses, the absolute losses are highest for the middle class (50-90% of the income distribution), but equally high for the bottom 50% relative to income. It is worth pointing out the high dollar amounts since these factor into aggregate spending. In 2013, a 20% house price drop raises the amount of system-wide negative equity to approximately 800 billion dollars.

What is the value of non-performing loans if all households with negative equity were to default on their mortgage payments? We also see the value at risk after a drop in house prices. Instead of adding up home equity, we take the sum of mortgages of households for which home equity has become negative after the drop in house prices.

It is important to stress that this number does not correspond to actual losses of the financial system and hence cannot be directly compared to loss-absorbing bank capital. However, the resulting amount of “problem loans” in the financial system if house prices dropped by 20% would climb to 3 trillion dollars in 2013, equivalent to 30 % of total income.

Figure 19: Stress testing household balance sheets



Notes: Share of households with negative home equity & mortgage payments-to-income >50% (20% drop in house prices, 50% drop in income), Value at risk after 20% drop in house prices & 50% drop in income. Households are assumed to be “at risk” if they have negative home equity and a debt-service-to-income ratio exceeding 50%.

6 Conclusions

This paper dissects the increase in household debt in the U.S. since WW2. Relative to income, household debt has risen by a factor of six. Yet empirically and theoretically the drivers of this process remain poorly understood. Our paper closes this gap of knowledge on the empirical side. For the first time, we use household level data to document the evolution of U.S. household debt, its composition and distribution.

Our findings provide new and important guiding principles for future theoretical research on household portfolio choices. The interaction between house prices and debt appears to be central to gaining a complete and nuanced understanding of the surge in household debt since WW2. At the same time, our study speaks to and quantifies the financial stability risks in highly-leveraged economies where even comparatively small changes in asset prices are magnified and can quickly turn into substantial losses for households and financial intermediaries.

References

- JORDÀ, Ò., M. SCHULARICK, AND A. M. TAYLOR (2017): “Macrofinancial history and the new business cycle facts,” *NBER Macroeconomics Annual*, 31(1), 213–263.
- KUHN, M., AND J.-V. RIOS-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(1).
- KUHN, M., M. SCHULARICK, AND I. U. STEINS (2017): “Wealth and Income Inequality in America, 1949-2013,” Discussion paper, CEPR Discussion paper 12218.
- MIAN, A., AND A. SUFI (2009): “The consequences of mortgage credit expansion: Evidence from the US mortgage default crisis,” *The Quarterly Journal of Economics*, 124(4), 1449–1496.
- (2011): “House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis,” *American Economic Review*, 101(5), 2132–56.