

THE GREAT AMERICAN DEBT BOOM, 1949-2013*

Moritz Kuhn[†] Moritz Schularick[‡] Ulrike I. Steins[§]

September 8, 2017

Abstract

The American economy experienced a dramatic increase in household debt since World War II. Relying on newly compiled archival micro data from historical waves of the Survey of Consumer Finances (SCF) going back to 1949, this paper makes the first systematic attempt to dissect the ascent of household debt in postwar America. We show that debt-to-income ratios have risen similarly across income groups and that debt growth since the 1970s occurred mainly on the intensive margin of housing debt. A quantitative assessment of household balance sheets demonstrates that financial vulnerabilities of different strata of the income distribution have risen substantially.

*We thank Lukas Gehring for outstanding research assistance, and Christian Bayer, Karen Dynan, Olivier Godechot, Felix Kubler, Atif Mian, Thomas Piketty and the participants of the Tipping Points Conference on June 22, 2017 for helpful comments and suggestions. Schularick gratefully acknowledges support from the Private Debt Project. Steins gratefully acknowledges financial support by the Wissenschaftsförderung der Sparkassen-Finanzgruppe. The usual disclaimer applies.

[†]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

In 1946, gross household debt stood at 15 percent of U.S. gross domestic product. It peaked at close to 100 percent on the eve of the 2008 financial crisis and has fallen slightly since then. This six-fold rise of household debt has become a major topic in the American public debate and an active research agenda for macroeconomists interested in household finance. In light of recent experience, understanding the interaction between household borrowing and the macroeconomic outcomes has become a central challenge for the discipline (Mian and Sufi 2016). However, puzzles still outnumber stylized facts in the rapidly growing research 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 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 (2017) 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 three core questions that have featured prominently in recent debates in economics and in the other social sciences.

The first set of questions relates to the cross-sectional evolution of household debt since WW2. We first ask if more households have taken on debt (growing financial inclusion) or if households have taken on more debt relative to their income (the intensive margin)? Has the increase in household debt been evenly distributed across the population, or have certain strata of the population have accumulated more debt, possibly connected to the increasing polarization of incomes? Note that aggregate statistics such as the well-known Flow of Funds Accounts (FFA) published by the Federal Reserve do not allow to study such distributional questions. Micro-level data are needed to understand the changing distribution of household debt over time.

In his 2009 best-selling book *Fault Lines*, Raghuram Rajan put forward the idea that rising debts and rising income inequality were closely linked. Households hit by stagnant real incomes increasingly relied on debt to finance consumption – be it out of sheer necessity or to “keep up with Joneses.” Kumhof and Ranciere (2015) recently presented a model where

higher leverage arises endogenously in response to a growing income share of high-income households. Such a nexus between socio-economic pressures and growing household credit has been an important research theme in other disciplines, too. Political scientists like Streeck (2011) and Krippner (2011) have linked the debt build-up to growing socio-economic pressures. In his history of household borrowing in America, historian Louis Hyman (2011) tied the growth of household debt in America to widening income disparities.¹

The household data show that from 1949 to the 1970s, the rise in household debt was mainly due to an increasing number of households accessing credit markets to purchase homes (financial inclusion at the extensive margin). After the 1970s, debt growth occurred mainly on the intensive margin of housing debt, not on the extensive margin. There is one exception. In line with Mian and Sufi (2013), the data point to an increase in the number of lower income households with mortgage debt over the past two decades (the “subprime boom”). But even for these poorer households, higher indebtedness appears quantitatively more important than the increase in the number of households with debt.

Has the distribution of debt between income groups changed over time, and, if so, how? We also use the joint information on income and debt in the SCF to study changes in indebtedness of different income groups since World War II. We do not find evidence for substantial shifts in the distribution of debt towards poorer strata of the population. The data show that debt-to-income ratios have increased at approximately the same rate over the past six decades across income groups. As income inequality has risen over the same period, the cross-sectional stability of debt-to-income ratios implies that the share of total debt owed by richer households has increased, not decreased. Put differently, on a household level, the correlation between debt and income has become more *positive* over time. We also construct a synthetic panel of birth-year cohorts that allows us to track the borrowing behavior of age groups over their life cycle. Using college education as a proxy for high (low) permanent income households, we also show that for individual birth cohorts the increase in debt-to-income ratios has been stronger for college households than for non-college households and that this is true both for housing and non-housing debt.

Our final point of inquiry is concerned with the implications of the household debt boom

¹Economic historians too 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).

for macroeconomic and financial stability. 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). In the last section of the paper 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 trends in debt-to-income ratios and decompose the increase in household debt by income group. In the third section, we construct synthetic birth-cohorts and study the debt profiles of various cohorts over time. In the fourth section, we discuss the effects of rising household debt on financial fragility. The fifth section discusses the implications of our results for theoretical models of household debt. The last section concludes.

2 Data

As a triennial survey, the SCF is a key resource for research on household finances. The data for the survey waves after 1983 are readily available for download from the website of the Federal Reserve. The comprehensiveness and quality of the SCF explain its popularity among

researchers (see Kuhn and Rios-Rull (2015) and references therein). Selected historical data for the period before 1983 such as the 1962 Survey of Financial Characteristics of Consumers (SFCC) and the 1963 Survey of Changes in Family Finances (SCFF) are also available from the Fed’s website. However, the first consumer finance surveys were conducted much earlier, namely as far back as 1948. Kuhn, Schularick, and Steins (2017) describe how to compile the Historical Survey of Consumer Finances (HSCF) data from the ”historical” and ”modern” waves of the SCF data. In this paper, we rely on this work. The data have hitherto been largely unexploited. We will briefly introduce the data and discuss how the survey data match the aggregate trends from the Flow of Funds before we explore the distribution of debt and its changes over time.

2.1 Historical SCF Data

The 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. While individual studies such as Malmendier and Nagel (2012) have used extracts of the data to address particular questions, to the best of our knowledge, the pre-1983 SCF data have not yet been systematically studied to track the increase in household debt.²

The HSCF contains all variables needed to construct long-run series for the evolution of debt. The SCF also provides additional information on age, sex, race, marital status, family size, and education levels. Based on the HSCF data, we construct the following variables. Total income that is defined 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. Total debt consists of housing and non-housing debt. Housing debt is calculated as the sum of debt on self-occupied homes and debt on other real estate. 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.³

²A detailed discussion of the challenges involved in harmonizing the historical and the modern data can be found in Kuhn, Schularick, and Steins (2017).

³Note 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.

Throughout the analysis, we exclude any outliers, that is those households with a debt-to-income ratio greater than 50, in order to not distort results. As shown in Figure I in the Appendix, these observations correspond to a negligible share of the sample and are mainly households who had an unusually low income in the survey year, typically due to unemployment spells.⁴

2.2 Matching Aggregate Trends


As a first step, 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). Household-level surveys often struggle to match aggregate data when the micro data is aggregated to the level of the macro-economy. In many cases, measurement concepts differ between micro surveys and macro data, explaining at least in part why even high quality micro data do not consistently correspond to aggregate data. For instance, Heathcote, Perri, and Violante (2010) compare NIPA income to the Current Population Survey (CPS). They explain that observed differences are due to the fact that NIPA income includes indirect capital income from pension plans, non-profit organizations and fiduciaries, as well as employer contributions for employee and health insurance funds. These positions are not measured in household surveys such as the CPS or the HSCF. With respect to the FoF, several wealth components of the household sector are measured as residuals obtained by subtracting the respective positions of all other sectors from the economy-wide total (see Henriques and Hsu (2013)). These residuals contain asset positions held by nonprofit organizations as well as domestic hedge funds which are not included in the HSCF. We demonstrate that despite the conceptual differences in measuring income and wealth, the historical SCF data closely matches trends in the aggregate data, effectively alleviating such concerns.

Figure 1 compares income and total debt of the HSCF with the corresponding NIPA and FoF values. Income 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.

⁴In 1949 0.02% and in 2013 0.05% of the respective sample are households with a debt to income ratio greater than 50.

For the base period 1983-1989, the HSCF data matches 84 percent of the data on income from NIPA and 86 percent of FoF debt. Figure 1 shows trends in income and debt for HSCF and aggregate data throughout the 1949-2013 time period. The aggregate data and the aggregated micro data show very similar trends. These similarities support our hypothesis that an exploration of the micro data can shed light on the distributional changes underlying macroeconomic trends. Contrasting the two data series also highlights the strong rise in U.S. household debt. Indexed to the 1980s, we find that income stood at 60 in 1950 and increased to slightly above 120 over the following six decades; debt stood at an index value of 30 in 1950 and increased to over 200 in 2013.

Figure 1: HSCF, NIPA, and FoF: 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 84% for income and 86% for total debt.

3 The Distribution of Household Debt

This section takes up where previous research left off due to the previous lack of micro-level data. By using the historical SCF data we can now track the historical evolution of the cross-section of household debt. We will look at total household debt as well as zoom in on two of its major components, housing debt and non-housing debt. In turn, we will also take a look at the intensive and extensive margin of the increase in household debt.

3.1 Total debt

First, we look at how debt is distributed among rich and poor households and how this distribution has evolved over time. In Figure 2, we sort households according to their income and compute the share of total debt that is owed by each income group. The upper graph of figure 2 divides the whole cross section into quintile groups. In the lower graph, we break the top 20% income households into four groups: households with income between the 80th to 85th, the 85th to 90th, the 90th to 95th and above the 95th percentile, i.e., the top-5 percent.


The upper panel of Figure 2 shows that debt shares, in the past and present, increase with income. In some sense, this might appear unsurprising as richer households have a greater capacity to carry debt. For the entire time period, debt of households in the bottom quintile corresponds to less than 5% of aggregate debt and that of households with income between the 60th and 80th percentile to about 24%. However, the data also show that these relative shares have been broadly stable over time. The top 20%'s share slightly increased over time. In 1950, it was about 45% and since 1992 the top-20 owe more than half of aggregate debt, more than twice as much as households in the quintile below. The lower panel of Figure 2 breaks down the share of debt owed by the top income households. The chart displays a rising share of total debt accounted for by rich households: the share of the top 20-15% households has always been below 15%, whereas that of the top 5% was about 15% in 1950 and has remained at about 20% since the 1990s.

Tables 1 and 2 show the shares in aggregate housing and non-housing debt, respectively. The distribution of non-housing debt has been roughly stable over time. Compared to the distribution of aggregate debt, however, the share of poor households in non-housing debt is greater: the bottom 40% of the income distribution owe about 20% of aggregate non-housing debt, but only about 10 % of aggregate debt.

The data show that shares in housing debt slightly decreased for income groups up to the 4th quintile. In contrast, the share of the top 20% income households increased from 41% in 1950 to 55% in 2013. This increase is mainly driven by the top 5% whose shares rose from 11 to 20% between 1950 and 2013. In other words, the increase in debt shares of top income households in Figure 2 is mainly driven by housing debt.

Figure 2: Shares in aggregate debt

(a) cross section



(b) top 20% income households




Table 1: Shares in Aggregate Non-Housing Debt

year	<20th	20-40th	40-60th	60-80th	>80th	80-85th	85-90th	90-95th	>95th
1950	7.1	15.4	18.7	17.3	42.0	8.4	6.7	10.9	15.9
1971	4.3	14.7	25.0	29.8	26.1	8.6	7.4	6.0	4.2
1992	7.8	13.3	19.6	23.9	36.1	7.1	6.5	6.9	15.6
2013	8.4	11.8	17.9	24.6	37.5	8.4	9.1	8.4	11.6

3.2 Debt and Income

One explanation for the changes in the distribution of debt are shifts in the distribution of income over time. In Figure 3, we track the evolution of debt-to-income ratios in different parts of the distribution to study the indebtedness of rich and poor households and how this

Table 2: Shares in Aggregate Housing Debt

year	<20th	20-40th	40-60th	60-80th	>80th	80-85th	85-90th	90-95th	>95th
1950	3.1	8.0	19.1	28.7	41.3	7.9	10.7	11.4	11.3
1971	2.5	6.8	17.0	30.3	43.8	10.8	9.0	11.2	13.0
1992	2.4	6.4	12.0	25.2	54.6	8.1	10.7	14.4	21.5
2013	2.3	6.6	12.5	24.5	54.6	9.0	11.5	14.0	20.0

has changed over time. Again, the upper graph divides the cross-section into quintiles, the bottom graph subdivides the top 20% into four subgroups. The blue and red bars indicate how high households are indebted in housing and non-housing debt, respectively.⁵

The indebtedness of households was relatively evenly distributed in 1950 with mean debt-to-income being less than 40% across all income groups. From 1950 to 1971 indebtedness of the bottom quintile stayed approximately constant, whereas that of income groups from the second quintile upwards increased. The highest growth was in the fourth quintile, where mean debt-to-income rose by more than 20 percentage points (pp). As a result, since 1971, mean debt-to-income ratios are the lowest for poor households and increase as income does up to the 4th quintile. The ratio of the 5th quintile is slightly lower than that of the 4th. This pattern is maintained until 2013.


Since the 1990s indebtedness has risen strongly across all income groups. However, the strongest increase did not take place in the bottom of the income distribution but rather in the middle and top. The lower graph of figure 3 shows that over time the increase in debt-to-income ratios was strongest for households between the 60th to 90th percentile. For these income groups, the mean debt-to-income ratios rose from less than 60% in 1971 to 140% in 2013. Once more this points to the upper part of the income distribution as the main driver of the debt increase.

If we dive deeper into the data and differentiate between housing and non-housing debt, we find that mean debt to income ratios of poorer households, i.e., those with income in the bottom 40%, rose due to rising housing and non-housing debt-to-income ratios. For the other income groups, housing debt was the key force behind the debt increase over time.


⁵The blue bars are mean housing debt to income ratios, the red bars indicate non-housing debt-to-income ratios so that the sum is approximately equal to the mean debt to income ratio.

Figure 3: Mean of Debt-to-Income Ratios

(a) cross section




(b) top 20% income households



Another way to express the cross-sectional relation of debt and income and track its evolution over time is to look at the correlation between debt and income in the different survey years. Figure 4 displays the correlation coefficient between income and debt for each survey year. The correlation is positive and even shows an upward trend over the past six decades. In other words, the positive relationship between debt and income has become stronger over time – the opposite of what one might expect if increasing income polarization had led to a rapid debt increase of lower income segments of the population.

The above discussion leads to three main results. First, rich households have always accounted for most of aggregate debt and their share has not declined since WW2. Second, the higher the income of a household, the higher the share of housing debt in its debt portfolio. Third, household indebtedness has increased strongly across all income groups since the 1970s and in turn there is no sign of differential trends in debt increase between various

Figure 4: Correlation of Total Debt and Income



Notes: The top 0.1% of the income distribution are excluded to account for the fact that income at the very top of the distribution is imputed in the historical surveys.

income groups.

3.3 Decomposition of Debt-to-Income Ratios

The previous section analyzed how the distribution of debt and households' indebtedness (debt-to-income ratios) evolved over time. However, we have not yet decomposed the increase in its extensive and intensive margin: Has the total number of indebted households increased, or have already indebted households taken on more debt? In this section, we decompose the rise of mean debt-to-income ratios across income groups in the following way: on the one hand, we determine changes in debt-to-income ratios that are due to the extensive margin, i.e., an increase in the number of households with debt. On the other hand, we back out the importance of the intensive margin, i.e., increases due to higher debt-to-income ratios of indebted households. The ex- and intensive margin effects are calculated separately for housing and non-housing debt.

More precisely, $\bar{d}_{i,t}$ 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 $\bar{d}_{i,t}^{H+}$ is the mean housing debt-to-income ratio of households with positive housing debt, i.e. the intensive margin. $s_{i,t}^{N+}$ and $\bar{d}_{i,t}^{N+}$ are the respective values of non-housing debt. The mean

debt-to-income ratio, $\bar{d}_{i,t}$ can be written as follows: $\bar{d}_{i,t} = s_{i,t}^{H+} \bar{d}_{i,t}^{H+} + s_{i,t}^{N+} \bar{d}_{i,t}^{N+}$. The percentage point change in debt-to-income ratios between period t and $t - 1$ is then calculated by:

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

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 \bar{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.

Overall, we find that half of the 78 pp increase in household debt relative to income was driven by the intensive margin of housing debt, about one third (25 pp) by the extensive margin of housing debt. Mortgage lending hence played the dominant role for the increase in household debt.

Table 3: Decomposition of the increase in aggregate debt-to-income between 1950 and 2013

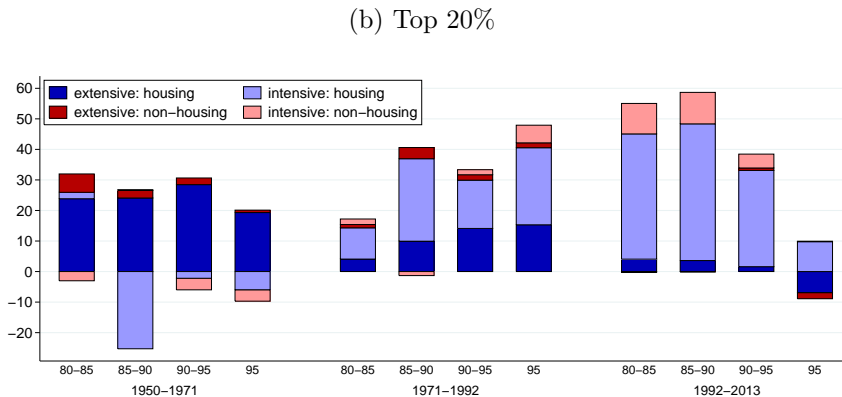
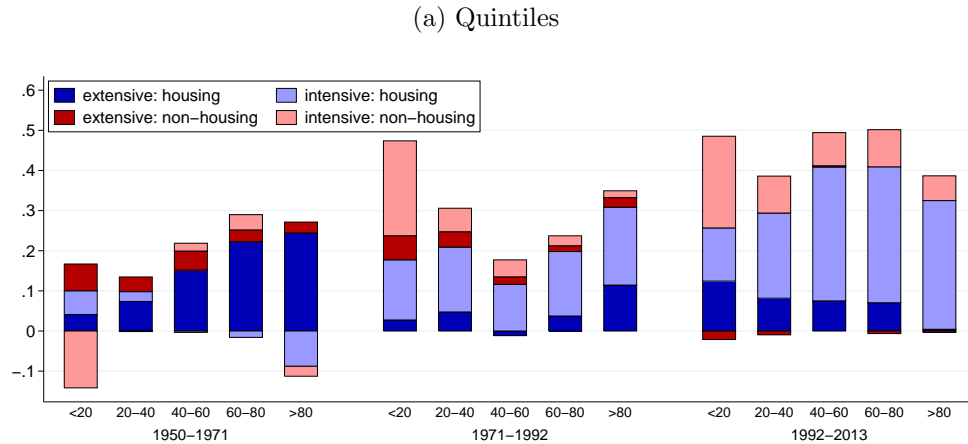
Extensive margin	housing debt	25.2
	non-housing debt	5.2
Intensive margin	housing debt	40.5
	non-housing debt	7.4
total increase		78.4

Notes: Percentage point change in aggregate debt-to-income between 1950 and 2013.

The results of this decomposition are presented in Figure 5. The dark blue and red bars are the extensive margin effects of housing and non-housing debt between two time periods. The light blue and red bars are the corresponding intensive margin effects. The sum of all four bars is the total percentage point change of mean debt-to-income ratios between two time periods. The upper graph shows results for quintile income groups and the lower graph further subdivides the top 20% into four groups. We see that household indebtedness of all

income groups except the bottom 20% increased between 1950 and 1971. The increase was highest for households between third and fourth quintile: their debt-to-income ratio rose by 30 pp. From 1971 to 1992 indebtedness again increased, including a pronounced increase among poor households. The highest rise in debt-to-income ratios was between 1992 and 2013. Except in the top 5%, indebtedness rose between 30 and 60 pp relative to income.

Figure 5: **Decomposition of Mean Debt-to-Income Ratios over Time**



Notes: The dark blue bars refer to the percentage point change due to variation in the extensive margin of housing debt. The dark red bars refer to a variation in the extensive margin of non-housing debt. The light blue and red bars refer to a variation in the intensive margin of housing and non-housing debt, respectively.

The decomposition in ex- and intensive margin effects reveals that different forces caused variations in debt-to-income ratios over time. The increase between 1950 and 1971 was mainly driven by higher extensive margins of housing debt, i.e., more people took on housing debt. In contrast, both the rise from 1971 to 1992 and 1992 to 2013 was mainly driven by

higher intensive margins of housing debt. An exception are poor households after 1992 – their increase in debt-to-income ratios was mainly due to higher intensive margins of non-housing debt.

In short, the rise in household indebtedness from 1950 to 1971 was mainly caused by an increase in the extensive margins of housing debt. However, growth in mean debt-to-income ratios both between 1971 and 1992 and between 1992 and 2013 were mostly due to an increase in the intensive margin of debt.

3.4 Decomposing the Rise of Housing Debt

In the previous section, we illustrated how household indebtedness has increased considerably over time. This has been mainly driven by rising indebtedness with respect to housing debt. However, we would like to evaluate whether this increase could be driven by another factor, i.e. by households buying more expensive houses. In order to do this, we decompose the housing debt-to-income ratio in period t in the following way:

$$\frac{HD_t}{I_t} = S_t^{HD} \cdot \frac{HD_t^+}{HV_t^+} \cdot \frac{HV_t^+}{I_t}.$$

with HD_t being mean housing debt and I_t mean income in period t . S_t^{HD} denotes the share of households with housing debt and HD_t^+ is mean housing debt of those households who have housing debt in period t . Finally, HV_t^+ denotes the mean housing value of home owners in period t . The growth rate (GR) of housing debt to income between period $t - 1$ and t is then approximately given by:⁶

⁶The growth rates are approximated by the log-difference, i.e.

$$\begin{aligned} \log\left(\frac{HD_t}{I_t}\right) - \log\left(\frac{HD_{t-1}}{I_{t-1}}\right) &= \log(S_t^{HD}) - \log(S_{t-1}^{HD}) \\ + \log\left(\frac{HD_t^+}{HV_t^+}\right) - \log\left(\frac{HD_{t-1}^+}{HV_{t-1}^+}\right) &+ \log\left(\frac{HV_t^+}{I_t}\right) - \log\left(\frac{HV_{t-1}^+}{I_{t-1}}\right) \end{aligned}$$

$$\text{GR} \left(\frac{\text{HD}}{\text{I}} \right) = \text{GR} (S^{HD}) + \text{GR} \left(\frac{\text{HD}^+}{\text{HV}_t^+} \right) + \text{GR} \left(\frac{\text{HV}^+}{\text{I}} \right) \quad (1)$$

In other words, the growth rate of the housing debt-to-income ratio is approximately equal to the sum of the growth rates of the share of home owners, the share of home owners with housing debt, the intensive margin of mortgage loan to value ratios and the intensive margin of housing value to income.⁷


Figure 6 shows the decomposition of changes in housing debt-to-income ratios according to Equation 1. The light red bars are the growth rates of the share of households with housing debt, S^{HD} . The dark red bars are the growth rates of the intensive margin of mortgage loan to value ratios and the blue bars are the growth rates of the intensive margin of housing to income, $\frac{\text{HV}^+}{\text{I}}$, $\frac{\text{HD}^+}{\text{HV}_t^+}$. The increase in housing debt-to-income ratios between 1950 and 1971 has been mainly due to a rise in the number of households who take on housing debt. This is true across all income groups. From 1971 to 1992 housing debt to income ratios increased mainly due to households having higher housing values relative to their income. In other words, households have become more in debt since they had to finance more expensive houses.

For the years from 1992 to 2013, the decomposition yields a more nuanced picture. For households in the bottom 90% of the income distribution, both higher housing to income ratios and higher mortgage loan to value ratios led housing debt-to-income ratios to increase. For the top 10%, housing debt-to-income ratios stayed approximately the same between 1992 and 2013.

Our calculations support the idea that household indebtedness was influenced in part by households buying more expensive houses. However, the impact of this phenomenon differed over time. Until the 1970s rising indebtedness was mainly due to improved access to finance. In other words, more households took on debt. From 1971 to 1992, housing debt to income ratios increased mainly because housing values rose faster than incomes. Continuing into the present, between 1992 and 2013 higher loan to value ratios served to perpetuate this trend.

⁷In contrast to section 3.3, we are now analyzing the percentage change rather than the percentage point change.

Figure 6: Decomposition of the Growth Rate of Housing Debt-to-Income Ratios



3.5 Debt and Assets


The trends in debt-to-income ratios discussed above are compatible with the idea that the richer strata of the population increased debt to purchase housing assets while the poorest Americans increased debt for consumption. This would point to very different debt dynamics when debt is scaled by the value of assets instead of income. Put simply, richer Americans who used their debt to purchase housing assets are likely to have seen stable (or even falling) leverage while poorer Americans (who did not acquire assets) debt-to-asset ratios would have risen sharply. We therefore need to better understand the dynamics of the debt increase at the top and the bottom of the income distribution.

Figure 7 starts our exploration of this topic. It compares the evolution of debt-to-income and debt-to-asset ratios over time. The visual contrast is very stark. Unlike debt-to-income ratios, which have increased at approximately the same rate across income groups, debt-to-asset ratios have surged for the poorest Americans. In the past 20 years, debt-to-asset ratios of the bottom 60% increased significantly, but stagnated at the top. As the chart illustrates, this increase is mainly due to higher intensive margins of non-housing debt. Housing debt plays a central role in the changes of the debt-to-income ratios, but is less central for changes in the debt-to-asset ratios.

A look at the composition of non-housing debt helps us to understand this phenomenon.

Figure 7: Percentage point change in debt-to-income and debt-to-asset ratios

(a) debt-to-income



(b) debt-to-assets




Figure 8 splits non-housing debt into its components. The left graph shows the shares of non-housing debt components for the bottom 20%, the right graph for the top 20%. In 1983, education debt was small and essentially non-existent. Its share was equal to 10% for the bottom 20% of the income distribution and less than 5% for the top 20%. While all income groups have significantly increased their share of student debt, this is particularly true for the bottom 20%. In 2013 education loans accounted for about 60% of the non-housing debt of this income group.

Figure 9 shows the percentage point change in debt-to-income and debt-to-asset ratios, excluding educational debt. In Figure 7 that included education loans, debt-to-asset ratios rose the most for lower income groups. Excluding education debt, the order is reversed, as can be seen in Figure 9. Therefore, rising student debt is chiefly responsible for the considerable increase in debt-to-asset ratios among the lower income segments of the population. Lacking major sources of income, students have always been relatively poor but students today take on substantially more debt than their ancestors to invest in their college education.

Figure 8: Shares of non-housing debt components by income groups


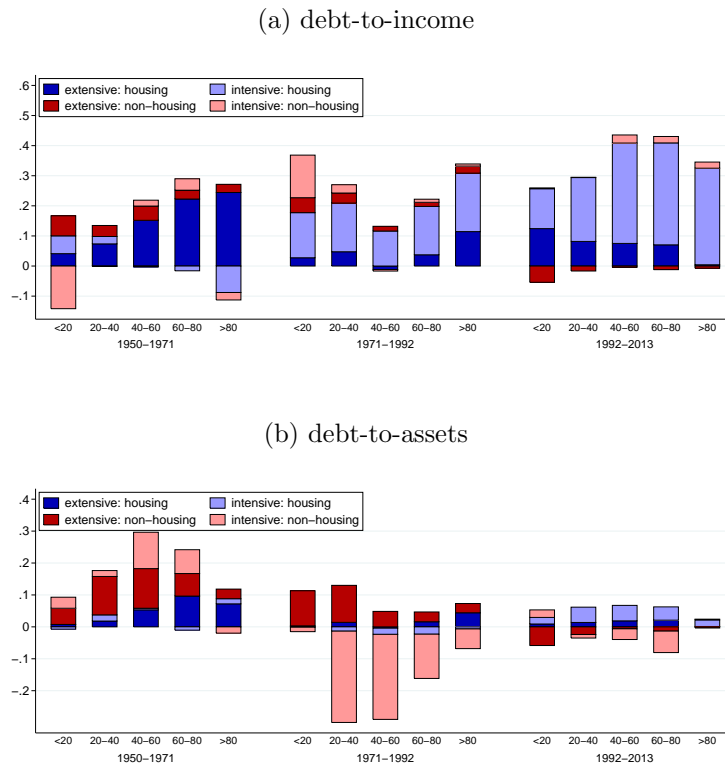


Figure 9: Change in debt-to-income and debt-to-asset ratios excluding educational debt



4 Cohorts


The analysis has thus far looked at the evolution of cross-sectional averages. The trends we observed could also be related to changes in the composition of the age structure of the

population. To shed more light on the forces underlying these trends, we organize the data in this section in terms of birth cohorts and track the financial status of these birth cohorts over time. Since the HSCF data is not a panel dataset, we build synthetic birth-year cohorts and observe the debt portfolios of several of these cohorts over their entire working life. We first track the life-cycle debt profiles of different cohorts.

4.1 Life-Cycle Profiles of Birth Cohorts

Figure 10 shows mean debt-to-income ratios for different birth cohorts. This graph reveals two main things. First, the younger the cohort, the higher their households are indebted throughout their whole life-cycle. Second, the life-cycle profiles changed substantially over time. Households that are part of the oldest cohort (being born between 1915 and 1924) experienced an increase in debt up to the age of 45. After this point, these household reduced their level of indebtedness. In contrast, the life-cycle profile of the following cohorts roughly became a flat line. In other words, these cohorts no longer reduced their indebtedness as they grew older. For the youngest two birth cohorts that we track (1945-1954, 1955-1964) mean debt-to-income ratios even increased with age.

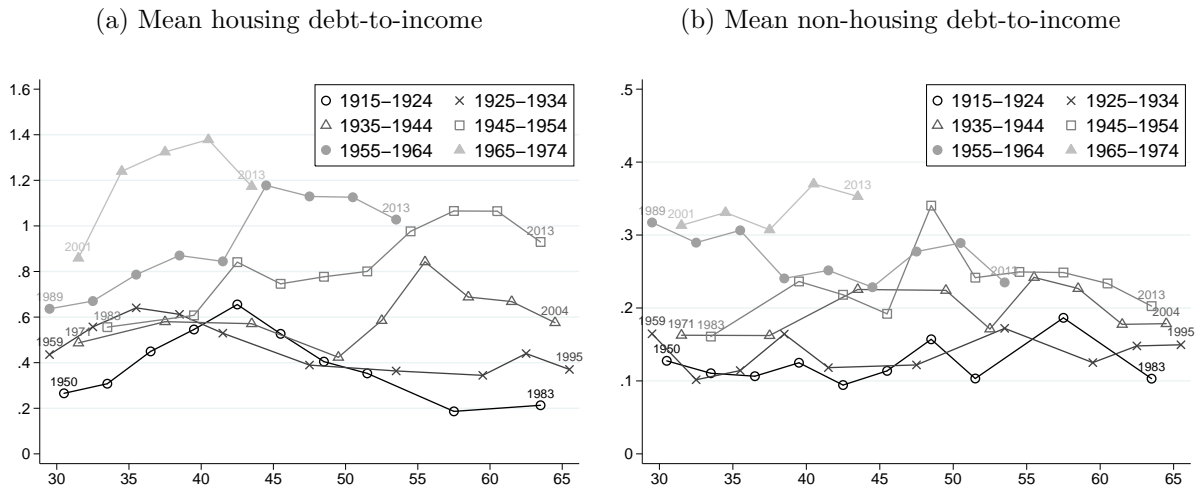
Figure 10: Mean debt-to-income of 10-year cohorts



Are these differences in the life-cycle profiles due to changing patterns of housing or non-housing debt accumulation? In Figure 11, total debt is subdivided into housing and non-housing debt. The evolution of life-cycle profiles with respect to housing debt is very similar to total debt. The younger the cohort, the higher the mean housing debt-to-income ratios


and the higher their indebtedness is at the end of their life-cycle. In contrast, the patterns of indebtedness with respect to non-housing debt differ. Mean non-housing debt-to-income ratios are higher for younger cohorts. However, for all cohorts the degree of indebtedness stays approximately constant over the life-cycle.

Figure 11: Mean housing and non-housing debt-to-income by 10-year cohorts



What about the path of loan to value ratios? It is possible that the life-cycle profile of gross housing debt has become steeper, but that of net housing equity has remained stable. The loan to value ratios are depicted in Figure 12. In contrast to the housing debt-to-income ratios shown above in Figure 11, the loan to value ratios decrease with age across all cohorts. Yet here there is also a trend towards flatter shapes for younger cohorts, as well as rising leverage over time.

Figure 12: Mean Mortgage Loan to Value Ratios of 10-Year Cohorts



4.2 High vs. Low Permanent Income Households

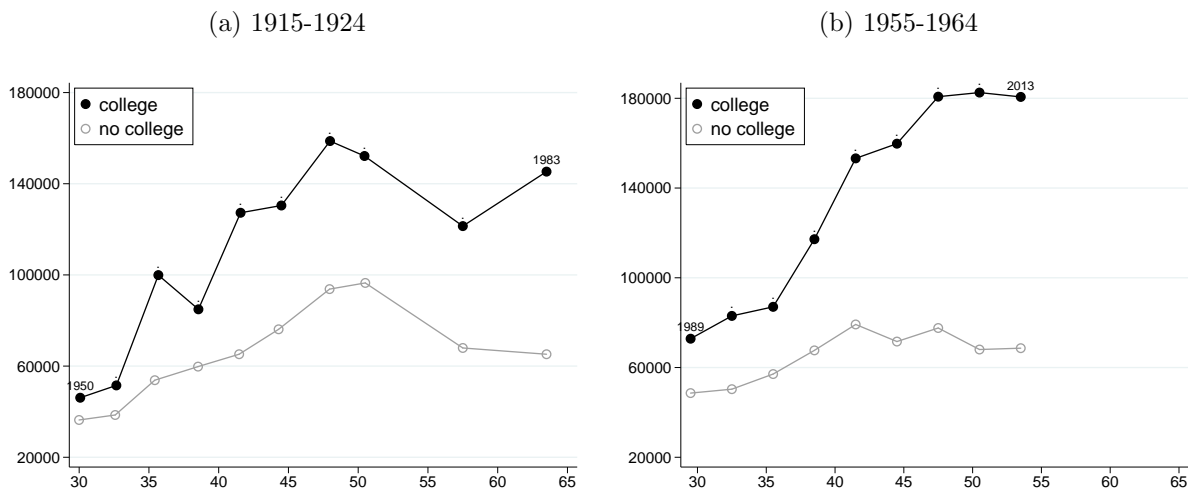
The cohort data also allows us to explore a new aspect of the question discussed in the previous section: is there evidence that poor households have accumulated more debt than rich households as income inequality widened? We use education levels as a proxy for increasing differences in life-time income. We will distinguish between households whose head has a college degree and those with lower education levels. As before, we divide households into 10-year age cohorts and track their debt levels over the time period from 1949-2013. Within each individual birth cohort, we then compare debt trends for college and non-college households.⁸ In the following, the terms college and high-permanent income households, as well as no-college and low permanent income households will be used synonymously.

In Figure 13, we demonstrate the widening income gap between college and non-college households that has been at the center of the debate about rising income inequality. The left graph refers to the birth cohort born between 1915 and 1924, the right to households born between 1955 and 1964. The black lines are the profiles of households with high permanent income and the gray lines are those of low permanent income households. Figure 14 shows the life cycle profiles of debt-to-income ratios for high and low income households. As in Figure 10 households in the younger cohorts are significantly more in debt across their entire life cycle. Moreover, indebtedness of the older birth cohort decreases with age while it increases for the younger cohorts. The life-cycle profiles are very similar for high and low income households

⁸Only households with at least two adults are included to control for changes in the household composition over time. We obtain similar results if we include single households.

in each cohort. For the younger cohort, indebtedness is highest for college households.

Figure 13: Mean income by 10-year education cohorts (in 2013 \$)




Next, in Figures 15 and 16 we look at the age profiles of housing and non-housing debt-to-income ratios for high and low permanent income households. The left graphs show the profiles for households born between 1915 and 1924, the right graphs correspond to those of households born between 1955 and 1964.

Overall, the life-cycle profiles of housing debt-to-income ratios are very similar to those of total debt. For the cohort born between 1915 and 1924 housing debt-to-income ratios increase up to their mid-40s and then decrease afterwards. Both the shape and the level of indebtedness over the life cycle are similar for high and low income households. College households born between 1955 and 1964 are significantly higher in debt than no college households. However, the pattern of housing debt accumulation is similar for high and low income households of this cohort: their housing debt-to-income ratio increases up to the age of 45 and stays constant afterwards. Taking a closer look at Figure 16, we see that indebtedness over the life cycle with respect to non-housing debt evolves somewhat differently compared to housing debt. Both for the cohort born between 1915 and 1924 as well as between 1955 and 1964 the age profiles for high and low income households roughly overlap and are approximately constant over the life cycle.

To sum up, the construction of synthetic panels yielded three main findings. First, the life cycle pattern of indebtedness has changed significantly over time. For cohorts born at the

Figure 14: Mean debt-to-income by 10-year education cohorts



beginning of the 20th century it was hump-shaped, but it has become increasingly flat for cohorts born after 1935. We found that this trend was mainly driven by housing debt levels. Second, the life-cycle profiles of loan to value ratios are continuing to decrease, but have simultaneously become flatter for younger cohorts. Finally, the life-cycle patterns of low and high permanent income households look similar.

Figure 15: Mean housing debt-to-income by 10-year education cohorts

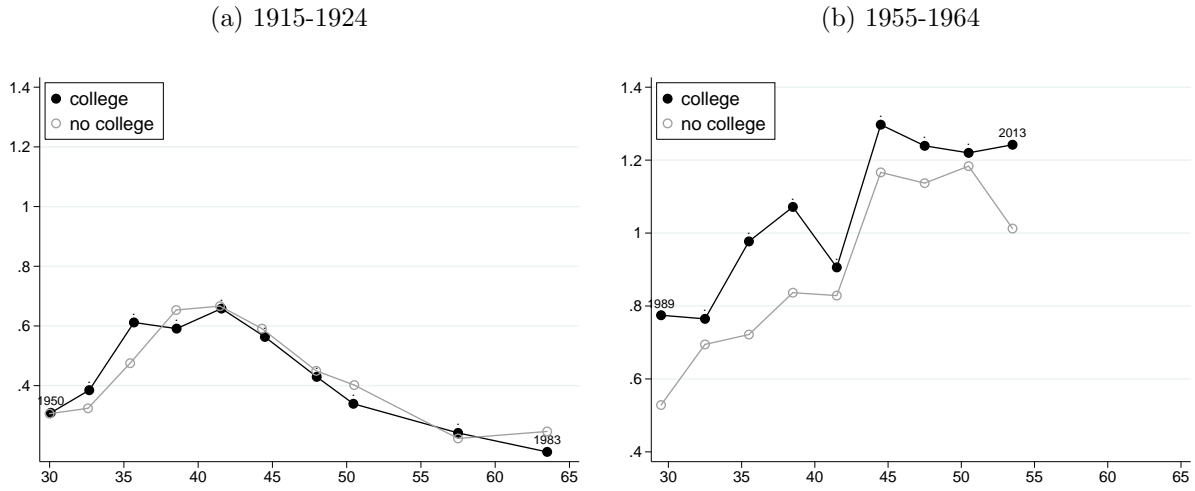



Figure 16: Mean non-housing debt-to-income by 10-year education cohorts



5 Household Debt and Financial Fragility

In the last part of the paper, we aim to understand the effects of the increase of household debt on macroeconomic and financial fragility. We will focus in particular on the question: if the sensitivity of household balance sheets to fluctuations in asset prices rose in lockstep with higher household debt and if so, how did they change? The 2008 financial crisis clearly demonstrated how a drop in house prices can lead to sizeable 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.

The empirical evidence for these effects are well documented. However, the route we are going to take in our analysis is different. The goal of this section is to highlight how the financial fragility of household balance sheets has grown over time. To do so, we propose an empirical exercise that is similar to a stress test for banks. We will “shock” household balance sheet with an exogenous decline in house prices and then track, for a given decline in house prices, the amount of negative equity and the share of negative net wealth households over time.

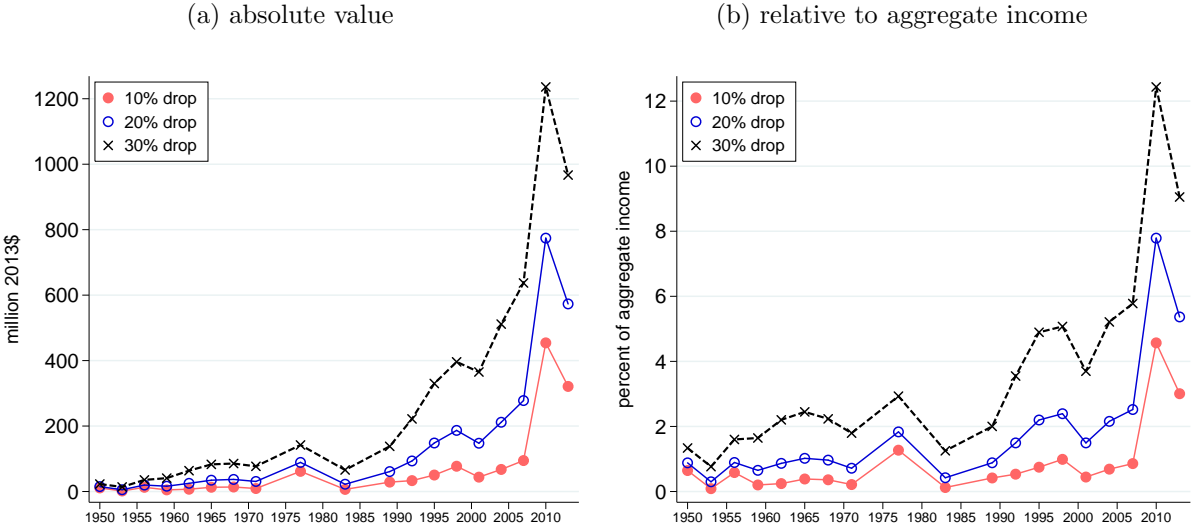
5.1 Negative Home Equity

We start the discussion in Figure 17 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. In the following Figures 18 and 19 we 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 bigger 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.

Figure 17: Home equity at risk



What is the value of non-performing loans if all households with negative equity were to default on their mortgage payments? In Figure 20 we 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 21 shows the value at risk if house prices dropped by 20% for different income groups, respectively.

Figure 18: Home equity at risk by age groups (20% drop in house prices)


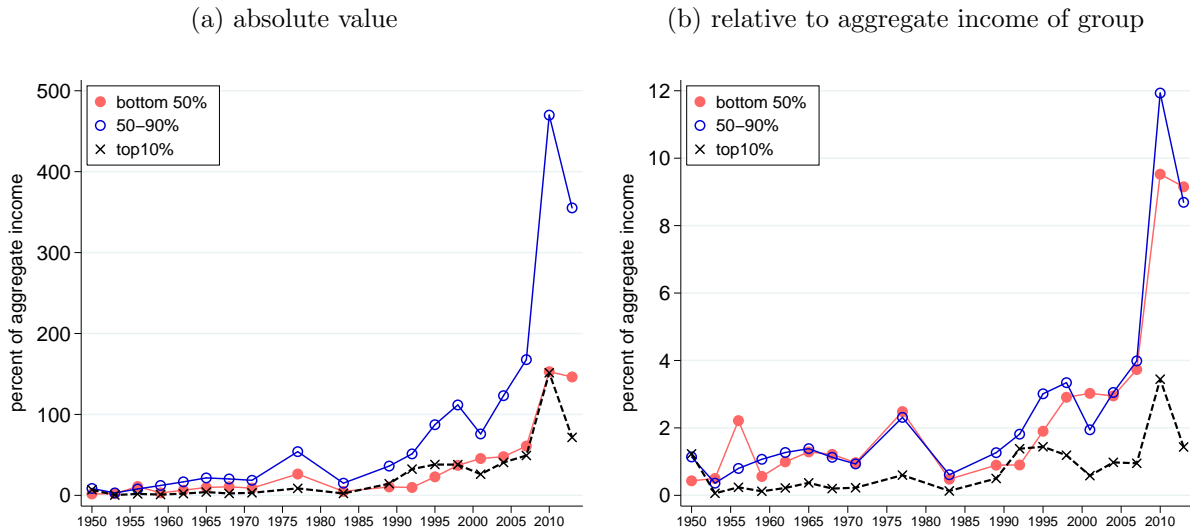


Figure 19: Home equity at risk by income groups (20% drop in house prices)



5.2 Share of Households with Negative Wealth

These stress tests of household balance sheets underline the growing sensitivity of household finances to house price fluctuations. What share of the households from different income groups would effectively end up with negative net wealth assuming a 20% house price drop?

Figure 20: Value at risk

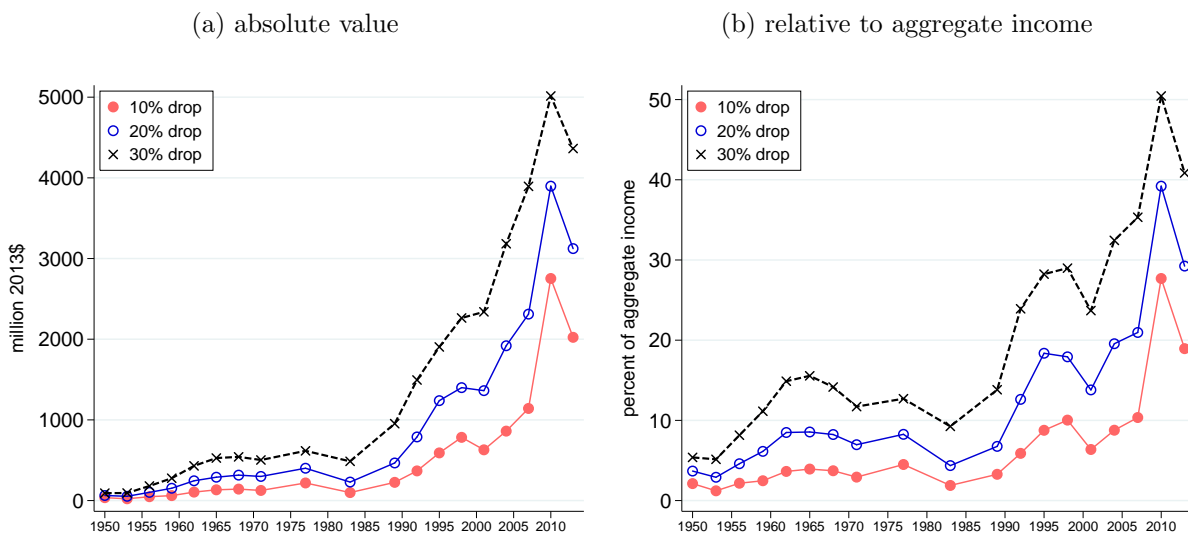


Figure 21: Value at risk by income groups (20% drop in house prices)

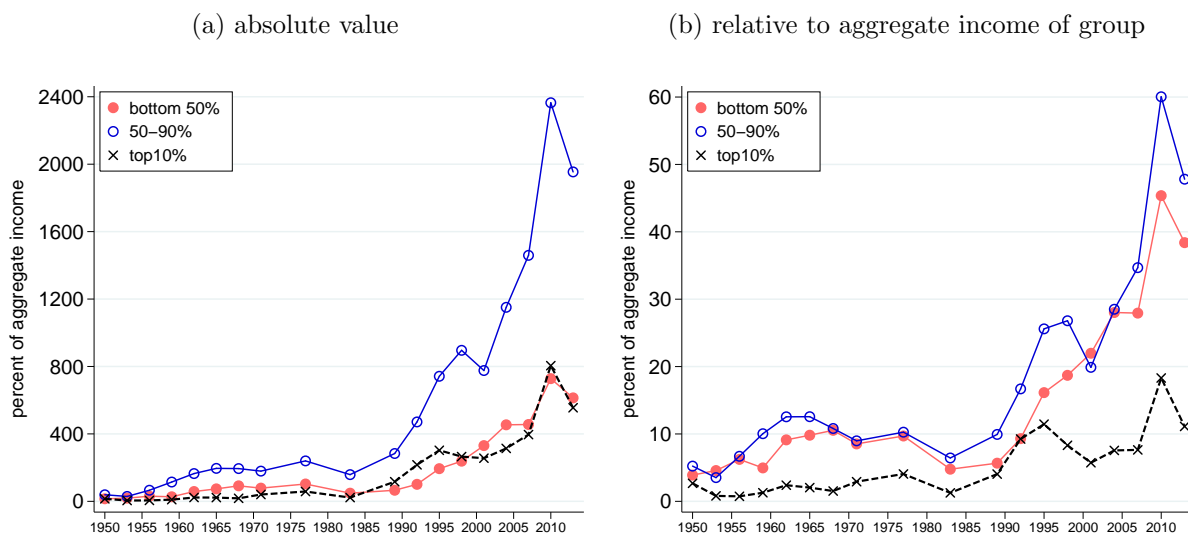


Figure 22 tracks the corresponding shares over time. In the last SCF survey in 2013, about 15% of households in the bottom income brackets would have negative net wealth after a 20% house price drop. Even in the upper half of the income distribution, 8% of households would lose all positive wealth.

It is not difficult to see why the sensitivity of household net worth has risen over time. The exposure of households to house price changes can be expressed as the elasticity of net wealth to house price changes. This elasticity is equal to $\frac{\text{Housing}}{\text{Net wealth}}$. For example, if house prices dropped by 10%, net wealth would decrease by $10\% \cdot \frac{\text{Housing}}{\text{Net wealth}}$.

The elasticity can be decomposed in the following way:

$$\frac{\text{Housing}}{\text{Net wealth}} = \frac{\text{Housing}}{\text{Assets}} \cdot \left(1 + \frac{\text{Debt}}{\text{Net wealth}} \right)$$

The first term is the share of housing in total assets which shows how much the asset portfolio of the household is diversified. The second term indicates how high the household is leveraged.

Increasing exposure of household net wealth to house prices is a function of higher leverage. As discussed above, loan to value ratios increased sharply across age cohorts since the 1980s. While higher leverage has magnified the wealth gains that middle class households made during the two decades of robust house price trends, the 2007 collapse of the housing market has exposed the dark side of these larger balance sheets.

Figure 22: Share of households with negative net wealth by income groups

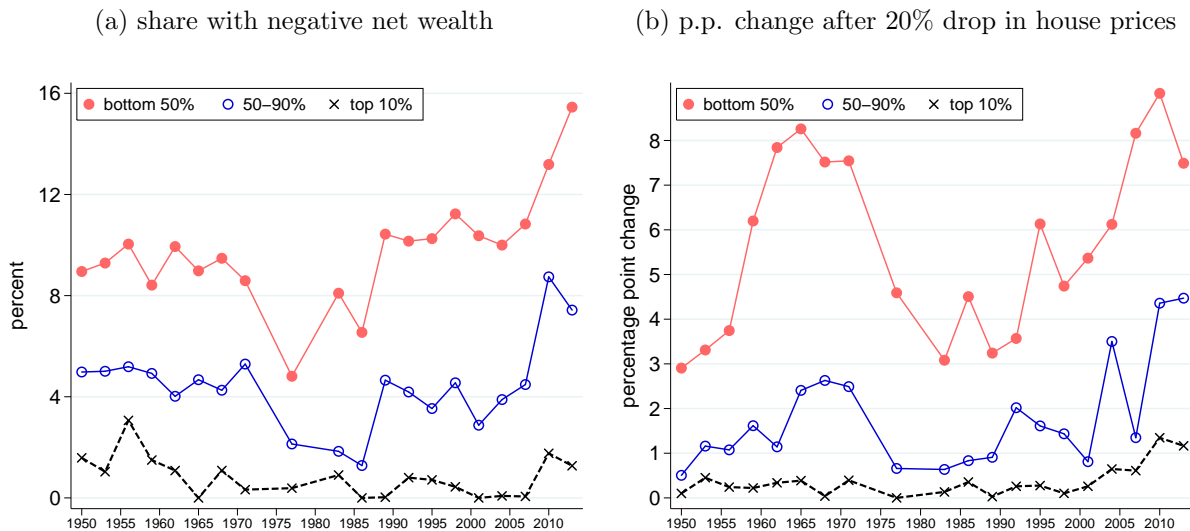


Figure 23: **Exposure to changes in house prices**



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 evaluation of U.S. household debt, its composition and distribution. In doing so, we addressed a number of hypotheses that have been put forward across the social sciences. Most of these do not closely match the new stylized facts we present, pointing to an urgent need for more research on household debt dynamics.

What do we know? First, we show that household debt has increased across income groups. There is very little in the data to support the idea that poorer strata of the population have borrowed more than richer households. If we were to single out an income group, it would be households between the 80th and 90th percentile that have borrowed the most relative

to their incomes. Second, we find that the increase in household debt occurred mainly on the intensive margin of housing debt. Third, we demonstrate to what extent higher leverage has made the American economy more financially fragile. Today, house price fluctuations have far more serious consequences on the health of the balance sheets of consumers – and of the banks who hold the mortgage loans. With higher leverage, asset price fluctuations have come to play a pivotal role in macroeconomic stability.

Our findings provide new and important guiding principles for future theoretical research on household portfolio choices. The interaction between financial deregulation, credit supply shocks and house prices 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 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.

A Data Availability

Table I: Survey of Consumer Finances 1948-1977

year	sample size	weighting scheme	unit	particularities
1948	3162	A1	SU	separate questionnaire for farmers
1949	3068	A2	SU	separate questionnaire for farmers
1950	3069	A3	SU	separate questionnaire for farmers
1951	3029	A3	SU	separate questionnaire for farmers
1952	2501	A3	SU	separate questionnaire for farmers
1953	2756	A2	SU	separate questionnaire for farmers; some DUs also interviewed in 1952
1954	2688	A2	SU	
1955	2805	A2	SU	
1956	2729	A4	SU	
1957	2770	A2	SU	
1958	2813	A2	SU	half of DUs had also been interviewed in 1957, all high income DUs reinterviewed
1959	2860	no	SU	
1960	2972	no	SU	
1961	1981	no	SU	1441 SUs also interviewed in 1960
1962	2117	no	SU	1478 SUs also interviewed in 1960, 1961
1963	2036	no	SU	
1964	1540	no	SU	
1965	1349	no	FU	
1966	2419	no	FU	
1967	3165	B	FU	
1968	2677	B	FU	
1969	2317	B	FU	
1970	2576	no	FU	
1971	1327	no	FU	
1977	2563	no	FU	

Notes: SU refers to spending unit, FU to family unit and DU to dwelling unit. Weighting schemes are described in table II. Farm and non-farm schedules differ with respect to questions on sources of income.

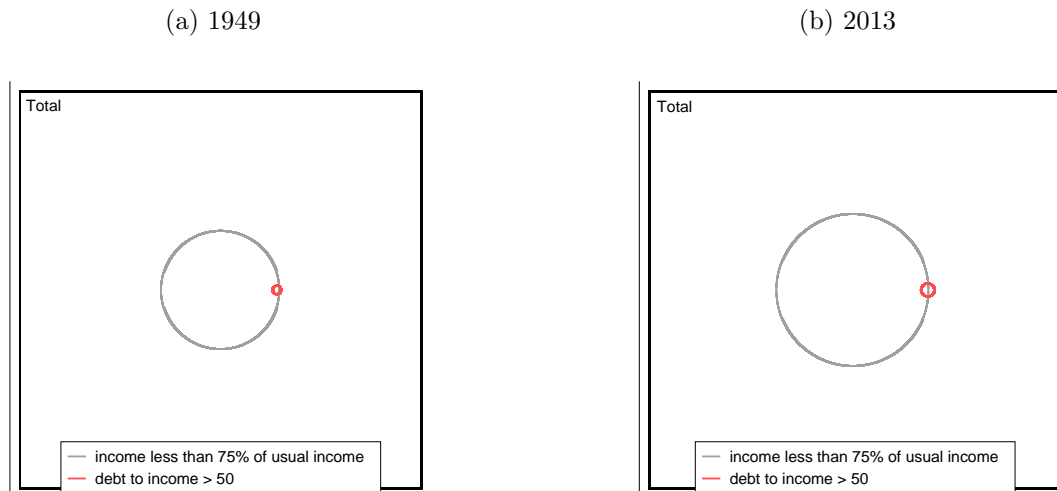
Table II: **Weighting schemes**

Households are classified into high medium and low value dwelling units based on census data on rent and property values and by interviewer's rating of the probable economic level of the residents.				
A	Sample rates	high value DU	medium value DU	low value DU
	A1	4 x basic rate	1 x basic rate	1 x basic rate
	A2	4 x basic rate	2 x basic rate	1 x basic rate
	A3	6 x basic rate	2 x basic rate	1 x basic rate
	A4	2 x basic rate	1 x basic rate	1/3 x basic rate
B	Weights to represent U.S. population			

B Sample selection

We abstain largely from any selection on the sample of households. For the analysis of household indebtedness, we look at debt-to-income ratios. Some households may have very low income in a particular survey year for transitory reasons like unemployment or business losses for self-employed. To avoid effects from outliers of such transitory events, we drop households with debt-to-income ratio larger than 50 from our sample. Figure I shows that dropped households overlap to a large extent with households who have large debt-to-income observations largely due to transitory reasons rather than due to their permanent state of indebtedness.

Figure I: Share of households with debt-to-income ratio > 50 in sample



References

- BELOTTI, F., P. DEB, W. G. MANNING, E. C. NORTON, ET AL. (2015): “twopm: Two-part models,” *Stata Journal*, 15(1), 3–20.
- 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(1), 15–51.
- HENRIQUES, A., AND J. W. HSU (2013): “Analysis of wealth using micro and macro data: A comparison of the Survey of Consumer Finances and Flow of Funds Accounts,” .
- KUHN, M., AND J.-V. RIOS-RULL (2015): “2013 Update on the US Earnings, Income, and Wealth Distributional Facts: A View from Macroeconomics,” .
- KUHN, M., M. SCHULARICK, AND I. U. STEINS (2017): “Wealth and Income Inequality in America, 1949-2013,” Discussion paper, CEPR Discussion paper 12218.
- YANG, Y. (2008): “Social inequalities in happiness in the United States, 1972 to 2004: An age-period-cohort analysis,” *American Sociological Review*, 73(2), 204–226.