Three Essays on Households' Finances and Their Macroeconomic Relevance

Inauguraldissertation

zur Erlangung des Grades eines Doktors

der Wirtschafts- und Gesellschaftswissenschaften

durch die

Rechts- und Staatswissenschaftliche Fakultät

der Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von

Alina Kristin Bartscher

aus *Lüdenscheid*

Bonn, 2021

Dekan:	Prof. Dr. Jürgen von Hagen
Erstreferent:	Prof. Dr. Moritz Schularick
Zweitreferent:	Prof. Dr. Moritz Kuhn
Tag der mündlichen Prüfung:	20.07.2021

Acknowledgements

I would like to thank all people who have accompanied and supported me while writing this thesis. First of all, I would like to thank my supervisors and co-authors Moritz Schularick and Moritz Kuhn for all their valuable feedback and many in-depth discussions, from which I have benefited greatly and learned a lot.

More generally, I am grateful to all the professors and colleagues at the University of Bonn and the University of Mannheim who encouraged me to take up and pursue a PhD, shared their insights and gave me helpful advice at many occasions along the way. Particularly, I would like to thank Sebastian Siegloch for his exceptional support in the job market phase.

Furthermore, I would like to thank my family and friends for giving me strength during the low points and celebrating the high points with me. Above all, I am indebted to Nils, with whom I have had innumerable conversations and discussions about economics and life. Furthermore, my special thanks go to Claudia, Tomasz, Dimo, Enrico, Michaela, Carla, Tamara, Sihong, Chui Yee, Ximi, Kai, Matthias and Gasper – thank you for all the conversations, lunches and coffees, sports sessions and other joint activities.

Finally, I would like to express my gratitude for the administrative assistance during my PhD studies and acknowledge the financial support from the Center for Doctoral Studies in Economics, the Bonn Graduate School of Economics and the German Research Foundation.

Contents

List of Figures

List of Tables

Introduction

How do households make their financial decisions, and how does this affect macroeconomic outcomes such as homeownership and borrowing rates, labor supply and socioeconomic inequality? This thesis contributes to the economic literature related to these questions by studying high-quality long-run microdata and combining them with economic theory. The first chapter studies the effects of relaxing income-based mortgage borrowing constraints on housing and labor market decisions of married couples. The second chapter sheds light on the evolution of U.S. household debt since the 1950s and its drivers, with a special focus on the role of housing debt and home equity extraction. The third chapter provides long-run evidence on the college wealth premium, and its relation to households' portfolio choices.

In CHAPTER 1, "It Takes Two to Borrow", I study the effects of the U.S. Equal Credit Opportunity Act (ECOA) on married couples' homeownership and mortgage borrowing choices. The ECOA prohibited the formerly common practice of discounting 50% or more of the wife's income in a couple's joint mortgage application. It thus provides a natural experiment to study the relaxation of income-related borrowing constraints. In the empirical part, I use data from the Panel Study of Income Dynamics (PSID) to estimate difference-in-difference regressions at the national level as well as state-level event studies. I find that the introduction of equal credit opportunity legislation had significant and sizable positive effects on couples' mortgage borrowing and homeownership, enabling around 1.4 million households to buy a home. As the relaxation of debt-to-income constraints worked via the creditable share of the wife's income, it plausibly also affected women's labor supply choices due to income and substitution effects. Employment incentives can be expected to be strongest for young women in recently founded households who still have to determine whether to buy a home and whether the wife should work. A married woman planning her financial future at the beginning of her household's life cycle in the post-reform world might hence make different labor supply decisions compared to a world in which the reform had not happened. In order to examine whether and to what extent this is the case, I build a simple life-cycle model of married households' homeownership and mortgage choices featuring debt-to-income and loan-to-value constraints, calibrate it to the early 1970s and simulate the reform. I find that the reform incentivized married women to join the labor force. Importantly, this has strong amplifying effects on homeownership and borrowing, magnifying the respective responses by a factor of two. Consistent with economic intuition, the effects are strongest for younger households with comparatively low male earnings. My results illustrate the relevance of income-based credit constraints, as well as the interconnectedness of housing, credit and labor markets. Furthermore, my findings imply that the financial inclusion of women can have important effects on household's asset and labor market decisions. Today, this may be particularly relevant in developing countries, in which women still face more restrictions in their access to financial markets.

CHAPTER 2, "Household Debt and Inequality in the United States, 1950-2019", which is joint work with Moritz Kuhn, Moritz Schularick and Ulrike Steins, studies the secular increase in U.S. household debt over the past seven decades. We use microdata from the "SCF+", which extends the modern Survey of Consumer Finances (SCF) back to the 1950s and covers the joint distribution of debt, income and wealth. From 1950 to 2010, the ratio of U.S. household debt to income quadrupled to about 120%. Around 80% of this debt is housing debt. We find that most

of the debt increase is accounted for by households from the 50th to 90th percentile of the income distribution, whom we refer to as the middle class. Furthermore, we find that the nature of the debt increase has changed over time. It was driven by rising numbers of borrowers in the first boom phase until 1965, when homeownership was expanding after World War II. By contrast, it was driven by larger debt amounts conditional on borrowing in the second boom phase from the early 1980s to the financial crisis of 2008. Over the 1970s, households had experienced windfall gains in home equity due to rising house prices and high inflation. This was particularly relevant for the middle class, for whom housing constitutes the largest share of their wealth. The economic life-cycle theory predicts that households would like to consume at least part of such gains if they are perceived as permanent. As houses cannot easily be divided, this means they have to borrow to extract equity from their home, unless they want to sell it. Moreover, house price increases relax collateral constraints, facilitating more borrowing. In the 1980s, falling mortgage interest rates provided incentives for cash-out refinancing, and tax reforms increased the attractiveness of mortgage borrowing. The spread of securitization and financial deregulation allowed lenders to invent new mortgage products and meet the rising demand for equity extraction. Using panel data from the PSID, we show that equity extraction can account for half of the increase in housing debt between 1980 and 2007, or about 43% of the increase in total household debt. Yet equity extraction, which typically occurs in the middle of the (adult) life cycle, induces pronounced changes in life-cycle debt profiles compared to a situation where households take out a mortgage early in life and repay it without re-augmenting the balance. We find that life-cycle debt-to-income profiles shifted and turned upward from cohort to cohort, implying that households entered retirement with substantially elevated debt levels. However, due to simultaneous rises in asset values, younger cohorts have similar net wealth over the life cycle as their predecessors. We further show that black households, who have much lower homeownership rates and house values than white households, have contributed relatively little to the debt increase over the past 70 decades, and particularly little to the extraction boom. In conclusion, the debt boom was mainly driven by white, middle-aged households from the middle class.

CHAPTER 3 "The College Wealth Divide: Education and Inequality in America, 1956-2016", which is joint work with Moritz Kuhn and Moritz Schularick, examines long-run trends in the income and wealth of U.S. households with and without a college-educated head, using microdata from the "SCF+". While the college income premium has been discussed extensively in the economic literature, the college *wealth* premium has received less attention. We document the emergence of a substantial college wealth premium since the 1980s, and find it to be about three times larger than the college income premium. The real wealth of U.S. households with a college-educated head has tripled over the past four decades, while that of non-college households has barely grown. Part of the wealth gap can be traced back to portfolio differences between college and non-college households, which give rise to different asset price exposures. Non-college households have small stock holdings and have thus profited much less from recent stock market surges. In particular, we find that a substantial part of the widening of the college wealth gap is driven by strong wealth gains in the top 10% of the wealth distribution, where stock holdings are concentrated. Capital gains can explain why the college wealth gap could expand much more than the college income gap. We also find that households with two college-educated spouses enjoyed particularly large increases in wealth. Yet this trend is not driven by assortative mating, but by the overall growth in college education, especially among women. Apart from stock ownership, we find an association between the college wealth premium and business ownership. We show that average financial literacy is higher among stock and business owners, and that business owners are able to generate larger returns to their financial and business wealth. While this is true for owners with and without a college degree, college-educated owners tend to possess much more valuable businesses, such that similar rates of return can translate into substantial level differences. Our results suggest that financial literacy and entrepreneurial skills play an important role in the evolution of the college wealth divide via their effects on returns and portfolio choice.

Chapter 1

It Takes Two to Borrow: The Effects of the Equal Credit Opportunity Act on Homeownership and Mortgage Debt of Married Couples

Until the 1970s, U.S. mortgage lenders commonly discounted the wife's income in couples' joint mortgage applications. This changed with the introduction of anti-discrimination legislation in the 1970s. The Equal Credit Opportunity Act of 1974 prohibited credit discrimination related to marital status and sex, providing a natural experiment to study the relaxation of income-related borrowing constraints. Using data from the Panel Study of Income Dynamics, I find positive effects on mortgage borrowing and homeownership rates of married couples with working wives. Difference-in-difference results at the national level are supported by event study regressions exploiting variation in hand-collected state laws. My results imply that the new legislation initially enabled more than a million households to move to their own home. I further built a life-cycle model of married households' homeownership and mortgage choices to explore the labor supply incentives for wives entailed by relaxing debt-to-income constraints via the creditable share of their income. The results show a positive effect on married women's labor force participation, which strongly amplifies the homeownership and borrowing effects.

1.1 Introduction

Access to credit is an essential part of the economic lives of millions of people in the United States and elsewhere. Yet only half a century ago, this access was restricted for a substantial share of the United States' citizens. Until the 1970s, discrimination in credit transactions based on characteristics like sex, marital status or race was widespread and institutionalized. Among women, especially those who were married faced severe restrictions. For instance, they could hardly obtain credit in their own name. Moreover, it was common to discount the wife's income by 50% when a couple jointly applied for a mortgage, or to even completely disregard it (Kendig 1973). This is of particular importance because mortgage debt is the dominant form of household debt, accounting for about 80% of total household borrowing (cf. Bartscher et al. 2020). Female labor force participation had increased substantially since the 1950s and kept rising, particularly among married women (Juhn and Potter 2006). With stable female employment becoming more

and more common, the established practice of income discounting increasingly failed to live up to married couples' reality of life.

Given the relevance of borrowing, and in particular mortgage debt, for economic participation and consumption smoothing, it is important to understand the effects of legislative attempts to lift restrictions in access to credit. In the early 1970s, not a single country in the world had a law that explicitly prohibited gender-based discrimination in credit access (see Figure 1.1). This changed when the Equal Credit Opportunity Act (ECOA) was passed in the United States in 1974, becoming effective in October 1975 (Smith 1977). The new law precluded discrimination in lending based on sex and marital status, including the practice of income discounting. It thus constitutes a natural experiment to study if and to what extent the relaxation of income-related borrowing constraints affects households' access to homeownership and mortgage credit.

While the macroeconomic literature has mainly focused on loan-to-value constraints to model frictional financial markets, income-related borrowing constraints have only recently gained more attention (Greenwald 2018). I apply difference-in-difference and event study techniques to show that the ECOA had sizable and significant positive effects on the homeownership and mortgage borrowing rate of married couples with working wives. These results provide clear evidence in support of the importance of debt-to-income constraints. Moreover, I use a simple quantitative life-cycle model to study the interactions between households' behavior in the mortgage market and women's behavior in the labor market. By increasing the creditable share of married women's incomes, the ECOA raised their return to working. I find that the labor supply incentives of the act were powerful enough to motivate women to join the labor force. The increased labor force participation in turn strongly amplifies the positive effects on homeownership and borrowing.

Although the ECOA required creditors to profoundly change their lending practices (Smith 1977), its effectiveness in increasing women's access to credit have remained an open question. Elliehausen and Durkin (1989) provide theoretical arguments why they think the ECOA did not increase credit availability "to anyone", whereas Ladd (1982) and Haurin and Kamara (1992) provide suggestive empirical evidence pointing to the contrary. However, none of these studies conducts a comprehensive econometric analysis of the topic.

In the empirical part of the chapter, I attempt to fill this gap by analyzing data from the Panel Study of Income Dynamics (PSID). In particular, I exploit the institutional setting to estimate difference-in-difference regressions at the federal and event study regressions at the state level. My main focus is on married households with working wives. Since the ECOA ended the practice of income discounting, my hypothesis is that households with a higher income contribution of the wife could benefit more from the new law. I therefore compare households with a higher to those with a lower female income contribution (first difference) before and after the reform (second difference).

My results show that subsequent to the ECOA, the increase in mortgage-to-income ratios was higher the more the wife's earnings had contributed to total household income in the pre-reform years. During the pre-reform period, there was no differential trend in mortgage-to-income ratios depending on the wife's income share. The effect on mortgage debt is mainly driven by the extensive margin, i.e. by a higher relative share of households holding a mortgage. Additionally, I find a relative increase in homeownership, which entails increases in housing-to-income ratios and home size. The relative increase in home size mostly stems from households moving from smaller rented to larger owned properties. My estimates imply that the introduction of legislation against gender-based credit discrimination allowed 1.4 million of married households to move to an own home, and 1.8 million to take out a mortgage.

The first congressional hearings on equal credit opportunity for women, which finally led to the passage of the ECOA, took place in 1972. Appendix Figure 1.A.4 provides graphical evidence on how attention for the topic "women and credit" surged from 1972 on, based on data from the

Figure 1.1: Laws prohibiting discrimination in access to credit based on gender

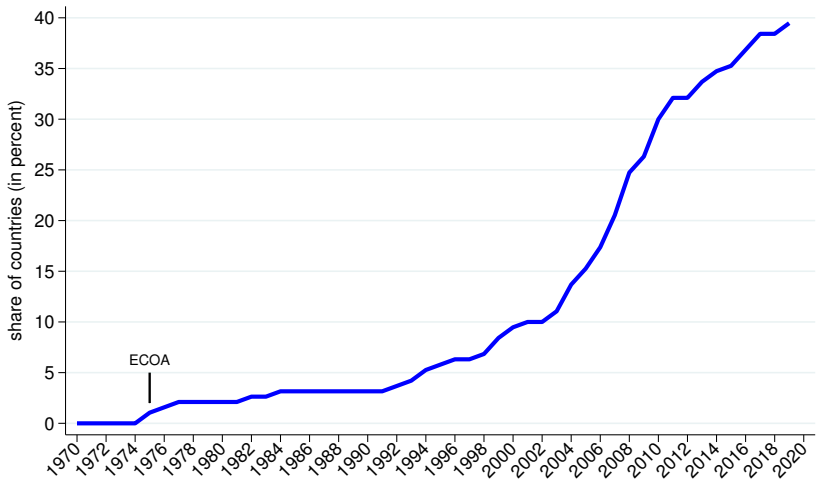

Notes: The graph shows the share of countries with explicit laws against discrimination in access to credit based on gender in a sample of 190 economies. Note that the ECOA became effective in 1975, and applied to Puerto Rico as well. Data source: World Bank Women, Business and the Law (WBL) database.

Google Books Ngram Viewer. Against this background, I use 1971 as the last pre-reform year in the federal-level difference-in-difference regressions. Another important aspect is that several states already introduced equal credit opportunity laws on their own initiative in the interim period between the first congressional hearings of 1972 and the effective date of the ECOA at the federal level in 1975 (U.S. Department of Labor 1975).

An advantage of this interim period is that it allows me to exploit state-level variation in an event study design. For this purpose, I collected information on the respective state laws. While the national-level difference-in-difference design relies on a comparison of households with higher versus lower pre-reform female income contributions within a state, the state-level event study design compares households with a working wife in states which have already implemented an anti-lending-discrimination law in a given year to similar households in states which have not done so yet. The results corroborate the national-level evidence based on state-level variation. While the structure of the data does not allow me to run the event-study regressions with mortgage debt as the outcome (see discussion in Section 1.3), I find positive and significant effects on homeownership, house size and housing-to-income ratios of married couples in treated relative to untreated states after the reform, whereas there were no differential trends during the pre-reform years.

An important strength of the state-level regressions is that they are unlikely to pick up other contemporaneous events, given that these would have to coincide geographically and temporally with the introduction of the state-level laws. I further conduct a battery of robustness checks for the national-level regressions to minimize the risk that the results are driven by confounding events, such as the return of veterans from the Vietnam War in the early 1970s, who could obtain advantageous mortgage conditions due to insurance by the Department of Veterans Affairs (VA) (Foote and Peterson 2008).

On impact, the new legislation was most beneficial for households in which the wife had already been working, as lenders were still allowed to take employment continuity into account when determining mortgage eligibility (Geary 1976, Cairns 1976). However, the ECOA may have changed the labor supply incentives for women in the subsequent years. There can be two opposing effects on female labor supply. On the one hand, a wife ceteris paribus had to work less after the reform to afford a mortgage of a given size. On the other hand, the return to

labor supply in terms of borrowing capacity increased with the reform, providing positive work incentives. These changes in labor supply incentives will not only affect existing households, but also and in particular couples who only form a household (and make their housing and labor market decisions) after the reform. For these households, it is not possible to rely on pre-versus-post-reform comparisons in the data. In order to examine the potential labor supply effects, I therefore build a life-cycle model of married households' homeownership and mortgage choices, drawing on previous work by Pizzinelli (2018), Attanasio et al. (2012), Druedahl (2015) and Bottazzi, Low, and Wakefield (2007). Couples face idiosyncratic income risk and can choose whether to rent or own a house. If they opt for ownership, they can borrow against their house up to the minimum of a loan-to-value and debt-to-income constraint. Mortgages are modeled as long-term debt. I calibrate the model to the early 1970s, and simulate it under the assumption that either 50% or 100% of the wife's labor income can be counted toward a mortgage.

The results show that the ECOA was powerful enough to encourage married women to join the labor force. Under the given calibration, an increase in the female income discounting factor from 50% to 100% entails an increase in the female labor force participation rate of 2 percentage points. The additional female labor supply approximately doubles the effect on married couple's homeownership. The effects are strongest for young households, consistent with the empirical fact that households typically buy their first home home before their mid-thirties. Moreover, the effects are stronger for households in which the husband earns less than the median male income. This is again in line with the data, where the effects on homeownership and debt are stronger the more the wife (and the less the husband) contributes to the household's income. While the new equal credit opportunity legislation already had sizable positive effects on married couple's homeownership upon its introduction, the model suggests even larger medium- to long-run effects due to the amplification via increased female labor force participation, because the number of households benefiting from the possibility to count the wife's income toward a mortgage increases with more women entering the labor force.

My work is related to different strands of literature. First, it contributes to the literature on women's financial rights and decisions. For instance, Hazan, Weiss, and Zoabi (2019) find that the extension of married women's property rights to movable property in the U.S. since the 1850s induced significant shifts in household portfolios. Goldsmith-Pinkham and Shue (forthcoming) show that even today, housing wealth is associated with gender differences in the U.S. They find that single women earn substantially smaller returns on housing compared to single men, while couples range in between. The authors name differences in market timing as an important explanation, but also point to the possibility of discrimination in negotiations.[1] My research shows that granting women the same rights as men in mortgage applications has important effects both on household balance sheets and wives' labor supply decisions.

The chapter also adds to the literature on home financing, the role of debt-to-income constraints and their interaction with female labor supply. Foote, Loewenstein, and Willen (2018) examine the computerization of U.S. mortgage lending in the 1990s, which permitted higher debt-to-income ratios to be accepted due to a stronger focus on credit scores, and find that this change in lending standards helped to raise homeownership rates. Greenwald (2018) points out that the importance of debt- and payment-to-income limits has remained understudied in macroeconomics, and highlights their importance for monetary policy transmission, especially via their interaction with loan-to-value constraints. These papers however do not study interactions between borrowing constraints and labor supply. A relaxation of debt-to-income constraints has different effects on labor supply incentives than a relaxation of loan-to-value constraints. Relaxing loan-to-value constraints allows households to borrow more against their house. Since

[1] However, Andersen et al. (2021) argue based on Danish data that the return difference between single men and women can be entirely explained by individual and property characteristics, which are not available in the U.S. data of Goldsmith-Pinkham and Shue (forthcoming).

additional borrowing can be used to cushion idiosyncratic income shocks (Braxton, Herkenhoff, and Phillips 2020, Herkenhoff, Phillips, and Cohen-Cole 2020), wives' labor supply incentives are reduced, as there is less need to provide insurance against husbands' negative income shocks (cf. Attanasio, Low, and Sánchez-Marcos 2005, Pruitt and Turner 2020). By contrast, a loosening of debt-to-income constraints has opposing income and substitution effects on labor supply, as discussed above.

Empirical papers have found negative effects of relaxed loan-to-value constraints and increased credit supply on female labor supply (Kumar and Liang 2019, Del Boca and Lusardi 2003, Dao Bui and Ume 2020). In contrast, women in households with larger mortgages, or households being closer to the debt-service constraint based on other than female labor income, have a higher probability to participate in the labor force, and tend to work more (see, e.g., Fortin 1995, Atalay, Barrett, and Edwards 2016 and Appendix Table 1.A.1). Causality can run in both directions, as housing, mortgage and labor supply choices may be determined jointly (Kohlhase 1986, Atalay, Barrett, and Edwards 2016). My results show that a relaxation of income-related borrowing constraints which directly operates on the creditable share of the wife's income has the potential to motivate women to join the labor force, and that the increased labor force participation amplifies the positive direct effects of relaxing the constraint on homeownership and borrowing.

The ECOA paved the way for similar laws in many countries around the world, which are collected in the World Bank's Women, Business and the Law (WBL) database. Figure 1.1 shows that today, 40% of the 190 countries in the WBL panel have explicit laws against gender-based credit discrimination. This however also implies that female access to credit is still not explicitly legally protected in more than half of the world's countries. While this does not necessarily mean that women suffer credit discrimination in all these countries, it still suggests there is scope for improvement. This may be particularly true for developing countries (see also Hyland, Djankov, and Goldberg 2020, Brock and De Haas 2020).[2]

Previous research has shown that strengthening women's financial rights in developing countries can have positive economic and social effects. For instance, Field et al. (2021) find that providing Indian women with access to their own bank accounts induces them to increase their labor supply, and exerts a positive influence on gender norms. My results suggest that developing countries can gain from improving women's inclusion in financial markets, encouraging women's labor force participation and allowing their household to build up wealth. Furthermore, discrimination in lending has recently regained attention also in developed countries due to the increased usage of algorithms and artificial intelligence in lending decisions. Such algorithms can be biased by human decisions in the process of their design, and their complexity can make it harder to detect violations of anti-discrimination laws, which puts new emphasis on the importance of equal access to credit (Morse and Pence 2020, Fuster et al. forthcoming).

The chapter is structured along the following outline. First, I will retrace the historical context of the ECOA in Section 1.2. After briefly describing the data in Section 1.3, I will present the empirical results in Section 1.4. Section 1.5 looks at the labor supply incentives of the act based on a life-cycle model. Section 1.6 concludes.

[2]Only seven OECD countries did not have explicit laws against gender-based credit discrimination by 2019. Out of these, five ranked below the average in the Economist's 2020 "Glass Ceiling Index", which compares 29 OECD countries with respect to women's chances for equal treatment at work. Among them are South Korea, Japan, Turkey and Switzerland, which occupy the lowest echelons in the ranking.

1.2 Historical context

I will start with a short recapitulation of the events leading to the passage of the ECOA. Moreover, I will give a brief overview of the ensuing debate about the law's effectiveness.

1.2.1 Women and credit in the U.S. until the 1970s

Until the 1970s, American women faced various difficulties if they wanted to borrow money. The obstacles they encountered are documented in the Report of the National Commission on Consumer Finance (1972). Single women were often unable to obtain credit, especially mortgages. Lenders commonly denied them credit in their own name, or at least required a (male) cosigner. For married women, it was even more difficult to borrow. As Cairns (1976, p. 967) puts it, married women had "greater difficulty in obtaining credit than [...] any other women".

Upon marriage, women could be required to re-apply for credit, often only under their husband's name. When a couple jointly applied for credit, it was a common practice to discount the wife's income. In certain cases, her income was not counted at all, e.g., if the marriage had lasted for less than five years, or if the couple was of young age (Kendig 1973). Income discounting was especially common if the wife was of "childbearing age" (see also Ladd 1982). Women could be required to sign an affidavit that they were practicing birth control and would not have any more children in order to get a mortgage (Kendig 1973, Cairns 1976). Lenders could even ask for a written confirmation of this from a medical practitioner, known as a "baby letter" (Geary 1976). Many creditors applied stricter standards for applications if the wife, instead of the husband, was the main wage earner (Cairns 1976).

It is hardly surprising that women perceived practices like the "baby letter" as a violation of their privacy. Moreover, women complained about economic disadvantages entailed by the described lending practices. For instance, mortgage credit can provide access to better jobs, education and healthcare facilities by providing the opportunity to move to a different neighborhood. The historian Louis Hyman (2012) emphasizes the importance of credit for women from the upper-middle class as an "indispensable foundation of their economic and social lives" (Hyman 2012, p. 191). Yet credit in the 1970s was not only important for well-off households, but rather a "necessity for all" (Cuomo 1981, p. 126).

While the prevailing lending practices entailed economic disadvantages for women and their families, the economic justifications for maintaining them were less clear. In the 1970s, persistent female employment had become much more prevalent than in the post-war years. The female labor force participation rate had increased from around 33% in 1948 to around 45% in the early 1970s, and kept rising (see Appendix Figure 1.A.1). This trend was mostly driven by married women (Juhn and Potter 2006). A large literature has identified important catalysts of this trend, including structural change and associated shifts in skill demand and skill premia, more favorable working conditions, legal and normative changes, as well as the increased availability of contraceptives and time-saving household appliances (see, e.g., Costa 2000, Juhn and Potter 2006, Greenwood, Seshadri, and Yorukoglu 2005 and references therein).

Of course, women still left the workforce due to pregnancy, but they also returned in increasing numbers (Lally 1974). Importantly, Lally (1974) points out that it would hardly be rational for a woman to leave the labor force to take care of her child if this would lead to a default on the mortgage with the consequent foreclosure.[3] In line with this reasoning, Fortin (1995) shows with Canadian data that women are more likely to work if their household would approach the debt-service constraint without their income. In Appendix Table 1.A.1, I show that a similar

[3] Diamond, Guren, and Tan (2020) show that foreclosures do not only entail substantial financial, but also non-pecuniary costs, and that these are disproportionately borne by the households losing their home.

pattern emerges for U.S. households in the 1970s. I also estimated event study regressions with annual hours worked and labor income as the outcome, and the purchase of a home after renting as the event. The results in Appendix Figure 1.A.2 suggest that married women even slightly increase their hours worked and labor income after a home purchase, and their average labor income follows a similar trajectory as that of their husbands.

Already at the time of the ECOA's passage, economic studies provided evidence that women are on average no worse, or even better, credit risks than men (see, e.g., Lally 1974, Cairns 1976). Kendig (1973, p. 1) concludes that there was "no economic justification for automatically discriminating against women applicants for mortgages". A group of 180 economists even signed a "Statement of Economists" against the "[a]rbitrary exclusion of persons who have the economic capacity to participate in the [mortgage] market place".

The Federal Housing Administration (FHA) had already decided to count the income of most working wives in support of a mortgage as early as 1965, acknowledging that more and more wives were participating in the labor market.[4] However, changes in norms and attitudes can take time. Fernández (2013) shows for the U.S. that the approval rate of married women going to work closely followed the S-shaped profile of married women's labor force participation over time. She develops a model in which intergenerational learning gives rise to cultural change, which initially evolves slowly and then accelerates. In line with this reasoning, it took time until all FHA field office personnel and local lenders became aware of the changes and implemented them (cf. Thurston 2018). Other important market participants, e.g., the Veterans Administration (VA), the Government-Sponsored Enterprises (GSEs) and commercial mortgage lenders, continued to commonly discount the wife's income until the 1970s (Kendig 1973).

In 1972, attempts toward more equal access to credit gained momentum. The Federal National Mortgage Association (Fannie Mae) had abolished previous guidelines that recommended to discount a working wife's income by 50% in December 1971 (Kendig 1973). The above-mentioned "Statement of Economists" was signed in March 1972, and the Equal Rights Amendment to the Constitution, which would have made discrimination against women in mortgage lending unconstitutional, was approved by the U.S. Senate in the same month.[5]

In May 1972, the National Commission on Consumer Finance held congressional hearings on the availability of credit to women. In response to the hearings and the commission's recommendations, several states already enacted laws banning discrimination on the basis of sex or marital status in the subsequent years (see Table 1.B.2). At the federal level, the ECOA was signed into law in October 1974 (Smith 1977). It required "that financial institutions and other firms engaged in the extension of credit make that credit equally available to all creditworthy customers without regard to sex or marital status" (Public Law 93-495, Title V, §502). The Board of Governors of the Federal Reserve System was mandated to write regulations for the implementation of the act, leading to the publication of "Regulation B" in October 1975, the month the ECOA had been scheduled to become effective. Importantly, the ECOA prescribed substantial penalties for violations[6], which were perceived as "adequate to sufficiently deter creditors from willful violations" (Smith 1977, p. 610). In 1976 , the act was amended to include discrimination for any purpose (Public Law 94-239), effective in March 1977 (Smith 1977).

[4] In 1964, Title VII of the Civil Rights Act (Public Law 88-352) had outlawed employment discrimination based on race, color, religion, sex and national origin. The market share of the FHA was around 16% in the 1960s, reached a short-lived peak of 24% in 1970, and then fell to around 7% in the mid-1970s (Golding, Szymanoski, and Lee 2014).

[5] However, it later was not ratified by a sufficient number of states, and finally failed in 1982 (Kendig 1973, Gladstone 2004).

[6] The penalties allowed for up to 10,000 dollars in punitive damages in individual and 500,000 dollars or one percent of the creditor's net worth in class actions, plus attorney costs and legal fees.

1.2.2 The debate about the ECOA's effects and effectiveness

After the enactment of the ECOA, a debate ensued about its effectiveness in extending women's access to credit. Some researchers argued that the act might actually exacerbate households' difficulties to obtain credit if lenders passed on the costs of compliance, for example costs for legal counsel, training of staff or the provision of new application forms (cf. Smith 1977). Smith (1977) reasons that if women are ceteris paribus better credit risks than men, the prohibition to use sex and marital status as predictors in lending models might *decrease* their chances to obtain credit (see also Elliehausen and Durkin 1989). However, Hyman (2012) documents that lending was still hardly based on statistical analysis in the early 1970s. Lenders would commonly apply point systems, assigning, e.g., one point for singles and two for married couples (Lally 1974). Hyman (2012, p. 191) writes: "While limited numeric systems existed, these were rarely based on detailed statistical analysis. Loan officers' everyday prejudices and assumptions more decisively determined credit eligibility." While some lenders already used some form of credit scoring, it was neither widespread nor very elaborate yet (cf. Cairns 1976, Exler and Tertilt 2020).

In the case of married couples jointly applying for a mortgage, some observers were worried that households relying on male and female income might be more risky borrowers than similar households relying on the husband's income alone (cf. Lally 1974). However, the ECOA only prohibited to discount the wife's income *merely* because she was female. It was still allowed to consider aspects like the probability of employment continuity (Geary 1976, Cairns 1976). If they had valid economic reasons to expect a couple to be more risky, lenders could demand higher interest rates (Agarwal et al. 2020, Hurst and Stafford 2004). Yet married couples with and without a working wife paid very similar interest rates after moving to a new home both before and after the ECOA. This remains true when controlling for mortgage and household characteristics (see Appendix Figure 1.A.3).[7]

Analyzing the act from a legal perspective, Cuomo (1981) concludes that it laid the ground for more equality. However, only a few researchers have examined effects of the ECOA empirically. Based on theoretical reasoning against the existence of discrimination in efficient markets and ex-post survey data, Elliehausen and Durkin (1989) conclude that "there is little evidence [...] that the act has increased credit availability to anyone". However, the survey data they use may not adequately capture whether women faced restrictions in credit availability prior to the ECOA, as the questions were asked several years after the first efforts against credit discrimination, and were by default addressed to the husband, as discussed in Appendix 1.E.

Ladd (1982) uses mortgage application data from California and New York for 1977 and 1978 to investigate the prevalence of discrimination against women after the ECOA. She concludes that the ECOA had a positive impact on female access to credit, whereas it was less successful for other protected groups like racial minorities. She suggests that the ECOA's relative success in providing credit access to women may be due to the "rapidly changing role of women in the labor market" (ibid., p. 170; cf. also Hyland, Djankov, and Goldberg 2020). Yet as she does not have pre-reform data, she cannot directly compare the extent of discrimination before and after the ECOA. Haurin and Kamara (1992) take a step in this direction with data from the National Longitudinal Surveys of Young and Mature Women from 1972/73 and 1982/83. They find that the homeownership rate increased between the two survey dates for married women and singles, conditional upon household characteristics, and argue that this is in line with positive effects of the ECOA.

In summary, there is suggestive evidence that the ECOA might have had important positive effects on married women's access to credit, but formal econometric evidence is missing. In the

[7] Based on post-1990 data, Tzioumis (2017) and Jakucionyte and Singh (2020) find performance differences between mortgages with and without co-borrowers, which are however not priced by lenders. However, they show that mortgages with a co-borrower have a significantly *lower* probability of default.

following, I will explore the effects of the ECOA in a difference-in-difference design. I will use data from the PSID, which allows me to track the same households over time, and to include pre-reform characteristics in the estimation.

1.3 Data

The PSID is a widely used representative survey of U.S. families. It was conducted at an annual frequency between 1968 and 1997, and has been continued at a biennial frequency since then. In order to obtain household-level data, I aggregated PSID families who are living together into one household (cf. Pfeffer et al. 2016). I use data from the PSID's "Survey Research Center (SRC) sample", which is nationally representative and tracks households from the first PSID wave over time, as well as the new households formed by former members of these households, e.g., adult children moving out.[8] Survey weights are used for the calculation of descriptive statistics (Hill 1991).[9] Specifically, I use the longitudinal PSID family weights and post-stratify them to match age, race and homeownership from the Current Population Survey (CPS), following the procedure of Kuhn, Schularick, and Steins (2020). All nominal variables were deflated with the CPI, obtained from the *Macrohistory Database* (Jordà, Schularick, and Taylor 2017), such that the reported results are in 2016 dollars.

Before the ECOA, women's access to credit was restricted both for secured and unsecured debt. Comprehensive information on wealth is available in the PSID only since 1984, and was initially only queried every five years. Questions on components of personal debt like student and credit card debt are only available since 2011. However, the majority of household debt in the U.S. consists of housing debt (Bartscher et al. 2020). Information on the outstanding mortgage balance is available in the PSID since 1969, as well as information on monthly mortgage payments. Unfortunately, both of these variables are missing for 1973-1975 and 1982. Information on the asset value of houses and their size is available without interruptions since 1968. Moreover, the survey provides information on the labor income of household heads and spouses, as well as total household income. Additionally, it contains a wide range of demographic variables, such as age, number of children and state of residence.

For some complementary results, I also use data from the "SCF+". This dataset combines the modern Survey on Consumer Finances (SCF) with its historical predecessors. Kuhn, Schularick, and Steins (2020) describe this data source in detail.

1.4 Effects of the ECOA

As discussed in Section 1.2.1, married women arguably had the greatest difficulties of all women to obtain access to credit prior to the ECOA, in particular regarding mortgages. Moreover, even if the share of single households living in an owned and mortgaged property increased after the mid-1970s, it was still comparatively small (see Figure 1.A.5 in the Appendix). Therefore, I will focus my analysis on married households. To guide the analysis, I give a brief theoretical motivation before proceeding to the empirical results.

[8] The full original 1968 PSID sample was a combination of two independent samples: the SRC sample, and an over-sample of low-income families, the "Survey of Economic Opportunity (SEO) sample".

[9] Solon, Haider, and Wooldridge (2015) discuss the pros and cons of using weights when estimating causal effects. I verified that my estimation results are similar with weights. Results are available upon request.

1.4.1 Some brief theoretical background

Let us consider a stylized model of married couples' household debt and housing choices over the life cycle. Households live from period $j = 1$ to J and consist of a male spouse m and a female spouse f who earn labor income y_j^m and y_j^f. Borrowing d_j is only possible if the household owns a house ($h_j > 0$), and is limited by a debt-to-income (DTI) constraint:

$$\phi(y_j^m, y_j^f) = \lambda^y(y_j^m + \lambda^d y_j^f). \tag{1.1}$$

λ^y is the DTI limit, and λ^d determines by how much the wife's income is discounted. Household thus have to maximize their utility subject to a budget constraint and the relevant borrowing constraint, which is zero if the household does not own a house, and otherwise given by equation (1.1), such that $\lambda^y(y_j^m + \lambda^d y_j^f) \geq d_{j+1}$.

Consequently, λ^d only matters for the household problem if the household owns a house ($h_j > 0$) and the DTI constraint is binding. As a first pass, let us abstract from potential changes in labor supply. In the short run, this assumption seems reasonable for two reasons. Firstly, adjustments in labor supply may require search time due to frictions in the labor market. Secondly, banks were still allowed to take employment continuity and income stability into account, such that short-term labor supply increases would not immediately translate into a relaxation of DTI constraints. Potential changes in labor supply are however likely to play a role in the medium to long run, and will be examined in Section 1.5. If λ^d is raised, households which were previously limited by the DTI constraint will be enabled to take out a mortgage and buy a home. The response will be stronger for households in which the wife contributes more to total household income y_j. To see this, consider a household at the DTI constraint, and take the derivative with respect to the income discounting factor λ^d:

$$\frac{\partial d_{j+1}}{\partial \lambda^d} = \lambda^y y_j^f = \lambda^y s_j^f y_j > 0.$$

s_j^f denotes the share of the wife's labor income in household income y_j. The derivative shows that the household will increase its debt if λ^d increases, and the response will be stronger the higher the wife's income share is. Based on these considerations, my hypothesis is that married households in which the wife contributed a larger share to household income have increased their housing debt by more in response to the reform. As I will show in the following, the PSID data provide clear evidence in support of this conjecture.

1.4.2 Descriptive evidence

As discussed in Section 1.2, female labor force participation had increased substantially by the 1970s and kept rising, such that more and more households relied at least in part on female labor income. Figure 1.2 stratifies married households by the wife's contribution to the household's total income. Around 50% of wives did not contribute any earnings to their household's income in the early 1970s.[10] This share steadily declined over the following years, until it leveled off at around 30% in the late 1990s. The share of households in which the wife's earnings make a moderate contribution of less than 30% of total household income increased slightly over the 1970s, stayed more or less constant until the early 1990s, and declined again thereafter. By contrast, the share of households in which the wife's labor income accounts for more than 30% of total household income has increased considerably over time, from less than 20% to 40%.

In 1986, the Reagan Tax Reform Act introduced major changes in mortgage markets (Bartscher

[10] This corresponds to about a third of all (married and non-married) households at the time.

Figure 1.2: Shares of married households by labor income share of wife

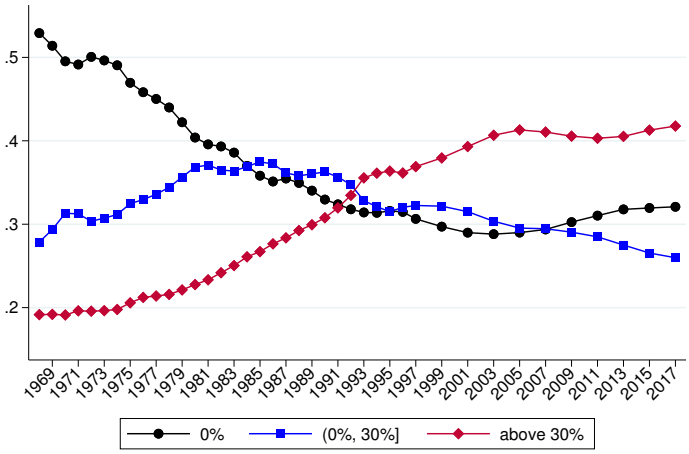

Notes: The graph shows shares among all married households, stratified by the wife's labor income contribution to total household income. The series were smoothed by taking a 3-year moving average.

et al. 2020), and potentially also labor markets (Juhn and Potter 2006, Bastian 2020). Therefore, I will focus on the period until 1985 in the following, which corresponds to a period of ten years after the ECOA's effective date in 1975. Single households are excluded from the sample, and the head is defined to be the male partner in all cases. I will refer to this sample as the baseline sample throughout.

Figure 1.3 shows the trajectory of housing debt for the three groups of married households from Figure 1.2. It reveals that married households with a working wife increased their housing debt more after the ECOA than married households with a non-working wife, for whom there was almost no change over the 1970s. The effect is stronger the more the wife contributes to total household income. In the right panel, I normalize housing debt with total household income to make sure that the divergence in debt does not simply reflect a divergence in incomes. While normalizing by income introduces some more noise, the qualitative patterns remain the same.

Figure 1.3: Housing debt of married households by female income contribution

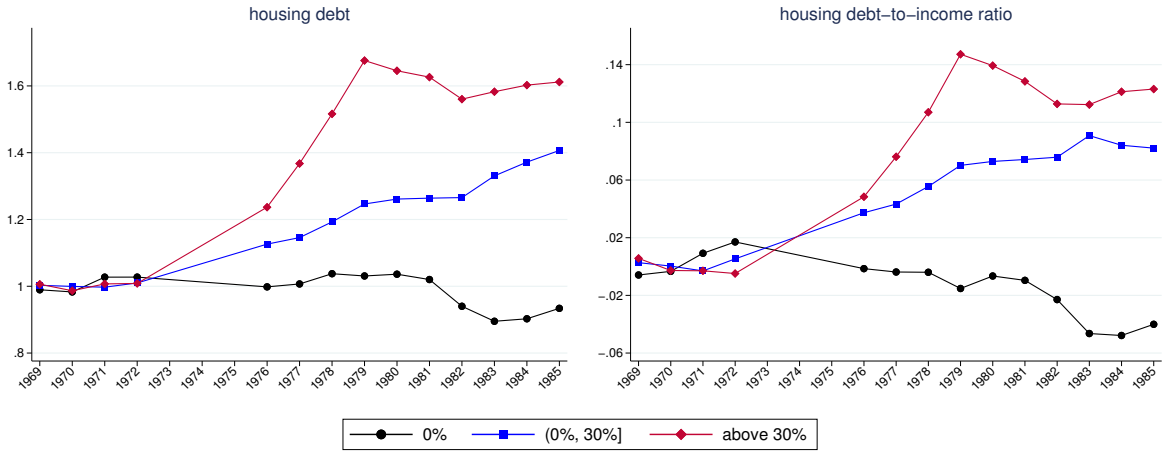

Notes: The graph shows housing debt of married households over time. It differentiates between households in which the wife's labor income accounts for different shares of total household income. All series were normalized with their average over the period 1969-1971. The left panel shows average housing debt, and the right panel shows the average housing debt-to-income ratio, after winsorizing at the 99th percentile within each year. The series were smoothed by taking a 3-year moving average.

The faster debt increase of married households with a working wife subsequent to the ECOA suggests that the possibility to (fully) count the wife's income toward a mortgage enabled married households to buy a home who had previously been restricted by debt-service constraints. After the reform, the extent to which married households relied on the wife's income for mortgage borrowing increased substantially. In the United States of the 1970s, banks viewed 25% of gross income as a critical threshold which mortgage payments should not exceed upon origination (Lally 1974, Gigot 1981). Figure 1.4 shows the shares of married households with a working wife whose debt-service-to-income (DSTI) ratio exceeds 25%. Additionally, it shows the share of households whose DSTI ratio *would be* above 25% if 50% or 100% of the wife's income were discounted.

Figure 1.4: Share of households with debt-service ratio above 0.25

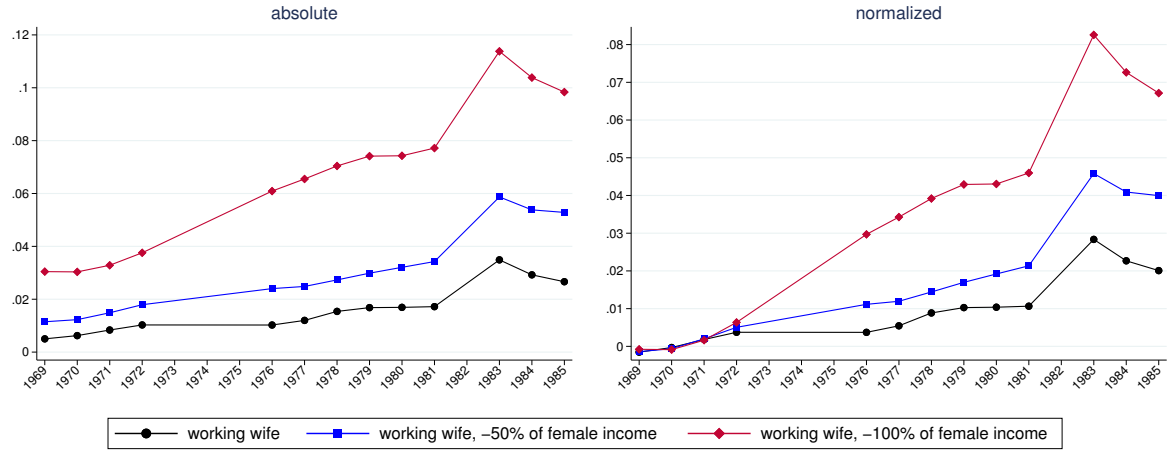

Notes: The left panel shows the share of households with a debt-service-to-income ratio above 25% among households with a working wife. Households with implausibly high ratios (above the 99.8th percentile) were excluded. The blue lines with squares (red lines with diamonds) show counterfactual shares based on debt-service ratios excluding 50% (100%) of the wife's income. The series were smoothed by taking a 3-year moving average. The right panel normalizes each series with its average over the period 1969-1971.

Households with a working wife already relied on female income before the ECOA, as the left panel of Figure 1.4 shows. When the wife's income is partially or fully disregarded, the share of households with a high debt-service ratio in 1969-1971 increases from around 1% to 1.5% or 3%, respectively. However, the reliance on female income grew substantially after the ECOA. This becomes clearer in the right panel of Figure 1.4, which normalizes the series from the left panel by subtracting their 1969-1971 averages. The share of households who would exceed the critical threshold without the wife's income surges rapidly after the reform, demonstrating the increased reliance on female income. More than 10% of married households with a working wife would have found themselves above the threshold without the wife's income in the early 1980s.[11] Looking at the share of households with a high DTI ratio leads to similar conclusions (see Appendix Figure 1.A.6).

A DSTI or DTI constraint typically only binds upon origination of a mortgage. If it is violated at a later point, e.g., due to a short-lived negative income shock, this is usually not sanctioned, as long as the household is still able to make the mortgage payments. Hence, there is no one-for-one mapping between the share of households violating the constraint and the share of households who are unable to service a mortgage. Yet Figure 1.4 suggests that many of the households who took out a mortgage in the period after the ECOA could not have borrowed as much without relying on the wife's income, unless the husbands would have had scope to work substantially more (see also Offer 2007).

[11] Note that some lenders began to relax debt service restrictions in the early 1980s (Gigot 1981).

1.4.3 Difference-in-difference estimates

To gain formal empirical evidence, I estimate simple difference-in-difference regressions. As discussed above, I will treat 1971 as the last pre-reform year, because the first actions toward equal access to credit for women were taken after this year. The empirical specification is as follows:

$$Y_{ist} = \sum_{t=1969,\ t\neq 1971}^{1985} \beta_t \cdot \delta_t \cdot share_i^{pre} + \alpha_{st} + \gamma_i + \Gamma' X_{ist} + \epsilon_{ist}. \tag{1.2}$$

Y_{ist} is the outcome variable, for instance the housing debt-to-income ratio of household i living in state s in year t. δ_t are year dummies and $share_i^{pre}$ is the average share of the wife's labor income in total household income over the period 1969-1971. Labor income includes wages and salaries, as well as the labor portion of other income, e.g., from businesses.[12] I use the average pre-reform income share, because it is predetermined with respect to the reform. A robustness check with contemporaneous income shares is presented in Appendix 1.B.

I control for state-year fixed effects α_{st} to absorb aggregate and state-specific trends, for example in house prices (cf. Offer 2007). Moreover, I exploit the panel dimension of the PSID by including household fixed effects γ_i to absorb all household characteristics that are constant over time. X_{ist} is a vector of time-varying demographic controls.[13] In particular, I include age group dummies and the number of children to control for life-cycle patterns, as well as income group dummies.[14] The inclusion of the state-year and age dummies implicitly controls for cohort membership. Standard errors are clustered at the household and state level.

The coefficients of interest are those on the interaction term, β_t. Given the previous discussion, it can be expected that households profited all the more from the reform the more the wife contributed to the household's income. The idea is therefore to use households with a lower pre-reform female income share as a control group for those with a higher share. One may also estimate a more coarse but slightly simplified version of equation (1.2), using a binary indicator as the interaction variable and thus comparing households with a positive to those with a zero pre-reform female income share. I will present results from this simplified version of the model below in Section 1.4.4, and use them to quantify the aggregate effects of the reform.

As a first outcome, I consider the housing debt-to-income ratio in Figure 1.5a, which plots the coefficients β_t over time.[15] Before 1971, housing debt-to-income ratios evolved similarly for all married households, no matter how much the wife's earnings contributed to total household income. Yet after the ECOA, households increased their housing debt-to-income ratios all the more the larger the wife's pre-reform income share was. The effect is persistent and remains significant at the 95% level until the mid-1980s (apart from one slightly less significant point estimate). Table 1.1 presents difference-in-difference point estimates obtained by replacing the year dummies in equation (1.2) with a dummy for the period after 1971. The point estimate for the DTI ratio in Table 1.1 implies that a household in which the wife's average pre-reform income share was one percentage point higher increased its housing debt-to-income ratio by around 0.21 percentage points (column 1) after the reform. Correspondingly, a thirty-percentage-point increase in the wife's income share translates into a six-percentage-point (30×0.21) increase in

[12] For simplicity, I will refer to the wife's income share in the following. This should always be understood as the share of the wife's labor income.

[13] Additionally, I examined the robustness of my results to the inclusion of an interaction between the predetermined level of the wife's education in 1971 and the year dummies. While this make the specification more demanding, it hardly changes the effects. Results are available upon request.

[14] Following a common classification, I group households into three income categories: bottom 50%, middle 40% and top 10%.

[15] Appendix Figure 1.B.14 shows that a similar pattern emerges with non-normalized housing debt as the outcome.

Figure 1.5: Housing debt

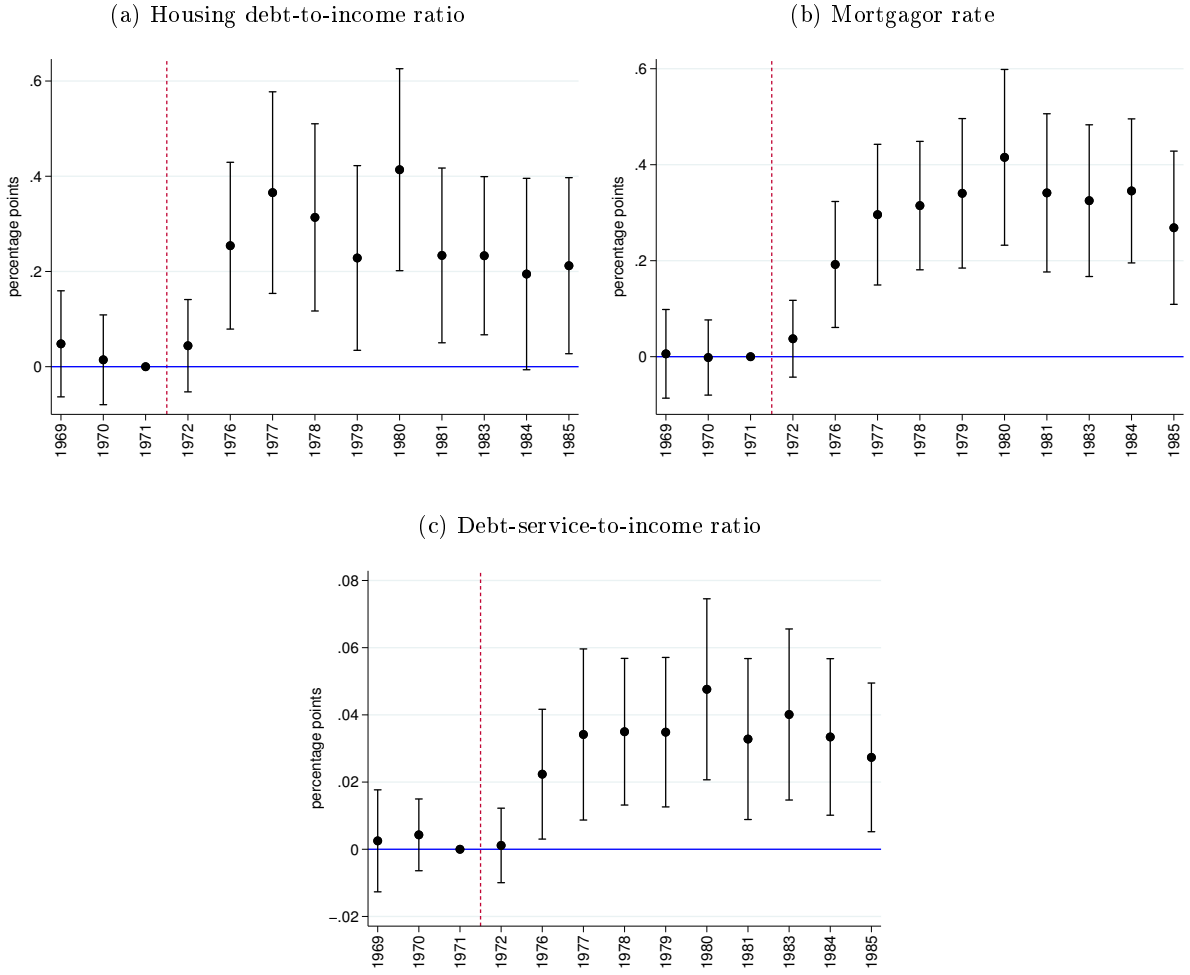

Notes: The graph presents the coefficients on the interaction term in equation (1.2). The base year is 1971. The wife's income share and the DTI and DSTI ratios are defined in percent. The DTI and DSTI ratios were winsorized at the 99th percentile within each year. The whiskers indicate 95% intervals.

debt-to-income, which is approximately 16% of the average in-sample debt-to-income ratio of 36.55%.

Home equity extraction only became popular and widespread in the 1980s (see Bartscher et al. (2020) and references therein). It is therefore likely that the relative increase in housing DTI ratios of households with a high female income contribution is at most partially driven by additional borrowing of incumbent owners against their home. Consistently, the effect operates through the extensive margin. This means that DTI ratios increase because more households take out mortgages, and not because existing borrowers increase their mortgage balance. The clear visual pattern from Figure 1.5b translates into a strongly significant point estimate in Table 1.1.[16] According to the estimate, a one percentage point higher female income contribution leads to a 0.3 percentage points higher probability of having housing debt (column 3). By contrast, the coefficient for log housing debt relative to income in column 2, which captures the intensive margin, is small, insignificant, and actually negative. As the relative share of households with positive housing debt increases among households with a high female income contribution, the relative share of households making mortgage payments increases as well. This leads to a highly significant point estimate for the debt-service-to-income (DSTI) ratio in column 4, while there is

[16] For a comparison of regressions with and without controls, see Table 1.B.1 and Figure 1.B.1.

Table 1.1: Housing debt

	DTI	DTI, int. (log DTI)	DTI, ext. (mortgagor)	DSTI	DSTI, int. (log DSTI)
Post 1971	0.210***	-0.002	0.003***	0.025***	-0.001
× Tot. Inc. Share Wife 71	(0.063)	(0.001)	(0.001)	(0.008)	(0.001)
Controls	yes	yes	yes	yes	yes
Household FE	yes	yes	yes	yes	yes
Time FE	yes	yes	yes	yes	yes
Mean	36.548	3.915	0.513	5.092	2.100
Observations	20,188	10,225	20,190	20,188	10,222

Notes: The table presents results for equation (1.2), after replacing the year dummies with a dummy for the period after 1971. The abbreviations "int." and "ext." refer to the intensive and extensive margin. Standard errors are given in parentheses and are clustered at the household and state level (* p<0.1, ** p<0.05, *** p<0.01). The debt-to-income (DTI) and debt-service-to-income (DSTI) ratios were winsorized at the 99th percentile within each year. The wife's income share and the DTI and DSTI ratio are defined in percent.

no effect at the intensive margin (column 5). Figure 1.5c confirms that also for the DSTI ratio, there were no pre-trends before the reform.

Absent home equity extraction, there are two other plausible reasons for increases in the extensive margin of housing debt. First, households may rely less on other sources of home financing, e.g., borrowing from their family (cf. Del Boca and Lusardi 2003). Second, the homeownership rate may increase. Figure 1.6b confirms that homeownership has indeed increased more for households with a high female income contribution. This relative increase at the extensive margin has lead to a corresponding relative increase in housing-to-income (HTI) ratios (see Figure 1.6a). Table 1.2 contains the related difference-in-difference point estimates. The average HTI ratio rises by around 0.3 percentage points more for every additional percentage point of the wife's average pre-reform income share (column 1), and the likelihood of homeownership increases by 0.2 percentage points (column 3).

Like the debt effect, the effect on housing mainly works through the extensive margin. The coefficient on the log housing-to-income ratio in the second column of Table 1.2 is positive, but very small and not significant. By construction, the loan-to-value (LTV) ratio is only defined

Table 1.2: Housing

	HTI	HTI, int. (log HTI)	HTI, ext. (homeowner)	LTV	number of rooms
Post 1971	0.327***	0.001	0.002***	-0.117***	0.004**
× Tot. Inc. Share Wife 71	(0.120)	(0.001)	(0.000)	(0.039)	(0.001)
Controls	yes	yes	yes	yes	yes
Household FE	yes	yes	yes	yes	yes
Time FE	yes	yes	yes	yes	yes
Mean	157.611	5.073	0.816	25.803	5.868
Observations	26,767	21,768	26,776	16,353	26,289

Notes: The table presents results for equation (1.2), after replacing the individual year dummies with a dummy for the period after 1971. The abbreviations "int." and "ext." refer to the intensive and extensive margin, respectively. Standard errors are given in parentheses and are clustered at the household and state level (* p<0.1, ** p<0.05, *** p<0.01). The housing-to-income (HTI) and loan-to-value (LTV) ratios were winsorized at the 99th percentile within each year. The wife's income share and the HTI and LTV ratio are defined in percent.

Figure 1.6: Housing

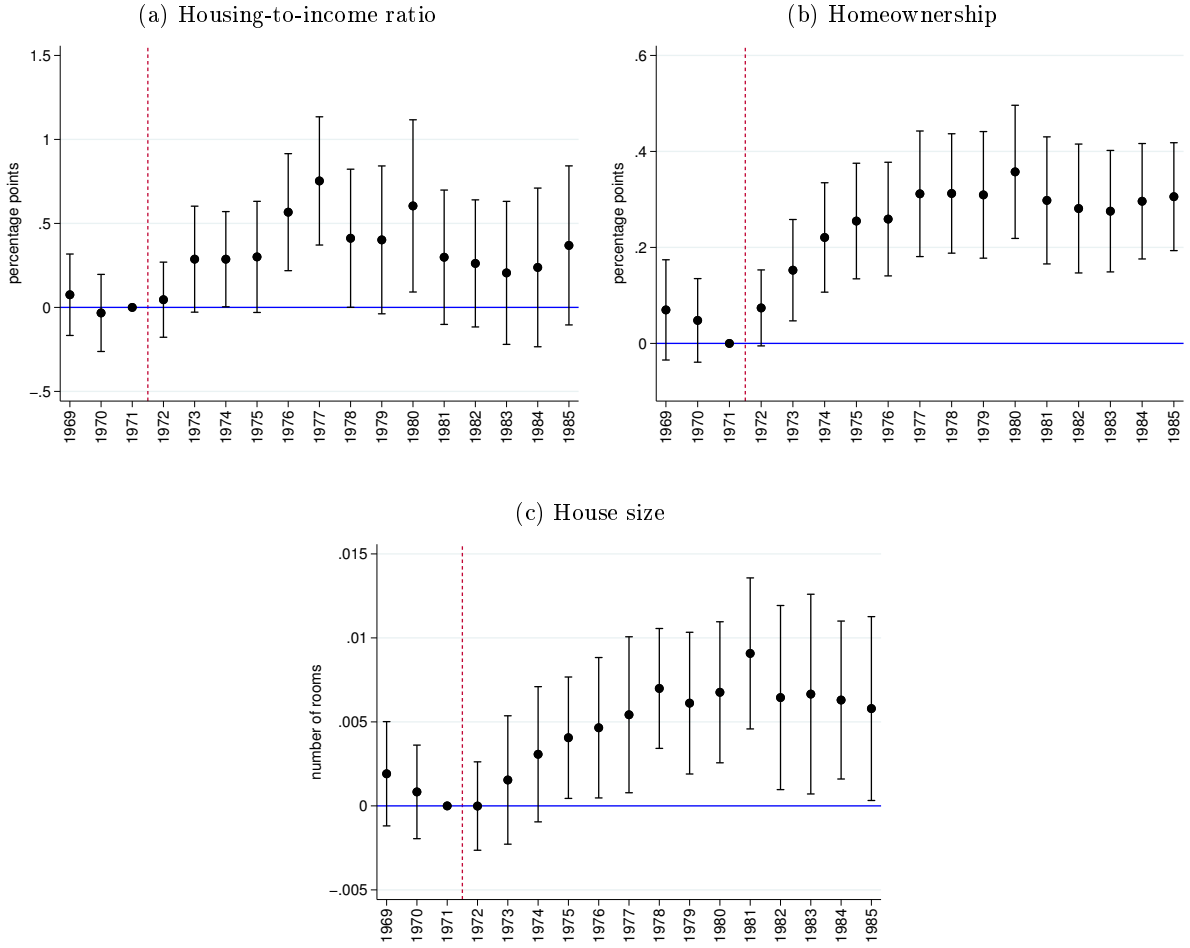

Notes: The graph presents the coefficients on the interaction term in equation (1.2). The base year is 1971. The wife's income share and the HTI ratio are defined in percent. The HTI ratio was winsorized at the 99th percentile within each year. The whiskers indicate 95% intervals.

for homeowners, and thus does not reflect changes at the extensive margin of housing. Together with the small negative effect at the intensive margin of housing debt, the small positive effect at the intensive margin of housing translates into a negative LTV effect (column 4). Finally, I find that households with a higher female income contribution increase the size of their home by more after the reform (see Figure 1.6c and column 5 of Table 1.2). This effect is mainly driven by households moving from a smaller rented to a larger owned property.[17] A few households also upgrade to larger homes, in line with the small positive intensive-margin effect.

The main focus of this chapter is on married households, because married women arguably had the greatest difficulties to obtain credit in the pre-reform era. However, singles may also have benefited from the reform, given that it precluded credit discrimination related to both sex and marital status. To investigate whether their access to credit improved as well, I estimated similar regressions for singles (see Appendix 1.C). Appendix Figure 1.C.3 shows that singles indeed borrowed more subsequent to the ECOA, compared to married households with a non-working wife. However, there is no evidence of an associated increase in homeownership. A possible interpretation is that singles could substitute away from other financial sources, such as loans or transfers from relatives (cf. Del Boca and Lusardi 2003), which are generally less

[17] In my baseline sample, rented homes have 4.9 rooms on average, compared to 6.1 for owned properties.

important for married couples. Appendix Figure 1.C.4 shows that the responses were similar for single men and single women.

1.4.4 Translation into aggregate effects

How do the estimated effects at the micro level translate into macro variables such as the homeownership rate? In order to quantify the effects on overall mortgage borrowing and homeownership, I replace $share_i^{pre}$ in equation (1.2) with a binary indicator equal to one if $share_i^{pre}$ is positive, and the year dummies with a dummy equal to one after 1971. This strategy allows me to directly compute the total effect on households with working wives.[18] Moreover, the resulting estimate can be compared more easily to the effect size under the alternative event study identification in Section 1.4.5. The point estimates for the borrowing and homeownership rates are summarized in Table 1.3, and the dynamic effects are shown in Appendix Figure 1.B.2. The point estimate of around 0.08 for the extensive margin of mortgage debt, multiplied by the average share of married households with a working wife in 1971 ($\approx 34.5\%$) and the total number of households in this year ($\approx$ 64,778,000) yields a number of 1.8 million additional borrowers. The analogous calculation for homeownership implies that 1.4 million households were enabled to buy their own home. This corresponds to a ceteris-paribus change of around 3.3 percentage points in the average homeownership rate of married households.[19]

Table 1.3: Point estimates with binary interaction

	mortgagor rate	homeowner-ship rate
Dummy ×	0.083***	0.062***
Post 1971	(0.021)	(0.014)
Controls	yes	yes
Household FE	yes	yes
Time FE	yes	yes
Mean	0.513	0.816
Observations	20,190	26,776

Notes: The table presents the results of estimating a binary version of equation (1.2), where the interaction term is replaced by the interaction between a dummy for whether the average pre-reform share of the wife's labor income was positive and a dummy for the post-reform period after 1971. Standard errors are given in parentheses and are clustered at the household and state level (* p<0.1, ** p<0.05, *** p<0.01).

As discussed by Cuomo (1981), the ECOA did not abolish all credit discrimination against women. For instance, married women without an own income could still not obtain credit without their husband's approval in separate property states (ibid.) However, he acknowledges that the ECOA provided more comprehensive protection to working wives earning their own income. My results show that this increased protection translated into sizable and significant positive effects on access to credit for working wives and their households. In particular, more married households with working wives could take out mortgages to become homeowners and increase the size of their home. In the following, I will examine the robustness of my baseline findings.

[18] Alternatively, I could use the point estimates from the third columns of Tables 1.1 and 1.2, and aggregate over households with different female income contributions. The resulting numbers are similar (1.3 million more homeowners and 1.6 million more borrowers).

[19] For comparison, the average homeownership rate of married households in the data changed by around 6.5 percentage points between 1971 and 1985.

1.4.5 State-level evidence and robustness

While the results of the national-level difference-in-difference regressions are consistent with plausible effects of the ECOA, they may still, at least partly, capture effects of other contemporaneous events that might have affected households with working wives differentially, such as the return of veterans from the Vietnam War. While I perform an extensive set of robustness checks in Section 1.4.5, the institutional setup allows me to provide even stronger evidence based on state-level variation. As described in Section 1.2.1, some U.S. states already introduced laws against discrimination before the ECOA became effective at the federal level. This setup opens up the possibility to exploit variation across states for identification. Such an approach is much more robust to potential confounding events, because these events would have to occur both in the same states and at the same time as the introduction of anti-lending-discrimination laws. For this reason, I use information on the relevant state laws and estimate event study regressions to assess whether my previous findings are consistent with those obtained from this alternative identification strategy. In Section 1.4.5, I then explicitly consider likely candidates for potential confounding events at the national level, and investigate the robustness of the federal-level results against the background of these events. Additional details and figures can be found in Appendix 1.B.[20]

State-level variation

According to the U.S. Department of Labor (1975, p. 384), 40 states had "legislation or regulations on some aspect of discrimination in credit based on sex and/or marital status" as of April 1975. Table 1.B.2 in the Appendix provides an overview over these laws. States which introduced anti-credit-discrimination legislation earlier might have witnessed effects on homeownership and housing debt earlier. To test this hypothesis, I estimate event study regressions similar to the specification in equation (1.2):

$$Y_{ist} = \sum_{j=\underline{j}}^{j=\overline{j}} \beta^j \cdot D_{st}^j + \alpha_s + \delta_{rt} + \gamma_i + \Gamma' X_{ist} + \varepsilon_{ist}. \tag{1.3}$$

α_s are state fixed effects, δ_{rt} are census-region-times-year fixed effects and the remaining notation is as in equation (1.2). The region-times-year fixed effects are included to capture differences in house price trends across census regions (cf. Bartscher et al. 2020).[21] The event time index j denotes periods relative to the event e_s in state s, covering an event window from $\underline{j}$ to $\overline{j}$. The treatment period e_s is chosen as the year when the law in state s became effective. For all states which had not introduced any state-level legislation against credit discrimination in home financing before the effective date of the ECOA in 1975, I choose 1975 as the treatment year. This includes seven states whose laws did not pertain to home financing (see Table 1.B.2). In order to achieve identification in this setup, I drop observations with $t \geq e_s + H$, where H denotes the difference between the latest and earliest event year in the data, as recommended by von Bismarck-Osten, Borusyak, and Schönberg (2020).[22]

Given the data availability, I focus on the housing outcomes.[23] I consider households in which the wife had a positive average labor income over the three periods prior to the event, and compare their outcomes between treated and untreated states before and after the event. Note that in

[20] Some robustness checks are not reported for the sake of space. All results are available upon request.
[21] State-level house price indices are only available since 1975.
[22] Alternatively, one may "bin" the end points together to identify the dynamic treatment effects when there are no never-treated units (Schmidheiny and Siegloch (2020)).
[23] Due to the timing of the reforms, I cannot use the mortgage-related variables, as I do not observe them between 1973 and 1975.

Figure 1.7: Event studies housing

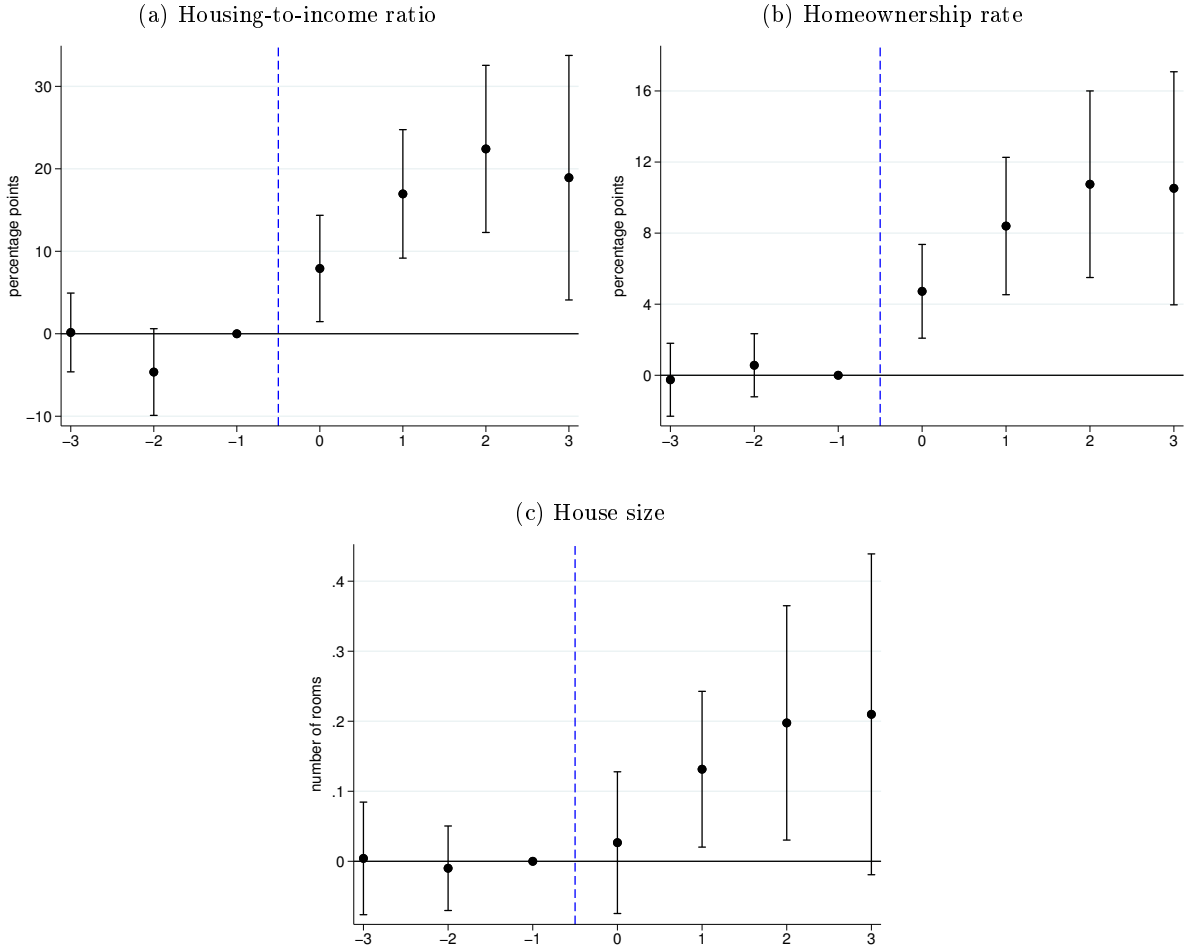

Notes: The graph shows the coefficients on the treatment indicators D^j_{st} from equation (1.3). The sample was restricted to households in which the wife had a positive average labor income over the three years prior to the event. Observations with $t \geq e_s + H$ are excluded, where H is the gap between the latest and earliest event year. Standard errors are clustered at the state level. The whiskers indicate 95% intervals.

the national-level difference-in-difference regressions, the effects are identified from a comparison between households with different female income contributions (or households with and without a working wife in the binary case) before and after the reform. However, one may be worried that households with a working wife are not fully comparable to households with a non-working wife for potentially unobservable reasons. While I include demographic controls and household fixed effects in my federal-level regressions, a further advantage of the event study design is its immunity to this concern, because the identification here comes from a comparison of households with a working wife in states which have already introduced anti-discrimination legislation in a given year to similar households with a working wife in states which have not introduced such legislation yet.[24]

The results are shown in Figure 1.7. Households with a working wife in the pre-reform period increased their housing-to-income ratio, homeownership rate, and house size significantly more

[24] Favara and Imbs (2015) mention the possibility that if a neighboring state has already implemented a reform, houses near the border in a no-reform state could be financed with a loan from a lender in the reform state. However, they find that such arbitrage only happens to a significant extent in immediate proximity to a border (less than 15 miles). The publicly available PSID data do not allow to observe if a household lives near a border. Arbitrage would however weaken the difference between reform and non-reform states, reducing the likelihood of finding a positive effect.

in treated states than in control states after the reform. By contrast, there were no differences across states in the pre-reform years. The homeownership effect stabilizes at slightly above ten percentage points, which is similar to the results from the binary difference-in-difference specification at the federal level shown in Appendix Figure 1.B.2b.[25] The dynamic pattern also resembles the one found in the difference-in-difference regressions, with a gradual build-up of the effect, consistent with the fact that housing is a slow-moving variable.

Recent work in econometrics has shown that for difference-in-difference regressions with variation in treatment timing, the estimation with two-way fixed effects may be biased in the presence of heterogeneous treatment effects (see Goodman-Bacon 2021, Sun and Abraham 2020 and the references therein). It should be noted that treatment effects may unfold dynamically over time under the assumption of homogenous treatment effects, but they may not differ across different treatment cohorts (Schmidheiny and Siegloch 2020, Sun and Abraham 2020).[26] Appendix Figure 1.B.11 shows that the results obtained with the "interaction-weighted estimator" of Sun and Abraham (2020), which is robust to heterogeneous treatment effects, are very similar to those obtained from estimating a comparable model with two-way fixed effects. The fact that the effects obtained from the alternative event study identification align well with those found in the difference-in-difference setup lends support to a causal interpretation of the national-level results.

Contemporaneous events

The event study results provide evidence in support of the hypothesis that the national-level results do not simply reflect the effects of other events in the reform period. In order to lend further credibility to my results, I examine plausible candidates for potential confounding events in more detail, and assess the robustness of the national-level difference-in-difference results to corresponding control strategies.

One potential concern is that the entrance of the baby boomer cohort born 1946-1964 into housing markets in the 1970s might be driving my results. However, I implicitly control for cohort membership via the age and time fixed effects. Moreover, I verified that the same patterns emerge when restricting the sample to households with a wife or head below 30 or 35 years of age in 1971. Furthermore, I investigated the robustness of the results to exploiting the cohort dimension in a triple-difference design. As noted in Section 1.2, the likelihood of a working wife's income being discounted was lower if she was no longer of childbearing age. It can thus be expected that women who were already older at the time of the reform were less affected. Based on this reasoning, I added the difference between younger and older cohorts to my regressions (see Section 1.B.2). Appendix Figure 1.B.3 shows that the triple-difference identification supports the results from the difference-in-difference estimations. The point estimates are even slightly larger than in the baseline, consistent with the hypothesis that the effects are stronger for younger households.

Leombroni et al. (2020) argue that the Great Inflation of 1965 to 1982 induced portfolio shifts from equity to housing in the 1970s.[27] I therefore want to make sure that my results are not driven by differential stock market exposure of married households with different female income contributions. Due to lacking information on financial assets in the data, I cannot directly control

[25] The same is true for the other two outcomes. Results are available upon request.

[26] For instance, heterogeneous effects may arise if treatment units select their initial treatment timing based on treatment effects. This would however not violate the parallel trends assumption, which is only violated if the timing is chosen based on the evolution of the outcome (Sun and Abraham 2020).

[27] The mechanism is that higher expected inflation lowers the expected after-tax real return of equity, but tends to reduce the user cost of housing, because housing capital gains are quasi untaxed in the U.S., whereas mortgage interest is tax-deductible (Poterba 1984).

for stock ownership. However, the SCF+ data show that the income-richest 10% of households held 70% to 80% of all stock wealth in the 1970s. At the same time, the effect of inflation on the user cost of housing via taxes is more pronounced for high-income households (Poterba 1991). I hence assessed the robustness to excluding top-10% households. Reassuringly, the patterns are very similar to the baseline.

In 1972, another equal opportunity law was passed: the Equal Employment Opportunity Act (EEOA). The EEOA authorized the Equal Employment Opportunity Commission of 1965 to initiate lawsuits on behalf of workers, and expanded the coverage of Title VII of the 1964 Civil Rights Act, which outlawed discrimination in wages and employment opportunities, to employers with more than 15 instead of 25 employees (Hill 1977). Yet the original version of Title VII became effective a whole decade prior to the ECOA, and the Equal Pay Act had already mandated equal pay for equal work of men and women in 1963 (Altonji and Blank 1999). Apart from that, I use the average female income share before 1972 as the interaction variable, which is independent of the EEOA. Theoretically, women with a higher pre-1972 income share could have been more likely to sue for higher wages after 1972, boosting their total household income. However, I control for household income and consider outcomes normalized by income to make sure that the effect is not driven by simple differential changes in available household income.

The year 1973 was also marked by a severe oil crisis. However, cars were of course not only used by working wives at the time, but also by homemakers, who had become an important clientele for car makers and were actively targeted in advertisements (Hill 2002). As a validation, I controlled for the household's total number of cars as a proxy for their exposure to oil prices, or the households' average commuting costs in 1970 and 1971. The results remain virtually unaltered.

The 1960s and 1970s also saw important changes in birth-control and divorce laws. "The pill" spread rapidly among married women after its introduction in 1960, facing a peak in 1967 (Goldin and Katz 2002). The diffusion among singles was delayed because their access to oral contraceptives was initially impaired by state laws on the age of majority (see also Bailey 2006). The abolition of these laws in the 1960s and early 1970s predominantly affected single women. However, one might imagine that the households of married women were enabled to borrow more because the wife had had early access to the pill before marriage, and could hence obtain a higher level of education (Goldin and Katz 2002) or gain more work experience (Bailey 2006), making her more creditworthy. I therefore used the data on state consent laws determining early access to the pill from Bailey (2006), and re-estimated the state-level event studies after including additional event dummies for these laws. Appendix Figure 1.B.6 shows that very similar patterns emerge.

With more liberal divorce laws, financially independent women might have a higher likelihood of divorce. Appendix Figure 1.B.4 however shows that divorce rates already started to increase around 1965. Chang (2018) shows that a liberalization of divorce laws puts downward pressure on married households' homeownership, as houses are indivisible and thus cannot be split easily upon divorce. If within the group of married households with working wives those without a house were more likely to get divorced and leave the sample, while the same was not true for those with a non-working wife, this might lead to a larger increase in homeownership among the former. Yet while I find a negative association between homeownership and the likelihood of divorce in the data, I find no evidence that this relationship would depend on the wife's earnings. Moreover, I restricted the sample to households who had entered the PSID at the latest in 1971, were still in the survey in 1985, and never reported a transition from married to single in between. Although this reduces the sample size by about a third, Appendix Figure 1.B.5 shows that the baseline pattern survives. Finally, I re-estimated the event studies after controlling for additional event dummies based on the state-level introduction of unilateral divorce from Gruber (2004), finding similar results to the baseline (see Appendix Figure 1.B.6).

Previous research has demonstrated the positive effects of inter- and intrastate bank branching deregulation on homeownership and mortgage credit (Favara and Imbs 2015, Tewari 2014). While interstate branching deregulation only occurred after 1994, intrastate branching restrictions were already lifted during the 1970s in six states[28]. I verified that the effects from both the national- and state-level regressions survive in a sample excluding these states. The 1970s also saw several housing and mortgage market reforms at the federal level, such as the Housing and Community Development Act of 1974, the Emergency Home Purchase Assistance Act of 1974 or the Emergency and Housing Act of 1975 (cf. Milgram 1994). There is no evident reason why any of these reforms should have differentially benefited households with different (pre-determined) female income contributions, in particular after controlling for total household income. However, as some of these reforms explicitly targeted low-income households and households facing temporary economic hardship, I verified that similar results emerge when only considering households between the 50th and 90th percent of the income distribution, or excluding households in which the head was temporarily unemployed in the mid-1970s.

Another important event in the 1970s was the end of the Vietnam War in 1975. Returning veterans could obtain cheaper loans due to insurance by the Department of Veterans Affairs (VA) (Foote and Peterson 2008). One could imagine that many of them were single, and that if they were married, the wife contributed more to the household's income while her husband was at war. However, most U.S. troops had already been withdrawn from Vietnam by 1972 (see Appendix Figure 1.B.7). As a robustness check, I re-estimated the regressions for married households on a sample excluding all households whose head had been in the armed forces any time between 1969 and 1975. Of course, this also excludes households whose head was in the armed forces without being a veteran. Nevertheless, Appendix Figure 1.B.8 shows that the results are still similar for this restricted sample.

Finally, Bastian (2020) argues that the introduction of the Earned Income Tax Credit (EITC) in 1975 induced around one million American mothers to join the labor force. The EITC is an earnings subsidy to working parents. In its original design, parents with annual nominal earnings of up to 4,000 dollars could obtain a maximum subsidy of 400 dollars. Reduced subsidies were available to households with earnings up to 8,000 dollars. It is important to note that households were eligible for the EITC independent of the wife's work status, as long as the household's total labor income did not exceed the thresholds, and at least one child lived in their home. In order to test if my results nevertheless simply capture effects of the EITC, I excluded all eligible households from my sample and only retained households with total nominal earnings above 8,000 dollars, or without children living at home, for the years after 1975. Appendix Figure 1.B.9 shows that the baseline results remain intact.

Alternative variable specifications

Lastly, I also tested the robustness of my results to the use of alternative outcome and interaction variables. As both income and housing or housing debt might change in response to the reform, I chose to use the wife's pre-reform income share as my interaction variable in the baseline. I verified that similar qualitative patterns emerge when using the wife's contemporaneous income share (see Appendix Figure 1.B.13). In Appendix Figure 1.B.14, I further use non-normalized housing debt and assets as the outcome. Both variables increase after the reform, although the effect on housing is estimated with less precision.

In the baseline analysis, I use the share of the wife's *labor* income in *total household income*, since comprehensive information on other income of the wife is only consistently available since 2005 in the PSID. However, labor income accounts for the largest share of most households'

[28] Maine (1975), New Jersey (1977), New York (1976), Ohio (1979), Vermont (1970) and Virginia (1978)

income, with a median of 83% between 1969 and 1985. As a robustness check, I used the wife's share in the total *labor* income of head and spouse. The results are reported in Table 1.B.3 and Figure 1.B.16 in the Appendix, and resemble the baseline. The effect is somewhat smaller, with an increase in the mortgage borrowing rate of 0.1 percentage points for a one percentage point higher average pre-reform share in the couple's total labor income. There are some households in which the wife earns a substantial share of the couple's total labor income, but still a small share of total household income. For instance, the wife may earn a small salary on a part-time job, whereas the head does not work for money at all, but receives substantive capital income. In such a case, the wife may contribute 100 percent to the household's *labor* income, but her contribution to overall income is still minor and thus will not substantially affect the household's borrowing capacity. A somewhat smaller effect can thus be expected.

1.5 Labor supply incentives

Upon its introduction, the ECOA primarily affected households in which the wife had already been working, given that banks were still allowed to take employment continuity and income stability into account. However, the ECOA might also have incentivized women to increase their labor supply in the subsequent years in order to benefit from the new lending rules. The theoretical effects on labor supply are ex ante ambiguous. In order to buy a house of a given quality and size, a woman can afford to work less if 100% instead of 50% of her income are considered. On the other hand, an additional dollar of income now translates into a higher borrowing capacity one-for-one, making labor supply more attractive. Consider the stylized example of a woman with an earnings potential of 30 dollars. Her husband currently earns 30 dollars, and their housing preferences are such that a suitable home would at least require a mortgage of 120 dollars. With a maximum DTI ratio of 2, they would need at least 60 dollars of income. As long as the couple can only borrow against half of the wife's income, i.e. 15 dollars, they will still not be able to buy their desired home even if the wife is working. Yet once her full income can be used, there is a strong incentive for her to take up work.

Given that employability declines with age, whereas the probability of already being a homeowner increases (and the empirical analysis has shown that the effects mainly work though the extensive margin, i.e. transitions from renting to ownership), positive employment incentives can be expected to be strongest for young women living in recently founded households. Yet the empirical setup, which relies on comparisons between the pre- and post-reform period, does not allow to study the behavior of households that are only formed after the reform. Would a woman at the beginning of her (economic) life cycle make different plans for her financial and professional future in the post-reform world compared to the pre-reform world? To answer this question, I use a structural model of households' homeownership and borrowing over the life cycle.[29]

Based on a life-cycle model calibrated to data from the British Household Panel Survey (BHPS) for the period 1991 to 2002, Bottazzi, Low, and Wakefield (2007) show that the empirical correlation between large housing debt and longer female hours worked can be generated by the requirement to meet current mortgage obligations. They further show that a tightening of the DTI constraint, which operates on both male and female income in their model, leads to delayed home purchases. They do not report effects on female labor force participation (FLFP), but note that the effects they found were small. Pizzinelli (2018) calibrates a similar two-earner model to BHPS data from 1991 to 2008, and allows the income of the secondary earner to influence the

[29] As a proxy, one may look at households that have already been formed, but do not own a home yet and the wife is not working or working very little. Results from an explorative event study are shown in Appendix Figure 1.B.12 and suggest a positive effect on labor force participation, but should be interpreted with caution due to the small sample size.

DTI constraint in a different way than that of the primary earner. In his specification, the DTI always depends on full-time earnings, i.e., the earnings that the secondary earner would have when working full time. He simulates a relaxation of this DTI constraint by increasing the multiplier on the secondary earner's full-time labor income, and finds a positive labor supply effect. His results further imply that the secondary earner's LFP response can amplify the increase in homeownership.

In the following, I will build a life-cycle model similar to the frameworks in these papers. After describing the model and my strategy to calibrate it to the early 1970s in the following subsections, I will use it as a laboratory to examine whether and to what extent the act had the power to change married women's labor supply incentives.

1.5.1 Model

My goal is to examine whether married women changed their financial and career planning at the beginning of the life cycle in response to the reform. To do so, I build a simple life-cycle model with borrowing constraints. Households are formed at age $j = 25$. Both spouses retire at age 65, and die at age 80.[30] They maximize their utility over consumption c_j, female labor supply n_j^f and housing h_j:

$$u(c_j, n_j^f, h_j) = \frac{c_j^{1-\sigma}}{1-\sigma} + \theta^f(j)\frac{(1-n_j^f)^{1-\psi}}{1-\psi}\mathbb{I}_{j\leq 65} + \left[\chi(j)\mu^h + \phi^h\frac{h_j - \underline{h}}{\overline{h} - \underline{h}}\right]\mathbb{I}_{h>0}.$$

The parameter σ determines the degree of risk aversion. $\theta^f(j)$ determines the strength of preferences for a working-age wife's time spent on other activities than market work, which can vary over the life cycle to reflect changes in the disutility of working induced by, e.g., the presence of small children in the household or the need to take care of grandchildren. ψ governs the elasticity of female labor supply. The strength of basic housing preferences depends on the parameter μ^h, and ϕ^h determines the preference for a larger house. Households can buy houses of a smaller size $\underline{h} = 1$ or a larger size $\overline{h} = 2$. Following a suggestion of Druedahl (2015), preferences for owning versus renting are allowed to vary over the life cycle, reflecting factors such as changing mobility preferences and changes in demand for space. I model these time-varying preferences in a reduced-form way, similar to Pizzinelli (2018), by pre-multiplying μ^h with an age-dependent factor $\chi(j)$ (see Section 1.5.2 and Appendix 1.D.2 for details).

Log hourly wages w_j^s are modeled as the sum of a deterministic function of age and an autoregressive process:

$$\begin{aligned} ln(w_j^s) &= \alpha_0^s + \alpha_1^s j + \alpha_2^s j^2 + \alpha_3^s j^3 + \alpha_4^s j^4 + ln(z_j^s), \quad s = m, f \\ ln(z_j^s) &= \rho^s ln(z_{j-1}^s) + \varepsilon_j^s, \quad \varepsilon_j^s \sim N(-\frac{\sigma_{\varepsilon^s}^2}{2}, \sigma_{\varepsilon^s}^2). \end{aligned} \tag{1.4}$$

Both spouses' wages are subject to idiosyncratic risk. Men always have a standard full-time contract, corresponding to 40 hours per week. Women can choose whether to work or not, and whether to work full-time or part-time. Specifically, they can choose to work 20, 30, 40 or 50 hours a week. In other words, they choose between discrete contracts in the set $\mathcal{N} = \{0, 20/T, 30/T, 40/T, 50/T\}$, which express n_j^f as hours worked relative to the total number of non-sleeping hours per week $T = 7 \times 16 = 112$.[31] Working less than 40 hours a week is associated

[30] The household's age is defined to be the age of the head.

[31] I experimented with different specifications of $\mathcal{N}$. Contracts with more than 50 hours are never chosen under plausible calibrations. This is in line with the data, where less than 1 percent of women report such high working hours in the early 1970s.

with a part-time penalty $\mathcal{P}$. After age 65, both spouses receive a retirement income based on the replacement rate b. The retirement income is calculated as b times the realization of the wage process in the last pre-retirement period.[32]

Households can save in a risk-free financial asset a_j at the interest rate $r = r^s$. They can also borrow $d_j = -a_j$ at a rate of $r = r^b$. This implies that households cannot hold both positive financial assets and mortgage debt at the same time, as in Attanasio et al. (2012) and Pizzinelli (2018).[33] Borrowing is limited by the minimum of a LTV and DTI constraint:

$$\phi(h_j, y_j^m, y_j^f) = min\left\{\lambda^h p h_j,\ \lambda^y(y_j^m + \lambda^d y_j^f)\right\}, \tag{1.5}$$

where λ^h is the LTV limit and p is the house price. This implies that there are three options: (a) the household does not own a house, in which case asset holdings must be positive; (b) the household owns a home, but the LTV constraint ($\lambda^h p h_j \geq d_{j+1}$) binds before the DTI constraint; or (c) the household owns a home, and the DTI constraint ($\lambda^y(y_j^m + \lambda^d y_j^f) \geq d_{j+1}$) binds before the LTV constraint.

Mortgages are modeled as long-term debt, meaning that households must only fulfill the LTV and DTI constraint upon origination of the mortgage (Attanasio et al. 2012, Kaplan, Mitman, and Violante 2020). An alternative modeling option would be to impose a DSTI instead of a DTI constraint. However, the DSTI and DTI ratio are directly proportionate for a mortgage with fixed duration and interest rate, as shown in Appendix 1.D.1.[34] An advantage of formulating the constraint in terms of debt rather than mortgage payments relative to income is that it can be compared directly to the LTV constraint, as illustrated in Appendix Figure 1.D.1.

I follow Attanasio et al. (2012) and require that households at least pay the interest $r^d d_j$ on their outstanding debt. As shown in Appendix 1.D.1, this is equivalent to the constraint $d_{j+1} \leq d_j$. Since the model is meant to describe the early 1970s, when home equity withdrawal was still uncommon, households are not allowed to extract home equity.[35] Households close to or in retirement (above 60 years of age) cannot take out new debt, but can continue to repay their pre-existing debt. This implies $\phi(\cdot) = 0\ \forall\ j \geq 65$.

Since it is not the interest of this chapter to model debt and homeownership in old age, I abstract from bequests. In the last period of life, households are required to pay all outstanding debt and consume what is left of their wealth. Households choose c_j, a_{j+1}, n_j^f, and h_j, given the current vector of states $X_j = [a_j, h_{j-1}]$. The model is in partial equilibrium. A household can either own a house ($h_j \in \{\underline{h}, \overline{h}\} > 0$) at prices $\underline{p}$ and $\overline{p}$, or rent one ($h_j = 0$) at cost q. Selling or buying a house is subject to transaction costs F^s and F^b, reflecting costs for real estate agents, legal fees and transaction taxes (cf. Attanasio et al. 2012).

In summary, households have to solve a utility maximization problem subject to a budget constraint and borrowing constraints which depend on whether the household is a renter, owns a "regular-size" house ($h_j = 1$) or owns a large house ($h_j = 2$). The wife's labor income depends on her labor supply choice $n_j^f \in \mathcal{N}$, where her wage is discounted by part-time penalties $\mathcal{P}$ if she chooses to work less than full-time. Mathematically, the problem can be summarized as follows:

[32] Women's retirement income is computed based on a full-time contract.

[33] As discussed by Druedahl (2015), this simplification comes at the cost of ruling out precautionary balance sheet expansions. He shows that the difference between the net and gross debt formulation of the model of Attanasio et al. (2012) is less important if there is an interest rate spread.

[34] Adjustable-rate mortgages were allowed only in 1982 (Garn–St. Germain Depository Institutions Act).

[35] Hardly any households had second mortgages in the early 1970s (Bartscher et al. (2020)). HELOCs only spread in the mid-1980s (Maki 2001). Cash-out refinancing was not yet common either. In the 1977 SCF, only 3% of the respondents stated they had ever refinanced their first mortgage.

$$V_j(X_j) = \max_{c_j, n_j^f, a_{j+1}, h_j} u(c_j, n_j^f, h_j) + \beta \mathbb{E}_j V_{j+1}(X_{j+1}) \quad s.t. \tag{1.6}$$

$$c_j + a_{j+1} + p h_j (1 + F^b \mathbb{I}_{h_j \neq h_{j-1}}) + q \mathbb{I}_{h_j = 0} = (1+r) a_j + y_j^m + y_j^f + p h_{j-1} (1 - F^s \mathbb{I}_{h_j \neq h_{j-1}})$$

$$y_j^f = w_j^f n_j^f \mathcal{P} T$$

$$a_{j+1} \geq \begin{cases} 0 & if\ h_j = 0 \\ min\{a_j, 0\} & if\ h_j > 0\ \wedge\ h_j = h_{j-1} \\ -\phi(h_j, y_j^m, y_j^f) & if\ h_j > 0\ \wedge\ h_j \neq h_{j-1} \end{cases}$$

$$h_j \in \{0, 1, 2\}$$

$$n_j^f \in \mathcal{N}.$$

1.5.2 Calibration

I calibrate the model to the early 1970s, and investigate the effects of relaxing the DTI constraint based on the creditable share of the wife's income. The LTV and DTI limits are set to $\lambda^h = 1$ and $\lambda^y = 2$, which corresponds to the 90th percentile of the LTV and DTI distributions for new homeowners in the baseline sample between 1969 and 1971. House values and net rental costs are chosen based on the same data. The rental cost is set to the median rental cost of around 5,300 dollars in the data. The value of the "entry-level" house is set to the average house value among households between ages 25 and 35, which is approximately 110,000 dollars. The value of the larger house is set to match the average for households aged 40 to 50 in the data (around 135,000 dollars), which is the period when life-cycle house values peak in the data. To calibrate the mortgage interest rate, I use the 1971 wave of the SCF+. Specifically, I compute the median mortgage interest rate of married homeowners aged 25 to 64 who bought a home within the same or previous year and have an interest-bearing mortgage, which is 6.3%. The savings interest rate is set to the annual three-month treasury bill rate published by the Board of Governors of the Federal Reserve System (FRED series TB3MS), which was 4.3% in 1971. The CRRA parameter σ is set to the common value of 2, and the discount factor β is chosen as 0.94. The pension replacement rate is set to 0.7, as in Attanasio et al. (2012). An overview of

Table 1.4: Externally calibrated parameters

name	value	definition	target/source
λ^h	1	LTV limit	PSID
λ^y	2	DTI limit	PSID
λ^d	0.5	income discounting factor	ECOA
q	0.53	net rental cost	PSID
$p_1 h_j$	11	house value $h_j = 1$	PSID
$p_2 h_j$	13.5	house value $h_j = 2$	PSID
r^s	0.043	interest rate on savings	Federal Reserve
r^b	0.063	interest rate on debt	SCF+
σ	2	CRRA parameter	
β	0.94	discount factor	
b	0.7	pension replacement rate	Attanasio et al. (2012)

Notes: The table summarizes the externally calibrated parameters (see text for details).

Table 1.5: Internally calibrated parameters

name	value	definition
ψ	1.64	labor elasticity parameter
μ^h	2.51e-3	basic housing preference parameter
ϕ^h	5.01e-3	house size preference parameter
F_b	0.04	buying cost
F_s	0.07	selling cost
$\theta^f(j)$	see Appendix 1.D.2	female leisure preference profile
$\chi(j)$	see Appendix 1.D.2	housing utility profile

Notes: The table summarizes the internally calibrated parameters (see text for details).

the externally calibrated parameters is given in Table 1.4.

The remaining parameters are chosen to minimize the average squared distance between moments from the data and the model. I target average homeownership, mortgagor and FLFP rates in ten-year age bins. A summary of the model and data moments is given in Appendix Table 1.D.2. Table 1.5 summarizes the internally calibrated parameters. The estimated value for the parameter ψ is 1.64. Based on a 30-hour contract and a time endowment of $T = 112$ non-sleeping hours, this implies a labor supply elasticity of $(1-30/112)/(1.64 \cdot 30/112) \approx 1.67$, which is within the range of estimates for the aggregate elasticity (intensive and extensive margin) over the life cycle (between 1.37 and 1.93) found by Attanasio et al. (2018). With values of 4% and 7%, the transaction cost parameters F_b and F_s are similar to the median selling and buying costs of 2.5% and 7% reported by Martin and Gruber (2004) based on the U.S. Consumer Expenditure Survey. I use polynomials to model the age-dependent preference parameters $\theta^f(j)$ and $\chi(j)$. The resulting profiles are plotted in Figure 1.D.2. Their shape aligns well with the timing of life-cycle events such as the arrival of children and grandchildren, as discussed in more detail in Appendix 1.D.2.

I estimate the wage processes from equation (1.4) based on my baseline sample of married households. Following Borella, De Nardi, and Yang (2018), I include households between 20 and 70 years of age when estimating the deterministic part to avoid end point problems, but compute the variances based on working-age households between 25 and 64. The wage processes are estimated separately for husbands and wives.[36] Men and women working below 500 hours per year are considered as not working in that year, such that their wages are set to zero. I regress log real hourly wages on a quadratic polynomial in household age, dummies for the number of children, as well as state and 5-year cohort fixed effects. For the wives, I apply a Heckman model, in which selection is based on "other" household income (total household income net of the wife's labor income), the number of children and the total household size. The intercepts α_0^s are determined by averages for the cohort born between 1945 and 1950, who were of typical "home-buying age" in the early 1970s. As women can work part-time in the model, I allow for part-time penalties $\mathcal{P}^{50}$ and $\mathcal{P}^{75}$ by including dummies for working 50% or 75%.

Following the previous literature, the parameter ρ^s is set to one for both spouses. The variances $\sigma^2_{\varepsilon^s},\ s = m, f$ are estimated according to the method of Heathcote, Perri, and Violante (2010).[37]

[36] In the model simulations, $\varepsilon^s_j,\ s = m, f$ are allowed to be correlated, following a standard approach in the literature. I employ a correlation coefficient of 0.25 based on the estimate of Hyslop (2001).

[37] This estimation is based on identifying autocovariance moments in levels, as the alternative identification based on moments in differences tends to overestimate the variance of the permanent shock, leading to an unrealistic growth of wage inequality over the life cycle (see also Daly, Hryshko, and Manovskii 2018).

Table 1.6: Parameters of wage profiles

names	values
α_0^m, α_1^m, α_2^m, α_3^m, α_4^m	-1.5381, 0.4036,-0.0132, 0.0002,-1.1439e-6
α_0^f, α_1^f, α_2^f, α_3^f, α_3^f	-0.7490, 0.2889,-0.0085, 0.0001,-4.8191e-7
$\sigma^2_{\varepsilon^m}$, $\sigma^2_{\varepsilon^f}$	0.0062, 0.0056
$\mathcal{P}^{50}$, $\mathcal{P}^{75}$	0.8301, 0.9368

Notes: The table shows the estimated coefficients of the log wage processes from equation (1.4) and the estimated part-time penalties $\mathcal{P}^{50}$ and $\mathcal{P}^{75}$. See text for details.

Table 1.6 shows the estimated coefficients, and Appendix Figure 1.D.1 plots the resulting income profiles under certainty. Additional details on the model solution are given in Appendix 1.D.2.

1.5.3 Experiment: relaxation of DTI constraint via female income

Figure 1.8 shows the homeownership rate, the share of households with mortgage debt, and the female labor force participation rate over the life cycle, computed from 10,000 simulations of the model. The model produces similar average homeownership, mortgagor and FLFP rates within ten-year age bins, which is a success of the calibration strategy (cf. Appendix Table 1.D.2). Besides, the model is also able to reproduce the overall shape of the empirical life-cycle profiles from the PSID, which are shown as dashed lines in Figure 1.8. It should also be noted that despite the simplification of a choice between discrete contracts, the model produces an intensive margin of labor close to the data, with an average of slightly above 28 compared to 31 hours.

Figure 1.8: Comparison to data

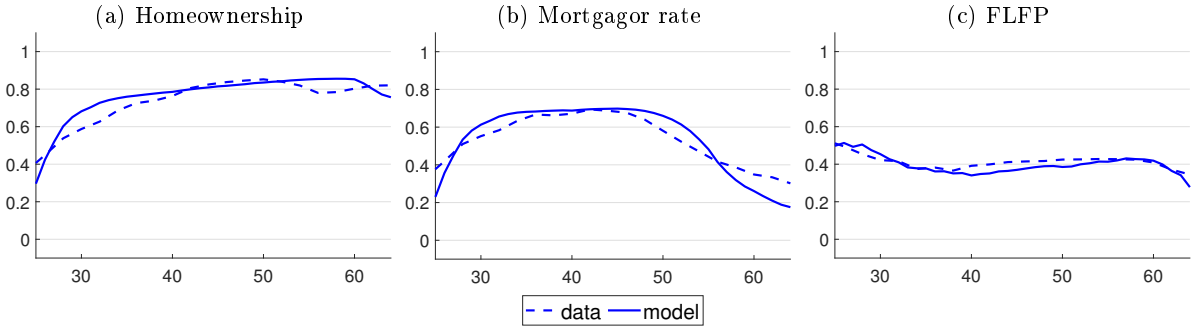

Notes: The graph compares the homeownership rate, the share of mortgagors and the FLFP rate over the life cycle from the model to the corresponding life-cycle profiles from the PSID for the period 1969-1971. The FLFP rate in the data refers to women working at least 500 hours per year. The data profiles were smoothed by taking a three-year moving average.

In the following, I will use the model as a laboratory to simulate the effects of relaxing the borrowing constraint via the female income discounting factor λ^d. Table 1.7 summarizes the results of my experiment. It shows how the shares of homeowners, mortgage holders and working wives change when the female income discounting factor λ^d is raised from 50% to 100%, such that households can count all of the wife's income toward a mortgage. In the cross section, the homeownership rate of married households between ages 25 and 64 increases by 9.5 percentage points. The surge in the share of households holding mortgage debt is similar, with 9 percentage points. These increases are not merely due to the fact that households can borrow against a higher share of what the wife would have earned anyway. The third row of Table 1.7 shows

Figure 1.9: Life-cycle effects of raising λ^d from 0.5 to 1

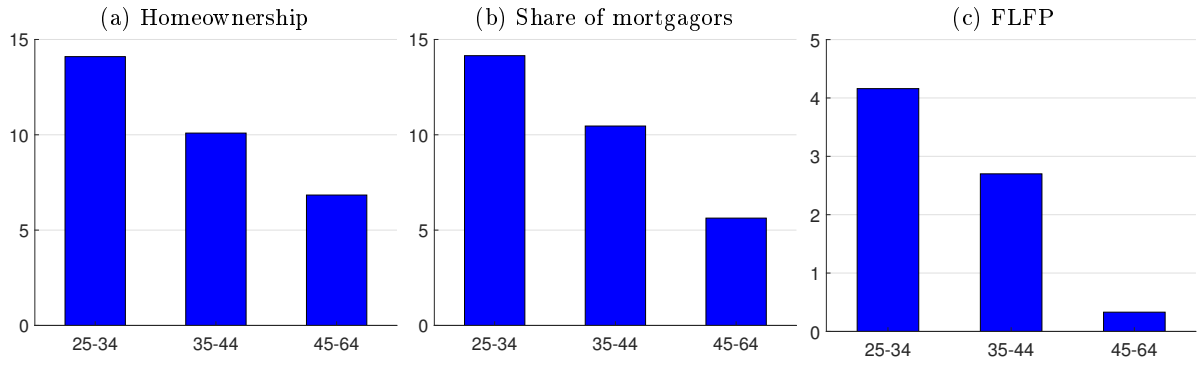

Notes: The graph shows the change in the homeownership rate, the share of mortgagors and the FLFP rate (in percentage points) over the life cycle when increasing the female income discounting factor λ_d from 0.5 to 1.

that the reform produces incentives for women to join the labor force, such that the FLFP rate increases by around 2 percentage points. To generate a similar increase in FLFP under the pre-reform scenario, average female wages would have to rise by around 4%. At the intensive margin, female labor supply changes little, with women working about half an hour less per week on average. This is because on the one hand, the women who would also have worked in the pre-reform scenario with $\lambda^d = 0.5$ choose to work slightly less (around 20 minutes per week). On the other hand, the women who join the labor force under the new scenario, but would not have worked in the old one, work somewhat less than the average (23.5 compared to 28 hours per week).

In order to get an impression of the relative importance of the change in female labor supply for the change in homeownership, the last column of Table 1.7 presents a counterfactual exercise. I compute conditional policy functions for each possible female labor supply choice under the post-reform scenario with $\lambda^d = 1$. In the simulations, I then use the realized trajectory of n_j^f from the simulation under the pre-reform scenario with $\lambda^d = 0.5$, and compute the asset and housing choices based on the conditional policy functions corresponding to these "old" labor supply choices in each period. In other words, I simulate a scenario in which women can now use their full income for the mortgage, but they cannot adapt their labor supply relative to the pre-reform scenario. Without additional women joining the labor force, the increase in

Table 1.7: Comparison: $\lambda^d = 0.5$ versus $\lambda^d = 1$

variable	difference	difference with fixed labor supply
homeownership	0.095	0.051
mortgagor rate	0.090	0.048
FLFP	0.019	0
hours (intensive margin)	-0.561	0

Notes: The first column shows the average changes in homeownership, the share of mortgagors and female labor force participation (in percentage points) as well as the change in average intensive-margin female hours worked (in hours) among households aged 25 to 64 when increasing the income discounting factor λ^d. from 0.5 to 1. The second column shows the respective differences when holding female labor supply choices constant at the optimal values from the pre-reform scenario with $\lambda^d = 0.5$.

homeownership and the share of mortgage holders would only be around 5 percentage points. This is a similar order of magnitude as the empirical estimate of 3.3 percentage points found in the data for women who had already been working (see Section 1.4.4). The model therefore suggests that the medium- to long-run effects of the ECOA on homeownership even exceeded the impact effects by encouraging more married women to join the labor force.

Figure 1.9 shows the effects of the reform over the life cycle. Consistent with the theoretical arguments discussed above, the effects are strongest for young households between 25 and 34 years of age, and decline as households become older. For households in the young group, who are of typical "home-buying age", the wives' LFP increases by around 4 percentage points. The ability to count twice as much of the wife's income toward the mortgage as before, combined with the active increase in FLFP, enables households to prepone homeownership, such that the homeownership rate in the youngest group increases most drastically.

Finally, Figure 1.10 further decomposes the results from Figure 1.9 by male income. In line with the empirical results and theoretical considerations, the reform bears the greatest advantages for households with lower-income husbands. Interestingly, the effect is much more persistent over the life cycle for households with relatively low male incomes. While young households benefit from the reform even if the husband has above-average earnings, the effects at older ages are almost exclusively driven by households with lower male earnings.

Figure 1.10: Heterogeneous effects: high versus low male income

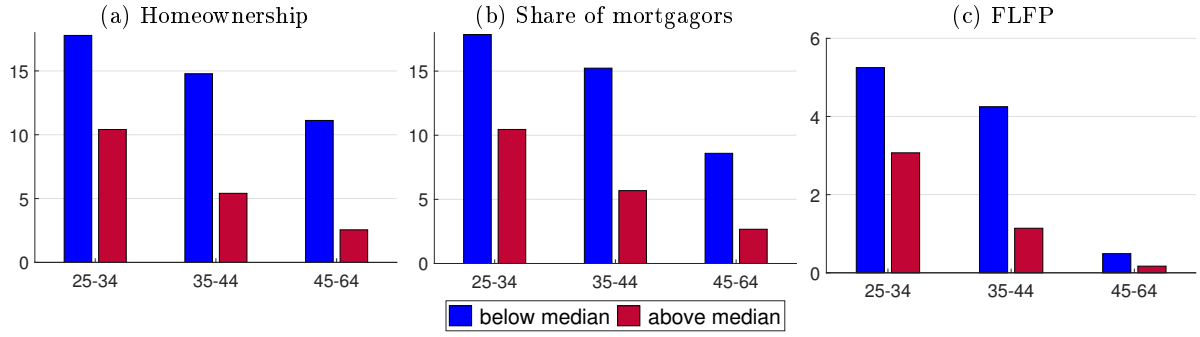

Notes: The graph shows the effects of the experiment in Figure 1.9, stratified by the husbands' incomes. Blue bars show the respective changes (in percentage points) for households with below-median male earnings, and red bars refer to households with above-median male earnings.

1.5.4 Discussion

The model shows that women who were planning their future in the early 1970s, based on the prevailing economic information, were more likely to embark on a working career. It further illustrates that additional female labor supply could substantially amplify the positive effects of equal credit opportunity laws on married couples' homeownership rates in the longer run, especially for younger couples and couples with lower male earnings.[38]

In an overlapping generations setup with intergenerational learning, the ECOA could generate further dynamic effects on labor supply by accelerating the pace at which women learn about the true cost of working (Fernández 2013). Furthermore, there might be spillovers to male employment and earnings in a setup with involuntary unemployment and search. In an important

[38] Of course, LFP and homeownership decisions in the data can also be influenced by factors such as changes in house prices, interests rates, household composition or the availability of alternative sources of financing. Incorporating these factors into the model would go beyond the scope of this chapter, but could be an interesting avenue for further research.

series of papers, Kyle Herkenhoff and coauthors have shown both empirically and theoretically that unemployed households take longer to find a job if they have more access to credit, but achieve higher earnings replacement rates, such that welfare improves (Herkenhoff 2019, Braxton, Herkenhoff, and Phillips 2020, Herkenhoff, Phillips, and Cohen-Cole 2020). As mentioned, home equity extraction became more popular from the 1980s on, and became even easier in the 1990s (Hurst and Stafford 2004). Hurst and Stafford (2004) have established that households experiencing an unemployment shock while having low liquid assets have a substantially higher propensity to refinance their mortgage and extract equity. Against this background, the ability to borrow against the home based on the wife's income if the husband becomes unemployed could allow the husband to search longer for a new job and obtain a better match.

Additionally, home equity lines of credit (HELOCs) were introduced on a large scale in the mid-1980s (Maki 2001). Although they typically require the household to meet a debt-to-income constraint upon application, it will not require another income assessment every time the line is drawn on (see also Braxton, Herkenhoff, and Phillips 2020). Therefore, both husbands and wives who were able to buy a house with the help of the wife's income and establish a HELOC gained the ability to flexibly borrow against their home in the case of future unemployment, allowing them to better smooth their consumption and adapt their job search behavior. I view the analysis of the potential effects of the ECOA on job search and within-household income and consumption smoothing as an interesting avenue for further research.

1.6 Conclusion

This chapter has shed light on the effects of the Equal Credit Opportunity Act on homeownership and mortgage borrowing of married households. The ECOA proscribed discrimination in mortgage lending based on sex and marital status. In particular, it prohibited the formerly common practice to partially or even fully discount the wife's income in joint mortgage applications of married couples. It therefore provides a natural experiment to study the relaxation of income-related borrowing constraints. Although the ECOA required profound changes in creditors' lending practices, its effects on female access to credit have long remained an open question.

Using data from the PSID, I find that married households with a working wife increased their mortgage borrowing, and with it homeownership and house size, subsequent to the passage of the ECOA. The results are supported by event study regressions exploiting state-level variation. Furthermore, the national-level results are robust to controlling for other contemporaneous events. The estimates imply that the new legislation enabled 1.4 million married households to move to an own home upon its introduction, and 1.8 million to take out a mortgage.

While the law initially mostly benefited households with a wife who had already been working, it also changed the labor supply incentives for women in future cohorts of new homeowners. I draw on a life-cycle model to explore potential changes in the incentives for married women's labor supply. The model shows that the incentives of the ECOA were powerful enough to increase the labor force participation rate of married women, thus amplifying the positive effects on married couples' homeownership and borrowing rates in the longer perspective.

Appendix 1.A Supplementary analyses

1.A.1 Female labor supply

Figure 1.A.1 shows the aggregate female labor force participation rate from the U.S. Bureau of Labor Statistics over time. By the early 1970s, it had already increased considerably compared to the 1950s, and its growth even slightly accelerated in the subsequent years.

Figure 1.A.1: Female labor force participation

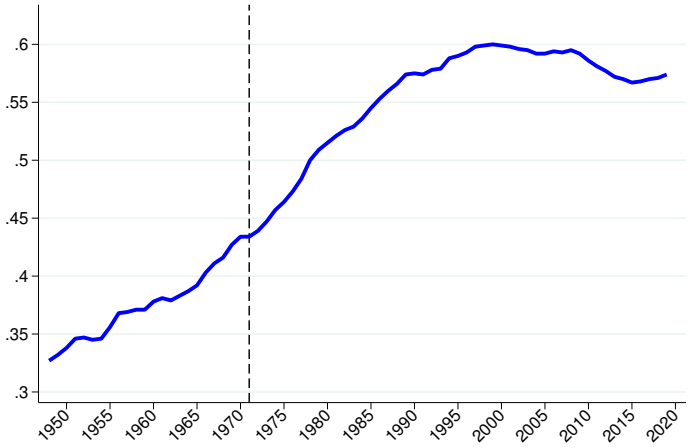

Notes: The graph shows the civilian labor force participation rate of women. Data source: U.S. Bureau of Labor Statistics.

In order to examine the labor supply decisions of married women around the time of a home purchase, I restrict my baseline PSID sample to households in which both spouses are below 60 years of age. As the effects on homeownership and housing debt in Section 1.4 mainly work via the extensive margin, I focus on renter-to-owner transitions. In a first step, I look at couples in which both spouses are working at the time of the home purchase. For couples that I observe over the 5 years subsequent to the home purchase (about 970 households between 1969 and 1985), I compute the number of those years during which the wife is still working.[39] Mortgage lenders who discounted the female's income were most concerned about debt-service capacity during the first years of a mortgage, which they considered as the most risky period (Thurston 2018; see also Foote, Loewenstein, and Willen 2018). In my data, only 10% of the wives worked for less than 50% of the period covering the home purchase and the 5 following years. Another 7% worked for half of that period, and 61% stayed employed over the full period.

These numbers do not suggest that discounting female income across-the-board was justified. Nevertheless, female employment continuity was still below that of men. In the same households, 96.6% of the husbands worked every year during the period covering the home purchase and the subsequent 5 years. Yet while women are more likely to stop working, often to take care of children, it would not be a rational choice for a woman to voluntarily leave the labor force for childcare reasons if this means jeopardizing her mortgage and risking foreclosure (see also Lally 1974). In line with this reasoning, Fortin (1995) shows that Canadian women are more likely to work if their household would be close to the debt-service constraint without their income.

Following the approach of Fortin (1995), I regress an indicator for female employment on indicators for whether the household's DSTI ratio without the wife's labor income is between 10% and

[39] I use the years after 1985 to calculate forward-looking variables, and then truncate the sample to end in 1985, as in the baseline.

Table 1.A.1: Female labor force participation and debt service obligations

	(1)		(2)	
DSTI^{wo} 10% - 17.5%	0.006	(0.011)	0.044***	(0.013)
DSTI^{wo} 17.5% - 25%	0.030***	(0.011)	0.091***	(0.015)
DSTI^{wo} > 25%	0.096***	(0.016)	0.177***	(0.023)
50% - 90%			0.133***	(0.013)
Top 10%			0.201***	(0.027)
1 Child			-0.059***	(0.011)
2 Children			-0.139***	(0.022)
3+ Children			-0.185***	(0.030)
House Value			-0.001	(0.001)
Mortgage Balance			-0.008***	(0.002)
FE	yes		yes	
Mean	0.656		0.656	
Observations	16,265		16,265	

Notes: The graph shows the results of regressions in which the outcome variable is a dummy for whether the wife works. DSTI^{wo} is the household's DSTI ratio without the wife's labor income. The regressions include age, state, household and region-year fixed effects. The sample was restricted to married homeowners below 60 years of age. Standard errors are clustered at the household and state level (* p<0.1, ** p<0.05, *** p<0.01).

17.5%, between 17.5% and 25%, or above 25%, restricting the sample to homeowners. The regression includes age, state, household and region-year fixed effects. In model (2), I additionally control for the income group, the number of children, the value of the house, and the remaining mortgage balance. Table 1.A.1 confirms that the patterns found by Fortin (1995) also apply to the PSID data.

To obtain a more direct impression of female labor supply and income around the time of a home purchase, I estimate event studies similar to equation (1.3), where the event in period 0 is defined as the purchase of a home after having been a renter. All right-hand-side variables

Figure 1.A.2: Female hours worked and income around home purchase

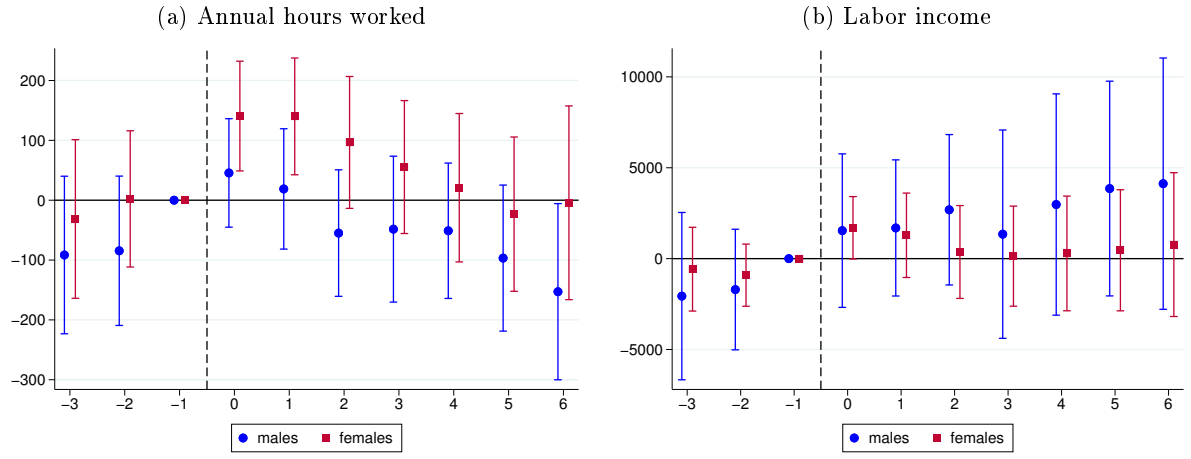

Notes: The graph shows the results of estimating event studies similar to equation (1.3), where the event is a home purchase after renting. All right-hand-side variables are interacted with a spouse dummy to allow a comparison between husbands and wives. Observations outside the estimation sample (covering 1969-1985) are used in computing the leads and lags. The sample was restricted to married couples below 60 years of age who both reported positive labor income in the year of the home purchase. Standard errors are clustered at the household and state level. The whiskers indicate 95% intervals.

are interacted with a spouse dummy to allow a comparison between husbands and wives. The results in Figure 1.A.2 show that women who reported a positive labor income in the period of the home purchase even slightly *increase* their average annual hours worked after buying a house. Their average labor income also increases slightly, and evolves very similarly to that of their husbands. It should be noted that I could not detect a differential probability for the birth of a first or additional child before and after a home purchase. Indeed, the same data give rise to substantial and persistent drops in female income subsequent to childbirth, in line with the evidence on "child penalties" by Kleven et al. (2019). The results are available upon request.

1.A.2 Debt service

Figure 1.A.3: Mortgage interest rates of recent owners

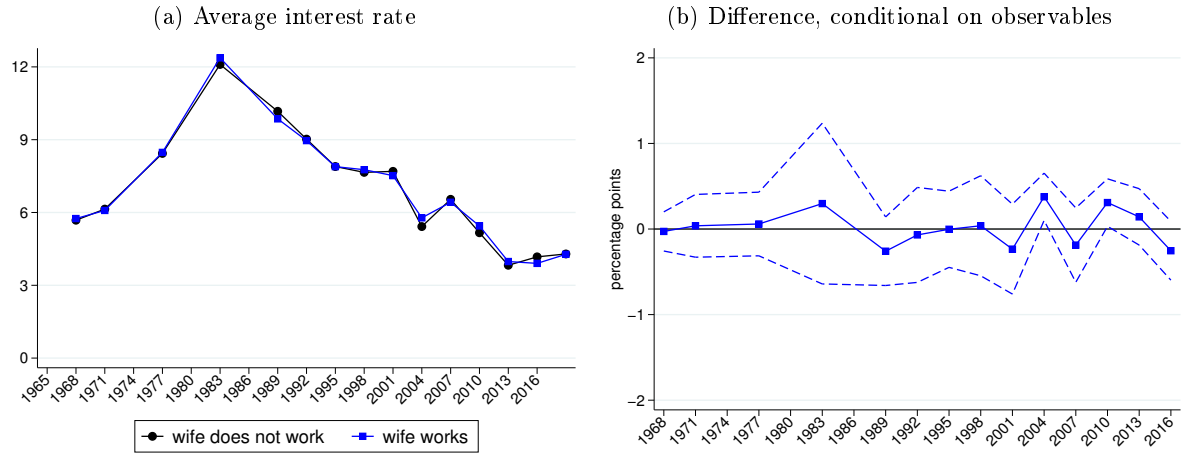

Notes: The left panel shows the average mortgage interest rate on the first mortgage of married recent homeowners (meaning they moved in the previous or current year) with and without a working wife over time. The right panel shows the coefficients $\beta_w + \beta_t$ from equation (1.A.1). The dashed lines show 95% CIs, based on robust standard errors taking multiple imputation into account.

Figure 1.A.3a shows the average interest rate on the principal residential mortgage of married recent homeowners from the SCF+. Recent means that the household has moved into its current residence during the current or the previous two years. The graph shows that there is no systematic difference in interest rates between married households with versus without a working wife. However, households with a working wife might have different characteristics. Therefore, I regressed mortgage interest rates on observable socioeconomic variables to control for potential confounders:

$$r_{it} = \beta_0 \tilde{y}_{it} + \beta_1 h_{it} + \beta_2 d_{it} + \beta_w works^w_{it} + \sum_{t=1971}^{2016} \beta_t \cdot \delta_t \cdot works^w_{it} + \delta_t + age_{it} + kids_{it} + black_{it} + college^h_{it} + college^w_{it} + \epsilon_{it}, \quad (1.A.1)$$

where $\tilde{y}_{it}$ denotes log total household income, h_{it} is the asset value of the house, d_{it} is the outstanding mortgage balance, $works^w_{it}$ is a dummy for whether the wife works, δ_t are time dummies, age_{it} is a set of age dummies, $kids_{it}$ are dummies for the number of children, $black_{it}$ is a dummy for being black, and $college^h_{it}$ and $college^w_{it}$ are dummies for whether head and wife have a college degree.[40] Figure 1.A.3b shows the coefficients $\beta_w + \beta_t$ over time. The results

[40] If not otherwise stated, demographics refer to the household head. Note that the SCF+ does not have a panel dimension, such that I could not include household fixed effects.

confirm that there is no systematic interest rate difference between the two groups.

1.A.3 Additional graphs

Figure 1.A.4 shows how mentions of the term "women and credit" in English books have evolved over time, using data from the Google Books Ngram Viewer. This online search engine displays the frequency of search strings (*n-grams*) in millions of digitized books (Michel et al. 2011). The mention frequency is zero until 1971, becomes positive in 1972, and sharply increases thereafter.

Figure 1.A.4: Google Books Ngram Viewer

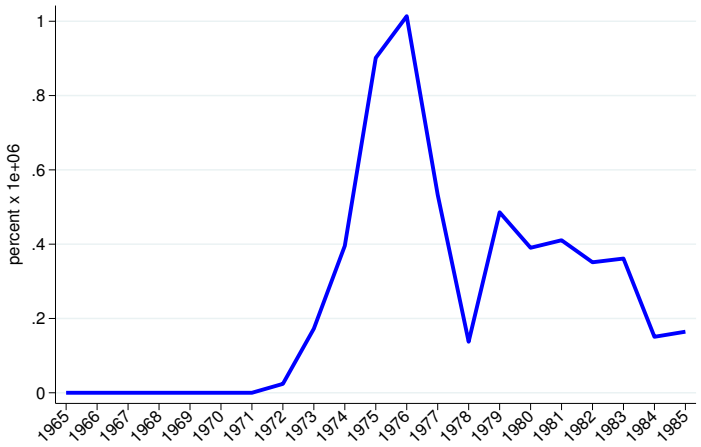

Notes: The graph shows how mentions of the 3-gram "women and credit" (case insensitive) have evolved over time. The figure is based on data from the Google Books Ngram Viewer. The y-axis shows the share of this 3-gram among all 3-grams contained in the Google sample of English books. The Google data are normalized with the total number of books published in each year.

Figure 1.A.5: Single and married mortgage holders

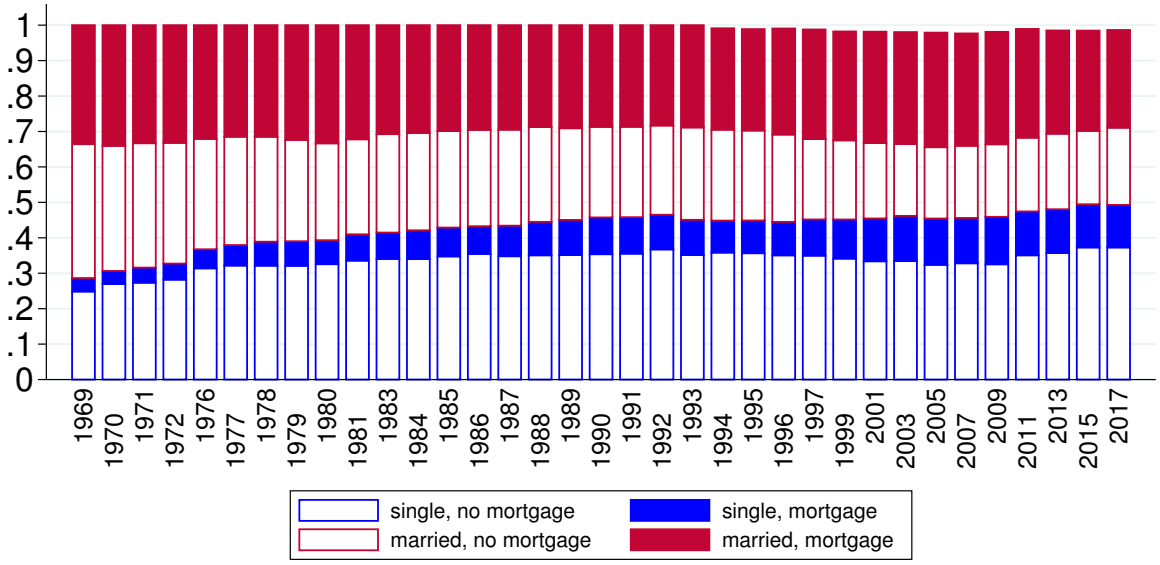

Notes: The graph shows the shares of single and married households with and without a mortgage over time.

Figure 1.A.5 shows the shares of single and married households with and without a mortgage in the PSID over time.

Figure 1.A.6 shows similar results as Figure 1.4 for DTI instead of DSTI ratios.

Figure 1.A.6: Household with debt-to-income ratio above 2

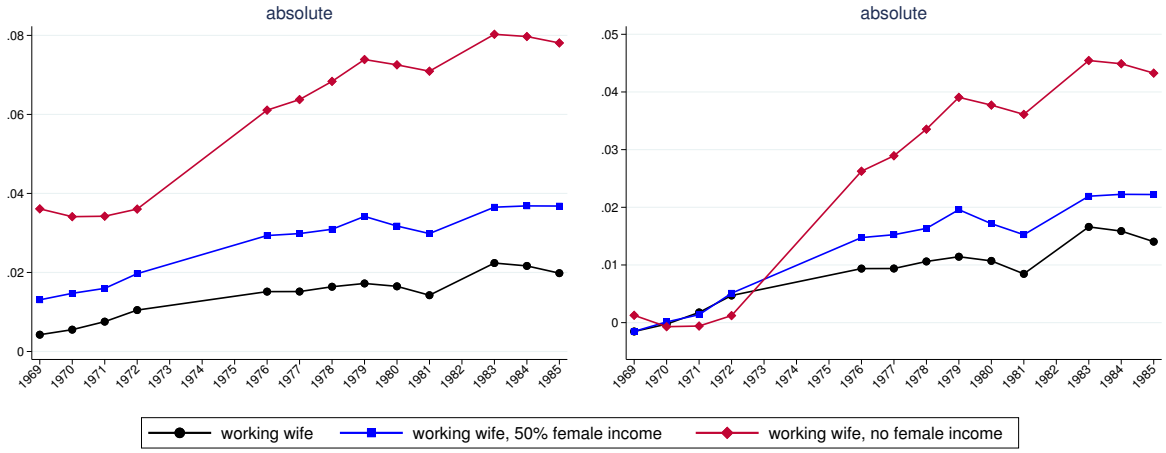

Notes: The left panel shows the share of households with a debt-to-income ratio above 25% among households with a working wife. Households with implausibly high ratios (above the 99.8th percentile) were excluded. The blue lines with squares (red lines with diamonds) show counterfactual shares based on debt-to-income ratios excluding 50% (100%) of the wife's income. The series were smoothed by taking a 3-year moving average. The right panel normalizes each series with its average over the period 1969-1971.

Appendix 1.B Robustness

1.B.1 Point estimates with and without controls

Table 1.B.1: Effect of controls: mortgagor rate

	(1)	(2)
Post 1971 × Tot. Inc. Share Wife 71	0.004*** (0.001)	0.003*** (0.001)
25-34		0.303*** (0.040)
35-44		0.406*** (0.047)
45-54		0.420*** (0.050)
55-64		0.337*** (0.051)
65+		0.221*** (0.068)
50% - 90%		0.054*** (0.013)
Top 10%		0.089*** (0.020)
1 Child		0.055*** (0.013)
2 Children		0.143*** (0.019)
3+ Children		0.150*** (0.025)
Household FE	yes	yes
Year FE	yes	no
State-year FE	no	yes
Mean	0.513	0.513
Observations	20,192	20,190

Notes: The table presents results for equation (1.B.1). The interaction variable is the average share of the wife's labor income in total household income over the pre-reform years. Standard errors are given in parentheses and are clustered at the household and state level (* p<0.1, ** p<0.05, *** p<0.01). The wife's income share is defined in percent.

One can simplify the model in equation (1.2) by replacing the individual year dummies with an indicator D_t equal to 1 for years after 1971:

$$Y_{it} = \beta_0 + \beta_1 \cdot D_t + \beta_2 \cdot D_t \cdot share_i^{pre} + \Gamma' X_{it} + \gamma_i + \delta_t + \epsilon_{it}. \tag{1.B.1}$$

The corresponding point estimates for the housing debt-to-income ratio as the outcome variable are summarized in Table 1.B.1.

Figure 1.B.1: Effect of controls

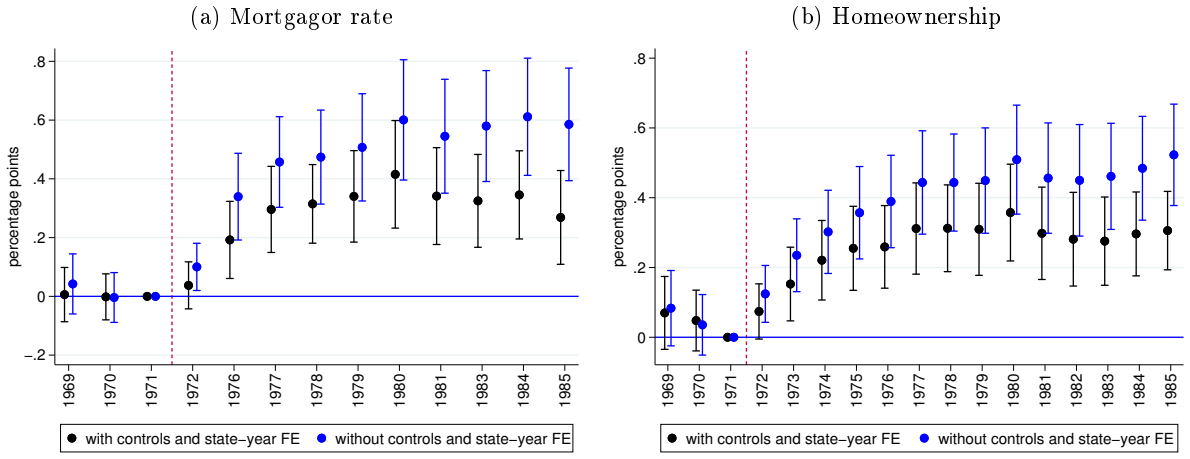

Notes: The graph compares the results from Figures 1.5b and 1.6b to versions without the controls X_{ist} and the state-year fixed effects. The whiskers indicate 95% intervals.

Figure 1.B.2: Dummy interaction with and without controls

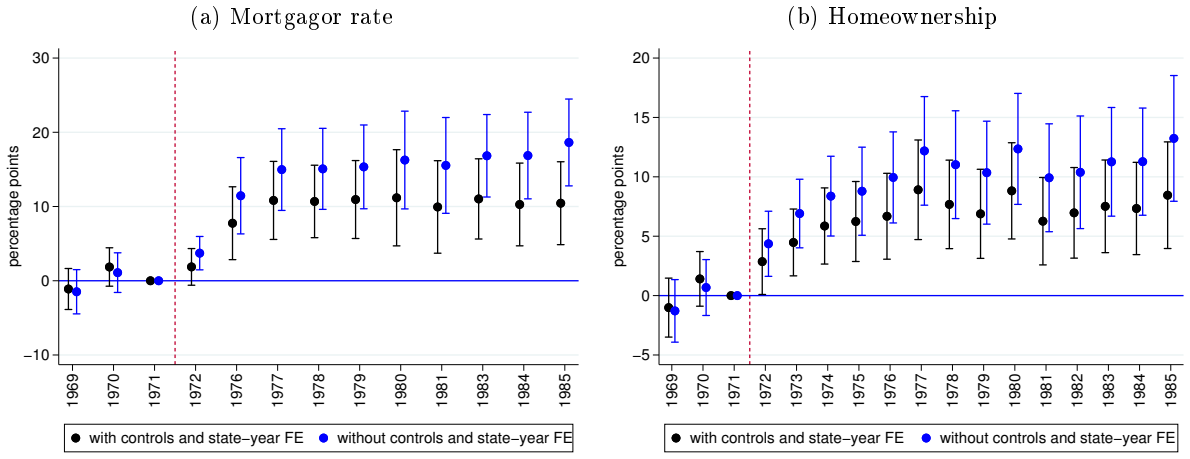

Notes: The graph shows the coefficients on the interaction term in equation (1.2), where $share_i^{pre}$ was replaced with a dummy for whether the average share of the wife's labor income in total household income over the pre-reform years (up to 1971) was positive. It also shows the results for the same regressions without the controls X_{ist} and the state-year fixed effects. The whiskers indicate 95% intervals.

1.B.2 Triple difference

As explained in Section 1.2, income discounting was especially common when the wife was of childbearing age. Therefore, one may use older women as an additional control group in a triple-difference regression:

$$Y_{ist} = \sum_{t=1969,\, t\neq 1971}^{1985} \left[\beta_t^1 \cdot \delta_t \cdot share_i^{pre} \cdot D_i^{<35\,pre} + \beta_t^2 \cdot \delta_t \cdot share_i^{pre} + \beta_t^3 \cdot \delta_t \cdot D_i^{<35\,pre}\right] + \Gamma' X_{ist} + \gamma_i + \alpha_{st} + \epsilon_{ist}, \tag{1.B.2}$$

where $D_i^{<35\,pre}$ is an indicator for whether the wife was below 35 years of age in 1971. Figure 1.B.3 plots the coefficients on the triple interaction term.

Figure 1.B.3: Triple difference with cohort

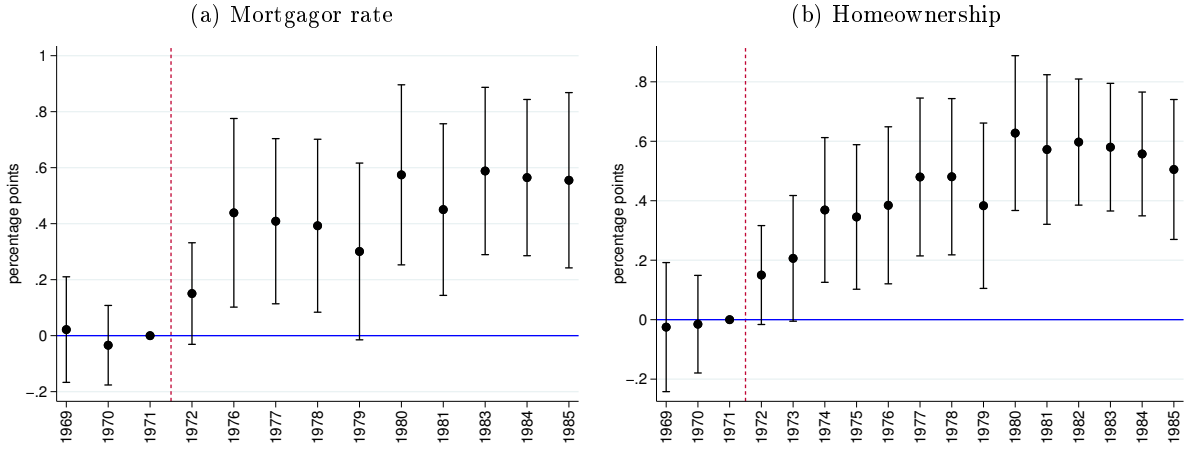

Notes: The graph presents the coefficients β_t^1 on the triple interaction in equation (1.B.2). The age dummies were omitted from the controls, as this regression already includes time and cohort dummies. The whiskers indicate 95% intervals.

1.B.3 Divorce rates

Figure 1.B.4: Divorce rates

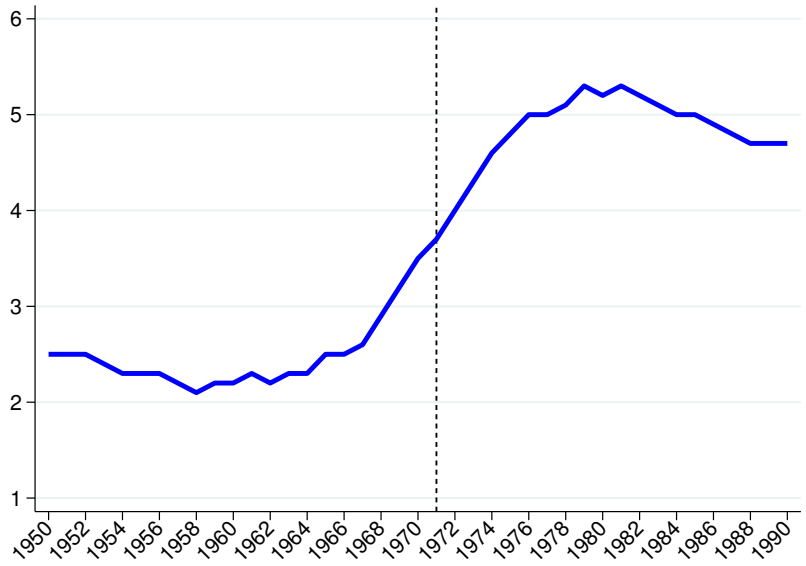

Notes: The graph shows the number of divorces per 1,000 inhabitants based on data from the National Center for Health Statistics (NCHS). The data were made available by Randal Olson at http://www.randalolson.com/wp-content/uploads/us-marriages-divorces-1867-2014.csv.

Figure 1.B.5: Excluding households who got divorced

(a) Mortgagor rate

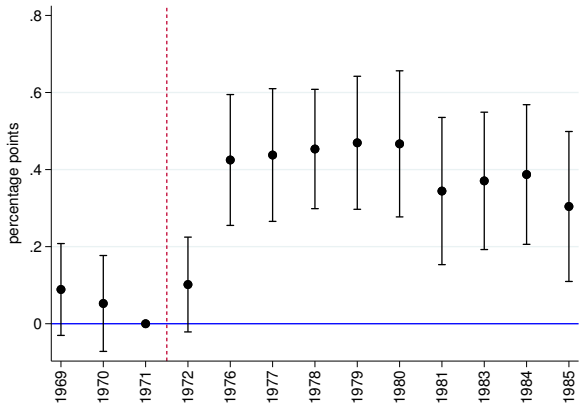

(b) Homeownership

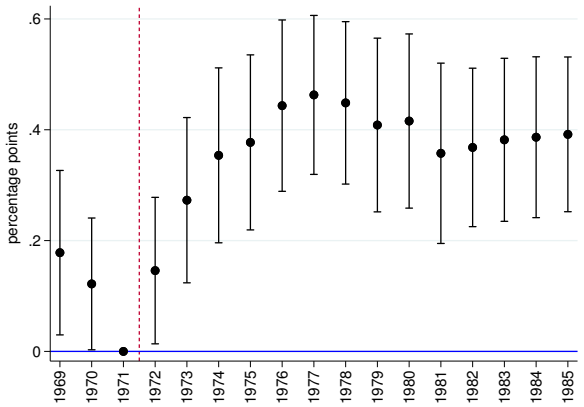

Notes: The graph shows a robustness check for Figures 1.5b and 1.6b, estimated on a sample restricted to households who had entered the PSID at the latest in 1971, were still in the survey in 1985, and never reported a transition from being married to being single in between. The whiskers indicate 95% intervals.

Figure 1.B.6: Event studies housing

(a) Housing-to-income ratio

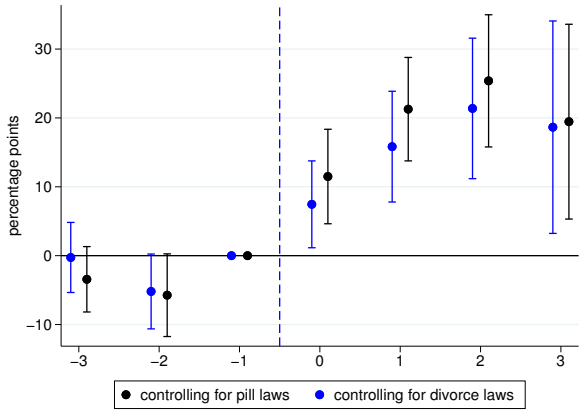

(b) Homeownership rate

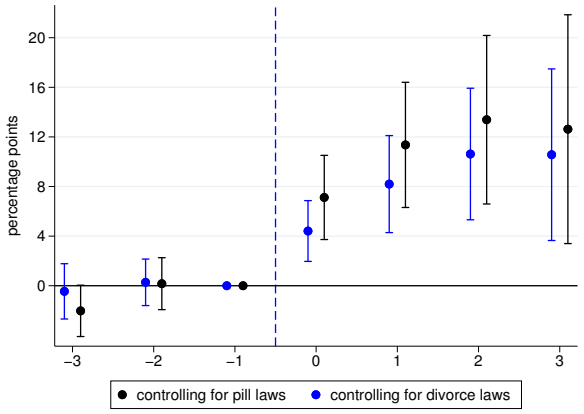

(c) House size

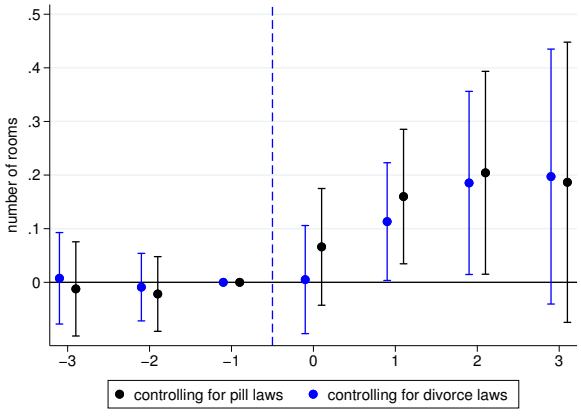

Notes: The graph shows a robustness check for Figure 1.7 in which I additionally control for event dummies indicating state consent laws determining early access to the pill from Bailey (2006) or state laws on unilateral divorce from Gruber (2004).

1.B.4 End of Vietnam War

Figure 1.B.7: Number of U.S. soldiers in Vietnam

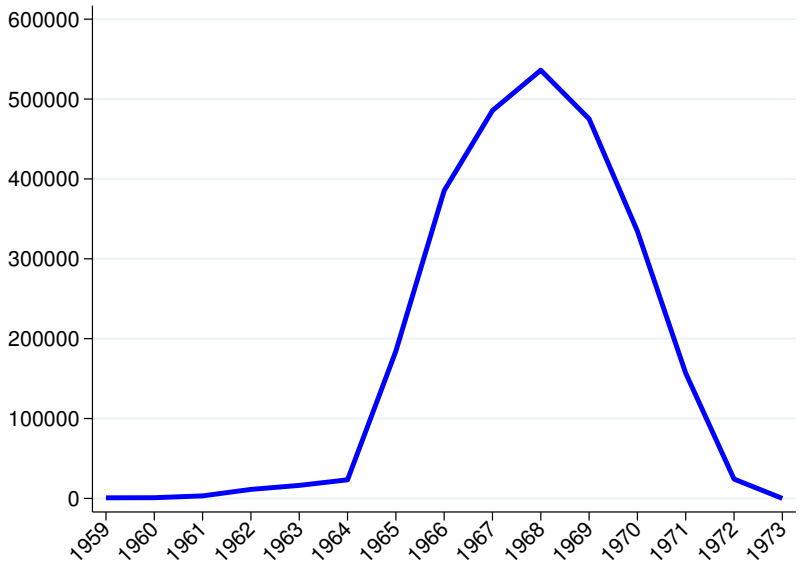

Notes: The graph shows the number of U.S. soldiers in Vietnam (source: https://www.americanwarlibrary.com/vietnam/vwatl.htm).

Figure 1.B.8: Excluding heads in armed forces

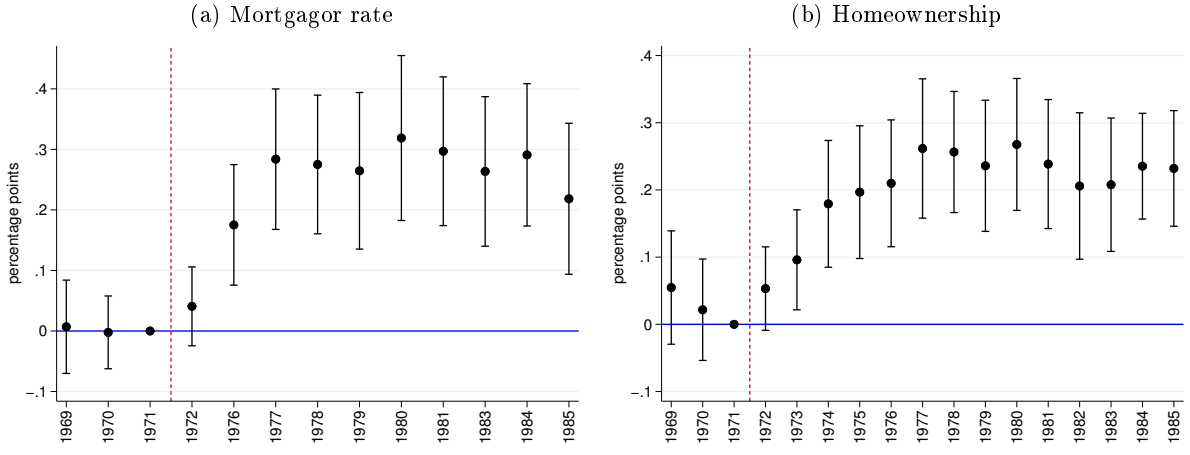

Notes: The graph shows a robustness check for Figures 1.5b and 1.6b, estimated on a sample which excludes all households whose head had been in the armed forces any time between the beginning of the sample in 1969 and the official end of the Vietnam War in 1975. The whiskers indicate 95% intervals.

1.B.5 Earned Income Tax Credit

Figure 1.B.9: Excluding households eligible for EITC

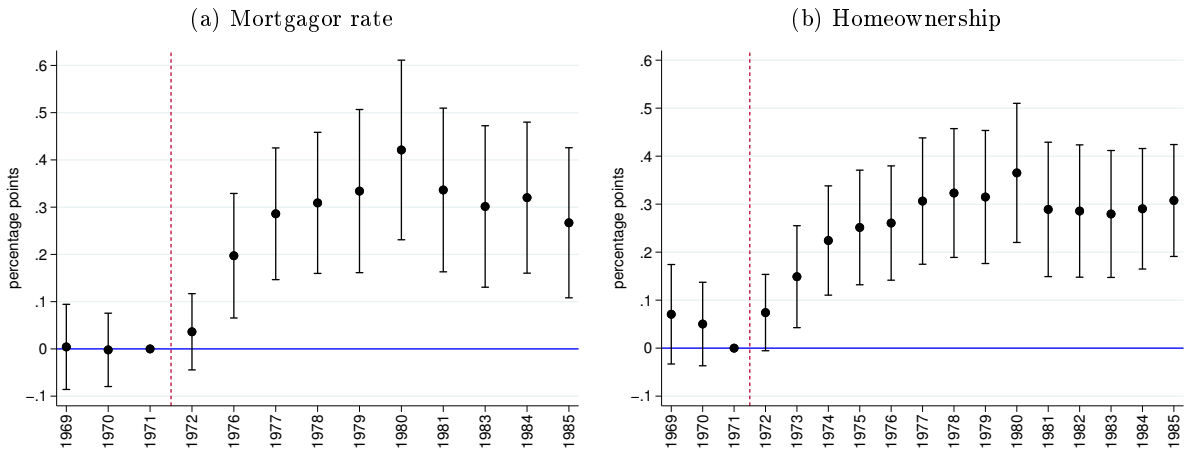

Notes: The graph shows a robustness check for Figures 1.5b and 1.6b, estimated on a sample which excludes all households who were eligible for the EITC after its introduction in 1975. The whiskers indicate 95% intervals.

1.B.6 State laws

Figure 1.B.10: Geographic location of states introducing laws

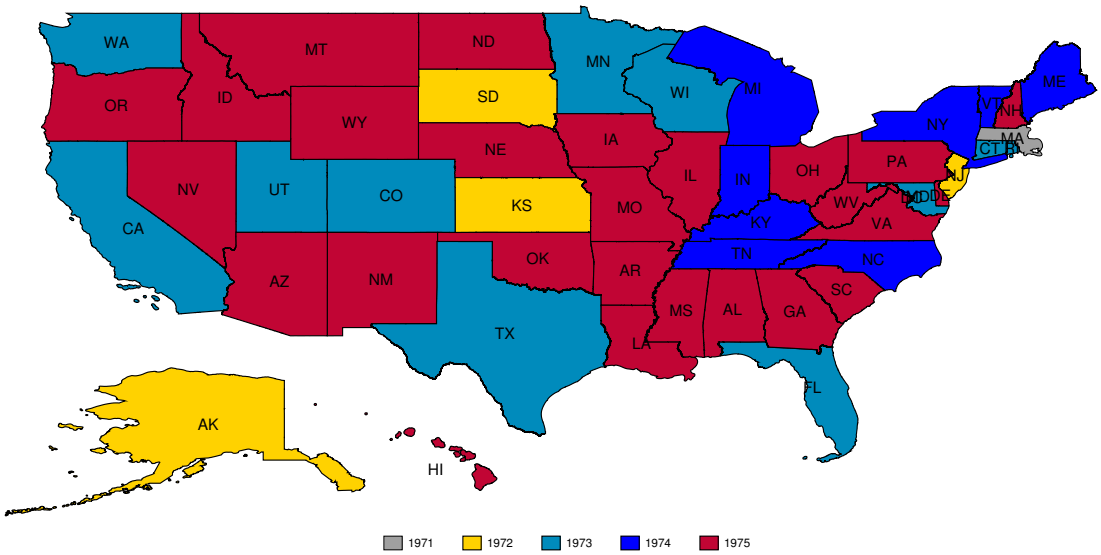

Notes: The map gives a geographical overview of the states introducing legislation against discrimination in access to credit for women in a given year.

Table 1.B.2: Overview: state laws

state	passed	effective	notes	source
Alaska	1972		does not include marital status	Gates (1974)
Arkansas*	1975		consumer credit only	Arkansas Code §4-87-104
California	1973/10		prohibits discrimination against *women*, not sex in general	Beckey (1974)
Colorado	1973/6			Beckey (1974)
Connecticut	1973/6	1973/10		Beckey (1974)
D.C.	1973/11			Beckey (1974)
Florida	1973/6	1973/10		Beckey (1974)
Georgia	1975/4			Acts and Resolutions of the General Assembly of the State of Georgia 1975, vol. 1, law no. 510
Hawaii	1975/5		only includes marital status	Session Laws of Hawaii. Act 109, H.B. no. 499
Illinois*	1973/9	1973/10	only applies to credit cards	Beckey (1974)
Indiana	1974			Bowdish (2010)
Iowa	1974	1975		Code of Iowa 1975, vol. 1, §601A.9
Kansas	1972		does not include marital status	Gates (1974)
Kentucky	1974/6			Bowdish (2010)
Louisiana	1975			Louisiana Revised Statutes, title 9, ch. 3
Maine	1974/2	1974/6		Bowdish (2010)
Maryland	1973/7			Gates (1974)

Table 1.B.2: Overview: state laws (ctd.)

state	passed	effective	notes	source
Massachusetts	1971		does not include marital status; more general law against credit discrimination since 1973/4, but applicability to mortgages unclear	Beckey (1974)
Michigan	1974/8			U.S. Office of Consumer Affairs (1975), Michigan Penal Code §750.147a
Minnesota	1973/5	1973/8		Beckey (1974)
Missouri*	1974/6		retail credit only	Laws of Missouri: Laws Passed by the Seventy-seventh General Assembly, Missouri Digital Heritage
Montana	1975/3			Laws and Resolutions of the State of Montana, vol. 1 1975, ch. 121
Nevada	1975			Nevada Revised Statutes, title 52, ch. 598b
New Jersey	1972			Beckey (1974)
New Mexico	1975			U.S. Advisory Commission on Intergovernmental Relations (1976)
New York	1974/6	1974/7		Beckey (1974)
North Carolina	1974/4			U.S. Office of Consumer Affairs (1975)
Ohio	1975/6			"Legislative Column" (1975)
Oklahoma*	1974		consumer credit only	Oklahoma Statutes §14A-1-109
Oregon*	1973/10		public accommodation only	Gates (1974)

Table 1.B.2: Overview: state laws (ctd.)

state	passed	effective	notes	source
Pennsylvania*	1969/6	1969/7	commercial property only	H.B. 567, Regular Session 1969-1970[†]
Rhode Island	1973/5			Gates (1974)
South Dakota	1972		does not include marital status	Gates (1974)
Tennessee	1974/4			Bowdish (2010)
Texas	1973/6	1973/8	does not include marital status	Gates (1974)
Utah	1973/5		does not include marital status; only state-regulated enterprises	Gates (1974)
Vermont	1974/2	1974/7		Gates (1974)
Virginia	1975			U.S. Office of Consumer Affairs (1976)
Washington	1973/3	1973/6		Beckey (1974)
West Virginia*	1973		public accommodation only	Gates (1974)
Wisconsin	1973/8			Gates (1974)

Notes: The table gives an overview of state laws against credit discrimination. If there is information that the dates of the law's passage and effectiveness differed, both are indicated. An asterisk (*) indicates laws which do not apply to home financing.
[†] The wording was changed from "commercial housing" to "housing accommodation or commercial property" by H.B. 141, Regular Session 1985-1986.

1.B.7 Alternative event study estimator

Figure 1.B.11 compares the estimated coefficients from a "conventional" event study with two-way fixed effects to those obtained by using the "interaction-weighted estimator" proposed by Sun and Abraham (2020), which is robust to heterogeneous treatment effects. In a setting without never-treated units, the approach of Sun and Abraham (2020) requires to drop time periods $t \geq \max(e_s)$, because each unit is treated from that point on, such that there is no well-defined control group that would allow to compute their estimator. Therefore, I exclude $\max(e_s) = 1975$ and all subsequent years in the estimation of both specifications. Moreover, I abstain from using additional controls in the two-way fixed effects model for consistency across the two strategies. The patterns produced by both estimation procedures turn out to be very similar.

Figure 1.B.11: Comparison to estimator of Sun and Abraham (2020)

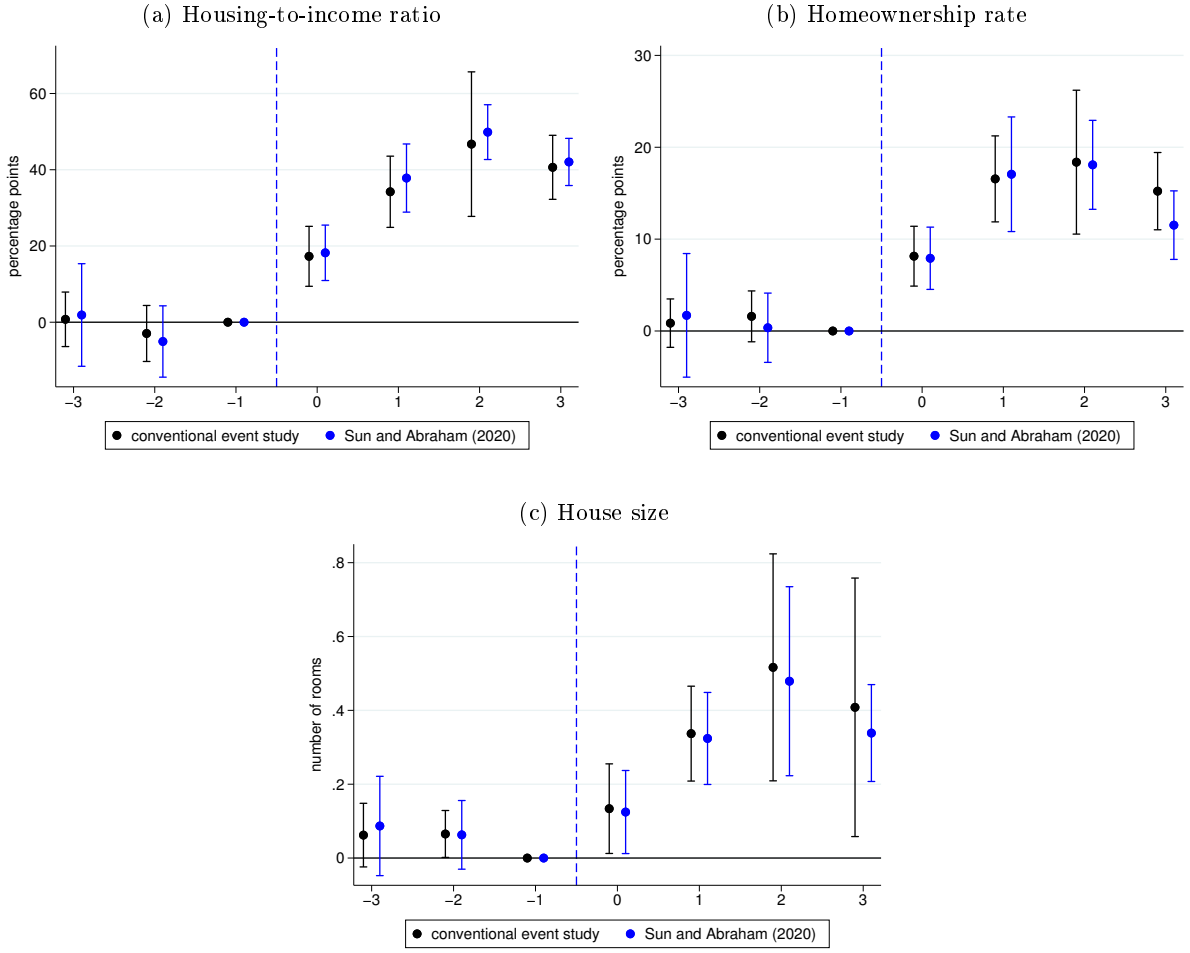

Notes: The graph presents a comparison of the coefficients obtained using a two-way fixed effects model with state and year fixed effects to those obtained with the "interaction-weighted estimator" of Sun and Abraham (2020). The sample is restricted to the years prior to 1975 and no further covariates are included, as required by the approach of Sun and Abraham (2020). Standard errors are clustered at the state level. The whiskers indicate 95% intervals.

1.B.8 Explorative regression: female labor supply

Figure 1.B.12 shows the results of an event study for the wife's probability to work more than 500 hours per year, which corresponds to a quarter of full-time employment. The sample includes young households (below age 45, the median age in my baseline sample in the early 1970s) who

were renters in the three pre-reform years, and whose wife was at most loosely attached to the labor force (i.e., working less than a quarter of annual full-time hours on average over this period). This group may be considered as a proxy for households who are at the beginning at their life cycle shortly after the reform, and have to decide whether to buy a home and whether the wife wants to work. Due to the small sample size (less than 200 households), these results should be taken as suggestive and interpreted with caution. While estimating the same regression on the full sample yields a zero effect, Figure 1.B.12 suggests that married women from the described group increased their labor supply after the reform in treated states, compared to their counterparts in untreated states.

Figure 1.B.12: Explorative event study: female labor supply

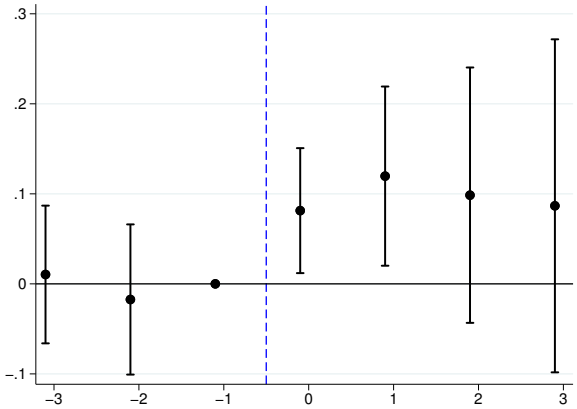

Notes: The graph shows the coefficients on the treatment indicators D_{st}^{j} from equation (1.3), using a dummy for whether the wife works more than 500 hours per year as the outcome. The sample was restricted to households who were renting and below age 45 prior to the reform, with a wife at most loosely attached to the labor force. Standard errors are clustered at the state level. The whiskers indicate 95% intervals.

1.B.9 Share in contemporaneous income

Figure 1.B.13: Housing debt: share of contemporaneous income

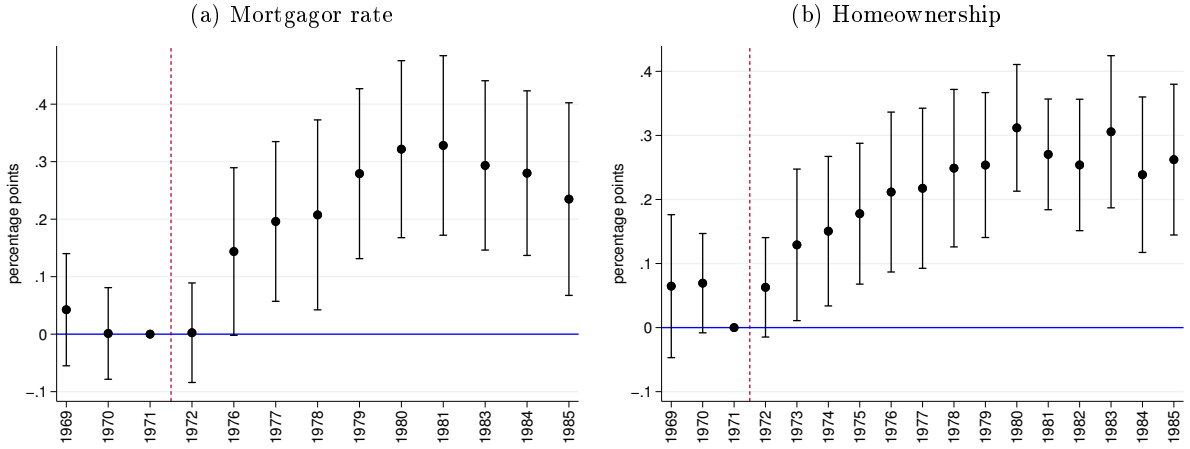

Notes: The graph presents the coefficients on the interaction term in equation (1.2), where $share_i^{pre}$ was replaced with the wife's contemporaneous income share (in percent). The sample consists of the same households as in the baseline. The whiskers indicate 95% intervals.

1.B.10 Non-normalized housing debt and housing assets

Figure 1.B.14: Housing debt and housing

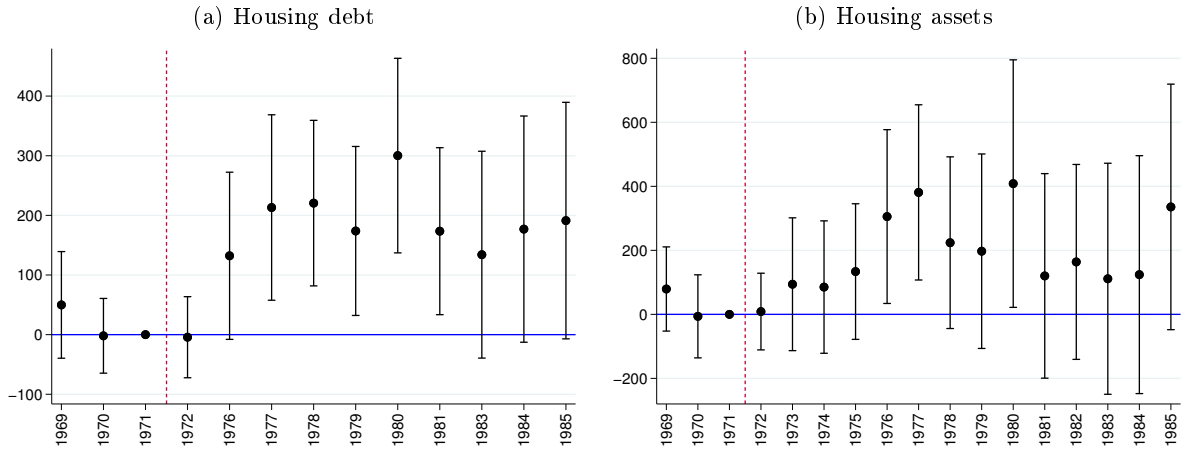

Notes: The graph presents the coefficients on the interaction term in equation (1.2). The wife's income share is defined in percent. The whiskers indicate 95% intervals.

1.B.11 Share of wife in joint labor income of both spouses

Figure 1.B.15: Shares of married households by labor income share of wife in *labor* income

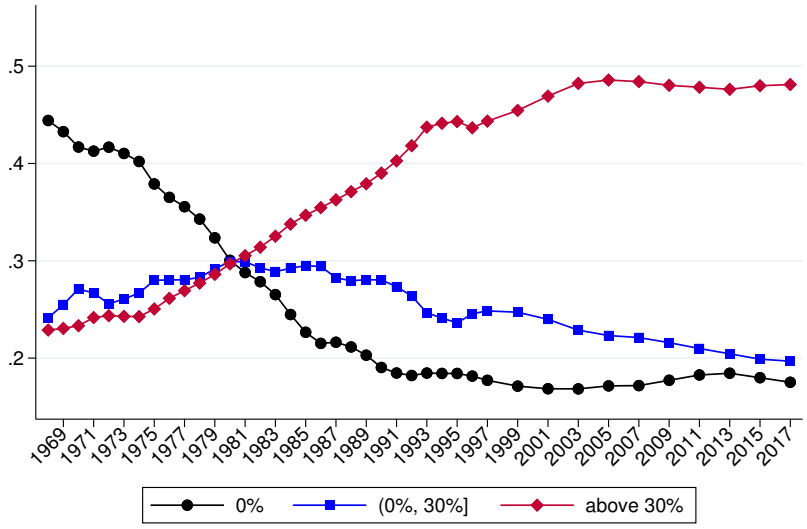

Notes: The graph shows shares among all married households with positive labor income, stratified by the wife's labor income contribution to the couple's total labor income. The series were smoothed by taking a 3-year moving average.

Table 1.B.3: Mortgagor rate: share in *labor* income

	(1)	(2)
Post 1971 × Lab. Inc. Share Wife 71	0.002*** (0.001)	0.001*** (0.000)
25-34		0.309*** (0.039)
35-44		0.412*** (0.047)
45-54		0.428*** (0.050)
55-64		0.346*** (0.053)
65+		0.225*** (0.070)
50% - 90%		0.053*** (0.013)
Top 10%		0.088*** (0.020)
1 Child		0.058*** (0.013)
2 Children		0.147*** (0.019)
3+ Children		0.156*** (0.025)
Household FE	yes	yes
Year FE	yes	no
State-year FE	no	yes
Mean	0.526	0.526
Observations	19,571	19,569

Notes: The table presents results for a version equation (1.B.1) in which the interaction variable is the wife's average share in both spouses' labor income over the pre-reform years. Standard errors are given in parentheses and are clustered at the household and state level (* p<0.1, ** p<0.05, *** p<0.01). The wife's income share is defined in percent.

Figure 1.B.16: Share in *labor* income of head and spouse

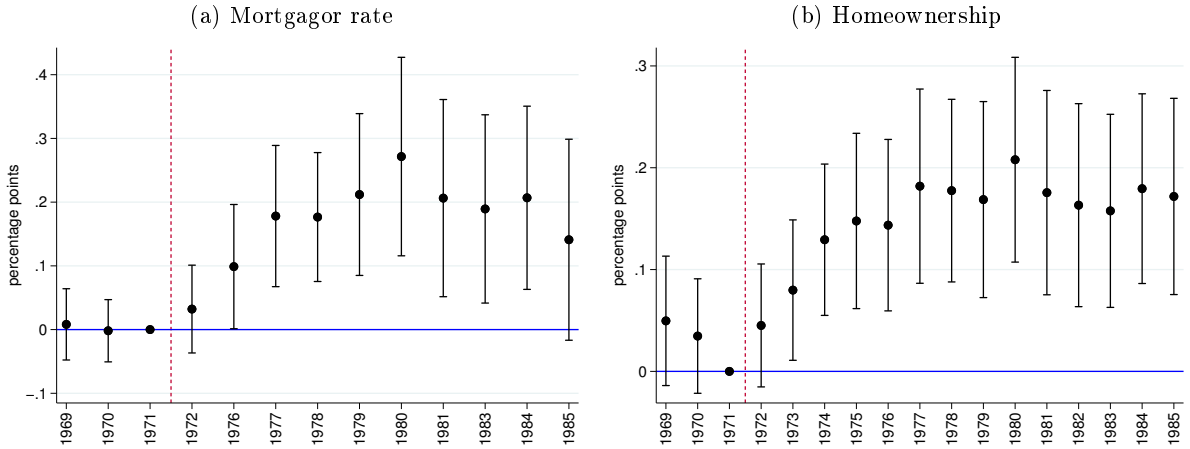

Notes: The graph presents the coefficients on the interaction term in equation (1.2), where $share_i^{pre}$ was replaced with the wife's average share in both spouses' labor income over the pre-reform years (up to 1971). The wife's income share is defined in percent. The whiskers indicate 95% intervals.

Appendix 1.C Results for single households

It should be noted that the ECOA does not prohibit creditors from asking about marital status in general. The only exception are applications for unsecured, separate accounts in separate property states or community property states allowing both spouses to manage and control the community property (Geary 1976). Yet as Geary (1976) points out, Regulation B limits the use a lender can make of marital status information.

One could imagine a scenario in which an increase in the share of singles with housing debt

reflects demographic changes, e.g., an increase in divorces. However, the divorce rates in Figure 1.B.4 have evolved smoothly over time. The same is true for the sample shares of single and married households, as Figure 1.C.1 shows.

Figure 1.C.1: Single and married homeowners

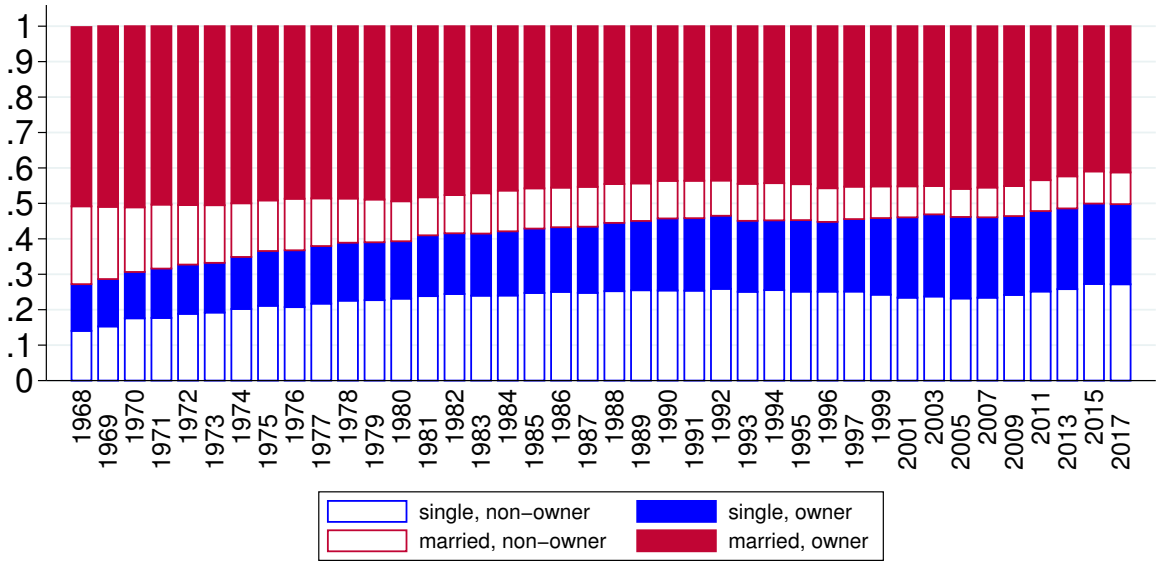

Notes: The graph shows the shares of single and married households with and without a house over time.

Figure 1.C.2 presents descriptive evidence, similar to Figure 1.3 in the main text. It shows that housing debt began to increase faster for singles than for married households with a non-working wife in the early 1970s.[41] A similar pattern emerges when normalizing with total household income.

Figure 1.C.2: Descriptive evidence, singles

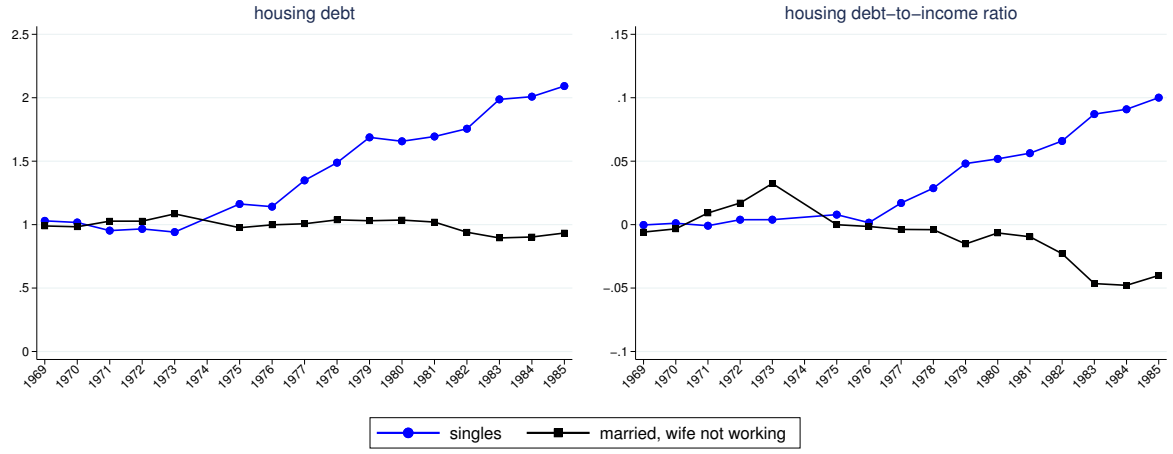

Notes: The graph shows housing debt from the PSID for singles and married households with a non-working wife over time. All series were normalized with their average over the period 1969-1971. The left panel shows average housing debt, and the right panel shows housing debt-to-income, after winsorizing at the 99th percentile within each year. The series were smoothed by taking a 3-year moving average.

[41] I excluded widowed households from the sample because they differ substantially from other single households with respect to demographic characteristics and homeownership.

I estimate a regressions similar to equation (1.2), replacing $share_i^{pre}$ by an indicator for being single in 1971, $single_{i,71}$. A male household head who is single in 1971 might for instance marry a working wife, such that the estimated effect might pick up effects on dual-earner couples if the treatment group was defined only according to single status in 1971. I therefore define the treatment group as single households who were already single in 1971. The control group is chosen as in the binary version of my baseline regression, i.e. married households in which the wife was not working in 1969-1971. I include the same set of controls as in the baseline.

Figure 1.C.3: Housing debt of singles

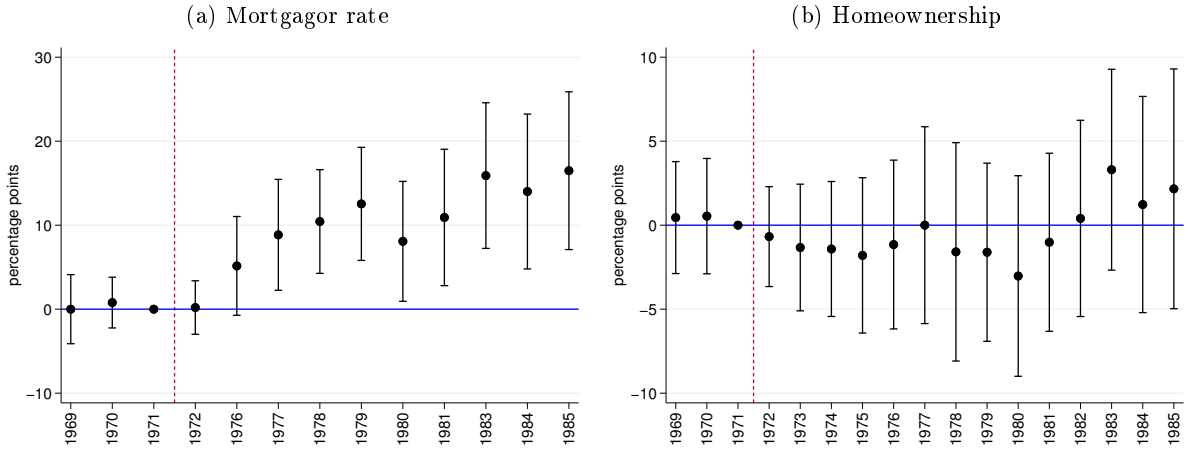

Notes: The graph presents the coefficients on the interaction term of the single dummy and the year dummy, as described in the text. The DTI ratio is defined in percent and was winsorized at the 99th percentile within each year. The whiskers indicate 95% intervals.

Figure 1.C.3 shows the coefficients on the interaction term of the single dummy and the year dummy for regressions with the extensive margin of housing debt and homeownership as the outcome variables. There is a pronounced increase in the extensive margin of singles' debt compared to married households with a non-working wife. However, there is no increase in homeownership. This would be in line with the interpretation that singles could substitute away from other financial sources, such as transfers or loans from their families (cf. Del Boca and Lusardi 2003).

The PSID asks the head about financial help he received from relatives outside the household in the last year. Information on similar transfers to the wife are only available since 1985. Likewise, information on transfers from non-relatives outside the household and loans from relatives are only available in more recent survey waves. I found that single heads had a lower probability of receiving financial support from relatives than heads from the control group in the post-reform period, whereas the difference was close to zero in the pre-reform years. Although the point estimates are not statistically significant at conventional levels, I interpret this as suggestive evidence in support of the substitution channel. The results are available upon request.

Figure 1.C.4 stratifies the singles by sex. It shows that there is no evidence of differential effects on single women compared to single men.

Figure 1.C.4: Housing debt of singles, by sex

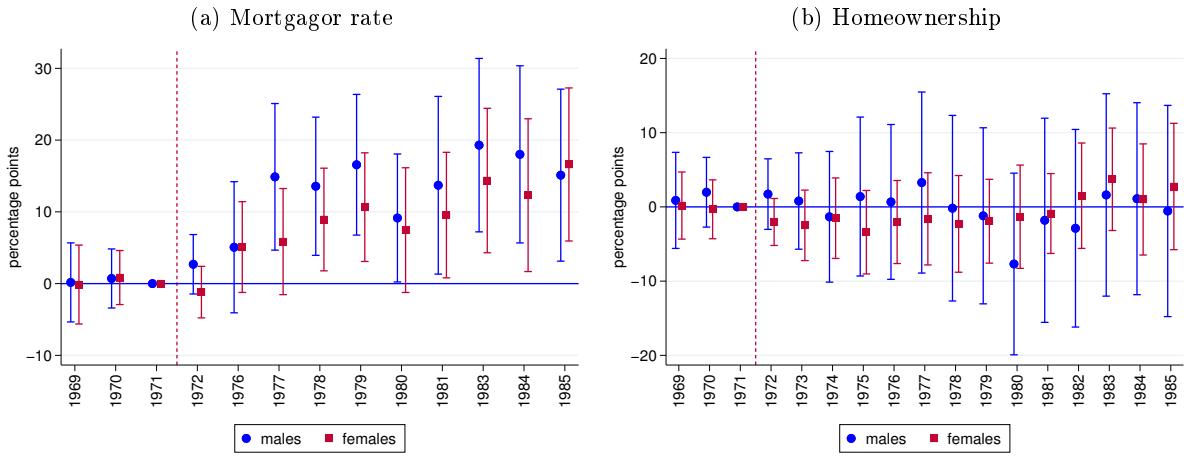

Notes: The graph presents the coefficients on the interaction term of the single dummy, the year dummy and a dummy for singles' sex. The DTI ratio is defined in percent and was winsorized at the 99th percentile within each year. The whiskers indicate 95% intervals.

Appendix 1.D Additional information on life-cycle model

This section provides additional details on the life-cycle model and its solution.

1.D.1 Supplementary calculations

$$d_{j+1} = (1+r)d_j - m_j,$$

where m_j are per-period mortgage payments. The constraint $d_{j+1} \leq d_j$ implies

$$(1+r)d_j - m_j \leq d_j \quad \Leftrightarrow \quad rd_j \leq m_j \ .$$

If the mortgage is scheduled to be amortized over T periods using constant payments m, a debt-to-income constraint upon origination is directly proportional to a debt-service-to-income constraint. To see this, let d_0 be the original mortgage upon origination:

$$d_T = (1+r)^T d_0 - \sum_{i=0}^{T-1}(1+r)^i m = (1+r)^T d_0 - \frac{(1+r)^T - 1}{r} m = 0 \quad \Leftrightarrow$$
$$\frac{d_0}{y_0} = \frac{(1+r)^T - 1}{r(1+r)^T}\frac{m}{y_0}. \tag{1.D.1}$$

1.D.2 Computational solution and life-cycle profiles

As explained in Attanasio et al. (2012), the structure of the optimization problem does not allow to rely on the existence of smooth first order conditions when solving the model. Therefore, I use value function iteration to find a solution. I compute "conditional" value functions for each current housing state ($h \in \{0, 1, 2\}$) based on corresponding "conditional" asset grids with 120 grid points. The solution is found recursively by iterating backwards from the end of the life cycle.

The wage processes from equation (1.4) are discretized using the method of Tauchen (1986), as adapted to the life-cycle setup and structure of the wage process by Attanasio et al. (2012). A choice of 14 grid points (both for men and women) yields a reasonably good approximation of the autoregressive processes. Figure 1.D.1a shows the income profiles of husbands and wives under certainty, obtained by multiplying the estimated wage profiles with 40 hours times 50 weeks for the husband, which corresponds to a typical contract in the data, and 20 hours times 50 weeks for the wife, corresponding to a typical half-time contract. Figure 1.D.1b illustrates how these profiles map into the DTI constraint when the household is allowed to count 50% or 100% of the wife's income toward the mortgage.

Figure 1.D.1: Income profiles and constraints

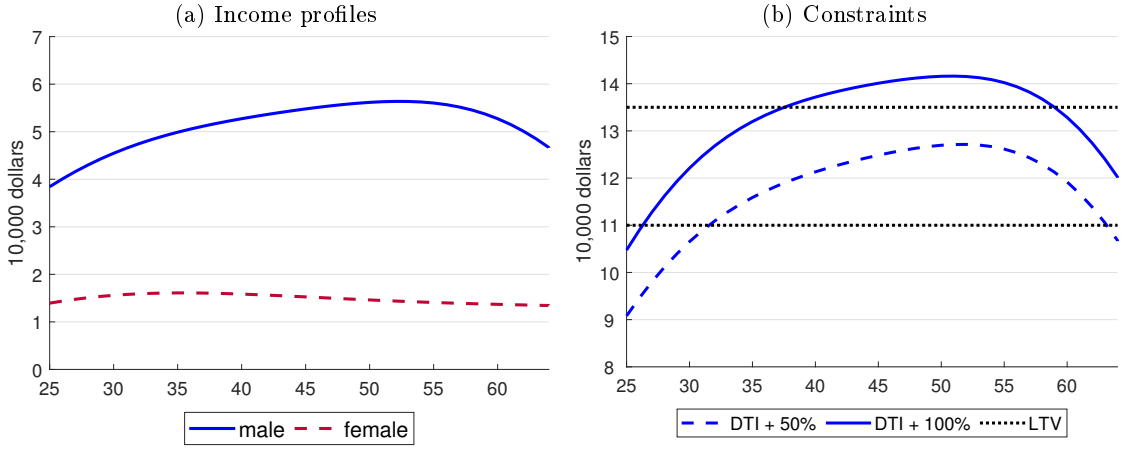

Notes: The left panel shows the income profiles from estimating equation (1.4) and abstracting from income risk, assuming a standard contract of 50 work weeks à 40 hours for the husband, and a half-time contract of 20 hours for the wife. The right panel shows the LTV constraints associated with the smaller and the bigger house, as well as the DTI constraints when the household is allowed to count 50% or 100% of the wife's income from the left panel toward the mortgage.

The initial distribution of net wealth is estimated from the SCF+ for the period 1969-1971, restricting the sample to married households with a head between ages 24 and 26. A few households report negative net wealth in the data, mainly due to personal debt. As my model does not include personal debt, I follow Pizzinelli (2018) and winsorize the distribution at 0. Moreover, I winsorize it at the 99th percentile to exclude a few households reporting exceptionally high net wealth. Households start their life cycle without owning a house.

The age-dependent female leisure parameter $\theta^f(j)$ is modeled as a third-order polynomial and the factor $\chi(j)$, which pre-multiplies the basic housing preferences μ^h, is modeled as the maximum of a second-order polynomial and one. The coefficients are reported in Table 1.D.1, and the resulting profiles are depicted in Figure 1.D.2. The female leisure preference is initially high and increases until the early thirties. Thereafter, the probability of additional children arriving decreases (see Figure 1.D.3a), and women's leisure preferences decrease. From the mid-fifties on, when retirement is approaching and the first grandchildren arrive, the leisure preference increases again. The housing preferences start at a high level and then decline almost linearly until the late fifties, which is the point in the life cycle when households sizes stabilize in the data (see Figure 1.D.3b).

Table 1.D.1: Polynomical coefficients

$\theta^f(j)$		$\chi(j)$	
θ_0	-0.5997	χ_0	6.7559
θ_1	0.0707	χ_1	-0.1035
θ_2	-0.0017	χ_2	0.0001
θ_3	1.2748e-5		

Notes: The table shows the coefficients of the polynomials used to model the age-dependent female leisure preferences $\theta^f(j)$ and housing preferences $\chi(j)$.

Figure 1.D.2: Age-dependent preference parameters

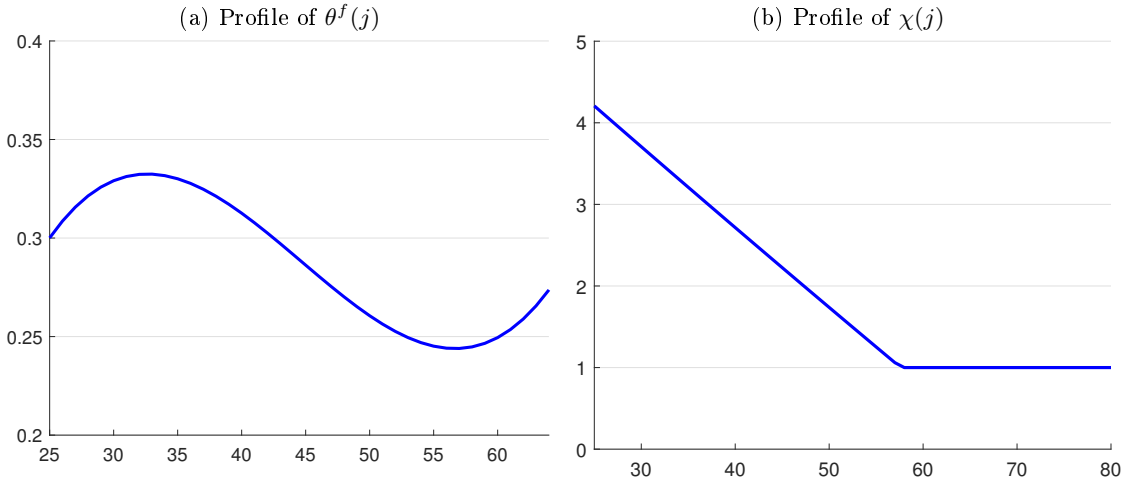

Notes: The graph shows the age-dependent leisure preferences $\theta^f(j) = \theta_0 + \theta_1 j + \theta_2 j^2 + \theta_3 j^3$ and the housing preferences $\chi(j) = \max(\chi_0 + \chi_1 j + \chi_2 j^2, 1)$.

Figure 1.D.3: Arrival of new children and household size over the life cycle

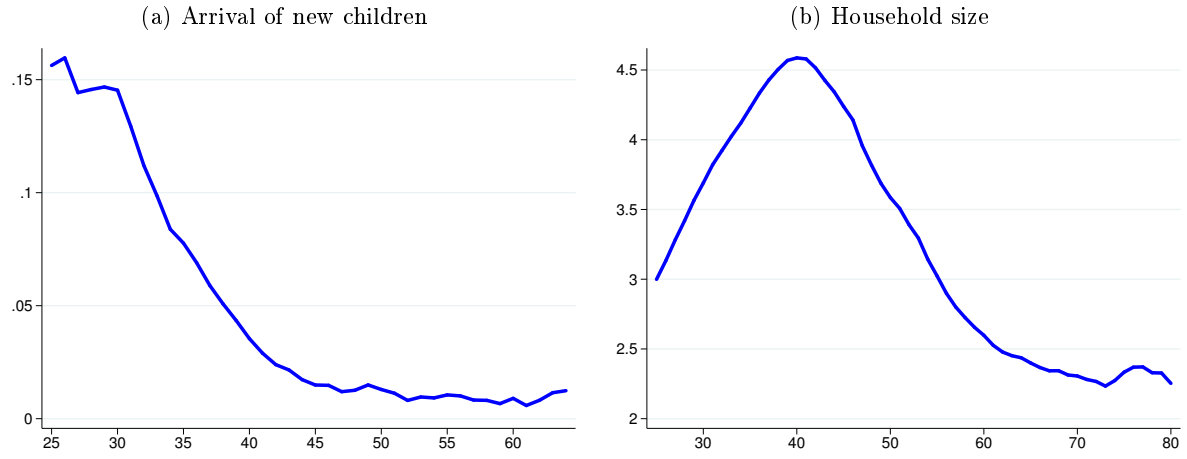

Notes: The left panel shows the share of households in the baseline sample reporting an increase in the number of children over the life cycle. The right panel shows the average household size over the life cycle. The data were smoothed by taking a three-year moving average.

Table 1.D.2: Targeted moments

	data	model
homeownership, 25-34	0.58	0.61
homeownership, 35-44	0.78	0.78
homeownership, 45-54	0.84	0.83
homeownership, 55-64	0.8	0.83
mortgagor rate, 25-34	0.55	0.54
mortgagor rate, 35-44	0.68	0.69
mortgagor rate, 45-54	0.57	0.65
mortgagor rate, 55-64	0.35	0.29
FLFP, 25-34	0.43	0.45
FLFP, 35-44	0.39	0.36
FLFP, 45-54	0.42	0.39
FLFP, 55-64	0.39	0.39

Notes: The table compares the homeownership, mortgagor and FLFP rates by age group in the model and the data. The FLFP rate in the data refers to women working at least 500 hours per year.

Appendix 1.E Questions related to the ECOA from the 1977 SCF

Elliehausen and Durkin (1989) draw on evidence from the Survey of Consumer Finances (SCF) of 1977, which included several questions related to the ECOA. In the following, I will summarize the answers to several of these questions. Indeed, the responses provide mixed evidence of self-reported discrimination, consistent with the findings of Elliehausen and Durkin (1989). However, there are several reasons why the responses may not adequately reflect the impact of the ECOA. First, the survey was conducted in August and September 1977, five years after the congressional hearings of 1972 that prepared the ground for the ECOA. Second, more than 99% of married household in the 1977 SCF have a male head (cf. Kuhn, Schularick, and Steins 2020). Questions addressed to the heads may conceal problems of their wives to obtain credit. Third, the question if a couple was unable to obtain the desired amount of credit does not capture whether they had adjusted their desired amount based on common lending practices. An important aspect of the ECOA was to educate the public about their new financial rights (Cairns 1976, Geary 1976), which may have changed households' reference points for what they can afford.

Singles, in particular women, reported larger perceived difficulties to obtain credit (see Table 1.E.1). Looking at the reasons in Figure 1.E.1, marital status and sex seem to be an important problem for single women, whereas these aspects are perceived as less problematic by single men, and hardly any married households report such problems. However, as the question asks for the respondent's credit experiences, it remains silent on potential problems of the spouse.

Single women also have above-average approval rates for the response options lack of job and income. This may in part reflect discriminatory practices, because lenders commonly did not accept alimony and child support as income before the implementation of the ECOA (see Cairns 1976). It may however also reflect justified economic reasons for single women's problems to obtain credit. Both male and female singles also report the type or duration of their job as a problem more often than married households. Credit records are the biggest problem for all groups, and especially for male singles. These results indicate that single households' difficulties to obtain credit were not only related to discrimination, but also to economic factors – even according to the self-perception of these households.

Table 1.E.1: Would it be difficult for people like yourself to get credit?

	single men	single women	married, wive no work	married, wive works	Total
no	64.70	53.80	80.59	84.70	74.31
yes	30.95	43.54	17.08	13.93	23.36
depends	4.36	2.66	2.33	1.37	2.32
Total	100.00	100.00	100.00	100.00	100.00

Figure 1.E.1: Why would it be difficult?

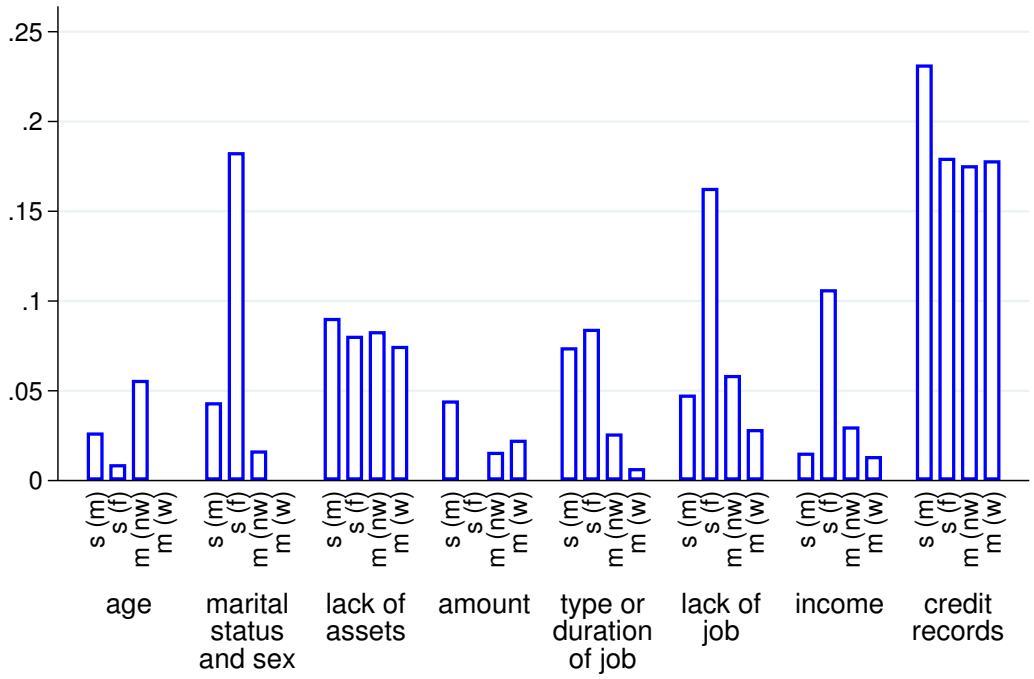

Notes: s(m) = single, male; s(f) = single, female; m(w) = married, wife s; m(nw) = married, wife does not work; 2 mentions possible

When asked whether they had been treated unfairly in their credit transactions, the approval rate of singles is actually *lower* than that of married households, and that of female singles even lower than that of male singles (Table 1.E.2). Among married households, the approval rate is slightly higher for households with a working wife, but the difference is small. Regarding the underlying problems, marital status and sex are again mainly mentioned by single women (Figure 1.E.2). Overall, the share of households mentioning this source of the problem is small. Interestingly, the share of households reporting they were refused all or part of the desired amount is higher for single women than single men, and higher for married households with a working as opposed to a non-working wife.

Table 1.E.2: Have you been treated unfairly in your credit transactions?

	single men	single women	married, wive no work	married, wive works	Total
no	76.67	81.52	75.19	72.44	75.93
yes	23.33	18.48	24.81	27.56	24.07
Total	100.00	100.00	100.00	100.00	100.00

Figure 1.E.2: What was the problem?

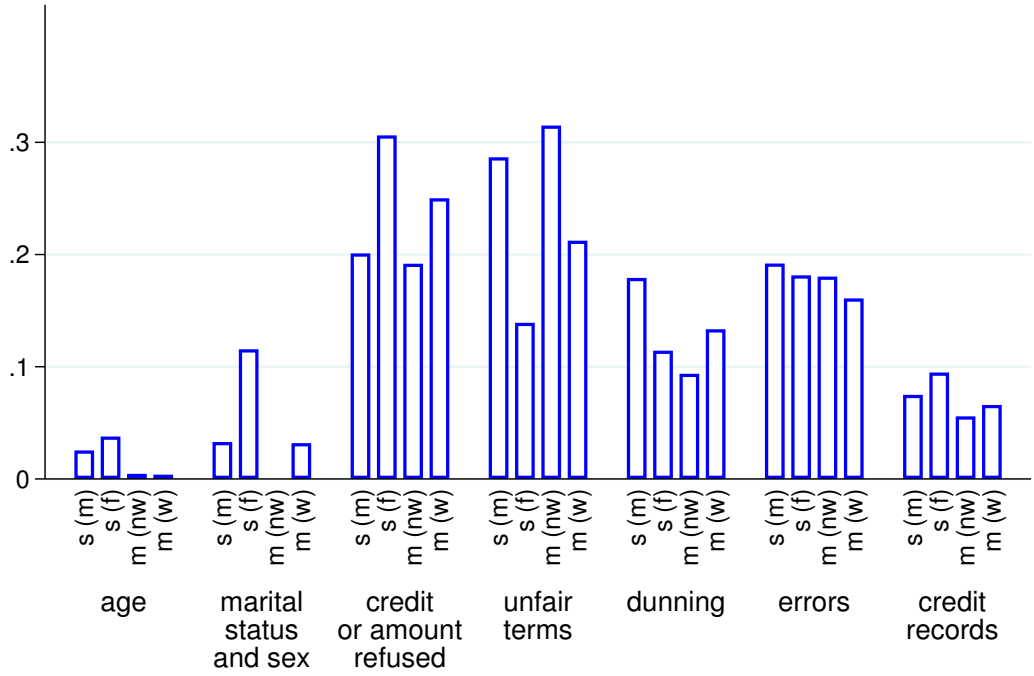

Notes: s(m) = single, male; s(f) = single, female; m(w) = married, wife works; m(nw) = married, wife does not work; 2 mentions possible

Consistently, the share of households reporting they were turned down for credit is higher among households with a working as opposed to a non-working wife (Table 1.E.3). However, it is lower for single women than for single men. Among the reasons, marital status and sex are once more mainly reported by single women, and do not play a big role (Figure 1.E.3). Again, credit records are the biggest problem, this time with the largest response rate among married households with a working wife.

Table 1.E.3: Have you been turned down for credit?

	single men	single women	married, wive no work	married, wive works	Total
no	78.03	82.33	86.98	82.93	83.54
yes	21.97	17.67	13.02	17.07	16.46
Total	100.00	100.00	100.00	100.00	100.00

Figure 1.E.3: Why were you turned down?

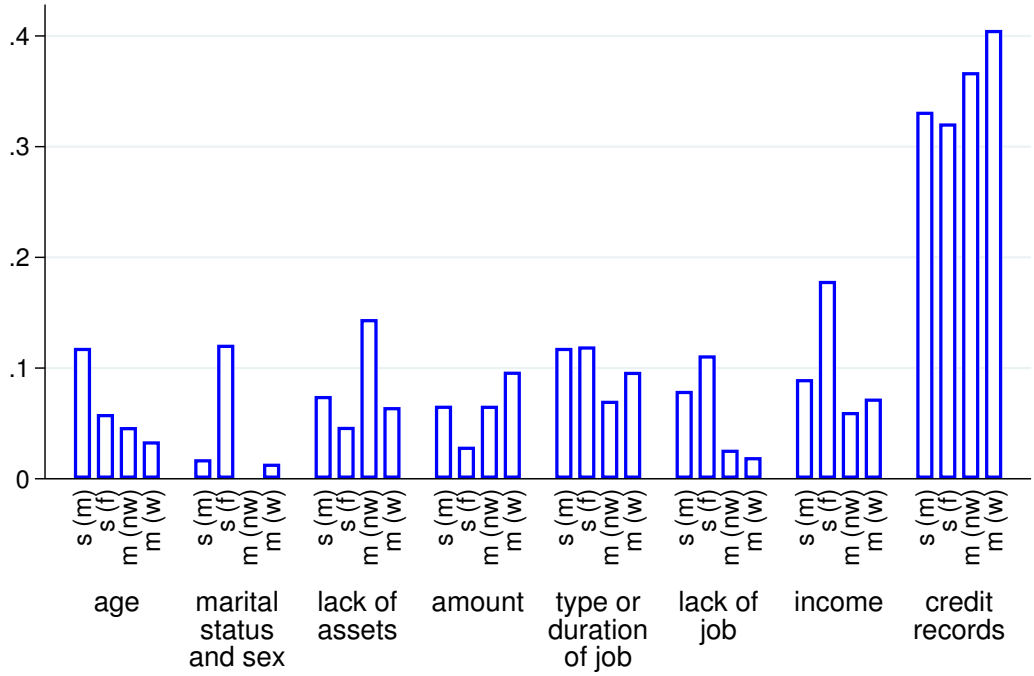

Notes: s(m) = single, male; s(f) = single, female; m(w) = married, wife works; m(nw) = married, wife does not work; 2 mentions possible

When asked specifically whether they had been unable to get as much credit as they wanted from a particular lender in the past few years, the share of affirmative answers is larger among married households if the wife is working, yet lower than among single men (see Figure 1.E.4). However, as discussed above, this question does not capture whether households have calculated the desired amount based on discriminatory market practices.

Figure 1.E.4: Were you unable to get as much credit as you wanted from a particular lender in the past few years?

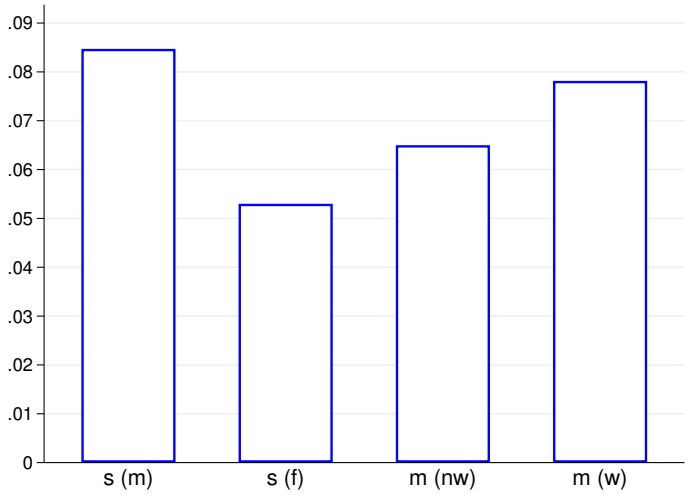

Notes: s(m) = single, male; s(f) = single, female; m(w) = married, wife works; m(nw) = married, wife not working

References

Agarwal, Sumit et al. (2020). *Searching for Approval*. Working Paper 27341. National Bureau of Economic Research.

Altonji, Joseph G. and Rebecca M. Blank (1999). "Race and Gender in the Labor Market." In: *Handbook of Labor Economics*. Vol. 3. Elsevier. Chap. 48, pp. 3143–3259.

Andersen, Steffen, Julie Marx, Kasper Meisner Nielsen, and Lise Vesterlund (Feb. 2021). "Gender Differences in Negotiation: Evidence from Real Estate Transactions." In: *The Economic Journal*.

Atalay, Kadir, Garry Barrett, and Rebecca Edwards (2016). *House Prices, Household Debt and Labour Supply in Australia*. Working Paper 266. AHURI Final Report.

Attanasio, Orazio P., Renata Bottazzi, Hamish W. Low, Lars Nesheim, and Matthew Wakefield (2012). "Modelling the Demand for Housing Over the Life Cycle." In: *Review of Economic Dynamics* 15.1, pp. 1–18.

Attanasio, Orazio, Peter Levell, Hamish Low, and Virginia Sánchez-Marcos (2018). "Aggregating Elasticities: Intensive and Extensive Margins of Women's Labor Supply." In: *Econometrica* 86.6, pp. 2049–2082.

Attanasio, Orazio, Hamish Low, and Virginia Sánchez-Marcos (2005). "Female Labor Supply as Insurance against Idiosyncratic Risk." In: *Journal of the European Economic Association* 3.2-3, pp. 755–764.

Bailey, Martha J. (2006). "More Power to the Pill: The Impact of Contraceptive Freedom on Women's Life Cycle Labor Supply." In: *The Quarterly Journal of Economics* 121.1, pp. 289–320.

Bartscher, Alina K., Moritz Kuhn, Moritz Schularick, and Ulrike I. Steins (2020). *Household Debt and Inequality in the U.S., 1950-2019*. Working Paper.

Bastian, Jacob (2020). "The Rise of Working Mothers and the 1975 Earned Income Tax Credit." In: *American Economic Journal: Economic Policy*.

Beckey, Sylvia L. (1974). *Women and Credit: Available Legal Remedies Against Discriminatory Practices*. Tech. rep. Washington, D.C.: Congressional Research Service.

Borella, Margherita, Mariacristina De Nardi, and Fang Yang (2018). "The aggregate implications of gender and marriage." In: *The Journal of the Economics of Ageing* 11. Macroeconomics of Aging, pp. 6–26.

Bottazzi, Renata, Hamish Low, and Matthew Wakefield (2007). *Why Do Home Owners Work Longer Hours?* Working Paper. Institute for Fiscal Studies.

Bowdish, Lawrence A. (2010). "Invidious Distinctions: Credit Discrimination against Women, 1960s–Present." PhD thesis. The Ohio State University.

Braxton, J. Carter, Kyle F. Herkenhoff, and Gordon M. Phillips (2020). *Can the Unemployed Borrow? Implications for Public Insurance*. Working Paper 27026. National Bureau of Economic Research.

Brock, Michelle and Ralph De Haas (2020). *Discriminatory Lending: Evidence from Bankers in the Lab*. Working Paper.

Cairns, John W. (1976). "Credit Equality Comes to Women: An Analysis of the Equal Credit Opportunity Act." In: *San Diego Law Review* 13.4, p. 960.

Chang, Minsu (2018). *A House Without a Ring: The Role of Changing Marital Transitions for Housing Decisions*. Working Paper.

Costa, Dora L. (2000). "From Mill Town to Board Room: The Rise of Women's Paid Labor." In: *Journal of Economic Perspectives* 14.4, pp. 101–122.

Cuomo, Andrew A. (1981). "Equal Credit Opportunity Act: How Much Can Women Expect." In: *Journal of Legislation* 8, p. 121.

Daly, Moira, Dmytro Hryshko, and Iourii Manovskii (2018). *Improving the Measurement of Earnings Dynamics*. Working Paper 22938. National Bureau of Economic Research.

Dao Bui, Kien and Ejindu S. Ume (2020). "Credit Constraints and Labor Supply: Evidence from Bank Branching Deregulation." In: *Economic Inquiry* 58.1, pp. 335–360.

Del Boca, Daniela and Annamaria Lusardi (2003). "Credit Market Constraints and Labor Market Decisions." In: *Labour Economics* 10.6, pp. 681–703.

Diamond, Rebecca, Adam Guren, and Rose Tan (2020). *The Effect of Foreclosures on Homeowners, Tenants, and Landlords*. Working Paper.

Druedahl, Jeppe (2015). *The Demand for Housing Over the Life-cycle under Long-term Gross Debt Contracts*. Working Paper.

Elliehausen, Gregory and Thomas Durkin (June 1989). "Theory and Evidence of the Impact of Equal Credit Opportunity: An Agnostic Review of the Literature." In: *Journal of Financial Services Research* 2, pp. 89–114.

Exler, Florian and Michèle Tertilt (2020). "Consumer Debt and Default: A Macro Perspective." In: *Oxford Research Encyclopedia of Economics and Finance*.

Favara, Giovanni and Jean Imbs (2015). "Credit Supply and the Price of Housing." In: *American Economic Review* 105.3, pp. 958–92.

Fernández, Raquel (2013). "Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century." In: *American Economic Review* 103.1, pp. 472–500.

Field, Erica, Rohini Pande, Natalia Rigol, Simone Schaner, and Charity Troyer Moore (2021). "On Her Own Account: How Strengthening Women's Financial Control Impacts Labor Supply and Gender Norms." In: *American Economic Review* 111.7.

Foote, B. Ellington and Meredith Peterson (2008). "VA-Home Loan Guaranty Program: An Overview." In: *Congressional Research Service, Library of Congress*.

Foote, Christopher L., Lara Loewenstein, and Paul S. Willen (2018). *Technological Innovation in Mortgage Underwriting and the Growth in Credit: 1985-2015*. Working Paper.

Fortin, Nicole M. (1995). "Allocation Inflexibilities, Female Labor Supply, and Housing Assets Accumulation: Are Women Working to Pay the Mortgage?" In: *Journal of Labor Economics* 13.3, pp. 524–557.

Fuster, Andreas, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther (forthcoming). "Predictably Unequal? The Effects of Machine Learning on Credit Markets." In: *Journal of Finance*.

Gates, Margaret J. (1974). "Credit Discrimination Against Women: Causes and Solutions." In: *Vanderbilt Law Review* 27, p. 409.

Geary, Anne J. (1976). "Equal Credit Opportunity — An Analysis of Regulation B." In: *The Business Lawyer* 31, pp. 1641–1658.

Gigot, Paul A. (1981). "Fading Dream." In: *The Wall Street Journal*.

Gladstone, Leslie W. (2004). "Equal Rights Amendments: State Provisions." In: *CRS Report for Congress*. Congressional Research Service, Library of Congress.

Goldin, Claudia and Lawrence F. Katz (2002). "The Power of the Pill: Oral Contraceptives and Women's Career and Marriage Decisions." In: *Journal of Political Economy* 110.4, pp. 730–770.

Golding, Edward, Edward Szymanoski, and Pamela Lee (2014). *FHA at 80: Preparing for the Future*. Tech. rep.

Goldsmith-Pinkham, Paul and Kelly Shue (forthcoming). "The Gender Gap in Housing Returns." In: *Journal of Finance*.

Goodman-Bacon, Andrew (2021). "Difference-in-Differences with Variation in Treatment Timing." In: *Journal of Econometrics*.

Greenwald, Daniel (2018). *The Mortgage Credit Channel of Macroeconomic Transmission*. Working Paper.

Greenwood, Jeremy, Ananth Seshadri, and Mehmet Yorukoglu (2005). "Engines of Liberation." In: *The Review of Economic Studies* 72.1, pp. 109–133.

Gruber, Jonathan (2004). "Is Making Divorce Easier Bad for Children? The Long-Run Implications of Unilateral Divorce." In: *Journal of Labor Economics* 22.4, pp. 799–833.

Haurin, Donald R. and Duewa A. Kamara (1992). "The homeownership decision of female-headed households." In: *Journal of Housing Economics* 2.4, pp. 293–309.

Hazan, Moshe, David Weiss, and Hosny Zoabi (2019). "Women's Liberation as a Financial Innovation." In: *Journal of Finance*.

Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. Violante (2010). "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967–2006." In: *Review of Economic Dynamics* 13.1. Special issue: Cross-Sectional Facts for Macroeconomists, pp. 15–51.

Herkenhoff, Kyle F (Feb. 2019). "The Impact of Consumer Credit Access on Unemployment." In: *The Review of Economic Studies* 86.6, pp. 2605–2642.

Herkenhoff, Kyle, Gordon Phillips, and Ethan Cohen-Cole (2020). *How Credit Constraints Impact Job Finding Rates, Sorting & Aggregate Output*. Working Paper.

Hill, Daniel D. (2002). *Advertising to the American Woman, 1900-1999*. Ohio State University Press.

Hill, Herbert (1977). "The Equal Employment Opportunity Acts of 1964 and 1972: A Critical Analysis of the Legislative History and Administration of the Law." In: *Industrial Relations Law Journal* 2.1, pp. 1–96.

Hill, Martha S. (1991). *The Panel Study of Income Dynamics: A User's Guide*. Vol. 2. SAGE Publications.

Hurst, Erik and Frank Stafford (2004). "Home Is Where the Equity Is: Mortgage Refinancing and Household Consumption." In: *Journal of Money, Credit and Banking* 36.6, pp. 985–1014.

Hyland, Marie, Simeon Djankov, and Pinelopi Koujianou Goldberg (2020). "Gendered Laws and Women in the Workforce." In: *American Economic Review: Insights*.

Hyman, Louis (2012). *Debtor Nation: The History of America in Red Ink*. Vol. 72. Princeton University Press.

Hyslop, Dean R. (2001). "Rising U.S. Earnings Inequality and Family Labor Supply: The Covariance Structure of Intrafamily Earnings." In: *American Economic Review* 91.4, pp. 755–777.

Jakucionyte, Egle and Swapnil Singh (2020). *Bowling Alone, Buying Alone: The Decline of Co-Borrowers in the US Mortgage Market*. Working Paper.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor (2017). "Macrofinancial History and the New Business Cycle Facts." In: *NBER Macroeconomics Annual* 31.1, pp. 213–263.

Juhn, Chinhui and Simon Potter (2006). "Changes in Labor Force Participation in the United States." In: *Journal of Economic Perspectives* 20.3, pp. 27–46.

Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante (2020). "The Housing Boom and Bust: Model Meets Evidence." In: *Journal of Political Economy* 128.9, pp. 3285–3345.

Kendig, Dennis (1973). "Discrimination Against Women in Home Mortgage Financing." In: *Yale Review of Law and Social Action* 3.2, p. 3.

Kleven, Henrik, Camille Landais, Johanna Posch, Andreas Steinhauer, and Josef Zweimüller (2019). "Child Penalties across Countries: Evidence and Explanations." In: *AEA Papers and Proceedings* 109, pp. 122–26.

Kohlhase, Janet E. (1986). "Labor Supply and Housing Demand for One- and Two-Earner Households." In: *The Review of Economics and Statistics* 68.1, pp. 48–57.

Kuhn, Moritz, Moritz Schularick, and Ulrike I. Steins (2020). "Income and Wealth Inequality in America, 1949–2016." In: *Journal of Political Economy* 128.9, pp. 3469–3519.

Kumar, Anil and Che-Yuan Liang (2019). *Labor Market Effects of Credit Constraints: Evidence from a Natural Experiment*. Working Papers. Federal Reserve Bank of Dallas.

Ladd, Helen F. (1982). "Equal Credit Opportunity: Women and Mortgage Credit." In: *The American Economic Review* 72.2, pp. 166–170.

Lally, Maureen Ellen (1974). "Women and Credit." In: *Duquesne Law Review* 12.4, pp. 863–890.

"Legislative Column" (1975). In: *Clearinghouse Review* 9.2, p. 96.

Leombroni, Matteo, Monika Piazzesi, Martin Schneider, and Ciaran Rogers (2020). *Inflation and the Price of Real Assets*. Working Paper 26740. National Bureau of Economic Research.

Maki, Dean M (2001). "Household Debt and the Tax Reform Act of 1986." In: *American Economic Review* 91.1, pp. 305–319.

Martin, Robert F. and Joseph W. Gruber (2004). *Does Housing Wealth Make Us Less Equal? The Role of Durable Goods in the Distribution of Wealth*. Econometric Society 2004 North American Summer Meetings 15.

Michel, Jean-Baptiste et al. (2011). "Quantitative Analysis of Culture Using Millions of Digitized Books." In: *Science* 331.6014, pp. 176–182.

Milgram, Grace (1994). *A Chronology of Housing Legislation and Selected Executive Actions, 1892-1992: Congressional Research Service Report Prepared for the Committee on Banking, Finance, and Urban Affairs and the Subcommittee on Housing and Community Development, House of Representatives, One Hundred Third Congress, First Session*. Vol. 4. US Government Printing Office.

Morse, Adair and Karen Pence (2020). *Technological Innovation and Discrimination in Household Finance*. Working Paper 26739. National Bureau of Economic Research.

National Commission on Consumer Finance (1972). *Consumer Credit in the United States*. Report. U.S. Department of the Treasury.

Offer, Avner (2007). *The Challenge of Affluence*. Oxford University Press.

Pfeffer, Fabian T., Robert F. Schoeni, Arthur Kennickell, and Patricia Andreski (2016). "Measuring Wealth and Wealth Inequality: Comparing Two U.S. Surveys." In: *Journal of Economic and Social Measurement* 41.2, pp. 103–120.

Pizzinelli, Carlo (2018). *Housing, Borrowing Constraints, and Labor Supply Over the Life Cycle*. Working Paper.

Poterba, James M. (1984). "Tax Subsidies to Owner-Occupied Housing: An Asset-Market Approach." In: *The Quarterly Journal of Economics* 99.4, pp. 729–752.

— (1991). "House Price Dynamics: The Role of Tax Policy and Demography." In: *Brookings Papers on Economic Activity* 2, pp. 143–203.

Pruitt, Seth and Nicholas Turner (2020). "Earnings Risk in the Household: Evidence from Millions of US Tax Returns." In: *American Economic Review: Insights* 2.2, pp. 237–54.

Schmidheiny, Kurt and Sebastian Siegloch (2020). *On Event Study Designs and Distributed-Lag Models: Equivalence, Generalization and Practical Implications*. Working Paper.

Smith, James F. (1977). "The Equal Credit Opportunity Act of 1974: A Cost/Benefit Analysis." In: *The Journal of Finance* 32.2, pp. 609–622.

Solon, Gary, Steven J. Haider, and Jeffrey M. Wooldridge (2015). "What Are We Weighting For?" In: *Journal of Human Resources* 50.2, pp. 301–316.

Statement of Economists (1972). *Federal Government's Role in the Achievement of Equal Opportunity in Housing: Hearings Before the Civil Rights Oversight Subcommittee (Subcommittee No. 4) of the Committee on the Judiciary, House of Representatives, 92d Congress*. U.S. Government Printing Office.

Sun, Liyang and Sarah Abraham (2020). "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects." In: *Journal of Econometrics*.

Tauchen, George (1986). "Finite state markov-chain approximations to univariate and vector autoregressions." In: *Economics Letters* 20.2, pp. 177–181.

Tewari, Ishani (2014). "The Distributive Impacts of Financial Development: Evidence from Mortgage Markets during US Bank Branch Deregulation." In: *American Economic Journal: Applied Economics* 6.4, pp. 175–96.

Thurston, Chloe N. (2018). *At the Boundaries of Homeownership: Credit, Discrimination, and the American State*. Cambridge University Press.

Tzioumis, Konstantinos (2017). "Mortgage (mis)pricing: The Case of Co-borrowers." In: *Journal of Urban Economics* 99, pp. 79–93.

U.S. Advisory Commission on Intergovernmental Relations (1976). *State Actions in 1975*.

U.S. Department of Labor (1975). *1975 Handbook on Women Workers. Bulletin 297*. Government Printing Office Washington, D.C.

U.S. Office of Consumer Affairs (1975). *State Consumer Action: Summary '74*.

— (1976). *State Consumer Action: Summary '75*.

von Bismarck-Osten, Clara, Kirill Borusyak, and Uta Schönberg (2020). *The Role of Schools in Transmission of the SARS-CoV-2 Virus: Quasi-Experimental Evidence from Germany*. Working Paper.

Chapter 2

Household Debt and Inequality in the United States, 1950-2019

with Moritz Kuhn, Moritz Schularick and Ulrike Steins

This chapter studies the secular increase in U.S. household debt over the past seven decades. We exploit a novel household-level dataset covering the joint distribution of debt, income and wealth. The data show that rising housing debt of middle-class families has played a central role in the debt increase. We further find that income and debt growth comoved closely until the late 1970s, but have diverged strongly thereafter. Since the 1950s, the vast majority of U.S. household debt has always been housing debt. In the postwar era, housing debt mainly increased because more households took on debt. Yet since the 1980s, it has mainly increased because households took on larger amounts of debt. We show that home-equity-based borrowing can account for about half of the housing debt increase over the boom period ending with the 2008 crisis. This has led to pronounced changes in life-cycle debt-to-income profiles, with higher levels of debt at older ages. Yet since asset values grew as well, life-cycle wealth-to-income ratios have remained similar across cohorts. Our data also show that black households, with their lower homeownership rates and housing assets, have only played a very limited role in the extraction boom and overall debt growth since the 1950s.

2.1 Introduction

The rising indebtedness of U.S. households is a much-debated phenomenon. The numbers are eye-catching. Between 1950 and the 2008 financial crisis, American household debt has grown fourfold relative to income. In 2010, the household debt-to-income ratio peaked at close to 120%, up from 30% at the end of World War II. Figure 2.1a shows the trajectory of this secular increase over the past seven decades. The figure shows an increase in debt-to-income ratios over the homeownership boom years after 1950. However, the increase is relatively moderate, and overall the debt-to-income ratio remained fairly stable until the late 1970s. Figure 2.1b confirms that over this period, average debt and income were growing at similar rates. Yet from the early 1980s on, we observe a decoupling of debt and income growth. While average income increased by less than 50%, average debt quadrupled between the 1970s and the financial crisis of 2008, resulting in a massive surge in the debt-to-income ratio.

Rising income inequality is frequently cited as a key driver of households' rising indebtedness. Piketty and Saez (2003) and Kuhn, Schularick, and Steins (2020) show that the share of the richest 10% of households in total household income increased from below 35% to above 45%

between 1950 and 2016. Rajan's (2011) influential book *Fault Lines* popularized the view that rising income inequality and higher indebtedness are two sides of the same coin. The idea is that households with stagnant incomes have increasingly relied on debt to finance consumption – whether out of sheer necessity or to "keep up with the Joneses" at the top of the income distribution, whose incomes were growing nicely (cf. Fligstein, Hastings, and Goldstein 2017). In recent work, Mian, Straub, and Sufi (2021b) discuss how rising income concentration at the top brought about a "savings glut of the rich" that supplied the funds for increased borrowing by non-rich households.

Yet we still know surprisingly little about the borrowers and their financial situation. From the borrowers' perspective, the financial history of the growth of U.S. household debt and its distribution remains largely unwritten. This chapter closes this gap. We study the evolution of household debt over the entire postwar period, asking which households borrowed so much more, and why. Without long-run household-level data for the joint distributions of income, debt, and assets, this task would be impossible. However, we can rely on a new dataset that combines historical waves of the Survey of Consumer Finances (SCF), going back to 1949, with the modern SCF that the Federal Reserve Board has administered since 1983 (see Kuhn, Schularick, and Steins 2020). This long-run "SCF+" makes it possible to follow the evolution of household borrowing across the income distribution over seven decades. Where a panel dimension is needed, we combine information from the cross-sectional SCF+ with data from the Panel Study of Income Dynamics (PSID), which provides information on housing assets and mortgages since 1969.

The data confirm that debt growth was concentrated among households with low income growth from the bottom 90% of the income distribution. In particular, middle-class households, defined here as households between the 50th and 90th percentiles of the income distribution, account for most (55%) of the total increase in household debt since 1950. By contrast, households in the bottom 50% of the income distribution account for a relatively small share of the total debt increase (15%). While their debt-to-income ratio has risen substantially, too, their share in total debt has fallen. Consequently, the American household debt boom of the past decades has first and foremost been a middle-class affair. However, we find that the decoupling of income and debt growth shown in Figure 2.1b is not just driven by one particular subgroup of the population. On the contrary, no matter whether we slice the data by income, education, race, marital status or age, we always find a close comovement until the late 1970s, and a strong divergence thereafter.

Figure 2.1: Debt-to-income ratio and decoupling of debt and income growth

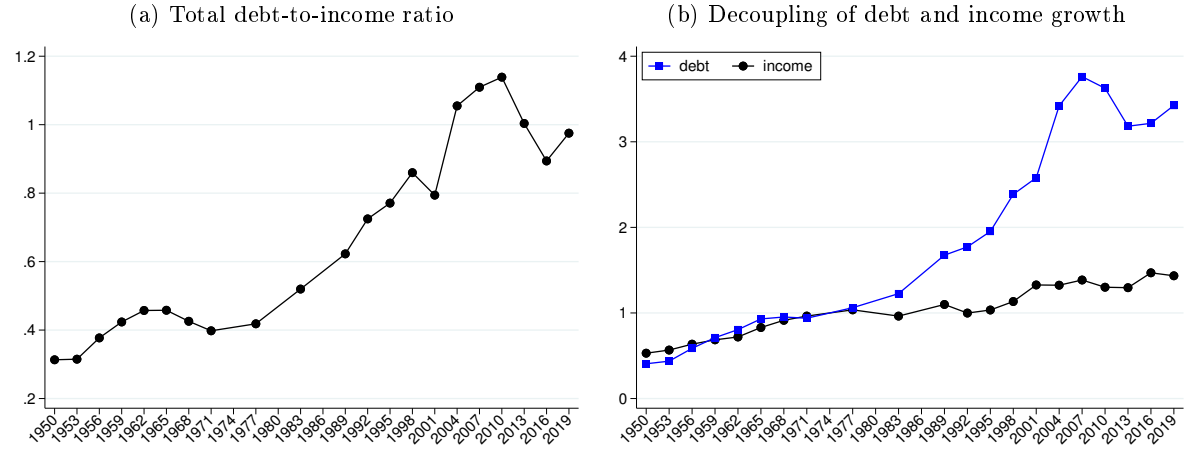

Notes: The left panel shows the ratio of average total household debt to average total household income over time. The right panel shows the growth of average total household debt and income, relative to their average over the 1970s.

Household debt in the U.S. has always been mostly housing debt since the 1950s, with a relatively stable share of around 80% of total household borrowing over time. However, the nature of borrowing has changed substantially since the postwar era. In the 1950s and early 1960s, increases in housing debt-to-income ratios were driven by the fact that more households were taking out mortgages in the course of the postwar homeownership boom. In other words, debt increased at the extensive margin. By contrast, when debt-to-income ratios started to surge again in the 1980s, the increase was due to the fact that conditional on borrowing, households were borrowing more. Previous research has identified home equity withdrawal (HEW) as an important factor for such debt increases at the intensive margin (Greenspan and Kennedy 2008, Klyuev and Mills 2007, Mian and Sufi 2011, Bhutta and Keys 2016).

We add to this literature by comparing the importance of equity extraction to that of other forms of borrowing. Moreover, we study equity extraction along socio-demographic lines. The PSID contains data on housing assets and mortgages that allow us to identify extractors and quantify the aggregate effects of home-equity-based borrowing since the 1980s. Using these data, we decompose the debt increase into additional debt incurred by different groups, in particular extractors, new homeowners, and upgraders moving to larger homes. We find that home equity extraction through refinancing, home equity lines of credit (HELOCs), and second mortgages can account for around 50% of the increase in housing debt since the 1980s, or around 43% of the increase in total household debt. From the early 1980s to the 2008 crisis, equity extraction alone pushed the housing debt-to-income ratio up by more than 30 percentage points.

Without equity extraction, the housing debt-to-income ratio would have stayed at around 50% of income until 2008. Home equity extraction averaged around 1.5% of annual income until the mid-1980s and rose to around 4.5% thereafter. Over a twenty-year period, the cumulative effects of additional equity extraction were substantial. Importantly, we find that home-equity-based borrowing was responsible for a significant fraction of the rise in U.S. household debt even before the extraction boom of the 2000s, which has been studied in prior work. We find that middle-class households are the largest contributors to the extraction boom. While their debt-to-income and extraction-to-income ratios are only slightly larger than those of households from the bottom 50%, their average debt and extraction amounts are much larger than for the bottom 50%, and therefore quantitatively much more important in the aggregate.

Previous research has shown that the propensity to extract equity increases when house prices rise and interest rates fall (Bhutta and Keys 2016, Andersen and Leth-Petersen 2021, Boar, Gorea, and Midrigan forthcoming). In the U.S., most mortgages are fixed-rate, such that a drop in interest rates provides the opportunity to refinance and "lock in" a lower rate. Moreover, rising house prices can lead to higher borrowing through different channels (Berger et al. 2018). The most important ones are a relaxation of collateral and liquidity constraints, and housing wealth effects. A growing literature stresses the importance of the former (e.g., Aladangady 2017, Cloyne et al. 2019, Aruoba, Elul, and Kalemli-Ozcan 2019, Andersen and Leth-Petersen 2021, Ganong and Noel 2020a, Chen, Michaux, and Roussanov 2020). Additionally, Berger et al. (2018) show that wealth effects may lead to increased household consumption in a life-cycle framework à la Modigliani and Brumberg (1954). In their model, a consumption response to housing wealth gains arises as soon as the strict assumptions underlying the model of Sinai and Souleles (2005) are relaxed.

If households would like to liquidate their capital gains without selling their home, they have to borrow against it. The intuition is straightforward: When homeowners make capital gains in the housing market, they are richer than they originally expected. As housing is indivisible, they need to tap into home equity if they wish to smooth consumption over time. In principle, they could also sell their house and buy a less expensive one. However, this would involve substantial transaction, search, and potentially also emotional costs (see Aladangady 2017), and few households do this in practice, as the PSID shows.

Based on matched microdata, Aladangady (2017) estimates a causal effect of house prices on consumption of around 5 cents per dollar increase of home value for homeowners (and zero for renters), corresponding to an overall marginal propensity to consume out of housing wealth of around 3 cents. Similar values are found, e.g., by Guren et al. (2021), Mian, Rao, and Sufi (2013), Andersen and Leth-Petersen (2021) and Christelis et al. (2019). Mian and Sufi (2014) explicitly consider the response of household debt to house price shocks. They exploit regional heterogeneity in the United States and also find substantial effects that can be rationalized in the context of recent models with liquidity-constrained consumers, such as Kaplan and Violante (2014). Guren et al. (2021) report substantial housing wealth effects even since the 1980s, exploiting systematic differences in city-level exposure to regional house price cycles. We find that our measure of equity extraction comoves with house prices at the regional level, and correlates positively with refinancing activity. After house price increases, middle-class households show the largest response in extraction relative to income, whereas interest rate reductions are particularly important for top-income households.

This history of U.S. household debt is compatible with the idea of a savings glut, arising either from global factors (Bernanke 2005) or growing income concentration at the top (Mian, Straub, and Sufi 2021b), which lowered interest rates, loosened borrowing constraints, and increased housing values. Other research has highlighted the role of house price expectations in driving up home values and borrowing (e.g., Kaplan, Mitman, and Violante 2020, Loewenstein 2018, De Stefani 2020). Our analysis does not speak to the initial trigger of this process.[1] Rising income inequality might well have played a role, as argued by Mian, Straub, and Sufi (2021a). The argument we make is that once the house price increase was under way, home-owning middle-class households made large wealth gains leading to home-equity-based borrowing, without a deterioration in net worth. Clearly, the fact that interest rates kept falling despite rising borrowing volumes meshes with the idea of a credit-supply-driven household debt boom. We discuss the importance of further enabling factors in Section 2.4.4, such as the invention of securitization and new mortgage products, financial deregulation and the 1986 tax reform, which maintained interest deductibility for mortgages and thereby created incentives to switch to home-equity-based products. Story (2008) describes how banks heavily advertised these new products in the 1980s with catchphrases such as "Now, when the value of your home goes up, you can take credit for it."

Equity extraction, which typically happens in the middle of the life cycle, leads to pronounced differences in debt profiles as opposed to a situation when a household takes out a mortgage at the beginning of the life cycle and then keeps repaying it without re-augmenting the debt balance. Our long-run SCF+ data allow us to study the life-cycle debt profiles of different age cohorts for the first time. While debt-to-income profiles used to fall over the life-cycle, we observe a shifting and turning of the profiles from cohort to cohort. The turning point coincides with the onset of the 1980s debt boom. By consequence, households enter retirement with substantially elevated debt levels (see also Lusardi, Mitchell, and Oggero 2018, 2020). While the prewar generations typically approached retirement with modest debt ratios of around 30% to 60% of income, households in the first baby boomer cohort (1945-1954) had debt ratios of almost 120% on average at the same age. The comprehensive balance sheet information in the SCF+ allows us to further examine whether this implies that households became poorer relative to previous cohorts. We find that net wealth-to-income profiles over the life cycle are actually similar across cohorts.[2] This is possible because increased borrowing went hand in hand with increases in asset values, in particular through strong house and stock price growth.

Increases in housing debt are, by definition, a phenomenon that is particular to homeowners.

[1] For a detailed overview on the drivers of house prices and their connection to credit markets, see Duca, Muellbauer, and Murphy (forthcoming).

[2] When we refer to wealth in the chapter, we mean net wealth throughout, i.e. assets net of debt.

Previous work has discussed the size and persistence of wealth gaps between black and white households, for which the black-white homeownership gap plays an important role (Kuhn, Schularick, and Steins 2020, Aliprantis and Carroll 2019, Bartscher et al. 2021). Our long-run data show that black and white households have had similar housing-to-income and debt-to-income ratios at the intensive margins over the past seven decades. However, the gap in homeownership and mortgage borrowing at the extensive margin has persisted. Moreover, the average value of black households' homes has remained lower than for white households, such that they have less housing wealth to borrow against. We find that black households have contributed relatively little to the debt increase over the past 70 decades, and particularly little to the extraction boom in the post-1980s. In summary, the extraction boom was therefore mainly driven by white, middle-aged middle-class households.

Literature: The analysis of household balance sheets and their importance for the business cycle and financial stability has become an active research field for macroeconomists (Mian and Sufi 2014, 2017, Zinman 2015, Jordà, Schularick, and Taylor 2013, Adelino, Schoar, and Severino (2018), Albanesi, De Giorgi, and Nosal 2017). A large empirical and theoretical literature has examined wealth and collateral effects due to house price increases and their consequences for household borrowing and consumption.[3] Empirical trends in household indebtedness have been discussed in Dynan and Kohn (2007) and Wolff (2010). Dynan and Kohn (2007) provide an early analysis of the 1990s debt boom and discuss potential sources for the rise in indebtedness of U.S. households. They likewise point to the important role of mortgage debt and document its comovement with house prices. Wolff (2010) provides a broader perspective on the change in household finances, which emphasizes the rise in middle-class debt since 1983.

Regarding house prices and credit conditions, several important papers have traced house price increases to regulatory changes since the 1980s (e.g., Hoffmann and Stewen 2019, Favara and Imbs 2015, Di Maggio and Kermani 2017). Recent research has also emphasized the link between rising inequality and household borrowing (De Stefani 2018, Mian, Straub, and Sufi 2021b). In their influential work, Mian and Sufi (2009, 2011) argue that household borrowing in low-income regions of the United States grew particularly strongly before the 2008 crisis, followed by severe output and employment losses. In a theoretical model, Kumhof, Rancière, and Winant (2015) show that higher savings of the rich may lead to a decline in interest rates, which leads to higher borrowing by low- and middle-income households and higher financial fragility. However, Coibion et al. (2020) find that low-income households face higher borrowing costs and reduced access to credit as inequality increases. Adelino, Schoar, and Severino (2016) and Albanesi, De Giorgi, and Nosal (2017) provide complementary evidence on the debt boom during the 2000s and highlight the important role of the middle class for the debt boom during these years. Adelino, Schoar, and Severino (2016) also conclude that the growth of middle-class debt played an important role. Similarly, Foote, Loewenstein, and Willen (2016) study debt growth in the early 2000s across the income distribution and discuss the implications for theoretical models of the debt boom.

The structure of the chapter is as follows. We first introduce and discuss the data. Second, we show that mortgage borrowing of households between the 50th and 90th percentiles of the income distribution accounts for the lion's share of the debt increase. Third, we show that equity extraction in response to higher housing wealth played a central role in the aggregate debt increase, and discuss characteristics, drivers and consequences of the extraction boom. Fourth, we examine the implications of changing debt patterns for life-cycle debt profiles. Finally, we study racial inequality in borrowing.

[3] Further examples, beyond those mentioned above, include Iacoviello (2005), Hurst and Stafford (2004), Calomiris, Longhofer, and Miles (2013), Guren et al. (2021), Campbell and Cocco (2007), and Kaplan, Mitman, and Violante (2020).

2.2 Data

The main data source for this chapter is the "SCF+", which allows us to track the financial situation of U.S. households since World War II by combining historical waves of the Survey of Consumer Finances (SCF) going back to 1949 with the modern waves available since 1983 (Kuhn, Schularick, and Steins, 2020). A key strength of the SCF+ data is that they provide joint information on income, debt, and asset holdings at the household level together with information on household demographics. Kuhn, Schularick, and Steins (2020) give a detailed description of the SCF+.

Our second main data source is the Panel Study of Income Dynamics (PSID). The PSID started in 1968 as a panel tracking U.S. households over time. Over the first 30 years, the PSID however provides only limited information on household asset and debt holdings. Yet though comprehensive balance sheet information is not available before 1984, the PSID still provides information on mortgage debt and housing values over almost the entire covered time period. As housing debt has played a key role for the U.S. household debt increase, we rely on this information for our analysis and exploit the PSID's panel structure to complement the cross-sectional information from the SCF+.

We further use long-run data on the consumer price index (CPI) from the *Macrohistory Database* (Jordà, Schularick, and Taylor 2017) to deflate nominal variables. If not explicitly stated otherwise, all presented results are in real terms, converted to 2019 dollars using the CPI.

2.2.1 Household debt in the SCF+

The SCF is a key resource for research on household finances. The modern surveys have been conducted every three years since 1983 by the Federal Reserve Board (see Bricker et al. 2017 for more details).[4] The comprehensiveness and quality of the SCF explain its popularity among researchers (see Kuhn and Rıos-Rull 2016 and the references therein). The historical predecessor surveys were carried out annually between 1947 and 1971 and then again in 1977. We follow Kuhn, Schularick, and Steins (2020) and use data since 1949, which is the first year in which all relevant variables are available, and pool the early waves into three-year bins to increase sample sizes. The SCF+ data are weighted with post-stratified cross-sectional weights that ensure representativeness along several socioeconomic characteristics, in particular race, education, age, and homeownership.

Of particular interest for our study is the coverage of household debt and its components, which we aggregate into housing and non-housing debt. For housing debt, we focus on debt for owner-occupied housing. This includes mortgages and home equity lines of credit. We treat investment in non-owner-occupied housing like business investment and use the net position to calculate wealth.[5] Non-housing debt includes car loans, education loans, and loans for the purchase of other consumer durables. Data on credit card balances become available after 1970 with the introduction and proliferation of credit cards. Note that the appearance of new financial products like credit cards or home equity lines of credit does not impair the construction of consistent data over time. Implicitly, these products are counted as zero for years before their appearance.

[4] The 1986 survey was designed as a panel survey to the 1983 survey but suffers from sample attrition and is therefore not included in our dataset.

[5] Several papers have stressed the importance of real estate investors for the debt boom prior to 2007 (Haughwout et al. 2011, Bhutta 2015, Mian and Sufi 2021, Albanesi, De Giorgi, and Nosal 2017, DeFusco, Nathanson, and Zwick 2017, De Stefani 2020). Real estate investors are defined as borrowers with multiple first-lien mortgages. While they accounted for a disproportionately large share of mortgage growth before 2007 compared to their relatively small population share, mortgage debt on the principal residence is on average eight times larger than that on other real estate (see Appendix Figure 2.A.1).

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

For most of the analysis, we abstain from any sample selection. One exception is the decomposition of changes in debt-to-income ratios in Section 2.3.3. Here we use household-level ratios and drop observations with extreme debt-to-income ratios. Our analysis in this part explicitly relies on individual ratios. Otherwise, we use ratios of averages instead of averages of ratios because of their greater robustness to outliers.

As discussed in Kuhn, Schularick, and Steins (2020), aggregated household surveys are not always easy to reconcile with macroeconomic data sources like the National Income and Product Accounts (NIPA) and the Financial Accounts (FA). Measurement concepts can differ, such that even high-quality microdata may not match aggregate data one-to-one. To judge the reliability of the SCF+ data, we compare the trends in average income and household debt in the SCF+ to data from the NIPA and FA in Appendix 2.A.2. We find that after accounting for conceptual measurement differences affecting levels, the aggregated microdata match macroeconomic trends closely so that they can be used to study underlying distributional changes over time. We find the alignment to be particularly close for house values and housing debt.

2.2.2 Panel data from the PSID

The key strength of the SCF+ is that it allows us to study the joint distribution of income, debt, and asset holdings at the household level over seven decades. However, the SCF+ data are repeated cross sections and thus do not allow us to track individual households over time. To explore how individual households change debt holdings over time, we complement the SCF+ with panel data from the PSID. Conceptually, the SCF+ reports data at the household level, whereas the PSID collects data at the family level. To account for these differences, we aggregate PSID families living together into one household for better comparability (cf. Pfeffer et al. 2016). Additional details are given in Appendix 2.B.

Following Kaplan, Violante, and Weidner (2014), we only use data from the *Survey Research Center (SRC)* sample. Post-stratified cross-sectional survey weights are only available for the waves between 1997 and 2003. We use the longitudinal PSID family weights and post-stratify them to match the same Census variables that are targeted in the post-stratification of the SCF+ waves. We verified that all reported results are similar when using the unweighted PSID data or the original longitudinal PSID weights without post-stratification. Appendix Figure 2.B.1 compares the PSID data and SCF+ data for housing assets, housing debt, and income. Overall, the two datasets align very well.

2.3 The American household debt boom, 1950-2019

We start by using the SCF+ to track the growth and distribution of household debt and its relation to income dynamics over the past seven decades. Which households have borrowed so much more, and for what reasons? We will first look at the distribution of debt among income groups over time and then establish that the middle class accounts for the largest part of both outstanding debt and debt growth. In a second step, we will decompose the overall debt increase into changes at the intensive and extensive margin of different debt components. Finally, we will use the identified patterns to classify the debt boom since World War II into four historical phases.

2.3.1 The distribution of household debt

How is household debt distributed among rich and poor households, and how has this distribution changed over time? To address these questions, we stratify households by income. Following standard practices in the literature, we divide the population into three groups according to their position in the income distribution (see Piketty and Saez 2003, Saez and Zucman 2016, and Alvaredo et al. 2018). The first group are households in the bottom 50% of the income distribution, and the second covers households between the 50th and 90th percentiles. We refer to this group as the "middle class" throughout the chapter. The third group consists of the top 10% of the income distribution. We will only occasionally talk about the top 1% to illustrate dynamics at the very top. Even very rich households owe considerable amounts of debt despite their high net wealth (with tax considerations likely playing an important role). Yet they are not central for trends in aggregate debt (but might have played an important role for the supply of funds, see Mian, Straub, and Sufi 2021b).

Before we study the evolution of debt shares and debt-to-income ratios of these different groups over time, it is important to recognize that the SCF+ is a repeated cross section. This means that households can move between income groups over time. Our groups are reasonably large so that inter-group mobility can be expected to be low, but we can use panel data from the PSID to test this assumption, along the lines of Díaz-Giménez, Glover, and Ríos-Rull (2011). The results are reassuring. The PSID shows that around 84% of households in the bottom 50% were already in this group two years ago (Table 2.A.1). The numbers for the 50%-90% and top 10% are 75% and 66%, respectively. When we extend the intervals to six years, the share of households who are in the same group six years later is still 77% for the bottom half, 68% for the middle class, and 53% for the top 10%. Moreover, households that change income groups tend to remain close to the "border" with the previous group. For instance, among households who changed into the middle-class group, 64% were no more than two deciles away from this group two years earlier. On average, households remain in the same income group for 77% of the periods in which we observe their income.[6]

Figure 2.2: Debt shares by income group

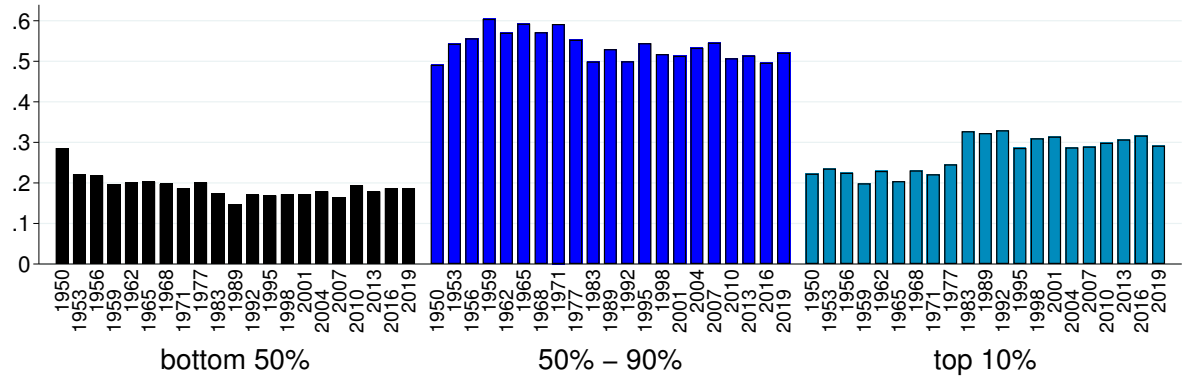

Notes: The figure shows shares in total debt for the different income groups over time.

Figure 2.2 shows the share of total debt owed by the three income groups. Debt shares have been rather stable over time. Over the entire postwar period, middle-class households have always accounted for the largest share of total outstanding debt, on average about 50% to 60%. Low-income households in the bottom half make up another 20%. The debt share of the top 10%

[6] Appendix Figure 2.A.4 presents additional evidence for income group stability. It shows income and housing debt, two key variables for our analyses, for households aged 25 to 55. We examine if the trends in debt look different depending on whether we sort households using their contemporaneous income, or the initial income at the beginning of a decade. The trends look very similar.

fluctuated around 20% before the 1980s and then increased to around 30%. Consequently, the upper half of the income distribution has always accounted for about 80% of total household debt outstanding.

It follows from the relative stability of the debt shares over the past seven decades that the middle class also played a dominant role in the growth of debt. Appendix Figure 2.A.5 confirms this visually. From 1950 to 2007, middle-class households accounted for 55% of the total debt increase, whereas households from the bottom 50% of the income distribution contributed only 15%, even less than the top 10% with almost 30%. This insight is important in itself. We see that 85% of the increase in U.S. household debt occurred within the upper 50% of the income distribution. The explanation for soaring household debt in the United States lies in the borrowing behavior of these income groups, and in particular of middle-class households (see also Adelino, Schoar, and Severino 2018).

A more comprehensive picture of the distributional dimension of the American household debt boom emerges from Figure 2.3. For different survey waves, the figure shows the evolution of debt-to-income ratios across the entire income distribution. The left panel shows total household debt relative to income, and the right panel shows housing debt-to-income ratios. Debt-to-income ratios were relatively constant in 1950, at less than 50% across most of the income spectrum. By 1983, debt-to-income ratios had increased somewhat, but were not far off their levels in the 1950s. Since then, indebtedness has risen strongly across all income groups, but soaring debt ratios of the middle class stand out. For households between the 50th and 90th percentiles, debt-to-income ratios approximately tripled between the early 1980s and 2007, driven by mortgage debt.[7]

Figure 2.4a shows the time evolution of debt-to-income ratios along the income distribution. For all income groups, debt-to-income ratios were fairly stable before 1977. The postwar increase is most pronounced for the middle class, and absent for the bottom 50%. After 1977, debt-to-income ratios surge for all households. However, the increase is much higher for households from the bottom 90%, who experienced only moderate or stagnant income growth over this period (Figure 2.4b). For both income groups within the bottom 90%, debt-to-income ratios rose from

Figure 2.3: Debt along the income distribution

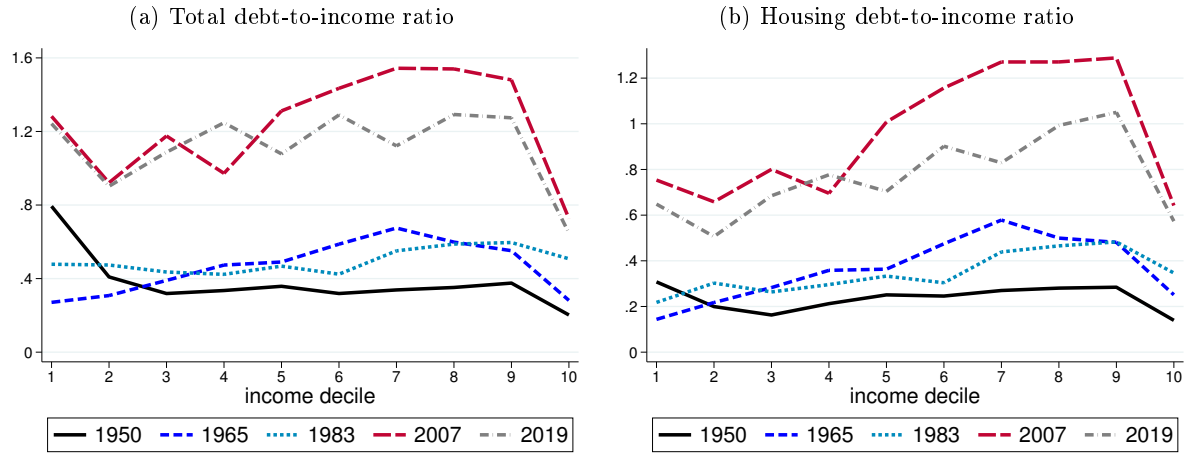

Notes: The graph shows the evolution of average total (left) and housing (right) debt-to-income ratios by deciles of the aggregate income distribution for the SCF+ waves 1950, 1965, 1983, 2007, and 2019. We excluded households with total income below 10% of the annual wage of a household with a single earner receiving the contemporaneous minimum wage.

[7] In Appendix Figure 2.A.9, we show that leverage has also increased most strongly for households from the middle of the income distribution.

around 40% in the early 1950s to close to 140% by 2007. For the top 10%, the increase is much more muted, even though they account for a higher share in total debt today compared to the 1950s. This is because their incomes have risen much more than for the bottom 90%. Appendix Figure 2.A.6 additionally shows debt-to-income ratios of the top 1%, compared to the bottom 90%. The chart underscores the divergent debt trajectories at the top and in the rest of the economy. It further illustrates that debt-to-income ratios at the top and bottom evolved in tandem until the 1980s and then sharply diverged as income concentration at the top increased.[8]

Figure 2.4: Debt and income along the income distribution

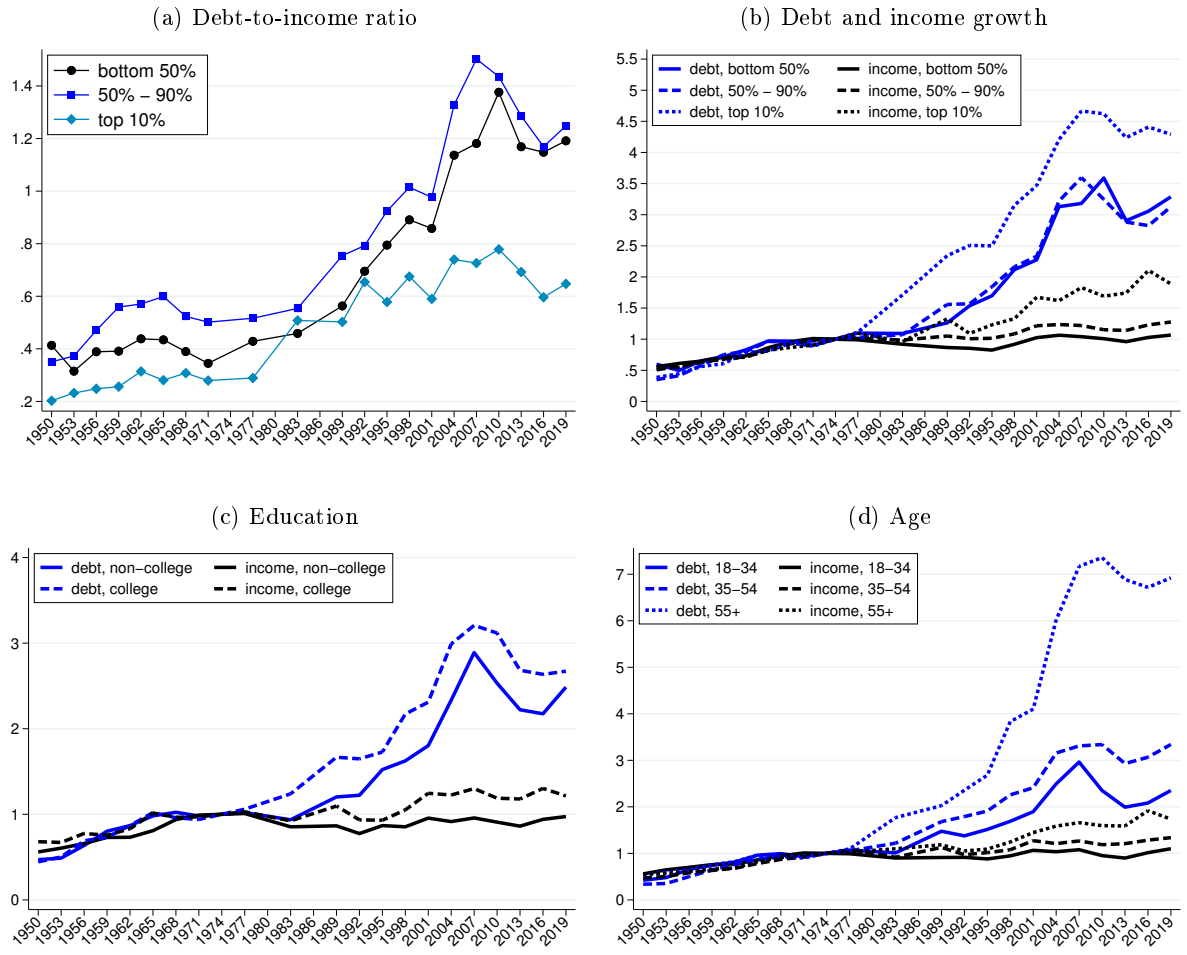

Notes: Figure 2.4a shows average debt-to-income ratios by income group. Figures 2.4b to 2.4d show the growth of average total housing debt and income by income group, education and age, relative to 1970s averages.

Figure 2.4b decomposes the debt and income time series from Figure 2.1b along the income distribution. It shows that the close comovement of debt and income growth until the late 1970s and subsequent divergence which we observed for the aggregate is actually not driven by one specific income group, but rather a phenomenon present along the whole income distribution. Given that we have shown that income groups are fairly stable over time, and households outside the top 10% had only meager income growth, it is unlikely that these households were anticipating a stronger future income growth against which they wanted to borrow. Alternatively, one may consider education as a proxy for lifetime income. Yet both college and non-college households substantially increased their borrowing from the 1980s on (Figure 2.4c). Moreover, remaining

[8] Appendix Figure 2.A.8 shows that the debt-to-asset ratio has equally stayed largely flat for high-income households. Both debt-to-income and debt-to-asset ratios have increased most for the middle class.

future lifetime income declines with age, yet we find that older households also borrowed more, and the debt increase was actually particularly strong for them (Figure 2.4d). In Appendix Figure 2.A.7, we slice the data along two additional important socioeconomic dimensions, namely race and marital status. The pattern is always the same: debt and income grew in tandem until the late 1970s, and show a pronounced divergence thereafter.

In order to get a better understanding of the underlying debt dynamics, we will first of all exploit our microdata to decompose the increase of debt-to-income ratios into the contributions of housing and personal debt, as well as changes at the extensive margin (How many households were borrowing?) and the intensive margin (Conditional on borrowing, how much did they borrow?).

2.3.2 The composition of household debt

Figure 2.3 illustrates the important role that housing debt plays for debt trends of households in the upper half of the income distribution. Adding information on the number of households with outstanding debt and the type of debt, we decompose the debt increase into its extensive and intensive margins. In other words, we answer to what extent the total number of indebted households has increased and to what extent indebted households have taken on larger amounts of debt. Additionally, we calculate the extensive- and intensive-margin effects separately for housing and non-housing debt.

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

$$d_{i,t} - d_{i,t-1} = \underbrace{(s_{i,t}^{H^+} - s_{i,t-1}^{H^+})\, d_{i,t-1}^{H^+}}_{\Delta \text{ extensive housing}} + \underbrace{s_{i,t}^{H^+}\, (d_{i,t}^{H^+} - d_{i,t-1}^{H^+})}_{\Delta \text{ intensive housing}} + \underbrace{(s_{i,t}^{N^+} - s_{i,t-1}^{N^+})\, d_{i,t-1}^{N^+}}_{\Delta \text{ extensive non-housing}} + \underbrace{s_{i,t}^{N^+}\, (d_{i,t}^{N^+} - d_{i,t-1}^{N^+})}_{\Delta \text{ intensive non-housing}}. \quad (2.1)$$

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, it captures by how much household indebtedness would have risen if only the share of households with 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, that is, changes in household indebtedness due to an increase in the level of debt, d_t^H, with the extensive margin of housing debt, $s_{i,t}^{H}$, constant at the level of period t and all non-housing debt components at the level of period $t-1$. The third and fourth parts are the respective effects for non-housing debt.

Table 2.1 shows the extensive- and intensive-margin effects of the increase in the average debt-to-income ratio between 1950 and 2019. Overall, we find that the intensive margin of housing debt accounts for 32 percentage points of the 73-percentage-point increase in the average household debt-to-income ratio. Another 20 percentage points are due to the extensive margin of housing debt. The remaining 21 percentage points are due to non-housing debt. This confirms that mortgage lending has played a dominant role relative to non-housing debt (e.g., credit cards or student loans) in the debt boom.

Figure 2.5 shows the intensive and extensive margins of indebtedness over time for both types of debt. The extensive margin in the left panel captures the share of households with positive

Table 2.1: Decomposition of the increase in debt-to-income ratios from 1950 to 2019

housing debt	intensive margin	31.9
	extensive margin	20.2
non-housing debt	intensive margin	13.0
	extensive margin	7.7
total		72.8

Notes: The table shows the percentage point change in the average debt-to-income ratio between 1950 and 2019, decomposed into extensive and intensive margin effects for housing and non-housing debt according to equation (2.1).

(non-)housing debt balances. A closer look at Figure 2.5 reveals that the extensive margin of housing debt closely tracks changes in the homeownership rate (dashed line). The intensive margin in the right panel is represented by the debt-to-income ratio for households with positive levels of (non-)housing debt. Overall, more households have personal debt than housing debt. In particular, the rollout of credit cards in the 1970s led to a substantial increase in the share of households with personal debt (Appendix Figure 2.A.10). Yet the amount that households owe is small compared to the average amount owed on housing debt, as the right-hand side shows.

2.3.3 Four phases of the postwar debt boom

From Figure 2.5, we identify four different phases of the postwar debt increase, which we will explore in more detail. To do so, Figure 2.6a decomposes the change in debt-to-income ratios into extensive and intensive margins, stratified by income. The figure shows two boom phases (1950-1965 and 1983-2007), followed by two periods of deleveraging (1965-1983 and 2007-2019).

Figure 2.5: Extensive and intensive margins of debt-to-income ratios

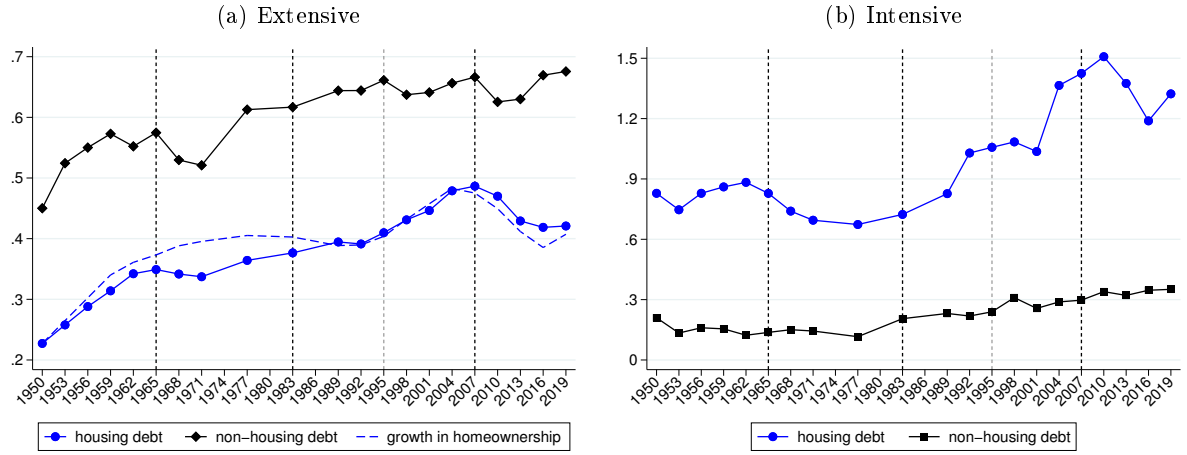

Notes: The left panel shows the share of households with positive housing debt (blue line with dots) and positive non-housing debt (black line with squares). Moreover, it shows the growth rate of the homeownership rate since 1950, normalized to extensive margin housing debt in 1950 for comparison. The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

Figure 2.6b shows a similar picture for loan-to-value ratios. There are substantial differences between the four periods.

The postwar homeownership boom, 1950-1965: The first period is characterized by the rise in homeownership after World War II until the mid-1960s, aided by public policies to increase homeownership (Fernández 2013; Fetter 2014). Housing debt-to-income ratios approximately doubled in this period, yet from a low baseline level of around 0.2 (Figure 2.A.2), mainly driven by the extensive margin of housing debt and the upper half of the income distribution. Likewise, average loan-to-value ratios increased, driven predominantly by the extensive margin and some intensive-margin effects in the middle class.

Stability, 1965-1983: The second period spans the years from roughly 1965 to 1983. It is characterized by almost stable debt-to-income ratios and a slight decline in intensive-margin housing debt of the middle class, with marginal increases at the extensive margin. At the top and in the bottom 50%, non-housing debt (car loans and credit cards) made a small but positive contribution to debt ratios. Loan-to-value ratios decreased across income groups.

Figure 2.6: Decomposition of changes in debt-to-income and loan-to-value ratios by income

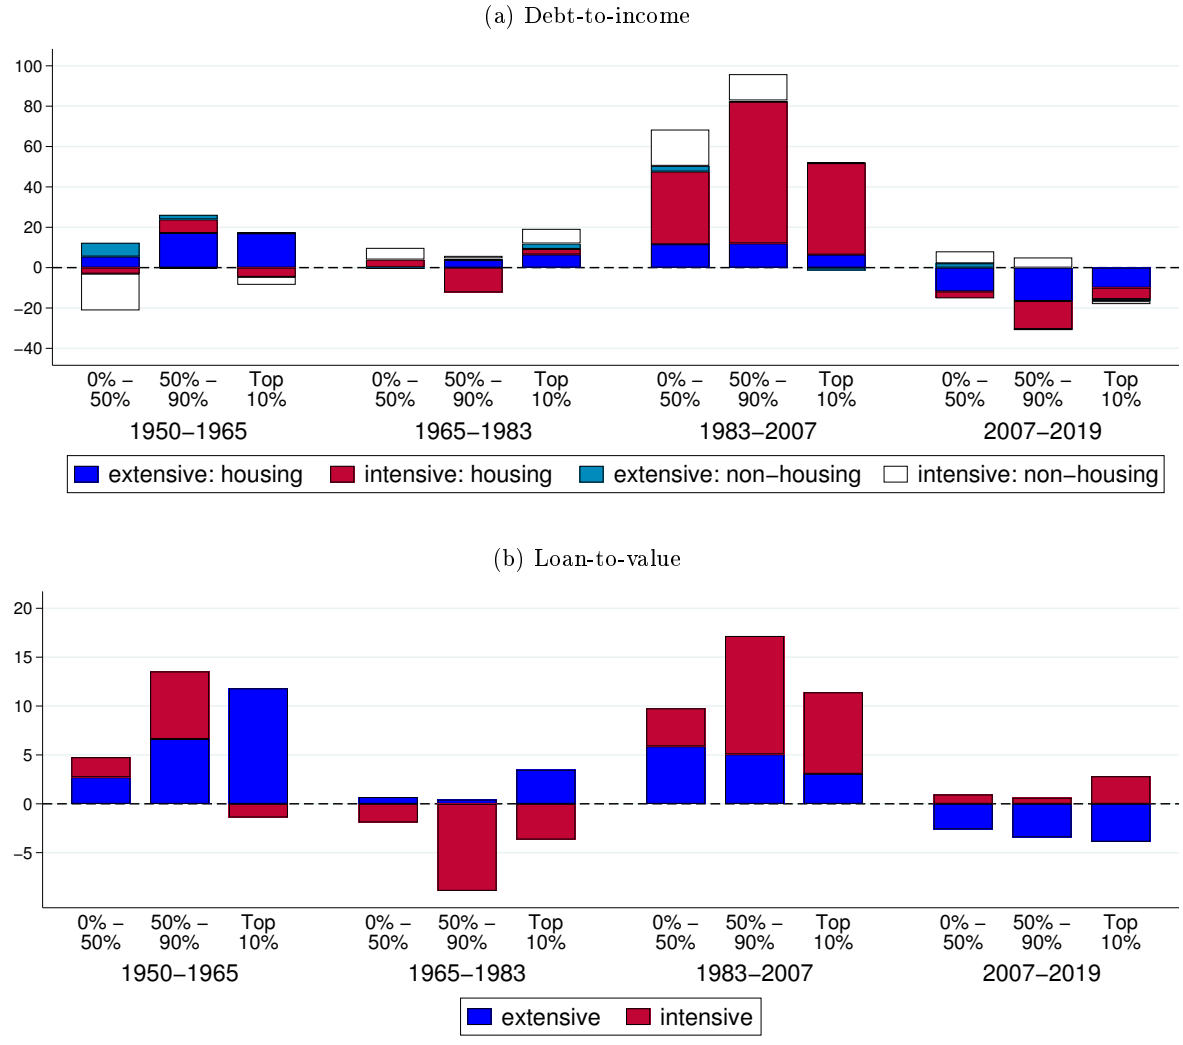

Notes: The upper panel shows the decomposition into extensive and intensive margin effects from equation (2.1) over the four phases of the debt boom, stratified by income. The lower panel shows an analogous decomposition of the loan-to-value ratio. Observations with debt-to-income ratios above 50 in absolute value were excluded.

The second debt boom, 1983-2007: Starting in the 1980s, the United States entered a second debt boom, which came to an end with the crisis in 2008. Debt-to-income ratios more than doubled within the 25 years between 1983 and 2007, from roughly 50% of income to above 110%. This time, the increase was mainly driven by higher intensive margins of housing debt, as Figure 2.6a shows. Overall, the extensive margin made a relatively small contribution, but the effect was larger in the 2000s, as we will see below. The boom was fueled by households from all parts of the income distribution, but the intensive-margin effect of the middle class (50%-90%) stands out, for both debt-to-income and loan-to-value ratios.

Crisis and deleveraging, 2007-2019: The final period covers the decade after the crisis and is marked by deleveraging. Overall, the debt-to-income ratio fell by about 30 percentage points. For the bottom 50%, non-housing debt, mainly education loans, showed positive growth. The middle class and the top 10% deleveraged at both margins but chiefly at the extensive margin. Homeownership rates have fallen across all income groups. The decline in LTVs was also mainly driven by a decline at the extensive margin.

Recently, the consequences of strongly rising student debt have received increased attention (see, for example, Looney and Yannelis 2015, Avery and Turner 2012). Rising student debt is visible in Figure 2.6a as a part of the intensive margin of non-housing debt. Since 1983, we find a significant contribution from this component, especially in the lower half of the income distribution. These increasing debt levels might shape the financial decision making of young generations of American households in the future. However, Figure 2.6a also shows that from a macroeconomic perspective, the contribution of student debt is much smaller than the increase in housing debt over the same period (see also Appendix Figure 2.A.10).

Figure 2.7: Housing debt during the second debt boom

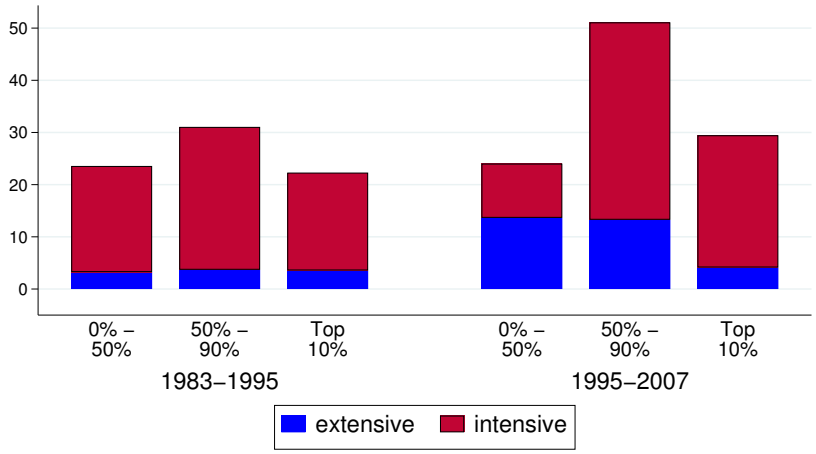

Notes: The graph repeats the analysis from Figure 2.6a for housing debt-to-income only, zooming in on the second debt boom.

Figure 2.7 zooms in on the post-1980 debt boom and focuses on housing debt. In its first phase, from 1983 to 1995, the housing debt increase was similar for all income groups, and intensive-margin housing debt played the central role. In the second phase, from 1995 to 2007, the quality of the debt boom changed considerably. The middle-class housing debt-to-income ratio grew twice as much as that of the other income groups. The significant increase in the housing debt ratio of the top 10% is also noteworthy, as it outpaced the increase in the bottom half of the income distribution. In the middle and lower half of the distribution, the extensive margin also made a substantial contribution to rising housing debt levels after 1995. This reflects the homeownership boom of the 2000s, partly driven by lending to households from the lower half of the distribution. Over the entire boom from 1983 to 2007, the middle-class housing debt-to-income ratio increased

by 82 percentage points, predominantly because of higher intensive-margin indebtedness.

2.4 The role of home equity withdrawal

We have seen that in contrast to the first debt boom, which was driven by the postwar homeownership boom and a coinciding increase in the extensive margin of debt, the second debt boom was also driven by a surge in housing debt, but this time at the intensive margin. Furthermore, we have seen that households in all parts of the population have increased their borrowing. A natural question is whether this surge in housing debt led to a reduction in home equity, i.e. housing net of debt. Figure 2.8 shows the change in average home equity relative to 1971, computed both from the SCF+ and the PSID.[9] We can see that home equity has increased substantially, despite the surge in debt. While it was relatively stable over the 1960s, home equity increased rapidly over the 1970s, and even more rapidly between the mid-1990s and the 2008 financial crisis.

Figure 2.8: Change in home equity, SCF+ and PSID

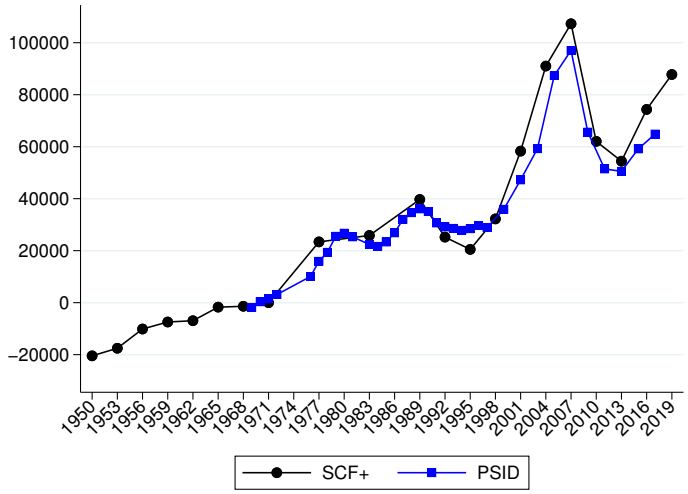

Notes: The graph shows the change in average home equity since 1971 from the SCF+ and PSID.

In order to better understand the contributors to the increase in home equity since the 1970s, we display counterfactual scenarios in Figure 2.9. The blue lines show the actual trajectory of average home equity in both panels. Figure 2.9a fixes the extensive margins of housing assets, housing debt, or both at their 1971 levels. We can see that extensive-margin changes hardly affect the evolution of home equity over the 1970s, and only have modest effects over the 1980s until the mid-1990s, when the homeownership rate started to rise. In Figure 2.9b, we conduct an analogous exercise for the intensive margin. Holding the intensive margin of debt constant has no effect before 1983. By contrast, if we fix the intensive margin of housing, we do not find any increase in home equity between the early 1970s and 1983. In other words, the increase in average home equity over the 1970s is entirely driven by the intensive margin of housing. The exactness of the timing is difficult to assess based on the SCF+ due to the triennial frequency of observation. However, both the pattern and timing of the described effects are strikingly similar in the PSID, as shown in Appendix Figure 2.B.3.

The PSID includes information on the state and Census region in which a household resides. Unlike at the state level, house price indices for the four Census regions exist for the whole period covered by the PSID. Figure 2.10a shows the change in home equity by region, and Figure shows

[9] As the SCF+ pools observations from 1969-1971, we use the same reference period in the PSID.

Figure 2.9: Decomposition of home equity, SCF+

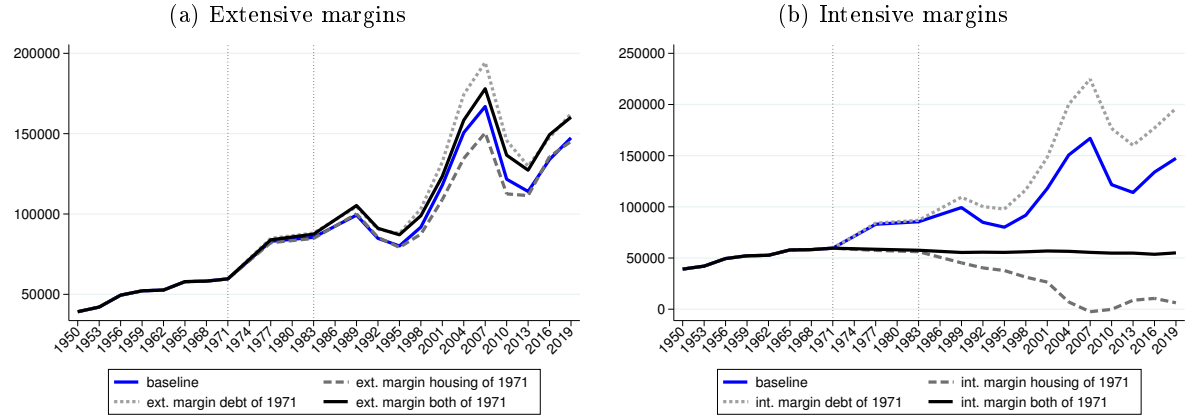

Notes: The graphs show average home equity from the SCF (blue), along with counterfactuals (gray). The counterfactuals in the left panel hold the extensive margins of housing, housing debt, or both at their 1971 values. The right panel does the same for the intensive margins. The dotted vertical lines indicate the years 1971 and 1983.

2.10b house price growth by Census region. We can see that home equity increased relatively little in the North Central and South regions until the 1990s. By contrast, it already surged in the mid-1970s in the West, and the early 1980s in the North East. Looking at Figure 2.10b, we can see that real house prices also increased substantially in the West from the mid-1970s on and in the North East from the early 1980s on, whereas the price changes were moderate in the other two regions. This evidence is in line with the work of Guren et al. (2021), who argue that exposure to regional house prices cycles has led to housing wealth effects already before the boom phase starting in the late 1990s.

When house prices rise, new homeowners have to take out larger mortgages to finance their purchase for a given LTV ratio, and if house prices rise faster than incomes, they will also have to rely on higher LTV ratios ceteris paribus. By contrast, existing owners experience capital gains, which may induce them to trade or extract home equity (cf. e.g., Mian and Sufi 2011,

Figure 2.10: House prices and home equity by Census region

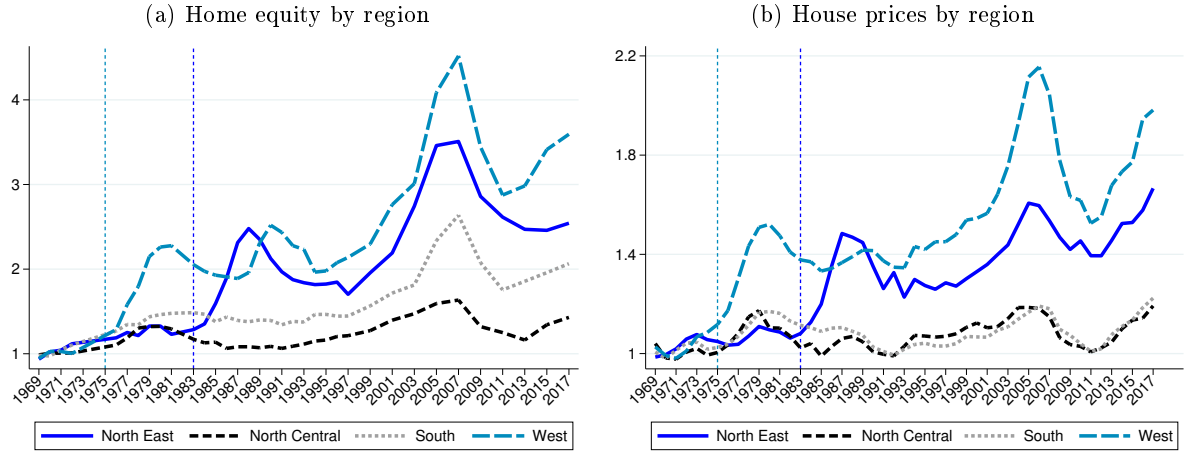

Notes: The left panel shows home equity growth from the PSID by Census region. The right panel shows real regional house price growth based on the U.S. Census Bureau's indices for new single-family houses sold, including lot value. The dashed vertical lines indicate the years 1975 and 1983.

Bhutta and Keys 2016). Appendix Figure 2.B.4 shows that the average homeowner in the regions North East, North Central and South had capital gains of between $15,000 and $20,000 from real house price growth over the 1970s, and the average homeowner from the West even $60,000. Moreover, the high inflation rates of the 1970s substantially reduced real debt burdens of existing homeowners (cf. Doepke and Schneider 2006). Using the PSID to calculate how much debt was "inflated away" from one year to the next, we estimate that the average real debt burden of homeowners was reduced by around $2,000 per year during the 1970s compared to a world without inflation, all else being equal.

What is more, mortgage interest rates fell substantially from the early 1980s on (see Appendix Figure 2.C.7), providing incentives to refinance. Previous work has shown that many households who refinance also cash out equity on the way, especially when house prices are rising (Bhutta and Keys 2016, Andersen and Leth-Petersen 2021). Moreover, new mortgage products like HELOCs were invented in the mid-1980s (Maki 2001), second mortgages were gaining popularity (Elia 1981; see also Appendix Figure 2.A.11b) and securitization allowed banks to expand their mortgage lending without the need to simultaneously increase deposits (Doepke and Schneider 2006). These financial innovations made it easier for households to tap into the home equity they had accumulated.

In the following, we will quantify the contribution of home-equity-based borrowing to the post-1980s debt increase, as well as the role of the middle class in this context. Subsequently, we will discuss the historical context in more detail, as well as the macroeconomic consequences of extraction at the household level.

2.4.1 The post-1980s equity extraction boom

To quantify the respective contributions of home-equity-based and other forms of borrowing to the overall housing debt increase, we complement the SCF+ data with panel data from the PSID. As discussed in Section 2.2.2, we use the SRC sample, which tracks the original households from the first PSID wave in 1968 over time, as well as the new households formed by former members of these households (e.g., adult children moving out). We will focus the analysis on housing debt as the largest component of household debt, as discussed in Section 2.3. Information on net wealth is available in the PSID only since 1984. However, information on housing is available since 1968, and on mortgage balances since 1969 (with the exceptions of 1973-1975 and 1982). The initial sample size was about 2,930 households in 1968 and increased to 5,601 by 2017. The PSID was conducted at an annual frequency until 1997 and every two years thereafter. To ensure consistency over time, we discard all even years from the sample.[10]

Several approaches have been made to quantify the importance of home equity extraction. Bhutta and Keys (2016) use the New York Federal Reserve Bank Consumer Credit Panel to calculate the amount of home equity withdrawal (HEW) based on home equity loans, HELOCs, second mortgages, and cash-out refinancings. According to their calculations, households on average extracted $40,000 between 1999 and 2010, their period of study. However, their data only cover a relatively short and recent time period, and do not include individual- or household-level demographics (although they can be linked to Census-tract-level information). Greenspan and Kennedy (2008) and Klyuev and Mills (2007) instead use aggregate data in order to compute the amount of home equity extraction over time. While this allows to consider a longer time span, the resulting measures are coarse, and aggregate data do not allow for demographic decompositions by construction. Moreover, none of these studies compare the relative importance of HEW and

[10] The only information we use from the even years is whether a household has moved over the last year. We use this information to construct a measure of whether the household has moved during the last two years, consistent with the data from the post-1997 waves.

other forms of mortgage borrowing. A more detailed overview of and comparison to the previous literature on HEW is given in Appendix 2.C.

To isolate the contribution of home equity withdrawal, we need to separate it from other channels that affect debt levels over time: transitions from renting to ownership and vice versa, upgrading to bigger or better homes, and downgrading. We employ the following definitions:

New owners are defined as households who (1) bought a house and (2) were not homeowners in the previous survey.

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

Downgraders are the mirror image of upgraders.[11] As we are interested in their contribution to the overall debt increase, we will focus on upgraders (downgraders) who increased (decreased) their mortgage in the following.

Extractors are defined following the approach of Bhutta and Keys (2016).[12] In particular, these are households who (1) did not purchase a new home and (2) increased their nominal mortgage balance by more than 5% from one survey to the next, with a minimum increase of 1,000 dollars.[13] The debt change is computed in real terms.

The sum of first and second mortgages is our outcome variable. Since 1996, the PSID provides detailed information on mortgage types. These reveal that on average, 92% of first mortgages are conventional mortgages, and 5% are home equity loans. Before 1994, the PSID only reports the remaining balance on first and second mortgages in one variable. However, the largest part of extraction happens via first mortgages, as the overall quantity of second mortgages is comparatively small (see Appendix Figures 2.A.11 to 2.A.13). Even at the peak of the boom in 2007, only 9% of households had a second mortgage according to the PSID, with an average balance of $4,600. By contrast, 47% had a first mortgage, with an average balance of about $77,000.

Figure 2.11 shows the extensive and intensive margins of the different groups over time. At each point in time, we report the share of households who extracted equity, upgraded, or bought a new home (extensive margin).[14] We see a pronounced increase in the share of extractors since the mid-1980s, whereas the shares of upgraders and new owners remained relatively constant over time.

The right-hand side of Figure 2.11 documents a surge in the amount by which households changed their debt conditional on extracting, upgrading, or changing from renting to owning (intensive margin). In the PSID, the average extraction amount is approximately $37,000 between 1999 and 2010.[15] This number is close to the estimate by Bhutta and Keys (2016) of $40,000 for this period. The SCF has had a question on equity extraction related to first mortgages since 2004. Despite some differences in mortgage classifications between the SCF and the PSID, the SCF also yields a similar average extraction amount of $39,000 between 2004 and 2010. Appendix

[11] The number of rooms was averaged across all years a household is living in a given house to avoid spurious classifications due to one-time misreporting. Households who increased (decreased) both the size and value of their house by more than 50% were defined as upgraders (downgraders) even if they did not explicitly indicate to have moved.

[12] See also Duca and Kumar (2014) for a similar approach.

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

[14] We focus on these groups because they will be most important for our following analysis. A full version with downgraders and households who sell their homes to become renters can be found in Appendix Figure 2.C.2.

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

Figure 2.11: Intensive and extensive margins by type

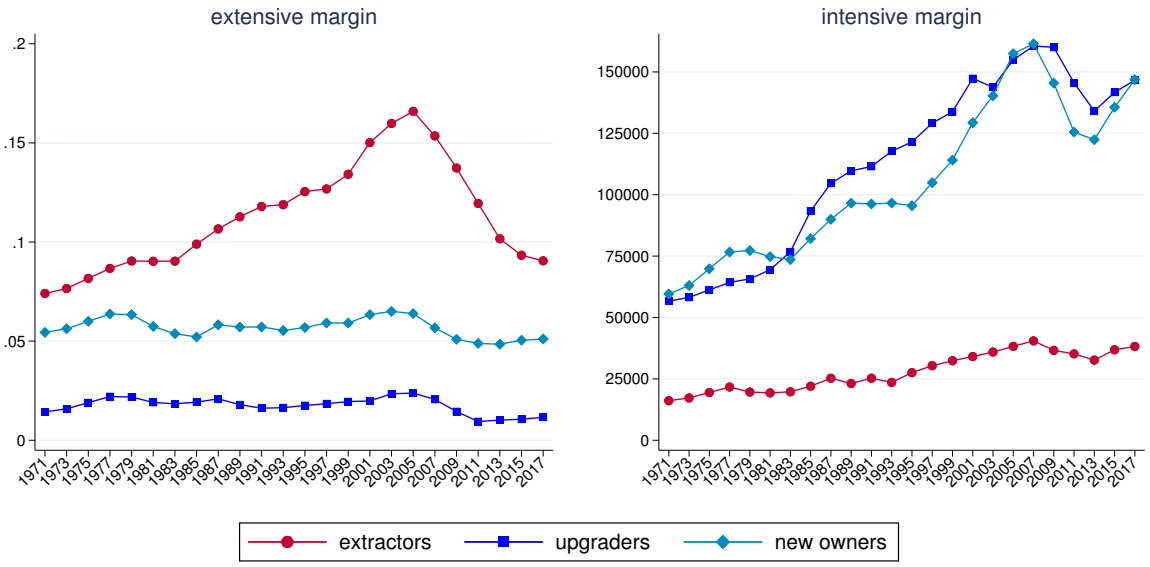

Notes: The left panel shows the share of households who extracted equity, upgraded, or bought a new home over time. The right panel shows the average debt increase of these households. The series were smoothed by taking a moving average across three neighboring waves.

Section 2.C discusses different estimates from the literature in detail and provides a comparison of equity extraction estimates from the PSID and SCF in Table 2.C.1.

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

$$D_t = D_{t-1} + D_{t-1}^+ - D_{t-1}^- - A_{t-1}. \tag{2.2}$$

Between the mid-1960s and early 1980s, the average debt stock was relatively constant (see Figure 2.A.3c). In other words, we had a situation in which $D_{t+1} - D_t \approx 0$, and therefore $D_t^+ \approx D_t^- + A_t$. For D_{t+1} to increase beyond D_t, we need to observe increases in D_t^+ or decreases in D_t^- or A_t.

As a specific example, consider a change in equity extraction. Two reasons account for additional debt due to equity extraction: First, there may be *more households* extracting equity (extensive margin). Second, conditional on extracting equity, households may extract *larger amounts* (intensive margin). Let b denote the base year, and let $\Delta_t D$ denote the average debt increase of households who extracted equity in period t (i.e., the intensive margin in Figure 2.12). Further let s_t denote the sample share of extractors in period t (i.e., the extensive margin in Figure 2.12). The additional debt due to increases in the share of extractors since the base year is $\Delta D_t^{ext} = \Delta D_t \times (s_t - s_b)$. The additional debt due to changes in the average amount by which households increase their debt at the time of extracting is $\Delta D_t^{int} = s_b \times (\Delta D_t - \Delta D_b)$. Adding these two numbers yields our estimate for the amount by which average housing debt would have been lower each period if the share and amount of extractors had stayed at their base-year levels. We cumulate these series to compute the amount by which the stock of housing debt would have been lowered over time in the absence of additional equity extraction. Analogous calculations are done for upgraders, downgraders, and new homeowners.

Figure 2.12 reports the results and plots the contribution of the different household types to the increase in housing debt relative to the base year. We consider data between 1981 and 2007 to cover the whole debt boom period since the 1980s. The dashed line in the figure shows the observed increase in housing debt since 1981.

Figure 2.12: Decomposition of the housing debt boom

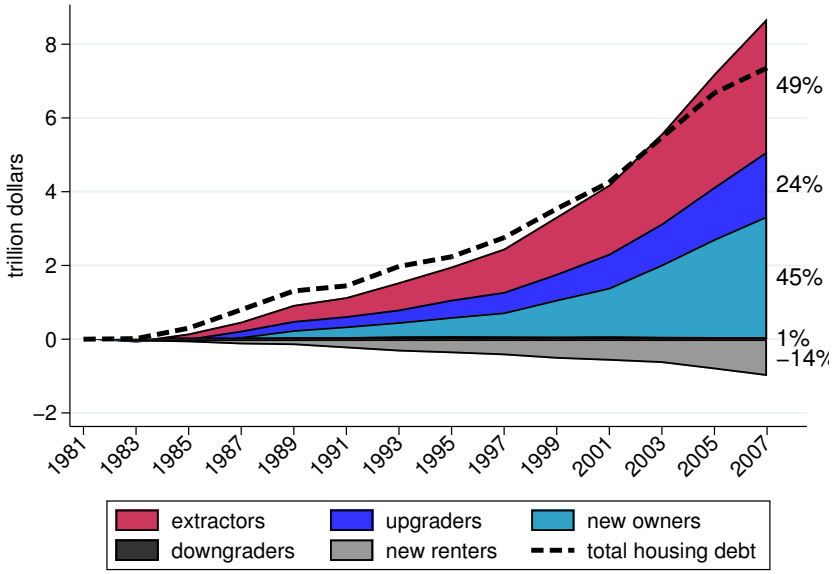

Notes: The graph shows the change in total housing debt since 1981 as a dashed black line, together with estimates of the change in the stock of housing debt due to HEW, upgrading, downgrading, new homeownership, and giving up homeownership. Please refer to the text for details on the construction of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

The first important observation is that our accounting framework closely matches the total housing debt increase between 1981 and 2007. The combined growth in debt across all individual groups accounts for almost the entire debt increase, with only a small residual. Going back to equation (2.2), this implies that no major changes took place in amortization behavior. The second important result is that home equity extraction has played the key quantitative role in driving the debt boom. It accounts for about 49% of the total increase in housing debt. In other words, about half of the increase in housing debt is driven by incumbent owners borrowing against their home equity. New owners account for a slightly smaller share, around 45%. Upgraders account for about 24%, while new renters contribute negatively to the total increase. The net contribution of downgraders was negligible over the considered period.

Together, upgrading and home equity extraction account for more than 70% of additional housing debt since 1981. This corroborates our previous finding that the intensive margin of housing debt is the key driver of the debt boom. Note that both extractors and upgraders tap into home equity for additional spending. Upgraders increase housing consumption by buying a larger house, while extractors may use the funds for home improvements or other consumption purposes.[16]

The relative contribution of new homeownership rose in the mid-1990s, reflecting the increase in homeownership rates prior to the 2008 crisis. While rising house prices bring capital gains to existing homeowners, they imply less purchasing power for prospective homeowners who have saved

[16] In the SCF, households are asked about the purpose for which they extracted home equity since 1995. Among the households who extracted equity, around one-third use the money for home improvements and repairs. Another 30% to 40% spend the money on consumption and the repayment of other debts. Other important purposes are the purchase of vehicles, vacation properties, and investments in other assets, with average response rates of around 5%-10% each.

for the down payment (see also Loewenstein 2018). With falling purchasing power, prospective homeowners have to accumulate more savings or rely on additional debt to finance their home purchase. As most households who change from renting to owning are young, this drove young households deeper into debt than in previous generations. Figure 2.C.3 shows that loan-to-value ratios of young homeowners increased from around 40% in 1950 to almost 80% by 2007. Yet the overall picture is dominated by incumbent homeowners and variations in their intensive margin of debt.

2.4.2 Equity extraction along the income distribution

How was the equity extraction boom distributed across income groups? Is there evidence that middle-class households played an active role in the process? Based on the PSID data, we answer these questions and show that, in particular, households between the 50th and 90th percentiles accounted for the dominant share of equity extraction. These middle-class households also exhibit higher extraction elasticities with respect to house price changes.

Figure 2.13: Extraction relative to income, by income group

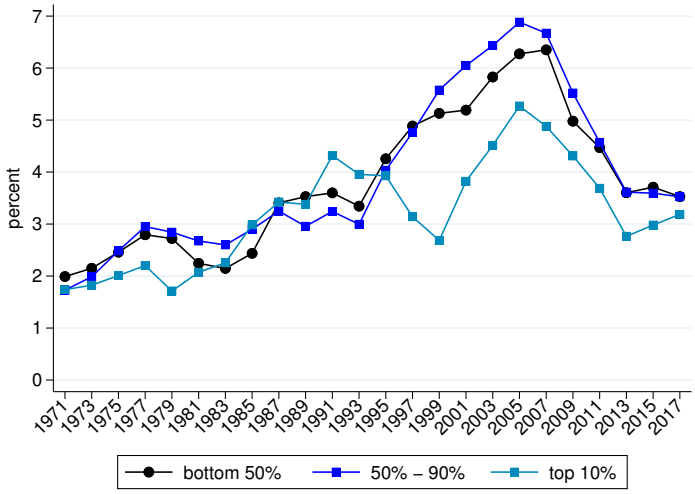

Notes: The graph shows total extraction relative to total income by income group. The series were linearly interpolated for 1973-1977, as mortgage information is not available for the years 1973-1975. Data series have been smoothed by taking a moving average over three neighboring waves.

Figure 2.13 shows total home equity extraction as a share of total annual household income for the bottom 50%, the 50%-90%, and the top 10% of the income distribution. We smoothed the data by taking a moving average across three neighboring waves. Before 1986, the ratio of extraction to income was similar for all three groups, at around 2%-3% for the bottom 90% and 2% for the top 10%. In the mid-1980s, we see an increase in extraction relative to income, which is particularly pronounced for the top 10%. This aligns with the evidence in Maki (1996, 2001), who shows that households reshuffled their debt portfolios from consumer towards housing debt after the Tax Reform Act of 1986. This reform maintained the interest deductibility for mortgage debt, but abolished it for personal debt, which arguably had a larger effect on households with higher incomes.

Over the 1990s, extraction rose from below 4% to more than 6% of annual income for households from the bottom 90%. After the crisis in 2008, it dropped to a level of around 3.5%, where it has remained since 2013. By contrast, extraction was falling over the 1990s for the top 10% and only rose again in the early 2000s. Even at the peak of the debt boom, it did not exceed around 5% of income.

For households with low income growth, additional extraction will translate almost one-to-one into higher debt-to-income ratios. To see this, let us reconsider equation (2.2) and divide by income Y_{t-1} on both sides:

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

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

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

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

$$d_t - d_0 = \sum_{i=0}^{t-1}\left[d^+_i - d^-_i - a_i\right]. \tag{2.3}$$

Until 1985, middle-class households on average extracted 2.6 percent of their annual income over a two-year period. For the period between 1986 and 2007, this figure increased by 2.3 percentage points to almost 5 percent of income. This additional 2.3-percentage-point annual extraction alone translated into a 23-percentage-point increase of the housing debt-to-income ratio.

2.4.3 Aggregate importance of middle-class equity extraction

In order to examine the macroeconomic relevance of middle-class equity extraction, we calculate the importance of equity extraction by the different income groups. To do this, we compute the average amount of additional debt due to extraction, as for Figure 2.12, and multiply it with the total number of households to obtain the aggregate effect. We then add up the resulting series to find the amount by which equity extraction contributed to the aggregate stock of housing debt each period. Finally, we subtract this estimate from total aggregate housing debt, which provides us with an estimate of how much debt would have increased absent the contribution from equity extraction.[17] The black line in Figure 2.14a shows the actual housing debt-to-income ratio from the PSID data.[18] The blue line shows the counterfactual housing debt-to-income ratio after subtracting our estimate of additional debt due to extraction.

Without equity extraction, housing debt would have increased by half as much over the 1981 to 2007 period. Debt-to-income ratios would have stayed at around 40% until 2001 and increased only during the boom of the 2000s, when new homeowners increased aggregate housing debt (see also Figure 2.12). Compared to the observed increase, the counterfactual increase would have been much more modest. By 2007, we estimate that the housing debt-to-income ratio would barely have exceeded 50% of income. Figure 2.14b highlights the role of the middle class in this development. Equity extraction of the middle class accounts for the lion's share of total equity extraction. While the debt-to-income ratios (Figure 2.4a) and extraction-to-income ratios (Figure 2.13) are only slightly smaller for the bottom 50% than for the middle class, the average debt and extraction *levels* are substantially smaller for the bottom 50%, such that their contribution to overall debt and extraction in the aggregate is limited.

We can also approximate the effect on total household debt based on the SCF+ data, which include comprehensive information on non-housing debt. If we assume that housing debt had

[17] This simple estimate abstracts from behavioral and general equilibrium responses.

[18] Note that the housing debt-to-income ratio has increased somewhat less in the PSID than in the SCF+, reaching 0.84 in 2007, compared to 0.92 in the SCF+ (Appendix Figure 2.B.2).

Figure 2.14: Quantitative importance of middle-class extraction

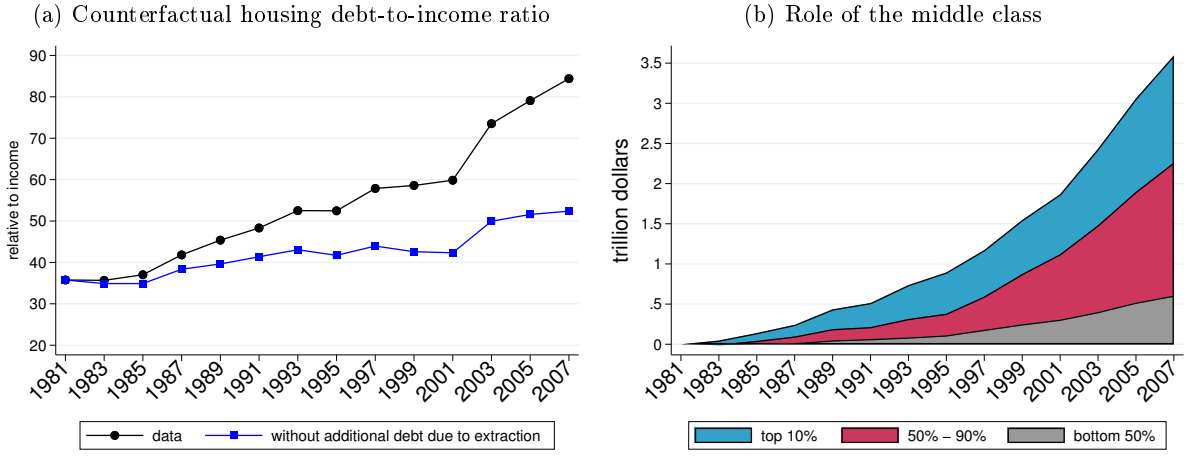

Notes: The left panel shows the housing debt-to-income ratio from the PSID. The blue line with squares shows actual housing debt minus additional debt due to extraction relative to income. The right panel shows additional debt due to extraction by income group.

increased by 50% less from 1983 to 2007 and that non-housing debt had not been affected by the slower increase in housing debt, total household debt relative to income would have peaked a third lower in 2007, at around 74% of income instead of close to around 110% (see Figure 2.A.2).[19] This implies that total household debt would have increased by 43% less between 1983 and 2007.

2.4.4 Drivers of equity extraction

The financial industry started to aggressively market new home equity borrowing products in the 1980s. In the mid-1980s, nearly half of the country's largest financial institutions spent more advertising dollars on these products than on anything else (Canner, Fergus, and Luckett 1988). For instance, Citibank advertised its new "Equity Source Account" by linking house prices to individual achievement: "Now, when the value of your home goes up, you can take credit for it" (Story 2008). Banks were successful in overcoming the negative connotation of second mortgage products, which were traditionally seen as a last resort for households in financial trouble. HELs were now branded as a cheap and convenient way to tap into home equity (Kowalewski 1987).

Within a few years in the 1980s, the HEL market grew from close to zero to $100 billion in volume (Story 2008). Mortgage-backed securities had been invented on the late 1970s and spread quickly from the 1980s on, allowing banks to finance mortgage lending with bonds instead of deposits (Doepke and Schneider 2006). Regulatory changes also facilitated equity extraction. For instance, the Truth in Lending Act had given consumers the right to rescind credit transactions secured by home equity within three days until its amendment in 1982. This had made second mortgage credit burdensome and expensive for the banks. Other Depression-era regulations on the mortgage market were also abolished during these years, allowing mainstream banks to sell secondary mortgage products (Story 2008).

In Appendix Figure 2.C.4, we show how mentions of the term "home equity loan" in American books have evolved over time. The data come from the Google Books Ngram Viewer, an online search engine that displays the frequency of search strings (*n-grams*) in sources printed until

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

2008 (see also Michel et al. 2011). The graph clearly mirrors the historical evidence: Until 1982, the term "home equity loan" was hardly mentioned at all. By 1983, the share of mentions starts to go up and then rises steeply in 1986. After reaching a plateau in the late 1980s, the share surges rapidly again in 1995, when house prices and homeownership rates went on a steep hike.

Another important catalyst were regulatory changes in taxation (see also Kovacs and Moran 2020 and references therein). An important change came with the Tax Reform Act of 1986, which limited the deductibility of interest on all sorts of debt to interest on debt secured by first and second homes. This change meant that homeowners could retain the tax deductibility of interest payments by shifting other debt to housing debt, for example, HELOCs (Kowalewski 1987). In addition, interest rates charged on HELOCs were considerably lower than for credit card debt (Canner, Fergus, and Luckett 1988). Maki (2001, 1996) shows how households took advantage of this reform and changed their debt portfolios from consumer debt toward housing debt after the abolition of the consumer interest rate deductibility.

From the mid-1990s on, house prices were surging particularly fast (see Figure 2.10b). Rising house prices can induce households to increase their mortgage borrowing by relaxing collateral constraints, and thereby liquidity constraints (Aladangady 2017, Cloyne et al. 2019, Aruoba, Elul, and Kalemli-Ozcan 2019, Andersen and Leth-Petersen 2021, Ganong and Noel 2020a, Chen, Michaux, and Roussanov 2020). Moreover, house price increases may lead to an increase in borrowing via wealth effects (Berger et al. 2018). The underlying idea is that rising house prices ceteris paribus increase the value of home equity on the household balance sheet. If households expect this increase in house prices to be persistent, they want to extract these capital gains to smooth their lifetime consumption. However, given that houses cannot easily be divided, the only option of doing so without having to sell the house and move to a less expensive home is to extract equity. In Appendix 2.E, we provide a detailed summary of the theoretical literature discussing under which circumstances such wealth effects may arise, and show in a stylized model framework how they translate into additional mortgage borrowing.

We examined the correlation of house prices and debt across space in our data. Appendix Figure 2.C.5 combines regional information from the SCF+, where we observe the state of residence until 1971, with the PSID. It shows a close comovement of housing and housing debt across regions. Appendix Figure 2.C.6 further confirms that our measure of equity extraction comoves with regional house prices.

While house prices were rising, conforming real mortgage interest rates fell from around 6% in the mid-1990s to 3% in the 2000s (Appendix Figure 2.C.7). This boom provided strong incentives for households to refinance, and many of them extracted home equity on the way via cash-out refinancing (Di Maggio, Kermani, and Palmer 2019). Bhutta and Keys (2016) show that cash-outs accounted for the largest share of equity extraction between the early 2000s and the crisis in 2008, followed by HELOCs and second mortgages. Correspondingly, our measure of equity extraction is correlated with refinancing, and the correlation increases in years which have been identified as periods of refinancing booms in the literature (see Appendix Figure 2.C.8).

In Appendix 2.C.3, we study the correlation of state-level house prices and average mortgage interest rates with equity extraction relative to income along the income distribution. Figure 2.C.9 shows that after a 1-percent increase in real house prices, a typical homeowner from the upper half of the income distribution extracts equity equal to about 0.3% of annual income over the following 8 years. Furthermore, a typical middle-class (top-10%) household extracts equity equivalent to about 2% (3%) of annual income after a 1-percentage-point reduction in mortgage interest rates, whereas households from the bottom 50% show no response to interest rate changes, and only a weak response to house price changes.

2.4.5 Aggregate consequences of equity extraction

Home equity extraction may be optimal from the perspective of an individual household who wants to smooth consumption over the life cycle (see also Boar, Gorea, and Midrigan forthcoming). However, it can also entail substantial risks at the level of the macroeconomy.[20] Previous research has already pointed out that increasing indebtedness at the household level leads to higher financial fragility at the macroeconomic level (Jappelli, Pagano, and Maggio 2013, Fuster, Guttman-Kenney, and Haughwout 2018, Kumhof, Rancière, and Winant 2015).

A full assessment of the welfare consequences of equity extraction would go beyond the scope of this chapter. However, our long-run microdata allow us to "stress test" household balance sheets over the whole period since World War II. The details are described in Appendix 2.F. Our results corroborate that there was a substantial increase in financial fragility coinciding with the post-1980 debt boom. Moreover, our results clearly show that there was no increase in fragility during the first boom phase after World War II, when debt-to-income levels were still much more moderate and home equity extraction did not play a role.

In the following section, we will explore how the evolution of debt over the life cycle has changed across cohorts. We will document a turning and shifting of life-cycle profiles going from older to younger cohorts. These patterns can also be traced back to the increased importance of intensive-margin debt and equity extraction.

2.5 Age structure and life-cycle profiles of borrowing

We have shown that equity extraction can account for about half of the increase in housing debt since the 1980s. An increase in equity extraction has important ramifications for how debt is distributed over the life cycle. Instead of paying down their debt and throwing a "mortgage-burning party" after 20 or 30 years, as previously common (Story 2008), households that extract equity increase their debt balance again in the middle of their life cycle. While the median "new owner", who takes out a mortgage to buy a home instead of renting, is 34 years old in the PSID data, the median extractor is 47 years old. This changes the age composition of U.S. household debt.

The left panel of Figure 2.15 shows that the shares of different age groups in total household debt were very stable over time until around 1990. Thereafter, the share of households above 45 years increases from around 40% to 60%, at the expense of younger households. The shares stabilized again after the Financial Crisis. Only part of the increase in the total debt share accounted for by older households is driven by changes in the age composition of the population. The right panel of Figure 2.15 tells us that the share of households aged 45 and older increased from around 50% in 1989 to 62.5% in 2019. If the age composition had stayed constant at the beginning-of-sample shares for each age group, households above age 45 would only have accounted for around 50% instead of 60% of total debt in 2019 (see Appendix Figure 2.D.1a). This implies that the shares in total debt increased faster than the population shares for households above 45, as illustrated in Figure 2.16.

Although older households got more indebted over time, and the share of older households increased, it is noteworthy that the pure age composition effect exerted *downward* pressure on average household debt. This is because even though they increased their average debt holdings over time, households above age 45 still have lower average debt levels than households below age 45 (see Figure 2.17a). Consequently, average household debt increased by less than what it

[20] Kovacs and Moran (2020) argue that equity extraction may reduce welfare even at the household level by increasing temptation, as it reduces the suitability of housing as a savings commitment device.

Figure 2.15: Shares of age groups in total debt and population

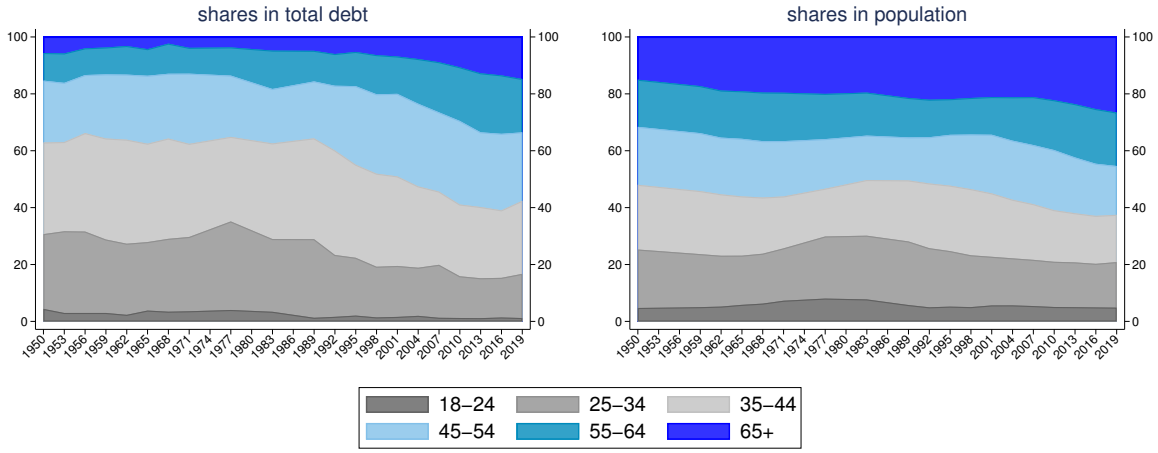

Notes: The left panel shows the share of each age group in total household debt. The right panel shows the population share of each age group among all households.

would have if the age composition had remained as at the beginning of the sample, all else being equal (see Figure 2.17b).

However, mid-life equity extraction and elevated debt levels at older ages imply stark differences in the evolution of debt over the life cycle. The long time period covered by the SCF+, in combination with the rich demographic information, allows us to gain new insights on how the life cycle of debt has changed over the long run. Most importantly, we will see that the slope of life-cycle debt-to-income profiles flattened substantially over time.

Since the SCF+ is not a panel, we construct synthetic birth cohorts. Households with heads born between 1915 and 1924 are our oldest cohort, and households with heads born between 1965 and 1974 are our youngest cohort. Correspondingly, our oldest cohort is on average 30 in 1950, and our youngest cohort is on average 50 in 2019. We estimate the life-cycle profiles of total debt-to-income ratios for each synthetic cohort by regressing individual ratios on six age group dummies. We focus on households between 25 and 85 years of age. The groups comprise

Figure 2.16: Debt share relative to population share

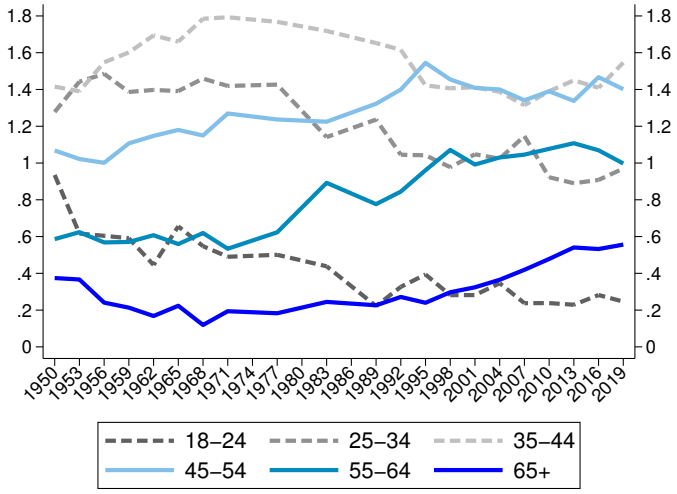

Notes: The graph shows the debt share of each age group relative to its population share.

Figure 2.17: Contribution of change in age composition to average debt

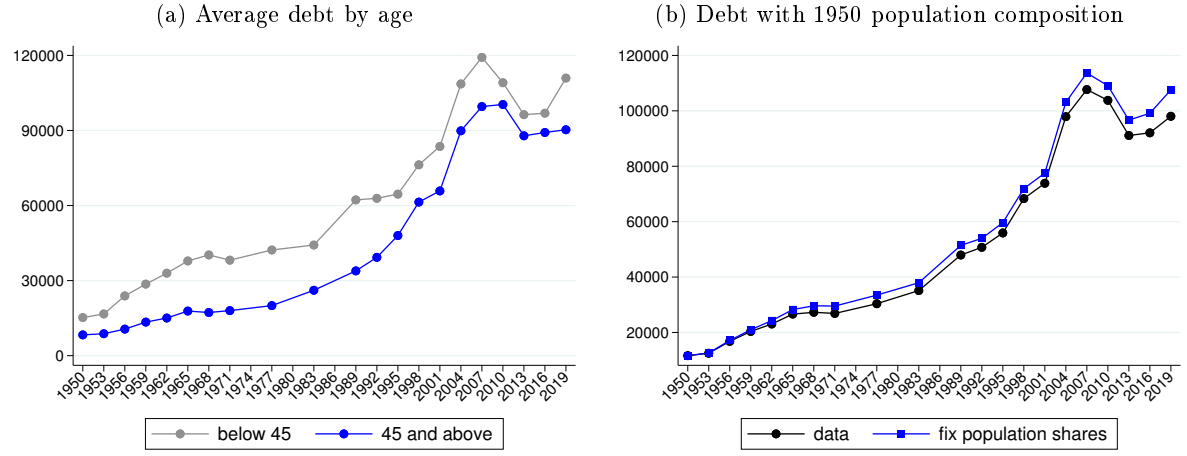

Notes: The left panel shows the average debt of households below and above 45. The right panel shows a counterfactual for average total household debt, keeping the population composition fix as in 1950.

households with a head of 25-34, 35-44, 45-54, 55-64, 65-74, and 75-85 years, respectively.

The resulting life-cycle profiles are shown in Figure 2.18a. Appendix Figure 2.D.3 shows the corresponding life-cycle profiles for housing debt-to-income ratios. Given that housing has always accounted for the largest share of total household debt, it is not surprising that the pattern is very similar. We observe a striking increase in debt-to-income ratios from one generation to the next, leading to an upward shift in the life-cycle profiles across cohorts. For instance, the generations born before World War II started with an average debt-to-income ratio of around 0.5. The debt-to-income ratios of the two baby boomer cohorts, born in the two decades after World War II, were slightly higher at the beginning of their (economic) life cycle. At age 30, they started with debt-to-income ratios between 0.5 and 0.6, possibly reflecting the effects of the postwar credit policies that encouraged homeownership (Fernández 2013).

Apart from the level shift, we also observe a turning of the life-cycle profiles. This upward rotation occurs when the average household from the 1915-1924 cohort is 60, the average household from the 1925-1934 cohort is 50, and the average household from the 1935-1944 cohort is 40. In other words, the turn coincides with the onset of the second debt boom around 1980. This pattern is also clearly visible when looking at life-cycle loan-to-value profiles (see Appendix Figure 2.D.2). These households reach retirement age with substantially elevated debt levels compared to previous cohorts (see also Lusardi, Mitchell, and Oggero 2018, 2020).

At age 60, the visual contrast is stark. The pre-war cohorts typically approached retirement with modest debt ratios of around 30% to 50% of income. Yet households in the first baby boomer cohort (1945-1954) had debt ratios of almost 120% on average at the same age. Generally, younger cohorts reach retirement age with considerably higher debt levels than before. We also note that the effect of the shift in the slope of the life-cycle profiles is considerably stronger than the upward shift in the profiles at the beginning of the life cycle.[21]

Given that the debt profiles of all cohorts are similar until the onset of the second debt boom around 1980, we would expect to see no major differences when assigning all households the average life-cycle debt-to-income ratios of the oldest cohort (born 1915-2924) instead of their actual debt-to-income ratios over this period. Figure 2.18b confirms this intuition. The blue line in this figure presents a counterfactual for which we multiply the average debt-to-income ratio of

[21] Appendix 2.D.2 shows that the same patterns are visible in the PSID data, which allow to follow actual instead of synthetic cohorts.

Figure 2.18: Changes in life-cycle debt dynamics

(a) Debt over the life cycle

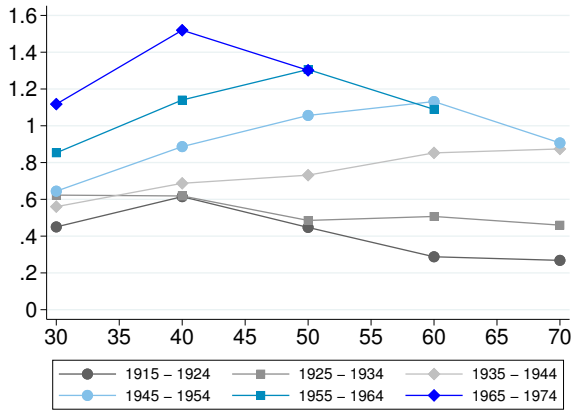

(b) Counterfactual: DTI of oldest cohort

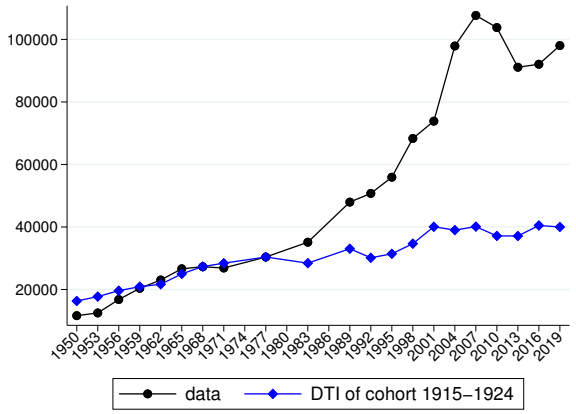

Notes: The left panel shows the life-cycle profiles of total debt-to-income (DTI) ratios for our synthetic cohorts. DTI ratios were winsorized at the 99th percentile within each year. The right panel shows average debt in comparison to a counterfactual assigning the average DTI ratio of the oldest cohort (1915-1924) to all households.

the oldest cohort at each age point in the life cycle with actual household income for all households of the respective age in all survey waves. Before 1980, the resulting counterfactual debt profile is very close to the actual average debt trajectory. Yet after 1980, there is a substantial divergence. If all households had maintained the life-cycle debt-to-income profile of the cohort cohort born between 1915 and 1924, average debt would barely have increased, reaching around 40,000 dollars. In reality, it climbed to a value 2.5 times as high.

Yet even though households have held higher amounts of debt over the life cycle from cohort to cohort, this did not lead to a decline in their net wealth. Figure 2.19a shows that not only the share of total debt held by older households has expanded considerably over time, but also their share in total net wealth. Appendix Figure 2.D.1b shows that even without any changes in the age composition, the share of households 45 and older would have increased from around 70% to around 80% (see also Bauluz and Meyer 2021). In the actual data, it even increased to

Figure 2.19: Wealth by age and over the life cycle

(a) Shares of age groups in total net wealth

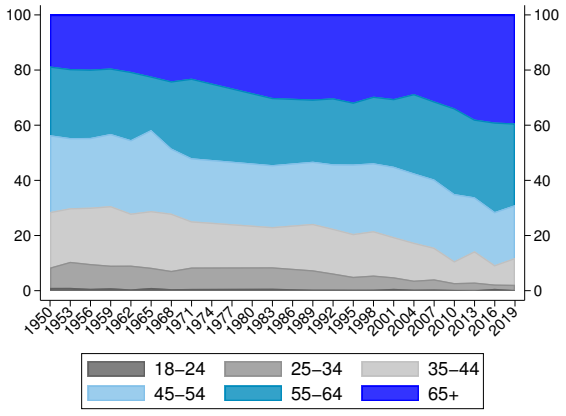

(b) Wealth-to-income ratios over the life cycle

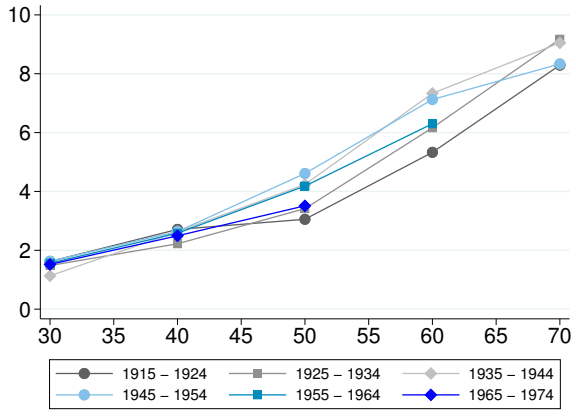

Notes: The left panel shows the share of each age group in total household net wealth. The right panel shows the life-cycle profiles of the wealth-to-income (WTI) ratio for each of our synthetic cohorts. WTI ratios were winsorized at the 99th percentile within each year.

almost 90% (Figure 2.19a). The increase was particularly strong between 2001 and 2010. By consequence, older households have accounted for an increasingly larger share of total debt over time, but also of total assets. Accordingly, higher debt-to-income ratios close to retirement have not led to systematically lower wealth-to-income ratios for younger cohorts, as we can see in Figure 2.19b.

2.6 Racial heterogeneity in debt

Previous research has shown that racial wealth gaps have been extremely sticky over the past 70 decades (see, e.g., Kuhn, Schularick, and Steins 2020, Aliprantis and Carroll 2019, Bartscher et al. 2021 and references therein). An important underlying factor is the black-white homeownership gap, shown in Figure 2.20a. Conditional on being a homeowner, the house of a typical black household is worth much less than that of a typical white household (see Appendix Figure 2.D.4). However, the average housing-to-income ratios of black and white households, conditional on owning a home (i.e., at the intensive margin), are very similar, as Figure 2.20b reveals.

Figure 2.20: Housing by race

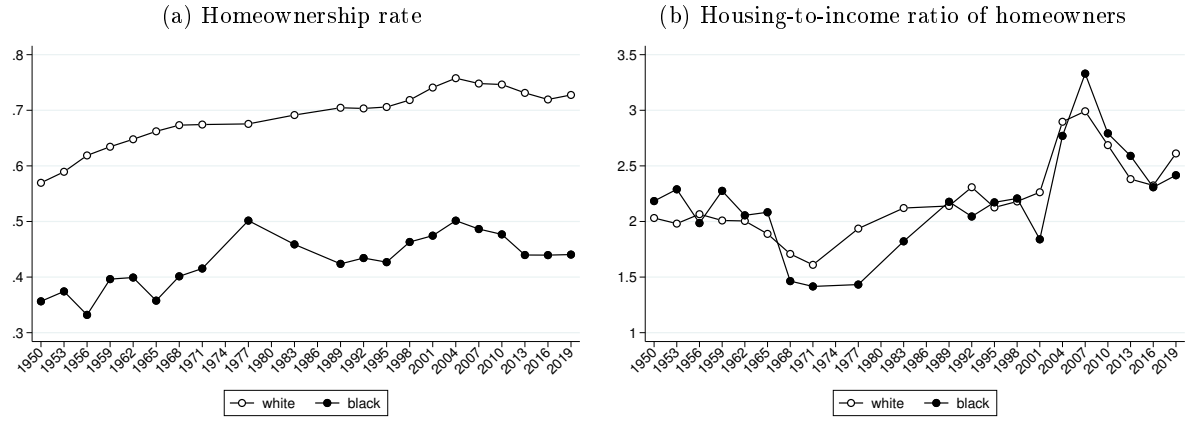

Notes: The graph shows the evolution of homeownership rates (left panel) and housing-to-income ratios of homeowners (intensive margin, right panel) of black and white households over time

The same is true for intensive-margin housing debt of black and white households, as shown in Figure 2.21. Appendix Figure 2.D.5 presents analogous results for non-housing debt, and shows that the intensive margin of non-housing debt has also been similar for black and white households throughout. Since the Great Recession, both housing and non-housing debt-to-income ratios have increased somewhat more for black households at the intensive margin, but overall the evolution for both groups has been similar over time. At the extensive margin, the gap in personal debt between black and white households has closed over time, with a very similar evolution for both groups since the late 1970s (Appendix Figure 2.D.5b). However, the gap in the extensive margin of housing debt has persisted (Figure 2.21b), reflecting the persistent homeownership gap.

Appendix Table 2.D.1 shows that housing accounts for a much larger share of the average black household's asset portfolio compared to the average white household, with around 60% as opposed to 40%. This means that black households have a comparatively high exposure to house price changes. Moreover, black households have lower liquid asset holdings, which may make equity extraction for consumption smoothing purposes more attractive. On the other hand, black households have much less housing assets to borrow against.

As a consequence, borrowing by black households has contributed relatively little to the total

Figure 2.21: Extensive and intensive margins of housing debt-to-income ratios, by race

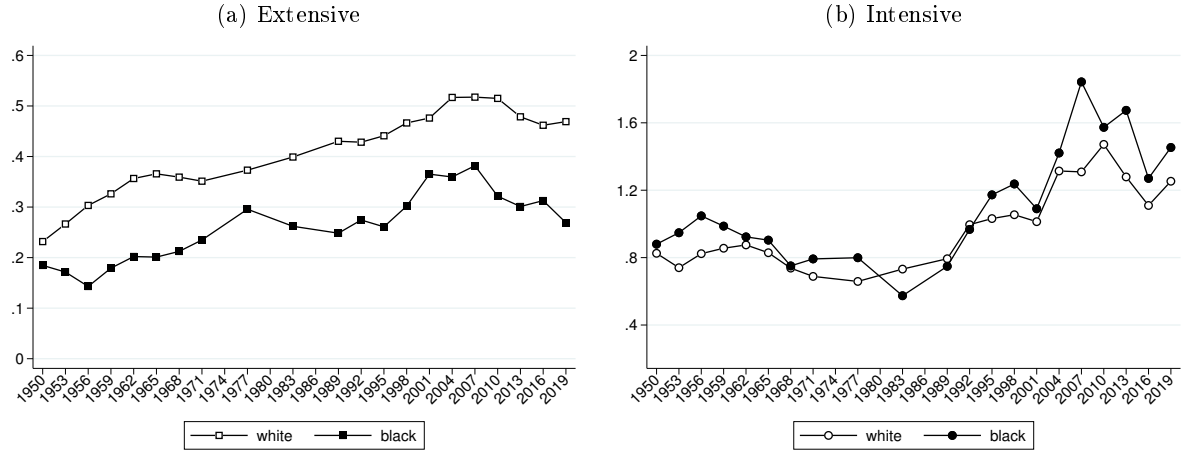

Notes: The graph shows the population shares (left) and housing debt-to-income ratios (right) of black and white households with positive housing debt.

increase in household debt since 1950, as shown in Figure 2.22a. Overall, black households only account for about 9% of the debt increase from 1950 to 2019, although their population share increased from around 9% to around 16% over this period. Moreover, Figure 2.22b repeats the decomposition from Figure 2.14b by race instead of income. It shows that the contribution of black households to the extraction-related debt accumulation since the 1980s was negligible.

Figure 2.22: Contribution of black households to debt boom and extraction boom

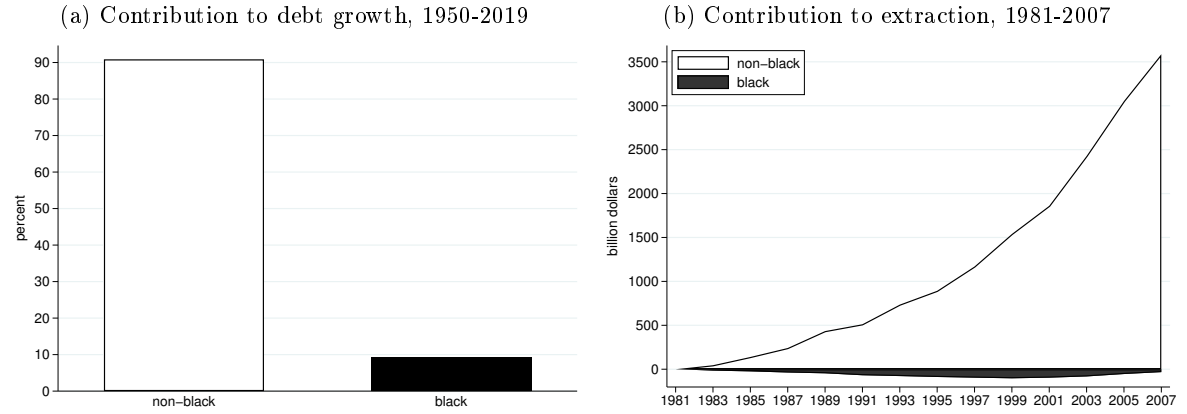

Notes: The left panel shows the share of black and white households in the increase in total household debt since 1950. The right panel shows the contribution of black and white households to additional housing debt due to extraction since 1981.

2.7 Conclusion

This chapter studies the increase in household debt in the United States since World War II. Relative to income, household debt has risen by a factor of four. Yet the financial history of the United States' postwar surge in household debt has remained unwritten. Using long-run household-level data from the SCF+, this chapter helps to close this gap. We document the growth of U.S. household debt, its composition and distribution, as well as its changing nature over time. The past seven decades saw two pronounced debt booms, one after World

War II and one from the 1980s to the 2008 financial crisis. The first boom was triggered by a homeownership expansion in the postwar era, and therefore mainly happened at the extensive margin. By contrast, regulatory reforms, rising house prices and falling interest rates created an environment favoring debt increases at the intensive margin during the second boom phase.

We emphasize the nexus between house prices, housing wealth, and equity extraction. House price increases led to a substantial increase in household wealth, to which households responded by extracting home equity by increasing their mortgage balances. Such home-equity-based borrowing accounts for about half of the increase in U.S. housing debt in the past four decades. We show that the white middle class was the largest contributor to the extraction boom, as well as the overall debt increase since the 1950s, whereas the contribution of black households to the debt increase was relatively small. Moreover, we show that equity extraction, which typically occurs in the middle of the life cycle, led to pronounced changes in the life-cycle profiles of debt. However, since the debt growth was backed by growth in asset values, in particular house values, household net wealth remained similar across cohorts.

Appendix 2.A Additional results

2.A.1 Debt on primary residence and other real estate debt

Figure 2.A.1 shows results for the amount of housing debt on primary residences and for other real estate debt. We find that the debt on principal residences is on average eight times larger than the debt on other real estate. The difference is particularly large in the second half of the sample after 1980.

Figure 2.A.1: Other real estate debt

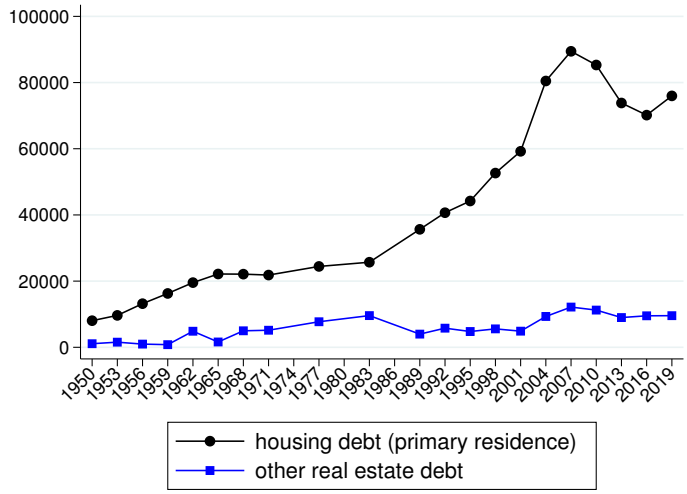

Notes: The graph shows housing debt on owner-occupied real estate in comparison to other real estate debt in the SCF+.

Figure 2.A.2: Total and housing debt-to-income ratios

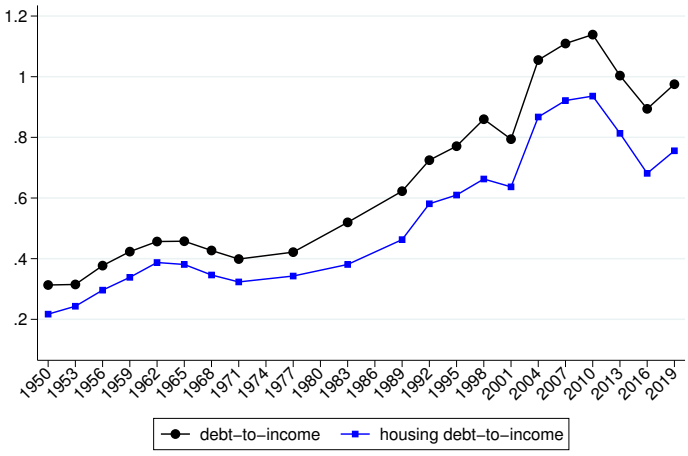

Notes: The graph shows the total and housing debt-to-income ratio from the SCF+ over time.

Figure 2.A.2 shows the evolution of debt-to-income ratios over the last seven decades. Debt-to-income ratios effectively quadrupled between 1950 and the 2007 crisis. They have fallen by about 20 percentage points since then.

2.A.2 Aggregate trends in SCF+ and NIPA

Figure 2.A.3: Income and debt in the SCF+ versus NIPA and FA

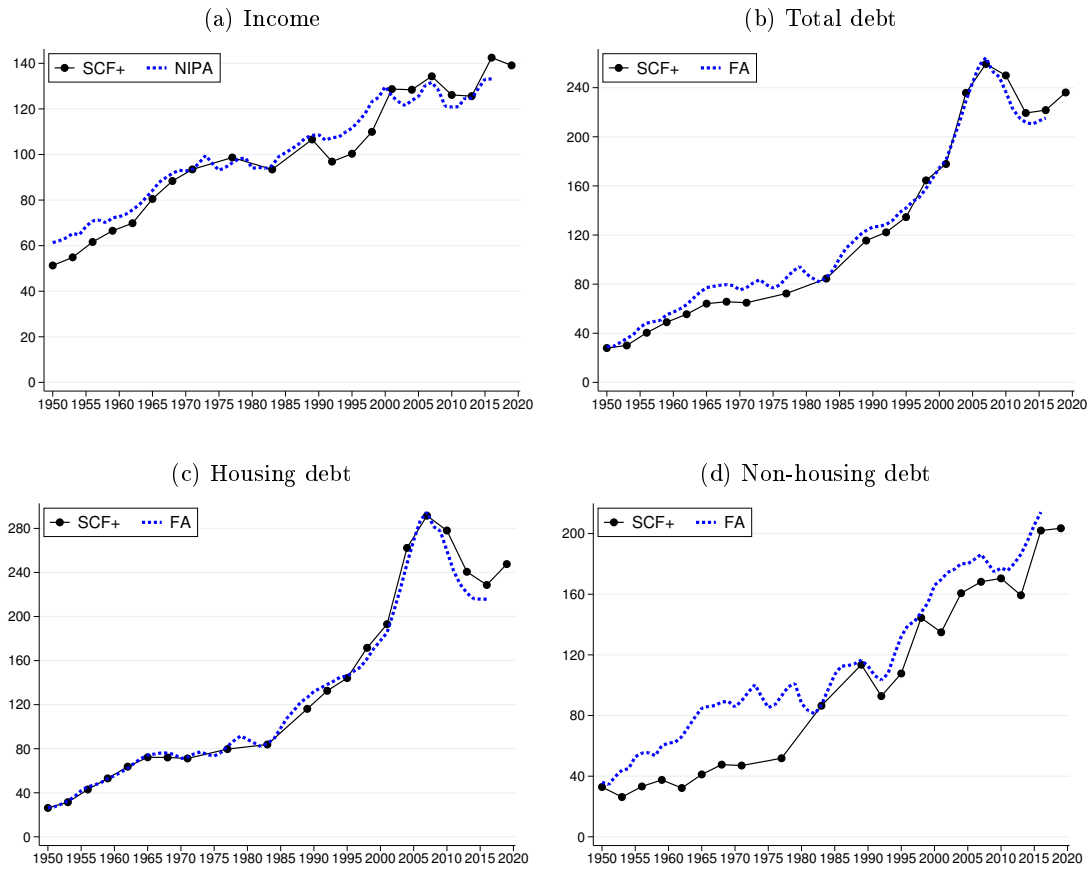

Notes: The figure shows average income and debt from the SCF+ (black lines with circles) in comparison to income per household from the NIPA and debt per household from the FA (dashed blue lines). All series have been indexed to 1983-1989 = 100. Over the index period, the SCF+ values correspond to 95% of NIPA income, 83% of FA total debt, 85% of FA housing debt, and 78% of FA non-housing debt.

We index the series to 100 in 1983-1989 to abstract from level differences that can be attributed to different measurement concepts and focus on comparing growth trends over time. During the base period 1983-1989, the SCF+ data correspond to 89% of NIPA income and 78% of FA debt in levels.[22]

Figure 2.A.3 shows the comparison of growth trends between the SCF+ and aggregate data for 1950 to 2019. Overall, the aggregate data and the aggregated microdata show very similar trends. With respect to housing debt, the SCF+ data and the FA match almost perfectly. Non-housing debt also aligns well with the FA data, albeit there is a certain discrepancy before the 1980s. All in all, the close alignment in growth trends effectively alleviates concerns that the microdata systematically miss parts of the distributional changes underlying the observed macroeconomic growth trends.

[22]The income NIPA components 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 FA debt components. Henriques and Hsu (2014) and Dettling et al. (2015) provide excellent discussions of the different measurement concepts between SCF, NIPA, and FA data.

2.A.3 Group stability over time

Table 2.A.1 documents the persistence within income groups in the PSID data. For households aged 25 to 55, Figure 2.A.4 compares the time series for household debt and income in the PSID when households are binned into income groups based on the contemporaneous income or their beginning-of-decade income. The age range is chosen based on the period of the life cycle when households typically work and buy houses. The SCF data do not have a panel dimension, so we can only sort households based on their contemporaneous income. Figure 2.A.4 demonstrates that the differences between the two sorting approaches are minor, owing to the high degree of persistence of income groups, as shown in Table 2.A.1.

Table 2.A.1: Income group stability

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

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

Figure 2.A.4: Sensitivity: housing debt and income by income group

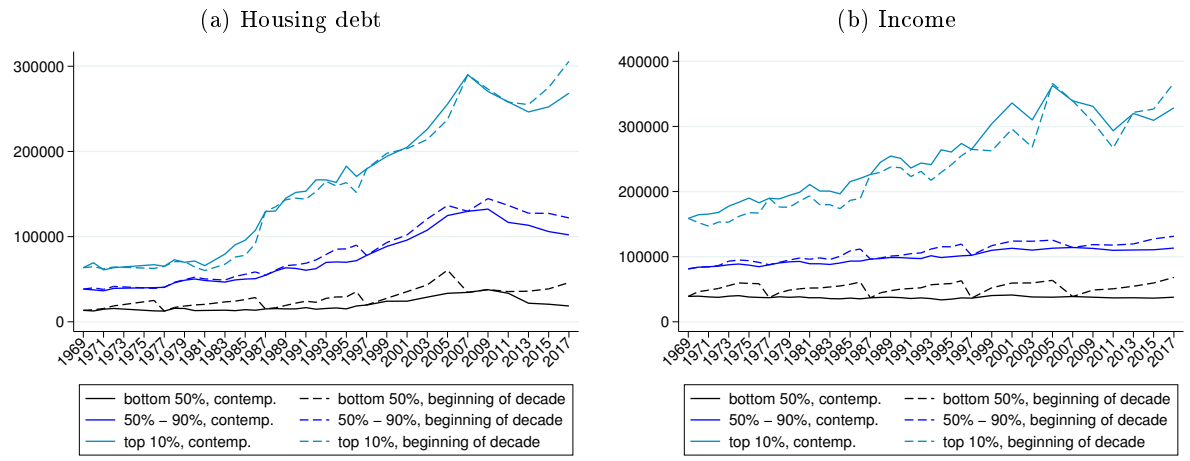

Notes: The graph shows average housing debt (left panel) and income (right panel) by income group for households between ages 30 and 55. We first sort households by their contemporaneous income and show the results as solid lines. For comparison, we sort households by their income at the beginning of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

2.A.4 Debt and income growth across income groups

Figure 2.A.5 shows the share of each income group in the total increase of household debt between 1950 and 2007.

Figure 2.A.5: Share of increase in debt, 1950-2007

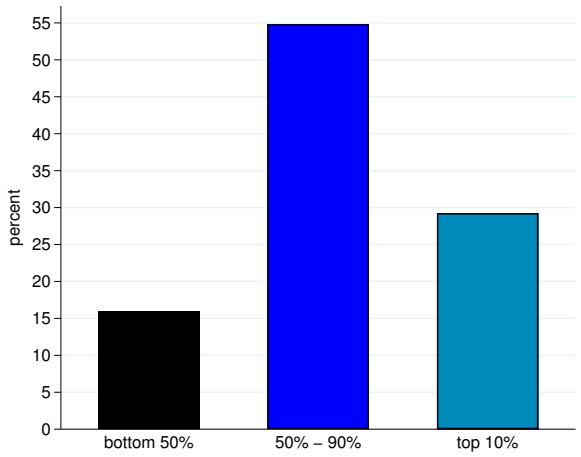

Notes: The graph shows the share of each income group in the total increase of household debt 1950-2007.

Figure 2.A.6 compares debt-to-income ratios of the bottom 90% and top 1%.

Figure 2.A.6: Debt-to-income ratios: bottom 90% and top 1%

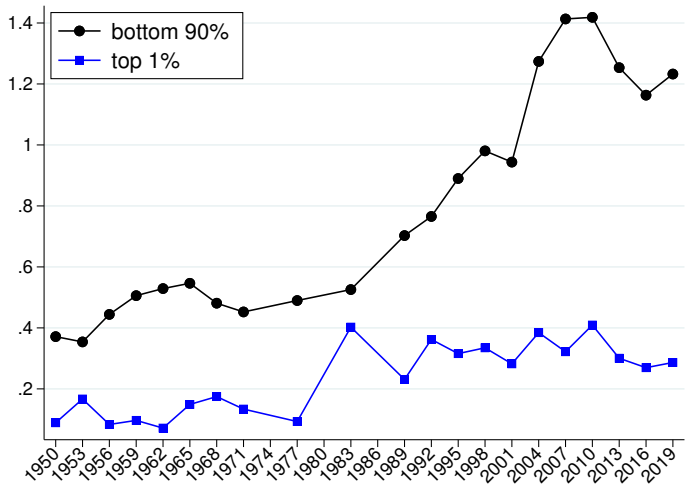

Notes: The graph compares debt-to-income ratios of the bottom 90% and top 1%.

Figure 2.A.7 documents income and debt growth along several socioeconomic lines. All time series are indexed to their averages over the 1970s. Before the late 1970s, all time series comove closely, but they strongly diverge thereafter.

Figure 2.A.7: Debt and income growth along socioeconomic lines

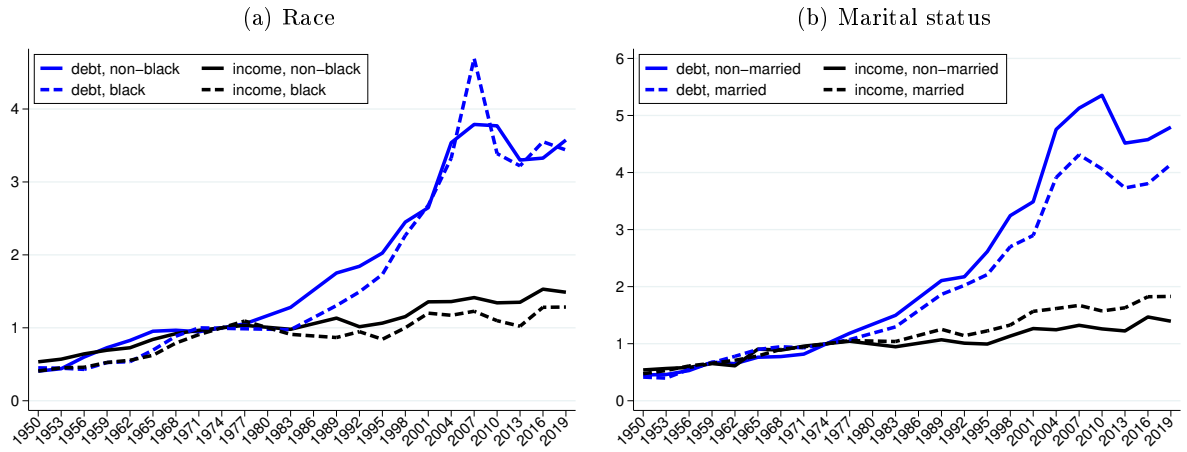

Notes: The graph shows the growth of average total housing debt and income by race and marital status, relative to 1970s averages.

2.A.5 Debt-to-asset ratios over time and by income

Figure 2.A.8 shows debt-to-asset ratios for different income groups over time. We observe a general upward trend in these ratios, but the overall increase is modest.

Figure 2.A.8: Debt-to-asset ratios

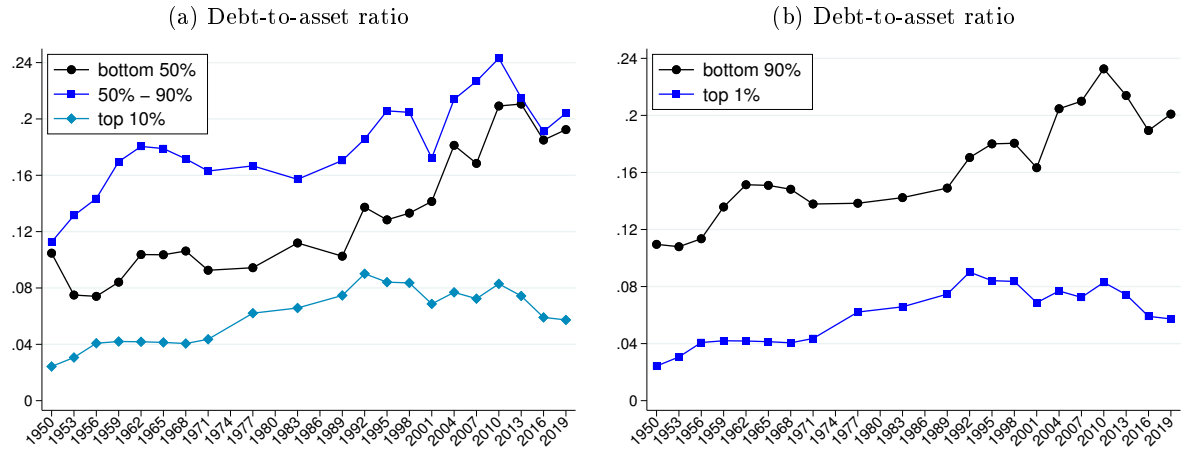

Notes: The left panel shows housing debt-to-asset ratios for the bottom 50%, 50%-90%, and top 10% of the income distribution. The right panel compares debt-to-asset ratios of the bottom 90% and top 1%.

Figure 2.A.9 shows the loan-to-value ratios and debt-to-asset ratios along the entire income distribution for selected SCF+ years. We observe a secular increase that happened along the entire income distribution. Figure 2.A.9a shows loan-to-value ratios along the income distribution. In 2007, LTVs along the whole income distribution exceeded those from the peak of the first debt boom in 1965. Like debt-to-income ratios, leverage has risen most strongly in the middle of the distribution. While middle-class debt-to-income ratios had decreased again in 2019, LTVs were still similar to 2007 because of the simultaneous decline in house values. Debt-to-asset ratios in Figure 2.A.9b also increased along the entire income distribution over time. The increase was more moderate and stronger toward the bottom of the income distribution.

Figure 2.A.9: LTV and debt-to-asset ratios along the income distribution

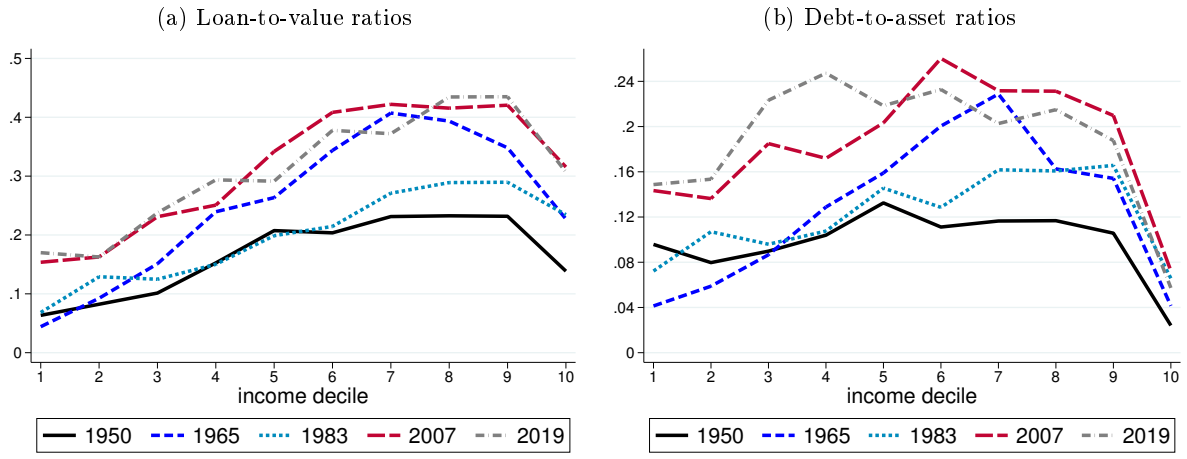

Notes: The left panel shows the evolution of average loan-to-value ratios by deciles of the aggregate income distribution for the SCF+ waves 1950, 1965, 1983, 2007, and 2019. The right panel shows the evolution of total debt to total assets. We exclude households with total income below 10% of the annual wage of a household with a single earner receiving the contemporaneous minimum wage.

2.A.6 Credit cards, education debt, and composition of mortgages

Figure 2.A.10 decomposes the extensive margin of personal debt over time. It shows the extensive margin for all non-housing debt, for the case when education debt is excluded, and for the case when education debt and credit card debt are excluded. We observe the largest effect on the extensive margin from excluding credit card debt. Excluding credit card debt reduces the share of households with personal debt by more than 10 percentage points after 1980. Without credit cards, we do not get an increase in the extensive margin of personal debt since 1970.

Figure 2.A.10: Personal debt, extensive margin

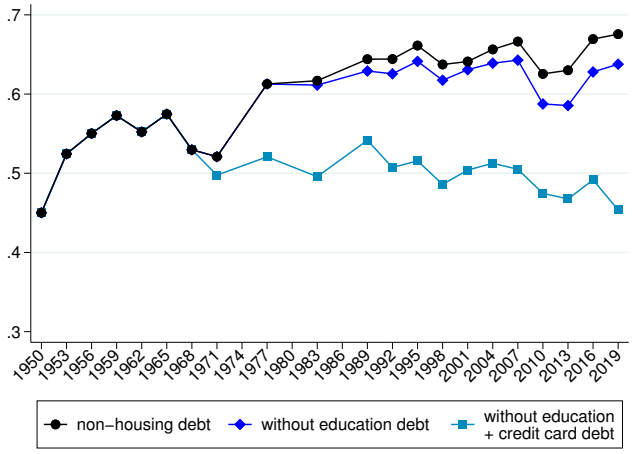

Notes: The graph shows the extensive margin of personal debt from Figure 2.5, together with counterfactuals in which credit card and education debt were set to zero.

Figure 2.A.11 decomposes housing debt into first and second mortgages. The SCF counts HELOCs separately, whereas the PSID counts them among the second (or if no other mortgage is held, even the first) mortgages. Therefore, we re-classify HELOCs, which are available in the modern SCFs since 1989, as first mortgages if no other mortgage is available and as second

mortgages if only a first mortgage is recorded. HELOCs were only introduced on a relevant scale in the mid-1980s (see Maki 2001). Figure 2.A.11 shows the average amount of debt in first and second mortgages in the SCF data since 1983. It also shows the extensive margin of the two types of mortgages, the share of households having first and second mortgages, respectively, which we observe since 1955 in the SCF data.

Figure 2.A.11: First and second mortgages, SCF+

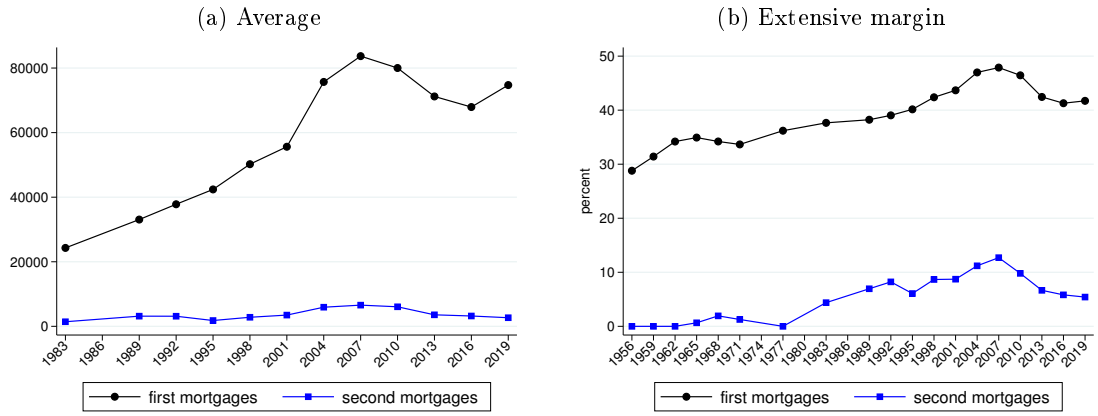

Notes: The left panel shows average first and second mortgages from the SCF. The right graph shows the share of households who have first or second mortgages. HELOCs are included (see text for details).

Figure 2.A.12 looks at the different types of first mortgages in the PSID data. Around 90% of all first mortgages in the PSID are traditional mortgages.

Figure 2.A.12: First mortgages, PSID

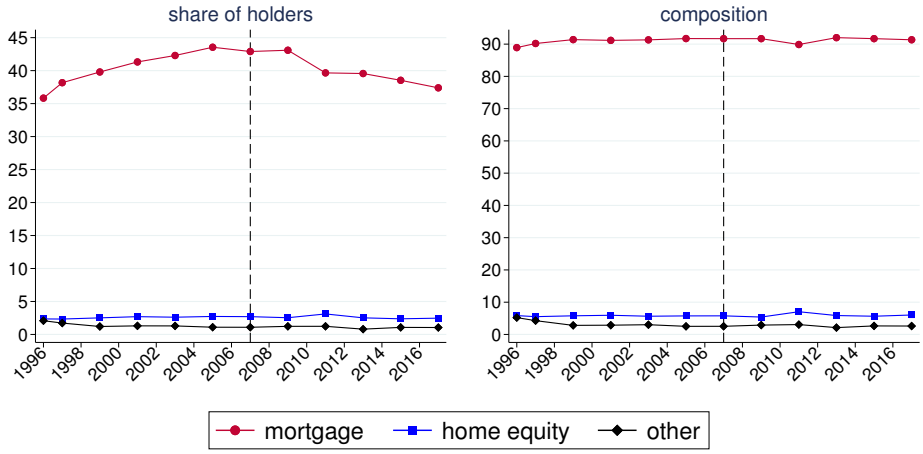

Notes: The left panel shows the share of households in the PSID who hold the respective type of mortgage. The right panel shows the share conditional upon having a first mortgage.

Figure 2.A.13 considers second mortgages in the PSID data that are observed from 1996 onward. The share of households with second mortgages is increasing over time, but even at the peak of the housing boom in 2007, not more than 9% of households had second mortgages. For the households with second mortgages, typically two-thirds were home equity loans, and the share remained quite stable over time.

Figure 2.A.13: Second mortgages, PSID

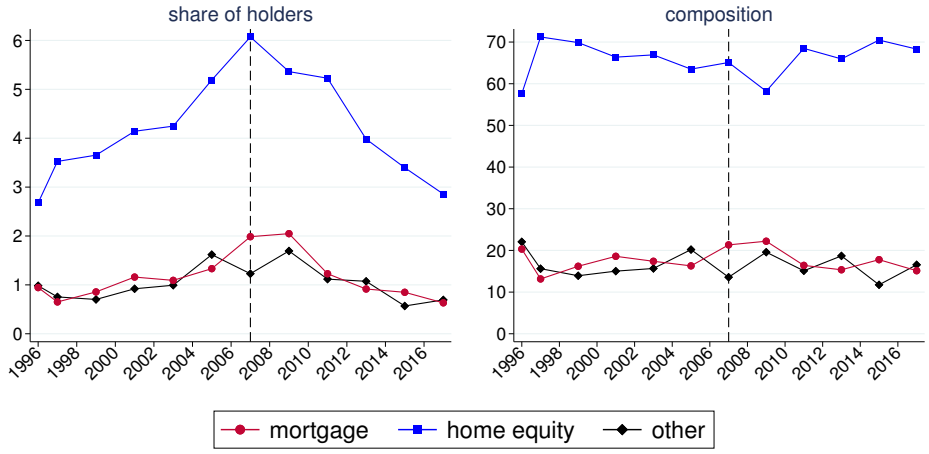

Notes: The left panel shows the share of households in the PSID who hold the respective type of mortgage. The right panel shows the share conditional upon having a second mortgage.

Appendix 2.B Comparison of PSID and SCF+ housing data

In this section, we compare the data on the two main variables of interest, housing and housing debt, from the PSID and the SCF+. The SCF+ collects data at the household level, whereas the PSID collects data at the family level. To make the data comparable, we aggregate PSID families living together into one household (cf. Pfeffer et al. 2016).[23] All variables are taken from the two surveys as they are, without further harmonization of income, asset, and debt concepts (cf. Pfeffer et al. 2016 for a comparison of the survey instruments with respect to wealth).

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

Figure 2.B.1: Comparison of average house value and housing debt: PSID vs. SCF+

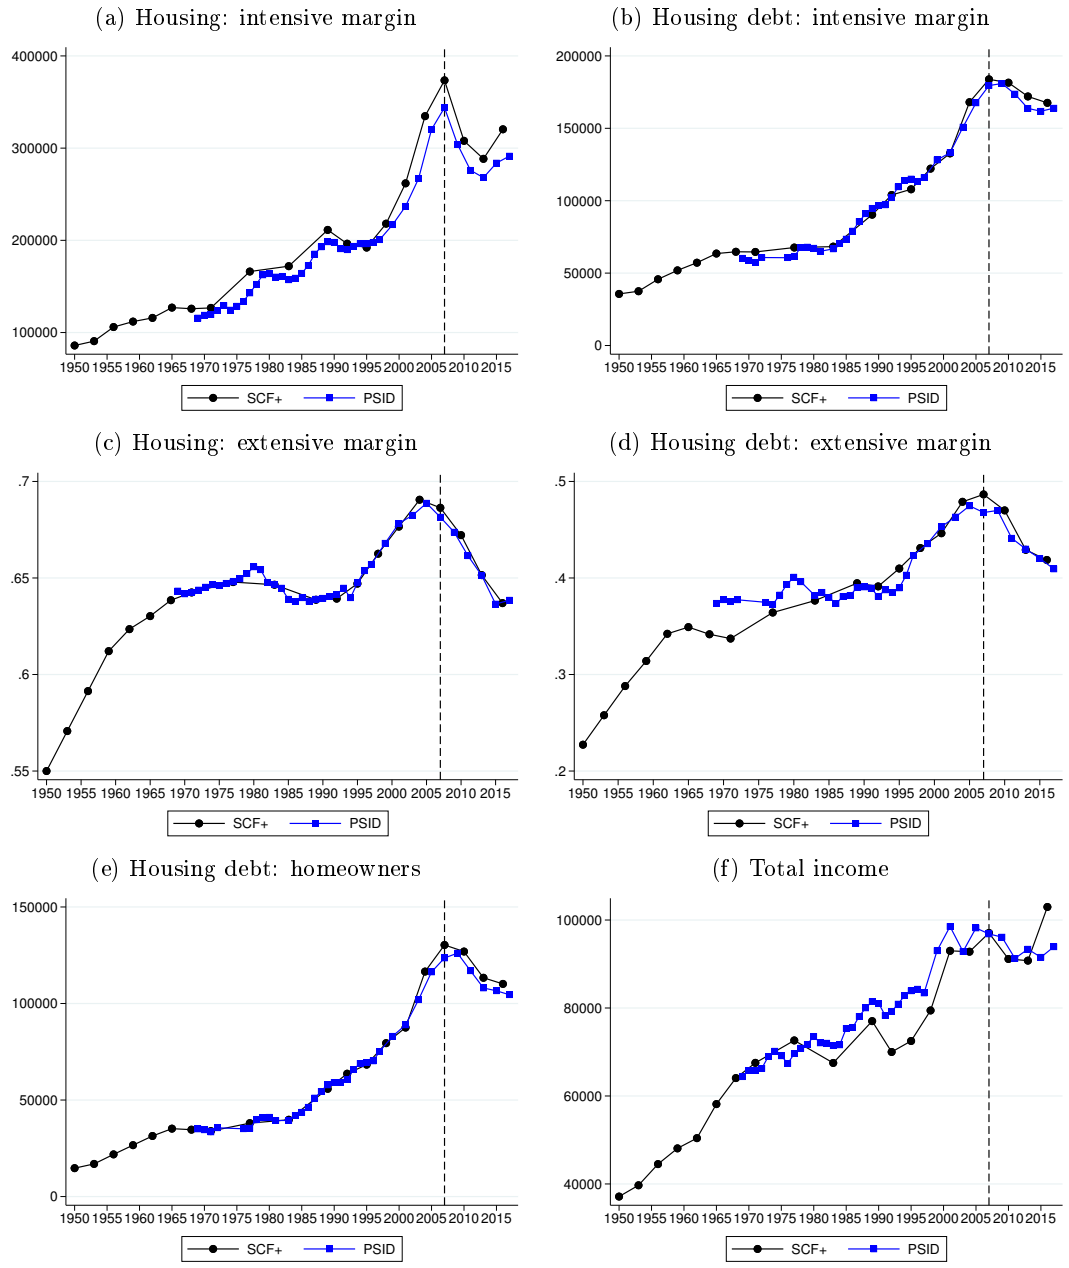

Notes: Panel (a) shows the average value of a house, conditional on being a homeowner. Panel (b) shows the average value of housing debt, conditional on having any housing debt. Panel (c) shows the homeownership rate. Panel (d) shows the share of households with positive housing debt. Panel (e) shows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

Figure 2.B.1 shows the intensive and extensive margins of housing and housing debt from the two data sources. We find that the two datasets yield very similar results at both margins. The intensive margin for housing is lower in the PSID, consistent with the fact that the SCF provides a better coverage of the right tail of the wealth distribution. The intensive margin of housing debt is matched very closely. There are some differences at the extensive margin for debt, especially during the 1970s and housing during the 2000s, consistent with the results of Pfeffer et al. (2016), who report several differences in asset ownership rates between the SCF and PSID. Overall, incomes align well between the two datasets.

Figure 2.B.2 shows debt-to-income ratios from the PSID and the SCF+. Both datasets show the secular rise in debt-to-income ratios in the aggregate and for the middle class over time. We find the increase to be slightly more pronounced in the SCF+ data at the aggregate and when focusing on the middle class.

Figure 2.B.2: Housing debt-to-income ratios in the SCF+ and PSID

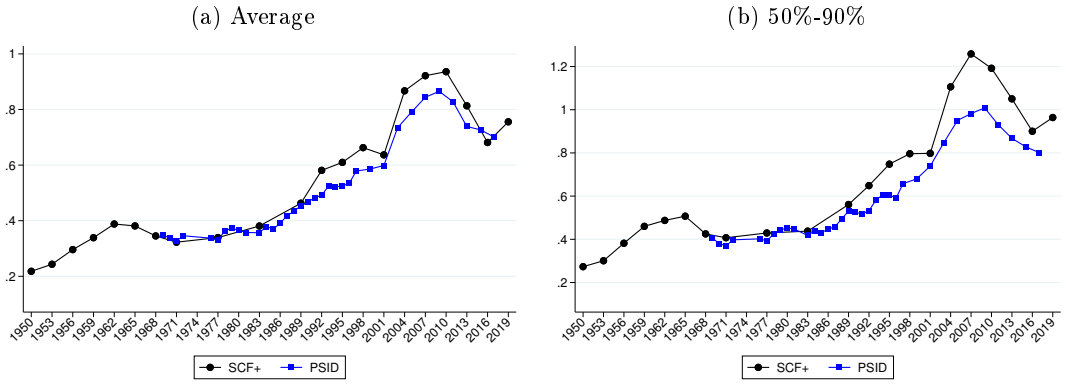

Notes: The graph shows the housing debt-to-income ratio in the SCF+ and PSID over time. The right panel shows results for households from the 50th to 90th percentiles of the income distribution only.

Figures 2.B.3 presents and analogue to Figure 2.9 from the main text based on PSID data.

Figure 2.B.3: Decomposition of home equity, PSID

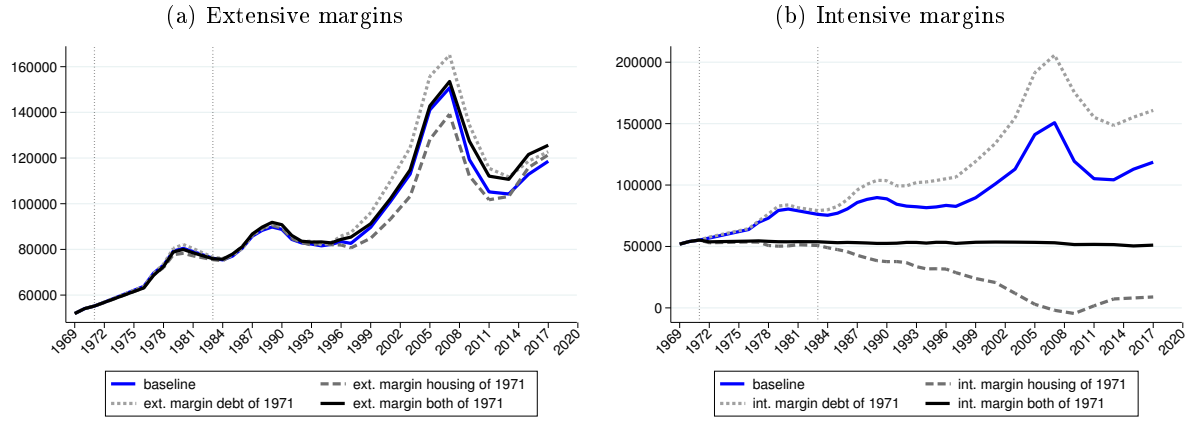

Notes: The graphs shows average home equity from the PSID (blue), along with counterfactuals (gray). The counterfactuals in the left panel hold the extensive margins of housing, housing debt, or both at their 1971 values. The right panel does the same for the intensive margins. The dotted vertical lines indicate the years 1971 and 1983.

Figures 2.B.4 shows the actual change in homeowners' home equity for each Census region since the early 1970s, together with a counterfactual derived from letting the average initial housing assets of homeowners in each region grow with the region's real house price level.

Figure 2.B.4: Counterfactuals: homeowners' equity due to house price growth

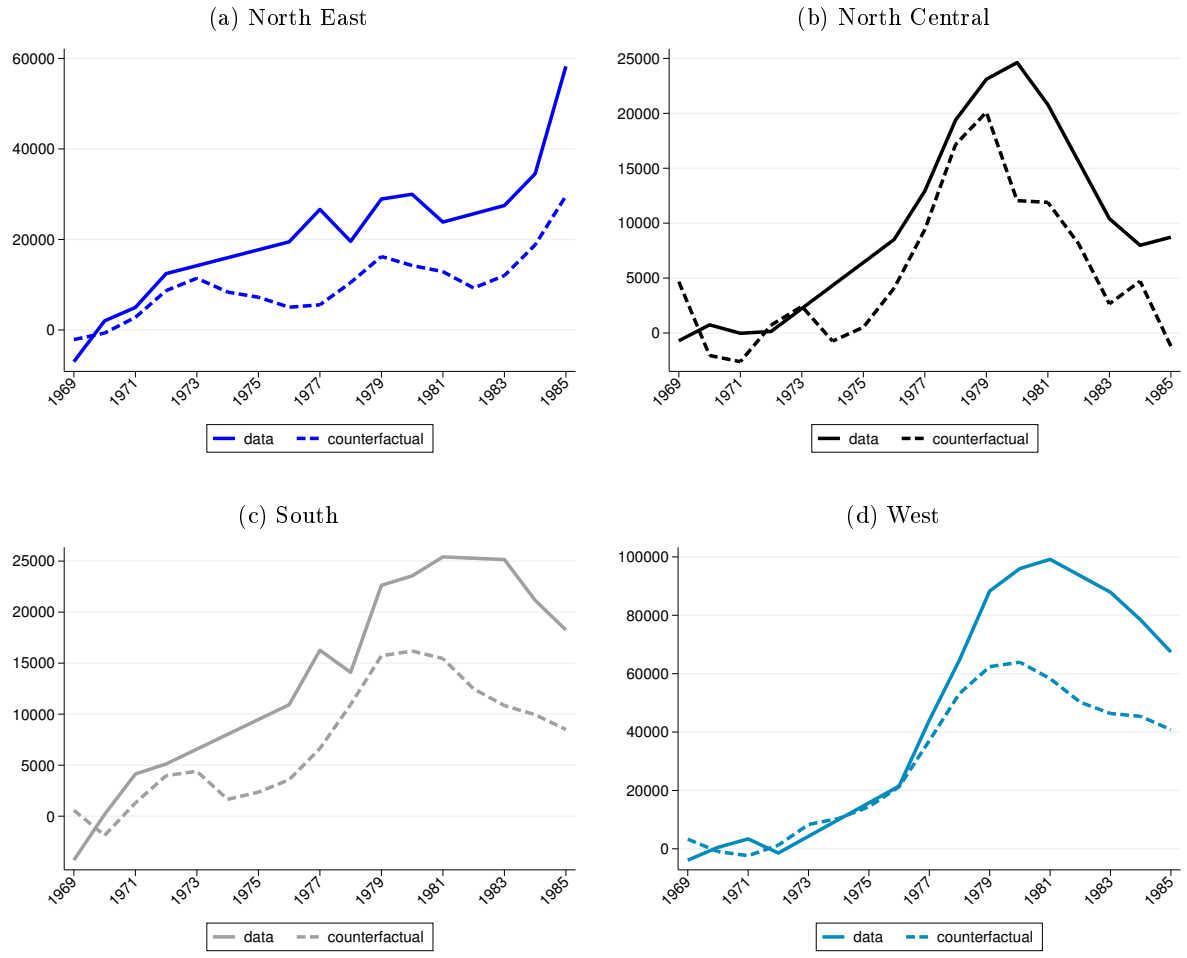

Notes: The panels show the change in homeowners' home equity by Census region relative to the early 1970s (1969-1971), together with counterfactuals derived from letting initial average house values grow with the respective Census region's house price index for new single-family houses sold (including lot value) from the U.S. Census Bureau, converted to real terms using the CPI.

Appendix 2.C Further evidence on home equity extraction

2.C.1 Discussion of previous literature on HEW

Several approaches have been made to quantify the importance of home equity extraction. Bhutta and Keys (2016) estimate that nearly $1 trillion of equity was extracted between 2002 and 2005 via home equity loans, HELOCs, second mortgages, and cash-out refinancings. They exclude the use of funds to move into a more expensive home or buy a second house. According to their calculations, households on average extracted $40,000 between 1999 and 2010. The share of extractors among households with positive mortgage debt holdings varied over time, from 8.5% in 1999 to 18.4% at the peak in 2003. Canner, Dynan, and Passmore (2002) estimate that around $132 billion was extracted via cash-out refinancings from 2001 to early 2002. They estimate that 16%-23% of households with mortgage debt were refinancing, out of which 45% extracted equity.

In the modern SCF, questions on equity extraction via cash-out refinancings and home equity loans have existed since 1995, and the amount has been elicited since 2004. Out of the households surveyed in 2004, 6.4% had extracted equity between 2002 and 2004, which amounts to 13.4%

Table 2.C.1: Average amount extracted

year	PSID	SCF+
1999	35059.19	.
2001	31430.86	.
2003	35829.86	.
2004	.	35185.82
2005	40528.03	.
2007	38408.70	47736.85
2009	42538.48	.
2010	.	34786.23
2011	28867.81	.
2013	34145.05	41825.63
2015	34917.97	.
2016	.	46413.52
2017	41514.09	.

Notes: The table reports the average amount extracted, conditional upon extracting, from the SCF and PSID in 2019 dollars. The SCF measure is based on first mortgages only and refers to households who extracted over the current and previous two years.

of all households with positive housing debt. Among those households who extracted between the last and the current SCF wave, the average extracted amount across all available years was \$41,200 (cf. Table 2.C.1). Extraction information in the SCF refers only to the first mortgage according to the SCF classification. While the PSID counts mortgages consecutively irrespective of their type, the SCF reports HELOCs in a separate variable. The year of origination is reported only for non-HELOC mortgages. Moreover, some households reported having a third mortgage without having a first or second mortgage. Therefore, a comparison of the extensive margin of extraction with the PSID is not straightforward. However, the extracted amount conditional on extracting is of a broadly similar magnitude in both surveys.

Greenspan and Kennedy (2008) take a broader perspective, taking into account existing home sales as well. They estimate that on average, HEW generated around \$590 billion of free cash per year between 1991 and 2006, out of which two-thirds were accounted for by existing home sales. However, their estimates are based on a so-called mortgage system, which was discontinued after 2008, as it did not adequately capture features of the housing market as experienced in the financial crisis of 2007 and 2008. Klyuev and Mills (2007) obtain slightly lower but similar

Figure 2.C.1: Comparison to FA measure of Klyuev and Mills (2007)

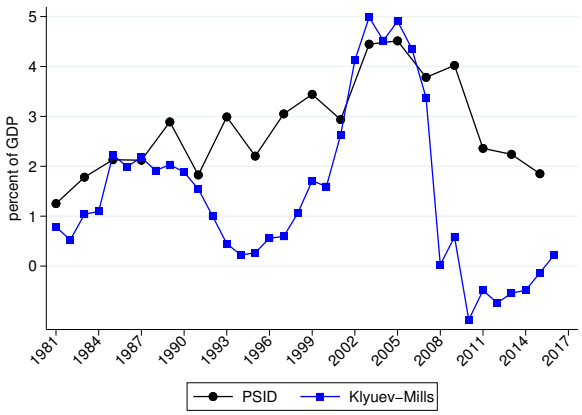

Notes: The figure shows the HEW measure proposed by Klyuev and Mills (2007) and the total amount extracted based on our computations with the PSID, both normalized by NIPA GDP.

estimates with a more simple method. They use the difference between all borrowing secured by dwellings (T_H) and the net acquisition of residential assets (T_{DH}) from the FA as a proxy. The FA mortgage transaction series T_{DH} includes all kinds of mortgages, except construction loans. The housing transaction series T_H includes gross fixed investment in residential structures, net of depreciation, as well as land sales from other sectors to the household sector. However, this "broad" HEW proxy is a somewhat coarse measure of equity extraction. For instance, if a household buys a new home for \$100, and takes out a mortgage for \$80, this measure would count it as *negative* equity extraction (equity injection) of \$20. We compare this measure to our PSID-based equity extraction measure in Figure 2.C.1.

2.C.2 Debt dynamics by household type

Figure 2.C.2 reports the extensive and intensive margins for all household types based on the observed debt dynamics in the PSID: extractors, upgraders, new owners, downgraders, and new renters.

Figure 2.C.2: Intensive and extensive margin by type

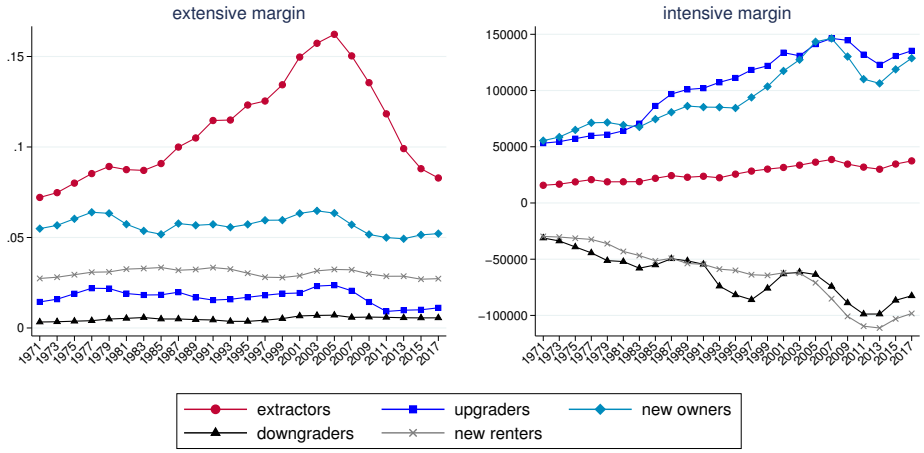

Notes: The left panel shows the share of households who extracted equity, upgraded, downgraded, bought a new home, or sold their home to become a renter. The right panel shows the average debt increase of these households. The series were smoothed by taking a moving average across three neighboring waves.

Figure 2.C.3: Loan-to-value ratios of young homeowners

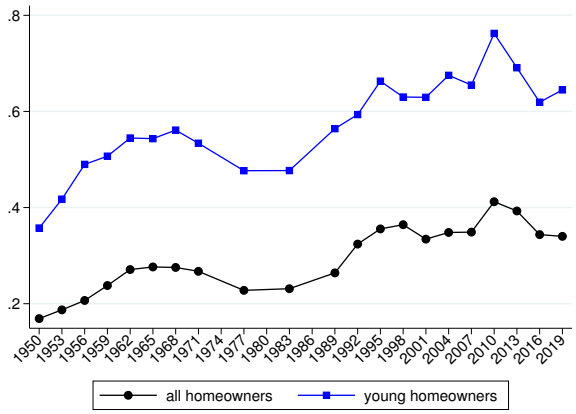

Notes: The graph shows average housing debt relative to average housing for all homeowners and homeowners with a head below age 35.

New owners are typically younger households. Figure 2.C.3 shows loan-to-value ratios of young homeowners (below 35) relative to all homeowners in SCF+ data. We find that LTVs are consistently and substantially higher for younger homeowners, but the difference to all homeowners remained relatively stable over time.

2.C.3 House prices, interest rates and extraction

Figure 2.C.4 provides text-search-based evidence on the proliferation of home equity loans. It reports the number of mentions of "home equity loans" from 1950 to 2007. While the phrase was virtually not mentioned before 1983, it increased dramatically afterward in line with a widespread proliferation of these products.

Figure 2.C.4: Google Books Ngram Viewer for "home equity loan"

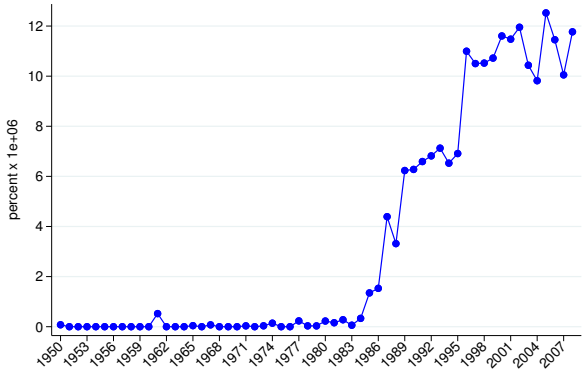

Notes: The graph shows how mentions on the 3-gram "home equity loan" have evolved over time. The figure is based on data from the Google Books Ngram Viewer. The y-axis shows the share of this 3-gram among all 3-grams contained in the Google sample of books written in English and published in the United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

Geographic variation in house values, debt and equity extraction

Figure 2.C.5: Housing and housing debt by Census region (intensive margin)

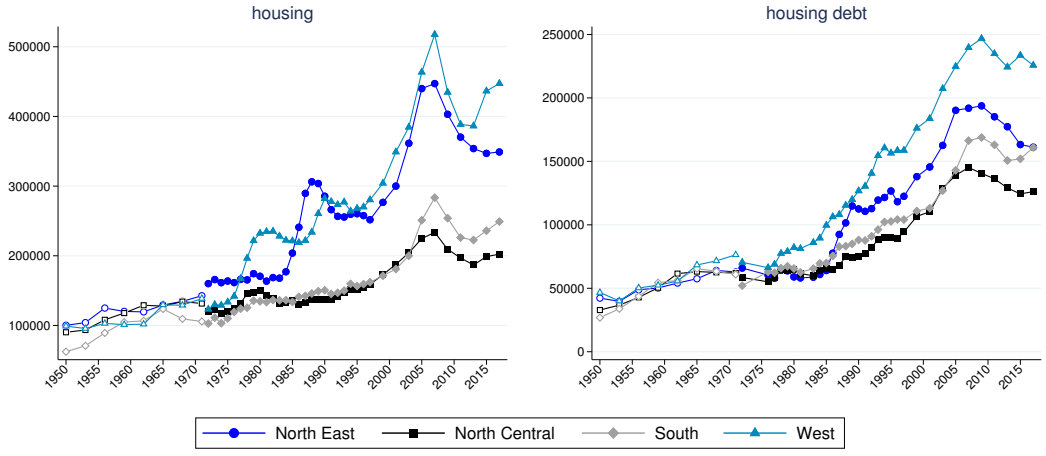

Notes: The graph shows the intensive margin of housing and housing debt by Census region. Filled markers show PSID data, and hollow markers show SCF+ data.

The left panel of Figure 2.C.5 shows the average value of housing assets for homeowners (intensive margin) for the four Census regions in the U.S. The right panel shows the corresponding mortgage debt levels. To construct the time series by Census region, we combine information from the SCF+ and the PSID data. We observe a comovement of housing assets and housing debt across regions. Figure 2.C.6 shows that our measure of equity extraction moves in tandem with regional house prices.

Figure 2.C.6: House prices and equity extraction by Census region

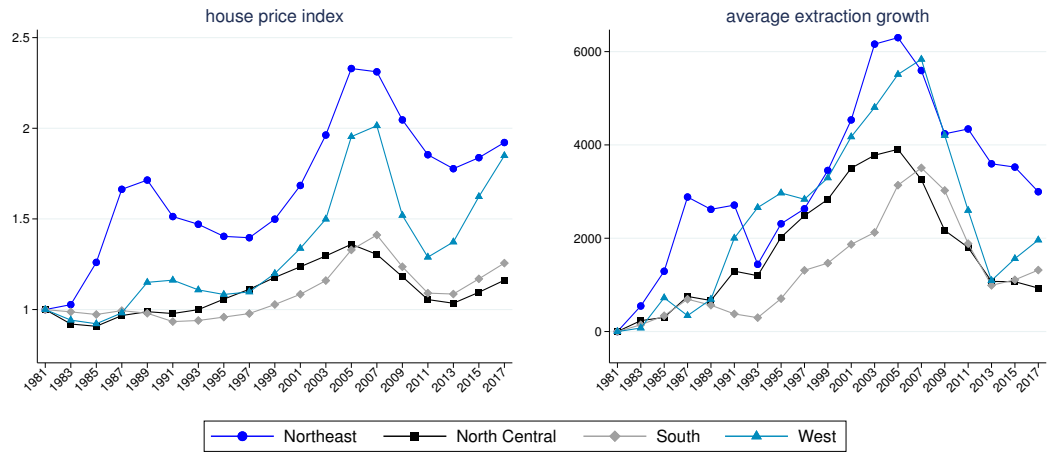

Notes: The left panel shows the growth of state-level FHFA house price indices since 1981, averaged by year and region. The right panel shows the average amount extracted by region, smoothed by taking a moving average across three neighboring waves and normalized by subtracting 1981 levels.

Mortgage interest rates

Figure 2.C.7: Mortgage interest rates (positive housing debt)

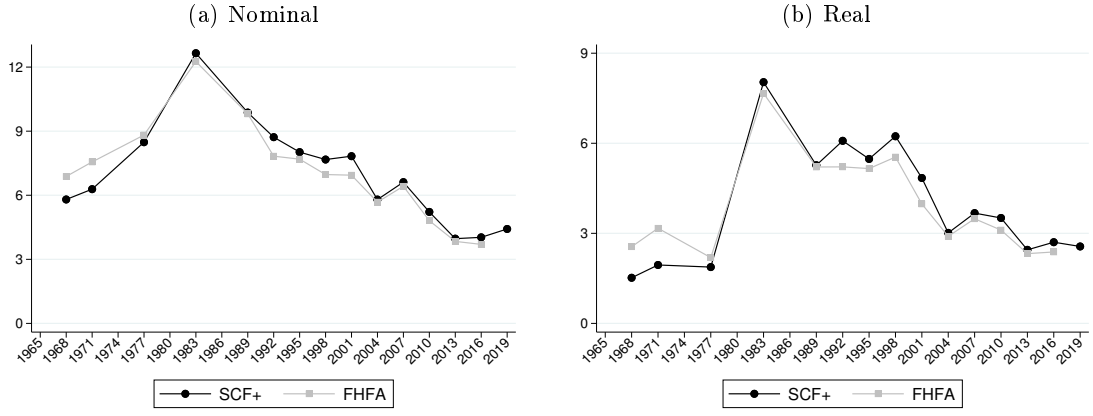

Notes: The graph shows average interest rates on first mortgages in the SCF+ among households who bought their home during the current or previous year. The left panel presents nominal interest rates i_t^m. Real interest rates in the right panel were calculated as $r_t^m = \left[(1+i_t^m)/(1+\pi_t) - 1\right] \cdot 100$, where π_t denotes year-on-year CPI inflation. The black lines with dots present the SCF+ data. As a comparison, the light gray lines with squares show the average interest rate on conventional non-farm single-family mortgages on new and previously occupied homes from the Monthly Interest Rate Survey of the FHFA. The survey excludes FHA-insured and VA-guaranteed loans, loans on multifamily buildings and mobile homes, as well as refinancing loans.

Figure 2.C.7 shows time series for average mortgage interest rates from the SCF+ data and the Federal Housing Finance Agency (FHFA). The estimates align closely and show a clear downward trend from close above 10% to below 5% over four decades.

Figure 2.C.8 shows the pairwise correlation of our indicator for equity extraction and the PSID indicator for refinancing of first mortgages, which is available since 1996. LaCour-Little, Rosenblatt, and Yao (2010) and Bhutta and Keys (2016) report extraction booms in 1998 and 2003. This is mirrored in a particularly high correlation around these years.

Figure 2.C.8: Extraction and refinancing

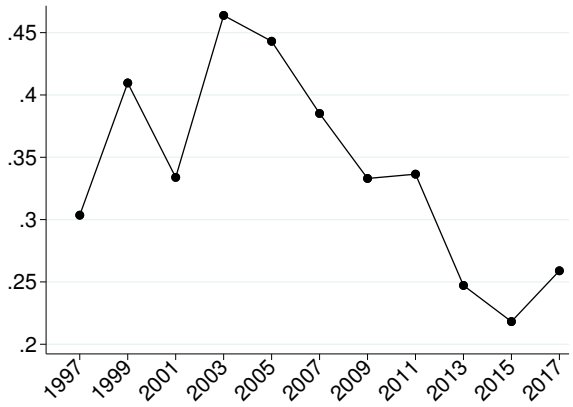

Notes: The graph shows the pairwise correlation of our indicator for equity extraction and the PSID indicator for refinancing of first mortgages over time.

Local projections and event study

We can study the correlation between state-level house prices and extraction in local projections (Jordà 2005). Previous research has stressed that house price exposure can vary considerably across geographies because of heterogeneity in house price developments (Bhutta and Keys 2016, Aladangady 2017, Fuster, Guttman-Kenney, and Haughwout 2018). We use the state-level version of the FHFA house price index[24], as well as the average interest rate on conventional non-farm single-family mortgages on new and previously occupied homes from the Monthly Interest Rate Survey of the FHFA, and estimate the following equation for different horizons h:

$$Y_{is,t+h} = \beta_0 + \beta_1\, g^p_{st} + \beta_2'\, g^p_{st} \cdot \delta_y + \beta_3\, g^r_{st} + \beta_4'\, g^r_{st} \cdot \delta_y + \beta_5'\, \delta_y + \Gamma' X_{ist} + \Phi' \gamma_i + \epsilon_{it}, \qquad (2.C.1)$$

where $Y_{is,t}$ denotes extraction relative to income for household i living in state s in year t. We focus exclusively on homeowners who do not move. The expression $Y_{is,t+h}$ denotes the cumulative extraction relative to income between period t and period $t+h$; g^p_{st} is the percent growth rate of the state-level FHFA house price index between two survey waves; g^r_{st} is the negative percentage-point change in the nominal mortgage interest rate; δ_y are dummies for our three income groups; and X_{ist} is a set of household-level demographic controls that are plausibly related to equity extraction.[25] The regressions also include household fixed effects γ_i to capture time-invariant household characteristics. We use data for the period from 1995 to 2007.

Figure 2.C.9 plots the estimated coefficients $\hat{\beta}_1 + \hat{\beta}_2$ (left panel) and $\hat{\beta}_3 + \hat{\beta}_4$ (right panel) from equation (2.C.1) for $h = 1, ..., 4$. The results imply that after a 10-percent increase in house prices, which roughly corresponds to one standard deviation of house price growth, a typical homeowner from the upper half of the income distribution extracts equity equal to about 3% of annual income over the following 8 years. Furthermore, a typical middle-class household

[24] The repeat sales index of the FHFA is designed to filter out changes in the average quality and size of homes (cf. Rappaport 2007).

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

Figure 2.C.9: Effect of house prices and interest rates over time

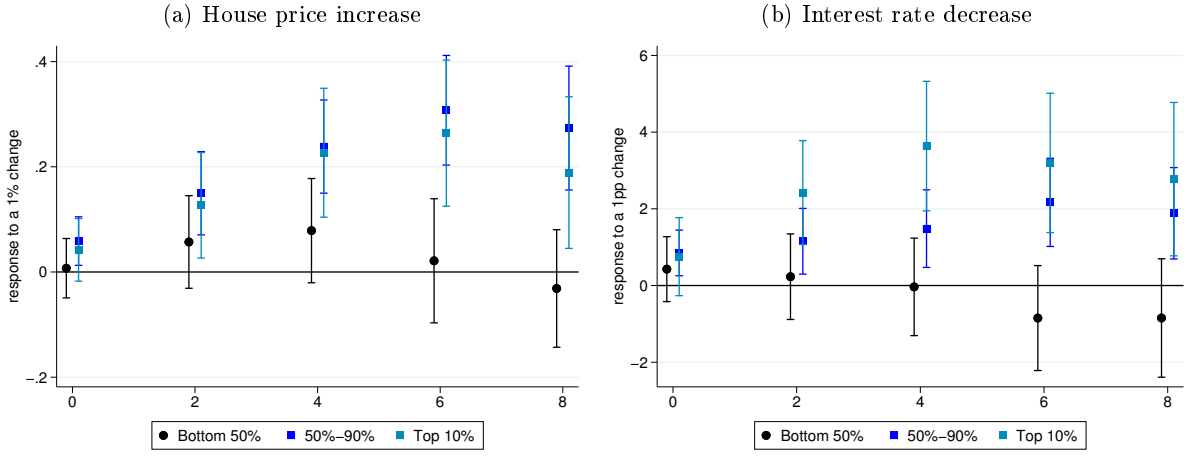

Notes: The left panel shows the estimates of the coefficients $\beta_1 + \beta_2$ at various horizons from equation (2.C.1). The right panel shows the estimates of $\beta_3 + \beta_4$. Observations with extraction amounts larger than twice the annual income or with negative income were excluded. The sample covers two-year periods from 1995 to 2007. Controls include dummies for age, children, and business ownership, as well as household fixed effects. Standard errors are clustered at the household level.

extracts equity equivalent to about 2% of annual income after a 1pp reduction in interest rates. The response is even stronger for households in the top 10%, with more than 3%, but completely absent for households from the bottom 50%.

We also estimated event-studies around the extraction date, using the reported value of a household's home as the outcome variable. The results show that the house values of extractors increased substantially more than those of non-extractors in the six years prior to extraction, consistent with the evidence from the local projections (Figure 2.C.10).

Figure 2.C.10: Event study: extraction

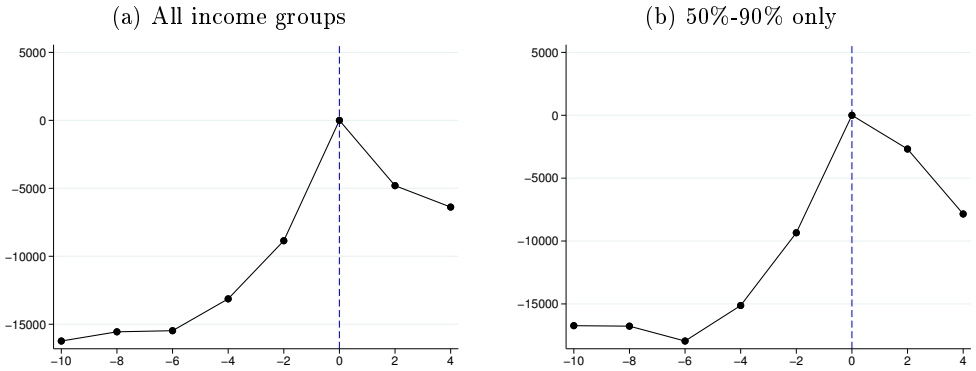

Notes: The graph shows the results of regressions of the house value on leads and lags of the extraction dummy. Zero is the period of extraction. Even years were discarded from the dataset to avoid a change in frequency. We focus on households that stay in their home upon extraction. The regressions include year and household fixed effects.

Appendix 2.D Supplementary results on debt and demographics

2.D.1 Debt and net wealth shares with age structure of 1950

Figure 2.D.1 shows the counterfactual shares of each age group in total debt and net wealth when fixing the age distribution at 1950 population shares for each age group.

Figure 2.D.1: Counterfactual debt and wealth shares with age composition of 1950

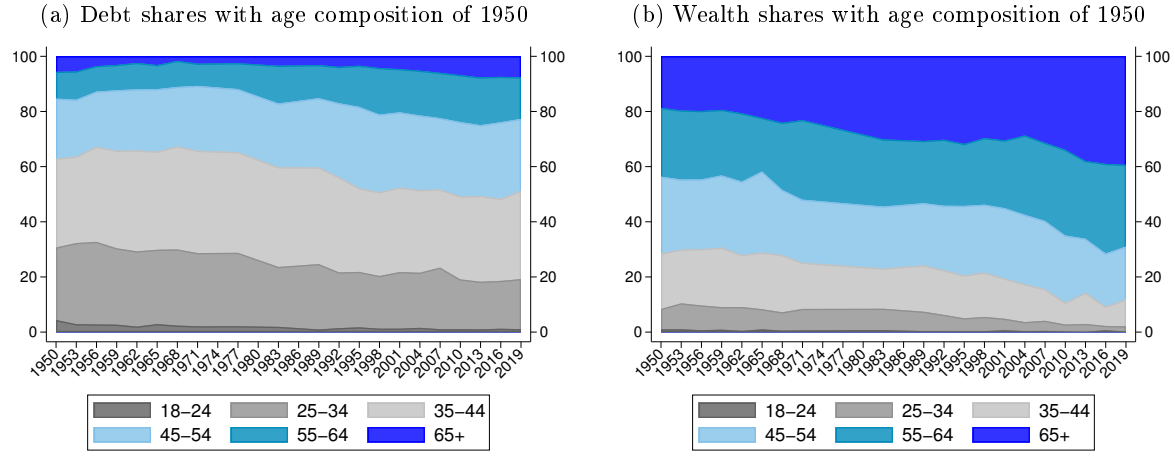

Notes: The graph shows counterfactual debt shares (left panel) and net wealth shares (right panel) of each age group when keeping the age composition of the population as it was in 1950.

2.D.2 Life cycle debt patterns in the PSID

Figure 2.D.2 shows life-cycle loan-to-value profiles obtained by regressing individual loan-to-value ratios on six age group dummies (25-34, 35-44, 45-54, 55-64, 65-74, and 75-85 years). The left panel is based on the SCF+ data. The middle panel shows PSID data treated analogously to the SCF+ data, and the right panel shows results that exploit the panel dimension of the PSID by including household fixed effects. Note that the SCF+ data start in 1950, whereas the PSID data only begin in 1969.[26]

Figure 2.D.2: Comparison of life-cycle loan-to-value ratios

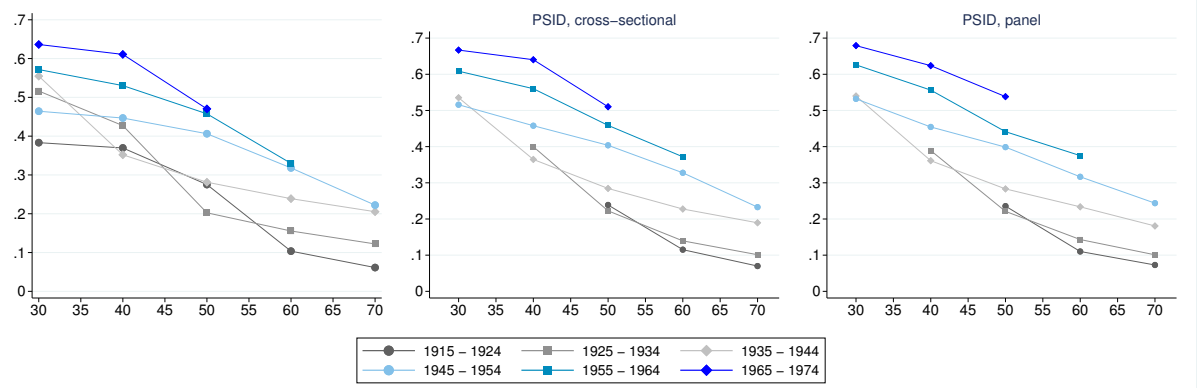

Notes: The graph shows life-cycle loan-to-value (LTV) profiles for different cohorts. The left panel shows the SCF+ data, the middle panel shows PSID data when treating the data as cross-sectional, and the right panel shows PSID data when exploiting the panel dimension by including household fixed effects. LTV ratios were winsorized at the 99th percentile within each year.

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

Figure 2.D.3 shows analogous results for the housing debt-to-income ratio. It also includes a fourth panel, in which we exploited the PSID's panel dimension to replace income by its three-year moving average (MA) within each household. This step helps to avoid extreme values due to temporary income fluctuations. The results are quantitatively and qualitatively similar across both datasets and all specifications. Housing debt-to-income ratios and leverage (loan-to-value) have both shifted and turned upward conspicuously. We see a shift in slopes around 1980 for all cohorts, no matter whether they were 40, 50, or 60 years at this point. The shift is most pronounced for households around age 40 in 1980. The results are very similar when controlling for household fixed effects in the PSID, which confirms that the results obtained with the SCF+ are not artifacts of working with synthetic cohorts.

Figure 2.D.3: Comparison of life-cycle housing debt-to-income profiles

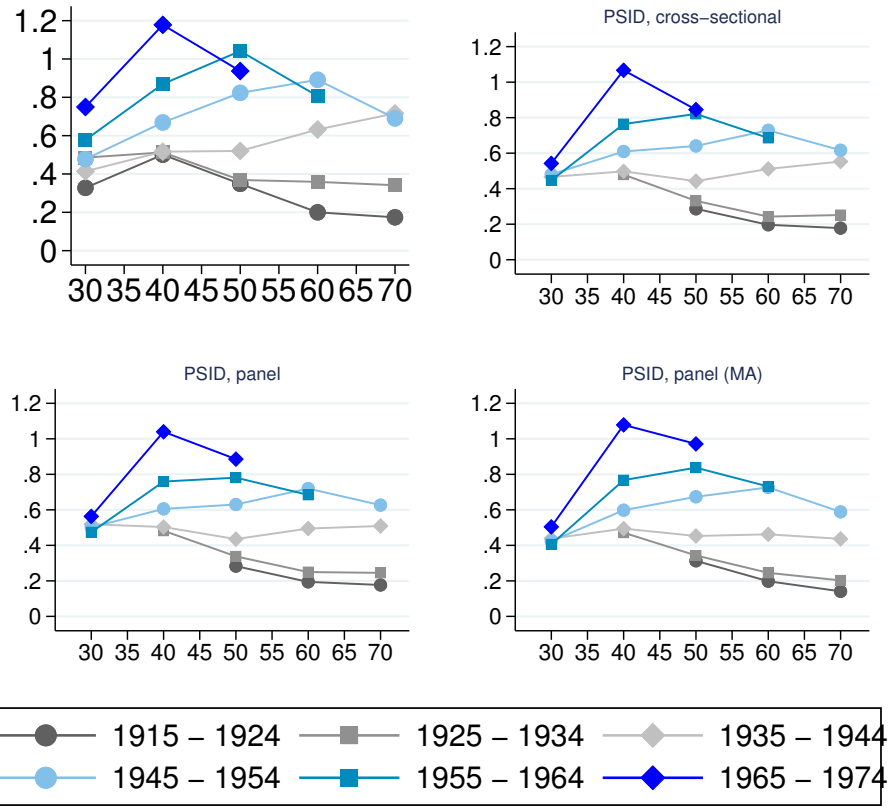

Notes: The graph shows life-cycle housing debt-to-income (HDTI) profiles for different cohorts. The upper left panel shows the SCF+ data, the upper right panel shows PSID data when treating the data as cross-sectional, and the lower left panel shows PSID data when exploiting the panel dimension by including household fixed effects. The lower right panel uses a three-year moving average of total household income in the denominator. HDTI ratios were winsorized at the 99th percentile within each year.

2.D.3 Debt by race

Figure 2.D.4 shows the average house values of black and white homeowners over time.

Figure 2.D.4: House values of homeowners, stratified by race

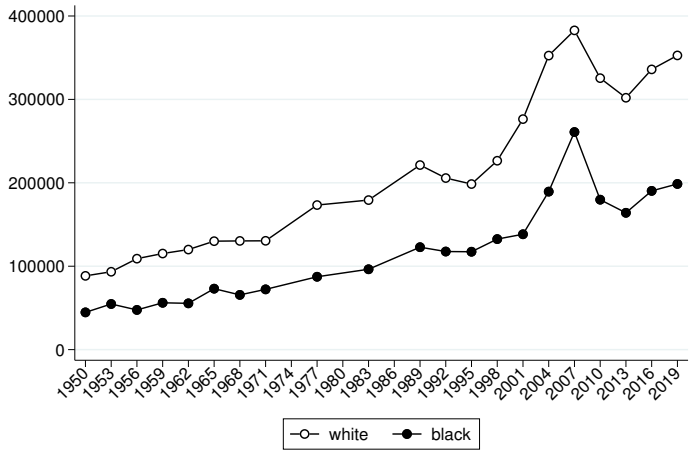

Notes: The graph shows the average house values of black and white homeowners over time.

Table 2.D.1 shows average portfolio shares of black and white households by decade.

Table 2.D.1: Portfolio shares by decade and race

decade	housing	equity + business	liq. assets + bonds	other other	housing debt	personal debt
			black			
1950	50.9	38.0	4.9	6.1	65.6	34.4
1960	59.6	28.6	5.6	6.2	70.3	29.7
1970	62.6	17.2	11.4	8.9	70.6	29.4
1980	65.1	6.6	6.9	21.4	62.9	37.1
1990	56.8	9.9	6.1	27.2	74.3	25.7
2000	58.8	11.6	5.3	24.3	77.7	22.3
2010	55.2	13.0	6.2	25.6	67.2	32.8
			white			
1950	31.4	55.6	10.2	2.7	77.8	22.2
1960	35.6	52.6	9.3	2.5	83.5	16.5
1970	42.9	43.4	10.5	3.1	81.9	18.1
1980	45.0	28.2	13.1	13.8	74.6	25.4
1990	40.2	30.4	10.1	19.2	78.9	21.1
2000	41.1	32.0	8.1	18.8	82.4	17.6
2010	36.4	34.7	7.7	21.2	80.4	19.6

Notes: The table shows the shares of several forms of assets in total assets (columns 2-5) and of debt in total debt (columns 6-7) by decade and race.

Figure 2.D.5 shows the ex- and intensive margins of non-housing debt-to-income ratios by race.

Figure 2.D.5: Ex- and intensive margins of non-housing debt-to-income ratios, by race

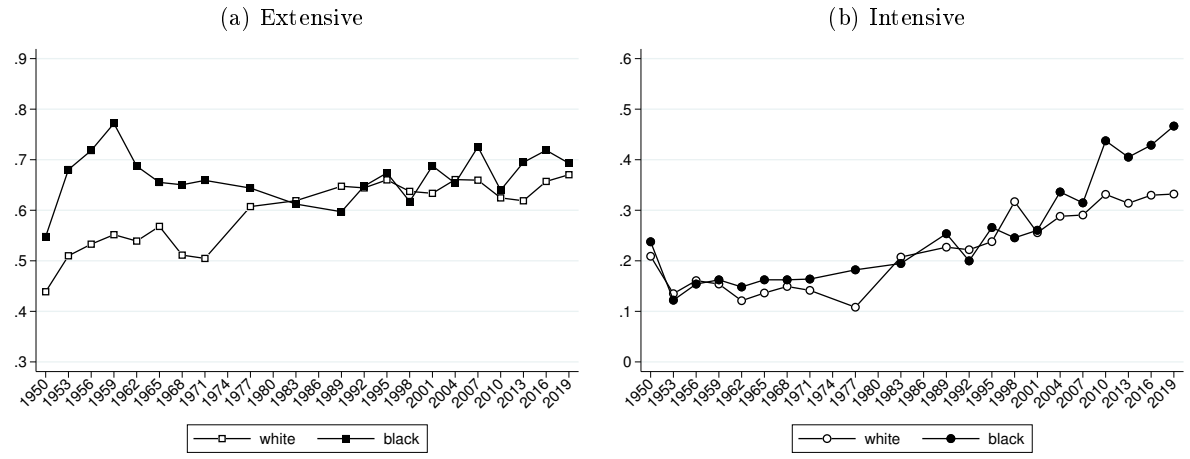

Notes: The graph shows the population shares (left) and non-housing debt-to-income ratios (right) of black and white households with positive non-housing debt.

Appendix 2.E Wealth effects in a simple life-cycle model

Two arguments against housing wealth effect are often put forward. The first is that housing is not wealth but a consumption good. If house prices rise, future housing consumption becomes more expensive and households effectively do not get wealthier. This intuition is derived by Sinai and Souleles (2005) in an infinite-horizon model with fixed housing consumption.[27] Key to their "neutrality" result for house price changes is the infinite housing tenure of agents. Introducing finite lifetimes, life-cycle variation in housing demand, contemporaneous ownership of housing by different generations, or imperfectly correlated local housing markets will imply that rising housing wealth triggers consumption responses of homeowners also in their model. This can be seen intuitively when taking a *Modigliani perspective* with a life-cycle model without bequests (Modigliani and Brumberg, 1954). In such a model, households will reduce housing consumption to zero at the end of their life, which implies that they will always realize capital gains from house price changes, so that the wealth effect arises naturally.

The second argument put forward against a large housing wealth effect is that under the permanent income hypothesis (PIH), marginal propensities to consume (MPCs) out of wealth are small. Again, it is crucial to take a *Modigliani perspective* for two reasons. First, housing makes up a large share of total wealth for the typical homeowner, so that house price changes have large effects on the lifetime budget constraint. Second, the shorter the remaining life span, the larger is the MPC out of wealth. The infinite-horizon abstraction is the polar case along both dimensions. First, human capital is large because of infinite working lives, so housing wealth accounts for only a small share of total household wealth. Second, the infinite horizon also implies that the remaining lifetime is infinite. In a life-cycle model with finite working lives and lifetime, MPCs out of wealth are substantially larger (Berger et al. 2018).

For our discussion, we rely on a life-cycle model adapted from Berger et al. (2018). To keep the model analytically tractable, we abstract from idiosyncratic income risk, borrowing constraints, and trading costs, but the consumption response to house price changes in our model is still quantitatively in line with the preferred estimate of Berger et al. (2018).

[27] This view is also prominently discussed in Case, Glaeser, and Parker (2000).

We consider the following economic environment. Households live for $J+1$ periods, have an exogenous income profile $\{y_j\}_{j=0}^{J}$, and aim at maximizing their lifetime utility from consumption. Instead of considering (net) household wealth, we split household wealth into housing assets h and financial assets (mortgage debt) d, of which households receive fixed initial endowments h_{-1} and d_0. Housing can be traded without frictions at price p_h each period and depreciates at rate δ. The mortgage interest rate is denoted by r. We abstract from bequests and assume that at the end of life, households sell their homes, repay their debt, and consume all available resources.

We assume that at each age j, households have a time-separable log utility function over a Cobb-Douglas composite of housing h_j and non-housing consumption c_j, $u(c,h) = \rho \log(c) + (1-\rho)\log(h)$. In this case, the household problem has a well-known solution with a constant expenditure share ρ for non-housing consumption and an optimal consumption path $c_j^* = c_0^*(\beta(1+r))^j$, where β denotes the time discount factor. The derivations are summarized at the end of this section. The level of the consumption path c_0^* is determined by total household wealth W, which is the sum of human capital Y, equal to the discounted incomes y_j at all ages j, and initial home equity E,

$$W = E + Y \qquad \text{with } Y = \sum_{j=0}^{J} y_j (1+r)^{-j} \text{ and } E = (1-\delta)p_h h_{-1} - (1+r)d_0, \tag{2.E.1}$$

multiplied by the MPC α and the optimal expenditure share ρ,

$$c_0^* = \alpha \rho W \qquad \text{with } \alpha = \frac{1-\beta}{1-\beta^{J+1}}. \tag{2.E.2}$$

It follows immediately that any change in home equity E from higher house prices p_h or lower debt levels d_0 will lead to an upward shift in the consumption profile. The optimal consumption dynamics, will however remain unaffected, as they only depend on the wedge between the time discount factor β and the interest rate r.[28]

Deriving the elasticity of the optimal consumption level c_0^* with respect to a persistent increase in house price p_h, we get a simple, intuitive expression (see equation (A.1) in Berger et al. 2018):

$$\frac{\partial c_0^*}{\partial p_h}\frac{p_h}{c_0^*} = \frac{(1-\delta)p_h h_{-1}}{W} = \theta_h, \tag{2.E.3}$$

where θ_h denotes the portfolio share of housing $(1-\delta)p_h h_{-1}$ in total wealth W. This elasticity for house price changes keeps human capital constant when changing the house price, thereby capturing a situation with rising house prices and stagnant incomes. Note that the formula applies to each point in the life cycle if lifetime J represents the *remaining* lifetime and the current period is interpreted as $j = 0$.

This simple expression for the elasticity states that the larger the exposure of household wealth W to house prices, the higher the elasticity of consumption with respect to house price shocks. Exposure to house prices in the model is determined by the portfolio share of housing in total wealth θ_h. This implies that the elasticity of consumption with respect to house price shocks is increasing in leverage, as leverage reduces W while leaving housing assets $(1-\delta)p_h h_{-1}$ unaffected so that θ_h increases.[29] Equation (2.E.3) therefore explains why households with a large exposure to the housing market responds more strongly to house price shocks.

When discussing the consumption response under the PIH, Berger et al. (2018) parametrize the expression for the consumption elasticity in equation (2.E.3) based on an infinite-horizon

[28] Key for this result is that we rule out potentially binding borrowing constraints.
[29] This may not extend to extreme cases such as underwater borrowers (Ganong and Noel 2020a).

abstraction that results in a small consumption response to house price shocks due to a very large value of human capital Y. When we consider finite working lives and a typical homeowner of age 51, we get an elasticity that is almost four times larger (0.18 vs. 0.05) with respect to house prices.[30] This consumption response is only slightly lower than the preferred empirical estimate of 0.23 by Berger et al. (2018). Key for the lower elasticity in our model is that we abstract from income risk and borrowing constraints that shape MPCs in the cross section. Berger et al. (2018) demonstrate that their richer model matches MPCs so that the theoretical consumption response approximates observed consumption responses to house price shocks very well.

These consumption dynamics induce corresponding debt dynamics according to the law of motion $d_{j+1} = (1+r)d_j - (y_j - c_j) + p_h(h_j - (1-\delta)h_{j-1})$. Future debt d_{j+1} depends positively on the current level of debt d_j, the current repayment $(y_j - c_j > 0)$ or extraction flow $(y_j - c_j < 0)$, and adjustments to the housing stock including depreciation $p_h(h_j - (1-\delta)h_{j-1})$. Iterating the law of motion forward, we get that, at any age $j+1$, the current debt level is simply the initial debt level d_0 plus the accumulated sum of repayment and extraction flows, housing adjustments, and accrued interest payments:

$$d_{j+1} = (1+r)^{j+1}d_0 + \sum_{s=0}^{j}(1+r)^{j-s}(c_s - y_s) + \sum_{s=0}^{j}(1+r)^{j-s}p_h(h_s - (1-\delta)h_{s-1}). \qquad (2.E.4)$$

It is important to acknowledge that our model is very stylized. In the following subsection, we discuss the simplifying assumptions made and how they might be relaxed. Yet despite its simplicity, this stylized model shows that the observed household behavior can be rationalized from a *Modigliani perspective*.

2.E.1 Discussion

In the model, households will reduce housing consumption after a positive house price shock, but housing wealth $(1-\delta)p_h h$ will increase nonetheless.[31] This result implies that our stylized model predicts that households will not upgrade to larger/better houses after a positive house price shock. A key reason is that the stylized model abstracts from borrowing constraints and adjustment costs.[32] In turn, the model predicts too much *downgrading*: households buy less/worse housing after a positive house price shock. Introducing trading and adjustment costs would allow us to more closely match the empirically observed patterns.

Moreover, the model abstracts from renters. Current renters constitute the pool of potential new owners who are affected by rising house prices. When house prices rise, households who switch from renting to owning have to pay more for a home of a given size. Hence, new homeowners will have to rely on additional debt to finance their home, buy a smaller house, or postpone homeownership. The data suggest that during the housing boom, many new homeowners relied on additional debt to finance their new home (Figure 2.C.3).

In our stylized environment, we do not consider ways in which extracted equity could be used other than for non-durable consumption. Empirical studies have found that home equity is also

[30] We follow Krebs, Kuhn, and Wright (2017) to estimate the human capital stock Y and focus on home equity as non-human wealth. We set the retirement age to 65, the mortgage interest rate to 5.5% (the average debt-weighted rate in 1992-2001 SCF data), and use 2001 SCF data as in Berger et al. (2018).

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

[32] Without borrowing constraints and adjustment costs, households react immediately to a positive shock to house prices and substitute away from housing. If, however, households are constrained, a shock that increases home equity slackens the constraint and allows them to upgrade. That upgrading households use (part of) their equity gain for the down payment of a new home has been discussed, for example, in Genesove and Mayer (1997) and Loewenstein (2018).

used for home improvements, the repayment of personal debt, or the foundation of a business (see Mian and Sufi 2011, Cloyne et al. 2019, Greenspan and Kennedy 2008). Finally, it should be noted that we abstract from other factors beyond house prices that have likely contributed to an increase in debt financing since the 1980s, such as lower mortgage interest rates and higher inflation, which raised the attractiveness of debt financing, falling mortgage transaction costs, the disappearing of mortgage prepayment penalties, or the rising costs of financing children's education (see, e.g., Bhutta and Keys 2016, Canner, Dynan, and Passmore 2002, Greenspan and Kennedy 2008, Cooper 2010).

2.E.2 Derivations

The agent's problem reads

$$\begin{aligned} \max_{\{c_j,h_j,d_{j+1}\}_{j=0}^{J}} \quad & \sum_{j=0}^{J} \beta^j \Big(\rho \log(c_j) + (1-\rho)\log(h_j) \Big) \\ s.t. \quad & c_j + p_h h_j - d_{j+1} = y_j - (1+r)d_j + (1-\delta)h_{j-1}p_h \\ & h_{-1},\ d_0 \qquad \text{given} \end{aligned} \tag{2.E.5}$$

First-order conditions deliver

$$\frac{1}{c_j}\rho p_h = (1-\rho)\frac{1}{h_j} + \beta\rho(1-\delta)p_h\frac{1}{c_{j+1}} \tag{2.E.6}$$

$$\frac{1}{c_j} = \beta(1+r)\frac{1}{c_{j+1}}. \tag{2.E.7}$$

From equation (2.E.7), we get the optimal path of consumption growth,

$$c_j = (\beta(1+r))^j c_0. \tag{2.E.8}$$

Using the Euler equation (2.E.7) in equation (2.E.6) delivers

$$\begin{aligned} \rho p_h &= (1-\rho)\frac{c_j}{h_j} + \beta\rho(1-\delta)p_h\frac{c_j}{c_{j+1}} \\ 1 &= \frac{1-\rho}{\rho}\frac{c_j}{p_h h_j} + \beta(1-\delta)(\beta(1+r))^{-1} \\ p_h h_j &= \frac{1-\rho}{\rho}c_j + \frac{1-\delta}{1+r}p_h h_j \\ p_h h_j &= \frac{1+r}{r+\delta}\frac{1-\rho}{\rho}c_j \end{aligned} \tag{2.E.9}$$

with the standard constant expenditure share result. Note that expenditures for housing are the user costs $\frac{r+\delta}{1+r}p_h h_j$. Combining equation (2.E.9) with the Euler equation delivers

$$p_h h_j = \frac{1+r}{r+\delta}\frac{1-\rho}{\rho}(\beta(1+r))^j c_0. \tag{2.E.10}$$

The law of motion for the debt level is

$$d_{j+1} = c_j - y_j + p_h h_j + (1+r)d_j - (1-\delta)h_{j-1}p_h. \tag{2.E.11}$$

Using this law of motion and plugging in recursively delivers

$$d_{j+1} = \sum_{s=0}^{j}(c_s - y_s)(1+r)^{j-s} + p_h h_j + \sum_{s=0}^{j-1} p_h h_s (r+\delta)(1+r)^{j-1-s} - (1+r)^j((1-\delta)h_{-1}p_h - (1+r)d_0). \tag{2.E.12}$$

For $j = J$, we get

$$d_{J+1} = \sum_{s=0}^{J}(c_s - y_s)(1+r)^{J-s} + p_h h_J + \sum_{s=0}^{J-1} p_h h_s (r+\delta)(1+r)^{J-1-s} - (1+r)^J((1-\delta)h_{-1}p_h - (1+r)d_0). \tag{2.E.13}$$

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

$$\begin{aligned}
d_{J+1}(1+r) - (1-\delta)p_h h_J &= (1+r)^{J+1}\left(\sum_{s=0}^{J}\frac{c_j - y_j}{(1+r)^s} + \frac{(1+r)p_h h_J - (1-\delta)p_h h_J}{(1+r)^{J+1}}\right. \\
&\quad \left. + \frac{1}{1+r}\sum_{s=0}^{J-1}\frac{p_h h_s}{(1+r)^s}(r+\delta) - \Big((1-\delta)h_{-1}p_h - (1+r)d_0\Big)\right) \\
\frac{d_{J+1}(1+r) - (1-\delta)p_h h_J}{(1+r)^{J+1}} &= \sum_{s=0}^{J}\frac{c_j - y_j}{(1+r)^s} + \frac{(r+\delta)p_h h_J}{(1+r)^{J+1}} \\
&\quad + \frac{r+\delta}{1+r}\sum_{s=0}^{J-1}\frac{p_h h_s}{(1+r)^s} - \Big((1-\delta)h_{-1}p_h - (1+r)d_0\Big) \\
\frac{d_{J+1}(1+r) - (1-\delta)p_h h_J}{(1+r)^{J+1}} &= \sum_{s=0}^{J}\frac{c_j}{(1+r)^s} - \overbrace{\sum_{s=0}^{J}\frac{y_j}{(1+r)^s}}^{=Y} \\
&\quad + \frac{r+\delta}{1+r}\sum_{s=0}^{J}\frac{p_h h_s}{(1+r)^s} - \underbrace{\Big((1-\delta)h_{-1}p_h - (1+r)d_0\Big)}_{=E} \\
\frac{d_{J+1}(1+r) - (1-\delta)p_h h_J}{(1+r)^{J+1}} &= \sum_{s=0}^{J}\frac{c_j}{(1+r)^s} + \frac{r+\delta}{1+r}\sum_{s=0}^{J}\frac{p_h h_s}{(1+r)^s} - (E+Y).
\end{aligned} \tag{2.E.14}$$

Under the optimal policy, it is always optimal that all resources are consumed in the last period, so that equity at the end of the life cycle is zero: $E' = (1-\delta)p_h h_J - d_{J+1}(1+r) = 0$. This implies that the left-hand side of equation (2.E.14) must be zero for the solution to be optimal, and we obtain

$$E + Y = \sum_{s=0}^{J}\frac{c_j}{(1+r)^s} + \frac{r+\delta}{1+r}\sum_{s=0}^{J}\frac{p_h h_s}{(1+r)^s}. \tag{2.E.15}$$

Now we plug in equations (2.E.8) and (2.E.10) and obtain

$$\begin{aligned}
\underbrace{E+Y}_{=W} &= \sum_{s=0}^{J}\frac{c_0(\beta(1+r))^s}{(1+r)^s} + \frac{r+\delta}{1+r}\sum_{s=0}^{J}\frac{\frac{1+r}{r+\delta}\frac{1-\rho}{\rho}(\beta(1+r))^s c_0}{(1+r)^s} \\
W &= c_0\sum_{s=0}^{J}\beta^s + \frac{1-\rho}{\rho}c_0\sum_{s=0}^{J}\beta^s \\
W &= c_0\frac{1-\beta^{J+1}}{1-\beta} + \frac{1-\rho}{\rho}c_0\frac{1-\beta^{J+1}}{1-\beta} \\
\underbrace{\frac{1-\beta}{1-\beta^{J+1}}}_{=\alpha}W &= \frac{1}{\rho}c_0 \\
\rho\alpha W &= c_0^*.
\end{aligned} \tag{2.E.16}$$

The law of motion from equation (2.E.4) follows directly from iterating equation (2.E.11):

$$d_{j+1} = \sum_{s=0}^{j}(c_s - y_s)(1+r)^{j-s} + \sum_{s=0}^{j}(p_h h_s - (1-\delta)p_h h_{s-1})(1+r)^{j-s} + (1+r)^{j+1}d_0. \tag{2.E.17}$$

Rearranging terms, we get the expression from equation (2.E.12) and plug in the result for the constant expenditure shares to obtain

$$\begin{aligned}
d_{j+1} &= \underbrace{\sum_{s=0}^{j} c_s(1+r)^{j-s}}_{\text{consumption costs}} - \underbrace{\sum_{s=0}^{j} y_s(1+r)^{j-s}}_{\text{income}} + \underbrace{p_h h_j}_{\text{current housing}} \\
&\quad + \underbrace{\sum_{s=0}^{j-1} p_h h_s \frac{r+\delta}{1+r}(1+r)^{j-s}}_{\text{user costs}} - (1+r)^j \underbrace{\Big((1-\delta)h_{-1}p_h - (1+r)d_0\Big)}_{\text{initial endowment}} \\
d_{j+1} &= \sum_{s=0}^{j} c_s(1+r)^{j-s} - \sum_{s=0}^{j} y_s(1+r)^{j-s} + p_h h_j - (1+r)^j(1-\delta)h_{-1}p_h \\
&\quad + \sum_{s=0}^{j-1} \frac{1-\rho}{\rho} c_s(1+r)^{j-s} + (1+r)^{j+1} d_0 \\
\underbrace{\frac{d_{j+1}}{(1+r)^j}}_{\substack{\text{present value} \\ \text{of debt}}} &= \overbrace{\sum_{s=0}^{j} \frac{c_s}{(1+r)^s} + \sum_{s=0}^{j-1} \frac{1-\rho}{\rho} \frac{c_s}{(1+r)^s}}^{\substack{\text{present value of} \\ \text{total expenditures}}} - \overbrace{\sum_{s=0}^{j} \frac{y_s}{(1+r)^s}}^{\substack{\text{present value} \\ \text{of income}}} \\
&\quad + \underbrace{\left(\frac{p_h h_j}{(1+r)^j} - (1-\delta)h_{-1}p_h\right)}_{\substack{\text{present value of} \\ \text{housing adjustments}}} + \underbrace{(1+r)d_0}_{\substack{\text{(present value)} \\ \text{initial debt}}} .
\end{aligned} \tag{2.E.18}$$

Appendix 2.F Financial fragility

In this section, we explore the consequences of the surge in debt-financed home equity extraction for financial stability. We will show that the sensitivity of households to income shocks has risen substantially as debt ratios have surged. The surge in home equity borrowing since the 1980s played an important role in this process. Financial fragility is a complex and multidimensional issue. We focus on household liquidity as one important dimension of financial risk. Liquidity has been emphasized in recent research as an important driver of household consumption decisions (Kaplan and Violante 2014) and mortgage defaults (Ganong and Noel 2020b). We quantify growing vulnerability using a *stress testing* approach.

Our analysis of the macroeconomic consequences of home-equity-based borrowing builds on the work of Mian and Sufi (2011) that explores the link between equity extraction and default rates in the crisis. In a similar spirit, Fuster, Guttman-Kenney, and Haughwout (2018) conduct a stress test for households based on Equifax CRISM data, shocking home equity positions. The latter paper focuses on a relatively short time period from 2005 to 2017. The long-run SCF+ data enable use to track the trends in financial fragility of the U.S. household sector over a long time period and demonstrate secular changes in macroeconomic financial fragility.

For our stress test scenario, we construct shocks that constrain the debt-servicing ability of households. Drops in income reduce liquidity and put households under *financial stress*. We employ the estimates on earnings losses following job displacement by Davis and von Wachter (2011). They document that earnings losses amount to 39% in the first year after displacement (Figure 4 in their paper). We thus let the income of the main wage earner of all households drop

by 39%.[33] We then consider a household to be under financial stress if the debt-service-to-income ratio exceeds 40% after the income shock.[34]

We report the share of households that will be under financial stress after receiving the shock and the *loan value at risk*, computed as the value of outstanding mortgage balances of all households under financial stress. Our results reflect a *shock potential* by applying the income shock to all households.[35] Figure 2.F.1 shows the value of outstanding mortgages which are at risk under our stress test scenario, relative to income. The right panel shows the corresponding share of households under financial stress. The figure shows that financial fragility was low and very stable until the onset of the post-1980 debt boom.

Figure 2.F.1: Stress testing household balance sheets

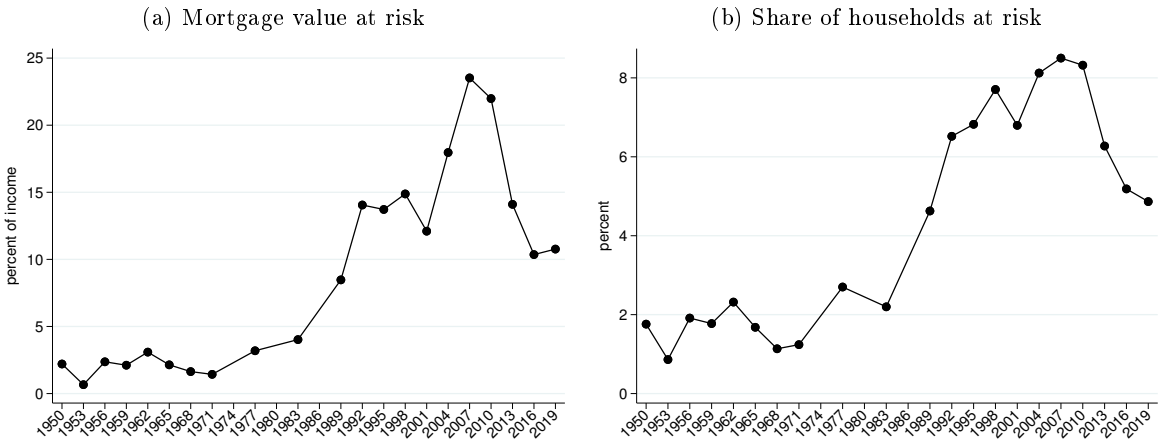

Notes: The left panel shows the value of mortgages at risk under our stress test scenario, relative to income. The right panel shows the share of households under financial stress.

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

[34] The value of 0.4 is used in the Financial Stability Reports of the Bank of England. We are grateful to Anil Kashyap for suggesting this source. The value of 0.4 also lies between the thresholds of 0.36 and 0.45 in the "eligibility matrix" used by the Federal National Mortgage Association (Fannie Mae).

[35] A quantification of the macroeconomic consequences will depend on general equilibrium effects and household default decisions, changes in consumption and saving behavior, and monetary and fiscal reactions to the initial shock. Characterizing and quantifying all these effects is beyond the scope of this exercise.

References

Adelino, Manuel, Antoinette Schoar, and Felipe Severino (2016). "Loan Originations and Defaults in the Mortgage Crisis: The Role of the Middle Class." In: *The Review of Financial Studies* 29.7, pp. 1635–1670.

— (2018). "The Role of Housing and Mortgage Markets in the Financial Crisis." In: *Annual Review of Financial Economics* 10.1, pp. 25–41.

Aladangady, Aditya (2017). "Housing Wealth and Consumption: Evidence from Geographically-Linked Microdata." In: *American Economic Review* 107.11, pp. 3415–46.

Albanesi, Stefania, Giacomo De Giorgi, and Jaromir Nosal (2017). *Credit Growth and the Financial Crisis: A New Narrative*. Working Paper.

Aliprantis, Dionissi and Daniel R. Carroll (2019). "What Is Behind the Persistence of the Racial Wealth Gap?" In: *Federal Reserve Bank of Cleveland Economic Commentary*.

Alvaredo, Facundo, Lucas Chancel, Thomas Piketty, Emmanuel Saez, and Gabriel Zucman (2018). *World Inequality Report 2018*. Harvard University Press.

Andersen, Henrik Yde and Søren Leth-Petersen (2021). "Housing Wealth or Collateral: How Home Value Shocks Drive Home Equity Extraction and Spending." In: *Journal of the European Economic Association* 19.1, pp. 403–440.

Aruoba, S. Boragan, Ronel Elul, and Sebnem Kalemli-Ozcan (2019). *How Big is the Wealth Effect? Decomposing the Response of Consumption to House Prices*. Working Papers 19-6. Federal Reserve Bank of Philadelphia.

Avery, Christopher and Sarah Turner (2012). "Student Loans: Do College Students Borrow Too Much–Or Not Enough?" In: *Journal of Economic Perspectives* 26.1, pp. 165–92.

Bartscher, Alina K., Moritz Kuhn, Moritz Schularick, and Paul Wachtel (2021). *Monetary Policy and Racial Inequality*. CESifo Working Paper Series 8877. CESifo.

Bauluz, Luis and Timothy Meyer (2021). *The Great Divergence: Intergenerational Wealth Inequality in the U.S. and France*. Working Paper.

Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra (2018). "House Prices and Consumer Spending." In: *The Review of Economic Studies* 85.3, pp. 1502–1542.

Bernanke, Ben S. (2005). *The Global Saving Glut and the U.S. Current Account Deficit*. Speech 77. Board of Governors of the Federal Reserve System.

Bhutta, Neil (2015). "The Ins and Outs of Mortgage Debt during the Housing Boom and Bust." In: *Journal of Monetary Economics* 76, pp. 284–298.

Bhutta, Neil and Benjamin J. Keys (2016). "Interest Rates and Equity Extraction during the Housing Boom." In: *American Economic Review* 106.7, pp. 1742–74.

Boar, Corina, Denis Gorea, and Virgiliu Midrigan (forthcoming). "Liquidity Constraints in the U.S. Housing Market." In: *Review of Economic Studies*.

Bricker, Jesse et al. (2017). "Changes in U.S. Family Finances from 2013 to 2016: Evidence from the Survey of Consumer Finances." In: *Federal Reserve Bulletin* 103.3.

Calomiris, Charles W., Stanley D. Longhofer, and William Miles (2013). "The Housing Wealth Effect: The Crucial Roles of Demographics, Wealth Distribution and Wealth Shares." In: *Critical Finance Review* 2.1, pp. 49–99.

Campbell, John Y. and Joao F. Cocco (2007). "How do House Prices Affect Consumption? Evidence from Micro Data." In: *Journal of monetary Economics* 54.3, pp. 591–621.

Canner, Glenn B., Karen E. Dynan, and Wayne Passmore (2002). "Mortgage Refinancing in 2001 and Early 2002." In: *Federal Reserve Bulletin* Dec, pp. 469–481.

Canner, Glenn B., James T. Fergus, and Charles A. Luckett (1988). "Home Equity Lines of Credit." In: *Federal Reserve Bulletin* 74, p. 361.

Case, Karl E., Edward L. Glaeser, and Jonathan A. Parker (2000). "Real estate and the macroeconomy." In: *Brookings Papers on Economic Activity* 2000.2, pp. 119–162.

Chen, Hui, Michael Michaux, and Nikolai Roussanov (2020). "Houses as ATMs: Mortgage Refinancing and Macroeconomic Uncertainty." In: *The Journal of Finance* 75.1, pp. 323–375.

Christelis, Dimitris, Dimitris Georgarakos, Tullio Jappelli, Luigi Pistaferri, and Maarten van Rooij (2019). *Wealth Shocks and MPC Heterogeneity*. Working Paper 25999. National Bureau of Economic Research.

Cloyne, James, Kilian Huber, Ethan Ilzetzki, and Henrik Kleven (2019). "The Effect of House Prices on Household Borrowing: A New Approach." In: *American Economic Review* 109.6, pp. 2104–36.

Coibion, Olivier, Yuriy Gorodnichenko, Marianna Kudlyak, and John Mondragon (2020). "Greater Inequality and Household Borrowing: New Evidence from Household Data." In: *Journal of the European Economic Association*.

Cooper, Daniel (2010). *Did Easy Credit Lead to Overspending? Home Equity Borrowing and Household Behavior in the Early 2000s*. Working Paper.

Davis, Steven J. and Till von Wachter (2011). "Recessions and the Costs of Job Loss." In: *Brookings Papers on Economic Activity*.

De Stefani, Alessia (2018). "Debt, Inequality and House Prices: Explaining the Dynamics of Household Borrowing Prior to the Great Recession." In: *Journal of Housing Economics*.

— (2020). "House Price History, Biased Expectations, and Credit Cycles: The Role of Housing investors." In: *Real Estate Economics*.

DeFusco, Anthony A., Charles G. Nathanson, and Eric Zwick (2017). *Speculative Dynamics of Prices and Volume*. Working Paper 23449. National Bureau of Economic Research.

Dettling, Lisa J., Sebastian Devlin-Foltz, Jacob Krimmel, Sarah Pack, and Jeffrey P: Thompson (2015). *Comparing Micro and Macro Sources for Household Accounts in the United States: Evidence from the Survey of Consumer Finances*. Working Paper.

Di Maggio, Marco and Amir Kermani (2017). "Credit-Induced Boom and Bust." In: *The Review of Financial Studies* 30.11, pp. 3711–3758.

Di Maggio, Marco, Amir Kermani, and Christopher J. Palmer (2019). "How Quantitative Easing Works: Evidence on the Refinancing Channel." In: *The Review of Economic Studies* 87.3, pp. 1498–1528.

Díaz-Giménez, Javier, Andy Glover, and José-Víctor Ríos-Rull (2011). "Facts on the Distributions of Earnings, Income, and Wealth in the United States: 2007 Update." In: *Federal Reserve Bank of Minneapolis Quarterly Review* 34.

Doepke, Matthias and Martin Schneider (2006). "Inflation and the Redistribution of Nominal Wealth." In: *Journal of Political Economy* 114.6, pp. 1069–1097.

Duca, John V. and Anil Kumar (2014). "Financial Literacy and Mortgage Equity Withdrawals." In: *Journal of Urban Economics* 80, pp. 62–75.

Duca, John V., John Muellbauer, and Anthony Murphy (forthcoming). "What Drives House Price Cycles? International Experience and Policy Issues." In: *Journal of Economic Literature*.

Dynan, Karen E. and Donald L. Kohn (2007). *The Rise in US Household Indebtedness: Causes and Consequences*. Working Paper.

Elia, Charles E. (June 1981). "Growing Use of Home Equity to Raise Cash Tends to Boost Interest Rates, Analyst Warns." In: *Wall Street Journal*, p. 53.

Favara, Giovanni and Jean Imbs (2015). "Credit Supply and the Price of Housing." In: *American Economic Review* 105.3, pp. 958–92.

Fernández, Raquel (2013). "Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century." In: *American Economic Review* 103.1, pp. 472–500.

Fetter, Daniel K. (2014). "The Twentieth-Century Increase in U.S. Home Ownership: Facts and Hypotheses." In: *Housing and Mortgage Markets in Historical Perspective*. University of Chicago Press. Chap. 10, pp. 329–350.

Fligstein, Neil, Orestes P. Hastings, and Adam Goldstein (2017). "Keeping up with the Joneses: How Households Fared in the Era of High Income Inequality and the Housing Price Bubble, 1999–2007." In: *Socius* 3.

Foote, Christopher L., Lara Loewenstein, and Paul S. Willen (2016). *Cross-sectional Patterns of Mortgage Debt During the Housing Boom: Evidence and Implications*. Tech. rep. National Bureau of Economic Research.

Fuster, Andreas, Benedict Guttman-Kenney, and Andrew F. Haughwout (2018). "Tracking and Stress-testing U.S. Household Leverage." In: *Economic Policy Review* 24-1, pp. 35–63.

Ganong, Peter and Pascal Noel (2020a). "Liquidity versus Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession." In: *American Economic Review* 110.10, pp. 3100–3138.

Ganong, Peter and Pascal J. Noel (2020b). *Why Do Borrowers Default on Mortgages? A New Method For Causal Attribution*. Working Paper 27585. National Bureau of Economic Research.

Genesove, David and Christopher J. Mayer (1997). "Equity and Time to Sale in the Real Estate Market." In: *The American Economic Review* 87.3, pp. 255–269.

Greenspan, Alan and James Kennedy (2008). "Sources and Uses of Equity Extracted from Homes." In: *Oxford Review of Economic Policy* 24.1, pp. 120–144.

Guren, Adam M., Alisdair McKay, Emi Nakamura, and Jón Steinsson (2021). "Housing Wealth Effects: The Long View." In: *The Review of Economic Studies* 88.2, pp. 669–707.

Haughwout, Andrew, Donghoon Lee, Joseph S. Tracy, and Wilbert Van der Klaauw (2011). "Real Estate Investors, the Leverage Cycle, and the Housing Market Crisis." In: *FRB of New York Staff Report* 514.

Henriques, Alice M and Joanne W Hsu (2014). "Analysis of wealth using micro-and macrodata: A comparison of the Survey of Consumer Finances and Flow of Funds accounts." In: *Measuring Economic Sustainability and Progress*. University of Chicago Press, pp. 245–274.

Hoffmann, Mathias and Iryna Stewen (2019). "Holes in the Dike: The Global Savings Glut, U.S. House Prices, and the Long Shadow of Banking Deregulation." In: *Journal of the European Economic Association*.

Hurst, Erik and Frank Stafford (2004). "Home Is Where the Equity Is: Mortgage Refinancing and Household Consumption." In: *Journal of Money, credit and Banking*, pp. 985–1014.

Iacoviello, Matteo (2005). "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle." In: *American Economic Review* 95.3, pp. 739–764.

Jappelli, Tullio, Marco Pagano, and Marco Di Maggio (2013). "Households' indebtedness and financial fragility." In: *Journal of Financial Management, Markets and Institutions* 1, pp. 23–46.

Jordà, Òscar (2005). "Estimation and Inference of Impulse Responses by Local Projections." In: *American Economic Review* 95.1, pp. 161–182.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor (2013). "When Credit Bites Back." In: *Journal of Money, Credit and Banking* 45, pp. 3–28.

— (2017). "Macrofinancial History and the New Business Cycle Facts." In: *NBER Macroeconomics Annual* 31.1, pp. 213–263.

Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante (2020). "The Housing Boom and Bust: Model Meets Evidence." In: *Journal of Political Economy* 128.9, pp. 3285–3345.

Kaplan, Greg and Giovanni L. Violante (2014). "A Model of the Consumption Response to Fiscal Stimulus Payments." In: *Econometrica* 82.4, pp. 1199–1239.

Kaplan, Greg, Giovanni L. Violante, and Justin Weidner (2014). "The Wealthy Hand-to-Mouth." In: *Brookings Papers on Economic Activity* 45.1, pp. 77–153.

Klyuev, Vladimir and Paul Mills (2007). "Is Housing Wealth an "ATM"? The Relationship between Household Wealth, Home Equity Withdrawal, and Saving Rates." In: *IMF Staff Papers* 54.3, pp. 539–561.

Kovacs, Agnes and Patrick Moran (2020). *Breaking the Commitment Device: The Effect of Home Equity Withdrawal on Consumption, Saving, and Welfare*. Working Paper.

Kowalewski, Kim J. (1987). "Home Equity Lines: Characteristics and Consequences." In: *Economic Commentary*.

Krebs, Tom, Moritz Kuhn, and Mark Wright (2017). "Under-insurance in Human Capital Models with Limited Enforcement." In: *Review of Economic Dynamics* 25. Special Issue on Human Capital and Inequality, pp. 121–150.

Kuhn, Moritz and José-Vıctor Rıos-Rull (2016). "2013 Update on the US Earnings, Income, and Wealth Distributional Facts: A View from Macroeconomics." In: *Federal Reserve Bank of Minneapolis Quarterly Review* 37.1.

Kuhn, Moritz, Moritz Schularick, and Ulrike I. Steins (2020). "Income and Wealth Inequality in America, 1949–2016." In: *Journal of Political Economy* 128.9, pp. 3469–3519.

Kumhof, Michael, Romain Rancière, and Pablo Winant (2015). "Inequality, Leverage, and Crises." In: *American Economic Review* 105.3, pp. 1217–45.

LaCour-Little, Michael, Eric Rosenblatt, and Vincent Yao (2010). "Home Equity Extraction by Homeowners: 2000–2006." In: *Journal of Real Estate Research* 32.1, pp. 23–46.

Loewenstein, Lara (2018). *Consumption of Housing During the 2000s Boom: Evidence and Theory*. Federal Reserve Bank of Cleveland Working Paper.

Looney, Adam and Constantine Yannelis (2015). "A Crisis in Student Loans?: How Changes in the Characteristics of Borrowers and in the Institutions They Attended Contributed to Rising Loan Defaults." In: *Brookings Papers on Economic Activity* 2015.2, pp. 1–89.

Lusardi, Annamaria, Olivia S. Mitchell, and Noemi Oggero (2018). "The Changing Face of Debt and Financial Fragility at Older Ages." In: *AEA Papers and Proceedings* 108, pp. 407–11.

— (Dec. 2020). *Understanding Debt in the Older Population*. Working Paper 28236. National Bureau of Economic Research.

Maki, Dean M (2001). "Household Debt and the Tax Reform Act of 1986." In: *American Economic Review* 91.1, pp. 305–319.

— (1996). "Portfolio Shuffling and Tax Reform." In: *National Tax Journal*, pp. 317–329.

Mian, Atif, Kamalesh Rao, and Amir Sufi (Sept. 2013). "Household Balance Sheets, Consumption, and the Economic Slump." In: *The Quarterly Journal of Economics* 128.4, pp. 1687–1726.

Mian, Atif, Ludwig Straub, and Amir Sufi (2021a). "Indebted Demand." In: *The Quarterly Journal of Economics*.

— (2021b). *The Saving Glut of the Rich*. Working Paper.

Mian, Atif and Amir Sufi (2009). "The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis." In: *The Quarterly Journal of Economics* 124.4, pp. 1449–1496.

— (2011). "House Prices, Home Equity-Based Borrowing, and the U.S. Household Leverage Crisis." In: *American Economic Review* 101.5, pp. 2132–56.

— (2014). *House Price Gains and US Household Spending from 2002 to 2006*. Tech. rep. National Bureau of Economic Research.

— (2017). "Household Debt and Defaults from 2000 to 2010: The Credit Supply View." In: *Evidence and Innovation in Housing Law and Policy*. Ed. by Lee Anne Fennell and Benjamin J. Keys. Cambridge University Press, pp. 257–288.

— (2021). "Credit Supply and Housing Speculation." In: *The Review of Financial Studies*.

Michel, Jean-Baptiste et al. (2011). "Quantitative Analysis of Culture Using Millions of Digitized Books." In: *Science* 331.6014, pp. 176–182.

Modigliani, Franco and Richard Brumberg (1954). "Utility Analysis and the Consumption Function: An Interpretation of Cross-section Data." In: *Post-keynesian economics* 1, pp. 338–436.

Pfeffer, Fabian T., Robert F. Schoeni, Arthur Kennickell, and Patricia Andreski (2016). "Measuring Wealth and Wealth Inequality: Comparing Two U.S. Surveys." In: *Journal of Economic and Social Measurement* 41.2, pp. 103–120.

Piketty, Thomas and Emmanuel Saez (2003). "Income Inequality in the United States, 1913-1998." In: *The Quarterly Journal of Economics* 118.1, pp. 1–39.

Rajan, Raghuram G. (2011). *Fault Lines: How Hidden Fractures Still Threaten the World Economy*. Princeton University press.

Rappaport, Jordan (2007). "A Guide to Aggregate House Price Measures." In: *Federal Reserve Bank of Kansas City Economic Review*, pp. 41–71.

Saez, Emmanuel and Gabriel Zucman (2016). "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data." In: *The Quarterly Journal of Economics* 131.2, pp. 519–578.

Sinai, Todd and Nicholas S. Souleles (2005). "Owner-Occupied Housing as a Hedge Against Rent Risk." In: *The Quarterly Journal of Economics* 120.2, pp. 763–789.

Story, Louise (2008). "Home Equity Frenzy Was a Bank Ad Come True." In: *The New York Times*.

Wolff, Edward N. (2010). *Recent Trends in Household Wealth in the United States: Rising Debt and the Middle-Class Squeeze – An Update to 2007*. Working Paper.

Zinman, Jonathan (2015). "Household Debt: Facts, Puzzles, Theories, and Policies." In: *Annual Review of Economics* 7.1, pp. 251–276.

Chapter 3

The College Wealth Divide: Education and Inequality in America, 1956-2016

with Moritz Kuhn and Moritz Schularick

Using new long-run microdata, this chapter studies wealth and income trends of college and non-college households in the United States since 1956. We document the emergence of a substantial *college wealth premium* since the 1980s, which is considerably larger than the *college income premium*. Over the past four decades, the wealth of American households with a college-educated head has tripled. By contrast, the wealth of non-college households has barely grown in real terms over the same period. Part of the rising wealth gap can be traced back to systematic portfolio differences between college and non-college households that give rise to different exposures to asset price changes. Non-college households have a lower exposure to the equity market and have profited much less from the recent surge in the stock market. We also discuss the importance of financial literacy and business ownership for the increase in wealth inequality between college and non-college households.

3.1 Introduction

It is a well-documented fact that the college wage premium has increased substantially since the 1980s (see e.g. Levy and Murnane 1992, Katz and Autor 1999, Goldin and Katz 2007). This trend can be traced back to differences in the growth of demand and supply of college-educated workers that are driven by skill-biased technical change, socio-demographic factors, and institutional features (Card and Lemieux 2001, Fortin 2006). Recent work has begun to analyze the relationship between college education and wealth inequality (Emmons, Kent, and Ricketts 2018, Pfeffer 2018), and demonstrated an increasing association between college education and wealth.

In this chapter, we take a long-run perspective on college and non-college income and wealth over almost the entire post-WWII period. We use a novel household-level dataset, the "SCF+", which combines the post-1983 Survey of Consumer Finances with data from historical surveys going back to 1949. To ensure consistent coding of education groups, our analysis starts in 1956. Kuhn, Schularick, and Steins (2020) have harmonized the data across the historical survey waves. The combined data provide long-run household-level information on income, assets, debt, and demographics. The SCF+ closes an important gap, as high-quality microdata were not available over longer time horizons before. For instance, the Panel Study of Income Dynamics (PSID), which is one of the most important sources of household-level wealth data in the U.S., only

includes questions on family wealth since 1984 (see Pfeffer et al. 2016).

Our analysis confirms a strong increase in the college income premium since the early 1980s. The average income of households with a college-educated head has increased by about 50% in real terms since then. However, the increase of the college income premium is dwarfed by that of the college wealth premium. The wealth of college households has increased by a factor of 3 between 1983 and 2016, while non-college wealth has barely grown at all. A substantial part of the widening of the wealth gap between college and non-college households is driven by strong wealth gains within the top 10% of the wealth distribution. Moreover, the share of non-college households making it to the top 10% of the wealth distribution has declined over time. We also document that households with two college-educated spouses have enjoyed particularly large gains in wealth. However, this trend is not driven by assortative matching, but by the overall growth in college education. Consistent with previous findings of Eika, Mogstad, and Zafar (2019), assortative mating appears to have *decreased* among college graduates over time.

An important question raised by our findings is why the ratio of college to non-college wealth has grown so much more than the income gap. The workhorse economic models of wealth accumulation imply a tight co-movement of income and wealth differences, as income is the sole determinant of wealth. We demonstrate that college and non-college households exhibit systematic differences not only in the size of their asset holdings, but also in the composition of their portfolios. College households own a higher share of stocks and mutual funds. As a consequence, college and non-college households are differentially exposed to asset price changes. College households reap disproportionately high capital gains during stock market booms. Importantly, such capital gains are unrelated to income. This is consistent with the de-coupling in the evolution of the college income and wealth advantage since the 1980s (see also Kuhn, Schularick, and Steins 2020).[1] We also find some indication that business ownership matters for the increase of the college wealth premium, especially since the late 1990s. The fact that the college wealth advantage is associated with equity holdings and business ownership may be related to higher levels of financial literacy and entrepreneurial skills among college households. We discuss the role of these factors and their potential to affect wealth via portfolio composition and differential returns.

While the chapter documents a sharply rising college wealth premium in recent decades, it is important to note that causality can run in both directions. College graduates may hold more wealth due to their higher educational attainment, but there is also evidence that it is easier to obtain college degrees when coming from a wealthy family. Wealthy families can afford more investment in their offspring's educational careers. The cost of college has increased considerably since the 1980s, which constitutes an obstacle for children from poorer households (see e.g. Haveman and Smeeding 2006). Beyond providing a basis for *inter vivos* transfers, wealth may have an insurance function as a "safety net" (Pfeffer 2018). In this sense, it can work as a "catalyzer", facilitating human capital investment in early life. This will typically lead to higher wealth, which may be augmented by gifts and bequests. The process can accumulate across generations, creating a succession of college-educated households with ever more wealth (Pfeffer and Killewald 2018).[2] Our SCF+ data consist of repeated cross sections, and therefore remain silent on inter-generational wealth links. At the same time, the wealth information in the PSID is less detailed than in the SCF+, and in particular the coverage of wealth at the top is not comprehensive (Pfeffer et al. 2016). New data sources are needed to address these questions.

[1] Note that total household income in the SCF+ is defined net of capital gains.

[2] Parental income and wealth may even affect educational outcomes beyond assistance to tap one's full skill potential. Looking at the income of U.S. households, Reeves and Howard (2013) find evidence of "glass floors" in educational outcomes: Children from high-income households tend to do better in terms of education and income than their skills would suggest. The authors stress that wealth is likely to play an important role beyond income. Certainly, the persistence of education and wealth across generations and their interaction are important topics for further research.

The chapter begins with a description of the data in Section 3.2. In Section 3.3, we present our empirical results. Section 3.4 focuses on the role of asset prices and business ownership for college wealth growth. Section 3.5 discusses potential transmission mechanisms, and Section 3.6 concludes.

3.2 Data

Our analysis is based on a newly compiled resource for inequality research, the SCF+. The modern Survey of Consumer Finances (SCF) is conducted every three years by the U.S. Federal Reserve Board (see Bricker et al. 2017). It is one of the most widely used datasets for the study of distributional issues in the United States. The modern waves cover the period since 1983. However, a predecessor of the modern surveys was conducted at an annual frequency by the Survey Research Center of the University of Michigan from 1947 to 1971, and again in 1977. Based on the original codebooks, Kuhn, Schularick, and Steins (2020) extract the historical data. They match and harmonize variables across the historic and modern waves to create rich microdata that allow to study the joint distribution of income and wealth, along with key demographic variables, over the period from 1949 to 2016. Bartscher et al. (2020) use the data to examine the post-war U.S. household debt boom. Following these papers, we pool the annual historic waves over three-year windows.

Missing data in the old waves were inferred by using multiple imputation methods like predictive mean matching (cp. Schenker and Taylor 1996), and historical data are re-weighted to account for non-response at the top of the income and wealth distribution. These adjustments are described in detail in Kuhn, Schularick, and Steins (2020). To assure representativeness along socio-demographic dimensions, the data were re-weighted to match demographic targets from the U.S. Current Population Survey (CPS) and the Decennial Census. Specifically, the data were post-stratified to match the age structure of the population, the share of households with a black household head, the share of households whose head has at least obtained some college education, and the homeownership rate.[3] For the new waves, the survey weights and data are the ones provided on the website of the U.S. Federal Reserve Board. The only amendment we made is to post-stratify the original 1983 weights to match the CPS homeownership rate. This is done for better consistency with the modern waves, which match the homeownership rate closely.

The key advantage of the dataset is that it combines rich information on economic and financial data with key socio-demographic variables. Kuhn, Schularick, and Steins (2020) exploit this feature of the data to study another key stratifying dimension of inequality, namely race. They find that income and wealth gaps between black and white households have hardly narrowed since the pre-civil rights era. The median black household only had about half the income and a tenth of the wealth of the median white household. By 2016, the income ratio of the median black to median white household has only increased by 10 percentage points, and the wealth gap has remained almost unchanged. We will abstain from a joint analysis of education and race due to low numbers of observations when slicing the data along both dimensions in the early years.

Total household income in the SCF+ data includes income from wages and salaries, professional practice and self employment, rental income, interest, dividends, business and farm income, as well as transfer payments. Assets comprise liquid assets (certificates of deposit, checking, saving, call and money market accounts), housing and other real estate (net of debt), bonds, stocks, mutual funds, corporate and non-corporate equity, and defined-contribution retirement accounts. Total debt sums housing debt on primary residences, car loans, education loans, loans for consumer durables, other non-housing debt, as well as credit card debt. Wealth is computed

[3] Throughout the chapter, demographic information will always refer to the household head, if not otherwise stated. In the case of married couples, the household head is typically male.

Figure 3.1: Comparison to census data

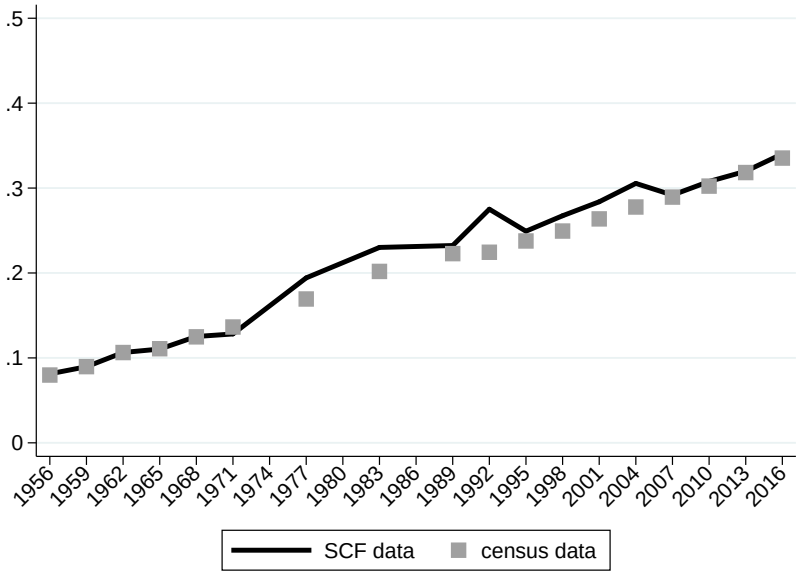

Notes: The figure shows the share of households with a college-educated head in the SCF+ data in comparison to the share obtained from the U.S. Current Population Survey (CPS) for the period from 1962 to 2016 and from the U.S. Decennial Census for 1950 and 1960. Intermediate data points were obtained by linear interpolation.

as total assets net of total debt. All monetary variables were transformed to 2016 dollars using the U.S. consumer price index (CPI) for all urban consumers from the Macrohistory Database (Jordà, Schularick, and Taylor 2017). This is also the source for the stock price data used in Section 3.4. Figure 3.1 compares the share of households headed by a college graduate in the SCF+ with data from the CPS, which are available since 1962. For the earlier periods, we rely on linearly interpolated data from the Decennial Census. Throughout the chapter, college households will be defined as households whose head has obtained at least a bachelor's degree. Householders with "some college" will be included in the group of non-college households.[4] The distinction between households with some college versus a college degree was not made in the earliest surveys, and there are notable differences in the portfolios and incomes of these groups (see table 3.B.1). Therefore, we decided to discard the first two three-year windows, and let our sample begin with the 1956 window (1954-1956).

While a close match to targeted census shares is a good test of the re-weighting procedure, it does not necessarily imply that the aggregated microdata match macroeconomic variables. Yet Kuhn, Schularick, and Steins (2020) demonstrate that the SCF+ data closely match aggregate trends in income, wealth, housing, financial and non-financial assets, as well as housing and non-housing debt. In addition, they demonstrate that the data exhibit a close fit to top income shares from Piketty and Saez (2003) using IRS tax data, and top wealth shares from Saez and Zucman (2016) using IRS data and the capitalization method.

3.3 Six decades of college income and wealth premia

Going beyond previous research, the SCF+ data allow us to document income and wealth differences between college and non-college households over the long run. As discussed above, much of the previous literature has focused on wage differences between college and non-college indi-

[4] As noted, the re-weighting was done based on the share of households with *at least some* college. A comparison of the SCF+ and census data with respect to this measure is provided in Figure 3.A.1.

viduals. Instead of looking at wages at the individual level, we consider total income and wealth at the household level, with a particular focus on the college wealth premium, the ratio of college to non-college wealth.

3.3.1 Income and wealth growth

Figure 3.2 shows the development of average household income and wealth for college and non-college households. The two groups evolved similarly until the 1970s, and have diverged afterwards. In Figure 3.2, we normalize the data to 1971 to track the divergence since the 1970s. The left panel reveals that income has grown at, by and large, similar rates for both groups until the 1970s. Afterwards, the real income of non-college households stagnated, while the real income of college households has risen by around 50%. In other words, our data confirm a secular rise of the college income premium.

The differential growth of college and non-college income is considerable, but it is dwarfed by the discrepancy in wealth. Just as income, wealth has evolved similarly for both groups until the 1970s, and stagnated for non-college households afterwards. The only exception is the period prior to the financial crisis in 2008, when non-college households increased their wealth to around 1.5 times its 1971 level. Consistent with the results of Kuhn, Schularick, and Steins (2020), this was mainly due to the short-lived effects of the house price boom in the 2000s. As we will document below, housing constitutes a particularly large share of total wealth for non-college household (see Figure 3.12). While non-college households were treading water in terms of wealth, college households have increased their net worth by a factor of three compared to 1971.

As average wealth has increased by more than average income for both college and non-college households, wealth-to-income ratios have expanded as well.[5] However, the increase has been much larger for college households due to their massive surge in wealth. Their wealth-to-income ratio has roughly doubled, from around 4.4 in 1971 to 8.5 in 2016, as can be seen in Figure 3.3. The corresponding growth for non-college households was only 26%, from around 3.7 to 4.7. Figure 3.A.2 in the Appendix shows that the gap between the two groups is somewhat reduced when excluding pension wealth. Yet the difference still remains substantial: While college households would still have experienced an increase in their average wealth-to-income

Figure 3.2: Wealth and income levels

Notes: The figure shows the average wealth and income of households with and without a college-educated head over time, normalized by each group's level in 1971.

[5]The ratios reported throughout the chapter are ratios of averages (as opposed to averages of ratios).

Figure 3.3: Wealth-to-income ratios

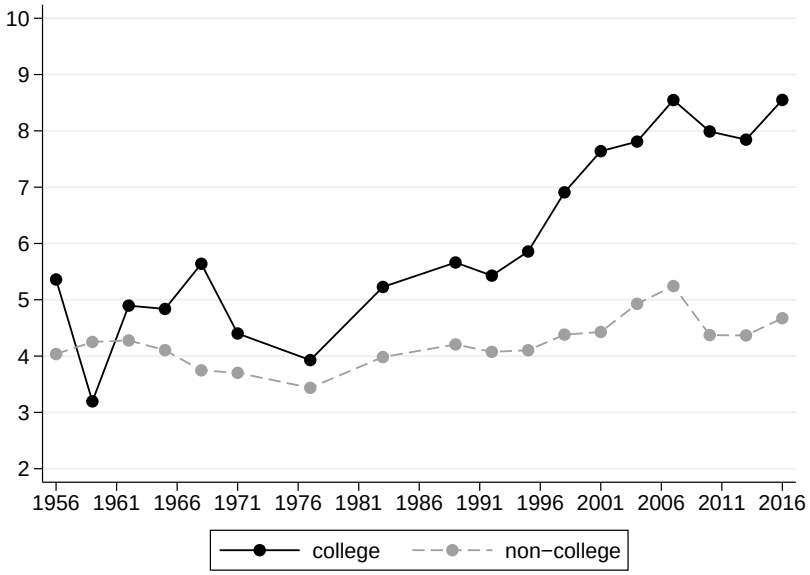

Notes: The figure shows the ratio of average net wealth to average income among households with and without a college-educated head over time.

ratio by a factor of approximately 1.6 (from 4.4 to 7.1) without pensions, non-college households would not have experienced any increase in wealth relative to income, apart from the house price boom period prior to 2007.

3.3.2 Decomposing wealth growth

In the next step, we explore wealth growth for three different wealth groups. Within each of these wealth groups, we distinguish between college and non-college households. In the middle group, which we define as the 50th to 90th percentile of the wealth distribution, college and non-college wealth has to co-move closely by construction, as it is limited both from below and above. What can change for this group is the share of college and non-college households who belong to the group. Figure 3.4 shows wealth growth for all three wealth groups, stratified by education. It reveals that the widening of the college wealth gap has been pronounced within the top 10% of the aggregate wealth distribution, whereas college and non-college households have evolved similarly within the bottom 50% and the middle 50-90%.[6] As Kuhn, Schularick, and Steins (2020) point out, the top 10% of the aggregate wealth distribution are more heavily invested in equity and business wealth. We will discuss the role of these factors in more detail in Sections 3.4 and 3.5. The graphs for the different parts of the aggregate wealth distribution in Figure 3.5 relate to Figure 3.2b via the following decomposition:

$$\frac{\overline{W}_{e,t}}{\overline{W}_{e,71}} = \frac{\sum_{i=1}^{3} s_{e,i,t}\overline{W}_{e,i,t}}{\overline{W}_{e,71}} = \sum_{i=1}^{3} s_{e,i,t}\frac{\overline{W}_{e,i,71}}{\overline{W}_{e,71}}\frac{\overline{W}_{e,i,t}}{\overline{W}_{e,i,71}} \qquad e = c, nc\,. \tag{3.1}$$

$\overline{W}_{e,t}$ denotes average wealth of education group $e \in \{c, nc\}$ at time t, where c means college and nc means non-college. The index $i \in \{1, 2, 3\}$ refers to the three groups of the aggregate wealth distribution, and $s_{e,i,t}$ is the share of households in education group e and wealth group i out of all households in education group e at time t. Hence, the widening of the college wealth gap

[6] By contrast, the college wealth gap has increased in all parts of the aggregate income distribution. The corresponding decomposition is available from the authors upon request.

Figure 3.4: Wealth growth along the wealth distribution

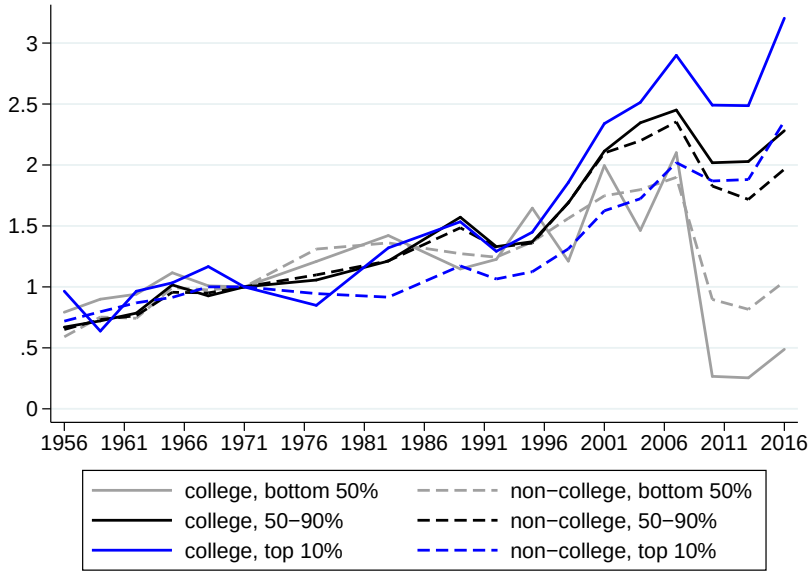

Notes: The figure shows the average wealth growth of college and non-college households in the bottom 50%, middle 50-90% and top 10% of the aggregate wealth distribution over time, relative to 1971.

in Figure 3.2b depends on three factors: initial conditions in the base period, the development of the share $s_{e,i,t}$ over time, and the wealth growth in each education-wealth group depicted in Figure 3.4, $\frac{\overline{W}_{e,i,t}}{\overline{W}_{e,i,71}}$. Based on this decomposition, Figure 3.5a presents three counterfactuals. The first one assigns the average wealth growth of non-college households in the top 10% of the aggregate wealth distribution, $\frac{\overline{W}_{nc,top10,t}}{\overline{W}_{nc,top10,71}}$, to their college counterparts. The second one holds the share of college households in each wealth group i fixed at its 1971 level, $s_{c,i,71}$. The third counterfactual combines the two previous counterfactuals. Figure 3.5b presents the analogous "converse" counterfactuals for non-college households. For college households, a substantial part

Figure 3.5: Wealth growth counterfactuals

(a) Counterfactuals: college

(b) Counterfactuals: non-college

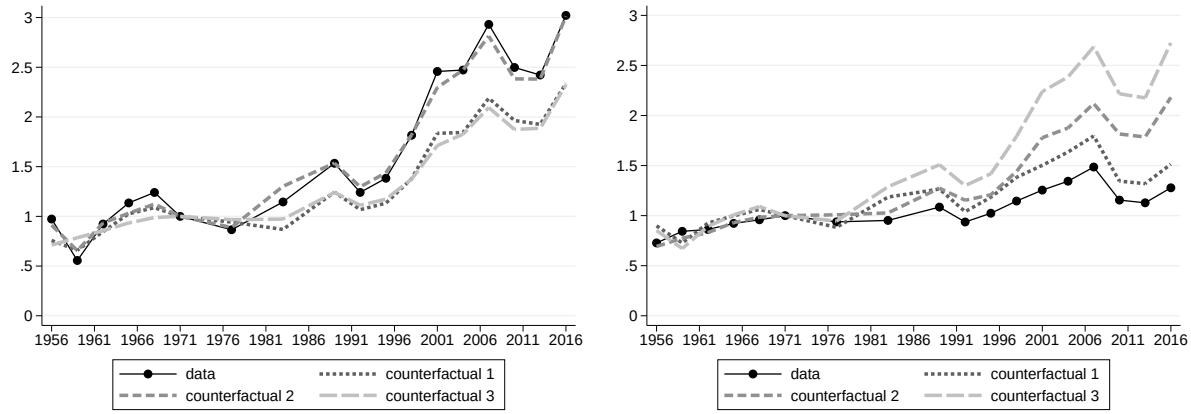

Notes: The left panel shows three counterfactuals for college households. In counterfactual 1, college households from the top 10% of the aggregate wealth distribution are assigned the average wealth growth of their non-college counterparts, $\overline{W}_{nc,top10,t} \,/\, \overline{W}_{nc,top10,71}$. In counterfactual 2, the share of college households in each wealth group i is fixed at its 1971 level, $s_{c,i,71}$. Counterfactual 3 combines the previous two counterfactuals. The right panel presents the analogous exercise for non-college households.

of their wealth growth was driven by faster wealth growth within the top 10% of the wealth distribution, whereas compositional effects across wealth groups barely mattered. By contrast, both compositional effects and wealth growth in the top 10% played an important role for lower wealth growth in the case of non-college households. Surprisingly at first glance, compositional effects across wealth groups only played a minor role in accounting for aggregate wealth growth. To understand why, Figure 3.6 shows the share of households belonging to the bottom 50%, the middle 50-90% and the top 10% of the aggregate wealth distribution within the group of college households (left panel) and non-college households (right panel). The finding that compositional changes contribute little to college wealth growth is reflected in the fact that shares have remained remarkably stable over time. Among college households, around 23% belong to the top 10% of the aggregate wealth distribution on average, 45% to the middle class, and 32% belong to the bottom 50%. We observe only small fluctuations around these time averages over time. For non-college households, the corresponding shares are 6%, 38% and 56%, but there is a visible trend towards a lower top 10% share and a larger bottom 50% share, which is reflected in the results of Figure 3.5b.

The relative stability of the shares within the group of college households also has important implications for the discussion of education as a means of financial mobility. Our results suggest that obtaining a college degree does not increase the probability of finding oneself in the upper parts of the wealth distribution. A college degree seems to help households to keep pace, but not to climb the wealth ladder.

While Figure 3.6 slices the data by educational attainment, Figure 3.7 shows the share of college and non-college households within each group of the wealth distribution, as well as the full cross section. The overall share of households with a college-educated head has quadrupled from 8.1% in 1956 to 34% in 2016. We find that this increase was distributed evenly across wealth groups, so that this trend is consistent with Figure 3.6. Between 1956 and 2016, the college share rose from 5% to 21% in the bottom 50% of the aggregate wealth distribution, from 9.1% to 39.5% in the middle class (50-90%), and from 19.8% to 76.8% in the top 10%. In other words, it has roughly quadrupled in each group, implying that obtaining a college degree does not necessarily go hand in hand with mobility towards the top of the wealth distribution. The college share is largest in the top 10% of the wealth distribution, but having a college degree is not a sufficient condition to reach the top.

Figure 3.6: Shares within education group

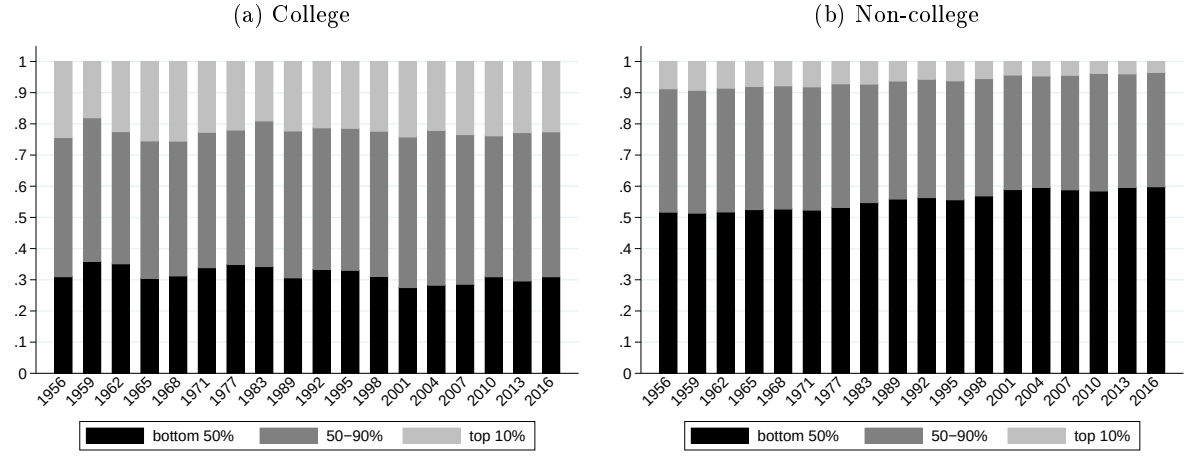

Notes: The figure shows the share of households belonging to the bottom 50%, 50-90% and top 10% of the aggregate wealth distribution for both college and non-college households over time.

It is noteworthy that the increase in average college income and wealth, which we have documented above, has taken place while the group of college households grew larger. Accordingly, the total cake has grown faster than the amount of people sharing it. These developments also imply that college households have appropriated larger and larger shares of total wealth and income over time. While non-college households still accounted for 78% of total wealth and 83% of total income in 1956, these shares have fallen to 26% and 39%, respectively, by 2016. As we will show next, the wealth and income advantages are particularly large if a household is formed by two spouses who both hold a college degree.

3.3.3 The role of marriage patterns for wealth and income growth

The share of households in which both partners hold a college degree has risen over time. Using Census data, previous research has investigated the importance of assortative mating for income inequality (Eika, Mogstad, and Zafar 2019, Greenwood et al. 2014, Greenwood et al. 2015). Positive (negative) assortative mating refers to a situation when people with the same level of education marry more (less) frequently than what would be expected if marriage patterns were random. The existing studies suggest that positive assortative mating helps to explain cross-sectional income inequality, but hardly contributes to changes of income inequality over time.

Figure 3.7: Shares of college and non-college households by wealth group

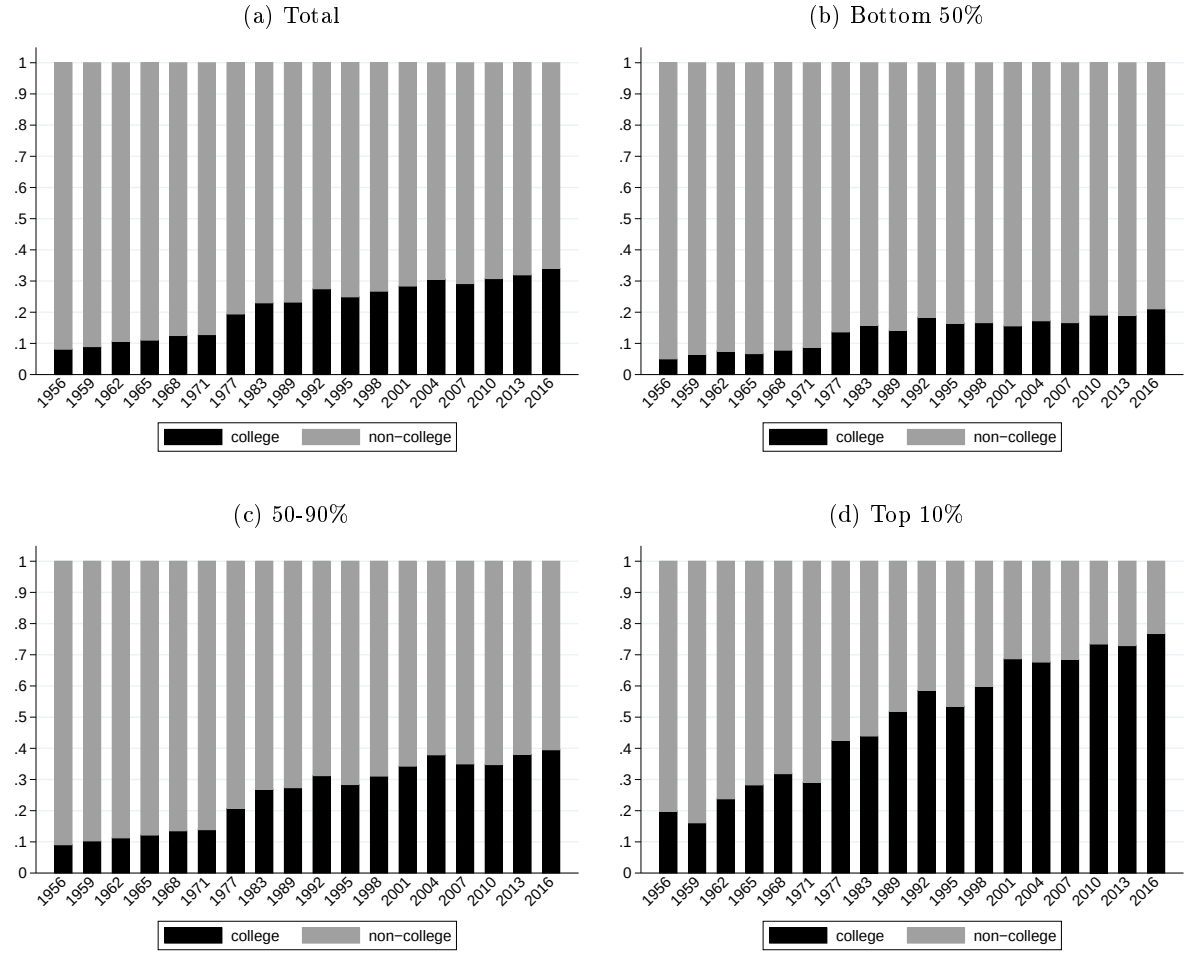

Notes: The figure shows the share of college and non-college households in the entire population and in each group of the aggregate wealth distribution over time.

Figure 3.8: Average income and wealth by education of spouses

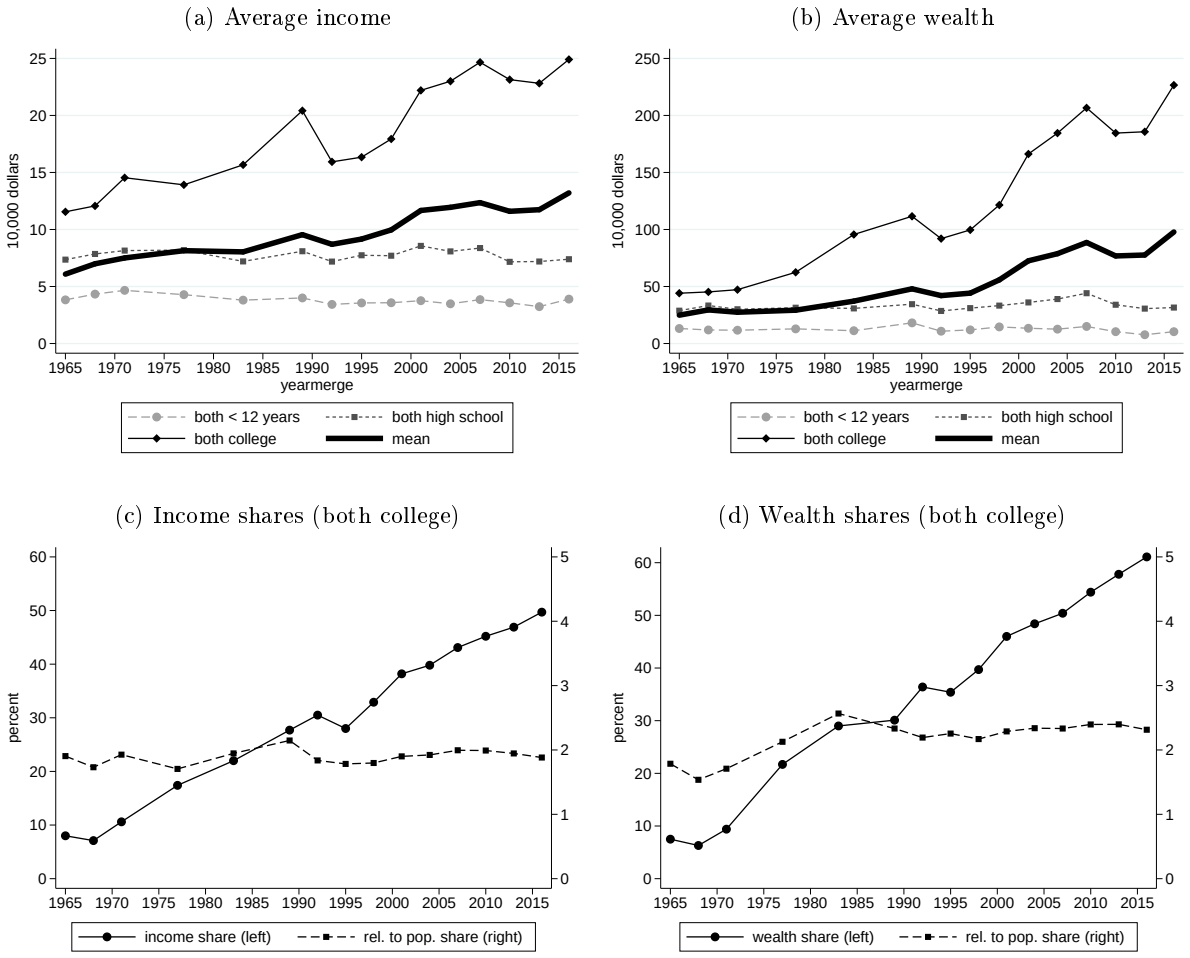

Notes: The upper left panel shows average income of households in which *both* partners have less than 12 years schooling, a high school degree and a college degree, respectively, over time. The thick black line depicts average income for all non-single (married or living with partner) households. The male partner was defined as the household head. The upper right panel shows analogous results for net wealth. The lower left panel shows the share of total income of non-single households appropriated by dual-college households, once in percent (black line, left axis), and once relative to this group's population share (gray line, right axis). The lower right panel shows analogous results for net wealth.

The SCF+ data allow to shed light on the role of marriage patterns and assortative mating for wealth in addition to income inequality. The data show that couples with two college-educated spouses have experienced particularly large increases in income and wealth. In Figure 3.8a, we see that dual-college households have increased their average income by a factor of around two between 1965 and 2016, while income stagnated for households in which both partners hold a high school degree, and decreased for households in which both spouses have completed less than 12 years of schooling.[7] A qualitatively similar, but quantitatively even more pronounced picture emerges for wealth in Figure 3.8b. Households in which both spouses have a college degree have more than quadrupled their wealth, while households without any college-educated spouse have experienced very meager wealth growth.

Figure 3.8c shows that while dual-college households only appropriated 8% of all non-single households' income in 1965, the figure has increased to 49.7% in 2016.[8] However, the population

[7] Information on the spouse's educational attainment is only available since 1965.

[8] Non-single households comprise marriage and cohabitation. We find very similar results when only considering married households.

Table 3.1: Marriage patterns in selected years – actual data versus random matching

	Data				Random			Ratio		
head/spouse	(1)	(2)	(3)	*sum*	(1)	(2)	(3)	(1)	(2)	(3)
		1965				*1965*			*1965*	
< 12 years (1)	**36.6**	14.6	0.4	51.6	**23.5**	24.9	3.2	**1.6**	0.6	0.1
high school (2)	8.6	**26.8**	1.6	37.0	16.8	**17.9**	2.3	0.5	**1.5**	0.7
college (3)	0.3	6.9	**4.2**	11.4	5.2	5.5	**0.7**	0.1	1.3	**5.9**
sum	45.5	48.3	6.2	100						
		1977				*1977*			*1977*	
< 12 years (1)	**20.2**	13.1	0.7	34.0	**9.7**	19.1	5.3	**2.1**	0.7	0.1
high school (2)	7.7	**32.7**	4.7	45.1	12.8	**25.3**	7.0	0.6	**1.3**	0.7
college (3)	0.5	10.3	**10.2**	21.0	6.0	11.8	**3.3**	0.1	0.9	**3.1**
sum	28.4	56.1	15.6	100						
		1989				*1989*			*1989*	
< 12 years (1)	**11.8**	10.7	0.7	23.2	**4.4**	14.4	4.4	**2.7**	0.7	0.2
high school (2)	6.6	**38.9**	5.2	50.7	9.6	**31.5**	9.5	0.7	**1.2**	0.5
college (3)	0.6	12.5	**12.9**	26.0	4.9	16.1	**4.9**	0.1	0.8	**2.6**
sum	19.0	62.1	18.8	100						
		1998				*1998*			*1998*	
< 12 years (1)	**6.9**	7.6	0.2	14.7	**1.7**	9.2	3.8	**4.0**	0.8	0.1
high school (2)	4.3	**44.4**	7.4	56.1	6.6	**35.1**	14.5	0.7	**1.3**	0.5
college (3)	0.5	10.5	**18.3**	29.3	3.4	18.3	**7.6**	0.1	0.6	**2.4**
sum	11.7	62.5	25.9	100						
		2007				*2007*			*2007*	
< 12 years (1)	**6.7**	6.3	0.3	13.3	**1.5**	7.8	4.0	**4.4**	0.8	0.1
high school (2)	4.4	**41.0**	8.2	53.6	6.1	**31.4**	16.1	0.7	**1.3**	0.5
college (3)	0.3	11.2	**21.6**	33.1	3.8	19.4	**10.0**	0.1	0.6	**2.2**
sum	11.4	58.5	30.1	100						
		2016				*2016*			*2016*	
< 12 years (1)	**5.4**	6.1	1.0	12.5	**1.2**	6.5	4.8	**4.5**	0.9	0.2
high school (2)	3.8	**35.2**	11.0	50.0	4.9	**26.0**	19.2	0.8	**1.4**	0.6
college (3)	0.5	10.7	**26.4**	37.6	3.6	19.6	**14.4**	0.1	0.5	**1.8**
sum	9.7	52.0	38.4	100						

Notes: The table shows the relative size (in percent) of marriage groups defined by education of head and spouse over time. The reference total are all non-single households (married or living with partner) with information on the educational attainment of both spouses. (1) means less than 12 years of schooling, (2) means that the person has a high school degree, and (3) that the person has a college degree. Rows refer to the head, columns to the spouse. The male partner was defined as the household head. The left part of the table shows each group's relative size in the data. The middle part shows the corresponding shares if matching was random. The right part shows the ratio of the shares in the data to the shares that would have been obtained with random matching. The counterfactual was computed from marginal frequencies. Reading example: In 1965, in 4.2% of households both head and spouse had a college degree. 6.2% of all spouses had a college degree, and 11.4% of all heads. The share of dual-college households was 5.9 times as large as it would have been with random matching.

share of this group has also increased from 4.2% to 26.4%, such that the income share of dual-college households relative to their population share has hardly changed. A similar result pertains to wealth (see Figure 3.8d). In this sense, the increasing share of total income and wealth accruing to (dual-)college households has not been disproportionate.

The data also document that dual-college households have appropriated larger shares of income and wealth over time, but that this has little to do with assortative mating. Table 3.1 compares actual marriage patterns with those that would have been observed under random matching based on marginal frequencies. We find that assortative mating has actually *decreased* for college-educated individuals, whereas it has increased for low-educated individuals. Our results are both qualitatively and quantitatively consistent with the findings of Eika, Mogstad, and Zafar (2019), who use U.S. data from the March Current Population Survey (CPS) for 1962-2013.[9] In other words, the fact that we see a larger share of dual-college households nowadays can mainly be attributed to increases in educational attainment, especially among females, rather than changes in preferences and sorting. While there were around 11.4% male college graduates in 1965, the share of college graduates among the female partners was only 6.2%. The shares have increased to 37.6% for males and 38.4% for females in 2016.

3.4 Portfolio composition and entrepreneurship

The previous section has presented evidence that college households have improved their wealth position substantially compared to non-college households, which we referred to as an increase of the college wealth premium. In this section, we will investigate potential drivers of this development in more detail.

Figure 3.9 contrasts the increase in the college wealth with the college income premium.[10] The ratio of college to non-college income was roughly stable until the late 1970s. The ratio of college to non-college wealth fluctuated somewhat more over this period, but did not show any trending behavior. In the early 1980s, both ratios embarked on a largely uninterrupted upward trend. The only exceptions were the early 1990s recession and the burst of the "dot-com bubble" in 2001. The wealth premium has increased considerably more than the income premium, namely by around 135% as opposed to an increase of 50% for the income premium between 1971 and 2016. The discrepancy in the development of income and wealth becomes even more explicit when we only consider the middle class (50-90%) of the *income* distribution, and average the data by decades.[11] The results are shown in Figure 3.10. Even for middle-income households who, by construction, had almost identical income paths, the increase of the college wealth premium since the 1980s stands out. While wealth has doubled for college households, non-college households with the same income trends saw their wealth only increasing by 25%. Which role do demographic shifts play for the observed phenomena? So far, we have looked at unconditional averages. To obtain an estimate of the "college wealth effect" net of potential confounders such as demographics, we estimate the following micro-level regression:

$$W_{it} = \beta_0 + \beta_1 c_{it} + \sum_{t>1956} \beta_{2,t} \mathbb{I}_{[year=t]} \cdot c_{it} + \sum_{t>1956} \beta_{3,t} \mathbb{I}_{[year=t]} + \Gamma' X_{it} + \xi_{it} \,. \tag{3.2}$$

W_{it} denotes wealth of household i in survey wave t, $\mathbb{I}_{[year=t]}$ are survey wave fixed effects for

[9] Eika, Mogstad, and Zafar (2019) demonstrate the robustness of these patterns to accounting for sorting by age and changes in the probability of marriage by education level, as well as to different measures of assortative mating.

[10] Note that Figure 3.3 showed wealth-to-income ratios.

[11] Note that our results in the following sections always refer to the full sample and not the "middle-class" sample, unless explicitly stated otherwise.

Figure 3.9: Wealth and income ratios: college/non-college

Notes: The graph shows the ratio of college to non-college wealth (black line) and income (gray dashed line) over time, relative to its level in 1971.

$t \in \{1959, 1962, ..., 2016\}$, and c_{it} is an indicator for whether the head has a college degree. The control vector X_{it} includes total household income, a full set of age dummies, a dummy for whether the household includes children, and an indicator for whether the head is married. As a baseline specification, we estimate this regression on the entire sample. As a robustness check, we also estimate it on a restricted sample that is limited to households in the 50-90% group of the aggregate income distribution. This restriction can be interpreted as an additional non-parametric way of controlling for income. Figure 3.11 illustrates the results, and the underlying coefficient estimates are summarized in Table 3.B.3. The college wealth effect ($\beta_1 + \beta_{2,t}$) is clearly visible from the 1980s on. The figure also illustrates that college wealth tends to be hit more

Figure 3.10: Middle-class wealth and income levels by decade

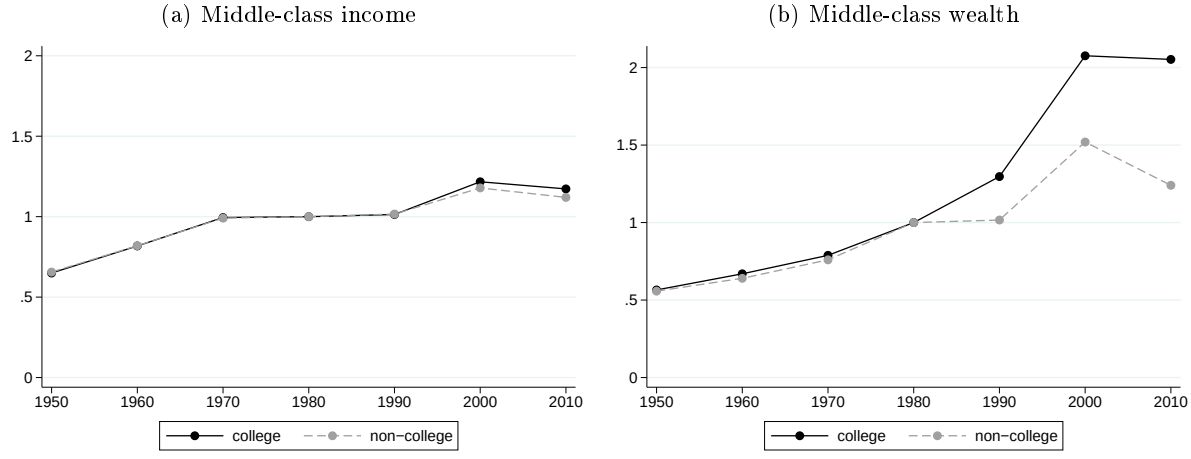

Notes: The figure shows average income and wealth for households from the 50-90% of the aggregate *income* distribution by education. Data were averaged across decades, and normalized to the 1980s.

severely in recessions, which tend to reduce the college wealth premium. For the middle-class income sample, the college wealth effect is smaller in size, but still clearly visible since the 1980s. Indeed, the college wealth effect is strongest for the top 10% of the income distribution, but also visible for the bottom 90% (see Figure 3.A.3). By contrast, the college income effect is much smaller, and entirely driven by the top 10% of the aggregate income distribution.

3.4.1 Portfolio heterogeneity

In workhorse models of wealth inequality following the early work of Huggett (1993) and Aiyagari (1994), wealth growth depends on the amount of savings, so that changes in income inequality translate into changes in wealth inequality. However, it is a well-established fact that wealth inequality exceeds income inequality. Economic theory highlights the role of the life cycle, bequests, entrepreneurship and differential returns on assets to explain this finding (see e.g. Cagetti and De Nardi 2008, De Nardi and Fella 2017, Benhabib, Bisin, and Luo 2017). Kuhn, Schularick, and Steins (2020) provide further empirical substance to the important role of differences in portfolio choice and associated returns. They illustrate how differences in household portfolios along the wealth distribution, combined with differential asset price growth, can lead to a "decoupling" of the growth of income and wealth. The SCF+ data allow us to examine the portfolio composition of households with different educational attainment. Figure 3.12 illustrates that the average portfolios of college and non-college households do not merely differ in size, but also in composition. In particular, the share of non-financial assets is substantially larger for non-college households. Table 3.B.2 in the Appendix shows that this high share is mainly accounted for by housing. For instance, the housing portfolio share of non-college households was 53% in 2007, compared to 37.3% for households with a college-educated head. In 2016, these shares had slightly decreased to 46.9% and 30.9%. By contrast, college households tend to hold larger shares of business wealth and equity than non-college households.

Kuhn, Schularick, and Steins (2020) show how portfolio differences give rise to differential expo-

Figure 3.11: Regression evidence

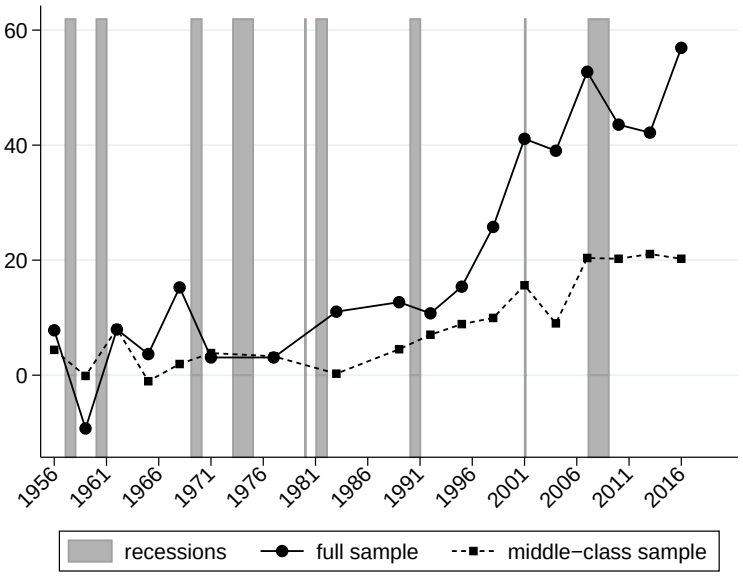

Notes: The graph shows the advantage of having a college degree ($\beta_1 + \beta_{2,t}$) over time. The solid line with dots refers to the baseline in (3.2). The dashed line with squares presents the results for a restricted sample including only households from the 50-90% of the aggregate income distribution. The shaded gray areas show NBER recessions.

sure to asset price changes. In the case of college households, their higher equity portfolio share allowed them to reap higher capital gains due to increasing stock prices over the past 30 years. Figure 3.13 shows that real equity prices, taken from the Macrohistory Database, have tripled since 1989. The figure also illustrates that the increase in stock prices has moved hand in hand with the ratio of college to non-college equity holdings over this period. As capital gains from asset price changes are unrelated to the development of income, they can help to explain why the college wealth premium has increased substantially more than the college income premium. Indeed, the estimated college wealth advantage is reduced when we control for stock market exposure in our micro regressions (see Figure 3.14). We measure stock price exposure via the portfolio share of equity s^e and the real average equity price P^e, included in levels, interacted with each other, and interacted with college. More precisely, we estimate the following regression:

$$W_{it} = \beta_0 + \beta_1 c_{it} + \sum_{t>1956} \beta_{2,t} \mathbb{I}_{[year=t]} \cdot c_{it} + \sum_{t>1956} \beta_{3,t} \mathbb{I}_{[year=t]} + \\ \beta_4 s^e_{it} + \beta_5 s^e_{it} \cdot c_{it} + \beta_6 s^e_{it} \cdot P^e_t + \beta_7 s^e_{it} \cdot P^e_t \cdot c_{it} + \Gamma' X_{it} + \epsilon_{it} \,. \tag{3.3}$$

Figure 3.14 shows that controlling for stock price exposure reduces the college wealth premium substantially, especially during the stock market booms in the 1960s and since the 1990s. The high correlation of college to non-college equity growth with stock price growth since the 1990s stock market boom, suggests that differential stock market exposure via direct stock and mutual fund holdings played a key role for the rapid increase in the college wealth premium since the 1980s.

3.4.2 Business ownership

While stock price exposure can account for an important share of the observed college wealth premium, there still remains an unexplained wealth growth differential between college and non-college households. In a second step, we explore the role of business ownership for the observed trends. Motivated by the fact that business wealth has gained importance in the portfolio of

Figure 3.12: Portfolio shares of college and non-college households

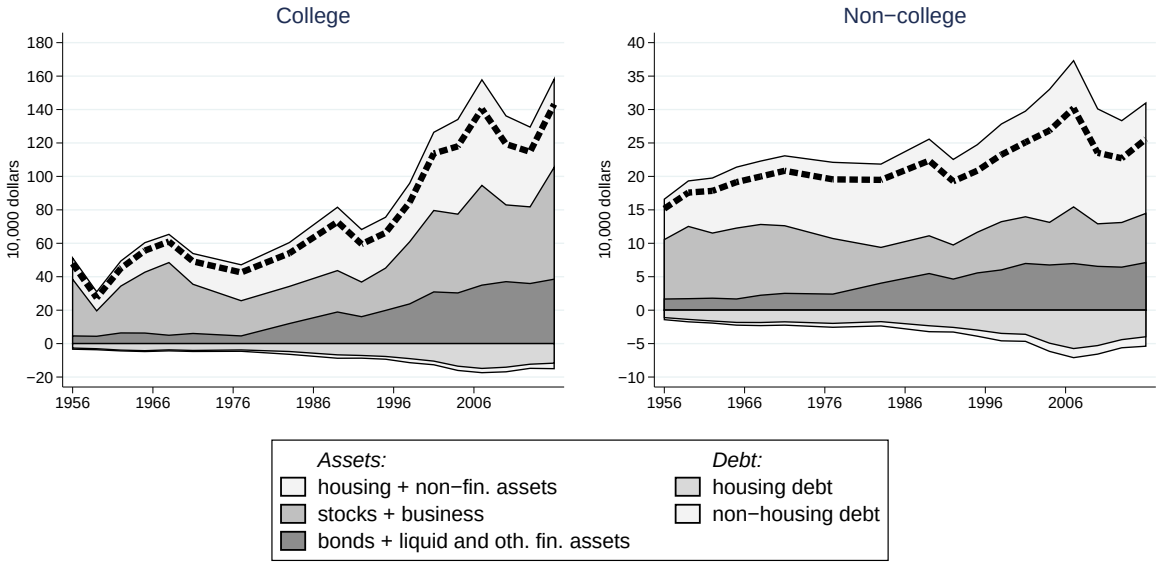

Notes: The left panel shows the portfolio composition of college households over time. The right panel shows the portfolio composition of non-college households. Housing includes non-residential real estate. Stocks include mutual fund holdings and other managed assets.

Figure 3.13: Stock market exposure and equity prices

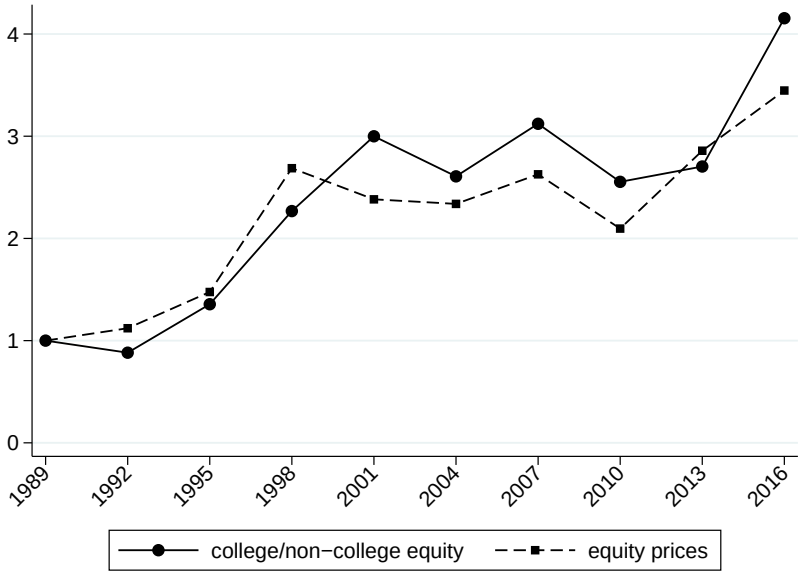

Notes: The solid line with dots shows the ratio of average college to non-college equity. The dashed line with squares shows the average stock price from the Macrohistory database, transformed to 2016 dollars. Both series were indexed to 1989. Stocks include mutual fund holdings and other managed assets.

college households in recent years, and that business assets, just like equity, are an asset class which is primarily held by the top 10% of the aggregate wealth distribution (Kuhn, Schularick, and Steins 2020), we look at the effects of controlling for business ownership.

In equation (3.3), we included the portfolio share of equity and mutual funds. Fagereng et al. (2020) point out that entrepreneurial skills may affect the whole portfolio via differential

Figure 3.14: Controlling for stock market exposure

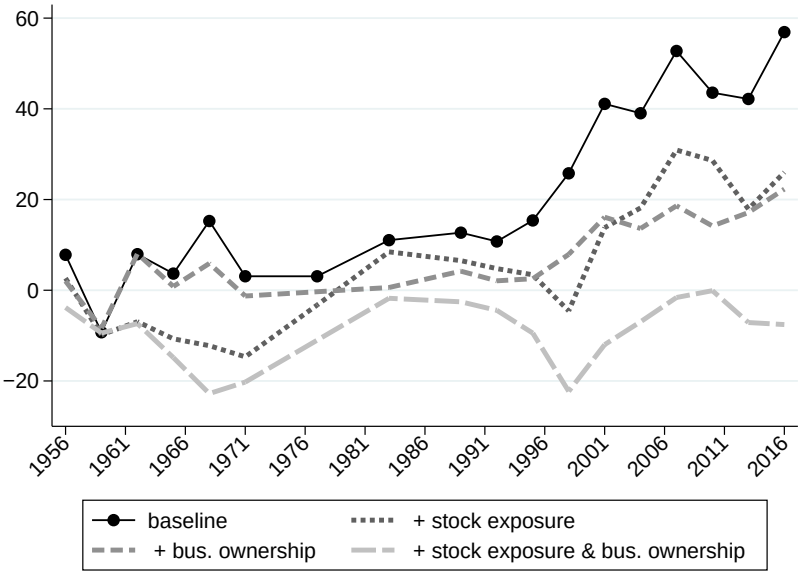

Notes: The graph shows the advantage of having a college degree ($\beta_1 + \beta_{2,t}$) over time. The solid line with dots repeats the baseline from (3.2) as a reference. The short-dashed, dark gray line presents the results for regression (3.3), the dashed, dark medium line presents the results for regression (3.4), and the long-dashed, light gray line includes the additional controls from both (3.3) and (3.4).

returns. Therefore, we include a more general dummy for business ownership, bus_{it}, instead of the portfolio share of business wealth in this specification. Since there is no general market price for business assets, we interacted it with year fixed effects to allow for variation across time. Apart from these slight changes, the specification follows that in equation (3.3):

$$W_{it} = \beta_0 + \beta_1 c_{it} + \sum_{t>1956} \beta_{2,t} \mathbb{I}_{[year=t]} \cdot c_{it} + \sum_{t>1956} \beta_{3,t} \mathbb{I}_{[year=t]} + \beta_4 bus_{it} + \beta_5 bus_{it} \cdot c_{it} + \sum_{t>1956} \beta_{6,t} \mathbb{I}_{[year=t]} \cdot bus_{it} + \sum_{t>1956} \beta_{7,t} \mathbb{I}_{[year=t]} \cdot bus_{it} \cdot c_{it} + \Gamma' X_{it} + \varepsilon_{it} \,. \tag{3.4}$$

The dashed, medium gray line in Figure 3.14 shows the resulting coefficients β_1 and $\beta_{2,t}$. Moreover, we estimate a specification of the regression in which we include the additional controls from (3.3) and (3.4) jointly. This specification is shown as the long-dashed, light gray line in Figure 3.14. The estimations suggest that also business ownership has contributed to the increase of the college wealth premium. Yet our regressions only show conditional correlations, and cannot lay claim to causality. In the following section, we will discuss potential underlying mechanisms for the observed correlations.

3.5 Financial literacy and returns on wealth

Our results suggest stock market exposure and business ownership as driving forces for the rise of the college wealth premium. In the following, we will explore potential reasons why these factors are important, pointing to promising directions for future research. In particular, we will ask which role financial literacy plays for portfolio composition and for differential returns on wealth.

3.5.1 Financial literacy and portfolio composition

One reason why college households hold different assets may be financial literacy. Previous research has consistently established that higher educational attainment is associated with higher levels of financial literacy (see Lusardi and Mitchell 2011, Lusardi and Mitchell 2014, Lusardi, Michaud, and Mitchell 2017).[12] Higher financial literacy can, for example, affect wealth growth through portfolio composition. Typically, financial literacy is measured via three questions that elicit the understanding of interest compounding, inflation, and risk diversification (see Lusardi and Mitchell 2011). Since the 2016 wave, these questions are also part of the SCF. The left part of Figure 3.15 shows the share of households who answered all three questions correctly, stratified by wealth and education. The figure reveals that financial literacy increases with wealth, and is clearly higher for college than for non-college households in all wealth groups. Moreover, the right part of Figure 3.15 shows that stock owners do on average have a higher level of financial literacy than non-owners. The same is true for business owners. In a study of Dutch household survey data, Von Gaudecker (2015) shows that low financial literacy leads to return losses due to under-diversification in financial assets (unless households seek financial advice). Indeed, the question on risk diversification is the financial literacy question which respondents find most challenging across a wide range of countries (see Lusardi and Mitchell 2014). In the 2016 SCF, the difference between college and non-college households is most pronounced for the risk diversification question, with an average share of correct answers of 75.9% for college

[12] Lusardi, Michaud, and Mitchell (2017) demonstrate that this can actually be an individually optimal outcome, as households with different levels of education have different life-cycle income paths, which entail different incentives to save. Given that financial knowledge helps to earn higher returns on savings, this creates different incentives to invest in financial literacy.

Figure 3.15: Financial literacy

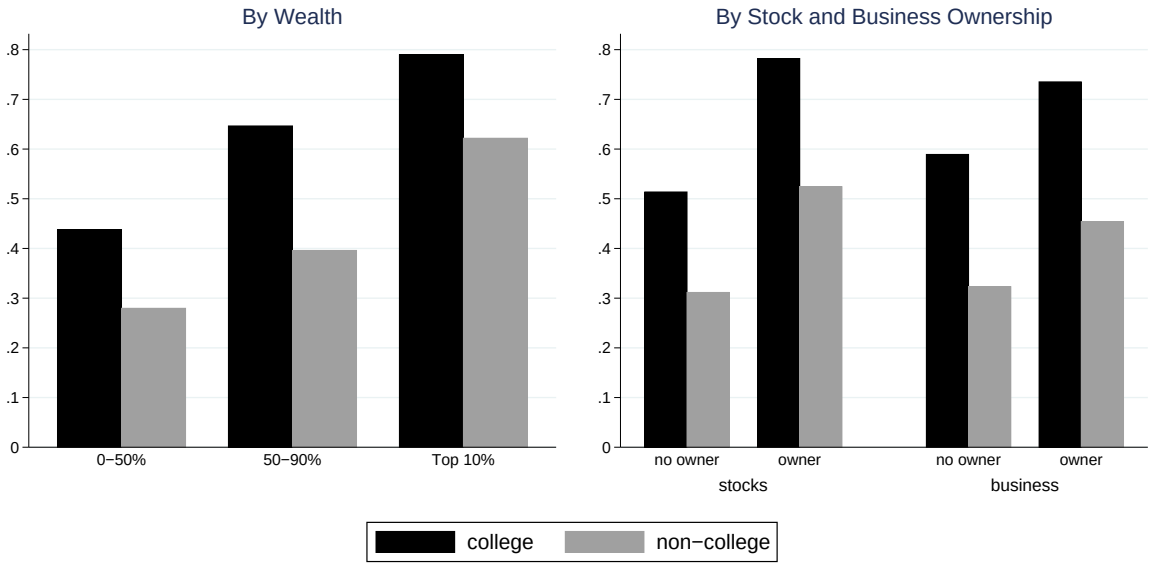

Notes: The left panel shows the share of households who answered all three questions on financial literacy correctly in the 2016 SCF by education and wealth group. The right panel shows this share stratified by stock and business ownership instead of wealth.

and 55.9% for non-college households.[13] Moreover, while we have only looked at direct stock market exposure, this may also impact indirect exposure via pension plans. Lusardi, Michaud, and Mitchell (2017) stress the heightened importance of financial literacy in the U.S. due to the movement from defined benefit to defined contribution plans, such as 401(k)s. This transition has started in the 1980s, coinciding with the timing of the widening college wealth premium.

To explore the role of financial literacy, we estimate a regression analogous to the baseline in equation (3.2) for 2016. When we included the SCF financial literacy measure (which equals one if all three questions were answered correctly, and otherwise zero), as well as its interaction with the college indicator, the estimated college effect was reduced by around 40%. Yet this result is only suggestive. Further research is necessary to investigate the robustness of the finding and potential transmission mechanisms.

3.5.2 Returns on wealth

College education might affect wealth accumulation not only through its effect on portfolio allocation across different asset classes, but also via higher returns *within* a given asset class. For instance, college households might be savvier in picking investments with high returns or low fees. Fagereng et al. (2020) demonstrate that persons with higher levels of education, and especially those with an economics-related college degree, earn higher returns on their wealth and financial assets even *conditional on* portfolio composition. This suggests that our estimate of the effect of differential asset price exposure from Section 3.4 might be conservative, given that we applied the average rate of return on stocks for all households.

Apart from financial savvy, Fagereng et al. (2020) also point to entrepreneurial skills as a source of differential returns. Moreover, borrowing constraints can induce entrepreneurs to save substantial amounts and thus become very wealthy (Cagetti and De Nardi 2006). There is evidence that business ownership is associated with higher education (see e.g. Hurst and Lusardi 2004).

[13] The corresponding share for the question on interest compounding are 86% and 73%, and for the question on real interest rates 86.7% and 72.9%.

Consistently, the share of business owners in the SCF+ is higher among college households than among non-college households in all waves (and conversely, the share of college households is disproportionately high among business owners).

Fagereng et al. (2020) use administrative individual-level data from Norway to construct a measure of returns to financial wealth. To this end, they add income from save and risky assets, and divide it by the average stock of financial and business assets.[14] Due to the panel structure of their data, they can use the average of beginning- and end-of-period assets as the denominator, in order to account for changes in the stock of assets over the current period. This is not possible with the SCF, as it consists of repeated cross sections. Moreover, income is reported for the year previous to the survey year. For these reasons, a similar measure constructed from SCF data

Figure 3.16: Returns to financial and business wealth

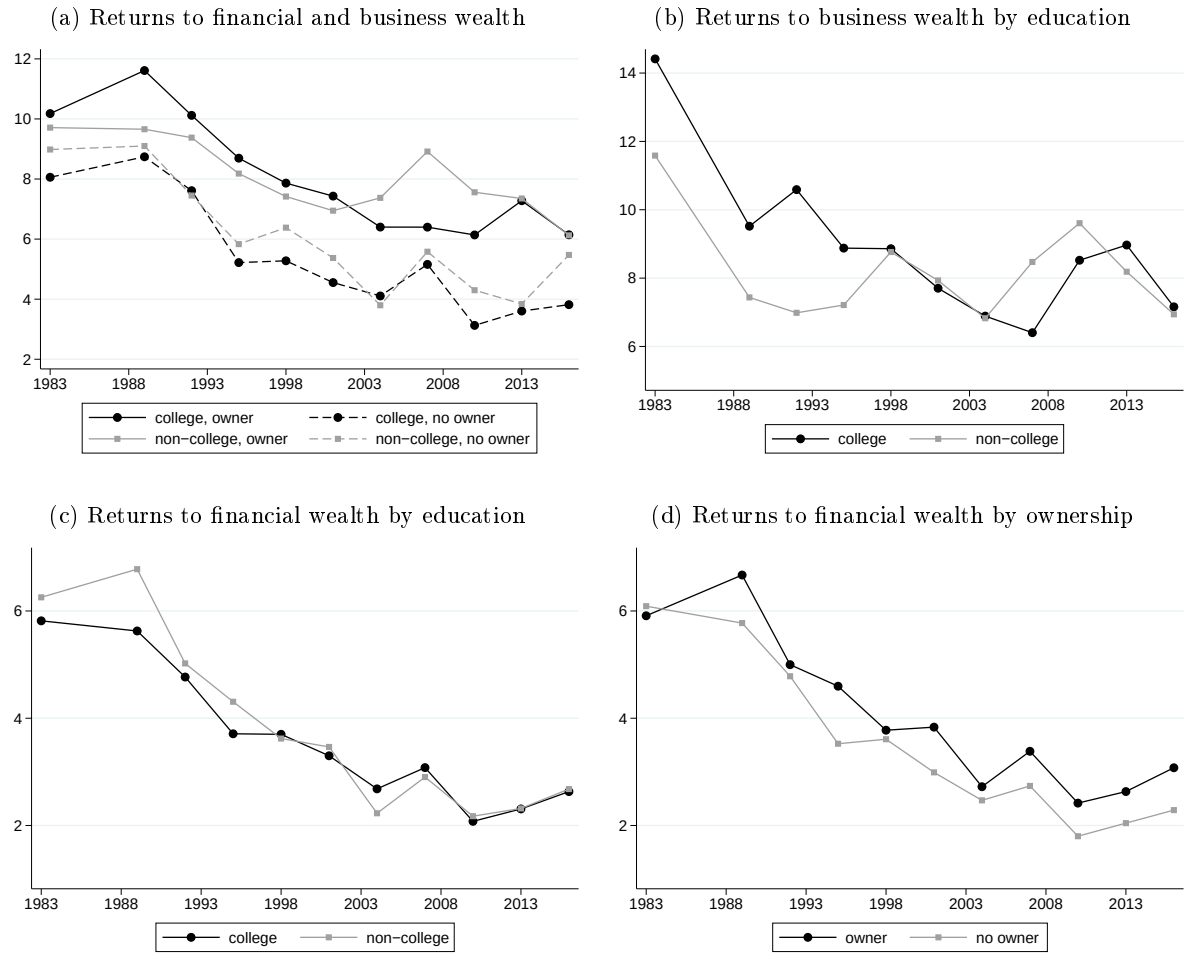

Notes: The upper left panel shows the proxy for returns to financial and business wealth by education and business ownership status. The upper right panel presents a similar proxy for business wealth only, stratified by education. The lower left panel presents analogous results for a proxy based on financial wealth only. The lower right panel presents the latter measure by business ownership status instead of education. See text for additional details.

is likely to include more measurement error. Keeping this in mind, we construct an analogous

[14] They consider the following income components: interest income earned on bank deposits and bond yields, yields from risky assets held abroad and outstanding claims and receivables, yields from mutual funds, yields from directly held listed shares (dividends and accrued capital gains), yields from all private equity holdings (distributed dividends and the individual share of retained profits). Financial wealth includes bank deposits, money market funds, bond mutual funds, government and corporate bonds, stocks and mutual fund shares, the value of shares in private businesses and other unlisted shares, and the value of risky assets held abroad and of outstanding claims and receivables.

proxy for returns based on the modern SCF data, which include detailed information on different components of income.[15] For the numerator, we use information on income from farming and business, income from other businesses, rents, trusts or royalties, income from non-taxable investments such as municipal bonds, dividend income, capital gains and losses, and other interest income. For the denominator, we add the amount of stocks, liquid assets and certificates of deposit, bonds, mutual funds, other managed and financial assets, the cash value of life insurances, defined contribution pension wealth, and business wealth.

Fagereng et al. (2020) exclude persons with less than 500 dollars in financial wealth, and winsorize the bottom and top 0.5% of the returns distribution. We also drop households with less than 500 dollars in financial wealth, and drop returns below the 0.5th and above the 95th percentile. The larger trimming region at the top was chosen to take into account that we only observe wealth at one point in time, and with a certain lag compared to income. If a household sold most of its financial assets in the year prior to the survey, it would have had a high capital income in that year, and a relatively low amount of financial assets when surveyed, which would lead to an upward bias in the returns proxy. As Fagereng et al. (2020) report that their results are insensitive to applying an age limit of 20 to 75 years, we include households of all ages. Like them, we use real variables before taxes. The results are presented in Figure 3.16a. On

Figure 3.17: College business ownership

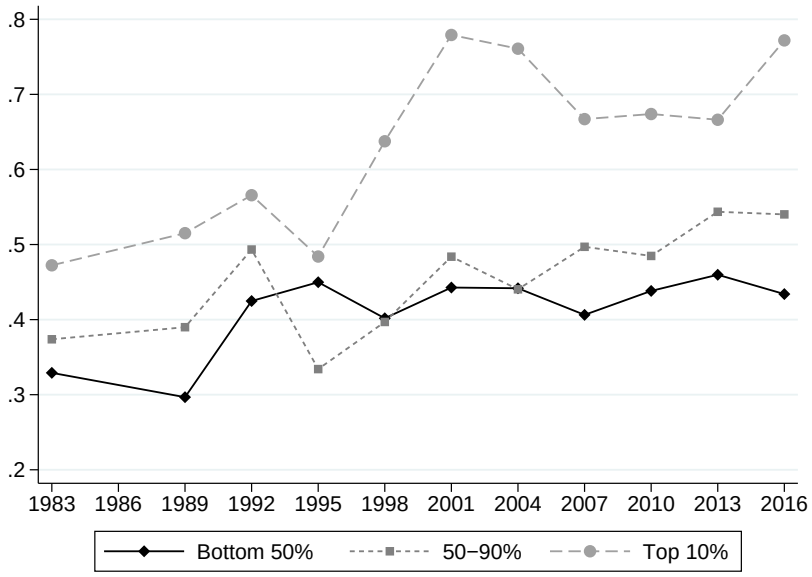

Notes: The figure shows the share of college households in the bottom 50%, middle 50-90% and top 10% of the business wealth distribution (conditional on owning a business).

average, business owners earn higher returns on their financial assets than non-owners. However, the difference between college and non-college households is limited. Figure 3.16b presents a similar returns proxy for business wealth only, i.e., with income from farming, business, other businesses, rents, trusts or royalties in the numerator, and business wealth in the denominator. While college households had higher returns to business wealth as measured by the proxy until the mid 1990s, the advantage disappears afterwards, and is even reversed after 2004.

The same comparison for financial wealth is presented in Figure 3.16c, using the complementary set of income measures in the numerator. Importantly, there is no advantage for college households with respect to this measure, and even a disadvantage in the 1980s. Finally, Figure 3.16d

[15] Fagereng et al. (2020) also construct a proxy for returns to net worth. As this measure requires information on interest payments on all debt, which is so far neither included in the SCF+ nor the readily available extracts of the modern SCF, we decided to focus on returns to financial and business wealth.

shows the proxy for returns to financial wealth by business ownership status instead of education. We find that business owners earn slightly higher returns on their non-business wealth as well, in line with the hypothesis of Fagereng et al. (2020) that entrepreneurs' "talent to manage and organize their business" (p. 5) enables them to generate higher returns in general.

However, while our return proxies are necessarily coarse due to the measurement issues described above, it appears that the return differences between college and non-college households are small. Yet it is important to keep in mind that a similar rate of return can translate into large level differences if the difference in the underlying asset values is large. Figure 3.17 shows the share of college households in the bottom 50%, middle 50-90% and top 10% of the business wealth distribution, conditional on owning a business. The college share is particularly high in the top 10% group, and has increased from slightly below 50% to almost 80% between 1983 and 2016.

Based on the existing literature and our explorative results presented in this section, it seems plausible that the interaction of educational attainment, financial literacy, and business acumen has played an important role in shaping the differential development of college as opposed to non-college wealth. However, portfolio composition, not differential returns between college and non-college households in the same asset class, appears to play the dominant role.

3.6 Conclusion

This chapter documents the evolution of U.S. college and non-college income and wealth over six decades using newly compiled long-run data at the household level. We corroborate that the college income premium has increased substantially since the 1980s. Yet though the college income premium has increased substantially, the college wealth premium has risen even more. Since the 1970s, college households have outpaced non-college households by a factor of 2.5 in terms of wealth growth. We provide evidence that especially households with two college-educated spouses could appropriate large amounts of wealth. However, we confirm previous evidence that this is not related to assortative mating, but rather to rising educational attainment.

We find that portfolio choices and the resulting exposure to asset price changes played a crucial role for the observed trends. Using the asset information in the SCF+, we uncover systematic differences in the size and composition of college versus non-college household portfolios. Building on insights from previous research, we study the combined role of portfolio choices and asset price changes for the evolution of the wealth distribution. Our results suggest that college households could reap large capital gains from stock market booms owing to the higher equity share in their portfolios. This explanation is consistent both with the fact that college wealth grew faster than non-college wealth, and that college wealth grew faster than college income, since capital gains from asset price changes are not directly related to other sources of income. Moreover, we provide suggestive evidence that the increase in the college wealth premium is related to business ownership.

In the last part of the chapter, we discuss potential reasons for the importance of differential asset price exposure and capital gains such as financial literacy and entrepreneurial skills. Both can affect wealth accumulation via portfolio choice and differential returns. These factors also interact with institutional features such as the change from defined benefit to defined contribution pension plans. Further research will be needed to disentangle different hypotheses for the rising college wealth premium and establish causal relationships.

Appendix 3.A Supplementary Figures

Figure 3.A.1: Comparison to census data

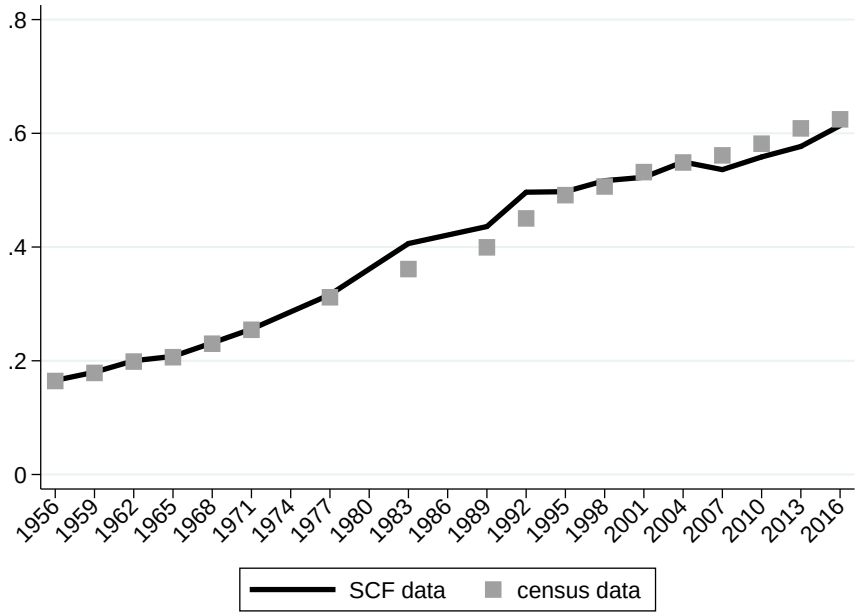

Notes: The figure shows the share of households whose head has at least obtained some college education in the SCF+ data in comparison to the share obtained from the U.S. Current Population Survey (CPS) for the period from 1962 to 2016 and from the U.S. Decennial Census for 1950 and 1960. Intermediate data points were obtained by linear interpolation.

Figure 3.A.2: Wealth-to-income ratios excluding pensions

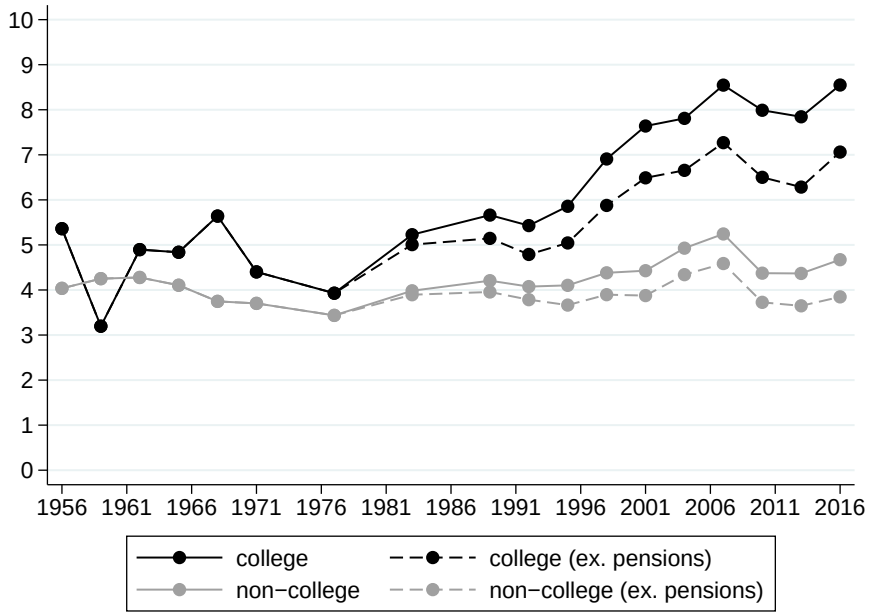

Notes: The figure shows the ratio of average net wealth to average income among households with and without a college-educated head over time. The solid lines replicate the baseline from Figure 3.3 for comparison. The dashed lines show average net wealth net of pensions relative to average income.

Figure 3.A.3: Regression coefficients: advantage of college within income groups

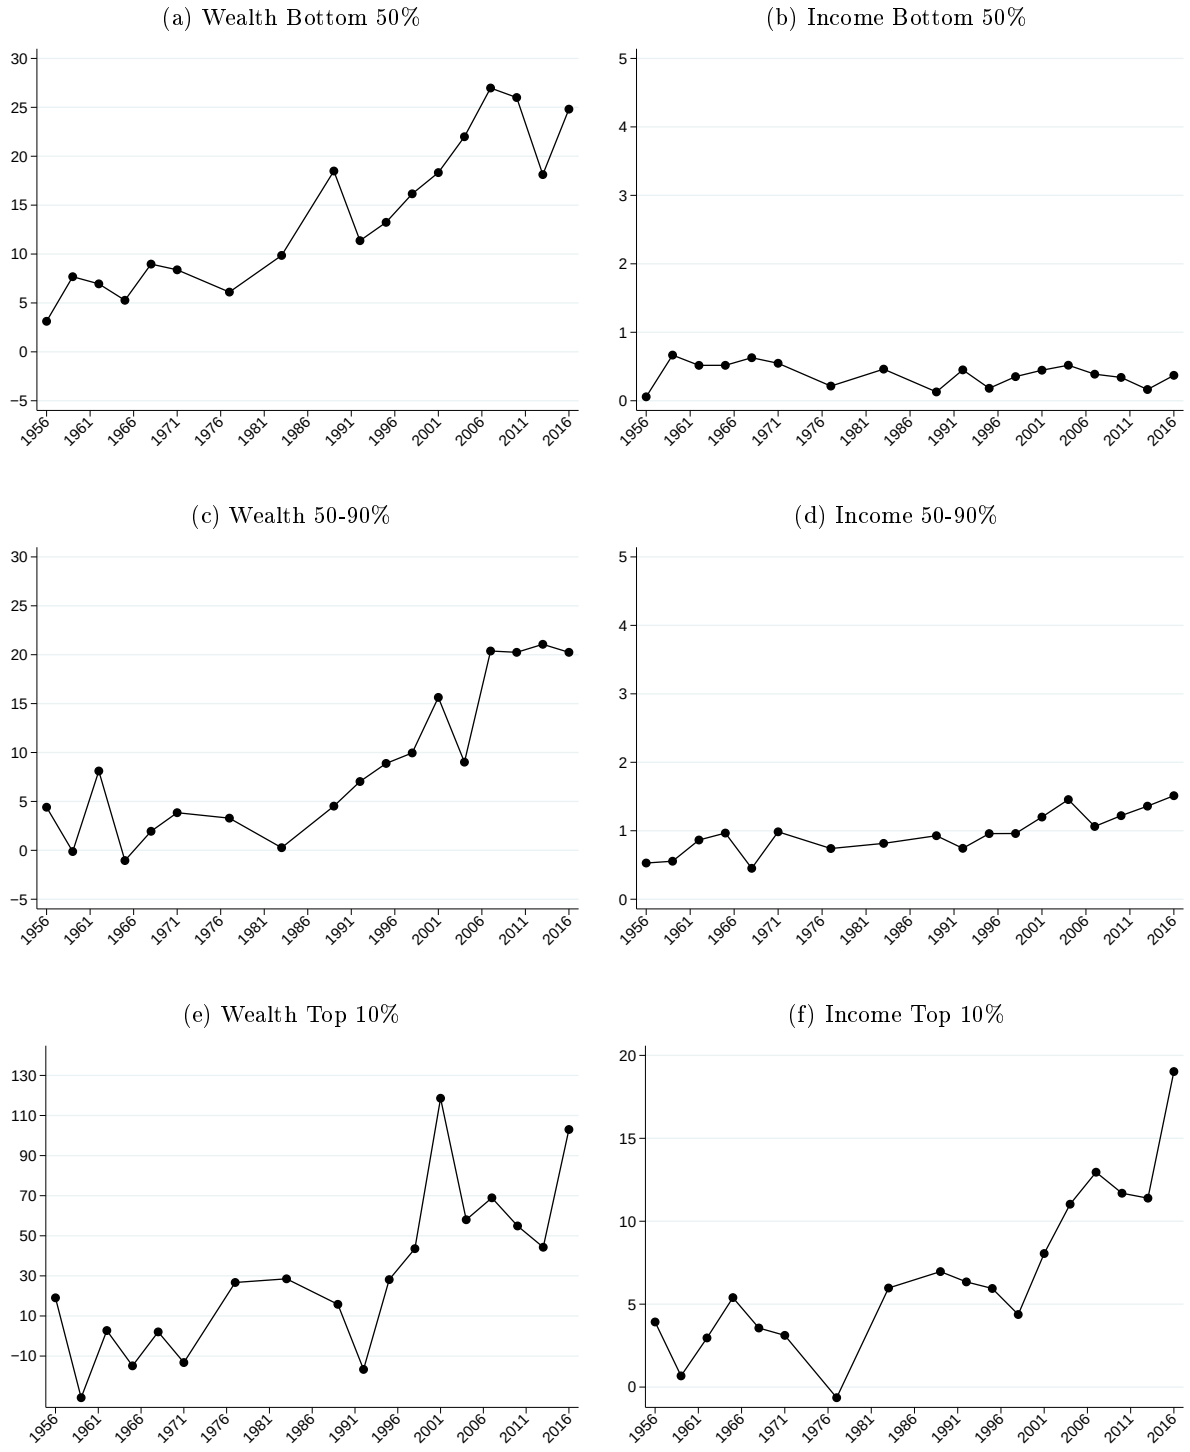

Notes: The figure shows the advantage of having a college degree ($\beta_1 + \beta_{2,t}$) within the bottom 50%, 50-90% and top 10% of the aggregate income distribution over time. The dependent variable is wealth for the left panels and income for the right panels.

Appendix 3.B Supplementary Tables

Table 3.B.1: Comparison of households with college versus some college

year	1956			1959		
college status	some college	college	% Δ	some college	college	% Δ
liquid assets + bonds	39740.9	45956.4	15.6	35207.0	44010.8	25.0
houses	98765.5	116740.2	18.2	94597.0	105928.8	12.0
other nonfin. assets	9340.6	10496.2	12.4	7366.8	8061.5	9.4
housing debt	23698.3	26688.1	12.6	25827.2	30477.1	18.0
non-housing debt	5806.7	6712.7	15.6	5471.3	6000.0	9.7
total income	74855.2	89300.1	19.3	72960.7	85590.3	17.3

Notes: The college status is college for households whose head has obtained at least a bachelor's degree, and some college for those whose head has at least attended college for a year. The columns "% Δ" show the difference between the two groups in percent.

Table 3.B.2: Portfolio shares non-college

year	oth. non-fin. ass.	housing	business	equity	liq. ass.+ bonds	oth. fin. ass.
(a) college						
1956	2.0	22.8	33.3	32.9	9.0	-
1965	1.3	28.0	26.1	34.2	10.5	-
1977	2.9	42.7	20.1	24.5	9.8	-
1989	4.6	41.8	20.9	9.5	10.7	12.5
1998	3.9	32.6	20.5	18.3	7.3	17.5
2007	2.7	37.3	22.5	15.3	6.4	15.8
2016	2.4	30.9	22.0	20.3	7.1	17.2
(b) non-college						
1956	3.2	33.2	32.4	21.1	10.1	-
1965	3.0	39.6	23.5	26.0	7.8	-
1977	4.1	47.4	32.4	5.2	10.9	-
1989	6.4	50.1	18.1	4.0	12.4	9.0
1998	6.5	46.0	14.7	11.2	8.6	13.1
2007	5.6	53.0	16.6	6.1	6.1	12.6
2016	6.4	46.9	17.1	6.6	6.6	16.4

Notes: The table shows the portfolio shares of other non-financial assets, housing, business assets, equity, liquid assets and bonds, as well as other financial assets for college and non-college households over time. Equity includes mutual fund holdings and other managed assets.

Table 3.B.3: Regression results: college effect ($\beta_1 + \beta_{2,t}$) by Year

year	baseline		middle-class sample		+ stock exposure		+ bus. ownership		+ stock & bus.	
1956	7.82	(0.126)	4.41*	(0.053)	2.66	(0.559)	2.02	(0.634)	-3.81	(0.292)
1959	-9.26***	(0.007)	-0.12	(0.943)	-9.83***	(0.003)	-8.37***	(0.003)	-9.43***	(0.001)
1962	7.95	(0.203)	8.11	(0.159)	-6.93	(0.218)	8.05	(0.145)	-7.35	(0.137)
1965	3.67	(0.695)	-1.05	(0.645)	-10.71	(0.210)	0.83	(0.897)	-14.85***	(0.005)
1968	15.26**	(0.049)	1.95	(0.304)	-12.22**	(0.037)	5.91	(0.391)	-22.78***	(0.000)
1971	3.09	(0.579)	3.85	(0.278)	-14.68***	(0.003)	-1.27	(0.759)	-20.25***	(0.000)
1977	3.08	(0.589)	3.29	(0.167)	-3.26	(0.541)	0.00***	(0.000)	0.00***	(0.000)
1983	11.05**	(0.015)	0.27	(0.868)	8.46*	(0.053)	0.61	(0.843)	-1.77	(0.543)
1989	12.69**	(0.028)	4.52	(0.141)	6.53	(0.233)	4.21	(0.274)	-2.55	(0.483)
1992	10.76**	(0.011)	7.03***	(0.001)	4.77	(0.218)	2.08	(0.502)	-4.42	(0.110)
1995	15.39***	(0.001)	8.88***	(0.000)	3.42	(0.408)	2.54	(0.478)	-9.45***	(0.003)
1998	25.78***	(0.000)	9.96***	(0.000)	-4.55	(0.309)	7.91**	(0.046)	-22.28***	(0.000)
2001	41.09***	(0.000)	15.64***	(0.000)	13.81**	(0.019)	16.11***	(0.002)	-11.99***	(0.005)
2004	39.02***	(0.000)	9.01***	(0.006)	18.15***	(0.002)	13.62***	(0.003)	-6.94*	(0.068)
2007	52.75***	(0.000)	20.38***	(0.000)	30.92***	(0.000)	18.62***	(0.000)	-1.60	(0.713)
2010	43.56***	(0.000)	20.24***	(0.000)	28.60***	(0.000)	14.17***	(0.003)	-0.09	(0.983)
2013	42.17***	(0.000)	21.06***	(0.000)	17.89***	(0.001)	17.14***	(0.000)	-7.11*	(0.051)
2016	56.91***	(0.000)	20.24***	(0.000)	26.01***	(0.000)	22.29***	(0.000)	-7.57*	(0.061)
N	89571		34297		86154		89571		86154	
R^2	0.255		0.135		0.264		0.270		0.278	

Notes: The dependent variable is net wealth. The controls include survey wave fixed effects, total household income, a full set of age dummies, a kids dummy, and an indicator for marital status. The "baseline" columns refer to the specification in (3.2), and "middle-class sample" presents the same regression for the middle 50-90% of the aggregate income distribution. The specification "+ stock exposure" includes the additional controls from (3.3), "+ bus. ownership" includes the additional controls from (3.4), and "+stock & bus." includes the additional controls from both (3.3) and (3.4). Multiply imputed observations were averaged for the regressions. p-values are given in parentheses (* p<0.1, ** p<0.05, *** p<0.01).

References

Aiyagari, S Rao (1994). "Uninsured Idiosyncratic Risk and Aggregate Saving." In: *The Quarterly Journal of Economics* 109.3, pp. 659–684.

Bartscher, Alina K., Moritz Kuhn, Moritz Schularick, and Ulrike I. Steins (2020). *Household Debt and Inequality in the U.S., 1950-2019*. Working Paper.

Benhabib, Jess, Alberto Bisin, and Mi Luo (May 2017). "Earnings Inequality and Other Determinants of Wealth Inequality." In: *American Economic Review* 107.5, pp. 593–97.

Bricker, Jesse et al. (2017). "Changes in U.S. Family Finances from 2013 to 2016: Evidence from the Survey of Consumer Finances." In: *Federal Reserve Bulletin* 103.3.

Cagetti, Marco and Mariacristina De Nardi (2006). "Entrepreneurship, Frictions, and Wealth." In: *Journal of Political Economy* 114.5, pp. 835–870.

— (2008). "Wealth Inequality: Data and Models." In: *Macroeconomic Dynamics* 12.S2, pp. 285–313.

Card, David and Thomas Lemieux (2001). "Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis." In: *The Quarterly Journal of Economics* 116.2, pp. 705–746.

De Nardi, Mariacristina and Giulio Fella (2017). "Saving and wealth inequality." In: *Review of Economic Dynamics* 26, pp. 280–300.

Eika, Lasse, Magne Mogstad, and Basit Zafar (2019). "Educational Assortative Mating and Household Income Inequality." In: *Journal of Political Economy* 127.6, pp. 2795–2835.

Emmons, William R., Ana Hernández Kent, and Lowell R. Ricketts (2018). "The Financial Returns from College across Generations: Large but Unequal." In: *Federal Reserve Bank of St. Louis Series "The Demographics of Wealth"*. no. 1.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri (2020). "Heterogeneity and Persistence in Returns to Wealth." In: *Econometrica* 88.1, pp. 115–170.

Fortin, Nicole M. (Sept. 2006). "Higher-Education Policies and the College Wage Premium: Cross-State Evidence from the 1990s." In: *American Economic Review* 96.4, pp. 959–987.

Goldin, Claudia and Lawrence F. Katz (2007). "Long-Run Changes in the Wage Structure: Narrowing, Widening, Polarizing." In: *Brookings Papers on Economic Activity* 38.2, pp. 135–168.

Greenwood, Jeremy, Nezih Guner, Georgi Kocharkov, and Cezar Santos (2014). "Marry Your Like: Assortative Mating and Income Inequality." In: *American Economic Review* 104.5, pp. 348–53.

— (2015). "Corrigendum to Marry Your Like: Assortative Mating and Income Inequality." In: *Correction*.

Haveman, Robert and Timothy Smeeding (2006). "The Role of Higher Education in Social Mobility." In: *The Future of Children* 16.2, pp. 125–150.

Huggett, Mark (1993). "The Risk-free Rate in Heterogeneous-Agent Incomplete-Insurance Economies." In: *Journal of Economic Dynamics and Control* 17.5-6, pp. 953–969.

Hurst, Erik and Annamaria Lusardi (2004). "Liquidity Constraints, Household Wealth, and Entrepreneurship." In: *Journal of Political Economy* 112.2, pp. 319–347.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor (2017). "Macrofinancial History and the New Business Cycle Facts." In: *NBER Macroeconomics Annual* 31.1, pp. 213–263.

Katz, Lawrence F. and David H. Autor (1999). "Changes in the Wage Structure and Earnings Inequality." In: *Handbook of Labor Economics*. Ed. by O. Ashenfelter and D. Card. Vol. 3. Handbook of Labor Economics. Elsevier. Chap. 26, pp. 1463–1555.

Kuhn, Moritz, Moritz Schularick, and Ulrike I. Steins (2020). "Income and Wealth Inequality in America, 1949–2016." In: *Journal of Political Economy* 128.9, pp. 3469–3519.

Levy, Frank and Richard J. Murnane (1992). "U.S. Earnings Levels and Earnings Inequality: A Review of Recent Trends and Proposed Explanations." In: *Journal of Economic Literature* 30.3, pp. 1333–1381.

Lusardi, Annamaria, Pierre-Carl Michaud, and Olivia S. Mitchell (2017). "Optimal Financial Knowledge and Wealth Inequality." In: *Journal of Political Economy* 125.2, pp. 431–477.

Lusardi, Annamaria and Olivia S. Mitchell (2011). "Financial literacy around the world: an overview." In: *Journal of Pension Economics and Finance* 10.4, pp. 497–508.

— (2014). "The Economic Importance of Financial Literacy: Theory and Evidence." In: *Journal of Economic Literature* 52.1, pp. 5–44.

Pfeffer, Fabian T. (2018). "Growing Wealth Gaps in Education." In: *Demography* 55.3, pp. 1033–1068.

Pfeffer, Fabian T., Robert F. Schoeni, Arthur Kennickell, and Patricia Andreski (2016). "Measuring Wealth and Wealth Inequality: Comparing Two U.S. Surveys." In: *Journal of Economic and Social Measurement* 41.2, pp. 103–120.

Pfeffer, Fabian T and Alexandra Killewald (2018). "Generations of Advantage. Multigenerational Correlations in Family Wealth." In: *Social Forces* 96.4, pp. 1411–1442.

Piketty, Thomas and Emmanuel Saez (2003). "Income Inequality in the United States, 1913-1998." In: *The Quarterly Journal of Economics* 118.1, pp. 1–39.

Reeves, Richard V. and Kimberly Howard (2013). "The Glass Floor: Education, Downward Mobility, and Opportunity Hoarding." In: *Research Report*.

Saez, Emmanuel and Gabriel Zucman (2016). "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data." In: *The Quarterly Journal of Economics* 131.2, pp. 519–578.

Schenker, Nathaniel and Jeremy M.G. Taylor (1996). "Partially parametric techniques for multiple imputation." In: *Computational Statistics & Data Analysis* 22.4, pp. 425–446.

Von Gaudecker, Hans-Martin (2015). "How Does Household Portfolio Diversification Vary with Financial Literacy and Financial Advice?" In: *The Journal of Finance* 70.2, pp. 489–507.