The Saving Glut of the Rich *

Atif Mian
Princeton & NBER

Ludwig Straub
Harvard & NBER

Amir Sufi
Chicago Booth & NBER

February 2021

Abstract

There has been a large rise in savings by Americans in the top 1% of the income or wealth distribution over the past 40 years, which we call *the saving glut of the rich*. Instead of financing investment, this saving glut has been associated with dissaving by the non-rich and dissaving by the government. An unveiling of the financial sector reveals that rich households have accumulated substantial financial assets that are direct claims on U.S. government and household debt. State-level analysis shows that the rise in top income shares has been important in generating the rise in savings by the rich.

*A previously circulated version of this study was entitled "The Saving Glut of the Rich and the Rise in Household Debt." We thank Ian Sapollnik for extraordinary research assistance throughout this project. Pranav Garg, Sebastian Hanson, and Bianca He also provided excellent research assistance at various stages. We are grateful to the following scholars who patiently answered questions on various conceptual and data issues: Jesse Bricker, Joseph Briggs, Jonathan Fisher, Fatih Guvenen, Jonathan Heathcote, David Johnson, Ralph Koijen, Eric Nielsen, Fabrizio Perri, Lukasz Rachel, Kamila Sommer, Alice Volz, Owen Zidar, and Gabriel Zucman. We also thank Heather Boushey, Benjamin Hébert, Greg Kaplan, Gianni La Cava, Lukasz Rachel, Moritz Schularick, Richard Thaler, Harald Uhlig, Stijn van Nieuwerburgh, Rob Vishny, and seminar participants at the Bank of England, the Bank of International Settlements, UC Berkeley, Brown, Chicago Booth, Columbia, Dartmouth, the Federal Reserve Bank of New York, the Federal Reserve Board of Governors, the IMF, the NBER Corporate Finance meeting, Princeton University, and the Reserve Bank of Australia. The replication kit (which is 16GB) for this study can be obtained by clicking here.

Contact info: Mian: (609) 258 6718, atif@princeton.edu; Straub: (617) 496 9188, ludwigstraub@fas.harvard.edu; Sufi: (773) 702 6148, amir.sufi@chicagobooth.edu

1 Introduction

A prominent idea in the history of economics is that an excess of savings relative to investment opportunities can present challenges for an economy (e.g., Hobson (1902), Keynes (1936), Eccles (1951)). A recent manifestation of this idea is the "global saving glut," or the argument that a rise in savings by certain countries is a primary culprit for the decline in interest rates to low levels around the world (e.g., Bernanke (2005), Summers (2014), Rachel and Smith (2017)).

This study shows an equally important saving glut of the rich. In the United States, the rise in savings by the top 1% of the income or wealth distribution since the 1980s has been substantial, it has not boosted investment, and it has been associated with dissaving by the government and the rest of the household sector. The saving glut of the rich has important implications for the concept of national savings, the evolution of the financial sector, and the demand for U.S. produced safe assets. It also warrants consideration in explanations of extremely low interest rates and high debt levels around the world.

The empirical analysis estimates savings across the distribution of U.S. households using two separate techniques. The first relies on the National Income and Product Accounts (NIPA) together with an estimation of income and consumption across the income distribution. The second relies on household wealth reported in the Financial Accounts of the Federal Reserve and the evolution of wealth across the wealth distribution over time. The results imply an increase in the savings by the top 1% of the distribution of 3 percentage points of national income annually when comparing the 1960s and 1970s with the post 2000 period.

To put this magnitude into perspective, the average annual savings by the top 1% of the distribution have been larger than average annual net domestic investment since 2000. The magnitude can also be compared to the increase in the average U.S. current account deficit since the 1980s, which represents savings entering the United States from other countries and is therefore connected to the global saving glut. The rise in savings by the top 1% of the distribution has been on the same order of magnitude as savings entering the United States from abroad.

The rise in savings by the top 1% has not been accompanied by a rise in net domestic investment. Instead the additional savings have been absorbed in a less conventional manner through dissavings by the government and the rest of the U.S. household sector. In particular, savings by the bottom 90% have fallen significantly since the early 1980s and the government has run larger deficits, especially after the Great Recession.

The rise in savings by the top 1% of the wealth distribution has been driven by a rise in the accumulation of financial assets. The rich have not been accumulating more real estate, nor have they been borrowing less. The dissaving of the bottom 90% has been driven by both a rise in borrowing and a decline in the accumulation of financial assets. The decline in the accumulation of financial assets has been the quantitatively more important driver of the dissaving by the bottom 90%.

To discover what financial assets the rich have been accumulating, we develop a methodology that "unveils" the financial sector. The basic idea is to remove the veil of financial intermediation to see who ultimately holds claims on financial assets such as household and government debt. This methodology allows for a quantification of how much of the rise in financial accumulation of the rich represents a claim on U.S. government and household borrowing. More generally, this methodology can be used to answer important questions of policy incidence such as who earns the debt service payments made by the government (e.g., Reinhardt (1945), Kalecki (1943), Hager (2016), Arbogast (2020)) or who lies behind the demand for U.S. dollar-denominated safe assets (e.g., Caballero et al. (2008), Caballero et al. (2017)). To the best of our knowledge, this unveiling process is novel to the literature, and potentially can be done for other asset classes and in other countries.

This methodology shows that almost two-thirds of the rise in financial asset accumulation of the top 1% of the wealth distribution since the 1980s has been a rise in the accumulation of claims on U.S. government and household debt. Furthermore, while the rest of the world has been an important financier of this debt, Americans in the top 1% of the wealth distribution have been almost equally important. From the 1980s until 2016, U.S. government and household debt increased annually by an average of 3 percentage points of national income more compared to the 1960s and 1970s. Half of this increase was financed by the rest of the world, and half was financed by the top 1%.

The unveiling exercise also allows for a calculation of net household debt positions across the wealth distribution, which we define as household debt held as a financial asset minus household debt owed as a liability. Net household debt positions clarify that rich Americans have increasingly financed the borrowing of non-rich Americans. From 1982 to 2007, the net household debt position of the bottom 90% fell by 39 percentage points of national income, while the net household debt position of the top 1% rose by 12 percentage points. That is, 30% of the rise in net household debt owed by the bottom 90% was financed by the top 1%.

The fact that rich Americans directly finance the borrowing of the non-rich and the government is important given a growing body of research showing that financial assets are not perfectly substitutable across investors (e.g., Greenwood et al. (2018), Koijen and Yogo (2019)). If financial investment is not perfectly fungible across assets, then the fact that the rich have accumulated substantial positions in household and government debt may have affected prices and quantities of such debt over time. In short, the asset positions of the wealthy matter.

It may be surprising that Americans in the top 1% of the wealth distribution hold so much U.S. government and household debt in their portfolios, but the unveiling process clarifies how this is in fact the case. For example, holdings of business equity by rich Americans represent a substantial claim on such debt. The reason is that non-financial businesses have increased their holdings of money market funds and time deposits substantially since the mid 1990s, and these time deposits and money market funds are claims on debt through the financial system. More generally, government and household debt have been financed by rich Americans through direct bond holdings, bond mutual fund holdings, business equity, time deposits, money market funds, and defined contribution pensions.

Analysis using a state-year level panel data set links the rise in savings by Americans at the top of the income distribution to the rise in top income shares. The state-level specification removes common aggregate patterns that occurred since the 1980s, and therefore brings us closer to the ideal thought experiment of examining economies with different shifts in top income shares while holding all else equal. The state-level specification allows us to control for other factors that may be related to the rise in savings by the rich, such as demographics, changes in the industrial shares of employment, and financial deregulation.

The core findings of this study are robust across a number of data sets. They are also robust to the various issues regarding the estimation of income and wealth shares that have been raised in the literature. The baseline analysis uses income shares from the Congressional Budget Office (CBO) and income and wealth shares from the Distributional National Accounts (DINA) microfiles made available publicly by Piketty et al. (2018). The baseline analysis using the DINA microfiles makes adjustments to the assumed interest rate earned by those in the top 1% of the wealth distribution (e.g., Bricker et al. (2018) and Smith et al. (2020)). The baseline analysis using the DINA microfiles also makes an adjustment to the imputation of pension wealth given issues raised by Auten and Splinter (2019). Wealth shares using the Distributional Financial Accounts (DFA), which are produced by the Federal Reserve and based on the Survey of Consumer Finances (SCF), are also used from 1989 to 2016 when such data are available. The estimates of savings by the top 1% using the DFA and the adjusted DINA microfiles are almost identical.[1] The rise in savings by the top 1% is a robust pattern that is not an artifact of a particular methodology or data set.

Implications The twin phenomena of rising debt and declining interest rates has occurred across many advanced economies since the early 1980s (e.g., Summers (2014), Jordà et al. (2016)). The analysis here suggests that the saving glut of the rich may be an important factor to consider when evaluating these trends.[2]

[1]The similarity of the results using the DFA and DINA microfiles is consistent with the observation in Bricker and Henriques (2020) that the rise in wealth inequality from 1989 to 2016 is similar when using the SCF and when using the original Saez and Zucman (2016) methodology with the DINA microfiles.

[2]This finding is related to the idea that rising inequality puts downward pressure on interest rates given high saving rates of the rich (e.g., Stiglitz (2016), Rachel and Smith (2017), Rachel and Summers (2019)). See Mian et al. (2021)

Furthermore, the rise in the demand for U.S. dollar-denominated safe assets has been attributed primarily to demand coming from other countries (e.g., Caballero and Krishnamurthy (2003), Gourinchas and Rey (2007), Caballero et al. (2008), Caballero and Krishnamurthy (2009)). This study shows that the top 1% within the United States have also increased their holdings of the debt claims that back these safe assets. Quantitatively, the demand for U.S. dollar-denominated safe assets comes almost as much from the rich within the United States as it does from the rest of the world.

The findings also imply that aggregated national measures of savings provide an incomplete portrait of how savings evolve within an economy. While national savings have fallen in the United States since the 1980s, savings by rich Americans have increased substantially. The aggregate savings rate is not evidence against the idea that there has been a saving glut of the rich.[3] The findings also call into question the idea in many macroeconomic models that a rise in savings automatically translates into additional capital formation. In the United States over the past 40 years, the substantial rise in savings by the top 1% has been associated with dissaving by the government and the bottom 90%, as investment actually fell.

This is related to the fundamental role of the financial sector. Since the 1980s, the financial sector have been growing (e.g., Philippon (2015)) while investment to national income ratios have been falling. The growth in the financial sector appears to be driven to a large degree by the channelling of savings from some households into borrowing by the government and borrowing by other households.

Related literature. Several studies in the wealth inequality literature conduct an exercise to estimate saving rates across the distribution (e.g., Saez and Zucman (2016), Kuhn et al. (2019), and Smith et al. (2020)). However, none of these studies use this exercise to answer the key question of this study: how has the total amount of savings by different parts of the income or wealth distribution evolved over time? This study is the first to place savings from different parts of the distribution into a national accounting framework and to explore where these savings ultimately settle. This study is also the first to directly link the financial asset accumulation of the rich to borrowing by the non-rich and to borrowing by the government.

There is also a growing literature focused on the rise in household debt in the United States. Most of this literature is focused on trends immediately before the Great Recession.[4] One exception is the recent working paper of Bartscher et al. (2020), which examines the rise in household debt since 1949 across the income distribution. Many of the results in Bartscher et al. (2020) are

---

for a formal model in which this is the case. Klein and Pettis (2020) argue that the rise in capital outflows by current account surplus countries is driven by a rise in savings by the rich in those countries, which suggests that the global saving glut and the saving glut of the rich are connected.

[3] A similar point is made by Pettis (2017).

[4] See, for example, Mian and Sufi (2015), Bhutta and Keys (2016), Mian and Sufi (2017), Adelino et al. (2018), Foote et al. (2016); and Albanesi et al. (2017)

complementary to the analysis here. For example, from 1983 to 2016, Bartscher et al. (2020) find no material change in the debt to income ratio of households in the top 1% of the income distribution, but a dramatic rise in the debt to income ratio of households in the bottom 90% of the income distribution. However, Bartscher et al. (2020) do not attempt to link the savings by the rich to the borrowing of the non-rich, which is a main focus of this study.[5]

Cynamon and Fazzari (2015) show evidence that the bottom 95% needed to borrow more after 1980 in order to keep consumption levels steady in the face of rising income inequality. A similar argument is made in Rajan (2011) and Bertrand and Morse (2016). In these studies, the emphasis is on an increase in credit demand by low income households because of lower income levels. Instead, this study emphasizes how an increase in credit supply coming from the top 1% may have contributed to higher debt levels of the bottom 90%, which helps explain why interest rates fell during this period. To motivate their model, Kumhof et al. (2015) show a number of stylized aggregate facts that are consistent with the idea that rising income inequality led to rising household debt in the years prior to the Great Recession. However, there is no attempt to directly link the rise in saving of the rich to the dissaving of the non-rich, as is done here. The state-level analysis linking top income shares to the rise in savings is novel to the literature.

2 Empirical Methodology

2.1 National accounting framework

The goal of the empirical analysis is to measure the contribution to aggregate savings from different parts of the distribution of U.S. households. The starting point of the methodology is national income.[6] Let Y be GDP, Z be National Income, C be personal consumption expenditures, G be government consumption, I be total gross domestic investment (which includes both government and private domestic investment), and $(X - M)$ be net exports. The standard GDP equation is:

$$Y = C + G + I + (X - M) \tag{1}$$

[5]The findings presented here are also related to the literature exploring consumption, income, and wealth inequality (e.g., Slesnick (2001), Krueger and Perri (2006), Blundell et al. (2008), Heathcote et al. (2010), Aguiar and Bils (2015), Attanasio and Pistaferri (2016), Meyer and Sullivan (2017), Guvenen et al. (2019b), Fisher et al. (2018), Wolff (2017), Bricker et al. (2018) Guvenen et al. (2019a), De Nardi et al. (2018), and Batty et al. (2019)). As shown below, this literature is an important input into the measurement of the consumption share of the top 1% over time.

[6]National income is preferred to GDP for measuring saving behavior because national income excludes the non-economic income item capital depreciation (or "consumption of fixed capital" as it is called in the national accounts). Furthermore, national income includes payments to U.S. owners of capital which is located abroad, and it excludes payments to foreign owners of capital which is located in the United States.

Let δ be consumption of fixed capital, and W be net income from abroad. The definition of national income is $Z = Y - \delta + W - \epsilon$. Then equation 1 can be written as:

$$Z - C = G + I^n + F - \epsilon \tag{2}$$

where $F = (X - M + W)$ is the current account of the United States, $I^n = I - \delta$ is net domestic investment and ϵ is the statistical discrepancy that equalizes Gross Domestic Income with Gross Domestic Product in the National Accounts. The term G is related to taxes T, transfers R, and government saving S^g through the government budget constraint $S^g = T - G - R$, which then allows us to write equation 2 as:

$$\Theta = Z - T + R - C = I^n + F - S^g - \epsilon \tag{3}$$

This is the definition of aggregate private savings (Θ): national income minus taxes plus transfers minus personal consumption expenditures. Notice that Account 6 of the NIPA gives us the equation: $S^p + S^\pi + S^g = I^n + F - \epsilon$. This gives us another definition of private savings:

$$\Theta = Z - T + R - C = S^p + S^\pi = I^n + F - S^g - \epsilon \tag{4}$$

This makes it clear that the definition of private savings (Θ) includes both personal savings (S^p) and business savings (S^π). As equation 4 shows, the notion of savings in national accounts used by NIPA does not recognize differences in savings across the distribution of households. If one household saves more by lending to another household, then there is no change in Θ.

This study measures Θ_{it}, which is the savings by group i of households in year t. By construction, $\sum_{i \in I} \Theta_{it} = \Theta_t$. There are two techniques used to measure Θ_{it}. The first approach uses the following definition:

$$\Theta_{it} = Z_{it} - T_{ti} + R_{it} - C_{it} \tag{5}$$

More specifically, Θ_{it} is pre-tax income minus taxes plus transfers minus consumption. This is referred to as the *income less consumption approach*.

The second approach relies on estimates of wealth and the consumer budget constraint that links savings to wealth accumulation, what we call the *wealth-based approach*. We start with the budget constraint:

$$Z_{it} - T_{it} + R_{it} - C_{it} = \sum_{j \in J} \left(P_t^j A_{it}^j - P_t^j A_{i,t-1}^j \right) \tag{6}$$

where $\Theta_{it} = Z_{it} - T_{it} + R_{it} - C_{it}$ is savings by group i and A_{it}^j is asset j held by the group i. There are a total of J asset types that households hold, with liabilities showing as negative values. Let $\pi_t^j \equiv \frac{P_t^j - P_{t-1}^j}{P_{t-1}^j}$ be asset price inflation for asset j. Then equation (6) simplifies to:

$$\Theta_{it} = \sum_{j \in J} \left(\Delta W_{it}^j - \pi_t^j W_{i,t-1}^j \right) \tag{7}$$

This approach relies on measures of wealth held by group i in asset class j in year t (W_{it}^j), and measures of pure asset price inflation π_t^j.

Appendix A.1 discusses the reasons why we do not use estimates of saving rates from survey data such as the Panel Study of Income Dynamics or the Survey Consumer Finances, and it also discusses Social Security fits into the estimates of Θ_{it}.

2.2 Data sets used

After-tax income shares are obtained from the Congressional Budget Office (e.g., CBO (2019)). Consumption shares are estimated using the Survey of Consumer Finances (SCF), following the methodology in Fisher et al. (2018) which is described in detail below in Section 2.3.

Data from the Distributional National Accounts (DINA) microfiles are also used in the analysis below (Piketty et al. (2018)). These microfiles rely on the yearly public-use tax return files available at the National Bureau of Economic Research, along with calculations to allocate national income that is not included as part of tax returns, as outlined by Piketty et al. (2018). The DINA microfiles are used to calculate shares of national income across the income distribution, and they are used to calculated wealth shares across the wealth distribution using the capitalization technique outlined in Saez and Zucman (2016) and Smith et al. (2020). Wealth shares for 1989 and afterward are also calculated using the Distributional Financial Accounts (DFA) produced by the Federal Reserve. However, the DFA are unavailable prior to 1989, which is a major limitation for the analysis here.[7]

There are two important adjustments that are made for the DINA microfiles relative to the methodologies outlined in Saez and Zucman (2016) and Piketty et al. (2018). The first relates to the distribution of pension income and pension wealth not captured on tax returns, an issue raised in Auten and Splinter (2019). The internal returns on undistributed pensions that are part of national income but are not reported on tax filings must be distributed across the income and wealth distribution. The baseline methodology in Piketty et al. (2018) uses realized taxable and non-taxable pension income as reported on tax returns to distribute this income.

However, Auten and Splinter (2019) argue that the non-taxable pension income reported on tax filings are actually rollovers from Individual Retirement Accounts, and are therefore not income but a rollover of wealth. As a result, the pension income of higher income Americans is over-estimated

[7]Results using the DFA for the post 1989 period are shown in Appendix Section D.

if one uses the non-taxable part of the pension income reported on tax filings. We follow Auten and Splinter (2019) and use only the taxable pension income reported on tax filings to estimate the undistributed pension income component of national income and pension wealth. This change affects both the distribution of national income and wealth, given that pension wealth is obtained by capitalizing the income earned on pensions. As shown in Auten and Splinter (2019), this adjustment brings the estimate of pension wealth across the income distribution in the DINA closer to the distribution in the SCF.[8]

The second adjustment relates to the capitalization of fixed income on tax returns to measure fixed income wealth across the wealth distribution (e.g, Bricker et al. (2018) and Smith et al. (2020)). Translating flows of income into stocks of wealth requires an assumption of the rate of return on assets, a process detailed in Saez and Zucman (2016). Research suggests that the baseline methodology in Saez and Zucman (2016) overstates the level of fixed income asset holdings of the top 1% given the assumption of a homogeneous rate of return on fixed income assets when estimating fixed income wealth from fixed income asset cash flows (e.g., Bricker et al. (2018) and Smith et al. (2020)).[9] This manifests itself in the assumed capitalization factor one uses to multiply the fixed income asset cash flows to obtain fixed income wealth.

Given this issue, the capitalization technique used in this study assumes that the top 1% of the wealth distribution earns a return on fixed income assets that is 1.3 times (1.3X) the overall rate of return on fixed income assets earned by U.S. households. This assumption is motivated by evidence from the SCF and estate tax filings in Bricker et al. (2018) and Saez and Zucman (2020). Appendix Figure A2 uses the SCF and shows the ratio of the top 1% of the wealth distribution's fixed income return to the fixed income return of all households from 2001 to 2019. The average is almost exactly 1.3X. The ratio is never above 1.6X in any year of the SCF.

We are confident that the results shown here are robust to issues related to the fixed income asset return of the top 1%. Appendix Table A8 shows the main results using both the DFA and the adjusted DINA microfiles for the post 1989 period for which both data sets are available. The results are similar. Recall that the DFA is based on the SCF and therefore does not have any issue regarding capitalization factors. Furthermore, Appendix Table A1 shows that the rise in savings by the top 1% would still be 2.5 percentage points of national income annually (as opposed to 2.9 percentage points) if a factor of 1.6X (which is the largest factor found in any year of the SCF from 2001 to 2019) were used instead of 1.3X.[10]

[8] Auten and Splinter (2019) make a number of other adjustments to the DINA methodology to obtain the striking result that the after-tax share of national income of the top 1% has not risen over the past 40 years. We have investigated each of these adjustments, and we have concluded that the pension adjustment is the only adjustment that is justified in our setting. The substantial rise in the after-tax share of income of the top 1% is also present in the Congressional Budget Office series.

[9] This point is acknowledged in Section IV.F of the original Saez and Zucman (2016) article.

[10] There is an active debate in the literature on whether the interest rate on fixed income assets for the top 1%

2.3 The income less consumption approach

Starting from equation 5, we use both the CBO and the DINA microfiles (adjusted for the pension issue as noted above) to measure the first three terms that define Θ_{it}: Z_{it}, T_{it}, and R_{it}. When constructing these variables using the DINA, we use the adult individual as the unit of observation and we split income equally among spouses. For the CBO income series, we take the CBO after-tax income shares reported and multiply the share by $Z_t - T_t + R_t$. As a result, the after-tax income amounts across the distribution for both the DINA and CBO add up to the same aggregate figure.[11]

The last component needed to measure savings by each income group (Θ_{it}) is consumption (C_{it}). Measurement of the consumption of the top of the income distribution is the most challenging aspect given the lack of a comprehensive data set focused on consumption of the rich. The approach taken here is to rely on two items: the share of consumption across the income distribution in a given baseline year, and an assumption of the evolution of the consumption to income ratio of the top 1% over time. We purposefully rely on conservative assumptions to generate the consumption share of the top 1% over time, given that the data are weakest on this particular item.

Three main groups are focus on in the analysis below: the top 1% of the income distribution, the next 9%, and the bottom 90%. Unfortunately, survey data sets typically used in the consumption literature such as the Consumer Expenditure Survey (CEX) or the Panel Study of Income Dynamics (PSID) do not measure the consumption of the highest income households in the economy well.[12]

Instead, we follow the analysis in Fisher et al. (2018), which uses the SCF to obtain consumption shares across the income distribution. The SCF has the advantage of having extensive coverage of high income and high wealth U.S. households. Since 2004, the SCF has also asked questions on expenditures on certain consumption categories. In particular, as Fisher et al. (2018) show, expenditures on food eaten at home, food eaten away from home, housing, new vehicle purchases, and used vehicle purchases can be measured using the SCF surveys from 2004 to 2016.

Fisher et al. (2018) use the CEX to show that the expenditure share of the goods reported in

should be assumed to be even higher (e.g., Smith et al. (2020), Saez and Zucman (2020)). Smith et al. (2020) use a methodology in which they assume that the interest rate earned by the top 0.1% of the fixed income distribution is the Moody's Aaa corporate rate, and the interest rate earned by 99.0 to 99.9th percentile of the fixed income distribution is the U.S. Treasury rate. In the analysis here, these assumptions would imply a fixed income return of the top 1% of the wealth distribution that is 2.0X the interest rate earned by the U.S. household sector. This is substantially higher than the highest factor ever observed for the top 1% of the wealth distribution in the SCF from 2001 to 2019, as shown in Appendix Figure A2. We follow the SCF because there are no issues measuring either wealth or the interest rate earned on fixed income assets for the top 1% of the wealth distribution in the SCF.

[11] Notice that the aggregate variable we are trying to match is $Z_t - T_t + R_t$, not national income. In Piketty et al. (2018), the post-tax shares of national income are designed to add up to $Z_t - T_t + R_t + G_t + S_t^g$, which given the government budget constraint adds up to Z_t.

[12] There is a large literature discussing the measurement of consumption by the rich in various surveys, but the CEX in particular. See for example, Aguiar and Bils (2015), Carroll et al. (2015), Attanasio and Pistaferri (2016), and Meyer and Sullivan (2017). The issue is two-fold: first, the very rich are typically not surveyed. Second, even conditional on being surveyed, the rich may under-report actual consumption more than the non-rich.

the SCF relative to total expenditures is stable across the income distribution, a result we have replicated. Although the CEX does not contain the richest U.S. households, the expenditure share on the categories reported in the SCF is stable even up to the richest households in the CEX. As a result, Fisher et al. (2018) use the consumption shares across the income distribution on goods reported in the SCF as the consumption share for all goods across the income distribution.[13]

The average consumption share of the top 1% of the income distribution using the Fisher et al. (2018) measure is 6.6% from 2004 to 2016, which we use for 2010 which is the mid-point of the years available from the SCF. How has this consumption share evolved over time? This is the most challenging variable to measure, as there is no data set that covers the consumption of rich U.S. households over the long time period of our sample, which is 1963 to 2016.

The consumption share of the top 1% over time is measured using the assumption that the consumption to income ratio of the top 1% has been constant over time, which would imply that consumption shares and income shares have increased at the same growth rate. This follows from the evidence in Aguiar and Bils (2015) that consumption inequality and income inequality have risen at a similar rate over time. It is important to recognize that this is a long-run assumption, as we are interested in measuring how consumption shares have evolved over a 50 year period.

The assumption of a constant consumption to income ratio over time is conservative, because it is likely that the consumption to income ratio falls in income even over long time periods (e.g., Straub (2019)).[14] If the consumption to income ratio of the top 1% has fallen over time, then the saving glut of the rich would be estimated to be even larger.[15] Figure A1 in the appendix plots the consumption share of the top 1% of the income distribution using this methodology. As it shows, the consumption share of the top 1% using this methodology has risen substantially over time, from 4 to 5% in the 1960s and 1970s, to 6 to 7% from 2010 to 2016.[16]

[13] Fisher et al. (2018) also use the SCF survey waves of 1989 to 2001 following a similar procedure. However, prior to 2004, the SCF survey waves did not contain a question on food purchased for consumption at home and food purchased for consumption away from home. Given the importance of spending on these goods, we use only the 2004 to 2016 waves in which spending on food at home and away from home can be measured.

[14] The average post-tax real income of the top 1% implied by their share of national income was $418 thousand in 1982 and $1.008 million in 2016 (in 2016 dollars). In contrast, the average post-tax real income of the bottom 90% increased from $29 thousand to $46 thousand. Given estimates in the literature, it is unlikely that the consumption to income ratio for the top 1% stayed constant given a rise in real income of 150%.

[15] We could also use the time series from the SCF from 2004 to 2016. The consumption share of the top 1% in the SCF is almost completely flat from 2004 to 2016, despite a rising in the share of income going to the top 1%. The assumption of a flat consumption share would significantly increase the size of the saving glut of the rich. However, we are hesitant to use this time series given the Great Recession occurred in the middle of this period. It is important to note that the assumed consumption to income ratios over time are meant to capture long-run trends as opposed to short-run changes due to cyclical factors. Heathcote and Perri (2018) show that such cyclical factors are important in explaining consumption to income ratios across the wealth distribution during recessions.

[16] It is possible to use the wealth-based approach described in the next subsection to calculate the implied consumption share of the top 1%. The wealth-based approach does not rely on any information on consumption to calculate savings, and it therefore acts a useful robustness exercise. The wealth-based approach implies a consumption share of the top 1% of the wealth distribution of 6.0% prior to 1983 and 6.8% afterward.

2.4 The wealth-based approach

Following equation 7, the wealth-based approach requires measurement of wealth held by group i in asset class j (W_{it}^{j}) and inflation for asset class j (π_t^j). For W_{it}^{j}, we first measure aggregate wealth in asset class j (W_t^j), and then we use an estimate of the share of aggregate wealth in asset class j held by group i in year t, which we call ω_{it}^{j}. Then, $W_{it}^{j} = \omega_{it}^{j} * W_t^j$.

The terms W_t^j come from the Table B.101 from the Financial Accounts by the Federal Reserve, which represents the aggregate balance sheet of the household and non-profit sector. The first two columns of Table A2 in the appendix show the name of each asset class j and its associated code in the Financial Accounts.[17] We exclude any wealth category that is exclusively associated with non-profits.[18]

Both the DINA microfiles and DFA are used to measure ω_{it}^{j}. This requires a mapping from the asset class categories in B.101 of the Financial Accounts to the DINA and the DFA. The full details of this mapping are in columns 5 and 6 of Appendix Table A2. The DFA data do not begin until 1989; as a result, the DINA micro files are used as the baseline to estimate ω_{it}^{j}. However, Section D of the appendix shows results on the level of savings for 1989 onward using the data from the DFA to estimate ω_{it}^{j}.

Finally, the wealth-based approach to measuring saving requires a π_t^j for each asset j. In theory, π_t^j refers to asset price inflation that is driven only by inflation or valuation effects. This is not the overall return on the asset, but only the pure asset price inflation. To use the consumer price analogy, π_t^j is the change in nominal value of the asset holding constant the "quality" and identity of the asset. For example, the asset should not change in terms of expected cash flows.

For housing assets, π_t^j is estimated using a repeat-sales house price index that controls for any changes in housing size or quality. The results shown in this paper use the Jorda-Schularick-Taylor Macrohistory Database for the house price index because of its longer coverage. However, the JST index is highly correlated with other repeat-sales indices, such as CoreLogic. In a robustness check, we also allowed π_t^j to vary by income cohort i by using an income-sorted zipcode-level house price index, but this did not change results materially.

For fixed income assets, π_t^j is equal to zero given the manner in which the Financial Accounts

[17] An alternative would be to use Table B.101.h which is exclusively the balance sheet of the household sector. However, there are two limitations associated with Table B.101.h. First, it does not start until 1987, and our analysis requires data back to 1963. Second, holdings of certain securities such as Agency GSE bonds by households versus non-profits cannot be separated, even in recent years. The ability to measure such holdings is crucial for the unveiling procedure described in Section 4.

[18] Conceptually, savings by non-profits are part of national savings and so should therefore be included. By using B.101, the methodology here implicitly assumes that the shares of non-profit assets held by households across the wealth distribution are the same as the shares of assets held directly by households. The non-profit share of financial assets is small and stable. From 1989 to 2016, the average ratio of financial assets held by non-profits to financial assets held by non-profits and households was 0.054, with a minimum of 0.042 and a maximum of 0.061.

are reported. However, in the case of debt, write downs must be taken into account, especially given the importance of debt write-downs during the Great Recession. Debt write-downs imply that π_t^j needs to incorporate a valuation gain for the borrower. In the absence of such an adjustment, the methodology would incorrectly conclude that borrowers saved part of their income to pay down debt. The likelihood of debt write-downs varies considerably by income group i, with lower income borrowers more likely to default and therefore experience a write-down. Therefore, π_t^j is calculated for home mortgages and consumer credit separately for the top 10% and bottom 90%, with the valuation terms being indexed as π_t^{ij}.

The terms $\pi_t^{ij} = 1 - WD_t^{ij}$ where WD_t^{ij} is the percentage of debt that is written down in a particular year for group i. WD_t^{ij} is estimated by first calculating net chargeoffs as a share of outstanding debt on bank balance sheets, separately for home mortgages and consumer credit. Since we know total outstanding debt in a given year, the net chargeoff ratio gives us the total amount of debt that is written down. We then distribute the written down debt to group i based on the fraction of total defaults accounted for by group i. This number is computed using zip code level data on defaults and average income of households living in a zip code.[19]

Finally, π_t^j must be estimated separately for corporate equities, assets which have within them other corporate equities (such as pensions), and non-corporate business equities. When estimating π_t^{equity}, the previous literature has used capital gains on equity. We depart from this assumption because π_t^{equity} is not the same as capital gains. The price of a share reflects savings done by the corporation on behalf of the shareholders. The appreciation in the price of the share that reflects such retained earnings should not be reflected in π_t^{equity}, as this is actual saving. Using capital gains as the measure of π_t^{equity} will mistake the rise in corporate savings as pure asset price inflation.

Furthermore, a share may generate a yield, which shows up as Z_{it} in equation 6, through either dividends or share buybacks. In the case of the latter, this is because buybacks change the quality of the asset – a share after a buyback cannot be considered equivalent to a share before the buyback. For these reasons, the typical share price gain is not the same as π_t^{equity}. An additional complication that arises in calculating π_t^{equity} is that the observed equity wealth, such as the one reported in the Financial Accounts of the Federal Reserve, is itself imputed using various valuation metrics. As such, what we really need is the "inflation" in these valuation metrics over time.

Given all these considerations, estimating π_t^{equity} is a challenge. However, given that we have estimated π_t^j for all other asset types, π_t^{equity} can be calculated as the residual pricing factor that ensures that aggregate private savings calculated using the wealth-based approach matches the aggregate private savings in NIPA ($S^p + S^\pi$ in equation 4 above). The resulting π_t^{equity} has a correlation with the equity capital gains series from the JST Macrohistory Database of 0.90. However, the mean of π_t^{equity} is 2.3 percentage points lower. The lower mean is exactly what should be expected given

[19]The complete details of this methodology are shown in Appendix Section A.4.

that capital gains includes corporate savings, which have been high during the sample period and should not be included in π_t^{equity}.

The wealth-based approach to calculation of savings across the distribution is also implemented in Saez and Zucman (2016), Kuhn et al. (2019), and Smith et al. (2020). There are two features of the methodology here that are different. First, by using total private savings from the NIPA to pin down π_t^{equity} instead of capital gains, the methodology appropriately captures corporate savings that accrue to households. This has a material effect on savings given the large increase in corporate savings in the United States over time (e.g., Chen et al. (2017)). Second, debt write-downs are modelled here as a valuation gain instead of as active saving. This more accurately captures the nature of debt write-downs, which are not "saving" in the sense of earned income being used to pay back debt.

The DINA micro-files are a repeated cross-section, not a panel of individuals. This does not present any issues for the income less consumption approach to calculation of savings of Section 2.3. For the income less consumption approach, savings in a given year are calculated as the savings by individuals that are in the top 1% in the same year. These need not be the same individuals over time.[20]

However, the repeated cross-section nature of the DINA files introduces an issue for the wealth-based approach, given that the wealth-based approach relies on changes in asset values from the past year to this year. For example, consider the group of individuals that are in the top 1% of the income distribution. For the top 1% in a given year t, the change in the assets held from year $t-1$ to t for this specific group is not possible to recover given the repeated cross-section nature of the data. As a result of this problem, the wealth-based methodology follows the literature (Saez and Zucman (2016), Kuhn et al. (2019), and Smith et al. (2020)) by sorting individuals by wealth instead of income. The logic of this decision is to try to reduce the amount of migration by individuals across groups; an individual's place in the wealth distribution is more stable than the individual's place in the income distribution over time.[21]

[20]There is no panel data set that covers the very top of the income or wealth distribution, and so using a repeated cross-section is the only option for the analysis here. This may be a concern if there has been a rise in transitory income shocks so that there is more across-group switching over time. However, Kopczuk et al. (2010) use the Social Security Administration data to show that "all of the increase in the variance of annual (log) earnings since 1970 is due to an increase in the variance of permanent earnings (as opposed to transitory earnings)." See also Guvenen et al. (2019b).

[21]In their Appendix C, Kuhn et al. (2019) use the PSID to explore how often households transition across wealth percentile groups, and they find that "the share of households that remain within their respective wealth group ... is generally high."

3 Top 1% Savings and their Absorption

3.1 Top 1% savings

Figure 1 shows average annual savings by the top 1% using the income less consumption approach with the DINA microfiles, the income less consumption approach with the CBO after-tax income shares, and the wealth-based approach. For the income less consumption approaches, savings by the top 1% of the income distribution are shown; for the wealth-based approach, savings by the top 1% of the wealth distribution are shown. Annual savings are scaled by aggregate national income, and the five-year averages are shown. The 1978 to 1982 period average savings is subtracted for each series. All three measures indicate a substantial increase in savings by the rich from the 1980s onward. The average increase is similar across all three measures.

Figure 1: Average Annual Savings by the Top 1% over Time

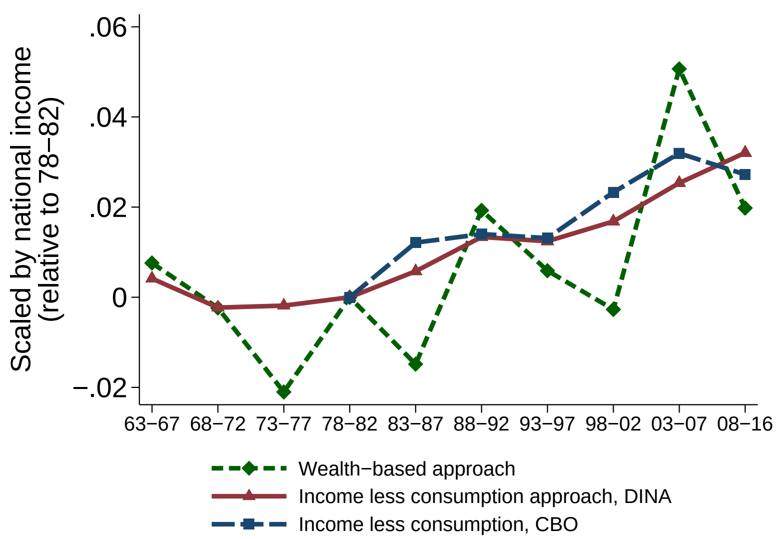

The DINA and CBO measures represent savings by the top 1% of the income distribution measured using the income less consumption approach, where income shares come from the Distributional National Accounts and the Congressional Budget Office, respectively. The wealth-based approach measures savings by the top 1% of the wealth distribution. For all three series, annual savings are scaled by annual national income, and five year averages are plotted. The 1978 to 1982 period is subtracted for all series.

Table 1 presents average savings by the rich over time. The breaks in time chosen in Table 1 capture the main macroeconomic episodes of the time period. The rise in top income shares and household debt began in the early 1980s; we choose 1983 as the initial breakpoint to avoid issues related to the recessions of 1980 and 1981-1982. The breakpoint in 1998 is meant to capture the period in which house price growth and household debt accelerated substantially. The breakpoint in 2008 captures the momentous Great Recession and its aftermath. For transparency, the full time series is always shown in addition to means by these four periods.[22]

Table 1: Savings by the Top 1%

	Levels			Relative to 63-82		
Period	DINA	CBO	Wealth-based	DINA	CBO	Wealth-based
63-82	0.053	0.036	0.026	0.000	0.000	0.000
83-97	0.064	0.049	0.034	0.011	0.013	0.007
98-07	0.075	0.064	0.054	0.021	0.028	0.028
08-16	0.085	0.063	0.050	0.032	0.027	0.024

The DINA and CBO measures represent savings by the top 1% of the income distribution measured using the income less consumption approach, where income comes from the Distributional National Accounts and the Congressional Budget Office, respectively. The wealth-based measure uses the wealth-based approach to measure savings by the top 1% of the wealth distribution. Annual savings are scaled by annual national income, and averages are shown for each period. For the CBO columns, the initial period includes only data from 1979 to 1982, as 1979 is the first year the data are available.

The level of savings by the top 1% differ across the three methodologies, but the rise in savings relative to the 1963-1982 period is similar. All three techniques imply a rise in annual savings scaled by aggregate national income of about 3 percentage points when comparing the 1960s and 1970s to the post 2000 period.[23]

3.2 Are savings absorbed through traditional channels?

Equation 4 above can be rewritten as follows:

$$\Theta_{top1,t} = I_t^n + F_t - S_t^g - \Theta_{next9,t} - \Theta_{bot90,t} \tag{8}$$

The saving glut of the rich could have been associated with a rise in net domestic investment I^n or it could have been associated with funds sent to other countries (F). It could also have been associated with dissaving by the government ($-S^g$) or dissaving by other households ($-\Theta_{next9,t} - \Theta_{bot90,t}$).

[22] The breakpoints are similar to those used in Bartscher et al. (2020), who call the years between 1965 to 1983 the "stability" period, the 1983 to 2007 period the "second debt boom", and the years between 2007 to 2016 the period of "crisis and deleveraging."

[23] Appendix Section A.6 shows that it is unlikely that population aging explains the rise in savings by the top 1%. A discussion of the implied saving rate of the top 1% under the income less consumption approach and the wealth-based approach is located in Appendix Section A.7. As we show there, the implied saving rate of the top 1% is consistent with the existing literature once missing forms of income with a high saving rate are taken into account.

Figure 2: Net Domestic Investment, Current Account Surplus, and Government Saving

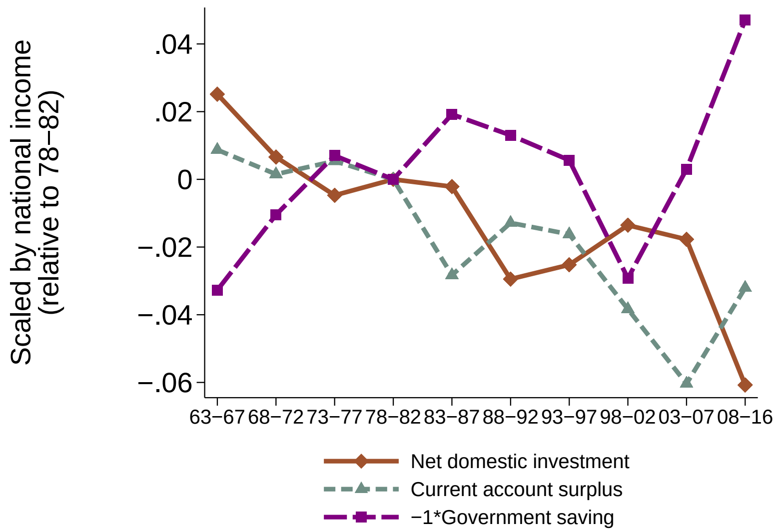

This figure plots net domestic investment, the current account surplus, and government saving. For government saving, the plot is multiplied by -1. For all three series, annual flows are scaled by annual national income, and five year averages are plotted. The 1978 to 1982 period is subtracted for all series.

Figure 2 plots net domestic investment, the current account surplus, and government saving over time. Government saving is multiplied by negative one and so an increase represents dissaving by the government ($-S^g$). As the figure shows, net domestic investment and the current account position have moved in the opposite direction as would have been required to absorb savings by the top 1%. Government borrowing has been more volatile. The government saved less from 1983 to the early 1990s, followed by a period in which the government began to run surpluses in the late 1990s. However, after the Great Recession, government saving has decreased substantially, representing almost 5 percentage points of national income less than the baseline period 1978 to 1982.

Table 2: Traditional Channels of Absorption

	Levels			Relative to 63-82		
Period	I^n	F	$-S^g$	I^n	F	$-S^g$
63-82	0.114	0.003	0.030	0.000	0.000	0.000
83-97	0.088	-0.020	0.051	-0.026	-0.023	0.022
98-07	0.091	-0.050	0.026	-0.022	-0.053	-0.004
08-16	0.046	-0.033	0.086	-0.068	-0.036	0.056

This table shows net domestic investment I^n, the current account surplus F, and government saving S^g. For government saving, $-S^g$ is shown. For all three series, annual flows are scaled by annual national income, and averages are plotted for the period in question.

Table 2 shows both the level and changes relative to the 1963-1982 period for these same three variables. Net domestic investment has fallen dramatically since 2008; from 2008 to 2016, net domestic investment has been 7 percentage points of national income lower than the pre-period. Table 2 also shows that the government has been dissaving 5.6 percentage points of national income from 2008 to 2016 compared to the pre-period.

A comparison of the evolution of the current account position in the fifth column of Table 2 with the evolution of top 1% savings in Table 1 allows for a direct comparison of the global saving glut (e.g., Bernanke (2005)) and the saving glut of the rich. Since 1982, the saving glut of the rich has been between 60 and 75% the size of the global saving glut.

3.3 Dissaving by the bottom 90%

The main variable that has moved contemporaneously with the rise in savings by the top 1% has been a decline in savings by the bottom 90%. This pattern is shown in Figures 3a, 3b, and 3c. Annual savings by households in the 90th to 99th percentile (what we call the next 9%) have been steady over time.

Figure 3: Saving across the Distribution

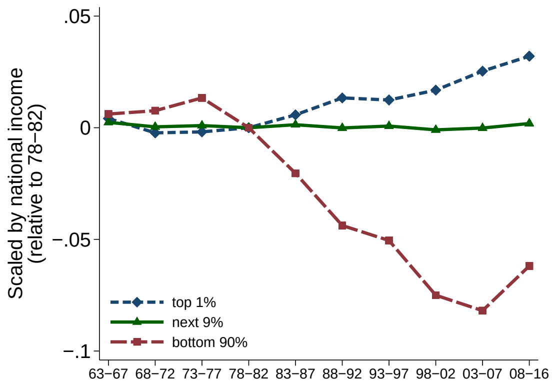

(a) Income less consumption, DINA

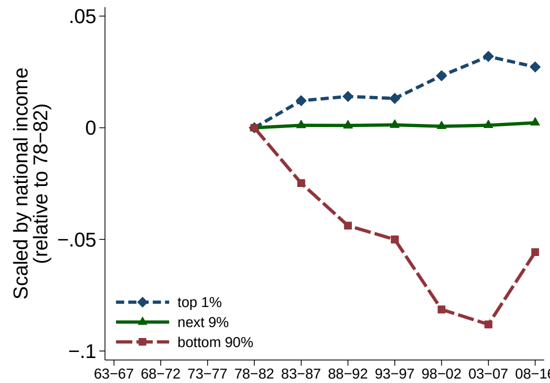

(b) Income less consumption, CBO

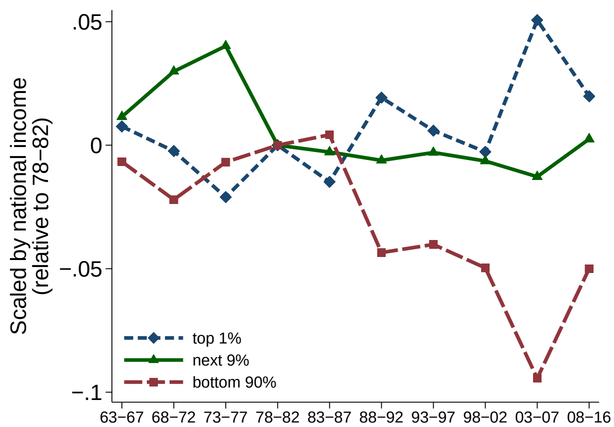

(c) Wealth-based approach

The DINA and CBO measures represent the savings across the income distribution measured using the income less consumption approach, where income shares come from the Distributional National Accounts and the Congressional Budget Office, respectively. The wealth-based approach measures savings across the wealth distribution. For all series, annual savings are scaled by annual national income, and five year averages are plotted. The 1978 to 1982 period is subtracted for all series.

In contrast, there has been a large decline in the savings by the bottom 90% of either the income or wealth distribution. Table 3 reports the average annual savings by the next 9% and the bottom 90%. As before, the DINA and CBO columns are based on the income less consumption approach, whereas the last column is based on the wealth-based approach. Across all three methodologies, the bottom 90% has saved substantially less post 2000 relative to the 1963 to 1982 period.

Table 3: Dissaving by the Next 9% and the Bottom 90%

Next 9%						
	Levels			Relative to 63-82		
Period	DINA	CBO	Wealth-based	DINA	CBO	Wealth-based
63-82	0.045	0.021	0.066	0.000	0.000	0.000
83-97	0.045	0.022	0.041	-0.000	0.001	-0.024
98-07	0.044	0.022	0.036	-0.001	0.001	-0.030
08-16	0.046	0.023	0.048	0.001	0.002	-0.018

Bottom 90%						
	Levels			Relative to 63-82		
Period	DINA	CBO	Wealth-based	DINA	CBO	Wealth-based
63-82	0.041	0.073	0.046	0.000	0.000	0.000
83-97	-0.004	0.033	0.028	-0.045	-0.040	-0.018
98-07	-0.045	-0.012	-0.017	-0.085	-0.085	-0.063
08-16	-0.028	0.017	0.005	-0.069	-0.056	-0.041

The DINA and CBO measures represent the savings by the top 1% of the income distribution measured using the income less consumption approach, where income comes from the Distributional National Accounts and the Congressional Budget Office, respectively. The wealth-based measure uses the wealth-based approach to measure savings. Annual savings are scaled by annual national income, and averages are shown for each period. For the CBO columns, the initial period includes only data from 1979 to 1982, as 1979 is the first year the data are available.

This decline in savings by the bottom 90% has been significantly larger than the increase in savings by the top 1%. This reflects the fact that both the global saving glut and the saving glut of the rich have increased substantially after 1982, and net domestic investment has actually fallen. Both the influx of foreign capital and the rise in savings by the top 1% have been associated with a large decline in savings by the bottom 90%.

Figure 4 accumulates all of these margins of absorption of savings by the top 1%. Starting with equation 8, we re-arrange to obtain:

$$\Theta_{top1,t} + \Theta_{bot99,t} + S_t^g - I_t^n - F_t + \epsilon_t = 0$$

For each of the 6 variables, we construct $\hat{V_t} = V_t - V_{pre}$, where V_{pre} is defined to be the average of variable V in the 10 years prior to 1983. Then for each variable we sum across all t to obtain $\bar{V} = \sum_{t=1983}^{2016} \hat{V_t}$ where $\bar{V}$ is the accumulation of the differences relative to the pre-period average. Therefore,

$$\bar{\Theta}_{top1} + \bar{\Theta}_{bot99} + \bar{S}^g - \bar{I}^n - \bar{F} + \bar{\epsilon} = 0 \tag{9}$$

Equation 9 implies that the accumulated savings by the top 1% ($\bar{\Theta}_{top1}$) must be absorbed by one of the other five terms.

Figure 4: Absorption of Accumulated Savings by the Top 1%

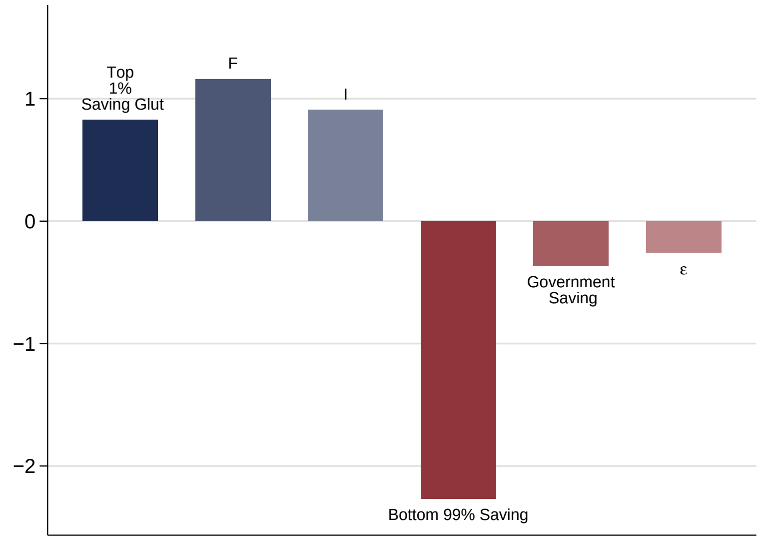

This figure presents the accumulated differences relative to the averages of the 1973 to 1982 levels in the equation: $\bar{\Theta}_{top1} + \bar{\Theta}_{bot99} + \bar{S^g} - \bar{I^n} - \bar{F} + \bar{\epsilon} = 0$. These terms represent, respectively, savings by the top 1%, the bottom 99%, and the government, in addition to net domestic investment, capital outflows, and the statistical discrepancy. All annual flow measures are scaled by contemporaneous aggregate national income before integration. The Θ_i terms come from the wealth-based approach to calculation of savings.

Figure 4 shows the accumulation of each of the six variables in equation 9. By construction, the bars sum to zero. The accumulated savings by the top 1% of the wealth distribution have been almost as large as aggregate national income from 1983 to 2016. Capital flows and investment have moved in the opposite direction as would have been needed to absorb some of these savings. To maintain the accounting identity, the combined savings by both the government and the bottom 99% must have fallen substantially. Figure 4 shows that most of the decline in saving was by the bottom 99%.

3.4 Decomposing savings

The wealth-based approach allows for a detailed look into the drivers of changes in savings across the wealth distribution. For Table 4, we start with equation 7 and let $\Delta NW_{i,t} = \sum_{j \in J} \Delta W^j_{i,t}$ be the annual change in net worth of wealth group i, $\Delta V_{i,t} = \sum_{j \in J} \pi^j_t W^j_{i,t-1}$ be the valuation effect, and $\Theta_{i,t}$ be savings. By equation 7, $\Theta_{i,t} = \Delta NW_{i,t} - \Delta V_{i,t}$. These three terms are calculated for each year, and scaled by national income in that year.

Table 4 shows the averages by period for these three terms. For the bottom 90%, the results show that strong valuation gains on assets (and in particular housing) masked a substantial decline in savings prior to the Great Recession. Average annual changes in net worth for the bottom 90% were relatively steady when comparing the 1983 to 1997 period with the 1963 to 1982 period. There was a slight decline in average annual changes in net worth from 1998 to 2007.

However, from 1998 to 2007, strong valuation gains in the housing market in particular masked a more dramatic decline in savings. The annual flow of savings by the bottom 90% fell by 6.3 percentage points of national income from 1998 to 2007 compared to the 1963 to 1982 period. When the housing market collapsed, so did the effect of valuation gains on net worth, which resulted in a much lower annual change in net worth during the 2008 to 2016 period relative to the pre-period.

These findings are consistent with those in Kuhn et al. (2019) who use a completely different data set (the SCF+) and find a similar result. Kuhn et al. (2019) write that "price effects account for a major part of the wealth gains of the middle class and the lower middle class," and that these price effects are driven by house price gains. They show that fixing house prices would lead to a substantial decline in the wealth share of households in the 50th to 90th percentile of the wealth distribution from 1989 to 2007, which is related to the finding here that high valuation gains masked a decline in savings for the bottom 90% during this period.

Table 4: Valuation Effects

Top 1%						
	Levels			Relative to 63-82		
Period	Θ	Δ NW	Δ V	Θ	Δ NW	Δ V
63-82	0.026	0.047	0.021	0.000	0.000	0.000
83-97	0.034	0.075	0.041	0.007	0.028	0.020
98-07	0.054	0.100	0.046	0.028	0.053	0.025
08-16	0.050	0.066	0.016	0.024	0.019	-0.005

Next 9%						
	Levels			Relative to 63-82		
Period	Θ	Δ NW	Δ V	Θ	Δ NW	Δ V
63-82	0.066	0.103	0.037	0.000	0.000	0.000
83-97	0.041	0.095	0.054	-0.024	-0.007	0.017
98-07	0.036	0.098	0.062	-0.030	-0.005	0.025
08-16	0.048	0.062	0.015	-0.018	-0.040	-0.023

Bottom 90%						
	Levels			Relative to 63-82		
Period	Θ	Δ NW	Δ V	Θ	Δ NW	Δ V
63-82	0.046	0.097	0.052	0.000	0.000	0.000
83-97	0.028	0.098	0.070	-0.018	0.001	0.019
98-07	-0.017	0.086	0.103	-0.063	-0.012	0.051
08-16	0.005	0.032	0.028	-0.041	-0.065	-0.024

This table decomposes average annual savings by each of the wealth groups ($\Theta_{i,t}$) into changes in net worth and valuation gains: $\Theta_{i,t} = \Delta NW_{i,t} - \Delta V_{i,t}$. Annual changes are scaled by national income, and then the average is taken for the period in question.

The wealth-based approach also allows us to explore the precise manner in which wealth groups altered savings over time. To do so, we start again with equation 7. We split asset classes j into three groups: financial assets (FA), real estate (RE), and debt (D). This allows us to decompose savings by each group i into: $\Theta_{it} = \Theta_{it}^{FA} + \Theta_{it}^{RE} + D_{it}$, where negative values of D indicate more borrowing.

Table 5 decomposes savings according to this equation for each wealth group. For the top 1%, the rise in savings is driven entirely by a rise in savings into financial assets (Θ^{FA}). Since 1982, the rich have not been paying down debt or accumulating real estate.[24]

[24]It is important to note that B.101 of the Financial Accounts places all tenant-occupied housing owned by the household sector into the business sector, which is owned by households through their holdings of business equity. As a result, any accumulation of tenant-occupied housing shows up as financial asset accumulation, not accumulation of

Table 5: Decomposing Savings

Top 1%								
	Levels				Relative to 63-82			
Period	Θ	Θ^{FA}	Θ^{RE}	D	Θ	Θ^{FA}	Θ^{RE}	D
63-82	0.026	0.025	0.002	-0.001	0.000	0.000	0.000	0.000
83-97	0.034	0.033	0.003	-0.003	0.007	0.008	0.001	-0.002
98-07	0.054	0.056	0.003	-0.004	0.028	0.031	0.001	-0.003
08-16	0.050	0.048	0.002	0.001	0.024	0.022	-0.001	0.002
Next 9%								
	Levels				Relative to 63-82			
Period	Θ	Θ^{FA}	Θ^{RE}	D	Θ	Θ^{FA}	Θ^{RE}	D
63-82	0.066	0.062	0.009	-0.006	0.000	0.000	0.000	0.000
83-97	0.041	0.041	0.009	-0.008	-0.024	-0.022	-0.001	-0.002
98-07	0.036	0.040	0.007	-0.012	-0.030	-0.022	-0.002	-0.006
08-16	0.048	0.045	0.005	-0.003	-0.018	-0.017	-0.004	0.003
Bottom 90%								
	Levels				Relative to 63-82			
Period	Θ	Θ^{FA}	Θ^{RE}	D	Θ	Θ^{FA}	Θ^{RE}	D
63-82	0.046	0.061	0.025	-0.040	0.000	0.000	0.000	0.000
83-97	0.028	0.056	0.018	-0.045	-0.018	-0.005	-0.007	-0.005
98-07	-0.017	0.038	0.018	-0.073	-0.063	-0.023	-0.007	-0.033
08-16	0.005	0.022	-0.002	-0.016	-0.041	-0.039	-0.026	0.024

This table decomposes the average annual savings by the top 1%, next 9% the bottom 90% through the equation $\Theta_{it} = \Theta_{it}^{FA} + \Theta_{it}^{RE} + D_{it}$. Θ_{it}^{FA} is asset accumulation in financial assets, Θ_{it}^{RE} is asset accumulation in real estate, and D_{it} is the change in debt (where a negative D represents more borrowing). Annual savings are scaled by national income, and then the averages for the period are calculated.

For the bottom 90%, the dissaving relative to the pre-period has been driven by both more borrowing (D) and less financial asset accumulation (Θ^{FA}). This is especially true for the 1998 to 2007 period. Comparing 1998 to 2007 to the 1963 to 1982 pre-period, the decline in savings by the bottom 90% coming from borrowing was only 1 percentage point of national income larger than the decline coming from reduced financial asset accumulation (3.3 versus 2.3 percentage points of national income annually). The rise in household debt during the 1998 to 2007 period has been the focus of a large body of research, and indeed it was quite substantial as shown in Table 5 (7.3 percentage points of national income annually). However, the results in Table 5 show that the dissaving by the bottom 90% relative to the 1963 to 1982 period was driven to a significant degree

real estate.

by lower financial asset accumulation.

The importance of lower financial asset accumulation by the bottom 90% can also be seen in the 2008 to 2016 period. Relative to previous years, the bottom 90% have borrowed much less per year in the last eight years of the sample, but they have continued to reduce financial asset accumulation substantially. To the degree that lower savings by the bottom 90% are a concern, the concern in recent years is not aggressive debt accumulation, but instead it is lower financial asset accumulation.

It is interesting to note that the financial asset accumulation patterns for the top 1% and bottom 90% are exactly the opposite. The top 1% are accumulating more financial assets while the bottom 90% are accumulating fewer financial assets. In addition, the top 1% may be directly financing the borrowing by the bottom 90% and the government.

4 Unveiling the Financial System and Saving in Debt

The top 1% of the wealth distribution has substantially increased financial asset accumulation since the early 1980s, while the bottom 90% and the government have been borrowing more. Has there been a direct link between the two patterns? This section describes an "unveiling" process that allows for a quantification of how much of the rise in financial asset accumulation of the rich has been a rise in accumulation of claims on household and government debt.

More formally, the exercise conducted in this section will allow us to split financial asset accumulation (Θ_{it}^{FA}) across the wealth distribution into accumulation of assets that are direct claims on household (Θ_{it}^{HHD}) and government debt (Θ_{it}^{GOVD}). Starting with $\Theta_{it} = \Theta_{it}^{FA} + \Theta_{it}^{RE} + D_{it}$, the goal is to further break down Θ_{it}^{FA} into:

$$\Theta_{it}^{FA} = \Theta_{it}^{HHD} + \Theta_{it}^{GOVD} + \Theta_{it}^{RSD}$$

where Θ_{it}^{RSD} is saving in residual financial assets not directly linked to household and government debt.

While a large body of research explores who owes debt as a liability, the methodology outlined here is designed to measure who saves in assets that are ultimately a claim on household debt and government debt. This is an important exercise given the growing body of research that demand for specific assets determines asset prices and asset quantities (e.g., Koijen and Yogo (2019)). This is closely related to the idea that financial markets are segmented (e.g., Greenwood et al. (2018)). Financial investment is not perfectly fungible across all asset classes, and so the specific holdings of rich Americans may matter for both the price and quantity of debt outstanding.

Another motivation is the large rise in both government and household debt in the United States. From 1980 to 2016, the combination of household and government debt in the United States scaled by GDP rose by almost 100 percentage points.[25] The reasons and implications for this expansion have been the focus of a large literature in macroeconomics and finance. Understanding who exactly owns the claims that are backed by U.S. government and household debt can help us answer important questions such as: who ultimately receives the interest and principal payments made on government debt (e.g., Reinhardt (1945), Kalecki (1943), Hager (2016), Arbogast (2020))? Or, who is behind the rise in the demand for U.S. dollar-denominated safe assets (e.g., Caballero et al. (2008), Caballero et al. (2017))?

4.1 Unveiling the financial system

In general, the rich do not directly lend to the non-rich. Instead, they hold a variety of assets which ultimately finance borrowing by others. For example, the rich are heavy investors in money market and mutual funds. These money market and mutual funds have sizable holdings of Agency Government-Sponsored Enterprise (Agency GSE) debt. Agency GSE debt is ultimately backed by home mortgages. For government debt, the rich directly hold U.S. Treasuries and municipal bonds; but they also lend to governments through other holdings of other assets. The unveiling exercise described here allows us to quantify how much U.S. government and household debt the rich hold through assets in their portfolio.

Overview

In the most general form, for each group i, the methodology explained here is designed to measure the total amount of an asset which is held by i at time t, or A_{it}. In total, there are I disjoint groups of households. Going forward, for simplicity, the time subscript t is dropped when writing specific elements of a matrix or vector, with the understanding that everything is measured for each year t. The time subscript t is kept when denoting entire matrices and vectors. The two assets unveiled in this study are U.S. government and household debt. The unveiling process is done separately for these two assets.

An asset can be held by households through various financial asset types, indexed by j. These are the same asset classes already discussed in Section 2.4 above. In particular, there are J total classes through which the asset is held, and these include pensions, mutual funds, time deposits, and GSE securities, to name a few examples. We call $F_{j,t}$ the total amount of the asset that is held through class j by households, which we calculate through an unveiling process, described in detail below. Group i's share of asset class j is ω_i^j (as in Section 2.4), and the assumption is made that group i's share of household debt held through asset class j is also ω_i^j.[26]

[25] See Figure A3 in the appendix.

[26] It may be the case that this is not true. For example, if top 1% households own riskier mutual funds than next 9%

Therefore,

$$\begin{bmatrix} A_1 \\ A_2 \\ \vdots \\ A_I \end{bmatrix} = \begin{bmatrix} \omega_1^1 & \omega_1^2 & \cdots & \cdots & \omega_1^J \\ \omega_2^1 & \omega_2^2 & \cdots & \cdots & \omega_2^J \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ \omega_I^1 & \omega_I^2 & \cdots & \cdots & \omega_I^J \end{bmatrix} \begin{bmatrix} F_1 \\ F_2 \\ \vdots \\ F_J \end{bmatrix}$$

or equivalently

$$\underset{I\times 1}{A_t} = \underset{I\times J}{\Omega_t} \times \underset{J\times 1}{F_t}. \tag{10}$$

The unveiling exercise can therefore be separated into two steps. The first step is to calculate the total amount of the asset in question held by households in each Financial Accounts asset class j to get the vector F_t. The second step is constructing the matrix Ω_t, which contains the shares of each asset j held by group i in year t. The second step is already described in Section 2.4 above, and so we focus only on the first step below.

Unveiling to obtain F_t

Figure 5 presents a graphical overview of the unveiling methodology. It focuses on household debt, but the process for government debt is similar. In the left most column, the figure starts with total household debt (home mortgages plus consumer credit) in the Federal Reserve Financial Accounts owed by the U.S. household sector as of 2005, scaled by national income. This happens to be almost exactly 100% of national income.

The end result in the right most column is the holdings of household debt as a financial asset by the U.S. household sector, by asset class (F_t, where t is 2005). The total amount held by the U.S. household sector is lower than the total amount owed by the U.S. household sector because the U.S. government and the rest of the world also hold household debt owed by the U.S. household sector. The unveiling process also allows us to measure how much the rest of the world and the U.S. government own of household debt.

Moving from left to right represents each round of unveiling. An arrow going into the box represents where the household debt comes from, and the arrows going out of the box represent where it goes to. The number in each box represents the total amount of household debt scaled by national income held in this category after taking into account all of the unveiling from previous rounds. As a result, the numbers add up to more than the total in the left most column. For example, depository institutions hold a large amount of pass-through debt, and so the 0.46 in this box reflects both household debt held directly and household debt held through holdings of agency GSE debt.[27]

households, then the share of mutual funds that is household debt might be higher for the next 9%.

[27]Figure 5 does not include every linkage, and it excludes the holdings of the rest of the world and the U.S. government. For these latter two categories, there should be arrows going out of each box that go to the rest of the world and to the government. All linkages are taken into account in the methodology, but for the sake of making the figure

Figure 5: Overview of Unveiling of Household Debt for 2005

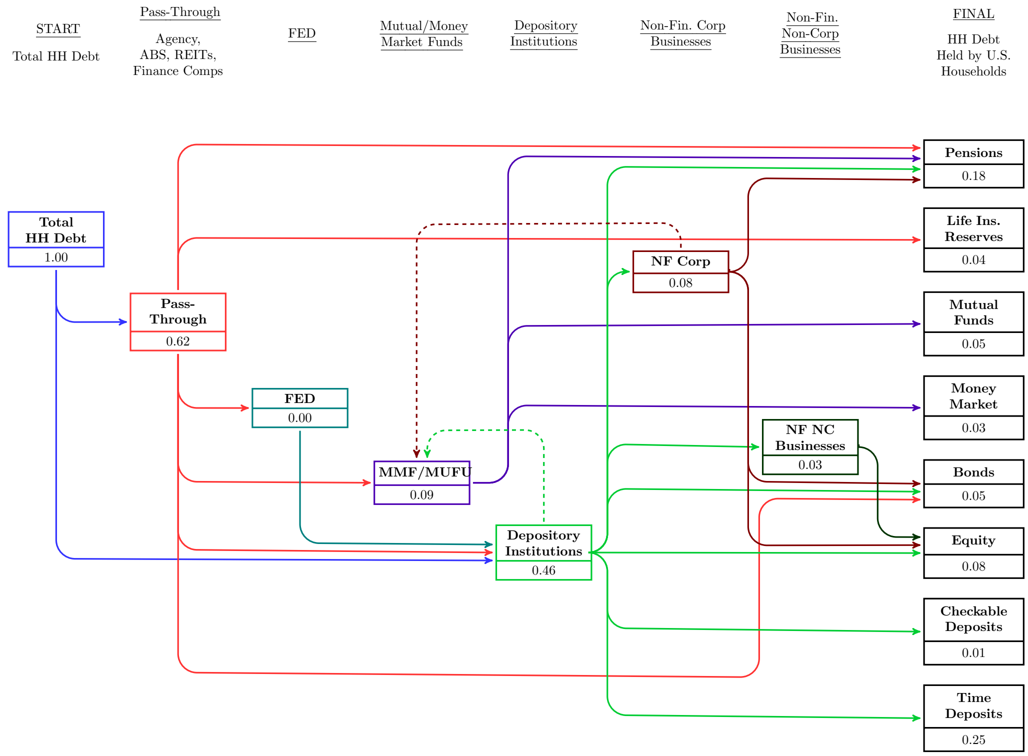

This is a visual representation of the unveiling of household debt for the year 2005. The left most column is total household debt for U.S. households, and the right most column reflects the household debt held as a financial asset by the U.S. household sector. The numbers in each box are total household debt for that box scaled by national income.

The first round places household debt into the two major initial intermediaries: pass-throughs (which include the Government Sponsored Enterprises and issuers of asset-backed securities) and private depository institutions. The second round then unveils the pass-throughs. Pass-through debt is held by a large number of intermediaries, and also by the household sector itself through its holdings of pensions, life insurance, and bonds.

The third round unveils the Fed, which is held primarily by the depository institutions. The Fed held no household debt in 2005, because at this point it had yet to buy any Agency GSE debt. But after 2009, the Fed held substantial household debt through its holdings of Agency GSE debt. The fourth round unveils mutual funds and money market funds.

The fifth round is quantitatively the most important. Depository institutions hold a huge amount of household debt through both their direct holdings (portfolio mortgages for example) and their purchases of pass-through debt. Post 2009, they also hold a large amount of household debt through

readable not all of these linkages are shown in Figure 5.

their reserves at the Fed. As the graph shows, the claim on the depository institutions is quantitatively meaningful for many players in the economy, including non-financial corporations, non-financial non-corporate businesses, and the household sector. This is primarily due to holdings of deposits, but it is also because of the equity claim on the depository institutions.

The final two rounds unveil the non-financial business sector. This is also a quantitatively important step, because the savings by the business sector have increased substantially after 1995. For example, non-financial businesses increased their holdings of deposits and money market funds by 10 percentage points of national income (see Appendix Figure A4). These are indirect claims on household debt through the holdings of liabilities in commercial banks and money market funds. As a result, the household sector held a substantial amount of household debt through their claims on the non-financial business sector.

How do we measure the arrows that come out of each box? We are able to do so because of the excellent data in the Financial Accounts that details the claims on any given institution by other institutions. For example, let us consider the Government-Sponsored Enterprises (GSEs), which are important immediate holders of household debt. Tables L.125 and L.126 in the Financial Accounts document the amount of home mortgages and consumer credit held as an asset by the GSEs. Table L.211 of the Financial Accounts documents the total debt issued by the GSEs, and the groups to which they owe these liabilities.

Take the orange arrow from the pass-throughs to mutual funds. Table L.211 lists the share of total agency GSE liabilities held by mutual funds. The main assumption used to create the arrow is a proportionality assumption: the share of the total liabilities of the Agency GSEs held by mutual funds is assumed to be the same as the share of the total amount of household debt held by mutual funds through their holdings of Agency GSE liabilities. Then, Table L.224 of the Financial Accounts lists the groups that own shares of mutual funds, which can then be further unveiled in the next round. This is an example of how the unveiling process works. It can be done for all intermediaries to determine ultimately who holds the household debt as a financial asset.

In addition to the vector F_t (the right most column in Figure 5), this process also produces the amount of household debt held by the government, the rest of the world, and a residual category.[28] Further details of the unveiling process are in Appendix Section B.2.

Once the vector F_t has been calculated, the household debt held by group i in year t can be calculated through the matrix Ω_t, which has already been described in Section 2.4.

[28] The residual component comprises a few factors. First, the holdings of some liabilities in the Financial Accounts are not detailed. These are gathered into one table in the Financial Accounts (L.231). Second, the unveiling process in general only moves in one direction. If there are claims against an entity by an entity that has already been unveiled, then the claim is put into a residual category. For example, loans by banks to non-financial businesses mean that some of the household debt held by non-financial businesses should be claimed by banks. But banks have already been unveiled when the round of unveiling for non-financial businesses is reached. These claims then go into the residual category.

The vector F_t over time

The information that is gained through the unveiling process can be seen in Figure 6. The left panel focuses on household debt held as a financial asset by the U.S. household sector, and the right panel focuses on on government debt held as a financial asset by the U.S. household sector. Each bar represents the change in the asset held through a given class J over time. This is the change in the vector F_t over time. The last year for household debt is 2007 given that this year was the peak of the household debt to national income ratio, whereas the last year is 2016 for government debt which was the peak for the government debt to national income ratio. The vector F_t is scaled by national income in year t before the first difference is taken.

Figure 6: Through What Asset Classes have U.S. Households Increased Debt Holdings?

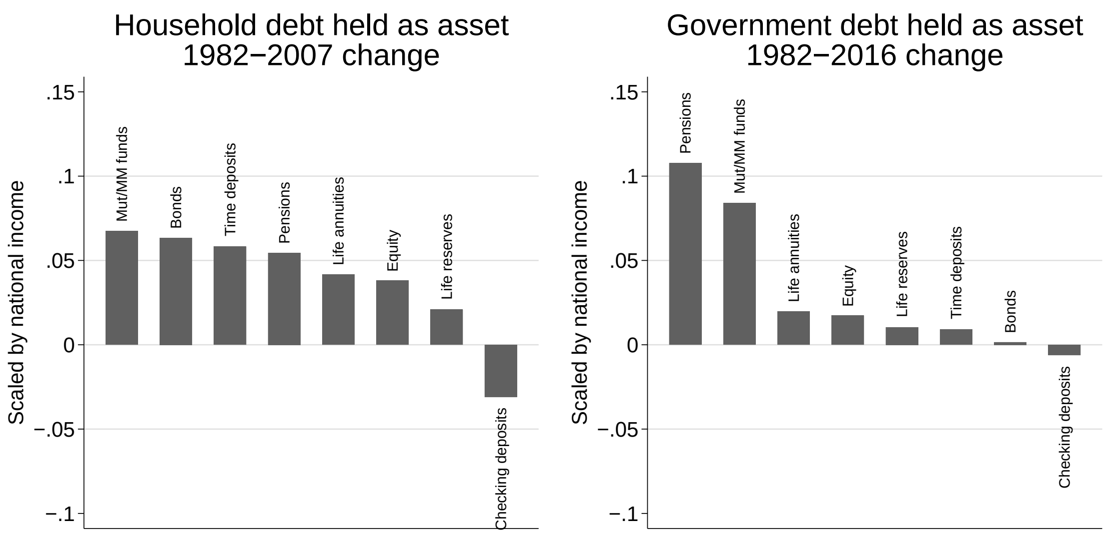

This figure shows the asset classes through which U.S. households have increased their holdings of household debt (left panel) and government debt (right panel) as financial assets.

The left panel of Figure 6 shows that that the largest rise in the holdings of U.S. household debt by the household sector have been through mutual and money market funds. Households have increased their holdings of household debt by 7 percentage points of national income through these funds from 1982 to 2007. The second largest increase has been through direct bond holdings, which has been driven by direct holdings of Agency GSE bonds and privately securitized bonds. The U.S. household sector has actually decreased its holdings of household debt through checking deposits.

The U.S. household sector has also substantially increased its holdings of household debt through its holdings of equity in corporate and non-corporate businesses. This may at first be surprising. However, it is important to remember that U.S. businesses have boosted their holdings of time deposits and money market funds considerably since 1995 (see Appendix Figure A4), and these deposits ultimately finance household debt through the financial system. This rise in business deposits means that the U.S. household sector has boosted holdings of household debt through their equity holdings of businesses. For U.S. government debt, the U.S. household sector has boosted holdings primarily through pensions and mutual and money market funds.

Figure 6 makes it clear that U.S. households have increased their holdings of household and government debt through several financial asset classes that are not limited to fixed income assets. Pensions, equity, and mutual funds increasingly contain within them large claims on government and household debt. This helps to explain why households even in the top 1% of the wealth distribution, who tend to hold a higher share of business equity and mutual funds, have been financing the rise in government and household debt. This is the focus of the next subsection.

4.2 Saving in debt across the wealth distribution

The unveiling procedure above produces the amount of household debt and government debt held as a financial asset by group i in year t, which we call HHD_{it} and $GOVD_{it}$, respectively. This allows for the construction of saving in household debt and government debt which is:

$$\Theta_{it}^{HHD} = \Delta HHD_{it} - \pi_t^{HHD} HHD_{i,t-1}$$
$$\Theta_{it}^{GOVD} = \Delta GOVD_{it} - \pi_t^{GOVD} GOVD_{i,t-1}$$

In the absence of debt write-downs, π_t^{HHD} and π_t^{GOVD} would be zero. Write-downs on government debt have been trivial during our sample period, whereas there have been some write-downs on household debt. Measuring π_t^{HHD} would be extremely challenging, as it would require an estimate of which asset classes holding household debt ultimately experienced losses. As a result, we make the simplifying assumption that $\pi_t^{HHD} = 0$.[29]

Table 6 shows the results of this exercise. More specifically, it shows total financial asset accumulation and the amount that is financial asset accumulation in household debt and government debt across the wealth distribution. Annual accumulation is scaled by contemporaneous national income, and then averages for the periods are shown.

The top 1% have accumulated a substantial amount of household and government debt after 1982. From 1963 to 1982, the top 1% invested 0.6 percentage points of national income annually in household and government debt. This number increased by a factor of four after 1982. Accumulation of household debt by the top 1% was especially large from 1998 to 2007, when rich households accumulated 1.8 percentage points of national income each year. The top 1% accumu-

[29]It is important to note that this assumption is conservative. Debt write-downs would translate into a negative π_t^{HHD} which would increase measures of Θ_{it}^{HHD}. Assuming zero debt write-downs on household debt makes it appear as if households are accumulating fewer claims on household debt than they actually are.

lated substantial amounts of government debt from 2008 to 2016, on the order of 1.8 percentage points of national income every year.

Overall, from 1983 to 2016, the top 1% accumulated financial assets at an average annual rate of 4.4 percentage points of national income, compared to 2.5 percentage points from 1963 to 1982. Of this 4.4 percentage points, 2.4 percentage points were invested in assets that are claims on household and government debt. More than half of the financial asset accumulation of the top 1% has been directed toward household and government debt since 1983.

Table 6: How Much Financial Asset Accumulation is Claim on Household and Government Debt?

Top 1%				
Time Period	Θ^{FA}	Θ^{HHD}	Θ^{GOVD}	$\Theta^{HHD} + \Theta^{GOVD}$
63-82	0.025	0.004	0.002	0.006
83-97	0.033	0.011	0.011	0.022
98-07	0.056	0.018	0.008	0.026
08-16	0.048	0.005	0.018	0.023
Next 9%				
Time Period	Θ^{FA}	Θ^{HHD}	Θ^{GOVD}	$\Theta^{HHD} + \Theta^{GOVD}$
63-82	0.062	0.014	0.009	0.023
83-97	0.041	0.015	0.014	0.030
98-07	0.040	0.016	0.003	0.018
08-16	0.045	-0.002	0.017	0.015
Bottom 90%				
Time Period	Θ^{FA}	Θ^{HHD}	Θ^{GOVD}	$\Theta^{HHD} + \Theta^{GOVD}$
63-82	0.061	0.015	0.011	0.026
83-97	0.056	0.014	0.013	0.027
98-07	0.038	0.013	0.002	0.014
08-16	0.022	-0.006	0.011	0.005

This table shows how much of the annual accumulation of financial assets (Θ_{it}^{FA}) is asset accumulation of claims on U.S. household debt (Θ_{it}^{HHD}) and U.S. government debt (Θ_{it}^{GOVD}). Annual accumulation is scaled by national income, and then the averages for the period is shown.

4.3 Net household debt positions

Households within the United States can borrow more, but they can also save more in financial assets that ultimately fund borrowing by other households. This suggests that an important variable in assessing the financial position of U.S. households is net household debt accumulation, which we define as $\Theta_{it}^{NHHD} = \Theta_{it}^{HHD} + D_{it}$, where Θ_{it}^{HHD} is the accumulation of household debt held as a financial asset and D_{it} is the accumulation of household debt owed as a financial liability (where negative values of D_{it} represent more borrowing).

Table 7 shows net household debt accumulation across the wealth distribution. The accumulation of household debt as a financial asset by the top 1% was offset only slightly by more borrowing from 1983 to 2007. This implies that the top 1% have seen a large increase in their net household debt position. From 1998 to 2007, the top 1% increased their claims on the borrowing of other U.S. households by 1.3 percentage points of national income per year.

Table 7: Net Household Debt Positions

	Bottom 90%			Next 9%			Top1%		
Period	D	Θ^{HHD}	Θ^{NHHD}	D	Θ^{HHD}	Θ^{NHHD}	D	Θ^{HHD}	Θ^{NHHD}
63-82	-0.040	0.015	-0.026	-0.006	0.014	0.008	-0.001	0.004	0.003
83-97	-0.045	0.014	-0.031	-0.008	0.015	0.007	-0.003	0.011	0.008
98-07	-0.073	0.013	-0.060	-0.012	0.016	0.004	-0.004	0.018	0.013
08-16	-0.016	-0.006	-0.022	-0.003	-0.002	-0.005	0.001	0.005	0.006

This table shows the net household debt accumulation which is defined to be $\Theta_{it}^{NHHD} = D_{it} + \Theta_{it}^{HHD}$, where a negative D is more borrowing. Annual accumulation is scaled by national income, and then the averages for the period are calculated.

The next 9% have borrowed more but they have also accumulated more in financial assets that are claims on household debt. As a result, the net household debt accumulation of the next 9% has been relatively steady as a percentage of national income. The bottom 90% have experienced a large decline in net household debt accumulation. They have been accumulating fewer financial assets that are claims on household debt, and they have substantially increased borrowing.

Net household debt positions make it clear that beginning in the 1980s, the top 1% began to increasingly finance borrowing by the bottom 90%. Figure 7 shows the stock of household debt held as a financial asset minus the stock of household debt owed as a liability across the wealth distribution over time. The net household debt position is scaled by national income, and the 1982 level is subtracted. At the peak in 2007, the net household debt position of the top 1% was 12 percentage points of national income, and the net household debt position of the bottom 90% was negative 39 percentage points of national income. Over 30% of the rise in net household debt owed by the bottom 90% was financed by the top 1%.

Figure 7: Net Household Debt across Wealth Distribution Relative to 1982

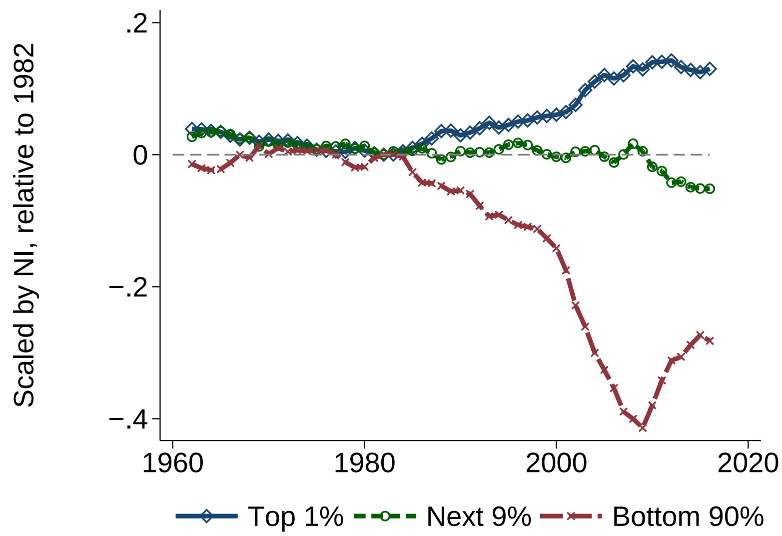

This figure shows net household debt outstanding by the U.S. household sector across the wealth distribution. Net household debt is defined household debt held as a financial asset minus household debt owed as a liability. All series are scaled by national income, and the 1982 level is subtracted.

4.4 Financing the rise in government and household debt

The unveiling exercise also allows us to quantify exactly how much of the additional accumulation of household and government debt shown in Figure A3 in the appendix has been associated with financial asset accumulation of the rich. Figure 8 presents the evidence.

To construct Figure 8, we start in 1982 and calculate the average annual rise in debt owed by the household and government sectors scaled by contemporaneous national income. We then subtract the annual average rise in debt in the pre-period, from 1963 to 1982. The orange bar on the left shows that on average the combination of household and government debt rose annually by almost 3 percentage points of national income more from 1982 to 2016 compared to 1963 to 1982.

Figure 8: Sources of Financing for Rise in Government and Household Debt

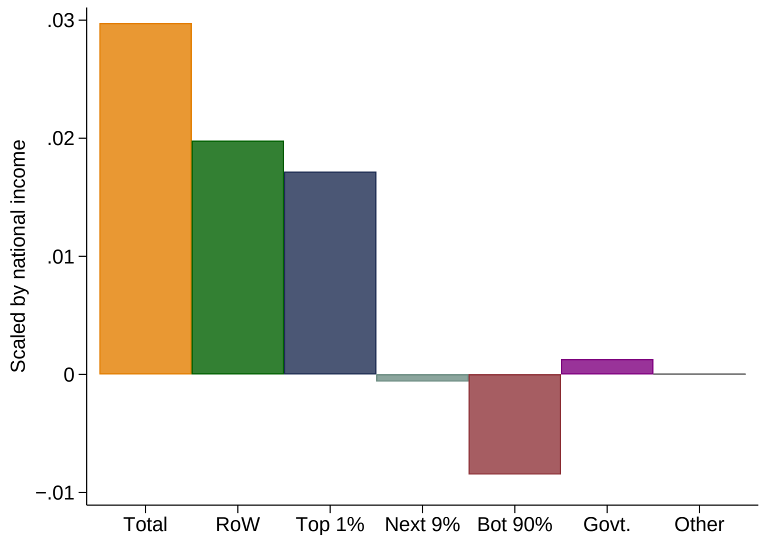

This figure shows the average annual rise in government and household debt from 1982 to 2016, where the annual rise is scaled by national income each year. The annual average annual rise from 1963 to 1982 is subtracted to isolate the difference between the two time period. It then shows how much of this rise in the latter period relative to the former period has been financed by different groups. The five bars on the right add up to the bar on the left.

Of the 3 percentage point average annual additional accumulation of household and government debt, 1.7 percentage points have been due to additional accumulation by the top 1% of the wealth distribution, and 2 percentage points have been due to additional accumulation from the rest of the world. The next 9% accumulated almost the exact same post 1982 and pre-1982, while the bottom 90% reduced their accumulation of household and government debt by almost a full percentage point of national income annually. The additional borrowing by households and the government has been financed almost equally by the top 1% and the rest of the world.

The pattern in Figure 8 is closely related to the literature exploring the rise in the demand for U.S. dollar-denominated safe assets. This literature has typically focused on demand coming from other countries. However, U.S. government and household debt are often the ultimate debt instrument backing such safe assets, and the unveiling process allows us to see that the top 1% within the United States have increased their holdings of these debt claims by almost as much as the rest of the world.

Figure 9: Safe Asset Demand: Who Holds U.S. Government and Household Debt?

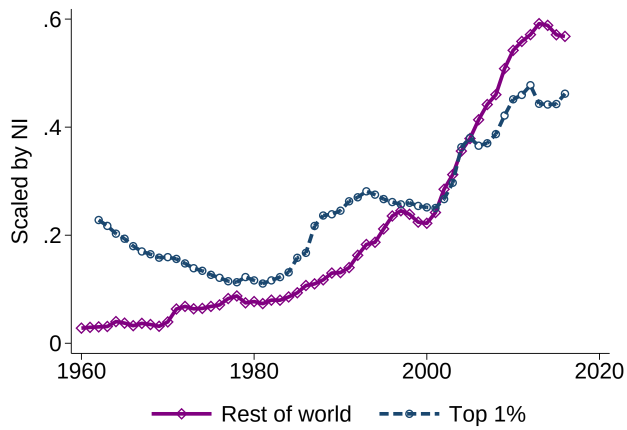

This figure shows the holdings of U.S. government and household debt by the rest of the world and by the top 1% of the U.S. wealth distribution, scaled by national income. These are holdings of assets after the financial sector is unveiled.

Figure 9 illustrates this point by showing the total amount of U.S. household and government debt held by the rest of the world and by the top 1% of the wealth distribution from 1963 to 2016. Both series are scaled by national income. The rest of the world has indeed boosted their holdings of U.S. government and household debt substantially. However, the rise in the holdings of the top 1% is on the same order of magnitude. This suggests that more attention should be focused on the rise in the demand for safe assets coming from within the United States by rich Americans.

5 Top Income Shares and the Saving Glut of the Rich

This section turns to potential explanations for the rise in savings by rich Americans. This requires us to go beyond aggregate data, as there have been several national trends in the United States since the early 1980s that may have affected the evolution of savings. One of these trends is the well-documented rise in income inequality.[30] The analysis in this section is based on a novel state-year level panel data set that allows for an investigation of whether there has been a direct connection between the rise in top income shares and the rise in savings by the rich. There was substantial state-level variation in the rise in top income shares from the 1980s to 2008, as shown in Figure 10. States like Florida, New York, and Nevada witnessed a larger increase in the top 1% share of income relative to states such as Michigan, Arizona, or California.

[30] See, e.g., Katz and Murphy (1992), Piketty and Saez (2003), Autor et al. (2008), Atkinson et al. (2011), Piketty et al. (2018), CBO (2019), and Smith et al. (2019). There is substantial evidence in the literature that the rise in top income shares reflected shifts in technology and globalization that began in the 1980s. This view is supported by the fact that the rise in the share of income of the top 1% is broad-based across many industries (e.g., Kaplan and Rauh (2013), Bakija et al. (2012)), and that much of these earnings are derived from human capital (Smith et al. (2019)).

Figure 10: Change in Top 1% Share of Income Across States

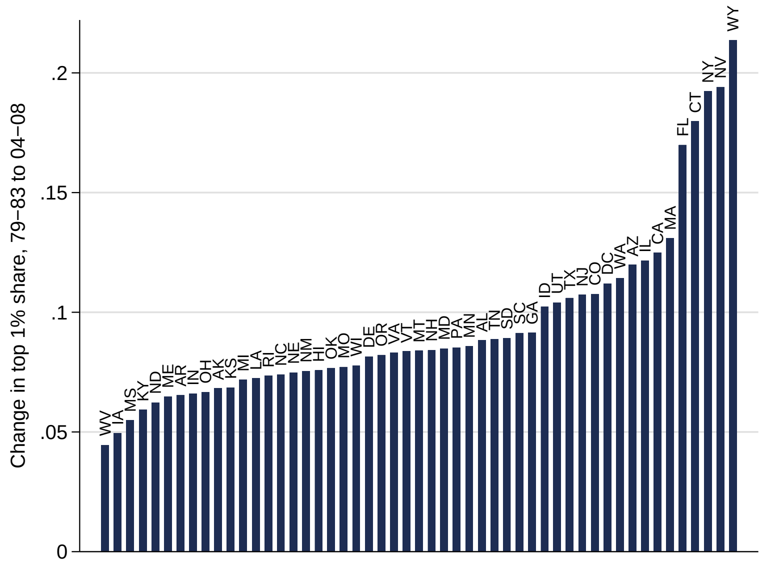

Data are from the World Inequality Database.

The state-level analysis tests whether states with a higher increase in top income shares experienced a greater rise in savings by those at the top of the income distribution.[31] State-level analysis brings us closer to an ideal experiment in which, all else equal, some economies experience a larger increase in top income shares than others. The analysis then tracks whether there is more asset accumulation in the economies with the larger increase in top income shares.

5.1 Measuring shares of wealth at the state level

The shares of asset class j held by group i at time t in state s, $\omega^j_{i,s,t}$, can be estimated much the same way as $\omega^j_{i,t}$ was estimated in section 2.4 given the availability of state identifiers in the DINA microfiles. However, there are a few limitations for the state level analysis.

First, state-level information is available from 1979 to 2008, which means the state-level analysis must end in 2008. Second, other than the year 1982, the state level identifiers are suppressed for any tax return with an adjusted gross income (AGI) above $200 thousand in nominal terms. In order to overcome this issue, the analysis uses state-level tax tables provided by the Statistics of Income Division (SOI) of the Internal Revenue Service. These tables contain the total number of returns, interest income, dividend income, capital gains, and taxable pension income at the state

[31]Given that financial markets are well-integrated across the United States, the state-level variation in the rise in top income shares is only related to asset accumulation by those at the top of the income distribution. The associated household borrowing and dissaving by the bottom of the income distribution may happen anywhere across the United States. As a result, state-level analysis is not able to detect whether the rise in top income shares is responsible for borrowing and dissaving by those outside the top of the income distribution.

level. Furthermore, these tables break down this information separately for filers with AGI above \$200 thousand. These data enable us to create a group in each state-year cell that contains all tax units that earn more than \$200 thousand in AGI. This means, however, that we must use tax units as the unit of observation instead of equal-split individuals as before.

Using information from the SOI, Figure A5 in the appendix plots the share of tax filings with AGI above \$200 thousand in each state by year. The shaded area covers the full variation across states in the share of units with AGI above \$200 thousand each year. Given the \$200 thousand limit, we cannot form a group for the top 1% in each state-year cell. However, the fraction of filers below \$200 thousand is always below 6% for every state-year observation. As a result, the top 6% is chosen as the main "top income" category in the state-level analysis.[32] The SOI data for the group of individuals in each state-year observation with AGI above \$200 thousand allows us to apply the capitalization methodology to estimate the wealth of the top 6% of the income distribution in each state.[33] More details are available in Section C of the appendix.

5.2 Estimation strategy

Let $\theta_{is} \equiv \frac{\Theta_{is}}{Y_s}$ denote saving of group i in state s relative to the state's share of national income. Suppose top-income inequality, τ_s, rises from 1979 to 2008. The goal of this section is to estimate the parameter,

$$\beta_i = \frac{\partial \theta_{is}}{\partial \tau_s}$$

There are two ways of estimating β_i. First, state-level panel data can be used to regress θ_{ist} on τ_{st} with appropriate controls. The wealth-based estimate of θ_{ist} is measured with considerable measurement error at the annual frequency; as a result, five-year state-level averages are taken. The estimation equation is:

$$\theta_{ist} = \alpha_s + \alpha_t + \beta_i * \tau_{st} + \Gamma * X_{st} + \varepsilon_{st} \tag{11}$$

where the log of income per capita is included in X_{st}.

The second method of estimating β_i uses the change in wealth at the state level for group i from

[32]The top 6% group for each state-year observation is the top 6% of tax units within the state, not tax units in the top 6% of the national distribution. Figure A6 in the appendix shows a strong relationship between the rise in the top 1% share of income in a state from 1982 to 2007 and the rise in top 6% share of income in the same year. The top 1% share of income at the state level is available from the World Inequality Database.

[33]Notice a key difference in the state-level analysis is that the wealth shares across the income distribution are used instead of the wealth shares across the wealth distribution. This is because of the issues related to missing state identifiers in the DINA and the need to use the SOI. The SOI does not include information on tax filings sorted by wealth.

1979 to 2008. Dividing equation 7 by state income Y_{st} leads to $\theta_{ist} = w_{ist} - \frac{1+\overline{\pi}_{st}}{1+g_{st}} w_{ist-1}$, where $\overline{\pi}_{st}$ is the average asset price inflation for wealth portfolio in state s, defined as $\overline{\pi}_{st} = \frac{\sum_j \pi^j_{st} W^j_{st-1}}{\sum_j W^j_{st-1}}$.[34] In this equation, $1 + g_{st} = \frac{Y_{st}}{Y_{st-1}}$ is the state-level income growth rate.

This equation links saving in each period to wealth accumulation. Summing this equation from $t = 0$ to $t = T$ allows us to relate wealth accumulated by cohort i in state s, Δw_{is} as a linear combination of per-period saving and initial wealth:

$$\Delta w_{is} = \sum_{t=1}^{T} \mathcal{R}_{st} \theta_{ist} + (\mathcal{R}_{s0} - 1)\, w_{is0} \tag{12}$$

where $\Delta w_{is} = \frac{W_{isT}}{Y_{sT}} - \frac{W_{is0}}{Y_{s0}}$, w_{is0} is the initial wealth to state income ratio, and $\mathcal{R}_t$ is the cumulative return on wealth relative to the size of the state economy from t until T, defined as $\mathcal{R}_{st} = \frac{1+\overline{\pi}_{st+1}}{1+g_{st+1}} \mathcal{R}_{st+1}$ with $R_T = 1$. Taking the derivative of equation 12 with respect to change in inequality $\Delta\tau = \tau_{sT} - \tau_{s0}$, we can see that $\frac{\partial \Delta w_{is}}{\partial \Delta \tau_s} = \sum_{t=1}^{T} \mathcal{R}_{st} * \beta_i$. There is therefore a proportional relationship between β_i in equation 11 and $\overline{\beta}_i$ in specification:

$$\Delta w_{is} = \alpha + \overline{\beta}_i * \Delta\tau_s + \Gamma * X_s + \varepsilon_s \tag{13}$$

where X includes initial wealth as a share of national income (w_{is0}), initial top-income share, initial log income per capita, and the change in log income per capita over the time period. We take five-year averages of the data in the beginning (1979-1983) and end (2004-2008) of the state-level sample period before taking the long difference Δ in equation 13. Both equations 11 and 13 allow for estimation of the effect of a rise in top income shares on savings by a specific income group. Results from both estimation equations are shown below. All regressions are weighted by total number of tax units in a state so that the coefficients have nationally representative economic magnitude.

5.3 State-level results

Figure 11 presents the scatter-plots of θ_{ist} and τ_{st}, following equation 11. The two variables are adjusted using state and year fixed effects, and the residuals are plotted. The left panel shows the bivariate relationship for the top 6% and the right panel shows the same for the bottom 94%. As the figure shows, there is a positive relationship between top income shares and savings by the top 6% of the income distribution, but no relationship between top income shares and savings by the bottom 94%.

[34] Asset price inflation for asset j in equation 7 is the same for all states except for housing and debt liabilities under the assumptions of an integrated financial market and that households across states hold similar portfolios on average within asset class j. For housing assets, state-specific house price indices from CoreLogic are used. For debt liabilities, write-downs are applied to the top 10% and bottom 90% of the distribution in the same way as in Section 2.4.

Figure 11: Change in Savings and the Rise in Top Income Shares

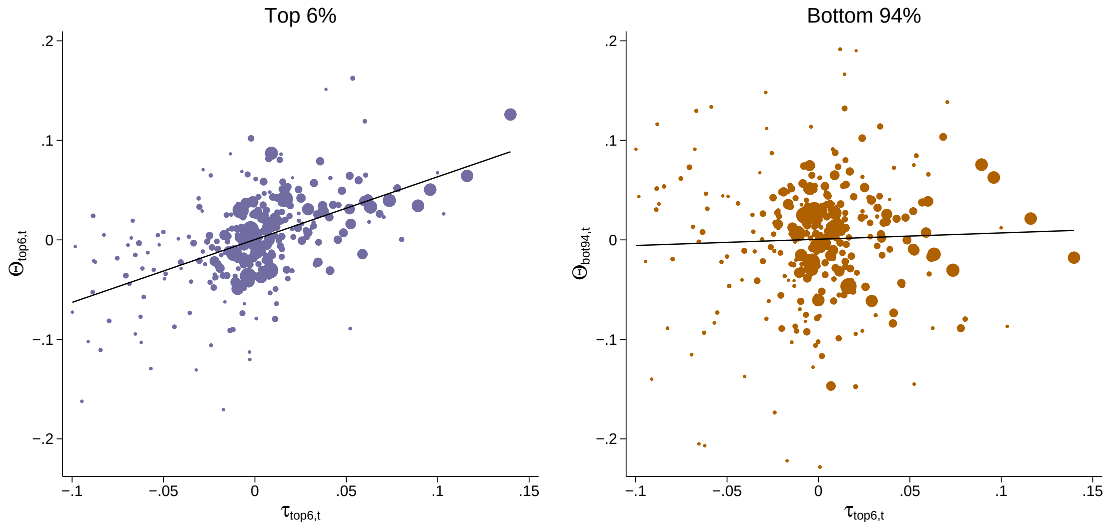

θ_{ist} reflects savings by group i in state s in year t scaled by state national income. i is the top 6% of the income distribution in the left panel and the bottom 94% in the right panel. τ_{st} is the share of income by the top 6% of the income distribution in state s in year t. All variables are adjusted using state and year fixed effects before plotting.

Figure 12 presents the scatter-plots of Δw_{is} and $\Delta \tau_s$, following equation 13. The left panel shows the bivariate relationship between the change in wealth to state-income for the top 6% of the income distribution, and the right panel shows the same for the bottom 94%. The contrast between the two figures illustrates that the entire increase in the accumulation of wealth as top income share rises is driven by top earners.

Figure 12: Change in Wealth to Income Ratios and Rise in Top Income Shares

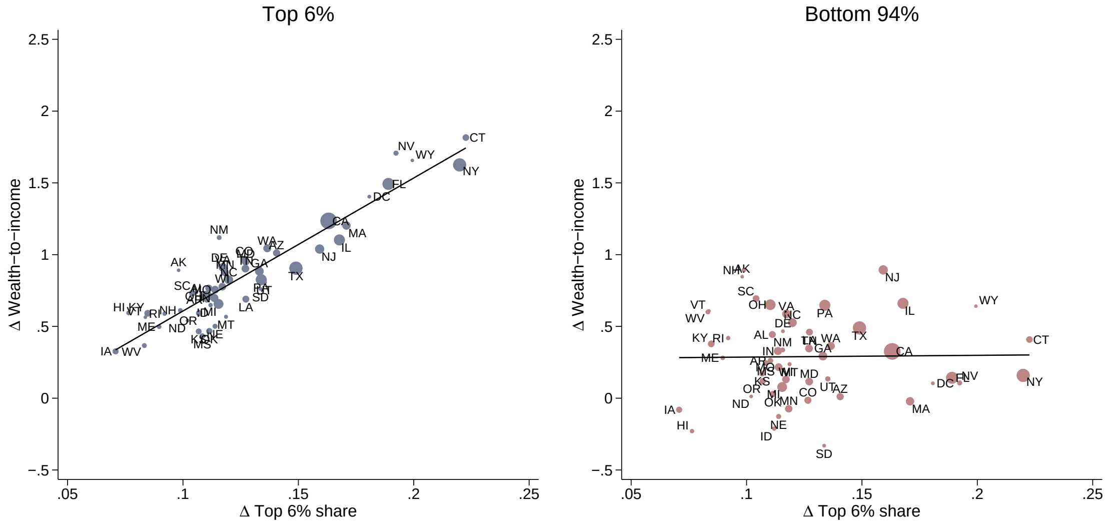

Δ Wealth-to-income is the change in net wealth by group i in state s scaled by state national income. i is the top 6% of the income distribution in the left panel and the bottom 94 % in the right panel. The Δ reflects the change in the variable in question using the average from 2004 to 2008 minus the average from 1979 to 1983.

Column 1 of Table 8 presents the estimates of equation 11 for all income groups collectively. The top 6% income share has a statistically significant positive effect on savings in a given state. Column 2 adds a control for per-capita income which has only a minor effect on the coefficient estimate. Columns 3 and 4 present results separately for savings of the top 6% of the income distribution and bottom 94%. Top income shares are associated with higher savings by the top 6% of the income distribution, but there is no effect on savings by the bottom 94%.

Table 8: Rise in Top Income Shares and Saving (panel regression)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Θ	Θ	Θ	Θ	Top 6% share	Θ	Θ	Θ
Top 6% share	0.698***	0.649***	0.645***	0.00365		0.716***	0.651***	0.0655
	(0.080)	(0.142)	(0.072)	(0.102)		(0.183)	(0.097)	(0.122)
Top 6% labour share					1.430***			
					(0.157)			
Log per capita income		0.0214	-0.00693	0.0272	0.0201	-0.00224	-0.00984	0.00648
		(0.047)	(0.029)	(0.028)	(0.014)	(0.057)	(0.034)	(0.033)
Group	All	All	Top 6%	Bot. 94%		All	Top 6%	Bot. 94%
Method	OLS	OLS	OLS	OLS	OLS	2SLS	2SLS	2SLS
R^2	0.16	0.16	0.59	0.43	0.96	0.17	0.59	0.42
Observations	306	306	306	306	306	306	306	306

The left hand side θ_{ist} reflects savings by group i in state s in year t scaled by state national income. The top 6% share is the share of income by the top 6% of the income distribution in state s in year t. Robust standard errors, clustered by state, are in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

One concern with the results in Table 8 is that they are partially mechanical given that capital income is both an input into the measurement of savings on the left hand side through capitalization, and the top income share on the right hand side. However, columns 5 through 8 show that this is not the case by using the top labor income share as an instrument in a two stage least square specification.

Column 5 shows that the top 6% share of labor income is strongly correlated with the top 6% income share that is used as the main RHS variable. Columns 6 through 8 repeat the exercise of columns 2 through 4, but use the predicted top 6% income share from column 5 in a 2SLS specification. This ensures that the estimated coefficients are not driven by the capital income component. The results indicate that the coefficients of interest are not materially affected.

Table 9: Rise in Top Income Shares and Wealth-to-Income Ratios

	(1) Δ wealth-to-income	(2) Δ wealth-to-income	(3) Δ wealth-to-income	(4) Δ wealth-to-income	(5) Δ Top 6% share	(6) Δ wealth-to-income	(7) Δ wealth-to-income	(8) Δ wealth-to-income
Δ Top 6% share	8.467***	8.706***	8.950***	-0.034		9.263***	8.019***	1.764
	(1.203)	(1.937)	(0.843)	(2.476)		(2.378)	(0.902)	(3.005)
Δ Top 6% labour share					1.432***			
					(0.165)			
Top 6% share 79-83		-2.366	4.551**	-4.834	0.221	-2.974	5.378***	-6.919
		(3.083)	(1.452)	(3.735)	(0.135)	(3.413)	(1.503)	(4.078)
Δ Log per capita income		1.897*	0.585	1.373	0.166***	1.806*	0.743*	1.076
		(0.721)	(0.307)	(0.832)	(0.023)	(0.714)	(0.308)	(0.874)
Log per capita income 79-83		1.527*	0.652*	1.116	0.072*	1.468*	0.761**	0.917
		(0.585)	(0.259)	(0.665)	(0.034)	(0.585)	(0.273)	(0.689)
Wealth-to-income 79-83		-0.427***	-0.756***	-0.419***	0.002	-0.432***	-0.729***	-0.436***
		(0.076)	(0.113)	(0.091)	(0.006)	(0.073)	(0.117)	(0.090)
Group	All	All	Top 6%	Bot. 94%		All	Top 6%	Bot. 94%
Method	OLS	OLS	OLS	OLS	OLS	2SLS	2SLS	2SLS
R^2	0.57	0.76	0.94	0.30	0.92	0.76	0.94	0.29
Observations	51	51	51	51	51	51	51	51

Δ wealth-to-income is change in overall wealth by group i in state s scaled by state national income. Robust standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Δ reflects the change in the variable in question using the average from 2004 to 2008 minus the average from 1979 to 1983.

Table 9 presents estimates of the second technique shown in equation 13. This specification relies on a single cross-section of long-differences. The results are similar: a rise in top income shares is associated with a substantial increase in the wealth to income ratio in a state over time. In terms of magnitude, a one standard deviation increase in the rise in the top 6% income share (0.035) leads to a 0.30 increase in the state-level wealth to income ratio, which is a 0.76 standard deviation of the left hand side variable.

Column 2 adds additional controls, which add statistical power but do not change the coefficient of interest materially. These controls move the specification closer to the ideal thought experiment of keeping income growth and initial conditions constant while changing the rise in top income shares.[35] The inclusion of the last control, the initial wealth to income ratio of income group i, ensures that the estimates are not driven by any mechanical "valuation effects."[36] Columns 3 through

[35]Table A6 in the appendix shows how the rise in top income share is correlated with these four controls.

[36]The specific concern is that the drop in long-term interest rates gives all holders of wealth a capital gain. Mechanically, states where the rich hold more initial wealth as a share of income will see a larger increase in wealth. But this change is entirely driven by the "valuation effect" of lower interest rates, and has nothing to do with savings. The

8 present specifications that are analogous to those reported in Table 8; the results are similar.

An advantage of the state-level analysis is that it allows us to control for other secular trends that may be responsible for both the rise in top income shares and the rise in wealth of the rich. Table A7 in the appendix reports results analogous to column 3 of Table 9 with additional controls for demographics; the share of employment in the financial, manufacturing, and construction sector; and measures of financial deregulation. The specifications reported in the appendix also add controls for measures of migration of high income individuals into a state to ensure that the effect of inequality on savings is not driven by in-migration of rich individuals into the state. Inclusion of these controls does not affect the estimate on top income shares.

How are the results from Tables 8 and 9 related? Given that $\frac{\partial \Delta w_{is}}{\partial \Delta \tau_s} = \sum_{t=1}^{T} \mathcal{R}_{st} * \beta_i$, we can define $\widehat{\frac{\partial \Delta w_{is}}{\partial \Delta \tau_s}} = \sum_{t=1}^{T} \overline{\mathcal{R}_t} * \widehat{\beta_i}$, where $\widehat{\beta_i}$ comes specifically from column 2 of Table 8 and thus equals 0.649 and $\overline{\mathcal{R}_t} = \frac{\sum_s ||s|| \mathcal{R}_{st}}{\sum_s ||s||}$. We define $\mathcal{R}_{st}$ by taking $t = 0$ as the average of 1979-1983 and T as the average of 2004-2008, with each year in between remaining independent. We can therefore compare $\widehat{\frac{\partial \Delta w_{is}}{\partial \Delta \tau_s}}$ to the coefficient in column 2 of Table 9, which is equal to 8.7. This calculation yields $\widehat{\frac{\partial \Delta w_{is}}{\partial \Delta \tau_s}} = 11.5$. Both techniques reveal similar quantitative effects of top income shares on the savings by the top 6%.

Tables A4 and A5 in the appendix break down the savings and wealth accumulation by the top 6% into different components. Both tables show that the rise in top income shares is associated with a large rise in financial asset accumulation as opposed to real estate accumulation or a reduction in debt. Furthermore, about 30% to 40% of the rise in savings associated with the rise in top income shares is driven by a rise in the accumulation of assets that are directly tied to household and government debt.

Overall, the state-level results help tie the rise in top income shares to the rise in savings by those at the top of the income distribution. The rise in the saving glut of the rich has been closely connected to the well-documented rise in income inequality in the United States since the 1980s.

6 Conclusion and Future Directions

There has been a large rise in savings by rich Americans since the 1980s. This rise in savings has been associated with a significant dissaving by other Americans and by the government. Furthermore, Americans in the top 1% of the wealth distribution have been accumulating substantial financial claims on household and government debt. State-level analysis points to the rise in top income shares as an important factor in understanding these patterns.

What are the broader implications of the saving glut of the rich? This question is at the heart

control for initial wealth to income removes this mechanical effect.

of the indebted demand model in a companion study (Mian et al. (2021)). The indebted demand model introduces non-homotheticity into the consumption-savings decision of households, and it explores the effects of a rise in income inequality and financial liberalization on interest rates and debt levels. As is shown there, a large rise in inequality generates a saving glut of the rich, which can push an economy into a debt trap characterized by low interest rates, high debt levels, and output below potential.

Another key question for future research is: why has the saving glut of the rich been associated with dissaving by other households and the government instead of higher business investment? There is a large body of research exploring the reasons behind weak investment in advanced economies. A related issue is whether the financial sector is directing the savings by the rich to other households because of government policy. For example, most of the rise in household debt since the 1980s has been associated with the rise in Agency GSE-backed mortgages, which carry either an implicit or explicit government guarantee.

Finally recent research on the economic consequences of the Covid 19 health crisis suggest that there has been a particularly large drop in spending by high income Americans, who also have experienced a smaller decline in labor income.[37] It is still too early to tell, but savings by the richest Americans may increase substantially as a result of the health crisis. The saving glut of the rich has already been quite large from the early 1980s to the end of our sample period in 2016; it may grow even larger in the years ahead.

[37]For evidence that higher wage workers have experienced a smaller decline in labor income in response to the Covid health crisis, see Adams-Prassl et al. (2020), Cajner et al. (2020), Mongey et al. (2020), and Hoke et al. (2020). For evidence that higher income individuals have cut spending by more, see Cox et al. (2020) and Chetty et al. (2020). A recent speech by Haldane (2020) shows a large rise in saving rates by high income individuals in the United Kingdom.

References

Adams-Prassl, Abigail, Teodora Boneva, Marta Golin, and Christopher Rauh, "Inequality in the Impact of the Coronavirus Shock: Evidence from Real Time Surveys," *Working Paper*, 2020.

Adelino, Manuel, Antoinette Schoar, and Felipe Severino, "Dynamics of housing debt in the recent boom and great recession," *NBER Macroeconomics Annual*, 2018, *32* (1), 265–311.

Aguiar, Mark and Mark Bils, "Has consumption inequality mirrored income inequality?," *American Economic Review*, 2015, *105* (9), 2725–56.

Albanesi, Stefania, Giacomo De Giorgi, and Jaromir Nosal, "Credit growth and the financial crisis: A new narrative," Technical Report, National Bureau of Economic Research 2017.

Arbogast, Tobias, "Who are these bond vigilantes anyway? The political economy of sovereign debt ownership in the eurozone," Technical Report, MPIfG Discussion Paper 2020.

Atkinson, Anthony B, Thomas Piketty, and Emmanuel Saez, "Top incomes in the long run of history," *Journal of economic literature*, 2011, *49* (1), 3–71.

Attanasio, Orazio P and Luigi Pistaferri, "Consumption inequality," *Journal of Economic Perspectives*, 2016, *30* (2), 3–28.

Auten, Gerald and David Splinter, "Income inequality in the United States: Using tax data to measure long-term trends," 2019.

Autor, David H, Lawrence F Katz, and Melissa S Kearney, "Trends in US wage inequality: Revising the revisionists," *The Review of economics and statistics*, 2008, *90* (2), 300–323.

Bakija, Jon, Adam Cole, and Bradley Heim, "Jobs and Income Growth of Top Earners and the Causes of Changing Income Inequality: Evidence from US Tax Return Data," Technical Report, Department of Economics, Williams College 2012.

Bartscher, Alina K, Moritz Kuhn, Moritz Schularick, and Ulrike I Steins, "Modigliani Meets Minsky: American Household Debt, 1949-2016," Technical Report, Mimeo 2020.

Batty, Michael, Jesse Bricker, Joseph Briggs, Elizabeth Holmquist, Susan McIntosh, Kevin B Moore, Sarah Reber, Molly Shatto, Tom Sweeney, and Alice Henriques, "Introducing the Distributional Financial Accounts of the United States," 2019.

Bernanke, Ben, "The global saving glut and the US current account deficit," Technical Report, Board of Governors of the Federal Reserve System (US) 2005.

Bertrand, Marianne and Adair Morse, "Trickle-down consumption," *Review of Economics and Statistics*, 2016, *98* (5), 863–879.

Bhutta, Neil and Benjamin J Keys, "Interest rates and equity extraction during the housing boom," *American Economic Review*, 2016, *106* (7), 1742–74.

Blundell, Richard, Luigi Pistaferri, and Ian Preston, "Consumption inequality and partial insurance," *American Economic Review*, 2008, *98* (5), 1887–1921.

Bricker, Jesse, Alice Henriques, and Peter Hansen, "How much has wealth concentration grown in the United States? A re-examination of data from 2001-2013," 2018.

_ and Alice M Henriques, "Wealth Concentration Levels and Growth: 1989-2016," Technical Report, Board of Governors of the Federal Reserve System (US) 2020.

Caballero, Ricardo J and Arvind Krishnamurthy, "Excessive dollar debt: Financial development and underinsurance," *The Journal of Finance*, 2003, *58* (2), 867–893.

_ and _ , "Global imbalances and financial fragility," *American Economic Review*, 2009, *99* (2), 584–88.

_ , Emmanuel Farhi, and Pierre-Olivier Gourinchas, "An equilibrium model of" global imbalances" and low interest rates," *American economic review*, 2008, *98* (1), 358–93.

_ , _ , and _ , "The safe assets shortage conundrum," *Journal of Economic Perspectives*, 2017, *31* (3), 29–46.

Cajner, Tomaz, Leland D Crane, Ryan A Decker, John Grigsby, Adrian Hamins-Puertolas, Eric Hurst, Christopher Kurz, and Ahu Yildermaz, "The U.S. Labor Market During the Beginning of the Pandemic Recession," 2020.

Carroll, Christopher D, Thomas F Crossley, and John Sabelhaus, *Improving the Measurement of Consumer Expenditures*, Vol. 74, University of Chicago Press, 2015.

Catherine, Sylvain, Max Miller, and Natasha Sarin, "Social Security and Trends in Inequality," *Available at SSRN*, 2020.

CBO, "The distribution of household income, 2016," *CBO Slides*, 2019, . (.), .

Chen, Peter, Loukas Karabarbounis, and Brent Neiman, "The global rise of corporate saving," *Journal of Monetary Economics*, 2017, *89*, 1–19.

Chetty, Raj, John N Friedman, Nathaniel Hendren, and Michael Stepner, "How Did COVID-19 and Stabilization Policies Affect Spending and Employment? A New Real-Time Economic Tracker Based on Private Sector Data," 2020.

Cox, Natalie, Peter Ganong, Pascal Noel, Joseph Vavra, Arlene Wong, Diana Farrell, and Fiona Greig, "Initial impacts of the pandemic on consumer behavior: Evidence from linked income, spending, and savings data," *University of Chicago, Becker Friedman Institute for Economics Working Paper*, 2020, (2020-82).

Cynamon, Barry Z and Steven M Fazzari, "Inequality, the Great Recession and slow recovery," *Cambridge Journal of Economics*, 2015, *40* (2), 373–399.

Devlin-Foltz, Sebastian, Alice Henriques Volz, and John Sabelhaus, "Are Disappearing Employer Pensions Contributing to Rising Wealth Inequality?," *FED Notes*, 2019, (2019-02), 01.

Dynan, Karen E, Jonathan Skinner, and Stephen P Zeldes, "Do the rich save more?," *Journal of political economy*, 2004, *112* (2), 397–444.

Eccles, Marriner, *Beckoning Frontiers: Public and Personal Recollections*, New York: Alfred A. Knopf, 1951.

Fisher, Jonathan, David Johnson, Timothy Smeeding, and Jeffrey Thompson, "Inequality in 3-D: Income, Consumption, and Wealth," *Board of Governors of the Federal Reserve System Finance and Economics Discussion Series 2018-001*, 2018.

Foote, Christopher L, Lara Loewenstein, and Paul S Willen, "Cross-sectional patterns of mortgage debt during the housing boom: evidence and implications," Technical Report, National Bureau of Economic Research 2016.

Gourinchas, Pierre-Olivier and Helene Rey, "From world banker to world venture capitalist: US external adjustment and the exorbitant privilege," in "G7 current account imbalances: sustainability and adjustment," University of Chicago Press, 2007, pp. 11–66.

Greenwood, Robin, Samuel G Hanson, and Gordon Y Liao, "Asset price dynamics in partially segmented markets," *The Review of Financial Studies*, 2018, *31* (9), 3307–3343.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song, "What do data on millions of US workers reveal about life-cycle earnings risk?," Technical Report, National Bureau of Economic Research 2019.

_ , Greg Kaplan, Jae Song, and Justin Weidner, "Lifetime incomes in the United States over six decades," Technical Report, National Bureau of Economic Research 2019.

Hager, Sandy Brian, *Public debt, inequality and power. The making of a modern debt state*, University of California Press, 2016.

Haldane, Andrew G, "The Second Quarter," Technical Report, Bank of England Speech 2020.

Heathcote, Jonathan and Fabrizio Perri, "Wealth and volatility," *The Review of Economic Studies*, 2018, *85* (4), 2173–2213.

_ , _ , and Giovanni L Violante, "Unequal we stand: An empirical analysis of economic inequality in the United States, 1967–2006," *Review of Economic dynamics*, 2010, *13* (1), 15–51.

Hobson, John Atkinson, *Imperialism: A Study*, New York: James Pott, 1902.

Hoke, Sinem Hacioglu, Diego Kanzig, and Paolo Surico, "Consumption in the time of Covid-19: Evidence from UK transaction-level data," Technical Report, Working paper 2020.

Jordà, Òscar, Moritz Schularick, and Alan M Taylor, "The great mortgaging: housing finance, crises and business cycles," *Economic Policy*, 2016, *31* (85), 107–152.

Kalecki, Michal, "Political aspects of full employment," *New York and London*, 1943.

Kaplan, Steven N and Joshua Rauh, "It's the market: The broad-based rise in the return to top talent," *Journal of Economic Perspectives*, 2013, *27* (3), 35–56.

Katz, Lawrence F and Kevin M Murphy, "Changes in relative wages, 1963–1987: supply and demand factors," *The quarterly journal of economics*, 1992, *107* (1), 35–78.

Keynes, John Maynard, *The General Theory of Employment, Interest, and Money*, New York: Harcourt, Brace, 1936.

Klein, Matthew C and Michael Pettis, *Trade wars are class wars: How rising inequality distorts the global economy and threatens international peace*, Yale University Press, 2020.

Koijen, Ralph SJ and Motohiro Yogo, "A demand system approach to asset pricing," *Journal of Political Economy*, 2019, *127* (4), 1475–1515.

Kopczuk, Wojciech, Emmanuel Saez, and Jae Song, "Earnings inequality and mobility in the United States: evidence from social security data since 1937," *The Quarterly Journal of Economics*, 2010, *125* (1), 91–128.

Krueger, Dirk and Fabrizio Perri, "Does income inequality lead to consumption inequality? Evidence and theory," *The Review of Economic Studies*, 2006, *73* (1), 163–193.

Kuhn, Moritz, Moritz Schularick, and Ulrike I Steins, "Income and wealth inequality in america," *Journal of Political Economy*, 2019.

Kumhof, Michael, Romain Rancière, and Pablo Winant, "Inequality, leverage, and crises," *American Economic Review*, 2015, *105* (3), 1217–45.

Meyer, Bruce D and James X Sullivan, "Consumption and Income Inequality in the US Since the 1960s," Technical Report, National Bureau of Economic Research 2017.

Mian, Atif, Amir Sufi, and Emil Verner, "How does credit supply expansion affect the real economy? the productive capacity and household demand channels," *The Journal of Finance*, 2020.

_ and _ , *House of debt: How they (and you) caused the Great Recession, and how we can prevent it from happening again*, University of Chicago Press, 2015.

_ , Ludwig Straub, and Amir Sufi, "Indebted Demand," *The Quarterly Journal of Economics*, 2021, *forthcoming*.

Mian, Atif R and Amir Sufi, "Household Debt and Defaults from 2000 to 2010: The Credit Supply View," *Evidence and Innovation in Housing Law and Policy*, 2017.

Mongey, Simon, Laura Pilossoph, and Alex Weinberg, "Which Workers Bear the Burden of Social Distancing Policies?," *Becker Friedman Institute for Economic White Paper*, 2020.

Nardi, Mariacristina De, Giulio Fella, and Gonzalo Paz-Pardo, "Nonlinear household earnings dynamics, self-insurance, and welfare," *Journal of the European Economic Association*, 2018.

Pettis, Michael, "Why a savings glut does not increase savings," Technical Report, Carnegie Endowment for International Peace 2017.

Philippon, Thomas, "Has the US finance industry become less efficient? On the theory and measurement of financial intermediation," *American Economic Review*, 2015, *105* (4), 1408–38.

Piketty, Thomas and Emmanuel Saez, "Income inequality in the United States, 1913–1998," *The Quarterly journal of economics*, 2003, *118* (1), 1–41.

_ , _ , and Gabriel Zucman, "Distributional national accounts: methods and estimates for the United States," *The Quarterly Journal of Economics*, 2018, *133* (2), 553–609.

Rachel, Lukasz and Lawrence H Summers, "On falling neutral real rates, fiscal policy and the risk of secular stagnation," in "Brookings Papers on Economic Activity BPEA Conference Drafts, March 7-8" 2019.

_ and Thomas D Smith, "Are low real interest rates here to stay?," *International Journal of Central Banking*, 2017, *13* (3), 1–42.

Rajan, Raghuram G, *Fault lines: How hidden fractures still threaten the world economy*, princeton University press, 2011.

Reinhardt, Hedwig, "On the Incidence of Public Debt," *Social Research*, 1945, pp. 205–226.

Saez, Emmanuel and Gabriel Zucman, "Wealth inequality in the United States since 1913: Evidence from capitalized income tax data," *The Quarterly Journal of Economics*, 2016, *131* (2), 519–578.

_ and _ , "Comments on Smith, Zidar, and Zwick," *Available at: http://gabriel-zucman.eu/files/SZ2020CommentsOnSZZ.pdf*, 2020.

Slesnick, Daniel T, *Consumption and social welfare: Living standards and their distribution in the United States*, Cambridge University Press, 2001.

Smith, Matthew, Danny Yagan, Owen Zidar, and Eric Zwick, "Capitalists in the Twenty-first Century," *The Quarterly Journal of Economics*, 2019, *134* (4), 1675–1745.

_ , Owen M Zidar, and Eric Zwick, "Top Wealth in the United States: New Estimates and Implications for Taxing the Rich," Technical Report, Working Paper 2020.

Stiglitz, Joseph E, "Inequality and economic growth," in "Rethinking Capitalism" 2016, pp. 134–155.

Straub, Ludwig, "Consumption, Savings, and the Distribution of Permanent Income," *Unpublished manuscript, Harvard University*, 2019.

Summers, Lawrence H, "US economic prospects: Secular stagnation, hysteresis, and the zero lower bound," *Business Economics*, 2014, *49* (2), 65–73.

Wolff, Edward N, "Household wealth trends in the United States, 1962 to 2016: has middle class wealth recovered?," Technical Report, National Bureau of Economic Research 2017.

Appendix for Online Publication

A Appendix for Sections 2 and 3

A.1 Survey-based saving rates and Social Security

An approach not used here would be to use an estimate of after-tax saving rates from survey data. The main drawback of such an approach is that the relevant income measure for calculating the contribution to national savings from any group must include all income, not just income reported in surveys. The claim on business savings (or undistributed corporate profits, S^{π} in equation 4 above) is one important example. Such savings have been rising over time (e.g., Chen et al. (2017)) and have represented 4.5% of National Income since the Great Recession. These savings would be missed in an approach using survey measures of saving rates.

Furthermore, Heathcote et al. (2010) show an average gap of 21 percentage points between the NIPA measure of personal income and the measure in the Current Population Survey. They show that most of the difference comes from the fact that NIPA includes employer contributions to pension and health care plans and the dividends and interest payments realized on pensions that are not distributed to households. The bottom line is that any approach using survey data to estimate the contribution to national savings from any group will be systematically underestimated given these important sources of savings that are missed in surveys.

Recent research suggests that expected Social Security benefits should also be counted as wealth (e.g., Catherine et al. (2020)). The national accounting framework used here takes fully into account taxes that are paid into Social Security (part of T), benefits received by beneficiaries (part of R), and any saving that the government does within the Social Security system (part of S^g). Beyond these flows (which are already taken into account in the analysis below), it would be incorrect in a national accounting framework to claim that individuals "save" through Social Security. In the national accounting framework used here, savings must be real resources that ultimately flow to some other sector in the economy. If the government promises its citizens payments in the future but does not invest in an asset that backs such a claim, then it is not savings in a national accounting framework.[38] The discipline of the national accounts is important because the government could

[38]Comparing the 1983 to 2016 period with the 1963 to 1982 period, Social Security increased its annual average savings by 0.9 percentage points of national income, which is already captured in the S^g above. As shown below, the bottom 90% decreased their annual average savings by between 4 and 6 percentage points of national income depending on the methodology used. Even under the most extreme assumption that all of the savings by Social Security is assigned to the bottom 90%, it would still be small compared to the decline in the savings of the bottom 90% over the same time period. We prefer not to assign Social Security's savings to the household sector, following the logic of the national accounts that S^g is separate from Θ. For example, while the Social Security system increased its savings, the government as a whole reduced its savings by 2.3 percentage points of national income annually; assigning even

make infinite promises of future benefit payments, but such promises would have zero effect on national savings unless the government put real resources behind such promises.[39]

A.2 Consumption share of the top 1% in the income less consumption approach

Figure A1: Consumption Share of the Top 1% of the Income Distribution

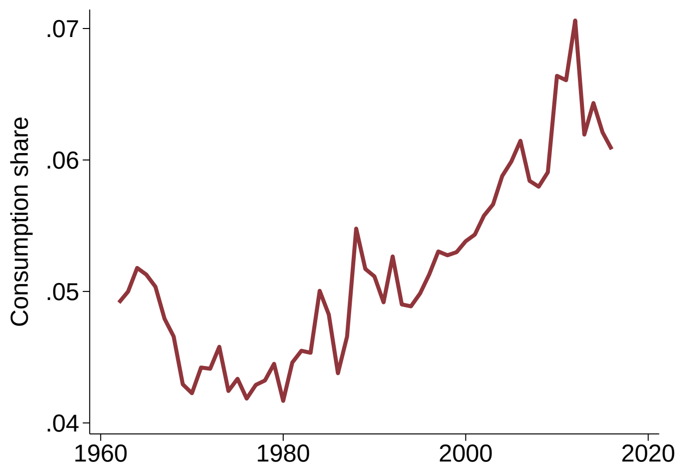

The average consumption share of the top 1% from 2004 to 2016 comes from Fisher et al. (2018). This average consumption share is used as the baseline in the year 2010. Then the time series is generated using the assumption of a constant consumption to income ratio.

half of this decline in savings to the bottom 90% would more than offset the rise in savings by Social Security.

[39] A similar argument applies to unfunded defined benefit pensions, which we also exclude when estimating the savings by households.

A.3 Fixed income capitalization factor for top 1% of wealth distribution

Figure A2: Capitalization Factor of Top 1% of the Wealth Distribution from SCF

This figure plots the average fixed income asset return for the top 1% of the wealth distribution divided by the average fixed income asset return for the entire population in the SCF. The average factor is almost exactly 1.3.

Table A1: Top 1% Savings for Different Factors of Top 1% Fixed Income Returns

Multiplicative Factor	1963-1979	2001-2016	Difference
1.0X	0.029	0.063	0.034
1.3X	0.027	0.056	0.029
1.6X	0.026	0.051	0.025

This table shows measures of average annual savings for the top 1% of the wealth distribution using the DINA microfiles and different assumed factors for the ratio of fixed income asset return for the top 1% to fixed income asset return for all households. The average annual savings scaled by national income is shown for the 1963 to 1979 period and the 2001 to 2016 period, and the difference is also shown.

A.4 More details on wealth-based approach to measuring savings

This section describes the data underlying the asset inflation measures, and explores alternative methods for calculating the synthetic savings by the cohorts from wealth, which is described in Section 2.4.

We first describe in detail the construction of the $\pi_t^{ij} = 1 - WD_t^{ij}$ for debt. We begin by constructing the net chargeoff rate on mortgage and non-mortgage debt for debt borrowed by top 10% and the bottom 90% separately.

Using Call report data, we calculate net chargeoff rate on mortgage and non-mortgage consumer debt. While not all household debt is held on banks balance sheets, we proceed with the assumption that household debt held outside of the banking sector has similar net chargeoff rate as debt held by banks directly. Debt held by non-bank entities such as GSEs is likely to be less risky and hence have lower net chargeoff ratio. However, there are other non-bank entities in the shadow banking sector, such as hedge funds, that are likely to hold the most risky debt and hence have a higher net chargeoff rate. We assume that these two factors cancel out and use bank-held debt net chargeoff rate as representative of overall net chargeoff rate.

We construct annual net charge off rate as net charge offs divided by the total outstanding debt using information in Call report data. This gives us a net chargeoff rate series for mortgage debt from 1991 to 2016, and for non-mortgage consumer debt from 1983 to 2016. Net charge off on mortgage debt is not available as a separate line item prior to 1991. We therefore impute net chargeoff rate on mortgage debt from 1983 to 1990 using non-mortgage consumer credit charge off rate and charge off rate on all loans issued by banks as predictors. In particular, we regress net chargeoff rate for mortgage debt between 1991 and 2016 on net chargeoff rate on non-mortgage consumer debt and net chargeoff rate on all bank loans. The R-sq of this regression is quite high at 0.75. We then use the predicted coefficient to predict net chargeoff rate on mortgage debt from 1983 to 1990.

Prior to 1983, call report data only allows us to construct an overall net chargeoff rate, i.e. chargeoff rate for all debt on banks balance sheets. We use this overall series to extend net chargeoff rate for mortgage and non-mortgage debt back to 1962 by regressing each of these two series (when available) directly on the overall net chargeoff rate series and using the predicted coefficients to predict net chargeoff rate back to 1962.

Once we have annual net chargeoff rate on mortgage and non-mortgage debt, we calculate how much of debt write down was on debt borrowed by the top 10% versus the bottom 90%. We do this using zipcode level data on consumer borrowing from Equifax and merging income data from the IRS. We first multiple total mortgage and total non-mortgage debt across zip codes in the U.S. to calculate the total dollar amount of debt written down every year. We then allocate the total written down amount to zip codes based on the share of total debt default that the zipcode has. We sort zip codes by their average income per capita (income measured by aggregate gross income)[40] and categorize zip codes into top 10% and bottom 90% by income (population weighted). Finally, we calculate the ratio of total written down debt amount to total outstanding debt within each income category and for both mortgage and non-mortgage debt separately.

[40]The IRS data is missing for certain years early on, in which case we use the latest available IRS data.

The above procedure allows us to compute debt write down rate WD_t^{ij} for j equal to mortgage and non-mortgage debt, and i equal to top 10% and bottom 90% from 1991 to 2016. There is no zip code level Equifax data prior to 1991. However, we can use Equifax-based estimates to impute WD_t^{ij} for year prior to 1991 by regressing WD_t^{ij} on US-level net chargeoff rate for mortgage and non-mortgage debt respectively for years 1991 to 2016. We then use the predicted coefficients and data on net chargeoffs at the US level to back-fill WD_t^{ij} from 1962 to 1990. Exact details of all of our procedure can be seen in the accompanying code that is made public.

Next we discuss the computation of π_t^j for equity. The starting point for this imputation is the wealth-implied private saving in the aggregate, $\Theta_t = \sum_{j \in J} \left(\Delta W_t^j - \pi_t^j W_{t-1}^j \right)$. The key observation is that we know Θ_t at the aggregate level from NIPA. We have also calculated π_t^j for all assets other than equity. We can therefore solve for π_t^j for equity assets.

A.5 Mapping for wealth-based approach

Table A2: Mapping for Wealth-Based Approach

Asset Class	**Flow of Funds Equivalent**	**Asset inflation**	**DINA Asset Class**	**DFA Asset Class**
Assets				
Real estate	LM155035015	Repeat-sales index (JST)	Owner-occupied housing	Real estate
Foreign deposits	LM153091003	0	Taxable bonds	Time deposits and short-term investments
Checkable deposits and currency	FL153020005-(FL893131573-FL543131503)*(Chk. deposits share, IRAs)	0	Currency	Checkable deposits and currency
Time and savings deposits	FL153030005-(FL893131573-FL543131503)*(Time deposits share, IRAs)	0	Taxable bonds	Time deposits and short-term investments
Money market fund shares	FL153034005-(FL893131573-FL543131503)*(MMF share, IRAs)	0	Taxable bonds	Money market fund shares
Treasury securities	LM153061105-(FL893131573-FL543131503)*(Treasuries share, IRAs)	0	Taxable bonds	US government and municipal securities
Agency- & GSE-backed securities	LM153061705-(FL893131573-FL543131503)*(GSE share, IRAs)	0	Taxable bonds	Debt securities
Municipal securities	LM153062005	0	Municipal securities	US government and municipal securities
Corporate & foreign bonds	LM153063005-(FL893131573-FL543131503)*(Bonds share, IRAs)	0	Taxable bonds	Corporate & foreign bonds
Loans & advances	FL153069005	0	Taxable bonds	Loans & advances
Mortgages	FL153065005-(FL893131573-FL543131503)*(Mortgages share, IRAs)	0	Taxable bonds	Mortgages
Corporate equities	LM153064105-(FL893131573-FL543131503)*(Equity share, IRAs)	Residual	Equity	Corporate equities and mutual funds
Mutual funds, equity portion	(LM153064205-(FL893131573-FL543131503)*(Mut. fund share, IRAs))*(Equity share, mut. fund)	Residual	Equity	Corporate equities and mutual funds
Mutual funds, munis portion	(LM153064205-(FL893131573-FL543131503)*(Mut. fund share, IRAs))*(Munis share, mut. fund)	0	Municipal securities	Corporate equities and mutual funds
Mutual funds, bonds portion	(LM153064205-(FL893131573-FL543131503)*(Mut. fund share, IRAs))*(Fix. inc. share, mut. fund)	0	Taxable bonds	Corporate equities and mutual funds
Life insurance reserves, equity portion	FL153040005*(Equity share, life ins. reserves)	Residual	Pensions	Life insurance reserves
Life insurance reserves, fixed income	FL153040005*(Fixed income share, life ins. reserves)	0	Pensions	Life insurance reserves
Pensions, equity portion	FL153050025*(Equity share, pensions)	Residual	Pensions	Pension entitlements
Pensions, fixed income portion	FL153050025*(Fixed income share, pensions)	0	Pensions	Pension entitlements
IRAs, equity portion	(FL893131573-FL543131503)*(Equity share, IRAs)	Residual	Pensions	N/A
IRAs, fixed income portion	(FL893131573-FL543131503)*(Fixed income share, IRAs)	0	Pensions	N/A
Equity in non-corporate business	LM152090205	Residual	Business	Equity in non-corporate business
Miscellaneous assets	FL153090005	Residual	Wealth minus housing	Miscellaneous assets
Liabilities				
Home mortgages	FL153165105	Write-down rate, home mort.	Owner-occupied mortgage debt	Home mortgages
Consumer credit	FL153166000	Write-down rate, cons. cred.	Non-mortgage debt	Consumer credit
Depository institution loans	FL153168005	Write-down rate, cons. cred.	Non-mortgage debt	Consumer credit
Other loans	FL153169005	Write-down rate, cons. cred.	Non-mortgage debt	Consumer credit
Deferred & unpaid life ins. premiums	FL543077073	Write-down rate, cons. cred.	Non-mortgage debt	Consumer credit

This table shows the mapping of asset classes between the Financial Accounts, DINA, and the Distributional Financial Accounts.

A.6 Population aging

An interesting area for future research is the interaction of population aging and the saving glut of the rich. Unfortunately, the data sets used in this study are not ideal for a comprehensive investigation of this question. The DINA microfiles, for example, only contain information placing individuals into one of two broad age groups for the entire sample period: people over and under 65.

Three facts give us confidence that population aging is unlikely to be an important driver of the rise in savings by the top 1%. First, while the population as a whole has been aging, the top 1% of the wealth distribution is almost identical in terms of the composition of the two age groups we can measure. Comparing the first 10 years of the sample (1963-1972) to the last 10 years of the sample (2007-2016), the fraction of all tax filers that are 65 or older has increased by 3 percentage points. However, within the top 1% of the wealth distribution, the fraction has actually fallen slightly by 0.8 percentage points. Second, among the top 1% of the wealth distribution, the saving rates of the two groups were similar prior to the rise in the saving glut of the rich. Third, research by Guvenen et al. (2019b) shows that the rise in the top 1% share of lifetime income is large even within-cohort (see the bottom right panel in Figure 16 of the January 2019 draft), which suggests that the rise in the top 1% share is not driven by demographic changes.

A.7 Implied saving rate for the top 1%

Table A3 shows the average saving rate out of own-group income for the top 1% across the different approaches used to measure savings in Sections 2.3 and 2.4. Recall that the DINA and CBO approach sort individuals by income, and the wealth-based approach sorts by wealth. The average saving rates for the top 1% vary between 0.38 and 0.63.

Table A3: Implied Saving Rates for the Top 1%

Period	DINA	CBO	Wealth-based
63-82	0.628	0.543	0.383
83-97	0.628	0.543	0.393
98-07	0.628	0.543	0.541
08-15	0.628	0.543	0.403

This table presents the average annual saving rates for the top 1% of the income distribution for the DINA and CBO columns, and the annual average saving rate of the top 1% of the wealth distribution for the wealth-based column. Income in the denominator is the total amount of national income earned by each group.

The best estimates of the saving rate of the top 1% of the income distribution come from the

Survey of Consumer Finances. For example, Dynan et al. (2004) use the 1983 to 1989 panel dimension of the SCF and estimate a saving rate of the top 1% of the income distribution of 0.5. The Fisher et al. (2018) study estimates an average consumption share of the top 1% from 2004 to 2016 in the SCF of 0.066. The average after-tax income share of the top 1% in the SCF during the same years is 0.173. The consumption to income ratio of the top 1% in the SCF can be calculated with these two numbers along with an adjustment factor for the levels of income and consumption, given that the consumption share of the top 1% divided by the income share of the top 1% multiplied by the aggregate level of consumption divided by aggregate level of income gives the consumption to income ratio of the top 1%. If one uses the ratio of aggregate Disposable Personal Income to Personal Consumption Expenditures estimate from the NIPA, it produces an average consumption to income ratio of the top 1% in the SCF from 2004 to 2016 of (0.066/0.173)*(0.91) = 0.35, which would imply a saving rate of the top 1% of 0.65.[41]

In general, the saving rate from survey data is not directly comparable to the saving rate calculated in this study because survey data miss important income that has a 100% saving rate. The most obvious example is undistributed corporate profits. Undistributed corporate profits represent saving by shareholders; such saving is captured in the methodology used in this study but missed in surveys. In addition, as mentioned in the text, Heathcote et al. (2010) show an average gap of 21 percentage points between the NIPA measure of personal income and the measure in the Current Population Survey. It is likely that a substantial amount of this missing income represents income with a high saving rate, such as employer contributions to pension plans or the internal dividends and interest payments earned by pensions.

Formally, suppose the saving rate of the top 1% estimated in survey data is $\hat{\phi_{top1}} = \frac{\Theta_{top1}}{Y_{top1}}$. Furthermore, let ψ_{top1} be income missing in surveys of the top 1% that has a 100% saving rate. Then the correct saving rate for the top 1% would be: $\phi = \frac{\Theta+\psi}{Y+\psi}$ where the subscript is removed in order to reduce clutter. The incorrect saving rate would have to be adjusted to:

$$\phi = \frac{\hat{\phi} + \frac{\psi}{Y}}{1 + \frac{\psi}{Y}} \tag{14}$$

The critical ratio ($\frac{\psi}{Y}$) is the amount of income missing in surveys that has a 100% saving rate relative to the amount of reported income in surveys. If this ratio is 0.3, for example, then the saving rate of the top 1% from surveys of 0.50 would imply a true saving rate of 0.62.

This ratio is difficult to estimate, given the lack of estimates of the income that accrues to the top 1% that is not included in surveys. However, as a lower-bound estimate, one can calculate the part of ψ that comes from the top 1% claim on undistributed corporate profits, which we know is

[41]We thank David Johnson for helping us with this calculation.

not included in surveys as income and has a 100% saving rate. From 1998 to 2016, $\frac{\psi}{Y}$ is estimated to be 0.25 using undistributed corporate profits alone, which would imply a true saving rate of 0.60 given a survey-reported saving rate of 0.50. This lower bound estimate convinces us that the saving rates shown in Table A3 are realistic for the top 1% once all sources of income are included.

B Appendix for Section 4

B.1 The large rise in debt

Figure A3: Government and Household Debt

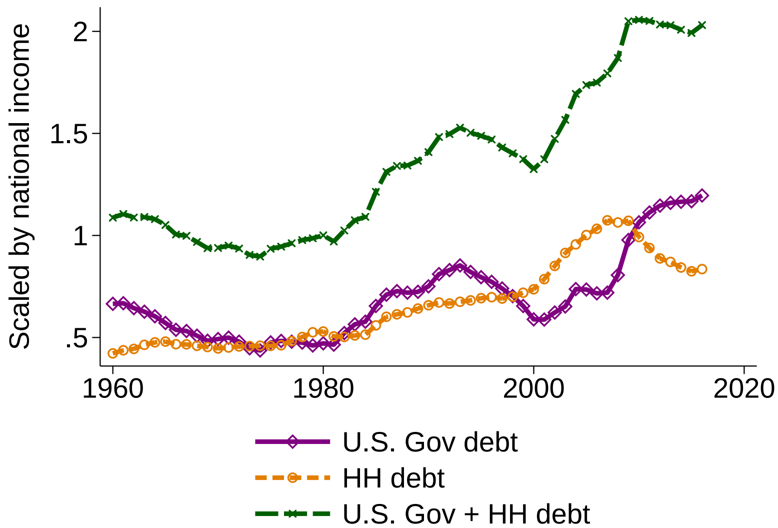

This figure plots government and household debt in the United States over time, scaled by national income.

B.2 Further details on unveiling

This section contains a few extra notes on the unveiling process. All code and data for the unveiling exercise are included in the replication kit. Below, we discuss specifically the unveiling process for household debt. But the government debt unveiling process is similar.

Formally, the unveiling procedure computes the entire set of household debt shares $\overline{F}_j$, including those asset classes j that are not directly owned by households. Denote by $\overline{J}$ the number of all such asset classes. The equation pinning down $\overline{F}$ is a "financial input-output network". Specifically, denote by $\eta_{j',j}$ the share of asset class j's liabilities that are owned by asset class j'; and denote by η_j^{HHD} share of household debt directly owned by asset class j. Observe that

$$\sum_{j=1}^{\overline{J}} \eta_j^{HHD} = 1 \quad \text{and} \quad \sum_{j'=1}^{\overline{J}} \eta_{j',j} \leq 1$$

where the latter inequality is strictly less than 1 for any asset class j that is partly owned by households. $\overline{F}$ must then satisfy

$$\begin{bmatrix} \overline{F}_1 \\ \overline{F}_2 \\ \vdots \\ \overline{F}_{\overline{J}} \end{bmatrix} = \begin{bmatrix} \eta_1^{HHD} \\ \eta_2^{HHD} \\ \vdots \\ \eta_{\overline{J}}^{HHD} \end{bmatrix} + \begin{bmatrix} \eta_{1,1} & \eta_{1,2} & \cdots & \eta_{1,\overline{J}} \\ \eta_{2,1} & \eta_{2,2} & \cdots & \eta_{2,\overline{J}} \\ \vdots & \vdots & \ddots & \vdots \\ \eta_{\overline{J},1} & \eta_{\overline{J},2} & \cdots & \eta_{\overline{J},\overline{J}} \end{bmatrix} \begin{bmatrix} \overline{F}_1 \\ \overline{F}_2 \\ \vdots \\ \overline{F}_{\overline{J}} \end{bmatrix} \tag{15}$$

In words, this equation captures that the household debt share of asset class c is equal to its directly owned share η_j^{HHD}, plus the indirectly owned share through other asset classes, $\sum_{j=1}^{\overline{J}} \eta_{j',j} \overline{F}_j$. The matrix product on the right hand side of this equation incorporates our assumption that the household debt owned by asset class j is attributed to its liabilities in proportion to their liabilities shares. In matrix notation, (15) can be expressed as

$$\overline{F} = \eta^{HHD} + H\overline{F}$$

which yields the solution

$$\overline{F} = (I - H)^{-1} \eta^{HHD}$$

The Leontieff inverse matrix, $(I - H)^{-1} = I + H + H^2 + \ldots$, captures any direct and indirect ownership of household debt after an arbitrary number of rounds of unveiling. As explained above, in our case at hand, seven rounds were sufficient to conduct the unveiling.

One important adjustment is made to the DFA shares based on defined benefit pensions. A substantial fraction of defined benefit pension wealth is unfunded. An unfunded DB pension cannot be a claim on household debt because there is no actual financial asset backing the unfunded part of the pension. We therefore exclude the unfunded portion of defined benefit pensions from the measure of wealth, and we re-calculate wealth shares for the top 1%, next 9%, and bottom 90%.

Another issue that is currently ignored in the unveiling process is the fact that financial asset shares of DB and defined contribution pension funds vary across the income distribution (e.g., Devlin-Foltz et al. (2019)). Data kindly shared to us by Alice Henriques Volz based on Devlin-Foltz et al. (2019) shows that from 1989 to 2016, the top 10% share of DC assets was 53% and the top 10% share of DB assets was 48%. When excluding unfunded DB pensions, the shares of overall pensions should be adjusted given that the claim of lower income households on unfunded pensions is larger than their claim on DC pension assets. We do not currently make an adjustment given the fact that the DFA as currently structured does not provide financial asset shares separately for DB and DC pension assets. The lack of this adjustment means that the current methodology overstates the amount of household debt held as a financial asset by the bottom 90% of the income distribution through pensions.

Another issue involves equity of private depository institutions. The Financial Accounts does not include an estimate of the equity of private depository institutions, which must be taken into account when distributing the household debt held by these institutions to other entities. The estimate of private depository institution equity comes from publicly traded banks through CRSP data.

Finally, the unveiling process currently ignores the equity holdings in other financial intermediaries such as the Agency GSEs and life insurance companies. Taking into account these equity holdings will boost the share of household debt held by the top of the income distribution, given that the top of the income distribution holds a larger share of equity than other asset classes.

B.3 Additional graphs for Section 4

Figure A4: Non-financial business deposits and money market fund holdings

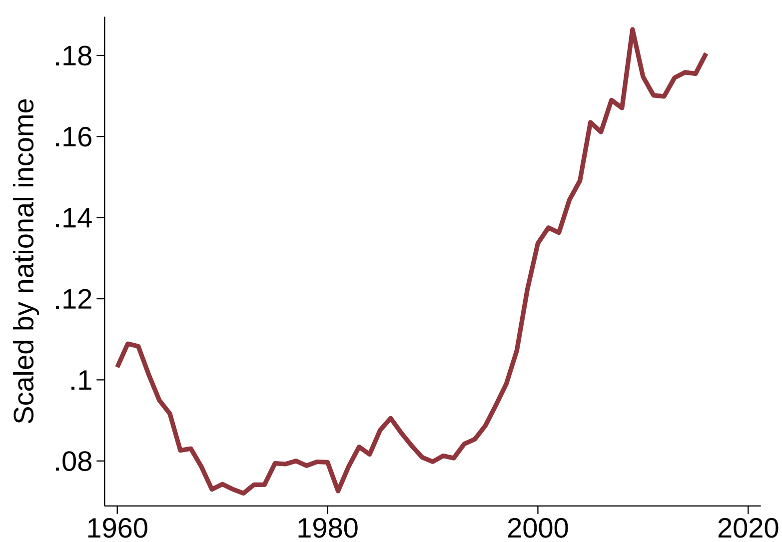

Data are from the Financial Accounts.

C Appendix for State-Level Analysis in Section 5

C.1 More details on state-level data

In this section we describe in detail our procedure for assigning tax returns with income above 200K to individual states for the years 1979-2008, excepting 1982. Recall that the household-level public use tax files files do not contain state identifiers previous to 1979 and after 2008. For 1982 we do not need to do this because the files contain state identifiers for all observations.

As mentioned above, we obtain the mean interest, dividend, and taxable pension income for units with AGI above $200,000 from the SOI aggregate data. In order to utilize these data in the Saez and Zucman (2016) capitalization technique, we also require the mean estate income and nontaxable pension income for these same units. To have data on all asset classes of interest, we additionally need the mean municipal bond and business wealth. To obtain these data, we rely on the US DINA microfiles made available by Piketty et al. (2018), in which we find these income variables and can directly construct the wealth variables with the Saez and Zucman (2016) technique. Given that state identifiers are missing for these top earners, we obtain state-level means by employing a probabilistic sampling approach.

Our key assumption in this approach is that for each state s and year y the distribution of income I for units with AGI above $200,000 is characterized by a Pareto distribution with probability density function $f_{sy}(I) = \frac{\alpha_{sy} 200000^{\alpha_{sy}}}{I^{\alpha_{sy}+1}}$ and mean $E_{sy}[I] = \frac{200000\alpha_{sy}}{\alpha_{sy}-1}$. We in fact do know $E_{sy}[I]$ thanks to the aggregate state-income group level data from the SOI - this is simply the mean AGI for units with income above $200,000. Thus, we can obtain

$$\alpha_{sy} = \frac{E_{sy}[I]}{E_{sy}[I] - 200000}.$$

Similarly, we obtain $\alpha_{US,y}$ using U.S.-level data. For each year, we assign each state a mean estate and nontaxable pension income, as well as a mean municipal bond wealth and business wealth for units with AGI above $200,000 by taking a weighted mean over all observations in the household tax return file with AGI above $200,000. The weights $w_{syi}(I)$ we use are the population weights multiplied by the relative likelihood that a household lives in a state. We calculate this relative likelihood as the ratio between $f_{sy}(I)$ and $f_{U.S.,y}(I)$. Thus in each year y, for each observation i with AGI I_i and population weight p_i, the weight assigned to that observation when constructing the mean for state s is

$$w_{syi}(I_i) = p_i \times \frac{f_{sy}(I_i)}{f_{U.S.,y}(I_i)} = p_i \times \frac{\alpha_{sy}}{\alpha_{USA,y}} \times 200000^{\alpha_{sy}-\alpha_{USA,y}} \times I_i^{\alpha_{USA,y}-\alpha_{sy}}. \tag{16}$$

Having done this, we assign each observation, representing all filers with AGI above $200,000 in a state, the appropriate population weight based on the number of returns filed by households with AGI over $200,000, as reported in the aggregate SOI data. We then have the mean business and municipal bond wealth for this income group in each state. We use the mean interest, dividend and taxable pension income from the SOI aggregates in conjunction with the mean estate income and nontaxable pension income obtained through this procedure to obtain the capitalized measures of fixed income, equity and pension wealth. Again, since the SOI aggregate data contains the total number of returns with AGI above $200,000 by state, knowing these means is sufficient to know the totals.

As a check, we can use the same procedure to obtain the sampled mean AGI for units above $200,000 by state. Doing this and comparing the values to the true SOI aggregate data, we obtain a correlation of 0.99 and a cross-sectional R^2 of 0.98 between the means in the SOI aggregates and in our sampling. This suggests that our sampling provides a close approximation to the true values of AGI. We make one final adjustment to ensure that our aggregate U.S. values match the true totals for all variables, by scaling as necessary without changing the distribution.

No imputation is required for earners below $200,000 and for all households in 1982 - data for these earners, with state identifiers, are contained in the public-use tax files. From these capitalized measures of total fixed income, equity, business and pension wealth, as well as their subcomponents, we construct a data set that contains, for various income groups in a state and year and for all the asset classes described in Section 2.4, that group's share of the U.S. total. With this, we apply the same unveiling process used at the national level to construct a measure of how much household and government debt is owned as an asset for various income groups in a state and year.

C.2 State-level analysis: additional tables and figures

Figure A5: Percentage of Filers with AGI Above $200,000

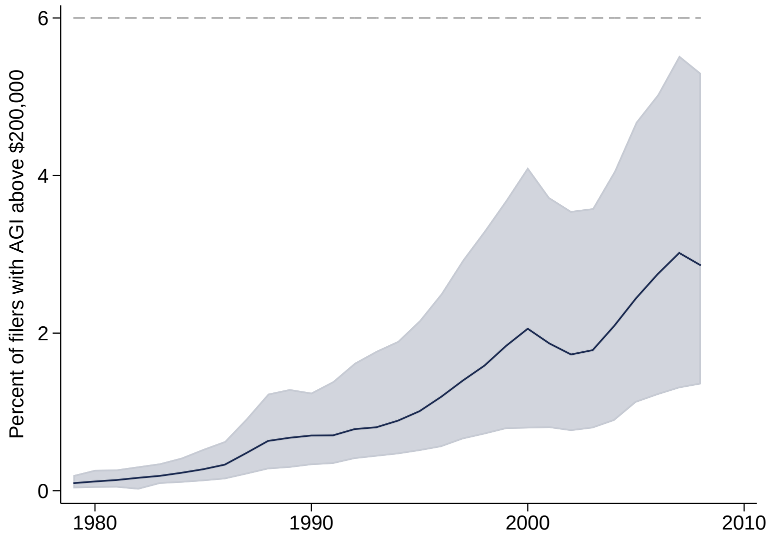

The solid line shows the percentage of tax filers in the U.S. with AGI above $200,000 over time. The shaded area represents the interval that contains this percentage for all states.

Table A4: Rise in Top Income Shares and Saving by wealth category (panel regression)

	(1) Overall wealth	(2) Financial assets	(3) Real estate	(4) Debt liabilities	(5) Household debt assets	(6) Government debt assets	(7) Non-debt financial assets
Top 6% share	0.645***	0.789***	-0.0941***	-0.0496***	0.158***	0.0915***	0.539***
	(0.072)	(0.072)	(0.021)	(0.007)	(0.012)	(0.008)	(0.056)
Log per capita income	-0.00693	-0.00391	-0.00689	0.00387	-0.00380	0.000986	-0.00110
	(0.029)	(0.028)	(0.007)	(0.003)	(0.005)	(0.004)	(0.020)
R^2	0.59	0.70	0.60	0.50	0.76	0.80	0.71
Observations	306	306	306	306	306	306	306

The left hand side θ_{ist} reflects savings by group i in state s in year t scaled by state national income. The top 6% share is the share of income by the top 6% of the income distribution in state s in year t. Robust standard errors, clustered by state, are in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table A5: Rise in Top Income Shares and Saving by wealth category

	(1) Overall wealth	(2) Financial assets	(3) Real estate	(4) Debt liabilities	(5) Household debt assets	(6) Government debt assets	(7) Non-debt financial assets
Δ Top 6% share	8.950***	8.815***	0.396	-0.488***	1.845***	0.877***	5.916***
	(0.843)	(1.082)	(0.384)	(0.087)	(0.257)	(0.082)	(0.706)
Top 6% share 79-83	4.551**	3.387	0.756	0.222	-0.136	0.162	3.602
	(1.452)	(3.096)	(0.711)	(0.187)	(0.462)	(0.212)	(1.826)
Δ Log per capita income	0.585	0.644*	-0.0732	0.0173	0.0478	0.0662**	0.523**
	(0.307)	(0.275)	(0.119)	(0.027)	(0.057)	(0.022)	(0.189)
Log per capita income 79-83	0.652*	0.377	0.182	0.0692*	0.00467	0.0242	0.355
	(0.259)	(0.281)	(0.100)	(0.026)	(0.062)	(0.027)	(0.192)
Wealth-to-income 79-83	-0.756***	-0.672**	-0.674***	-0.263***	0.508	-0.250	-0.866***
	(0.113)	(0.236)	(0.084)	(0.044)	(0.357)	(0.180)	(0.183)
R^2	0.94	0.94	0.79	0.66	0.91	0.92	0.95
Observations	51	51	51	51	51	51	51

The dependent variable in column 1 is Δ wealth-to-income, the change in net wealth of top 6% in state s scaled by state national income. Columns 2 through 8 break down net wealth into various sub-components. Robust standard errors in parentheses. Δ reflects the change in the variable in question using the average from 2004 to 2008 minus the average from 1979 to 1983. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table A6: Relationships Between Controls and Rise in Top Income Share

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Top 6% share	1.577***						1.111***
	(0.327)						(0.210)
Δ Log per capita income		0.253**					0.0996*
		(0.080)					(0.046)
Log per capita income 79-83			0.291***				0.0513
			(0.059)				(0.089)
Wealth-to-income 79-83				0.0413***			0.0124
				(0.012)			(0.009)
Skill sh 79-83					0.312***		0.000962
					(0.076)		(0.073)
Farm/Agg sh 79-83						-0.743***	-0.321**
						(0.191)	(0.113)
R^2	0.55	0.29	0.38	0.34	0.22	0.33	0.81
Observations	51	51	51	51	51	51	51

Robust standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table A7: Additional Controls on Wealth-to-Income Ratios

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Δ Top 6% share	8.950***	9.047***	9.010***	8.703***	8.703***	9.305***	8.206***	9.007***	9.477***	8.861***
	(0.843)	(0.831)	(0.838)	(0.849)	(0.768)	(0.702)	(0.733)	(0.779)	(0.905)	(0.909)
Top 6% share 79-83	4.551**	4.635*	4.018	5.284***	5.725***	3.013*	3.201*	3.899**	3.509*	4.572**
	(1.452)	(2.167)	(2.013)	(1.490)	(1.529)	(1.174)	(1.224)	(1.288)	(1.488)	(1.452)
Δ Log per capita income	0.585	0.452	0.655	0.628	0.416	0.923***	0.692***	0.518	0.404	0.544
	(0.307)	(0.360)	(0.437)	(0.340)	(0.334)	(0.256)	(0.185)	(0.278)	(0.309)	(0.305)
Log per capita income 79-83	0.652*	0.384	0.529	0.748*	0.822*	0.319	0.271	0.704**	0.445	0.659*
	(0.259)	(0.445)	(0.499)	(0.321)	(0.337)	(0.265)	(0.220)	(0.216)	(0.299)	(0.259)
Wealth-to-income 79-83	-0.756***	-0.742***	-0.726***	-0.766***	-0.851***	-0.760***	-0.662***	-0.674***	-0.668***	-0.744***
	(0.113)	(0.146)	(0.137)	(0.125)	(0.119)	(0.104)	(0.102)	(0.114)	(0.133)	(0.122)
Control 79-83		0.549	0.458	0.628	-1.916***	-0.774*	3.538**	2.046	0.0818*	0.00727
		(0.490)	(1.326)	(0.506)	(0.411)	(0.372)	(1.105)	(1.047)	(0.033)	(0.017)
Δ Control		0.265	-0.695	1.610	-1.821	0.351	-0.738	4.520		
		(0.790)	(1.135)	(1.049)	(0.957)	(0.652)	(0.886)	(2.435)		
Control	None	Skill sh	High skill sh	Old dep. ratio	Young dep. ratio	Manufac. sh	Fin. sh	Cons. sh	Migration	Dereg
R^2	0.94	0.94	0.94	0.95	0.95	0.96	0.96	0.95	0.95	0.94
Observations	51	51	51	51	51	51	51	51	51	51

This table presents a specification that is analogous to column 3 of Table 9 with additional controls. The controls, in order, are the skilled and high skilled shares of labour (at least 1 year of post-secondary education and at least 4 years of post-secondary education, respectively), the old and young dependency ratios, the manufacturing, finance and construction shares of employment, the deregulation measure from Mian et al. (2020), and a measure of rich in-migration. Data come from the US Census Bureau, the Bureau of Economic Analysis and the Current Population Survey. For columns 2 to 8, we control for both the baseline level of the measure and the change in the measure. Columns 9 and 10 control for the level of the measure. Robust standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Figure A6: Change in Top 1% Share Against Change in Top 6% Share

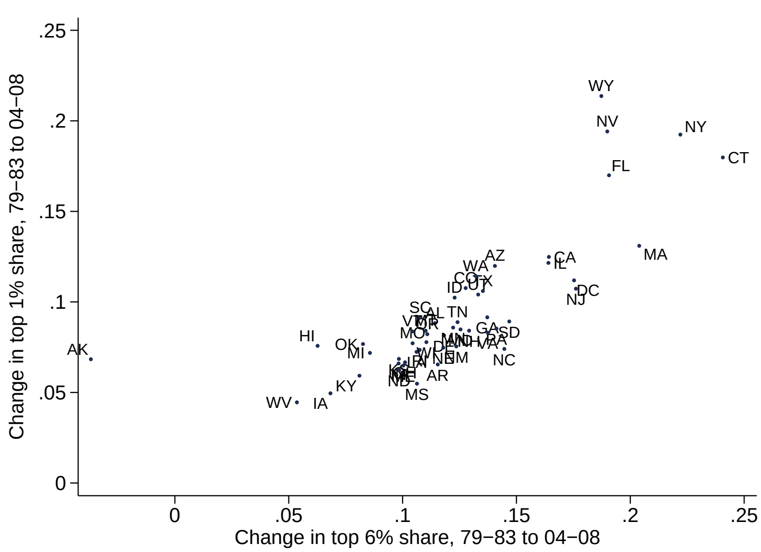

D Results using the DFA

The wealth-based methodology for the calculation of savings can also be done using the DFA wealth shares (ω_{it}^{j}) instead of the DINA wealth shares. The main issue with the DFA is that the data are only available for 1989 onward. As a result, it is impossible to measure the change in savings relative to the 1963 to 1982 period, which is the basis for most of the central results in this study.

Nonetheless, it is possible to compare the levels of savings from 1989 to 2016 implied by the DFA and DINA wealth shares, which is done in Table A8. As before, annual savings are scaled by contemporaneous national income, and the averages for 1989 to 2016 are shown. For the sake of comparison, Table A8 also shows averages using the DINA wealth shares for the same period in question.

Table A8: Results using DFA: 1989 to 2016

	DINA			DFA		
	Top 1%	Next 9%	Bottom 90%	Top 1%	Next 9%	Bottom 90%
Θ	0.048	0.042	-0.001	0.047	0.045	-0.003
ΔV	0.034	0.042	0.066	0.037	0.044	0.061
ΔNW	0.082	0.085	0.065	0.084	0.090	0.058
Θ^{FA}	0.049	0.045	0.035	0.045	0.049	0.034
Θ^{RE}	0.002	0.005	0.009	0.004	0.004	0.006
D	-0.002	-0.007	-0.044	-0.002	-0.009	-0.043
Θ^{HHD}	0.011	0.009	0.005	0.008	0.010	0.008
Θ^{GOVD}	0.011	0.010	0.007	0.009	0.011	0.009
$\Theta^{HHD} + \Theta^{GOVD}$	0.022	0.019	0.013	0.016	0.020	0.017
Θ^{NHHD}	0.009	0.002	-0.039	0.006	0.001	-0.035

This table shows measures of average annual savings across the wealth distribution using the DFA wealth shares for the 1989 to 2016 period for which the DFA data are available. It also shows results average annual savings using the DINA wealth shares over the same time period. Annual savings are scaled by contemporaneous national income, and the averages are shown.