The Saving Glut of the Rich and the Rise in Household Debt [*]

Atif Mian
Princeton & NBER

Ludwig Straub
Harvard & NBER

Amir Sufi
Chicago Booth & NBER

November 2019

Abstract

Rising income inequality since 1980 in the United States has generated a large increase in saving by the top of the income distribution, which we call *the saving glut of the rich*. The saving glut of the rich has been on the same order of magnitude as the global saving glut, and it has not been associated with an increase in investment. An examination of the distribution of income and wealth reveals that a large fraction of the rise in household borrowing by non-rich households has been financed by rich households through this saving glut. Analysis using variation across states in the rise in top income shares shows that income growth at the top of the income distribution can explain 75% of the accumulation of household debt held as a financial asset by households in the United States. After the Great Recession, evidence suggests that the saving glut of the rich has been financing government deficits to a greater degree.

1 Introduction

Rising income inequality since 1980 in the United States has generated a large increase in saving by the top of the income distribution, which we call *the saving glut of the rich*. This saving glut has been closely linked to the substantial dis-saving and large accumulation of household debt of the bottom 90%. Since 1980, the savings of the top of the income distribution increasingly has financed borrowing by the rest of the population. The rise in income inequality and the rise in household debt over the past 40 years have been two sides of the same coin.

This study illustrates this pattern in three empirical settings: (1) an analysis of saving in the National Income and Product Accounts (NIPA), (2) an analysis of the stock of wealth and debt in the Financial Accounts of the Federal Reserve, and (3) in a state-by-year panel in which cross-sectional variation across states in the top income shares is related to the stock of household debt held as a financial asset by individuals in the state.

To measure the saving glut of the rich, the analysis relies on the NIPA combined with measures of the distribution of both national income and consumption. The after-tax shares of national income from Piketty et al. (2017) and the CBO (2019) along with estimates of the consumption share of the top 1% are important inputs to the methodology. The methodology developed here produces an estimate of the contribution to national saving that comes from the top 1% of the income distribution.

This saving contribution of the top 1% of the income distribution was steady from 1970 to 1980 at an average of 4.5 percentage points of national income annually. This increased to 5.5 percentage points annually in the 1980s, to 8.3 percentage points annually from 2000 to 2010, and to 9.5 annually percentage points from 2010 to 2014. To put the size of this saving glut in perspective, since 2000, the annual saving coming from the top 1% of the income distribution has been higher on average than total annual net domestic investment in the United States.

Many scholars have cited the global saving glut as a reason behind the decline in real interest rates to low levels across many advanced economies (e.g., Bernanke (2005)). However, in the context of the United States, the rise in the saving glut of the rich was on the same order of magnitude as the increase in the inflow of capital from overseas from 1980 to 2014. The analysis here suggests that the saving glut of the rich should also be recognized as a potential reason for low real interest rates.

From a national accounting perspective, the saving glut of the rich must have been absorbed by some other part of the economy. In a closed economy, one natural place to look would be net domestic investment or a rise in government borrowing. However, investment actually declined from 1980 to 2014, and government deficits were stable until the Great Recession. In an open economy framework, it is also possible for some of the savings to find its way overseas. But, as is

well known, the net foreign account moves in the opposite direction. The United States as a whole borrowed more from the rest of the world during this time period.

This leaves only one remaining margin: the rest of the U.S. household sector must have reduced saving substantially. This is what the analysis finds. Saving by the bottom 90% of the income distribution fell significantly over this time frame. The rise in saving of the top 1% was associated with a substantial dis-saving by the bottom 90%.

The second part of the study focuses on the evolution of debt and wealth in order to assess whether the accumulation of savings by the rich has been an important source of financing of the debt accumulation of the non-rich. To explore this question, the analysis first allocates household debt owed by U.S. households to three potential providers of capital: the rest of the world, the government, and the U.S. household sector. This allocation is done using the extensive information on the linkages within the financial sector that are detailed in the Financial Accounts of the United States produced by the Federal Reserve. The allocation process is best viewed as an exercise seeking to remove the veil of financial intermediation: the financial linkages between institutions are used to uncover who ultimately holds U.S. household debt as a financial asset. To the best of our knowledge, this unveiling process is novel to the literature, and can be potentially done for other asset classes and in other countries.

This unveiling process reveals that the primary source of financing for debt owed by U.S. households is the U.S. household sector. The methodology then allocates the household debt held as a financial asset by the U.S. household sector across the wealth distribution, using financial asset shares reported by Saez and Zucman (2016) and the Distributional Financial Accounts of the Federal Reserve (explained in detail by Batty et al. (2019)). The final product from the methodology allows us to quantify exactly how much of household debt in the United States represents a financial asset held by the top of the wealth distribution.

The results show that the rise in household debt owed as a liability was driven by the bottom 90% of the income distribution, whereas the rise in household debt held as a financial asset was driven by the top of the wealth distribution. This suggests that a better measure of household debt claims across the wealth distribution is net household debt owed, which is defined as gross household debt owed minus household debt held as a financial asset. Net household debt positions clarify that the post 1980 period was one in which rich Americans increasingly financed the borrowing of non-rich Americans. The net household debt position of the top 1% fell by 20 percentage points of national income since 1980 which reflected their accumulation of household debt held as a financial asset. In contrast, the net household debt position of the bottom 90% increased by 40 percentage points.

The aggregate patterns suggest a link between the rise in top income shares and the rise in household debt held as a financial asset by the wealthy. The third part of the study constructs a novel state-level data set that allows for a tighter statistical test of this link. In particular, the state-

level data allow us conduct an analysis closer to the ideal thought experiment: do economies that experienced the largest increases in top income shares also experience the largest rise in the amount of household debt held as a financial asset? There was substantial variation across states in the rise in top income shares from the 1980s onward. The long-difference specification at the state-level removes common aggregate patterns that occurred since the 1980s. This 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 analysis confirms the close association of the rise in top income shares and the rise in household debt held as a financial asset by top income earners. The magnitude is substantial. Applying the coefficient estimate from the state analysis to the aggregate implies that the rise of top income shares explains 75% of the rise in household debt held as a financial asset by the household sector in the United States. Furthermore, the ability of the top income shares in a state to predict the state's rise in overall household debt held as a financial asset comes uniquely from its ability to predict the rise in household debt held as a financial asset by top income earners. There is no correlation between the rise in the top income shares in a state and household debt held as a financial asset by non-top income earners.

After the Great Recession, the saving glut of the rich continued to accumulate, but the net saving position of the bottom 90% increased. So where did the saving glut settle? The answer lies in government borrowing. Both government deficits and government debt rose substantially after the recession. Furthermore, the government substantially increased their lending to the household sector. Such lending occurred through Federal Reserve purchases of Agency mortgage-backed securities and Federal Government issuance of student loans. After the Great Recession, the government played a more important role in intermediating the savings of the rich and the borrowing of the non-rich.

Overall, a central implication of the findings presented here is that a single factor–a rise in top income shares–could potentially explain two common patterns witnessed in many advanced economies since the 1980s: a substantial decline in interest rates and a large rise in household debt (e.g., Jordà et al. (2016)). A companion study (Mian et al. (2019)) incorporates non-homothetic preferences over saving into an otherwise standard deterministic two-agent macroeconomic model, and it finds that a rise in income inequality generates more saving by the wealthy, more borrowing by the non-wealthy, and a decline in interest rates. The patterns shown here are consistent with these predictions.

Another implication is that aggregate measures of national saving should be treated with caution when evaluating the importance of a saving glut. For example, some have pointed to the decline in the aggregate personal saving rate as evidence against the idea that there was a saving glut generated by the rise in income inequality in the United States. The analysis done here shows that such an

argument is incorrect: a focus on the top 1% of the income distribution provides evidence in favor of the view that the rise in top income shares generated a substantial increase in saving. This saving, however, was transformed into borrowing by the bottom 90%, thereby eliminating any response of the national saving rate. In general, savings in the NIPA are by definition equal to net domestic investment and net savings overseas: any saving and borrowing within the household sector is masked in aggregate saving rates.[1]

Finally, the findings give a different perspective on the growth in the size of the financial sector in the United States since the 1980s (e.g., Philippon (2015)). Traditional models view the primary role of the financial sector as channeling the savings of the household sector into investment by the business sector. Since the 1980s, business investment and productivity growth have been constant or even declining, while the financial sector and household debt have grown considerably. The growth in the financial sector since the 1980s appears to be driven to a large degree by the channeling of savings by some households into borrowing by other households.

A note on causality. The rise in top income shares in the United States and world-wide is well-documented (e.g., Katz and Murphy (1992), Piketty and Saez (2003), Autor et al. (2008), Atkinson et al. (2011), Piketty et al. (2017), CBO (2019), and Smith et al. (2019a)). 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. (2019a)). For these reasons, we treat the rise in inequality that began in the early 1980s as the initial shift in the economy, and we speak of other aggregates as responding to this shift. The state-level analysis supports this interpretation. However, we acknowledge up front that there is no specific source of exogenous variation in the rise in inequality used in this study.

Related literature. Several studies have detailed the evolution of wealth inequality in the United States (e.g., Saez and Zucman (2016), Wolff (2017), Bricker et al. (2018), Batty et al. (2019), Kuhn et al. (2019), and Smith et al. (2019b)). This study is the first to our knowledge to focus on the holdings of household debt as a financial asset by the wealthy. The argument that a rise in inequality generates important dynamics for wealth and interest rates is developed in Straub (2019), who also emphasizes how aggregate saving rates can be misleading when discussing how rising income inequality affects wealth accumulation.[2]

There is also a growing literature focusing on the rise in household debt in the United States.

---

[1]A similar point is made by Pettis (2017).
[2]Related arguments are made by Kaymak and Poschke (2016) and Auclert and Rognlie (2017).

Most of this literature is focused on trends immediately before the Great Recession.[3] One exception is the recent working paper of Bartscher et al. (2019), which examines the rise in household debt since 1949 across the income distribution. Many of the results in Bartscher et al. (2019) are complementary to the analysis here. For example, from 1983 to 2016, Bartscher et al. (2019) 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. (2019) do not attempt to link the saving of the rich to the borrowing of the non-rich, which is the main focus of this study.

The findings presented here are also related to the literature exploring consumption, income, and wealth inequality more generally (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. (2017), Fisher et al. (2016), Guvenen et al. (2019), and De Nardi et al. (2018)). As shown below, this literature is an important input into the measurement of the consumption share of the top 1% over time.

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% contributed to higher debt levels of the bottom 90%, which helps explain why interest rates fell during 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 dis-saving of the non-rich, as is done here. The state-level analysis linking top income shares to increased holdings of household debt is novel to the literature.

The findings here are related to the secular stagnation literature (e.g., Summers (2014)). A number of studies in the literature suggest that rising inequality is a factor putting downward pressure on interest rates given high saving rates of the rich (e.g., Stiglitz (2016), Rachel and Smith (2017), and Rachel and Summers (2019)). The findings of this study support this argument. As mentioned above, the model in Mian et al. (2019) with non-homothetic preferences over saving has the implication that a rise in income inequality can simultaneously explain a rise in household debt and lower interest rates.

---

[3]See, for example, Mian and Sufi (2014), Bhutta and Keys (2016), Mian and Sufi (2017), Adelino et al. (2017), Foote et al. (2016); and Albanesi et al. (2017)

2 Measurement Using National Income Data

The empirical analysis seeks to measure how the rise in the income earned by the top of the U.S. income distribution since 1980 has affected saving in the U.S. economy. This section explains the methodology used to measure this concept, which we call the saving glut of the rich.

2.1 National income and national saving

The starting point of the measurement exercise is national income. 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. It is therefore the correct starting point from which to measure total savings of U.S. households.

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:

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

Let δ be consumption of fixed capital, and W be net income from abroad.[4] 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 net foreign position 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 and transfers to the household sector through the government budget $S^g = T - R - G$, 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: $S^p + S^\pi + S^g = I^n + F - \epsilon$, which illustrates that the concept of savings used here (Θ) includes

---

[4]More specifically, W comes from the Foreign Transactions Current Account (Account 5) and is defined as income and transfer receipts from the rest of the world minus income payments and transferst to the rest of the world.

both personal savings (S^p) and business savings (S^π). The term S^π is important when measuring the saving behavior of the wealthy.

A problem with the notion of saving in national accounts is that the NIPA do not recognize differences in saving across the distribution of households. This problem is clear in equation 3. If one part of the household sector saves more than before, then private saving only increases in the aggregate if the savings of that sector are invested (I^n), sent abroad (F), or borrowed by the government ($-S^g$).

A more accurate portrayal of what occurred in the United States after 1980 requires that private saving be split by income group:

$$\Theta_{top1} + \Theta_{next9} + \Theta_{bot90} = I^n + F - S^g - \epsilon \tag{4}$$

As we will show, the saving of the top 1% of the income distribution (Θ_{top1}) significantly increased from 1980 onward. Focusing only on the right hand side of equation 4 risks missing an important margin of adjustment in the economy: a rise in saving of the top 1% may be absorbed through dis-saving of the bottom 90%.

The main measurement exercise in this section is to measure Θ_i, which is the savings of income group i. There are at least two techniques to conduct such measurement. The first would be to try to measure actual savings across the distribution, for example by using household level survey data such as the Survey of Consumer Finances or the Panel Study of Income Dynamics.

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 above) is one important example. Such saving represented 4.2% of National Income from 2012 to 2015, and would be missed in a survey-based approach.

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 saving from any group will be systematically underestimated given these important sources of saving that are missed in surveys.

As a result, we take a different approach. This approach attempts to measure directly the components of Θ_i across the income distribution. More specifically, we estimate $Z - T + R - C$ for three groups: the top 1% of the income distribution, the next 9%, and the bottom 90%. The next two sections explain how this is done.

2.2 Shares of national income across income distribution

For the first three terms that define Θ_i, Z_i, T_i, and R_i, the approach here uses the Piketty et al. (2017) (PSZ) measures of the share of national income across the income distribution.[5] The benefit of the PSZ approach is that it allocates all of national income across the income distribution, not just fiscal income reported on tax filings.

The after-tax income shares ($\widehat{\alpha}_i$) start with pre-tax income shares (z_i), subtract taxes (t_i), and add back government consumption expenditures (g_i), transfers (r_i), and the share of the government surplus for each group (s_i). All lower case letters reflect the nominal amounts scaled by national income (Z).

Formally, the after tax income shares are defined as:

$$\widehat{\alpha}_i = \frac{z_i * Z - T_i + R_i + G_i + S_i}{Z} = z_i - t_i + r_i + g_i + s_i \tag{5}$$

The first three terms are straight-forward, but some confusion may arise with the terms g_i and s_i. PSZ add these terms back to the after-tax income shares to ensure that the sum of the after-tax income shares add up to national income. This follows from the government budget equation: $S^g = T - R - G$. Given this budget equation:

$$\sum_{i=1}^{N} \widehat{\alpha}_i * Z = Z - T + R + G + S^g = Z$$

However, as is clear in equation 3, we instead want to measure $z_i - t_i + r_i$. That is, we want to ignore the g_i and s_i terms in the PSZ definition of after-tax income shares in order to capture the household saving decision independent of what the government does with its own spending and borrowing policies. As a result, the final measures of after-tax income shares used in this study are:

$$\alpha_i = \frac{z_i * Z - T_i + R_i}{Z} = z_i - t_i + r_i = \widehat{\alpha}_i - g_i - s_i \tag{6}$$

Figure 1 plots the after-tax share (α_i) for the top 1%. As it shows, the after-tax income share of the top 1% of the income distribution increased by 3.2 percentage points from 1980 to 1988, by 5.6 percentage points by 2005, and by 7.5 percentage points by 2012.

---

[5]As shown in Section 3.4, the saving glut of the rich is even larger if the after-tax income shares of the CBO (CBO (2019)) are used.

Figure 1: Top 1% After-tax Shares of National Income
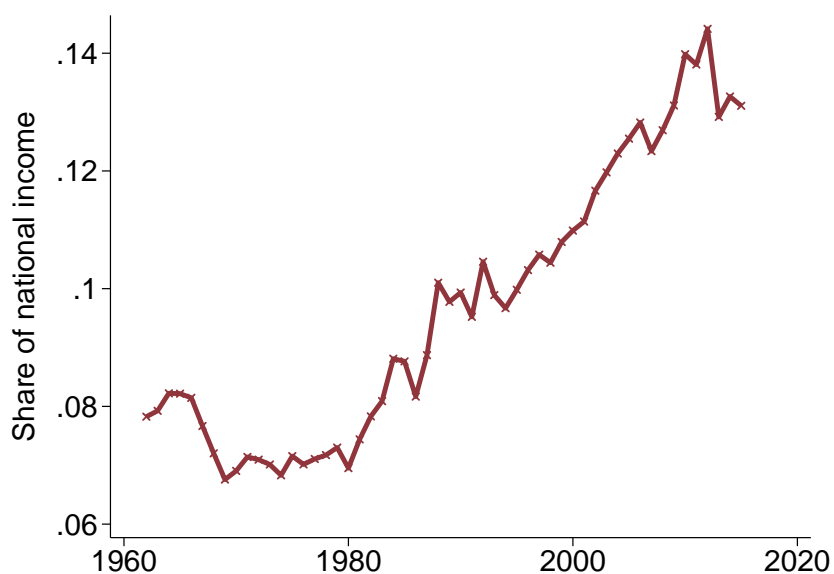
Data are from Piketty et al. (2017). This represents pre-tax income minus taxes plus transfers for the top 1%, scaled by national income.

2.3 Consumption expenditures across the income distribution

The last component needed to measure savings of each income group (Θ_i) is consumption (C_i). Measurement of the consumption of the top of the income distribution is a challenge 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 consumption share of the top 1% in a given baseline year, and an assumption of the evolution of the consumption to income ratio of the top 1% over time.

As a baseline, the consumption share of the top 1% is measured using the the Panel Study on Income Dynamics (PSID) for the years 2011 and 2013. The average consumption shares of the three main groups are calculated for these two years, and this is the consumption share used for the baseline year 2012.[6]

The consumption share of the top 1% in 2012 using this technique is 3.7%. However, there is widespread evidence that the consumption share of the top of the income distribution is under-reported in survey evidence.[7] To account for this under-reporting, in the baseline measurement

---

[6]The measure of consumption in the PSID was significantly expanded in 2005, and consumption of the top 1% is available for every other year from 2005 to 2013. The last two of these years are chosen to avoid dynamics surrounding the Great Recession (e.g., Heathcote and Perri (2018)). However, all results are similar if the average is calculated for all years and 2009 is used as the baseline year.

[7]There is a large literature discussing the potential under-reporting of consumption by the rich in various surveys, but the Consumer Expenditure Survey in particular. See for example, Aguiar and Bils (2015), Carroll et al. (2015),

of the saving glut of the rich, we assume that the consumption of the top 1% is under-reported by 50%. This yields a consumption share of the top 1% in 2012 of 5.6 percentage points. The total consumption that is added to the consumption of the top 1% to correct for under-reporting is subtracted from the next 9% and the bottom 90% proportional to after-tax shares of national income of the two groups.

The second critical input into the calculation of consumption shares is the assumption of the consumption to income ratio over time. The most extreme approach would be to assume that the consumption to income ratio of the top 1% has been constant over time, which would imply that the consumption shares and income shares have increased at the same growth rate. This is one assumption made in the analysis below. Under this assumption, the top 1% and next 9% are assumed to have a constant consumption to income ratio over time at the 2012 level, and the bottom 90% are assigned the residual consumption. The assumption of a constant consumption to income ratio is extreme in the sense that it would imply a large increase in the consumption share of the top 1% since 1980.

This assumption, in our view, is too extreme given evidence in the literature. The average post-tax real income of the top 1% implied by their share of national income was \$344 thousand in 1980 and \$1.0 million in 2014 (in 2014 dollars). In contrast, the average post-tax real income of the bottom 90% increased from \$31 thousand to \$43 thousand. Given estimates in the literature, it is implausible that the consumption to income ratio for the top 1% stayed constant given a rise in real income of a factor of three. For example, Straub (2019) finds an elasticity of consumption to changes in permanent income of 0.7. This is for *permanent income*, not the actual income used here which includes both a permanent and transitory component. Furthermore, recent evidence in Guvenen et al. (2019) and De Nardi et al. (2018) shows that high income earners are more likely to experience positive transitory shocks to their income process, bolstering the viewpoint that the consumption to income ratio almost assuredly declines with high income realizations.

Following the evidence in Straub (2019), we use an alternative assumption for the evolution of the consumption to income ratio of the top 1%:

$$\frac{C_{it}}{\bar{y}_t} = K * \left(\frac{y_{it}}{\bar{y}_t} \right)^{\beta} \tag{7}$$

where y_{it} is real post-tax income of the top 1% in year t and $\bar{y}_t$ is average real post-tax income across all groups in year t. The scaling of all variables helps ensure that average changes in income over time do not induce changes in the consumption to income ratio. The constant K is chosen so that the equation holds in the benchmark year, here $t = 0$: $K = \frac{C_{i0}}{\bar{y}_0} * \left(\frac{y_{i0}}{\bar{y}_0} \right)^{-\beta}$.

The year $t = 0$ is the benchmark year, which in this setting is 2012. The critical parameter is β,

Attanasio and Pistaferri (2016), and Meyer and Sullivan (2017).

which reflects the elasticity of consumption with respect to income, which Straub (2019) estimates to be 0.7 when the income measure is permanent income. In this specification, the consumption to income ratio is constant if $\beta = 1$, but declines in income if $\beta < 1$. The baseline specification uses the assumption that $\beta = 0.7$, but results are also shown for $\beta = 0.5$ and $\beta = 1$. When assuming $\beta = 0.7$ or $\beta = 0.5$ instead of $\beta = 1$, the extra consumption removed from the top 1% is added to the next 9% and the bottom 90% proportional to their after-tax income shares.[8]

The solid red line in Figure 2 plots the consumption shares of the top 1% of the income distribution where the 2012 level is adjusted for under-reporting and where β is assumed to be 0.7 in equation 7. This is the baseline specification used going forward in the analysis. For the sake of completeness, the figure also plots the consumption shares for specifications where the under-reporting correction is not implemented, and for specifications with $\beta = 0.5$ and $\beta = 1$.

Figure 2: Consumption Share of the Top 1% National Income Earners
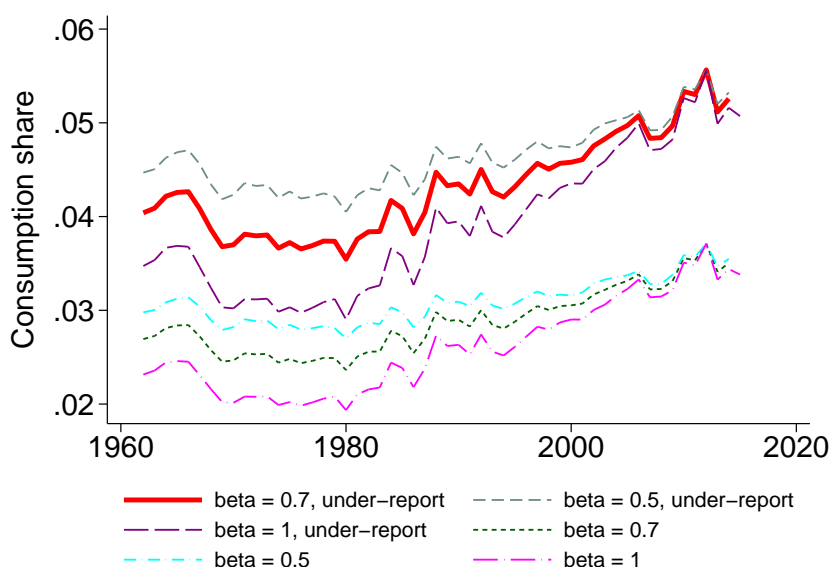
The consumption share of the top 1% is calculated using the PSID for baseline year 2012, and then the time series is generated using different assumptions on the evolution of the consumption to income ratio of the top 1%. Beta = 1 is a constant consumption to income ratio, and lower levels of beta reflect a steeper decline in the consumption to income ratio based on income. Under-report refers to the fact that the baseline 2007 consumption share is inflated by 50% to account for under-reporting of consumption in the PSID. The solid red line is the baseline specification used going forward. Please see text for more details.

Under the baseline assumption, the consumption share of the top 1% increased from 3.5% in

---

[8]The assumed consumption to income ratios overtime 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.

1980 to 5.3% in 2014. As should be expected the consumption share increases by more if the assumption of a constant consumption to income ratio is made (purple line), and it increases by less under the assumption that consumption to income ratios decline even more in income (teal line). Under the most extreme assumption of correction for under-reporting and a constant consumption to income ratio, the consumption share of the top 1% increased by 2.3 percentage points from 1980 to 2014, from 2.9% to 5.2%.

3 Saving Glut of the Rich

3.1 Magnitude

Following equation 6 above, the saving glut of the rich over time is defined as:

$$\theta_{top1,t} = \frac{(Z - T + R - C)_{top1,t}}{Z_t} = \alpha_{top1,t} - c_{top1,t}$$

This captures the total amount of savings generated by the top 1% of the income distribution, scaled by national income in order to help interpret magnitudes. Figure 3 plots the saving glut of the rich under the various assumptions on the evolution of consumption of the top 1% discussed in Section 2.3 above. As Figure 3 shows, the rise in the saving glut of the rich is large under any of the assumptions.

Table 1 presents the magnitude saving glut of the rich over time. Under the baseline assumptions, the saving glut of the rich increased by an average of 5.0 percentage points of national income annually from the 1970s to the 2010 to 2014 period. The saving glut increased almost linearly across decades, representing 1 percentage point more of national income annually during the 1980s, 3.8 percentage points more during the years 2000 to 2010, and 5.0 percentage points more during the 2010 to 2014 period. Post 2000, the top 1% saved an additional 4.3 percentage points of national income annually relative to the 1970s. To put this into perspective, the average annual net domestic investment from 2000 onward was 6.7 percentage points of national income. The annual rise in saving of the rich was 60% average net domestic investment from 2000 onward.

Table 1 also reports the saving glut of the rich under alternative assumptions regarding the evolution of the consumption to income ratio over time. For all three columns, the consumption of the top 1% as of 2012 is assumed to be under-reported by 50% in the PSID. Even under the most conservative assumption of a constant consumption to income ratio of the top 1%, the rise in the saving glut of the rich is only 0.4 percentage points of national income smaller than the baseline assumption of $\beta = 0.7$. From this point forward, the analysis focuses on the measure of the saving glut of the rich using the 50% under-reporting assumption and $\beta = 0.7$.

12

Figure 3: Saving Glut of the Rich
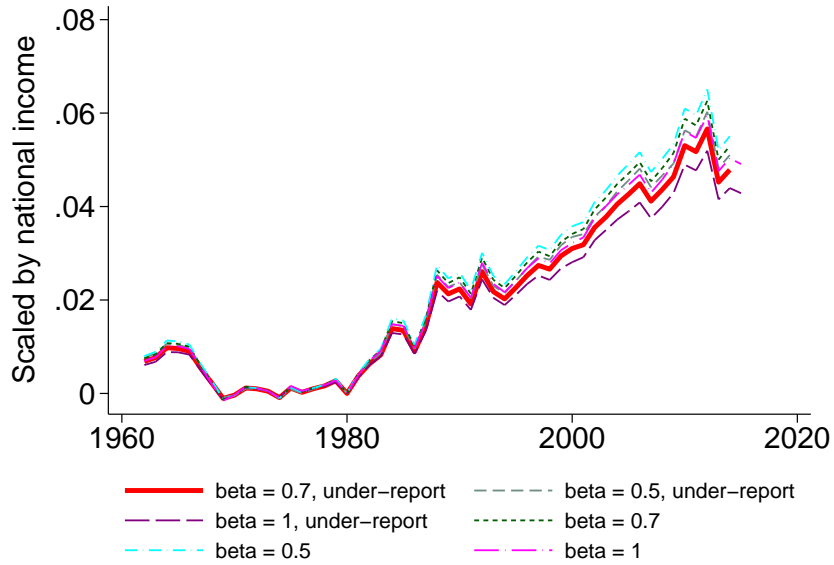
The saving glut of the rich is defined to be the after-tax income of the top 1% of the income distribution minus personal and government consumption of the top 1% of the income distribution, scaled by national income. All series are relative to 1980. Different series reflect different assumptions on the evolution of the consumption to income ratio of the top 1%. Beta = 1 is a constant consumption to income ratio, and lower levels of beta reflect a steeper decline in the consumption to income ratio based on income. Under-report refers to the fact that the baseline 2012 consumption share is inflated by 50% to account for under-reporting of consumption in the PSID. The solid red line is the baseline specification used going forward.

Table 1: Saving Glut of the Rich Over Time

Decade	Baseline	$\beta = 0.5$	$\beta = 1$
1970s	0.045	0.041	0.049
1980s	0.055	0.052	0.059
1990s	0.068	0.066	0.071
2000s	0.083	0.082	0.085
2010s	0.095	0.094	0.095

The saving glut of the rich is the after-tax income of the top 1% of the income distribution minus personal consumption of the top 1% of the income distribution, scaled by national income. The baseline estimate assumes a 50% under-reporting of consumption by those in the top 1% of the income distribution in 2012, and a consumption to income ratio that declines in income with the power of $\beta = 0.7$ (see equation 7). Estimates using $\beta = 0.5$ and $\beta = 1$ are in columns 2 and 3, respectively. Averages for the decades are shown.

3.2 Do the savings finance traditional activities?

Where did the substantial rise in savings of the rich ultimately settle? National accounting provides for a simple decomposition exercise. Re-writing equation 4 above with all terms scaled by national income, we have:

$$\theta_{top1,t} = \frac{I_t^n + F_t}{Z_t} - \theta_{next9,t} - \theta_{bot90,t} - \frac{S_t^g}{Z_t} \quad (8)$$

The saving glut of the rich could have financed net domestic investment I^n or it could have been invested in other countries (F). If neither of these happened, then a rise in the saving glut of the rich must have increased net borrowing by other households or by the government ($-S^g$).

Figure 4: Net Domestic Investment and the Foreign Account
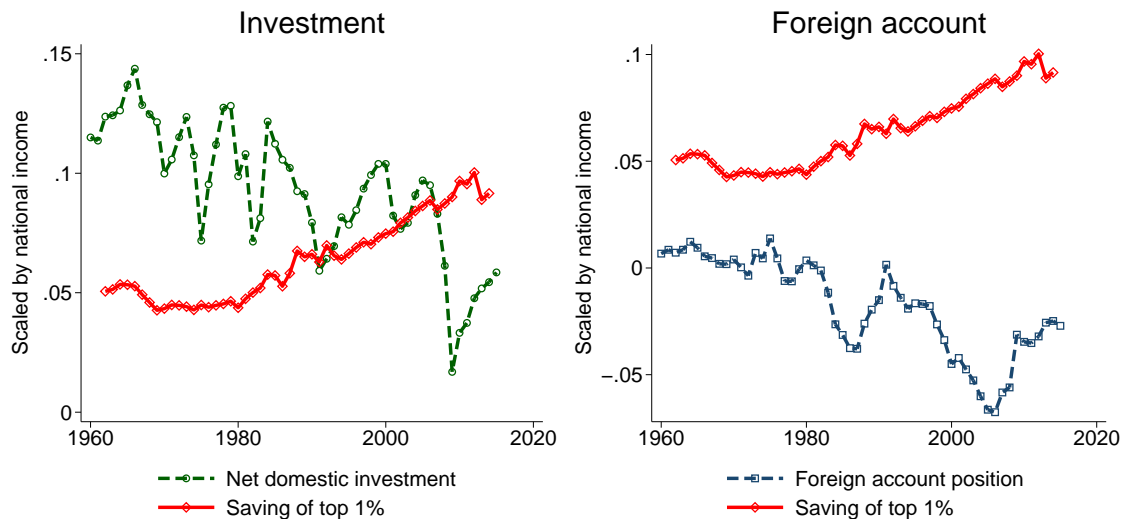
The saving of the top 1% is defined to be the after-tax income of the top 1% of the income distribution minus personal consumption of the top 1% of the income distribution, scaled by national income. Net domestic investment comes from the national accounts, and includes both government and private investment. The foreign account is net exports adjusted for net income flows based on the difference in how GDP and National Income account for net income and transfers to foreigners. All series are scaled by national income.

As the left panel of Figure 4 shows, net domestic investment moved in the opposite direction. As the saving glut of the rich increased, net domestic investment fell. In fact, the pattern is so strong that after the Great Recession, the rich were saving a higher percentage of national income than total net domestic investment. The right panel of Figure 4 shows that the foreign account position of the United States also moved in the opposite direction. As is well known, the United States borrowed

more from the rest of the world from 1980 to 2014 rather than investing in it.

This latter pattern has been called the global saving glut, first highlighted by Bernanke (2005). This is the idea that there has been an influx of foreign savings that have been transformed into borrowing by governments, firms, and households in many advanced economies. Using the foreign account position of the United States, we can directly compare the global saving glut and the saving glut of the rich.

To do so, the global saving glut in the United States is measured as the foreign account position multiplied by -1. Figure 5 plots both saving gluts. As it shows, the global saving glut and the saving glut of the rich are on the same order of magnitude. There have been periods in both the 1990s and 2010s in which the saving glut of the rich has exceeded the global saving glut.

Figure 5: Saving Glut of the Rich and the Global Saving Glut
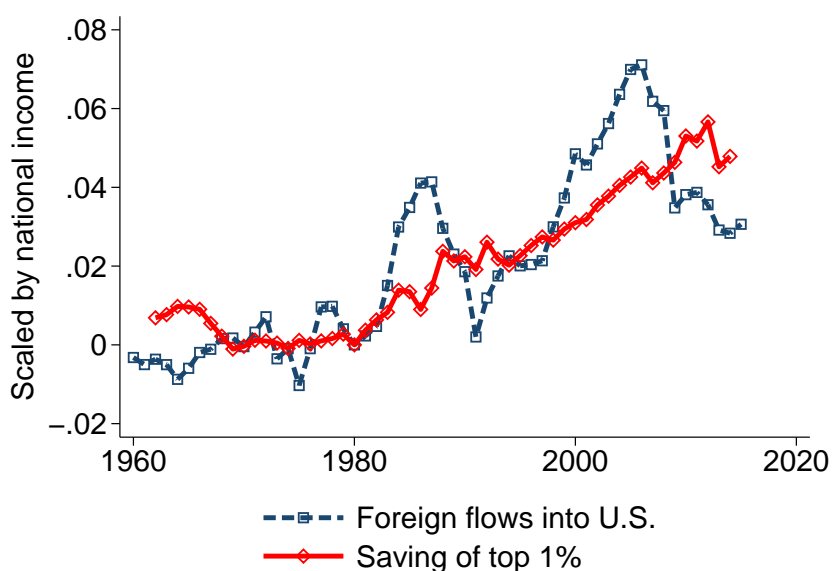
The saving of the top 1% is defined to be the after-tax income of the top 1% of the income distribution minus personal consumption of the top 1% of the income distribution, scaled by national income. Foreign flows into the U.S. is -1 multiplied by the foreign account position of the United States. All series are scaled by national income and the 1980 level is subtracted for each series, respectively.

Table 2 summarizes these results. Net investment has fallen substantially during the period in which the saving glut of the rich accelerated. Relative to the 1970s, the average annual increase in the saving glut of the rich exceeded the annual average global saving glut in the 1990s (0.068-0.045 compared to 0.017+0.002) and in the 2010 to 2014 period (0.095-0.045 compared to 0.030+0.002). During the 1980s, and the first decade of the 2000s, the saving glut of the rich was between half and two-thirds of the global saving glut.

15

Table 2: Traditional Channels of Absorption

Decade	Saving glut	Investment	Rest of world
1970s	0.045	0.109	0.002
1980s	0.055	0.099	-0.019
1990s	0.068	0.081	-0.017
2000s	0.083	0.079	-0.053
2010s	0.095	0.047	-0.030

The saving glut of the rich is the after-tax income of the top 1% of the income distribution minus personal consumption of the top 1% of the income distribution, scaled by national income. Investment is net domestic investment, rest of world is net exports adjusted to take into account the residence of the owners of capital. All series are scaled ny national income. Averages for the decades are shown.

## 3.3 Dis-saving of the bottom 90% and government

Given that investment and the foreign account position of the United States did not increase from 1980 onward, Equation 8 implies that either the government or households in the bottom 99% of the income distribution must have reduced saving significantly. The household side is shown in Figure 6. Saving of households in the 90th to 99th percentile of the income distribution (what we call the next 9%) increased slightly from 1980 onward, representing 4.2 percentage points of total national income from 1990 onward.

In contrast, there has been a large decline in the saving of the bottom 90% of the income distribution. Table 3 reports the saving of the top 1%, next 9%, and bottom 90%, all scaled by national income. From 2000 to 2010, the bottom 90% saved about 10 percentage points of total national income less than they did in the 1970s.

This decline in saving of the bottom 90% is significantly larger than the increase in saving of the top 1%. This reflects the fact that both the global saving glut and the saving glut of the rich increased substantially from 1980 onward, and net domestic investment actually fell. Both the influx of foreign capital and the rise in saving of the top 1% have been associated with a large decline in the saving of the bottom 90%.

The final column of Table 2 examines the government deficit, which is the final margin of adjustment available to absorb the saving glut of the rich. As it shows, the government deficit was steady from the 1970s to the 2000 to 2010 period, before increasing by a factor of two after the Great Recession. This was the same time period in which the net saving position of the bottom 90% increased by two percentage points of national income. The post Great Recession evidence suggests that the government stepped in to absorb the saving glut of the rich when consumption of

Figure 6: Saving Glut of the Rich and Saving of the non-Rich
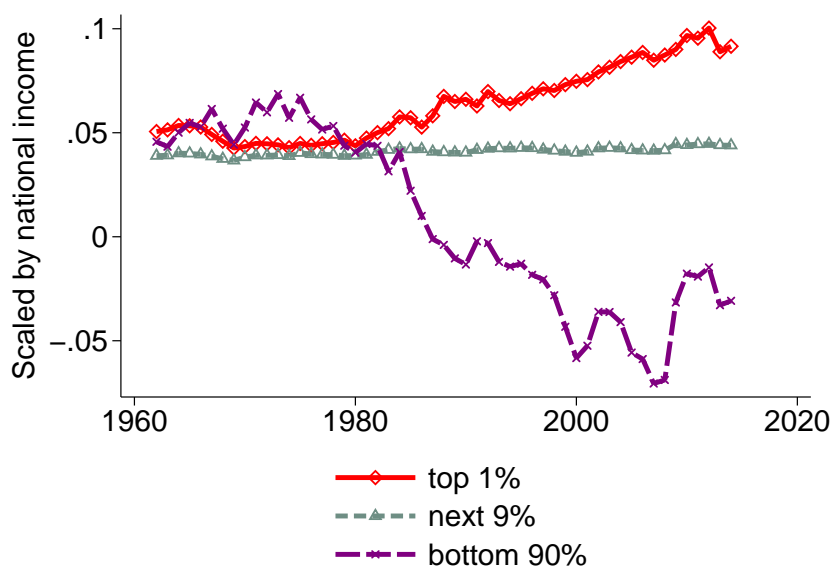
The saving of the top 1% is defined to be the after-tax income of the top 1% of the income distribution minus personal consumption of the top 1% of the income distribution, scaled by national income. Net saving of the other two groups is similarly defined.

Table 3: Absorption by the Bottom 90%

Decade	Top 1%	Next 9%	Bottom 90%	Gov Saving
1970s	0.045	0.039	0.057	-0.041
1980s	0.055	0.041	0.022	-0.051
1990s	0.068	0.042	-0.017	-0.040
2000s	0.083	0.042	-0.051	-0.046
2010s	0.095	0.044	-0.023	-0.085

The saving of the top 1% is defined to be the after-tax income of the top 1% of the income distribution minus personal consumption of the top 1% of the income distribution, scaled by national income. Net saving of the other two groups is similarly defined. Averages for the decades are shown.

the bottom 90% fell. We return to this point in Section 6 below.

Figure 7 accumulates all of these margins of absorption of the saving glut of the rich. Starting with equation 8, we re-arrange to obtain:

$$\theta_{top1,t} + \theta_{bot99,t} - \left(\frac{I^n}{Z}\right)_t - \left(\frac{F}{Z}\right)_t + \left(\frac{S^g}{Z}\right)_t + \epsilon_t = 0$$

For each of the 6 variables, we construct

$$\hat{V}_t = V_t - V_{1980}$$

with V standing in for any of the variables. Then for each variable we sum across all t to obtain

$$\bar{V} = \sum_{t=1981}^{2014} \hat{V}_t$$

where $\bar{V}$ is the accumulation of the differences relative to 1980. Therefore,

$$\bar{\theta}_{top1} + \bar{\theta}_{bot99} - \left(\frac{\bar{I}^n}{Z}\right) - \left(\frac{\bar{F}}{Z}\right) + \left(\frac{\bar{S}^g}{Z}\right) + \bar{\epsilon} = 0 \tag{9}$$

Equation 9 implies that the accumulated saving glut of the rich ($\bar{\theta}_{top1}$) must be absorbed by one of the other five terms: dis-saving of the bottom 99% ($\bar{\theta}_{bot99}$) or the government $\left(\frac{\bar{S}^g}{Z}\right)$, a rise in investment $\left(\frac{\bar{I}^n}{Z}\right)$, a rise in capital outflows $\left(\frac{\bar{F}}{Z}\right)$, or the statistical discrepancy ($\bar{\epsilon}$).

Figure 7 shows the accumulation of each of the six variables in equation 9. By construction, the bars sum to zero. The accumulated saving glut of the rich was on the same order of magnitude as national income from 1980 to 2014. The foreign account and investment move in the opposite direction as would be needed to absorb some of the saving glut of the rich, as already noted above. To maintain the accounting identity, the combined savings of both the government and the bottom 99% must have fallen substantially. Figure 7 shows that most of the decline in saving was by the bottom 99%. The accumulated dis-saving of the bottom 99% from 1980 to 2014, relative to the 1980 level, was almost twice national income. The saving glut of the rich was associated with a substantial dis-saving of non-rich households.

Figure 7: Absorption of the Accumulated Saving Glut of Rich
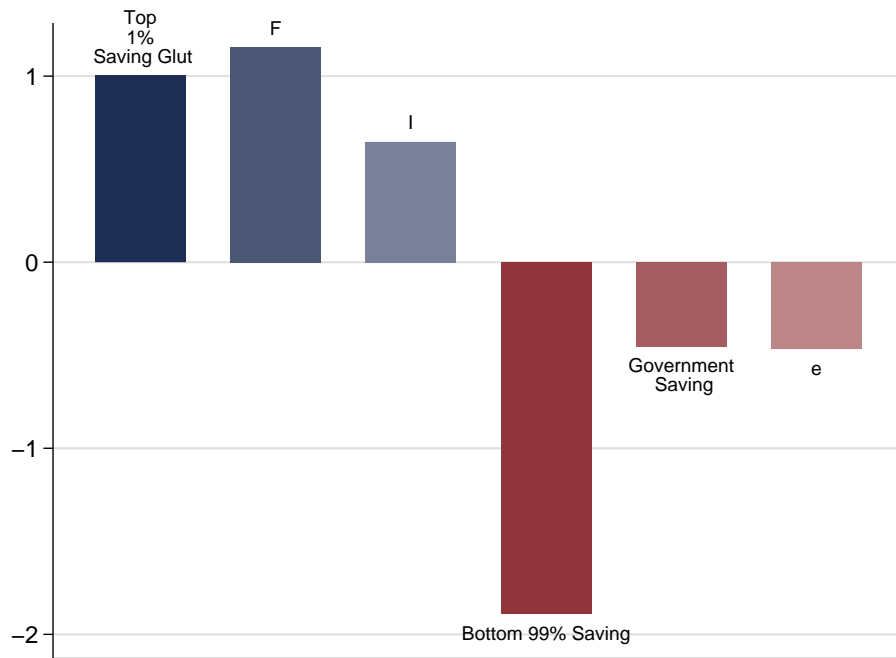
This figure presents the accumulated differences relative to 1980 levels in the equation: $\bar{\theta}_{top1} + \bar{\theta}_{bot99} - \left(\frac{\bar{I^n}}{Z}\right) - \left(\frac{\bar{F}}{Z}\right) + \left(\frac{\bar{S^g}}{Z}\right) + \bar{\epsilon} = 0$. These terms represent saving of the bottom 99% ($\bar{\theta}_{bot99}$) and the government $\left(\frac{\bar{S^g}}{Z}\right)$, investment $\left(\frac{\bar{I^n}}{Z}\right)$, capital inflows $\left(\frac{\bar{F}}{Z}\right)$, and the statistical discrepancy ($\bar{\epsilon}$).

3.4 Robustness

3.4.1 Alternative income shares

The size of the saving glut of the rich is slightly larger using the after-tax income shares reported by the Congressional Budget Office (CBO (2019)). The income shares used by the CBO do not add up to national income, which is why the Piketty et al. (2017) series is used as the baseline. However, the CBO income shares are supposed to capture $z_i - t_i + r_i$ in equation 6 above. As a result, it is straight-forward to compare the saving glut of the rich according to equation 6 for both the CBO measures of income shares and the measures from PSZ. The same consumption to national income ratio is subtracted from both, and so the difference in the measurement of the saving glut of the rich is driven exclusively by the measurement of the top 1% income share.[9]

Table 4 reports the saving glut of the rich using the baseline measure and the alternative measure based on the CBO income shares. The CBO measure leads to a saving glut of the rich during the first decade of the 2000s that is 1.5 percentage points higher than the Piketty et al. (2017) measures. In fact, for the years 1980 onward, the CBO measure leads to a saving glut of the rich that is on average 0.4 percentage points higher than the Piketty et al. (2017) measures. The use of the Piketty et al. (2017) income shares leads to a conservative estimate of the size of the saving glut of the rich relative to use of the CBO income shares.

Table 4: Robustness of the Saving Glut of Rich

Decade	Baseline	CBO income share
1970s	0.045	0.047
1980s	0.055	0.068
1990s	0.068	0.077
2000s	0.083	0.098
2010s	0.095	0.089

In column 2, the top 1% share of income is taken from the CBO (2019) instead of Piketty et al. (2017).

3.4.2 Consumption of the rich

The baseline consumption share of the top 1% of the income distribution used in the calculation of the saving glut of the rich is 5.6% in 2012. Unfortunately, few existing research studies report

---

[9]Appendix Figure A1 plots the after-tax income shares of the top 1% from PSZ and CBO and the saving glut of the rich under both measures for the full time series.

an estimate of the consumption share of the top 1% of the income distribution, which makes comparisons difficult. However, our reading of the literature is that there is no systematic evidence that this top 1% consumption share is too low. In fact, our reading of the literature suggests that it may be on the high end.

The PSID data set used by Heathcote and Perri (2018) yields a similar estimate of the consumption share of the top 1% before the under-reporting correction of 50%. This gives us confidence that the baseline PSID consumption share of 3.7% (before the under-reporting correction) used here is robust.

Since 2014, the Consumer Expenditure Survey has reported the consumption share of the top 10% of the income distribution, which they report in 2014 as 23.5%. Using this estimate, the top 1% consumption share of 5.6% implies that the top 1% consumes 3 times the average consumption of the next 9%. Fisher et al. (2016) report a consumption share of the top 5% of the income distribution using the Survey of Consumer Finances of 17.1% in 2013. Using this estimate, the top 1% consumption share of 5.6% implies that the top 1% consumes twice as much as the average consumption of the next 4%.

The baseline consumption share of the top 1% implies a ratio of personal consumption to income of the top 1% of 0.30. This ratio, however, is not comparable to standard measures of the consumption to income ratio used in the consumption and saving literature because the denominator is a different income concept. More specifically, the income concept used in the existing literature does not include important sources of national income that should be included in income when trying to estimate the contribution of the top 1% to national saving.

Two adjustments to the denominator must be made to do the comparison with the existing literature. First, the claim on business savings (or undistributed corporate profits) must be excluded. To do so, we exclude the top 1% claim on business profits, where the claim of the top 1% is assumed to be proportional to the pre-tax share of equity income of the top 1%. Second, the denominator must be adjusted for the 21 percentage point gap between survey based measures of personal income used in the literature and the NIPA measure of personal income (Heathcote et al. (2010)). For this last adjustment, the total amount of missing income in the aggregate is calculated as 21% of personal income, and the top 1% claim on this missing income is proportional to the pre-tax share of pension income of the top 1%. This reflects the fact that a large fraction of the missing 21% is related to pension income.

With these adjustments to the denominator, the consumption to income ratio of the top 1% in 2012 was 0.51, implying a saving rate of 0.49. Dynan et al. (2004) report a saving rate of 0.51 for the top 1% when using data from 1984 to 1989 (see their Table 3, column 2).

4 Financing the Rise in Household Debt

The mechanism proposed in this study is that the rise in top income shares generated a saving glut of the rich, which then translated into significant dis-saving by the bottom 90% of the income distribution. The rest of this study explores whether the saving glut of the rich led to substantial accumulation of household debt as a financial asset in portfolios of wealthy individuals. This would be the case if the rise in top income shares reflected a strong permanent income component, if inter-generational mobility across the income distribution did not substantially increase, and if a large amount of savings of the rich were invested in instruments that ultimately finance household debt. The first two of these factors has strong support in existing research.[10] The focus here is on the third factor.[11]

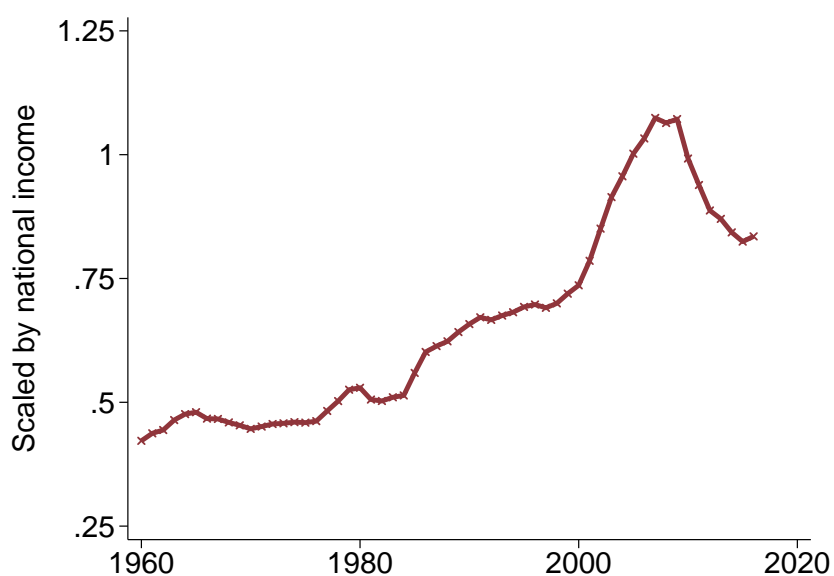## 4.1 Who holds household debt as a financial asset?

Figure 8 shows the rise in household debt to national income in the United States. From 1980 to 2007, the household debt to national income rose by 54 percentage points. Who ultimately financed the rise in household debt? The Financial Accounts of the United States allow for a detailed decomposition to answer this question. The process described in this sub-section is best thought of as an attempt to remove the veil of financial intermediation. Financial intermediaries are the immediate holders of household debt as a financial asset, but who holds the financial securities of the financial intermediaries? This question can be answered given the extensive information in the Financial Accounts on holdings of securities by different groups.

The methodology consists of a series of rounds of unveiling, and in each round a proportion-ality ratio is assumed to allocate the household debt to the remaining groups. 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.

---

[10]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. (2017). Using different data sets, Lee and Solon (2009) and Chetty et al. (2014) show that the large rise in inequality has not been associated with a decline in intergenerational mobility in the income distribution. We explore the consequences of a change in intergenerational mobility for asset accumulation of the wealthy in our companion study (Mian et al. (2019)).

[11]If all financing were perfectly fungible, then it would not be important to show that the saving glut of the rich directly financed the rise in household debt. However, recent research demonstrates that financial assets should not be viewed as perfect substitutes (e.g., Koijen and Yogo (2019)). This imperfect substitution in the demand for assets implies that holdings of specific asset classes by wealthy Americans could matter for aggregate financing patterns.

Figure 8: Household Debt to National Income Ratio

Data are from the Financial Accounts

As of 2014, the GSEs held \$5.7 trillion in household debt as an asset, which was 77.8% of all assets held by the GSEs. Mutual funds owned \$670 billion in GSE securities as assets. The assumption made in the methodology is that mutual funds therefore hold (0.778*\$670 billion=) \$521 billion of household debt through their holdings of GSE securities. 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.

The goal of the process of removing the veil of intermediation is to assign all household debt in the U.S. economy as a financial asset held by one of three groups: the government, the rest of the world, and the U.S. household sector. That is, the goal is to be able to document who among these three groups ultimately holds U.S. household debt as a financial asset.

This process of removing the veil of intermediation requires detailed explanation, and this explanation is relegated to the appendix. Furthermore, all Stata code and data will be made available for researchers to examine how this process is done. This sub-section contains an overview of the process, seeking to highlight the most important aspects.

There are in total six rounds of unveiling. The first round of unveiling assigns household debt to those that one could call immediate holders. These are primarily GSEs, the issuers of asset-backed securities, and private depository institutions. The second round removes the veil for the GSEs, issuers of asset-backed securites, financial companies, and real estate investment trusts (REITS),

which are almost pure pass-throughs. These intermediaries are primarily funded with bonds and commercial paper that are directly backed with the household debt they hold as an asset.

The critical assumption is that the household debt held in these intermediaries is distributed proportionally to other sectors of the economy according to their holdings of bonds and commercial paper issued by these institutions (which is detailed in Tables L.208, L.209, L.211 and L.213 of the Financial Accounts). This step also requires a process of netting out claims, as issuers of ABS securities for example also own some agency GSE mortgage securities.

The third round of unveiling focuses on money market funds and mutual funds, which is one of the most straight-forward steps. These funds hold substantial amounts of household debt through their holdings of the securities issued by pass-throughs (in particular the GSEs). Tables L.206 and L.224 of the Financial Accounts allow for a straight-forward assignment of the household debt held by these funds to other sectors of the economy.

The fourth round of intermediation tackles private depository institutions. As in the second round, there is a netting out procedure that must happen in this step which primarily involves bank holding companies. There are ultimately six liabilities issued by private depository institutions that allow for the unveiling of their holdings: time deposits, checkable deposits, bonds, commercial paper, holding company debt, and equity. The Financial Accounts does not include an estimate of equity for depository institutions, and so this data item is pulled from CRSP.

The fifth and sixth round of intermediation focuses on non-financial corporations and non-financial non-corporate businesses. These entities hold a significant amount of household debt, primarily through their holdings of deposits at financial institutions. The ownership of these institutions is assigned based on Tables L.213, L.214, L.216, L.223, and L.229 in the Financial Accounts.

The outcome of this process yields household debt held by households, the government, the rest of the world, and a residual category that we are unable to assign. The household debt held by households can be broken down into eleven asset categories, which are listed here in the order of largest to smallest based on amount as of 2014: time deposits, pensions (both private and public), business equity (including equity of non-financial corporations, non-financial non-corporate businesses, and private depository institutions), life insurance annuities, mutual funds, bonds (including GSE bonds), life insurance reserves, checkable deposits, money market funds, household claims on property and casualty insurers, and direct holdings of mortgages.

4.2 Household holdings of household debt

Figure 9 shows the ultimate holders of household debt owed by U.S. households. The four lines in the left panel add up to total household debt in the United States. As the left panel shows, the lion's share of household debt is held as a financial asset by U.S. households. This is true even post 2000,

when a rising share of total household debt has been held by the rest of the world. Most household borrowing in the United States is ultimately provided by the household sector.

Figure 9: Who Holds Household Debt as a Financial Asset?
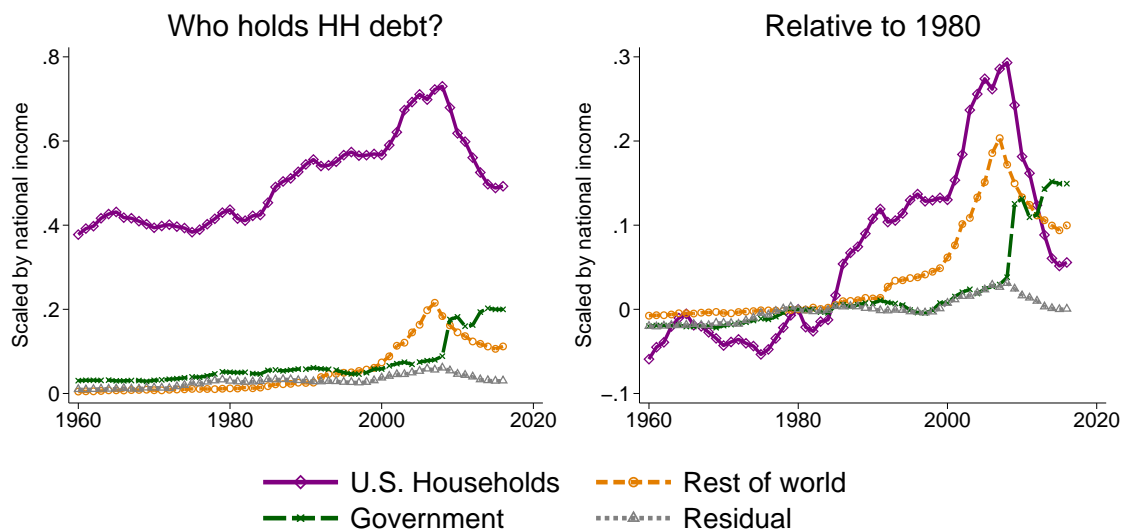
This is the result of the unveiling exercise described in detail in Section 4.1. All series are scaled by national income.

Furthermore, as the right panel shows, the rise in household debt has been primarily financed by U.S. households. This is particularly the case from 1980 to the late 1990s. From 1980 to 1996, the rise in the overall household debt to national income ratio was 16.8 percentage points. The U.S. household sector accounted for 13.6 percentage points of this increase. After 1996, the rest of the world began to provide substantial amounts of funding for U.S. household borrowing, primarily through purchases of agency GSE securities. However, even post 1996, a substantial fraction of the increase in U.S. household borrowing was financed by U.S. households.[12]

4.3 Household debt holdings of the wealthy

Once the holdings of household debt as a financial asset are allocated to U.S. households, it is straight-forward to ask the question of who in the wealth distribution holds the household debt. There are two data sources used to allocate household debt across the wealth distribution: the wealth shares calculated using the capitalization of IRS income flows in Saez and Zucman (2016), and the wealth shares calculated by the Distributional Financial Accounts conducted by the Federal Reserve (Batty et al. (2019)).

---

[12]Appendix Figure A2 shows the precise asset classes through which households held household debt as a financial asset over time.

More specifically, for asset class c, both SZ and the DFA produce a set of shares $\omega_{c,p,t}$, which are the shares of asset class c held by group p in year t. Here, the focus is on the top 10% and the bottom 90% of the wealth distribution.[13] The unveiling exercise described in the above sub-section produces the total amount of household debt held as a financial asset in asset class c. Given these inputs, the allocation of the household debt held as a financial asset across the wealth distribution is given by:

$$A_{p,t}^{HHD} = \sum_c \omega_{c,p,t} * A_{c,t}^{HHD}$$

where $\omega_{c,p,t}$ come from SZ and DFA and $A_{c,t}^{HHD}$ is the total amount of household debt held by the household sector through asset class c. By construction, the total amount of household debt held as a financial asset across the groups p add up to the total amount of household debt held as a financial asset by the U.S. household sector.[14]

Figure 10 is the central result of this exercise. It is the same as Figure 9 above except that the holdings of household debt by the U.S. household sector are split between the top 10% of the wealth distribution and the bottom 90%. The left panel uses the Saez and Zucman (2016) wealth shares to allocate household debt, whereas the right panel uses the Distributional Financial Accounts of the Federal Reserve. Qualitatively, the results are similar when using either the SZ or DFA wealth shares.

Figure 10 makes two important points. First, a substantial amount of the borrowing by the U.S. household sector is financed with saving from the top 10% of the wealth distribution. This was true even before the rise in household debt began in 1980. As of 1980, using the SZ wealth shares, 53% of total household debt was financed by the top 10% of the wealth distribution.

Second, a substantial fraction of the rise in household debt from 1980 to 2007 was financed by the top 10% of the wealth distribution. Households in the top 10% of the wealth distribution increased their holdings of U.S. household debt by 27 percentage points of national income from 1980 to 2007, which is almost half of the overall rise in household debt during this period. Prior

---

[13]The primary focus is on the top 10% of the wealth distribution instead of the top 1%. This is for two reasons. First, there is substantial disagreement on the wealth share of the top 1% versus the next 9% in the literature, but there is more agreement on the top 10%. To ensure that none of the results are sensitive to this distinction, the top 10% is used. Second, individuals who at some point earn enough to enter the top 1% of the income distribution may not end up in the top 1% of the wealth distribution. For example, recent evidence in Guvenen et al. (2019) and De Nardi et al. (2018) shows that high income earners are more likely to experience positive transitory shocks to their income process. If such individuals save a large fraction of their transitory positive shock, then such individuals may hold a substantial amount of household debt as a financial asset but not be in the top 1% of the wealth distribution. The state-level analysis in Section 5 more directly links top income earners to asset accumulation.

[14]The appendix describes in more detail how the asset class c from the Financial Accounts is mapped to the asset classes on which $\omega_{c,p,t}$ from SZ and DFA are based. For the DFA, this mapping is excellent given that the DFA is based on the same asset classes as in the Financial Accounts.

Figure 10: Who Holds Household Debt across the Wealth Distribution?
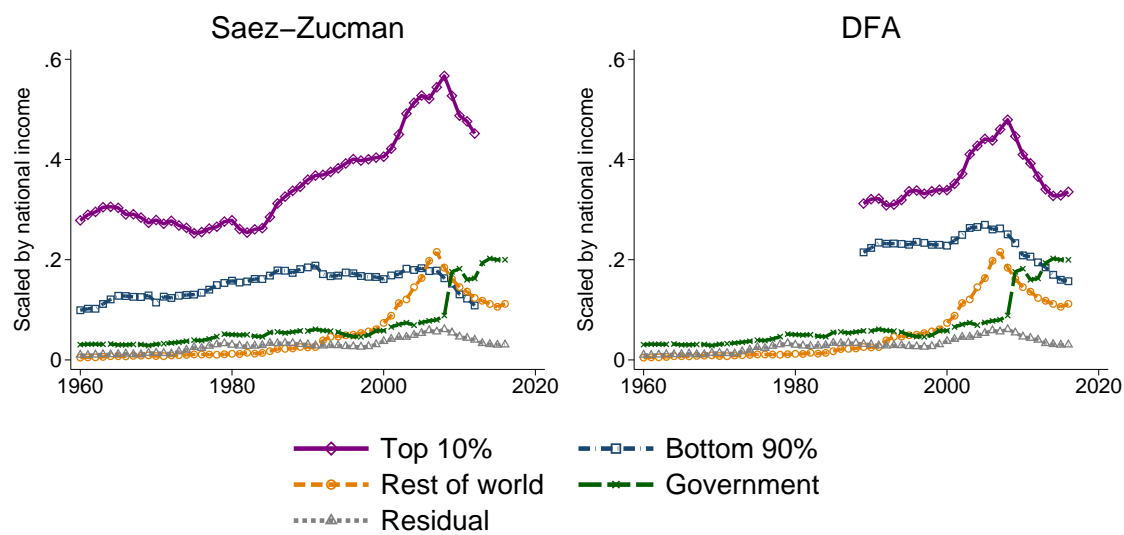
This figure decomposes the holdings of household debt by the U.S. household sector across the wealth distribution. The left panel uses the financial asset shares from Saez and Zucman (2016) and the right panel uses the financial asset shares from the DFA (described in Batty et al. (2019)). All series are scaled by national income.

to 1996, the fraction is significantly higher. The total rise in the household debt to national income ratio from 1980 to 1996 was 17 percentage points, and the top 10% of the wealth distribution accounted for 12 percentage points of the total increase. The saving glut of the rich was financing the rise in household debt in the United States before the global saving glut was.

The results utilizing the DFA and SZ wealth shares are qualitatively similar, but the overall rise in the holdings of the top 10% using the SZ wealth shares is larger when focusing on the post-1989 for which shares are available from both data sets. The results using the DFA shares imply an increase in household debt holdings of the top 10% of the wealth distribution of 15 percentage points of national income from 1989 to 2007. The SZ shares imply an increase of 20 percentage points over the same time period.

The fact that households at the top of the wealth distribution hold a substantial amount of household debt as a financial asset suggests that a more appropriate calculation of intra-household sector borrowing should examine the net debt position of each wealth group. That is, the gross household debt positions of different groups within the household sector may give a misleading impression of the net borrowing and lending position of these groups. The net debt position of any group within the household sector is the gross amount of household debt owed minus the amount of household debt held as a financial asset.

Figure 11: Net Household Debt across Wealth Distribution Relative to 1980
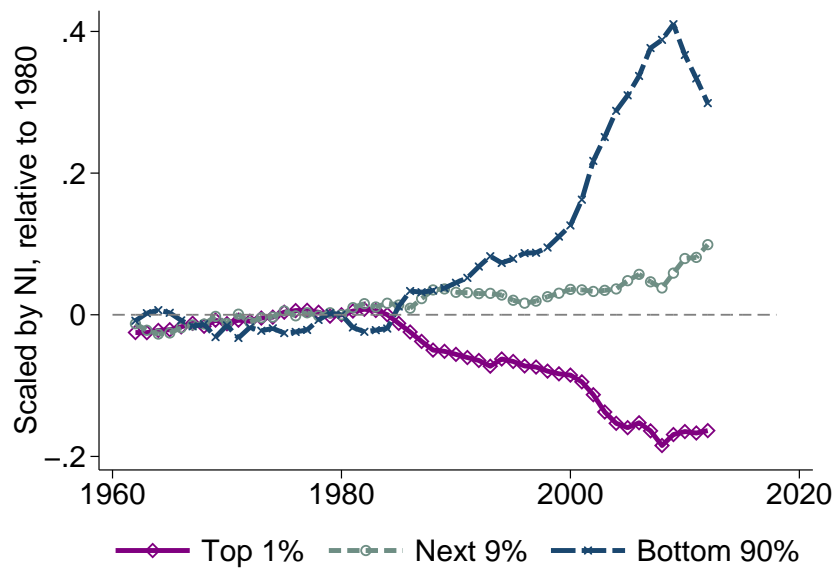
This figure shows net household borrowing by the U.S. household sector across the wealth distribution. It uses the debt and asset shares from Saez and Zucman (2016) to construct net household borrowing. Net household borrowing is defined as gross household borrowing minus household debt held as a financial asset. All series are scaled by national income, and the 1980 level is subtracted.

Figure 11 shows the net amount of debt owed by the top 1%, next 9%, and bottom 90% using the SZ wealth shares, and subtracting the 1980 level. The net debt positions make clear that the top 1% of the wealth distribution helped finance the rise in household debt for the bottom 90% of the distribution. Their net debt position fell significantly since 1980. The next 9% was somewhere in the middle. They increased borrowing after 1980 but not nearly to the same degree as the bottom 90%. The net debt position of the bottom 90% increased substantially. Research by Bartscher et al. (2019) shows that the gross household debt to income ratio increased substantially for the bottom 90% from 1980 onward. The net household debt position of the bottom 90% also increased substantially during this time period.

It may be surprising that the top 10% held such a large share of household debt, given that the popular view is that the portfolios of the top 10% are focused primarily on corporate and non-corporate business equity. Appendix Figure A3 uses the DFA allocation shares and shows the five largest asset classes through which the top 10% held household debt as of 2007, and the five largest classes in terms of the rise in debt holdings of the top 10% from 1992 to 2007. As of 2007, the top 10% held a substantial amount of household debt through time deposits at private depository institutions (17 percentage points of national income). Time deposits were also the main source through which the top 10% increased their holdings of household debt from 1992 to 2007. As of 2007, the top 10% of the wealth distribution held 61 percent of all time deposits held by U.S. households, up from 52 percent in 1992.

Equity was the second most common instrument through which the top 10% held household debt. Part of this reflected equity ownership of private depository institutions, but an even larger fraction comes from the fact that non-corporate and corporate businesses had large balances in deposits and money market funds. From 1992 to 2016, non-financial businesses increased their holdings of deposits and money market funds by 15 percentage points of National Income (see Appendix Figure A4). The rich, as the primary equity-holders in businesses, held substantial amounts of household debt through the business deposit channel.

The results in this section show that a substantial amount of the rise in household debt was held as a financial asset by the top 10% of the wealth distribution. This is consistent with the view that a large fraction of the rise in household debt from 1980 to 2007 was financed by the saving glut of the rich. The appendix presents several additional results and robustness tests for this basic finding.[15]

---

[15]The appendix shows household debt held as a financial asset by the top 1% and next 9% (Appendix Figure A5). There is also a detailed discussion of the advantages and disadvantages of the SZ and DFA approaches. Finally, it shows that the core findings of this section are robust to the critique of the capitalization approach to obtaining wealth estimates that is made by Bricker et al. (2018) and Smith et al. (2019b) (Appendix Figures A6 and A7).

5 State-level Analysis

At the national level, the rise in savings generated by the top 1% was closely linked to the rise in household debt held as a financial asset by wealthy Americans. State-level analysis can help more directly tie the rise in top income shares to the rise in household debt held as a financial asset by the rich. There was large variation at the state-level in the rise in top income shares from 1982 to 2008, as shown in Figure 12. 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.

Figure 12: Change in Top 1% Share of Income Across States
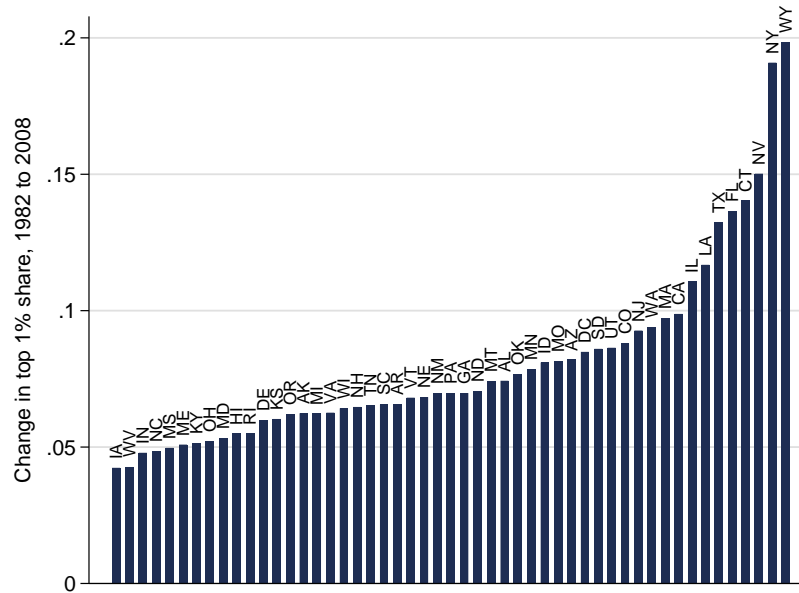
Data are from the World Inequality Database.

The state-level analysis in this section tests whether states with a higher increase in top income shares experienced a greater increase in the holdings of household debt as a financial asset. And if they did, it tests whether such acccumulation of household debt as a financial asset was primarily driven by those at the top of the income distribution.[16]

The advantage of the state-level analysis is that it brings us closer to an ideal experiment in which, all else equal, some economies experience a larger increase in top income shares than others, and we can track whether there is more asset accumulation in the economies with the larger increase in top income shares. This helps address concerns with the aggregate analysis that there were

---

[16]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. The associated household borrowing can of course happen anywhere across the United States, and it does not need to take place in the specific state that is experiencing the rise in top income shares.

important changes from 1980 to 2007 other than the rise in top income shares that could have independently led to wealthy Americans holding more household debt as a financial asset. These include demographic changes, the decline in aggregate investment, the rise in the current account deficit, the rise in government deficits, and valuation effects coming from the decline in interest rates. The first-difference specification at the state-level differences out these other national trends, and therefore provides a more compelling estimate of the effect of the rise in top income shares on the rise in household debt holdings of the rich. The first-difference specification also ties the rise in top income shares more directly to the rise in wealth accumulation of top income earners.

5.1   State-level data

The key dependent variable in the analysis requires an estimate of total household debt held as a financial asset by income group i in a given state s in a given year t, which we label A_{ist}^{HHD}. To construct this measure, the capitalization technique of Saez and Zucman (2016) is used. However, unlike the approaches in the existing literature which calculate wealth shares of each income group at the national level, a novel data set allows us to construct measures of asset ownership at the state-income group level.

Two sources of tax data are used in the capitalization process. First, state-level tax record data from 1979 through 2016 are collected from the Statistics of Income Division (SOI) of the Internal Revenue Service. For the years 1989 through 2016, these tables contain the total number of returns, AGI, interest income, dividend income, capital gains, and taxable pension income at the state level, as well as broken down further by income brackets (e.g. earners with AGI above 200K). For years 1979 through 1988, these tables only have information on total number of returns and AGI at the state level, and for years 1979 through 1981 these two variables are further divided by income brackets.[17]

Second, we use the yearly public-use individual level tax return files available at the NBER. These files have the state identifier for the tax-filer for years 1979 through 2008 whenever the tax filer has AGI less than \$200K. For tax-filers with AGI above \$200K in any year, the state identifier is missing. However, the year 1982 is an exception when the state identifier is available even for filers with AGI above \$200K.

While individual-level tax information for filers with AGI above \$200K is missing at the state level, we can fill in this information for earners above \$200K as a group at the state level using the state-level SOI files described above. In particular, for years 1989 through 2008, we can pull in the average AGI, interest income, dividend income, capital gains, and taxable pension income for filers in a state with AGI above \$200K. We then give this grouped tax data the appropriate population

---

[17]For 1983 to 1989, dividend and interest income are also reported at the state level.

weight.

For years 1989 through 2008, we still require other data not found in the state-level SOI files, which we can obtain from the individual-level files. To do this, we adopt the following statistical procedure to assign tax filers with AGI above $200K at the national level to specific states. We assume that income for tax filers with AGI above $200K follows a Pareto distribution. Then, using information from 1989 onwards when we directly observe mean AGI for filers above $200K in the state-level SOI data, we estimate the parameters of the distribution and impute the relative likelihood of a tax filer with AGI above $200K residing in a particular state. We then assign filers with AGI above 200K in a state an average AGI, interest income, dividend income, capital gains and taxable pension income based on a sampling using this relative probability. The complete details, including the relative importance of this procedure in the final allocation of household debt ownership, are described in Appendix A.5.

Because we have to combine all tax-filers with AGI above $200K into a single group for each state every year, it limits how narrow a top-income group we can consistently track over time. This can be seen in Figure 13 that plots the share of tax filers with AGI above $200K in each state by year. The shaded area illustrates the cross-state variation in the share of filers with AGI above $200K each year. Obviously the share of filers with AGI above $200K increases over time, but this number stays below 6% for all states until 2008. We therefore pick the top 6% as our main "top income" category as we can track it consistently over a long period of time.

Figure 13: Percentage of Filers with AGI Above $200,000
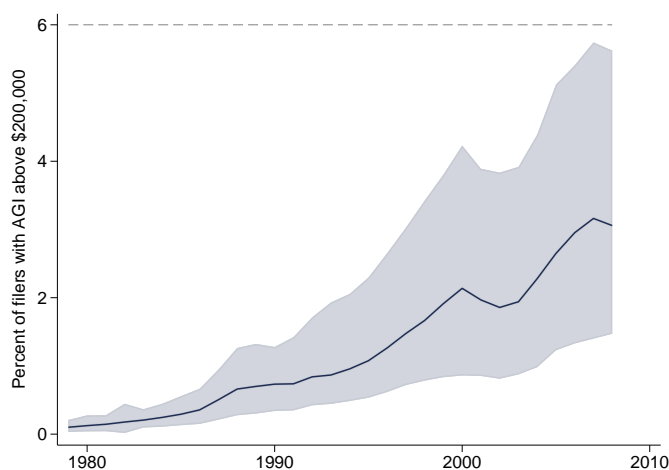
The solid line shows the percentage of tax filers in the U.S. with AGI above $200,000 over time. The shaded area represents the interval that contains this percentage for all states.

We use the income capitalization method introduced by Saez and Zucman (2016) to estimate

wealth for top 6% and bottom 94% income earners separately for each state-year.[18] In particular, the capitalization procedure allows us to separately estimate fixed income assets, equity assets, business assets, pension assets, and debt liabilities across the income distribution. The estimation of the separate asset classes also allows for an estimate of $\pi_{c,i,s,t}$, which is the share of asset class c owned by income group i in state s in year t. Following the method discussed in Section 4.3, this allows for construction of:

$$A_{ist}^{HHD} = \sum_c \pi_{c,i,s,t} * A_{c,t}^{HHD}$$

which is household debt owned as a financial asset by income group i in state s in year t. As in Section 4.3, $A_{c,t}^{HHD}$ comes from the unveiling of household debt holdings in the Financial Accounts described in Section 4.1.

5.2 Top income shares and household debt holdings of the rich

The first difference specification focuses on the 1979 through 2008 period, as it is the longest span of time over which we observe both holdings of assets and income shares. We pick 1982 as the base year because this is the year for which we have state identifiers for the entire income distribution, and we therefore do not have to impute tax information for filers with income above \$200K.[19] We use the average of 2004 to 2007 as the "post" period. Because we track the top 6% income earners over time, we use the change in top 6% income share as the main right hand side measure of a rise in inequality.

Let $Y_{ist} = \frac{A_{ist}^{HHD}}{AGI_{st}}$, which is the ratio of holdings of household debt as a financial asset by income group i in a state to the total adjusted gross income of everyone in that state. The dependent variable is the change in Y_{ist} between 1982 and 2004-2007: $\Delta_{82,07}Y_{is} = \frac{A_{is,2004-07}^{HHD}}{AGI_{s,2004-07}} - \frac{A_{is,1982}^{HHD}}{AGI_{s,1982}}$. The primary estimation equation is:

$$\Delta_{82,07}Y_{is} = \alpha + \beta_i * \Delta_{82,07}Top6Share_s + \Gamma * X_s + \varepsilon_s \tag{10}$$

where X_s are potential control variables. Given that $\Delta_{82,07}Y_s = \sum_i \Delta_{82,07}Y_{is}$, the sum of β_i across all income groups within a state reflects the total contribution of the rise in top-income shares on the rise in household debt holdings as a share of state income.

Column 1 of Table 5 presents the estimates of equation 10 for all income groups collectively. The rise in the top 1% income share has a statistically significant positive effect on the rise in

---

[18]Please see the Appendix for full details.

[19]We choose 1982 as the base year for expositional clarity only. For example, the results are similar if we use the 1979-1982 four year average as the base period.

household debt holdings in a given state. Columns 2 and 3 split the total effect into the effect coming from accumulation of household debt holdings by the top 6% versus the rest of the population. The entire increase in the holdings of household debt associated with the rise in top-income share is driven by the top 6%. This is a powerful test indicating that the rise in the income share of the top of the income distribution in particular leads to a rise in household debt accumulation of the rich.

Table 5: Effect of Change in Top Income Share on Household Debt Holdings

	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta_{82,07}$ Top 6% Share	1.929***	2.323***	-0.394	2.697***	2.367***	0.328
	(0.509)	(0.228)	(0.399)	(0.564)	(0.278)	(0.422)
Top 6% Share 1982				0.357	0.832	-0.330
				(1.130)	(0.468)	(0.762)
$\Delta_{82,07}$ Log Per Capita AGI				-0.336*	-0.0472	-0.292*
				(0.138)	(0.061)	(0.117)
Log Per Capita AGI 1982				0.0108	0.00694	0.00169
				(0.164)	(0.092)	(0.121)
Group Household Debt Holding 1982				-0.426***	-0.519**	-0.399**
				(0.113)	(0.150)	(0.118)
Group	All	Top 6	Bot. 94	All	Top 6	Bot. 94
R^2	0.20	0.66	0.02	0.55	0.73	0.45
Observations	51	51	51	51	51	51

Dependent variable, $\Delta_{82,07}Y_{is}$, is the change in household debt held as a financial asset by group i in state s scaled by state income. Robust standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

The magnitude of the saving glut of the top 6% is large. The rise in the share of the top 6% income share is 13.9% at the U.S. level. The estimated coefficient in column 2 implies that a 13.9% rise in the share of top 6% income in a state is associated with a 32.3 percentage point increase in the holdings of household debt as a share of state income. For comparison, the total rise in household debt held as a financial asset by households as a share of AGI at the U.S. level between 1982 and 2007 is 43.5 percentage points.[20]

The addition of control variables in columns 4 through 6 do not change the results qualitatively. 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.[21] The inclusion

---

[20]The state-level analysis scales debt by AGI as opposed to the scaling by national income in the nation-level analysis. The total rise in household debt held as a financial asset by U.S. households was 31 percentage points of national income from 1982 to 2007 at the national level. The total rise in household debt to national income from 1982 to 2007 was 57 percentage points. Recall that some of the rise in household debt is financed by the rest of the world and the government.

[21]Table A2 in the appendix shows how the rise in top income share is correlated with these four controls. Figure

of the last control, the initial holdings of household debt by income group i as a share of total state income, ensures that the estimates are not driven by any mechanical "valuation effects."[22]

Figure 14 summarizes the core finding graphically. The left panel shows the bivariate relationship estimated in column 2 for the top 6% of earners, and the right panel shows the same for the bottom 94%. The contrast between the two figures illustrates that the entire increase in the holdings of household debt as a financial asset as top income shares rise is driven by top earners.

Figure 14: Change in Household Debt Holding Against Rise in Top Income Share
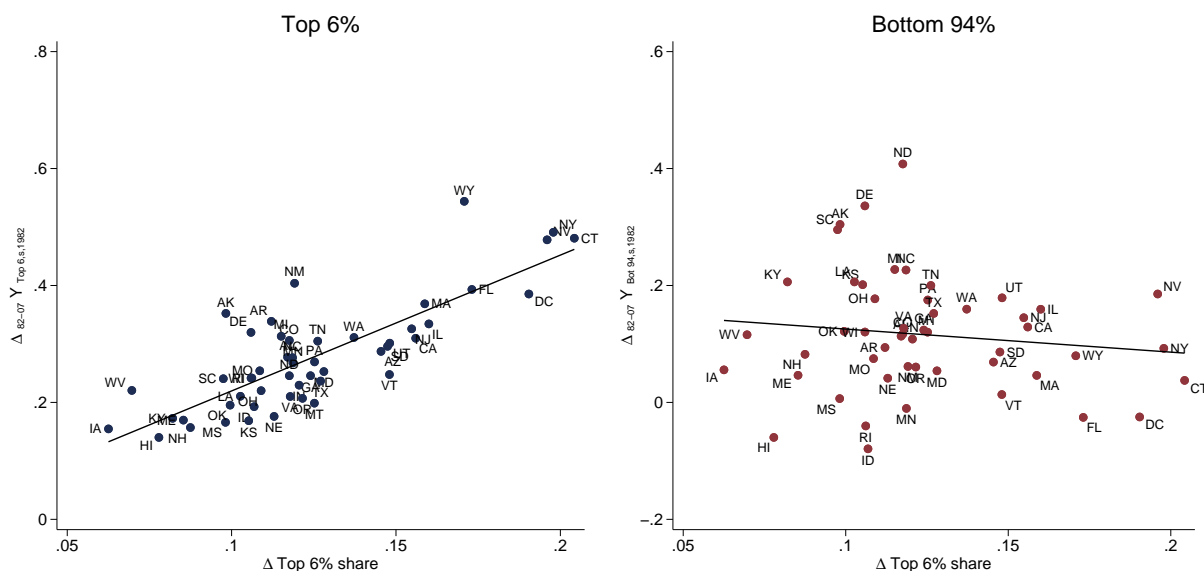
$\Delta_{82,07}Y_{is}$ is the change in household debt held as a financial asset by group i in state s scaled by state income. i is the top 6% in the left panel and the bottom 94 % in the right panel.

As discussed above, the top-income group is defined as the top 6% because we must ensure that no individual with above \$200K in AGI is outside the top income group for that state-year. However, we can narrow the top-income group further at the cost of losing the richer states for whom the top-income AGI threshold would fall above \$200K. Table A3 in the appendix progressively narrows the top-income group from the top 6% to the top 2%. Going to the top 2% leaves us with only 7 states and therefore little statistical power. However, we can reasonably say that the overall increase in

---

A8 in the appendix shows the strong relationship between rise in top 1% share of income and rise in top 6% share of income. The rise in the share of top 1% income is 11.8% at the U.S. level, which is associated with a 14.0 percentage point increase in the holdings of household debt as a share of state income.

[22]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 the saving glut of the rich. The addition of initial holdings of household debt as a control variable mitigate this concern.

holdings of household debt is concentrated even further within the top 6%, and likely comfortably within the top 3%.

Overall, the state-level analysis provides a cleaner estimate of the relationship between the rise in top income shares and the holdings of household debt as a financial asset by the rich. The analysis allows us to control for demographics and initial wealth levels, and the first-difference specification mitigates the concern that other secular trends in the U.S. economy explain the close association between these two variables at the aggregate level.

6 Post Great Recession: The Rise of Government Debt

Using the flow of income, the saving glut of the rich continued to rise after the Great Recession (see Figure 3 and Table 3). However, the net saving of the bottom 90% actually increased during the 2010 to 2014 period. Furthermore, evaluating the stock of household debt, the 27 year secular rise in the household debt to national income ratio reversed starting in 2007 (see Figure 8).

In contrast, after 2010, government saving fell substantially, with the annual government deficit doubling from 2010 to 2014 relative to the previous decade (see fourth column of Table 3). The government debt to national income ratio increased by 60 percentage points from 2007 to 2016. These facts suggest that the saving glut of the rich began financing government borrowing to a larger extent after the recession. We are investigating this pattern in ongoing work, where a similar unveiling process is done for government debt.

As government debt increased relative to household debt, the government itself increasingly financed household borrowing. It appears that the government is now acting as a major intermediary channeling the savings of rich Americans into the borrowing of non-rich Americans.

This pattern is shown in the left panel of Figure 15, which focuses exclusively on the post 2007 period, where the 2007 level is subtracted to highlight the changes over time. The pullback in the household sector's financing of household debt was substantial after the Great Recession. This pullback occurred across many asset classes, including pensions, direct holdings of bonds, money market funds, and time deposits (see Figure A2). Throughout the financial system, the institutions through which households were previously holding a large amount of household debt as a financial asset shed those assets.

The government helped make up for this shortfall. The right panel breaks out government lending to households into three sub-components: lending by the Federal government, lending by state and local governments, and lending by the Federal Reserve. As it shows, the sharp increase in government financing of household debt after the Great Recession was driven by both the Federal Reserve (through purchases of mortgage-backed securities) and the Federal government.

The fact that the Federal Reserve purchased a large amount of Agency GSE mortgage backed

Figure 15: What Happened after Great Recession
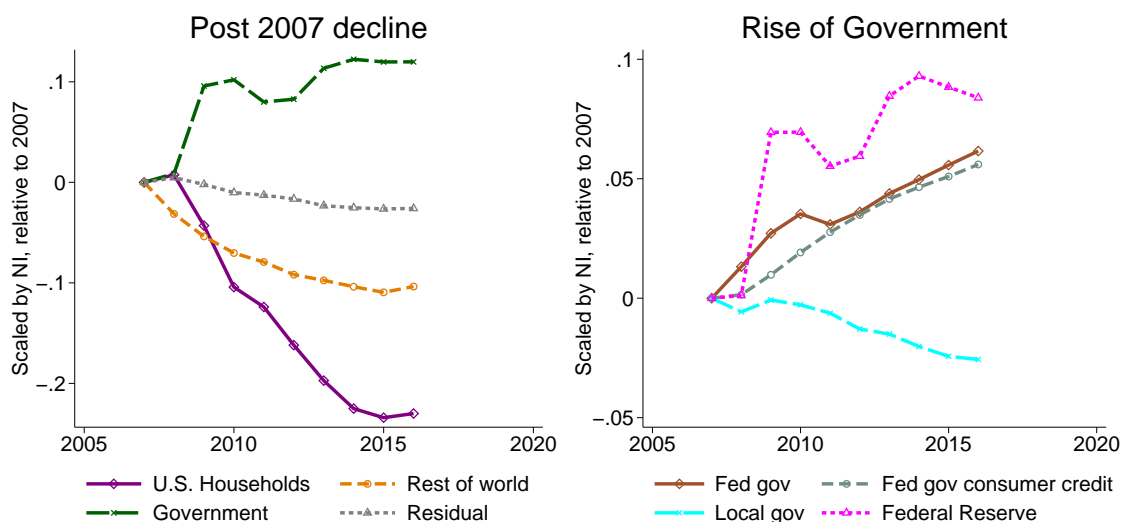
This figure explores the post 2007 decline in the holdings of household debt. All series are scaled by national income, and the 2007 level is subtracted.

securities is well known. But what drove the increase in Federal government lending? It turns out that student loans were an important driver. Student loans are included in the Financial Accounts variable that tracks consumer credit held as an asset by the Federal Government. As the right panel of Figure 15 shows, the rise in Federal government lending to households closely tracks the rise in consumer credit held as an asset by the Federal government.

This example illustrates what may be occurring in many advanced economies. High wealth households may have pulled back from financing household debt through the financial sector, and they may instead increasingly be financing transfers and lending by the government directly. This is an interesting question for future research.

7 Conclusion

The rise in top income shares in the United States since 1980 generated a saving glut of the rich. The evidence in this study suggests that much of this savings was channeled into borrowing by the non-rich. Using aggregate evidence from the NIPA and the Financial Accounts, this study shows a close association between the rise in top income shares, dis-saving of the bottom 90% of the income distribution, and the accumulation of household debt in the economy. The state-level analysis provides support to the view that the rise in top income shares led to an accumulation of household debt as a financial asset in the portfolios of the wealthy.

The results linking the saving glut of the rich and the rise in government deficits after the Great Recession shown in this study are preliminary. We are currently "unveiling" the holdings of government debt to see whether the holdings of government debt by the wealthy have increased substantially since the Great Recession. This is an important question for many countries, given the rise in government debt in many advanced economies since 2010.

The findings of this study support the model developed in our companion study (Mian et al. (2019)). The model in the companion study shows how non-homotheticity over the consumption-saving decision can help explain how a rise in the inequality of permanent income earned across households can increase household debt and lower interest rates. The model also has important implications for the effectiveness of monetary and fiscal policy in the presence of non-homotheticity and high household debt burdens.

Figure 16: Top income shares and rising household debt across countries
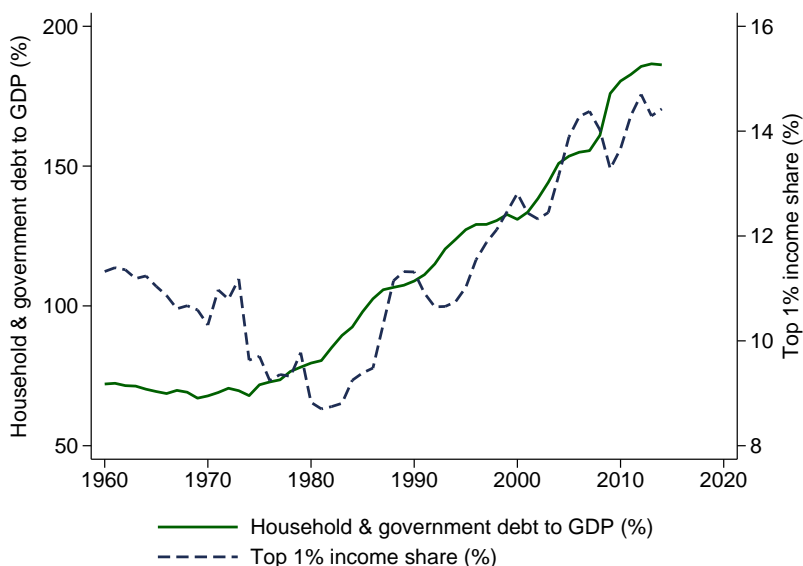
Series are cross-country averages, weighted by real GDP in 1970. The countries in the sample are Australia, Canada, Finland, France, Germany, Italy, Japan, New Zealand, Norway, Portugal, Spain, Sweden, United States and United Kingdom. Data come from the World Inequality Database, IMF Global Debt Database, the Jorda-Schularick-Taylor Macrohistory Database and the New Zealand Treasury. See Mian et al. (2019) for more details.

This study focuses on the United States, but the findings may be applicable to other countries. Australia, Canada, and the United Kingdom, for example, have all witnessed a substantial increase in top income shares and large increases in household debt. Figure 16 shows cross-country averages in the share of income earned by the top 1% and the household debt to GDP ratio for 14 advanced countries. This study focuses on the United States; however, Figure 16 shows that the basic pattern holds across a large number of countries. The saving glut of the rich may be linked to the rise in household debt worldwide. We look forward to further research on this topic.

References

Adelino, Manuel, Antoinette Schoar, and Felipe Severino, "Dynamics of Housing Debt in the Recent Boom and Great Recession," Technical Report, National Bureau of Economic Research 2017.

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.

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.

Auclert, Adrien and Matthew Rognlie, "Aggregate demand and the top 1 percent," *American Economic Review P&P*, 2017, *107* (5), 588–92.

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 2019.

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.

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

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

Chetty, Raj, Nathaniel Hendren, Patrick Kline, Emmanuel Saez, and Nicholas Turner, "Is the United States still a land of opportunity? Recent trends in intergenerational mobility," *American Economic Review*, 2014, *104* (5), 141–47.

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.

Fisher, Jonathan, David Johnson, Timothy M Smeeding, and Jeffrey P Thompson, "Inequality in 3-D: Income, Consumption, and Wealth," 2016.

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.

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 2017.

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.

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

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.

Kaymak, Barış and Markus Poschke, "The evolution of wealth inequality over half a century: The role of taxes, transfers and technology," *Journal of Monetary Economics*, 2016, *77*, 1–25.

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, 1949-2013," Technical Report, CESifo Working Paper 2019.

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

Lee, Chul-In and Gary Solon, "Trends in intergenerational income mobility," *The Review of Economics and Statistics*, 2009, *91* (4), 766–772.

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 and Amir Sufi, "House price gains and US household spending from 2002 to 2006," Technical Report, National Bureau of Economic Research 2014.

_ , Ludwig Straub, and Amir Sufi, "Indebted Demand," Technical Report, Working paper 2019.

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.

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*, 2017, *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.

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.

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 M Zidar, and Eric Zwick, "Capitalists in the Twenty-first Century," Technical Report, National Bureau of Economic Research 2019.

—   **, Owen M Zidar, and Eric Zwick**, "Top Wealth in the United States: New Estimates and Implications for Taxing the Rich," Technical Report, Working Paper 2019.

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.

A Appendix

A.1 Saving glut of rich using CBO top income share

Figure A1: Saving glut of rich using CBO top income share
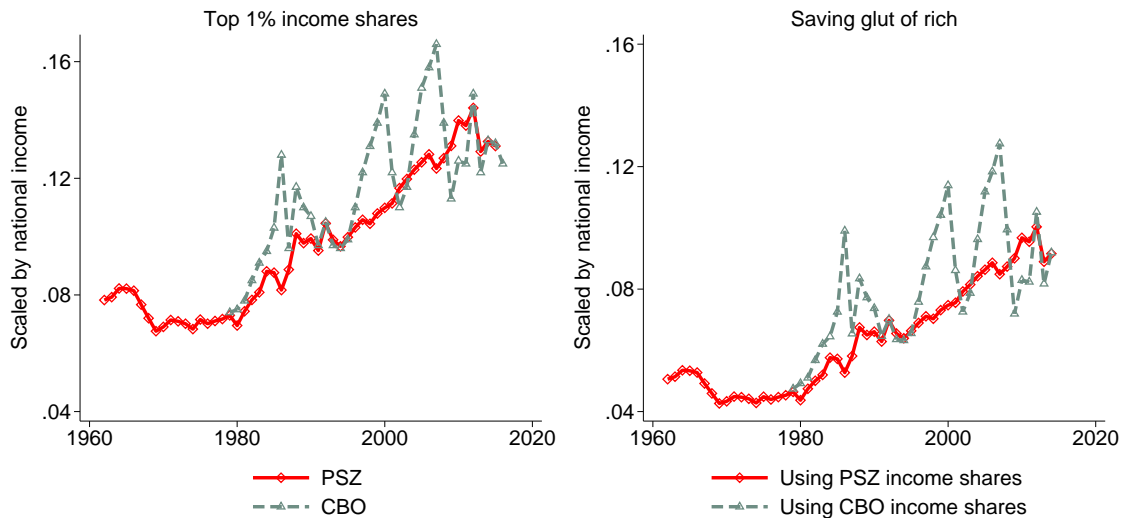
The left panel plots the top 1% after-tax income shares from the CBO and PSZ. The right panel plots the saving glut of the rich using the two different income shares.

A.2 Details on unveiling process of Section 4

Further details and Stata code that unveils who holds household debt in the financial accounts is coming soon.

A critical step to obtain the amount of household debt owned as a financial asset across the wealth distribution is matching the aggregate asset classes through which household debt is held to the asset classes in SZ and DFA. That is, A_c^{HHD} must be matched with a $\omega_{c,p}$ from Section 4.3. For the DFA, this process is almost a perfect match, as the DFA is based on the same categories in the Financial Accounts that are utilized in the unveiling process described in the previous subsection. 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.

There are four asset classes in SZ: equity, pensions, fixed income, and business assets. The various categories of assets from the unveiling exercise must be matched to these five asset classes. For the SZ asset classes, household debt held as a financial asset through bonds, deposits, money market funds, and direct mortgages are matched to the SZ wealth shares for fixed income assets. Household debt held as a financial asset through mutual funds and equity of non-financial and financial corporations are matched to the SZ wealth shares for equity. Household debt held as a financial asset through non-corporate business equity are matched to the SZ wealth shares for business assets. Household debt held as a financial asset through life insurance reserves, life insurance annuities, and any pension plan are matched to the SZ wealth shares for pensions.

A key assumption of this methodology is proportionality of holdings of household debt by asset class across the wealth distribution. The methodology first calculates the total amount of household debt owned in a given financial asset class c, and then it distributes this household debt across the wealth distribution according to the SZ and DFA wealth shares of this asset class. This assumes that the holdings of a given asset class across the wealth distribution are similar enough that household debt held as a financial asset can be allocate across the wealth distribution according to these shares.

A.3   Additional graphs for Section 4

Figure A2 provides details on the assets through which the U.S. household sector holds household debt. The left panel plots more equity-like instruments, and the right panel plots more fixed-income instruments. The most traditional channel through which households would lend to other households would be through bank deposits. Since 1980, however, this has not been an important source of the overall rise in household lending to other households. Instead, the most important channels through which households increasingly lend to other households are pensions, mutual funds, annuities (most of which are variable annuities sold by life insurance companies), equity, and bonds.

Figure A2: Instruments through which Household Debt Held by Households
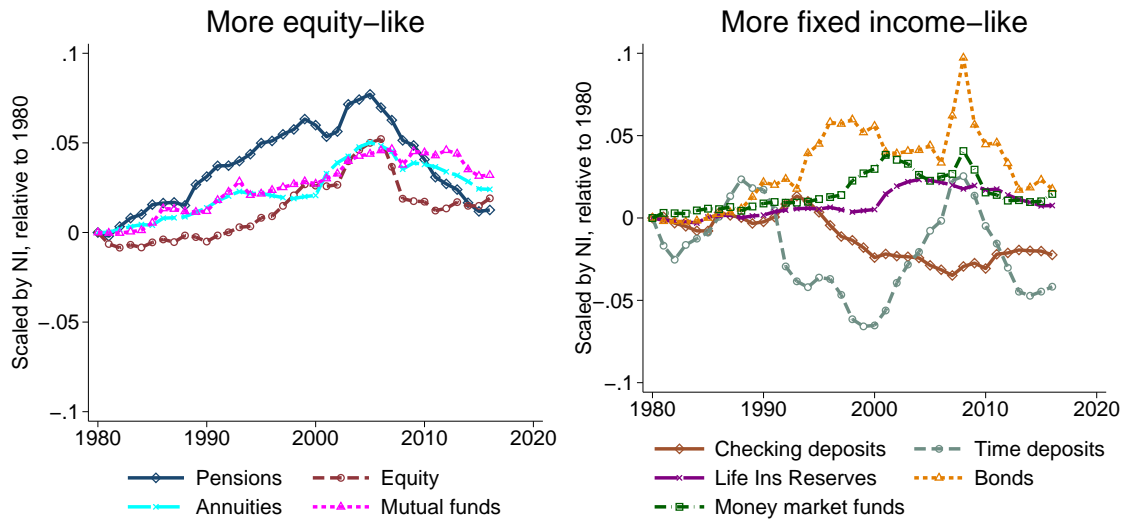
The two panels plot the asset class through which households hold household debt as a financial asset. Bonds include Agency GSE bonds, and bonds issued by financial and non-financial firms. Equity includes the equity of private depository institutions, and both corporate and non-corporate businesses. All series are scaled by national income.

In fact, the asset class for which lower wealth households in the United States traditionally hold a high share is checking deposits. As the right panel shows, holdings of household debt by households through checking deposits has actually declined substantially. This is an indication that the rise in household debt has not been financed by lower wealth households. The distribution of holdings of household debt in the United States is the central focus of the next sub-section.

Appendix Figure A3 uses the DFA allocation shares and shows the five largest asset classes through which the top 10% held household debt as of 2007, and the five largest classes in terms of the rise in debt holdings of the top 10% from 1992 to 2007. From 1992 to 2016, non-financial businesses increased their holdings of deposits and money market funds by 15 percentage points of National Income (see Appendix Figure A4).

## A.4   DFA, SZ, and capitalization factor critique

What are the advantages and disadvantages of the SZ approach versus the DFA? The main disadvantage of the DFA is that it only goes back to 1989. This is a big disadvantage for this study, as the rise in household debt from 1983 to 1989 was substantial, and the DFA do not allow for an analysis of the distribution of holdings for this time period. There are two disadvantages of the SZ approach. First, there are only four asset classes in SZ, and so the match with the asset classes from

Figure A3: Through What Instruments Does Top 10% Hold Household Debt?
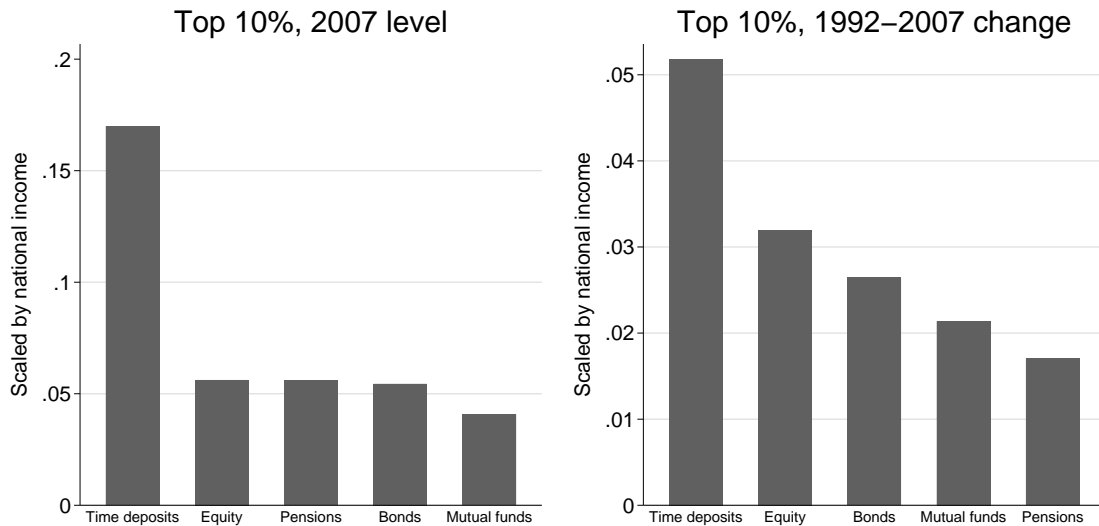
This figure shows the top 5 financial instruments through which the top 10% hold household debt as of 2007 (left panel), and the top 5 financial instruments through which the top 10% increased their holdings of household debt from 1992 to 2007 (right panel).

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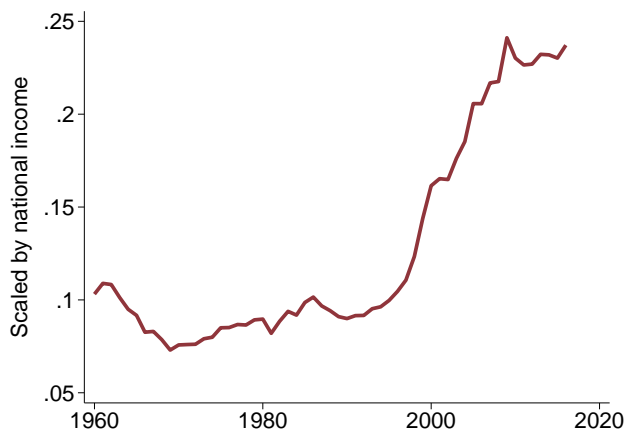

Data are from the Financial Accounts.

the Financial Accounts is less precise. Second, the assumption of a constant capitalization factor for fixed income assets likely overstates the wealth share of the top 1% of the wealth distribution (e.g., Bricker et al. (2018) and Smith et al. (2019b)).

Figure A5 provides more detail on the holdings of household debt across the wealth distribution. This is the breakdown of the results shown in the main text in Figure 10. For both the level and trends, the use of DFA versus SZ wealth shares produces substantial differences when focusing on the top 1% and next 9% holdings of household debt. The SZ wealth shares imply a larger increase in holdings of the top 1%, but a more modest increase in the holdings of the next 9%.

Figure A5: Who Holds Household Debt across the Wealth Distribution? Details
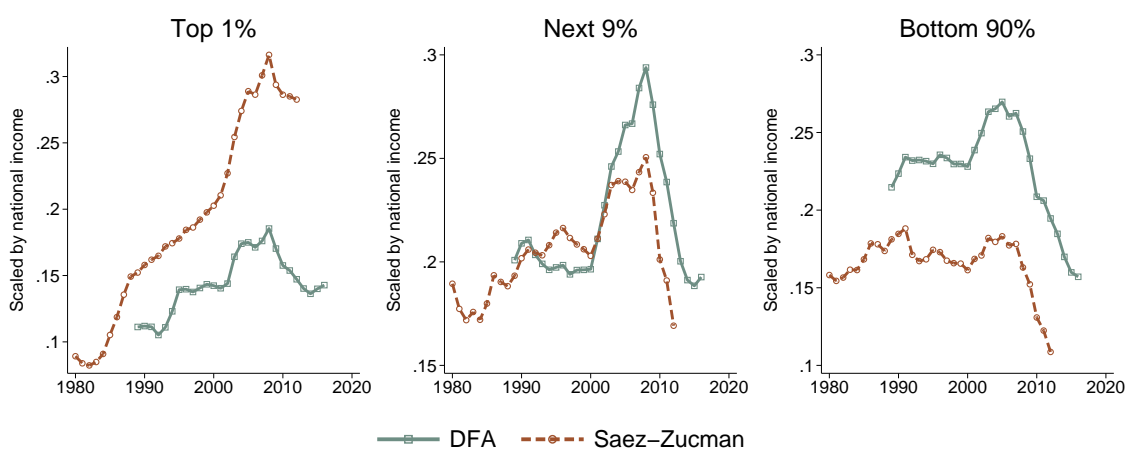
This figure decomposes the holdings of household debt by the U.S. household sector across the wealth distribution. The left panel uses the financial asset shares from Saez and Zucman (2016) and the right panel uses the financial asset shares from the DFA (described in Batty et al. (2019)). All series are scaled by national income.

However, trends for the bottom 90% are similar regardless of the wealth shares used. The DFA wealth shares imply a higher level of household debt held by the bottom 90%, but the trend over time is similar to household debt holdings implied by the SZ wealth shares. The main difference in the two approaches is more about whether the top 1% or the next 9% financed the large increase in household debt from 1989 to 2007. But both measures imply that a substantial fraction of the overall rise in household debt was financed by the top 10% of the wealth distribution.

Recent research suggests that the SZ methodology overstates the level of wealth of the top 1% given the assumption of a constant capitalization factor when estimated fixed income wealth from fixed income asset cash flows (e.g., Bricker et al. (2018) and Smith et al. (2019b)). This is one reason why results above are shown for both the SZ and DFA financial asset shares. To assess whether this issue is material for the pattern in Figure 10 in the main text, we use the wealth shares estimated by

Smith et al. (2019b) (SZZ). The wealth shares from Smith et al. (2019b) were obtained from Figure A.2 in the July 2019 draft of the study. The Smith et al. (2019b) study adjusts the wealth shares relative to SZ on a number of dimensions, one of which is allowing the capitalization factor to be substantially lower for higher income individuals.

Unfortunately, the SZZ shares for the different asset classes and for debt are not available, which prevents us from implementing the exact same methodology using these alternative share estimates. As a rough approximation, holdings of household debt across the distribution are estimated using the SZ and SZZ wealth shares, where the total amount of household debt held by the U.S. household sector is multiplied by the shares to obtain how much household debt is held by each group. Formally, the rough approximation is:

$$HHDebtHoldings_{p,t} = \omega_{p,t} * HHDebtHoldings_t$$

where $\omega_{p,t}$ is the share of wealth of group p according to the SZ and SZZ wealth shares.

Figure A6 shows the results of this exercise. The left panel plots the wealth shares using the adjusted technique for SZ, and the right panel plots the wealth shares using the SZZ wealth shares. As the figure shows, the rise in household debt held as a financial asset for the top 10% is similar.

Figure A6: Robustness of Methodology using Smith et al. (2019b) Wealth Shares
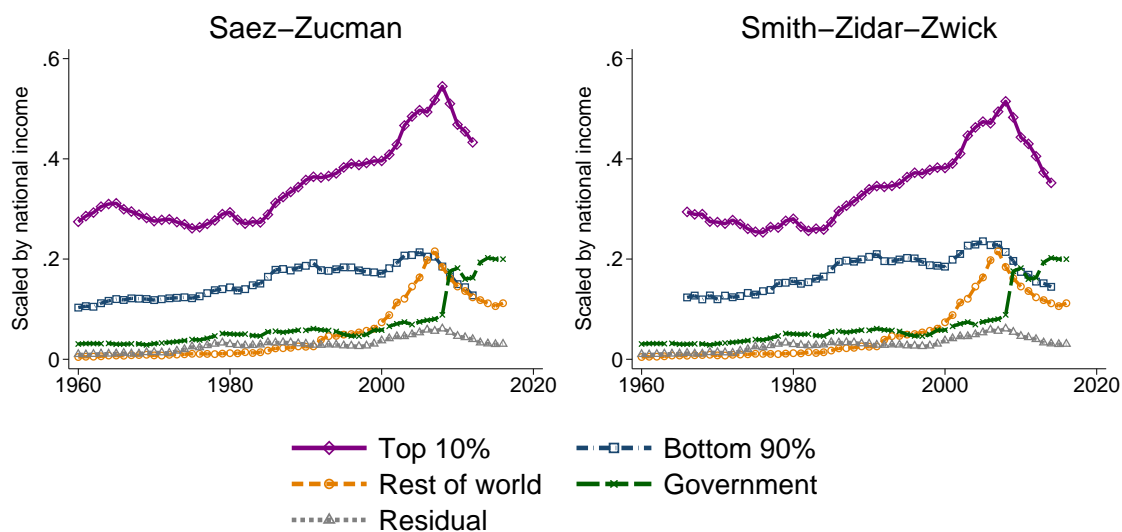
This figure decomposes the holdings of household debt by the U.S. household sector across the wealth distribution. The left panel uses wealth shares from Saez and Zucman (2016) and the right panel uses wealth shares from Smith et al. (2019b). All series are scaled by national income.

The reason the two panels are so similar can be seen in Figure A7, which plots the wealth shares of the top 1% and top 10% for the two studies. For both panels, there is an important level

difference, with the SZZ wealth shares implying a lower share of wealth for both the top 1% and top 10%. However, prior to 2008, there is almost no difference in the trend for either the top 1% or top 10% wealth shares. As a result, use of the SZZ wealth shares implies a lower amount of household debt held by the top 10% of the wealth distribution, but the trend is almost identical before 2008.

Figure A7: Wealth Shares from Saez and Zucman (2016) and Smith et al. (2019b)
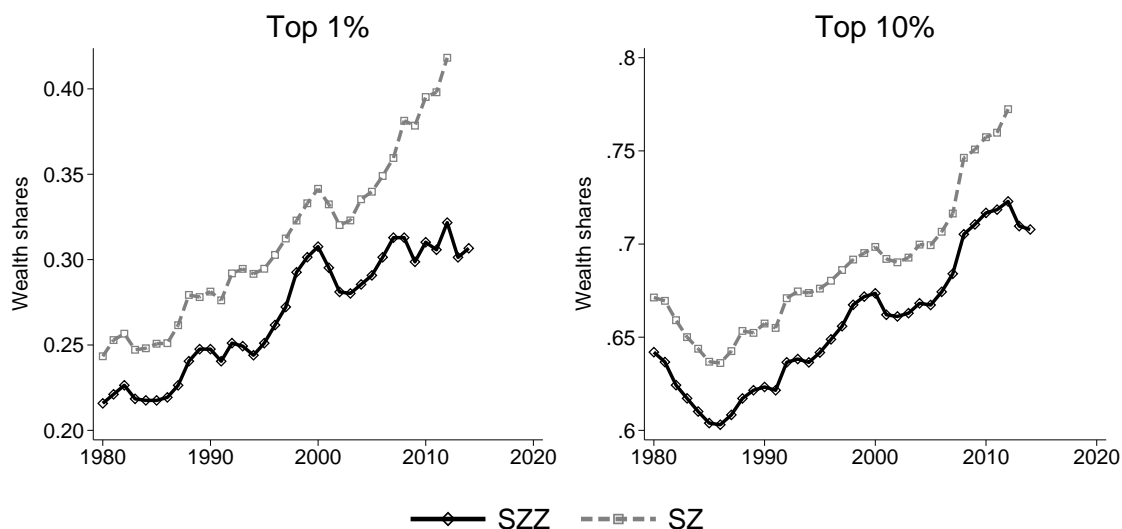
The figure plots wealth shares from Saez and Zucman (2016) and Smith et al. (2019b) for the top 1% of the wealth distribution and the top 10%.

A.5   Further details on state-level analysis

In this section we describe in detail our procedure for assigning tax returns with income above 200K to individual states for the years 1989-2007. Recall that for 1982 we do not need to do this because the individual-level public use tax files contain state identifiers for all observations.

As mentioned above, we obtain the mean interest, dividend and taxable pension income for filers 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 filers. 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 yearly public-use individual level tax return files available at the NBER, 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 filers 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 filers 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 filers with AGI above \$200,000 by taking a weighted mean over all observations in the individual tax return file with AGI above \$200,000. The weights $w_{syi}(I)$ we use are the individual population weights multiplied by the relative likelihood that an individual 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

$$w_{syi}(I_i) = p_i \times \frac{f_{sy}(I_i)}{f_{U.S.,y}(I_i)} = p_i \times \frac{\alpha_{sy}}{\alpha_{USA,y}} \times 200000^{\alpha_{sy}-\alpha_{USA,y}} \times I_i^{\alpha_{USA,y}-\alpha_{sy}}. \tag{11}$$

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 individuals 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 bond, 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.

We can use the same procedure to obtain the sampled mean AGI for filers above \$200,000 by state. Doing this and comparing the values to the true SOI aggregate data, we obtain a correlation of 0.995 and a cross-sectional R^2 of 0.9914 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.

No imputation is required for earners below \$200,000 and for all individuals in 1982 - data for these earners, with state identifiers, are contained in the public-use tax files. From these capitalized measures of total bond, equity, business and pension wealth, we construct a dataset that contains, for various income groups in a state and year, 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

debt is owned as an asset in a state and year, as well as by different income groups therein. Our main groups of focus are the top 6% of earners and the bottom 94%.

Table A1 shows which income variables are used in the Saez and Zucman (2016) capitalization process, as well as the source(s) and to what extent we rely on each variable in allocating household debt ownership. The sampling procedure is used whenever a variable is obtained from the individual data for filers with AGI above $200,000. The final column shows the weight of each variable & source in the final allocation of household debt ownership for the top 6% from 2004 to 2007. In total, 73.5% of household debt allocated to this group relies on state-level data, while the other 26.5% relies on the sampled individual data.

Table A1: Weight of Capitalized Income Variables in Final Household Debt Ownership Allocation

Type of Wealth	Wealth Weight in Allocation of Household Debt Owned	Underlying Capitalized Income Variables	Level of Source Data	Sampling Used	Fraction of Total Assets	Source Weight in Allocation of Household Debt Owned
Bond	0.525	Taxable interest income	State	No	0.195	0.357
			Individual	No	0.031	0.056
		Estate & exempt interest income	Individual	No	0.007	0.013
			Individual	Yes	0.054	0.100
Business	0.033	Partnership, estate, sole prop., & royalty income	Individual	No	0.023	0.005
			Individual	Yes	0.131	0.028
Equity	0.141	Dividend income	State	No	0.239	0.093
			Individual	No	0.035	0.014
		Estate & S-Corp income	Individual	No	0.007	0.003
			Individual	Yes	0.082	0.032
Pension	0.300	Taxable pension & annuity income	State	No	0.033	0.051
			Individual	No	0.046	0.070
		Estate & nontaxable pension income	Individual	No	0.048	0.075
			Individual	Yes	0.069	0.105

This applies only to 2004-2007. Source for all variables in 1982 is the individual-level data. Sampling procedure is utilized whenever the individual-level data is used for filers with AGI above $200,000. No sampling is used when relying on state-level data, or when calculating wealth for filers in the top 6% of a state who earn less than $200,000. All numeric columns sum up to 1. In total, 26.5% of household debt owned is allocated through data that rely on the sampling procedure.

Table A2: Relationships Between Controls and Rise in Top Income Share

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Top 6% Share 1982	0.416						-0.0739
	(0.297)						(0.252)
$\Delta_{82,07}$ Log Per Capita AGI		0.102**					0.120**
		(0.032)					(0.037)
Log Per Capita AGI 1982			0.115*				0.147*
			(0.054)				(0.055)
Top 6% Household Debt Holding 1982				0.0739			0.0795
				(0.084)			(0.106)
Skilled labor share 1980					0.277***		0.0118
					(0.070)		(0.113)
Farms & agriculture labor share 1982						-0.301*	0.00117
						(0.131)	(0.157)
R^2	0.07	0.21	0.16	0.01	0.20	0.14	0.46
Observations	51	51	51	51	51	51	51

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

Table A3: Effect of Change in Top 6% Share on Household Debt Holding For Different Top Income Groups

	Panel A: No Controls				
	(1)	(2)	(3)	(4)	(5)
	Top 6 %	Top 5 %	Top 4 %	Top 3 %	Top 2 %
$\Delta_{82,07}$ Top 6% Share	2.323***	2.257***	2.363***	2.384***	1.145
	(0.228)	(0.319)	(0.361)	(0.477)	(0.643)
R^2	0.66	0.60	0.62	0.61	0.31
Observations	51	47	43	36	7
	Panel B: Controls				
	(1)	(2)	(3)	(4)	(5)
	Top 6 %	Top 5 %	Top 4 %	Top 3 %	Top 2 %
$\Delta_{82,07}$ Top 6% Share	2.367***	2.234***	2.046***	2.023***	1.636*
	(0.266)	(0.283)	(0.304)	(0.374)	(0.620)
R^2	0.57	0.55	0.51	0.48	0.43
Observations	51	47	43	36	7

Panel A estimates column (2) of Table 5 for different top income groups. Panel B estimates column (5) of Table 5 for these same groups. Robust standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Figure A8: Change in Top 1% Share Against Change in Top 6% Share