NBER WORKING PAPER SERIES

THE SAVING GLUT OF THE RICH

Atif R. Mian
Ludwig Straub
Amir Sufi

Working Paper 26941
http://www.nber.org/papers/w26941

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2020, revised July 2025

This is a heavily revised draft of an earlier version that circulated as "The Saving Glut of the Rich and the Rise in Household Debt". We thank Ian Sapollnik, Pranav Garg, Sebastian Hanson, Keelan Beirne, Tudor Schlanger, Veronica Backer-Peral, Bianca He and Amil Mumssen for excellent research assistance. We are grateful to the following scholars who patiently answered questions on various conceptual and data issues: Jesse Bricker, Joseph Briggs, Jonathan Fisher, Fatih Guvenen, Jonathan Heathcote, Martin Holm, David Johnson, Ralph Koijen, Ben Moll, Eric Nielsen, Fabrizio Perri, Lukasz Rachel, Kamila Sommer, Alice Volz, Owen Zidar, and Gabriel Zucman. We also thank Heather Boushey, Benjamin Hébert, Greg Kaplan, Gianni La Cava, Lukasz Rachel, Moritz Schularick, Richard Thaler, Harald Uhlig, Stijn van Nieuwerburgh, Rob Vishny, and many seminar participants at various places. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

The Saving Glut of the Rich
Atif R. Mian, Ludwig Straub, and Amir Sufi
NBER Working Paper No. 26941
April 2020, revised July 2025
JEL No. D31, E21, E44, G51

ABSTRACT

This paper develops a new unveiling methodology to trace distributional household saving through the financial system and applies it to tax records from 1963 to 2019. Since the 1980s, saving by the top 1% has surged, creating a saving glut of the rich comparable in scale to the global saving glut and equally influential in driving the demand for safe assets. Unveiling the financial sector reveals that this glut financed much of the rise in middle-class borrowing before 2008 and the expansion of federal debt thereafter. Our framework helps illustrate how rising inequality reshapes financial intermediation and debt dynamics.

Atif R. Mian
Princeton University
and NBER
atif@princeton.edu

Ludwig Straub
Harvard University
Department of Economics
and NBER
ludwigstraub@g.harvard.edu

Amir Sufi
University of Chicago
Booth School of Business
and NBER
amir.sufi@chicagobooth.edu

A Replication kit available here is available at
https://www.dropbox.com/s/a9bjxhfiax14r5y/MSS2021Febreplicationkit.tar.gz?e=1&dl=0

1 Introduction

The decision of how much to consume versus save across households, and how this evolves over time, is central to understanding financial markets. The financial system exists to transform household saving into productive uses through a complex process of intermediation. A household's saving might ultimately fund the government's fiscal deficit, finance corporate investment, or support another household's debt-financed consumption. Importantly, when saving rates differ across the income distribution, shifts in income inequality can profoundly affect the financial sector—altering not only the total amount of saving in the economy but also the nature of financial intermediation - shaping everything from credit availability to interest rates.

This paper introduces a new *unveiling methodology* that systematically traces household saving through the financial system to identify the ultimate assets they finance. The modern financial sector is a dense network of intermediaries—banks, mutual funds, insurance companies, and pension funds—linked by a complex web of cross-holdings that obscure the final economic use of household saving. To lift this veil, we develop a framework that unwinds these interconnections, mapping the flow of funds from households through financial intermediaries to their final destinations, whether in government debt, corporate investment, or household borrowing.

Applying the unveiling methodology to tax records, we track how household saving evolves across the wealth distribution and identify its ultimate economic use. Our findings reveal the emergence of *the saving glut of the rich*: after remaining stable from 1963 to the mid-1980s, the share of private saving coming from the top 1% quadrupled as a share of national income. Unveiling reveals that the saving glut of the rich is primarily absorbed by non-productive debt. It fueled the rapid expansion of household debt, particularly among the middle class till 2008, and switched to financing rising fiscal deficits post-2008. Overall, the saving glut of the rich is very large, and as important a source for the demand of "safe assets" as the global saving glut.

The unveiling methodology consists of three steps. First, using Federal Reserve flow of funds data, we construct a *direct ownership matrix* that maps cross-holdings between financial intermediaries and asset owners—akin to an input-output matrix in production networks. Following the release of our initial manuscript, Batty et al. (2023) expanded and refined this matrix, which we incorporate into our analysis. Second, drawing on insights from production networks, we apply an operation similar to the Leontief inverse to unveil indirect ownership, producing an *unveiled ownership matrix* that fully traces asset holdings, whether direct or intermediated. Third, we map financial assets to primary assets that fund final use economic expenditures of household, corporate, government, or rest of the world sectors. Our methodology thus untangles the knotted layers of intermediation and links household saving to the final economic uses it finances.

The unveiling procedure is efficient in uncovering indirect holdings and attributing saving to ultimate primary assets. As an example, only 5.3% of Government-Sponsored Enterprises (GSEs) assets are held directly by households in flow of funds data in 2006. But after unveiling, it turns out that households own a total of 64.0% of GSE assets via other intermediaries. Only 1.8 percent of total U.S. assets are left in the residual asset category that cannot be attributed to a known final use.

We map household wealth, recorded in the Federal Reserve's flow of funds data, to annual tax microfiles as in Piketty et al. (2018), allowing us to unveil financial assets across the wealth distribution. Since financial wealth reflects both accumulated savings and asset price appreciation, we construct annual asset price appreciation for fixed income securities, housing, and equities to isolate net saving for each wealth cohort. To ensure accuracy, we carefully estimate equity price inflation so that total private saving remains consistent with national accounts. Maintaining this consistency is crucial to prevent errors that could distort distributional analysis.

The unveiling of the U.S. financial system reveals a steady rise in the "financialization" of the economy until the Great Recession, after which it plateaus and slightly reverses. The share of financial assets held indirectly through intermediaries by U.S. households grew from 45% in 1960 to a peak of 70% in 2007—an increase of over 100% of national income. The length of intermediation chains also expanded by 50% over this period before declining modestly post-crisis.

Our distributional analysis shows that the top 1% save at an exceptionally high rate, averaging well over 40% of their disposable income. The saving rate drops to 20% for the next 9%, falls to 12% for households in the 51st to 90th percentile, and is effectively zero for the bottom 50%, who live hand-to-mouth. While these estimates are derived from repeated cross-sectional tax records, they are similar when using a Survey of Consumer Finances (SCF) panel that is nationally representative and follows the same households from 1983 to 1989. In fact, the saving rate estimate for the top 1% from tax data is conservative relative to the top 1% saving rate estimate in the SCF.

Saving rates at the top have become increasingly skewed, reflecting strong non-homothetic preferences. Between 1963-1982 and 1983-2019, the top 1%'s saving rate rose from 43% to 54%, while the rest saw theirs fall from 14% to 7%. It is this divergence, combined with the rising income share of the wealthy that drives the saving glut of the rich. After 1982, the annual *flow* of saving from the top 1% increases by 2.9pp of national income—over 680 billion in 2024 dollars per year. Crucially, this surge does not boost investment or capital formation but instead funds dissaving elsewhere, primarily among middle-class households.

The unveiling methodology allows us to track what the saving glut of the rich finances by measuring the net debt position of each wealth cohort—the difference between how much they lend and borrow. Between 1963 and 1982, net debt positions remained stable, but then diverged sharply. The top 1% accumulated an additional 20pp of national income in net debt claims, largely against the American middle class from the 51st to 90th percentile of wealth.

The saving glut of the rich mirrors the global saving glut in both scale and economic impact. In particular, the bottom 99% financed expenditures equivalent to an additional 1.3% of national income annually via new borrowing between 1982 and 2019. The top 1% financed this borrowing by increasing their holding of U.S. household debt, considered "safe assets", on a scale similar to the global saving glut. Traditional analysis using national accounts aggregate household saving at the country level, thus obscuring critical within-country shifts. Our study on the other hand highlights the importance of *domestic* saving gluts in an era of rising inequality, showing that their macroeconomic consequences can be as significant as cross-border imbalances.

The saving glut of the rich in the United States identified in this paper is likely to be a global phenomena. Following up on the initial draft of this paper, Bauluz et al. (2022) unveil the financial system for China and Europe as well, and find the emergence of a similar saving glut. There is a large increase in saving at the top of the wealth distribution, and a decline in saving for the rest as we see in the United States.

One symptom of the global rise in the saving glut of the rich is the increase in corporate saving Chen et al. (2017). Corporate ownership is highly concentrated— far more than income— so a rise in corporate saving naturally reflects a rising saving glut at the top. Using OECD data and following Chen et al. (2017), we show that since 1980, corporate saving as a share of national income has risen, while personal saving has declined, reinforcing the view that the saving glut of the rich is a global phenomenon.

The link between corporate and top-end personal saving is partly driven by tax incentives. Many countries favor corporate over personal saving, making it tax-efficient for the wealthy to accumulate wealth inside firms. As a result, some capital gains reflect new saving rather than pure valuation effects. Our methodology, detailed in Section 2, adjusts for this by imputing asset inflation while ensuring total private saving remains consistent with national accounts, preventing accumulated savings from being misclassified as asset price appreciation.

We illustrate the importance of this approach using Norway, which has highly detailed income and wealth data. Relying on public estimates separately from Iacono and Palagi (2023) and Fagereng et al. (2025), we show that properly imputing asset price inflation for equities and real estate— while ensuring consistency with national accounts— produces a distributional net saving rate pattern that is remarkably similar to that of the U.S.

Our baseline analysis relies on wealth shares from the Distributional National Accounts (DINA) microfiles provided by Piketty et al. (2018). We refine these estimates by adjusting the assumed interest rates for the top 1% (Bricker et al. (2018); Smith et al. (2020)) and correcting pension wealth imputations based on concerns raised by Auten and Splinter (2019). To further validate our findings, we compare them with wealth shares from the Distributional Financial Accounts (DFA) produced by the Federal Reserve (1989–2016). Saving estimates for the top 1% from the DFA and adjusted DINA microfiles align closely.[1] Thus, the rise in saving by the top 1% is a robust trend, independent of the specific data source or methodology used.

Relationship with literature. This paper connects shifts in inequality to the size and nature of the financial sector, emphasizing how differences in saving rates across the distribution shape financial markets. While prior work (Saez and Zucman (2016), Kuhn et al. (2019), Smith et al. (2020)) estimates saving rates using wealth data, we go further by unveiling the financial sector and tracking distributional saving over time. This approach highlights how the saving glut of the rich has fundamentally reshaped financial markets.

Our paper deepens understanding of long-run macro-financial linkages. Using the unveiling procedure and distributional tax records, we show how rising inequality expands credit supply from the wealthy while increasing debt demand from lower-wealth households and the government. This dynamic helps explain the long-term rise in debt-to-GDP while exerting downward pressure on interest rates to maintain debt sustainability in general equilibrium as emphasized in Mian et al. (2021a).

Kumhof et al. (2015) also provide a theoretical framework linking rising inequality to increasing household debt. On the empirical side, Bartscher et al. (2020) examine household debt since 1949 and show that between 1983 and 2019 debt-to-income ratios remained stable for the top 1% but rose sharply for the bottom 90%. However, they do not focus on unveiling wealth to trace the financing of debt of different wealth cohorts. Papers such as Cynamon and Fazzari (2015), Rajan (2011), and Bertrand and Morse (2016) emphasize the use of credit by lower-income households in sustaining consumption growth.

The rise in demand for U.S. safe assets is often attributed to foreign demand (Caballero and Krishnamurthy (2003), Gourinchas and Rey (2007), Caballero et al. (2008), Caballero and Krishnamurthy (2009)). This paper shows that domestically, the saving glut of the rich—driven by rising inequality and the high saving rates of the top 1% is an equally significant force.[2]

Our distributional analysis highlights how aggregate national accounts can mask major shifts within the economy that potentially shape financial sector size, total credit, interest rates, and asset prices. In fact, despite declining U.S. net saving since the 1980s, the saving glut of the rich has surged. This underscores the need to incorporate heterogeneity, and a focus on more granular data across the distribution for understanding macroeconomic trends.

Our findings raise fundamental questions about the role of the financial sector and its expanding size since the 1980s (Philippon (2015)). Despite doubling as a share of GDP, its growth has

[1] This consistency echoes Bricker and Henriques (2020), who find similar trends in wealth inequality from 1989 to 2019 when comparing the SCF and the original Saez and Zucman (2016) methodology with the DINA microfiles.

[2] A related argument is made by Klein and Pettis (2020), who link rising capital outflows in surplus countries to increased saving by the wealthy.

primarily financed consumption rather than productive investment, and investment's share of GDP has declined. Understanding the incentives driving this allocation within financial markets is an important question for future research.

Finally, our paper is also related to the recent literature on "demand side asset pricing", that emphasizes the role market segmentation and asset demand elasticity in determining equilibrium asset prices and quantities (e.g. Greenwood et al. (2018) and Koijen and Yogo (2019)). One important area where this has implications is fiscal policy. If rising inequality generates a saving glut of the rich, then that expands fiscal space by increasing the demand for safe assets as explained in Mian et al. (2024).

The rest of the paper is organized as follows. Section 2 introduces the unveiling methodology, tracing household saving to its ultimate use. Section 3 estimates distributional saving, emphasizing the rise of the saving glut of the rich. Section 4 examines how this saving is allocated, unveiling its final use and the evolution of net debt claims across wealth cohorts. Section 5 extends the analysis to international contexts, and Section 6 concludes.

2 Unveiling the Financial Accounts

At the core of the modern financial system lies a dense web of interconnections between financial intermediaries. These linkages channel household saving either toward other households or governments taking on debt, or toward firms seeking to invest. Building on the macroeconomic literature on production networks, this section introduces a new methodology that "unveils" these financial interconnections. By tracing saving through the system, we identify the ultimate "primary assets" it finances, such as household debt, government debt or corporate equity.

2.1 Methodology

The flow of funds data of the Federal Reserve categorizes financial intermediaries into $k = 23$ sectors[3]. These sectors are connected via a complex web of cross-holdings. For example, claims issued by mortgage-backed securities may be held by US-chartered depository institutions, which may, themselves, issue claims held by money market funds, which may in turn hold claims on

[3]The twenty-three categories are: Agency Government Sponsored Enterprises (12), Mortgage Real Estate Investment Trusts (18), Equity Real Estate Investment Trusts (8), US-Chartered Depository Institutions (27), Credit Unions (22), Other Banks in US Affiliated Areas (20), Finance Companies (10), Holding Companies (13), Non-Financial Businesses (5), the Monetary Authority (i.e. the Federal Reserve Bank (16)), Security Brokers and Dealers (26), Foreign Banking Offices (11), Property-Casualty Insurance Companies (25), Life Insurance Companies (15), Private and Public Pension Funds (24), Mutual Funds (19), Money Market Funds (17), Asset Backed Security Funds (14), Closed-End Funds (21), Exchange-Traded Funds (9), Other Financial Businesses (23), Publicly Held Non-Financial Corporations (also called Corporate Non-Financial Business (6)) and Privately Held Non-Financial Corporations (also called Non-Corporate Non-Financial Business(7)). Numbers here and for owners are used for classification in A2.

mortgage-backed securities. All assets ultimately belong to $m = 4$ types of owners: U.S. households (1), the U.S. federal government (2), U.S. state and local governments (3), and the rest of the world (4).

We often label intermediary groups and owners with indices i, j running from $1, \dots, m + k$, where the first $m = 4$ entries correspond to the four owners in the order listed above. The next three columns represent real estate, corporate business, and non-corporate business intermediary sectors, respectively. We choose this ordering to simplify matrix notation when we unveil to primary assets later on.

We first "unveil" the web of financial cross-holdings, attributing each asset held by an intermediary to their ultimate owners by unwinding the chain of intermediation in between. Let $s_t(i, j)$ denote the share of total liabilities issued by an intermediary or owner i that is held directly by intermediary or owner j at time t. We stack these shares $s_t(i, j)$ into a $(m+k) \times (m+k)$ matrix, $\bar{\mathbf{M}}_t$, with i indexing rows and j indexing columns. We call this the *direct ownership matrix*. The first m entries in the first column of $\bar{\mathbf{M}}_t$ represent the share of direct household ownership of the liabilities issued by each of the m owners, and the next k entries represent the share of direct household ownership of the liabilities issued by each financial intermediary sector. By construction, the row sums of $\bar{\mathbf{M}}_t$ are one, that is, $\sum_{j=1}^{m+k} s_t(i, j) = 1$ for all i. $\bar{\mathbf{M}}_t$ represents the intermediation network at time t.

The matrix $\bar{\mathbf{M}}_t$ only captures direct holdings, yet indirect holdings are crucial to understand the workings of the financial system. To give a concrete example, consider the "Government-Sponsored Enterprises" (GSE) financial intermediation sector. This sector is important because it holds most mortgages issued in the United States. However, only 5.3% of GSE claims are held *directly* by U.S. households in 2006. Instead, one suspects that a much larger share of GSE claims are held *indirectly* by U.S. households via the financial intermediation network. Banks[4], money market funds and mutual funds all hold significant shares of GSE liabilities, and U.S. households hold direct shares in each of these intermediaries.

To investigate how much of each intermediary an ultimate owner owns, directly or indirectly, we need to attribute the assets of each intermediary to its direct owners. As a starting point, we make a simple proportionality assumption and attribute the intermediary's assets exactly in line with the direct ownership shares of its liabilities. Based on this assumption, intermediary (or owner) j owns a share $s_t(i, j)$ of i's liabilities directly, and a share $\sum_{\tilde{j}=m+1}^{m+k} s_t(i, \tilde{j}) s_t(\tilde{j}, j)$ indirectly through another intermediary $\tilde{j}$.

To compute ultimate ownership shares for arbitrarily long chains of indirect holdings, we define the matrix $\mathbf{M}_t$ as equal to $\bar{\mathbf{M}}_t$, except that we set the first m rows of $\mathbf{M}_t$ to zero. We can then obtain the indirect holdings by taking powers of the matrix $\mathbf{M}_t$. For instance, the first column of

[4] Classified as "US-chartered depository institutions" as one of k sectors by the Federal Reserve.

$(\mathbf{M}_t)^2$ records shares held by U.S. households in each intermediary k indirectly via a single other intermediary. Likewise, $(\mathbf{M}_t)^3$ contains the shares held indirectly through an intermediation chain with two other intermediaries, and so forth. Accounting for arbitrarily long intermediation chains, we compute

$$\bar{\boldsymbol{\Omega}}_t \equiv \mathbf{M}_t + (\mathbf{M}_t)^2 + (\mathbf{M}_t)^3 + (\mathbf{M}_t)^4 + \cdots = \mathbf{M}_t \left(\mathbf{I} - \mathbf{M}_t\right)^{-1} \tag{1}$$

This matrix is closely related to the Leontief inverse in the literature on production networks. The entries in row $i > m$ tell us how much any intermediary j holds of i's liabilities, directly and indirectly. The first m rows of $\bar{\Omega}_t$ are zero, by construction.

We are particularly interested in the first m column entries of rows $i > m$. We collect these entries in the $k \times m$ sub-matrix Ω_t. These entries exactly capture the distribution of ultimate ownership for each intermediary i. Notably, each row of Ω_t sums to one.

We call Ω_t the *unveiled ownership matrix*. Ω_t unveils the direct ownership patterns given by $\bar{\mathbf{M}}_t$. Continuing with the example above, Ω_t suggests that U.S. households held another 64.0% of GSEs indirectly via the intermediation network in 2006. The unveiling process also provides information on how many layers of intermediation a given asset goes through before it is owned by its ultimate owners. To the extent there are important agency issues at each layer of delegated intermediation, the unveiling data is also useful for questions in corporate governance and finance.

Mapping to net worth There is a simple mapping between the direct ownership matrix $\bar{\mathbf{M}}_t$ and total household net worth W_t. Let $\bar{\mathbf{M}}_{:1,t}$ refer to the first column of $\bar{\mathbf{M}}_t$ which records the shares of each of the $(m+k)$ sectors' liabilities held by the household sector. Let $\mathbf{L}_t$ be a $1 \times (m+k)$ vector containing the total liabilities in dollars issued by the m owners and k intermediaries. Then total household net worth is given by $\mathbf{L}_t \cdot \bar{\mathbf{M}}_{:1,t} - \mathbf{L}_{1,1t}$. In general, total net worth in dollars of the $(m+k)$ sectors is given by $\mathbf{W}_t$,

$$\underset{1\times(m+k)}{\mathbf{W}_t} = \underset{1\times(m+k)}{\mathbf{L}_t} \cdot \underset{(m+k)\times(m+k)}{\bar{\mathbf{M}}_t} - \underset{1\times(m+k)}{\mathbf{L}_t} \tag{2}$$

where the sum of net wealth across all sectors is zero, i.e. $\sum_{i=1}^{m+k} W_{i,t} = 0$.

Mapping to primary assets The unveiling of the financial network allows us to map total saving via financial intermediaries to the $n = 8$ "primary assets" that this saving finances. Primary assets finance the ultimate consumption and investment expenditures in the real economy that economists care about. The eight primary assets are: (i) household debt that funds consumption or investment expenditure via the household inter-temporal budget constraint, (ii) federal government debt that funds government expenditure and investment via the government's inter-temporal budget constraint, (iii) state and local government debt, (iv) credit extended to the rest of the world, (v) claim on non-financial real estate assets[5], (vi) claim on the private corporate sector that funds investment, (vii) claim on private non-corporate sector, and (viii) residual assets that reflect any unveiled financial assets.

To formalize the mapping to primary assets, denote by $\mathbf{P}_t$ the $n \times k$ matrix whose rows reflect the shares of each primary asset held *directly* by the various intermediaries. Further, denote by Λ_t the $n \times m$ matrix whose rows reflect the shares of each primary asset *directly* held by the m ultimate owners. The row-sums of $\mathbf{P}_t$ and Λ_t sum to 1. Then, the $n \times m$ matrix $\mathbf{A}_t$ that captures the shares of each primary asset that is, directly or indirectly, held by the various ultimate owners can be constructed as

$$\underset{n\times m}{\mathbf{A}_t} = \underset{n\times k}{\mathbf{P}_t} \times \underset{k\times m}{\mathbf{\Omega}_t} + \underset{n\times m}{\mathbf{\Lambda}_t} \tag{3}$$

The first column of $\mathbf{A}_t$ records how much of each primary asset is held by the household sector as a whole. The unveiling of intermediation network via $\mathbf{\Omega}_t$ is central to estimating $\mathbf{A}_t$, along with information from the flow of funds data on how primary assets are held by each intermediary to obtain $\mathbf{P}_t$.

2.2 Implementation

Constructing $\bar{\mathbf{M}}_t$ We have formally described the full unveiling process that takes us from the financial assets held by households, governments and the rest of the world via intermediaries, to the primary assets they finance. The entire procedure starts from the intermediation network represented by $\bar{\mathbf{M}}_t$. How can $\bar{\mathbf{M}}_t$ be constructed from real-world data?

The U.S. Financial Accounts provide highly detailed data on assets and liabilities, which we utilize to populate every cell of the matrix $\bar{\mathbf{M}}_t$. For example, suppose that the i^{th} row of $\bar{\mathbf{M}}_t$ refers to Agency GSEs and that the j^{th} column refers to the Federal Reserve. To populate cell $\mathbf{M}_t[i, j]$, we refer to Table L.109 of the U.S. Financial Accounts, which indicates the dollar amount of GSE-backed security assets held by the Federal Reserve in each year. To convert this into the share of total liabilities issued by Agency GSEs, we scale it by Government-Sponsored Enterprises' total liabilities given in Table L.125 of the Financial Accounts. We can repeat this for every column and every row of the matrix. In cases where an institution issues liabilities to another institution through multiple financial instruments, we consider the total amount of liabilities issued. The values are then scaled by the total amount of liabilities issued by each borrower ensuring that the rows always sum to unity.

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

We put together this information using the "from where to where" (FWTW) spreadsheets published by the Federal Reserve. Following the initial public release of our manuscript and building upon our methodology, Batty et al. (2023) produce an augmented version of the matrix $\bar{\mathbf{M}}_t$ which explicitly identifies the data fields used to construct each cell of the matrix. FWTW spreadsheets list the flow of funds variable that records the liability issued by an intermediary or ultimate owner, i, and held by another intermediary, or ultimate owner j.[6] In fact, the information here is broken down further by 34 instruments that represent the cross-holding between the liability and asset side of i and j. We construct $\bar{\mathbf{M}}_t$ by summing across all the instruments,

$$\bar{\mathbf{M}}_t = \left(\frac{1}{L_t}\right)^\top \odot \left(\sum_{i=1}^{\ell} M_{i,t}\right) \tag{4}$$

where ℓ=34[7] Each matrix $M_{i,t}$ is a $(k+m) \times (k+m)$ matrix identifying the total liabilities issued by each borrower and held by each lender through instrument i. $\left(\frac{1}{L_t}\right)^\top$ and the operator $\odot$ scales each row of the summed matrices by the total liabilities issued by each row, thus converting the final matrix in share terms, with each row summing to one. Batty et al. (2023) fill the cells of $M_{i,t}$ with exact values when they are known, and infer the unknown cells via a proportionality assumption as in our earlier version. Batty et al. (2023) also use an algorithm to fill $M_{i,t}$ when some but not all of the cells are known, such that the implied sum of the rows and columns is consistent with the known row and column totals. Going forward, we adopt this methodology and expand it to include corporate equities, an instrument which is not included in the Batty et al. (2023) data set, and separately identify privately and publicly owned corporations.

2.3 What unveiling tells us about the changing nature of financial system

The unveiling process is helpful for understanding changes in the network of financial intermediation. Figure 1 plots the share of U.S. financial assets held by households[8] that are held indirectly through intermediary chains and hence need to be unveiled. The unveiled share of financial assets held by households rises from 45% in 1960 to a peak of about 70% on the eve of the Great Recession in 2007, before falling back down to about 60% in recent years. There is thus a clear

[6]The federal reserve refers to both intermediaries and ultimate owners as "sectors" in its terminology.

[7]These financial instruments include: U.S. Official Reserve Assets and Special Drawing Rights (SDR) Allocations, SDR Certificates and Treasury Currency, US Deposits in Foreign Countries, Net Interbank Transactions, Checkable Deposits and Currency, Time and Saving Deposits, Money Market Fund Shares, Federal Funds and Security Repurchase Agreements, Open Market Paper, Treasury Securities, Agency and GSE Backed Securities, Municipal Securities, Corporate and Foreign Bonds, Depository Loans Not Elsewhere Classified, Other Loans and Advances, Home Mortgages, Multifamily Residential Mortgages, Commercial Mortgages, Farm Mortgages, Consumer Credit, Corporate Equities, Mutual Fund Shares, Trade Credit, Life Insurance Reserves, Pension Entitlements, Taxes Payable by Business, Equity in Non-Corporate Business, Direct Investment, and a Miscellaneous category.

[8]The figure excludes non-financial, rest of world, and residual assets.

rise in "financialization" of assets in the U.S. since 1960 that is reversed with the Great Recession. Figure 1 also shows the total value of unveiled U.S. financial assets held by households as a share of national income. The total size of U.S. financial assets behind the veil of intermediation rises from around 150% of national income in the 60's to well over 250% in recent years.

Figure 1: Indirect Household Ownership of U.S. Financial Assets

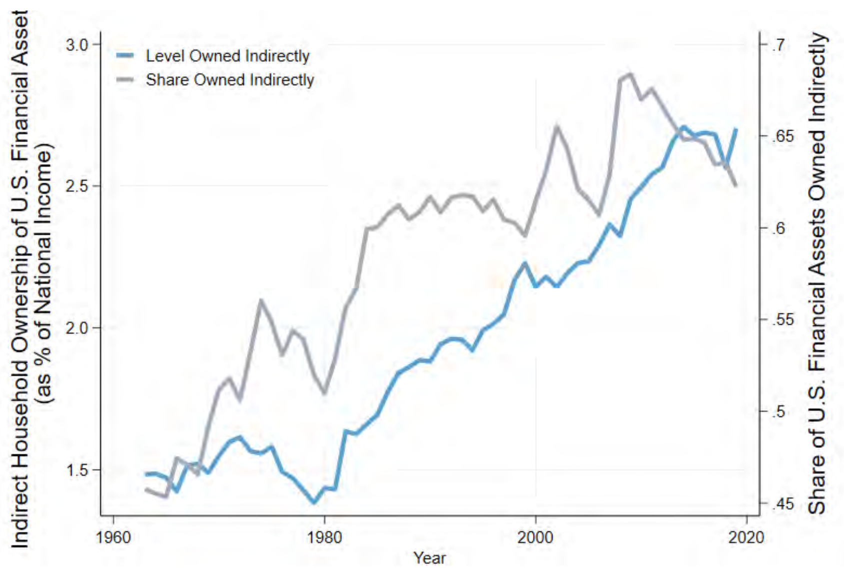

This figure plots the total dollar amount to national income (left axis) and share (right axis) of U.S. Financial assets held indirectly by households through intermediaries.

The unveiling procedure is highly effective in unveiling almost all of the intermediated U.S. financial assets to primary assets. Only less than one percent of U.S. assets remain unveiled on average and are put in the "residual" primary asset category. The unveiling process in equation (1) also allows us to see the number of intermediation chains an asset goes through before being assigned to an ultimate owner. Figure A1 in the appendix plots the dollar-weighted average intermediation chain length, and shows that the average length rises by 47% between 1960 and 2007, before falling - a pattern similar to the size and fraction of assets intermediated over time. The rise of intermediation volume and length reflects rising distance between savers and the ultimate users of finance. As a result, there are more layers of delegation for managing assets, and it is worth investigating the possible financial allocation consequences of such shifts. Figure A2 in the appendix visually displays the change in network structure, and the rise in connections within intermediaries, from 1960 to 2007[9].

There are also interesting differences in the composition of ownership of primary assets between

[9]To avoid too much clutter in the graph, we only show top-5 "most notable" connections for any given issuer or holder - represented by a node.

U.S. households and the rest of the world (ROW). Figure 2 shows that while there has been a steady rise in the share of U.S. federal debt financed by ROW, U.S. household funding is more skewed toward household debt, and state and local government debt.

Figure 2: U.S. Household and Foreign Ownership of Primary Assets

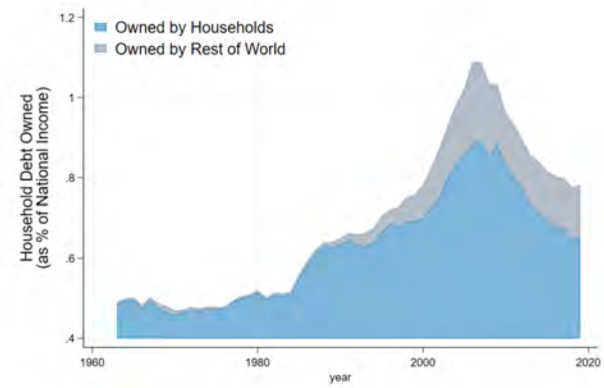

(a) Household Debt

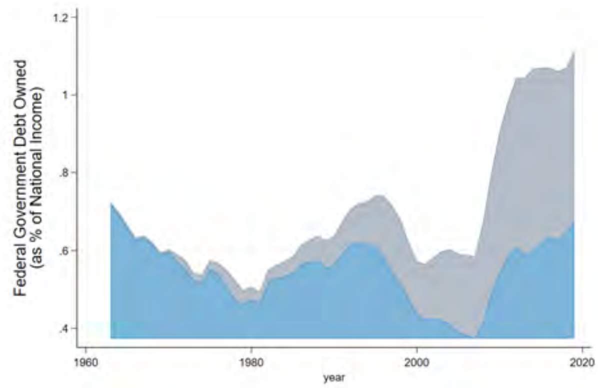

(b) Federal Government Debt

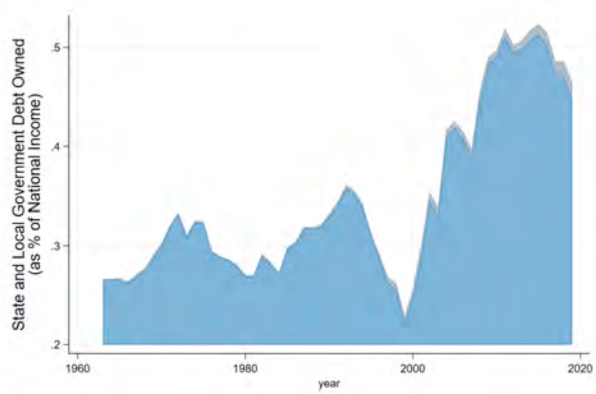

(c) State and Local Government Debt

This figure plots household, federal government, and state plus local government debt held by households and the rest of the world. All series are scaled by national income.

Unveiling also reveals the composition of funding of primary assets by intermediaries. Figure 3 shows how household debt, federal government debt, and state/local government debt are held across various intermediaries[10]. One advantage of stripping out non-business financial assets from the business sector is that we can compare the "pure" market value of business assets with their book value as reported in BEA data and calculate average Tobin's Q. Figure A3 in the appendix

[10]The intermediaries are categorized in the figure as follows: Mutual, Closed-End, and Exchange-Traded Funds comprise Investment Funds. Banks denote Credit Unions, U.S.-Chartered Depository Institutions, Foreign Banking Offices in the U.S, Banks in U.S.-Affiliated Areas, and Money Market Funds. Property-Casualty Insurance Companies and Life Insurance Companies comprise Insurance. Finance Companies, Security Brokers and Dealers, Holding Companies, Mortgage Real Estate Investment Trusts, Other Financial Business, Equity Real Estate Investment Trusts, and Issuers of Asset-Backed Securities are put in the "Other" category.

plots market to book value of real estate and corporate / non-corporate sectors.

The holding of financial assets, as a share of national income, by non-financial firms has risen steadily starting the mid-70s after being flat in the prior period (see Figure A4 in the appendix). Overall, financial assets held by non-financial firms have increased by over 100% of national income. This extremely large accumulation of financial assets by non-financial firms reflects two phenomena. First, corporate saving has become an increasingly larger share of total private saving. Second, the accumulation of financial assets shows that the high corporate saving rate is not entirely devoted to real investment, but is instead partly used to accumulate financial assets. As we show, this has become an increasingly important channel for the very rich to save through the corporate veil, ostensibly for tax reasons.

Figure 3: How Is The Ownership of Primary Assets Intermediated?

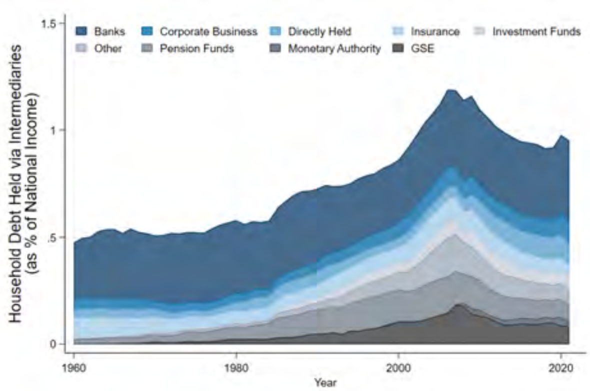

(a) Household Debt

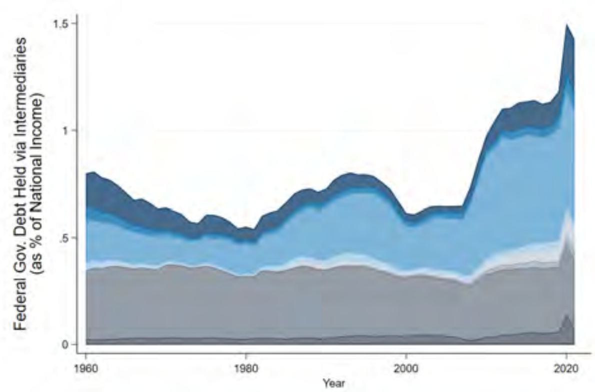

(b) Federal Government Debt

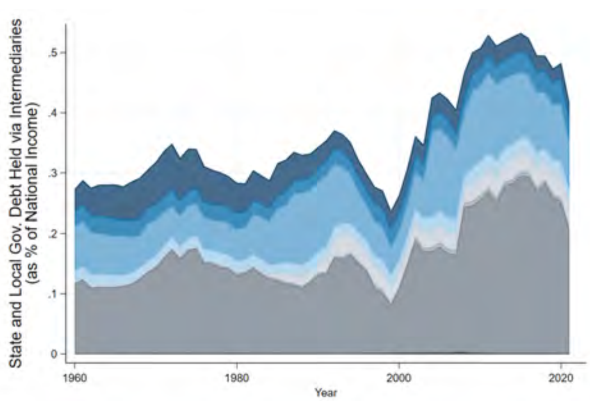

(c) State and Local Government Debt

This figure illustrates the levels of household and government debt held by owners through each financial intermediary, scaled by national income. For this figure, we group the 23 intermediaries from the unveiling exercise into 9 categories.

2.4 Mapping wealth to saving

The primary assets unveiled by our process connect the saving of households, governments and the rest of the world to consumption and investment expenditures that this saving finances. Since the flow of funds data used in the unveiling process is based on the same national income accounts framework, we can formally map unveiled wealth to saving using the national income and product accounts (NIPA) accounting identities.

Let Z^d be disposable income, i.e. net income post taxes and transfers that is available for saving and consumption. $Z^d = Z - T_{RoW} + R_{RoW}$, where Z is national income[11] and T_{RoW} and R_{RoW} are net taxes and transfers vis-a-vis the rest of the world respectively. Z^d is broken down by three sectors in national accounts as $Z^d = Z^{d,p} + Z^{d,\pi} + Z^{d,g}$, where $Z^{d,p}$ is personal household disposable income, $Z^{d,\pi}$ is corporate disposable income, and $Z^{d,g}$ is government disposable income. Total private saving of the household sector, Θ, can be written as,

$$\Theta = S^p + S^\pi = Z^{d,p} + Z^{d,\pi} - C \tag{5}$$

where C is total private consumption. S^p refers to the direct personal saving of households, while S^π refers to household saving done via the corporate veil. The share of private saving done via corporations has increased drastically since the 1990s, with the share rising from 28.2% from 1970-1994 to 47.4% from 1995-2019. Since corporate ownership is highly skewed toward the very rich, attributing saving done via corporations to households is increasingly important for understanding the evolution of distributional saving. The private saving of the household sector is connected to net investment, government saving, and the current account via the saving-investment identity,

$$\Theta = I^n + F - S^g \tag{6}$$

where $S^g = Z^{d,g} - G$ is government net surplus/deficit, $I^n = I - \delta$ is net domestic investment, and F is the current account. Total private saving adds to household wealth, over and above any change in wealth due to asset revaluation. Formally, $\Theta_t = \Delta W_t - \pi_t W_{t-1}$, where W_t is market value of net household wealth, and π_t is asset inflation for wealth held last period.

[11] $Z = Y - \delta + W^{ROW}$, where Y is GDP, δ is consumption of fixed capital, and W^{ROW} is net income from abroad. Z^d is preferred to GDP because it excludes the non-economic income item, capital depreciation or "consumption of fixed capital" as it is called in the national accounts. The same is also excluded in the calculation of net saving.

3 Saving Glut of the Rich

By combining tax records with national financial accounts, we track how wealth accumulation translates into net saving over time. This section shows that since the 1980s, the top 1% has steadily increased its share of private saving, while the bottom 90% has experienced a sharp decline. The shift in distributional saving, which we term the *saving glut of the rich*, has significant implications for financial markets and the economy. We begin by outlining our methodology for estimating saving across the wealth distribution and then present the key results.

3.1 Estimating the distribution of saving

We estimate Θ_i separately for cohort i across the wealth distribution. Let W_{it}^j be the value of asset j held by household cohort i, and $\pi_t^j \equiv \frac{P_t^j - P_{t-1}^j}{P_{t-1}^j}$ be asset price inflation for asset j. Household saving Θ_{it} is given by,

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

The wealth-based approach to estimating distributional saving in (7) requires measurement of W_{it}^j and π_t^j. We use the random sample of individual tax filing data (the DINA microfiles) made available by Piketty et al. (2018) to estimate the share of aggregate wealth in asset class j held by group i in year t, ω_{it}^j. We then compute $W_{it}^j = \omega_{it}^j \cdot W_t^j$, where W_t^j comes from Table B.101 in the Financial Accounts by the Federal Reserve. Total household wealth here matches B.101 except we remove wealth belonging to the non-profit sector, as well as unfunded pensions entitlements[12].

There may be a concern that since we remove wealth belonging to the non-profit sector in calculating saving, our total saving Θ_t may not be consistent with total private net saving in national accounts, as total private saving in national accounts includes the non-profit sector. However, the national accounts report saving separately for the non-profit sector from 1992 onwards, and figure A5 in the appendix shows that this type of saving is essentially zero on average.

The mapping between wealth asset classes in DINA microfiles and the Flow of Funds data from Financial Accounts used in the unveiling procedure is provided in Table A1 of the appendix. The table also provides mapping to wealth asset classes in the DFA data which starts in 1989. Section C.2 of the appendix shows robustness results for saving from 1989 onward using the DFA data to estimate ω_{it}^j.

The wealth-based approach to measuring saving requires a π_t^j for each asset j. In theory, π_t^j refers to asset price inflation that is driven only by inflation or valuation effects. This is not the overall return on the asset, but only the pure asset price inflation. To use the consumer price analogy,

[12] Unfunded pension entitlements are not backed by a financial asset, and as such not funded directly via an active saving flow.

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

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

$\pi_t^{ij} = 1 - WD_t^{ij}$, where WD_t^{ij} is the percentage of debt that is written-down in a particular year for cohort i. WD_t^{ij} is estimated by first calculating net chargeoffs as a share of outstanding debt on bank balance sheets, separately for home mortgages and consumer credit. Since we know total outstanding debt in a given year, the net chargeoff ratio gives us the total amount of debt that is written down. We then distribute the written down debt to cohort i based on the fraction of total defaults accounted for by cohort i. This number is computed using zip code level data on defaults and average income of households living in a zip code.[13]

Estimating π_t^{equity}

The procedure above gives us π_t^j for real estate and fixed income, and the only remaining π_t^j is the one for equity. When estimating π_t^{equity}, the previous literature has used capital gains for equity. However, this is not a correct choice for estimating saving. The reason is corporations also save on behalf of shareholders via retained earnings, and while this saving shows up in capital gains, it should be stripped away when measuring π_t^{equity}. In particular, using capital gain as π_t^{equity} typically leads to an overestimation of the valuation component of wealth and, hence, and underestimation of saving. The rise of corporate saving means that this issue has become only more important over time. We therefore depart from the assumption that capital gain is the same as π_t^{equity}, and estimate π_t^{equity} independently.

Our approach for estimating π_t^{equity} takes advantage of the fact that we have already estimated

[13] The complete details of this methodology are shown in Appendix Section C.4.

the remaining asset inflation terms for other assets. With only one π_t^{equity} remaining, we can back it out as the residual pricing factor that ensures that aggregate private saving from the wealth-based approach matches the aggregate private saving in NIPA (i.e. that $\sum_i \Theta_i = S^p + S^\pi$). This approach also has the additional benefit that it ensures that the distributional saving estimates are national-account consistent, and aggregate up to total private saving in national accounts. As we show in Section 5.2 in the case of Norway, when the imputed π_t^j are inconsistent with national accounts in the aggregate, it can lead to serious errors in estimating distributional saving[14].

The resulting π_t^{equity} from our approach has a strong correlation with the equity capital gains series from S&P500 of 0.89 (see Figure A6 in the appendix). However, the mean of π_t^{equity} is 2.4 percentage points lower - exactly what one would expect given that capital gains includes corporate saving. Thus, on average, about 2.4 percentage points of capital gain a year reflect valuation gains from accumulated corporate saving.

3.2 National Account Consistent Distributional Saving Rates

Ensuring that saving rates estimated from micro data are consistent with national accounts is essential. This alignment guards against estimation errors and captures important components of saving that micro data alone may miss. One such component is *residential investment* which counts toward personal saving. For instance, when homeowners expand their houses (e.g., by adding a room), this form of saving is recorded in national accounts (often via building permit data) but may be absent from standard micro datasets.

Another critical component, especially for wealthy households, is *corporate saving*. Corporations—both public and private—retain a substantial share of earnings on behalf of their shareholders. Shareholders often prefer this form of saving due to tax deferral advantages. Without explicit adjustments, micro data typically overlook saving done via the corporate sector. For these reasons, our methodology ensures that micro-data-based distributional saving rates aggregate to national account totals. Formally, we define three saving rates using the net private saving equation (7):

[14] Section B of the appendix maps the national accounting identities used in this paper for the United States to the UN System of National Accounts (SNA) to facilitate the implementation of our broader methodology in other OECD economies that have all adopted SNA.

$$s_{it}^{G} = \frac{\Theta_{it} + \pi_t W_{i,t-1}}{Z_{it}^{d}} \tag{8a}$$

$$s_{it}^{\kappa} = \frac{\pi_t W_{i,t-1}}{Z_{it}^{d}} \tag{8b}$$

$$s_{it}^{N} = \frac{\Theta_{it}}{Z_{it}^{d}} \tag{8c}$$

s_{it}^{N} denotes the *net saving rate* for wealth cohort i, s_{it}^{G} is "gross" saving rate that includes *capital gains* (or asset price inflation π_t), and s_{it}^{κ} is the saving coming in the form of asset price inflation, or capital gains. The three saving rates are linked via the accounting identity $s_{it}^{G} = s_{it}^{\kappa} + s_{it}^{N}$. The gross saving s_{it}^{G} can be decomposed into a net saving "active" flow, $\frac{s_{it}^{N}}{s_{it}^{G}}$, that results from consumption less than income, and a capital gains portion, $\frac{s_{it}^{\kappa}}{s_{it}^{G}}$. The denominator, Z_{it}^{d}, is total disposable income for cohort i, defined as:

$$Z_{it}^{d} = Z_{it}^{d,p} + Z_{it}^{d,\pi}$$

where $Z_{it}^{d,p}$ is personal disposable income, and $Z_{it}^{d,\pi}$ is corporate disposable income (retained earnings) attributable to cohort i. To accurately reflect saving behavior across the wealth distribution, both corporate saving and corporate income must be included in the numerator and denominator, respectively.

To ensure consistency with national accounts, the income-weighted sum of cohort-level saving rates must match respective aggregates from national accounts:

$$\sum_{i} \left(s_{it}^{G} \cdot Z_{it}^{d} \right) = \Delta W_t \ , \quad \forall t \tag{9a}$$

$$\sum_{i} \left(s_{it}^{\kappa} \cdot Z_{it}^{d} \right) = \Delta W_t \ - \ S_t^{p} - S_t^{\pi} \ , \quad \forall t \tag{9b}$$

$$\sum_{i} \left(s_{it}^{N} \cdot Z_{it}^{d} \right) = S_t^{p} + S_t^{\pi} \ , \quad \forall t \tag{9c}$$

Equations (9a) and (9c) define the key consistency conditions: micro-data-based saving rates s_{it}^{G} and s_{it}^{N} must match national accounts in terms of changes in household net wealth (ΔW_t), personal saving (S_t^{p}), and corporate saving attributed to households (S_t^{π}).

Calculating S_t^π, corporate saving belonging to the household sector

The U.S. national accounts separate out corporate saving that accrues to private households as part of total net private saving. However, in some countries this is not directly reported, and government may also own a significant fraction of corporate saving in countries with sovereign wealth funds. There may also be additional issues in small open economies with substantial holding of foreign equities. We outline how S_t^π can be calculated using sectoral account data in these countries using the example of Norway in section D.1 of the appendix. Section 5.2 uses this estimate to plot net saving rate across the wealth distribution for Norway.

3.3 Saving rates in repeated cross-section vs. panel data

The measurement of saving in DINA data comes from individual-level repeated cross-section random sample of tax returns. One concern with measuring saving via repeated cross-section is that the measured saving rate might be biased by individuals who enter or exit the sample from one period to the next. For example, consider the calculation of saving for the top 1% cohort from $t-1$ to t. An individual in the top 1% cohort in year $t-1$ might drop out of the top 1% cohort in year t. Exit from the top 1% can happen for reasons such as mean reversion. Since individuals are sorted by estimated wealth, there is likely mean reversion for individuals with temporarily high wealth estimates. In general, the exit and entry of individuals from the top 1% can lead to either under- or overestimation of the true saving rate for the top 1% cohort, depending on the relevant wealth of a new entrant that replaces the individual exiting from top 1%. For our analysis, we care more about the *change* in the saving behavior of the rich versus the non-rich over time, and it is more difficult to think of reason why that would be biased. Nonetheless, we address this concern carefully here.

We test for a potential bias in saving rate estimates due to the repeated cross-sectional nature of the DINA data by using an individual-level panel data for the years 1983 and 1989 available from the Survey of Consumer Finances (SCF). The SCF is a nationally representative survey of U.S. households, done every three years. While the SCF is also a repeated cross-section, it made an exception in 1989 and resurveyed participants from the 1983 wave, including those who may have moved. We thus have a complete panel of individual wealth levels between these two years. For each individual, we use their wealth data between 1983 and 1989 to estimate saving rates. The data on the wealth and income of households interviewed in 1983 is taken from the dataset produced by Kuhn et al. (2020), while for 1989 we produce these same statistics using code provided by the Fed, which is available for 1989 and the subsequent surveys.

We estimate the average annual net saving rate, s_i^N, for household i across the 1983-1989 period as,

$$s_i^N = \frac{W_{i,1989} - \prod_{t=1984}^{1989}(1+\overline{\pi}_{i,t})W_{i,1983}}{Y_{i,1988}^D + \sum_{k=0}^{4}\left(Y_{i,1987-k}^D \prod_{t=1989-k}^{1989}(1+\overline{\pi}_{i,t})\right)} \tag{10}$$

where $\overline{\pi}_{i,t}$ is the average capital gain for household i[15], and Y^D denotes the disposable income[16] for household i. We sort households into percentiles based on their average wealth between 1983-1989, and divide them into four groups- top 1%, next 9%, next 40%, and bottom 50%.

Figure 4 plots estimated saving rates for the 1983-1989 period from both the SCF individual-panel, as well as the individual repeated cross-section DINA microfiles. The saving rates are similar in the two data sets in terms of their trend across percentile groups. In both data sets, the top 1% save at a much higher rate (40% to 50% out of disposable income) than the rest, and the saving rate declines across the distribution with the bottom 50% of the population having zero saving rates in both data sets.

Figure 4: Saving Rates Using Individual Panel (SCF) and Repeated Cross-section (DINA)

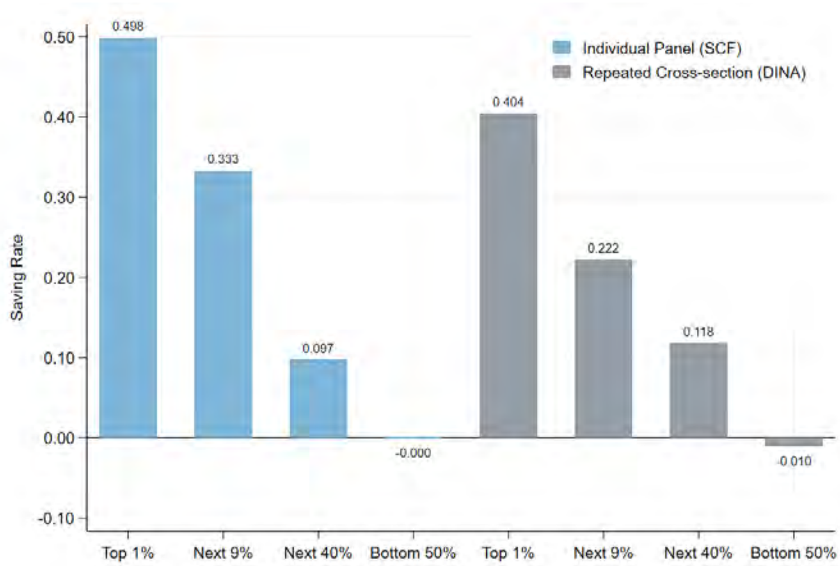

This figure compares saving rates across the wealth distribution for the period 1983-1989, using individual panel data from the Survey of Consumer Finances (left) and repeated-cross section data from DINA (right).

Figure A7 in the appendix reproduces the saving rate in 4 from our main DINA tax data, but also adds capital gains to the net saving rate. Inclusive of capital gains, the difference in saving rates is even larger across wealth groups. For example, the saving rate for the top 1% exceeds well over 60%, while saving rate for the 51-90 percentile, and 1-50 percentile groups is only around 20% and marginally above zero respectively. The increased gap in saving rate inclusive of capital gains is partly driven by the fact that wealth is more concentrated than disposable income.

[15] $\overline{\pi}_{i,t}$ is calculated as the weighted mean of $\pi^j_{i,t}$, using the share of household i's wealth in each asset j as the corresponding weight. The $\pi^j_{i,t}$ come from the procedure already outlined.

[16] The SCF Survey reports only pre-tax income, so we adjust each household's income using the ratio of disposable to pre-tax income, by wealth percentile, from the DINA microfiles. The saving rates from the SCF Survey, measured relative to pre-tax income, are 32.3 %, 26.9 %, 8.5 %, and 0 % for the top 1 %, next 9 %, next 40 %, and bottom 50 %, respectively.

There are two important takeaways from the above analysis. First, saving rates rise with wealth, especially at the very top end for the top 1% and the next 9%. The rise in saving rates at the top end of the wealth distribution reflects non-homotheticity of preferences with respect to the saving decision. Second, since these results are based on individual level panel data, statistical concerns from the repeated cross-section data do not apply here. These results also suggest that the rise in inequality will generate an increase in saving at the top end of the distribution, which we explore in greater detail next.

3.4 The rise of the saving glut of the rich

Saving rates estimated using household-level repeated cross-section tax returns show no particular bias when compared with saving rates estimated using household panel data over the same time period. Both the repeated cross-section DINA data, and the SCF panel data also have the advantage of being randomly drawn, and are nationally representative with over-sampling at the very top of the wealth distribution. The over-sampling at the top-end is very useful for us since it means that we can estimate saving rate at the top-end with precision. In the rest of analysis, we use the DINA microfiles to estimate saving across the distribution and over time, since DINA data is available consistently from 1962 till 2019. It therefore provides us with the best possible understanding of the shifting patterns of saving across the U.S. households, and its possible consequences.

Figure 5 plots the estimated total private saving (Θ_{it}) as a share of aggregate national income for the top 1% and the bottom 99% of households (sorted by wealth). The saving estimate at the annual frequency has a lot of noise, driven by the difficulty of measuring wealth and asset inflation π accurately over shorter horizons. We therefore present the 10-year moving average of the estimates in Figure 5. The moving average also has the advantage that by design it picks up the slow moving but persistent shifts in saving behavior over long-horizon, which is the principal objective of this paper. The first observation in the figure starts from 1972, which reflects the average saving from years 1963 through 1972.

Figure 5 subtracts the 1982 level of saving as a share of national income from each line so it is easier to focus on the magnitude of the relative shifts over time. The 1982 saving levels are mentioned in the figure note for ease of understanding the absolute levels as well. We combine the entire distribution below the top 1% into a single bottom 99% category in order to emphasize the most important sub-group that starts to deviate from the rest of the U.S. population: the top 1% is the most relevant and important group for understanding the biggest shift in saving distribution in the United States.

Our underlying data allows us to estimate saving rates at more granular levels within the bottom 99%, and we present some of that breakdown in the light dashed lines in Figure 5, which plots saving as a share of national income separately for the next 9%, 51-90, and the bottom 50% within the bottom 99%. The largest decline in saving as a share of national income comes from the bottom 90%, and within it the 51-90 wealth percentile group. It is also important to note that we are making public all of the underlying raw data and code. So interested readers can investigate other distributional cuts of the data as they wish going forward. Moreover, all of our data - the unveiling, as well as distributional saving - can be updated going forward as the underlying data sets are updated publicly.[17]

Figure 5 summarizes one of the key findings of our paper: there has been a dramatic shift in the relative magnitude of saving from the top 1% compared to the rest of the population. Total net saving coming from the top 1% and bottom 99% remained relatively constant between 1962 and 1982, at 2.5% and 10.5% of national income annually, respectively. However, starting in the early 1980s, there is a dramatic divergence between the two. The 10-year moving average of saving by the top 1% increases steadily starting in the mid-80s, peaking at 7.1 percentage points of national income above its 1982 level. This amounts to an over three-fold increase in annual total saving of the top 1% as a share of the national income at its highest. In fact, at its highest, almost the entire net saving of the household sector was coming from the top 1%. We refer to the large increase in saving of the wealthiest top 1% of the population as the *saving glut of the rich*.

Figure 5: Saving by the top 1% and bottom 99%

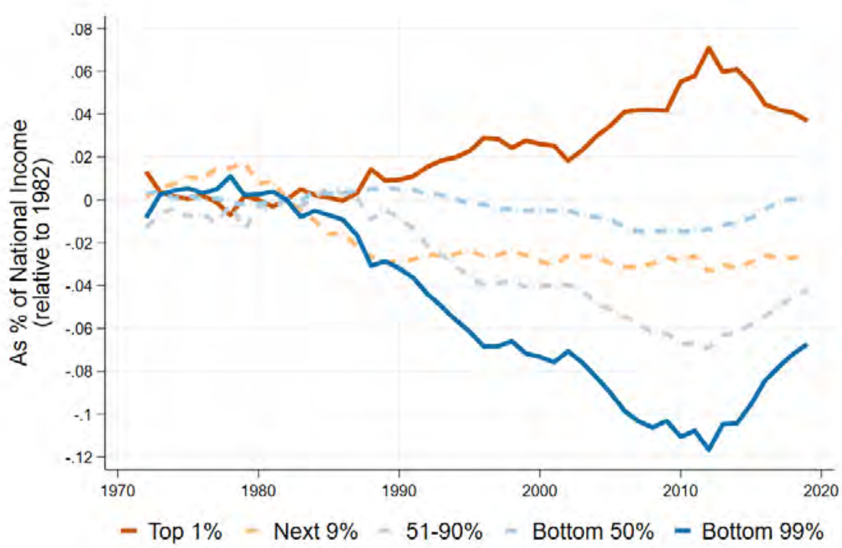

Note: This figure plots the 10-year moving average of saving for the top 1% and bottom 99% of households, scaled by national income and relative to 1982. Saving by the top 1% and bottom 99% in 1982/83 were 2.0% and 11.5% of national income, respectively (10-year MA). Saving by the next 9%, 51-90%, and bottom 50% are illustrated using the dotted lines.

[17]The main constraint is the DINA microfiles which are currently available till 2019.

The annual saving from the top 1% increased by 2.9 percentage points of national income on average, relative to 1982. In 2024 dollar terms, this represents an additional saving flow coming from the wealthiest top 1% of over 680 billion dollars *every year* from 1982 onward. An increase in the saving flow of this magnitude must be absorbed elsewhere in the economy according to the accounting identity,

$$\Theta_{top1} + \Theta_{bot99} = I^n + F - S^g \tag{11}$$

Figure 5 shows that the large rise in saving of the top 1% was more than compensated by the fall in saving of the remaining bottom 99%. The fall in saving of the bottom 99% was not only large enough to absorb the rise in saving of the top 1%, but also the rise in global saving glut (F). Overall, annual saving from the bottom 99% decreased by 6.7 percentage points of national income after 1982. The decline in saving from the bottom 99% represents a drop of 60% relative to the average saving of 11.1% of national income before 1982.

A more than tripling of saving from the wealthiest top 1% and a drop in saving of sixty percent from the remaining 99% represents a massive shift in saving composition after 1982. One likely factor behind the rise in saving of the top 1% is the rise in the share of disposable income going to the top 1%. Given the very high rate of saving for the top 1% relative to the rest of the population documented in Figure 4, a rise in the top income share will increase the contribution of saving from the top 1%. However, the increase in income inequality is not the only reason for the rise in saving glut of the rich.

Figure 6 estimates saving rates by wealth percentile separately for the periods 1963-1982 and 1983-2019. The underlying microdata is oversampled at the upper end of the wealth distribution; therefore, we bin percentile categories into a single bin for 0 to 40th of wealth, and then bin every single percentile from then on. As seen earlier, saving rates rise non-linearly with wealth, with the top 1% having a much higher saving rate than the rest. Figure A8 in the appendix plots the full time series of saving rates separately for the top 1%, next 9%, the next 40%, and the bottom 50%. The largest decline in saving rate comes from the bottom 90%.

Figure 6: Saving Rate by wealth percentile over time

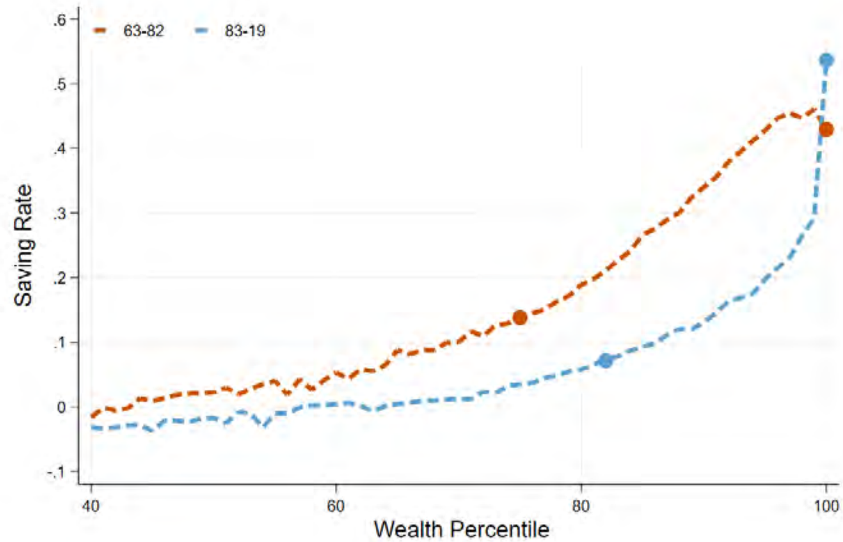

This figure plots average saving rates across the wealth distribution before and after 1982. Dots denote the disposable income-weighted mean saving rate for the top 1% and bottom 99% within each period.

The key new finding in Figure 6 is that the saving rate for the top 1% relative to the rest of the population became even more skewed after 1982. For example, the top 1%'s saving rate rose from 42.9% during 1963-1982 to 53.6% afterwards, while average saving rate for the rest of the population declined from 13.8% during 1963-1982 to 7.1% afterwards. Overall, the rise in the saving glut of the rich and the big shift in the composition of aggregate saving after 1982 is a result of both the rise in top-income inequality and greater divergence in saving rates.

The saving rate and total saving estimated in this paper are determined in equilibrium. For example, the accounting identity in (11) has to hold at all times in equilibrium. Thus, if total saving flowing into the financial system from the top 1% rises, prices and incentives will naturally adjust in the broader economy to make sure that there remains an overall balance between the sources and uses of saving. As an example, interest rates might fall to encourage more people to borrow the "excess saving". Asset prices may rise as a result of falling rates, and partly facilitate the absorption of excess saving by allowing people to borrow more against the rising value of collateral. The primary objective of this paper is only measurement though. We leave the important questions of which margins are more likely to adjust, and why, for subsequent research.

3.5 Compositional shift in distributed saving

The wealth-based approach also provides a detailed look into the drivers of changes in saving across the wealth distribution. Starting with equation (7), $\Delta W_{i,t} = \sum_{j \in J} \Delta W_{i,t}^{j}$ is the annual change in net worth of wealth cohort i, $\Delta V_{i,t} = \sum_{j \in J} \pi_t^j W_{i,t-1}^j$ is the valuation effect, and $\Theta_{i,t}$ is saving. Then, $\Theta_{i,t} = \Delta NW_{i,t} - \Delta V_{i,t}$. Table A2 reports these three terms as a share of national income for various wealth groups over time.

Figure A9 plots saving, change in net wealth and change in valuation components separately for the top 1% and bottom 99%. The top 1% not only increase their saving as already seen, but also experience a large increase in net wealth due to additional valuation gains. The bottom 99%, on the other hand, experience a decline in net wealth as a share of national income despite early valuation gains. The reason for this is the large decline in saving identified earlier. In other words, the bottom 99% dissave or borrow against the rising valuation of existing assets, deteriorating their overall net wealth position - especially when the levered real estate assets decline in value post the Great Recession.

Figure 7 plots household debt as a share of national income separately for top 1% and bottom 99%. We again normalize household debt to national income to its 1982 level for comparison across the two groups. The two groups have a stable ratio for household debt from 1962 till 1984. However, starting in the mid-80s, the bottom 99% start borrowing at a stronger pace, which accelerates during the 1998 to 2007 period. It is remarkable that there is no significant movement in household debt in comparison for the top 1%.

Figure 7: Household debt

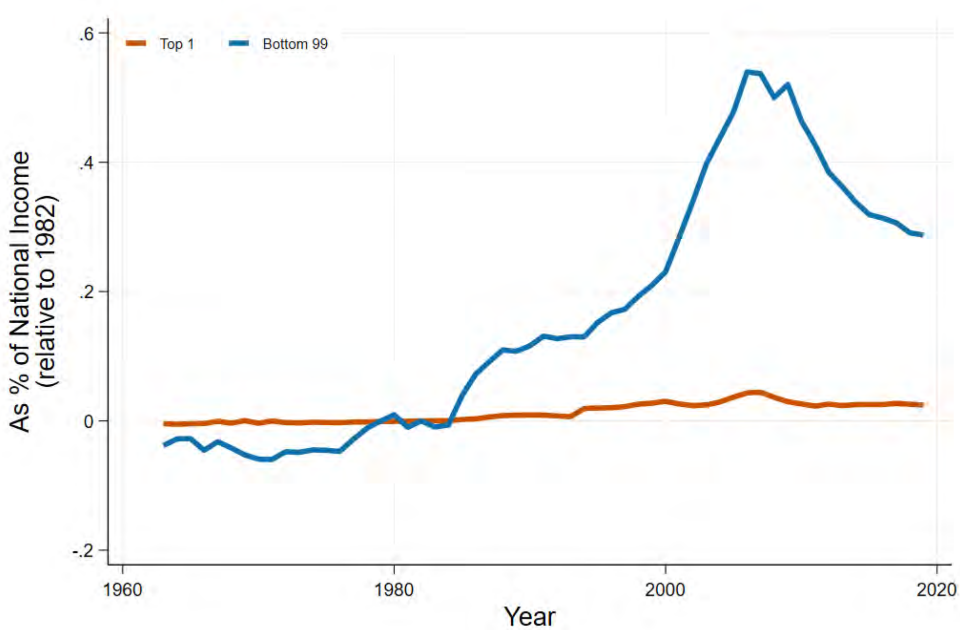

This figure plots total liabilities for the top 1% and bottom 99% of households. All series are scaled by national income and expressed relative to 1982.

The single most important factor that explains the large decline in saving of the bottom 99% is the rise in household debt: borrowing against rising valuations of existing assets. As already emphasized, the rising valuation of assets and rising borrowing against it could be equilibrium responses to deeper underlying structural shifts as in Mian et al. (2021a). When the housing market collapsed, so did the effect of valuation gains on net worth, which resulted in a much lower annual change in net worth during the 2008 to 2019 period relative to the earlier period.

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

The wealth-based approach also allows us to explore the precise manner in which wealth cohorts altered saving over time. For example, starting again with equation 7, we split asset classes j into three groups: financial assets (FA), real estate (RE), and debt (D). This allows us to decompose saving by each group i into: $\Theta_{it} = \Theta_{it}^{FA} + \Theta_{it}^{RE} + D_{it}$, where negative values of D indicate more borrowing. Table A2 in the appendix decomposes saving according to this equation for various wealth groups. The results show that while increased borrowing is the most important reason for decline in saving below the top 1%, part of the dissaving also reflects sales of existing financial assets to the top 1%.

3.6 Robustness of main result

The main finding of this study is the rise of the saving glut of the rich: there is a large increase in the flow of saving coming from the top 1% starting in the 1980s. We used the tax files for documenting the main result because tax data is the most comprehensive and long-running data set available. The data starts in 1963, is representative of tax filers, and over-samples the very top end of the distribution. Section C in the appendix shows robustness of the main results, to alternative data sources, and to variations of the wealth capitalization technique.

The methodology thus far relies on estimating saving via the wealth-based approach that relies on measuring wealth over time, and uses equation (7) to back out saving. Section C.1 of the appendix uses an alternative, the income-less-consumption approach, to estimate saving. This method subtracts consumption from after-tax income to measure saving. The main challenge in implementing this approach is getting good measures of consumption, especially by wealth distribution. Using measures of consumption from existing literature, Section C.1 in the appendix shows that results from the income-less-consumption approach are also in line with the wealth-based approach: there is a substantial increase in saving by the rich from the 1980s onward. Thus the rise in saving by the top 1% is a robust pattern that is not an artifact of a particular methodology or data set.[18]

Evidence from the Survey of Consumer Finances (SCF) also corroborates the main findings. Section 3.3 showed that the saving rate for the top 1% is much higher than the rest of the population in the SCF panel between 1983 and 1989. Mian et al. (2021b) use the repeated cross-section SCF+ survey to estimate saving for the top 10% and bottom 90% for each birth cohort separately. The analysis confirms that saving of the top 10% rose significantly post 1980s while falling for the rest of the population. Since the SCF+ analysis is within-birth-cohort, it also confirms that the saving glut result is not driven by demographic trends.

Our baseline analysis uses wealth shares from the Distributional National Accounts (DINA) microfiles made available by Piketty et al. (2018). We also make additional adjustments to the assumed interest rate earned by those in the top 1% of the wealth distribution (e.g., Bricker et al. (2018) and Smith et al. (2020)), and to the imputation of pension wealth given issues raised by Auten and Splinter (2019). We discuss the impact of these varying assumptions on the rise of the saving glut of the rich in Section C.3 of the appendix.

Section C.2 of the appendix cross-checks our results using wealth shares in the Distributional Financial Accounts (DFA) produced by the Federal Reserve from 1989 to 2016, when such data are available. The estimates of saving by the top 1% using the DFA and the adjusted DINA microfiles are almost identical. The similarity of the results using the DFA and DINA microfiles is consistent with the observation in Bricker and Henriques (2020) that the rise in wealth inequality from 1989 to 2019 is similar when using the SCF and when using the original Saez and Zucman (2016) methodology with the DINA microfiles.

4 What does the saving glut of the rich finance?

Section 2 unveiled the financial sector so household saving invested with intermediaries is traced down to the ultimate primary asset it finances. Section 3 estimated saving across the wealth distribution and over time using wealth data imputed from tax records. We now discuss how wealth categories in tax records can be matched with wealth categories in the flow of funds data, allowing us to link distributional saving to primary assets. This will provide important insights into the nature of expenditures that saving from different cohorts finance.

[18] Appendix Section C.5 shows that it is unlikely that population aging explains the rise in saving by the top 1%. A discussion of the implied saving rate of the top 1% under the income less consumption approach and the wealth-based approach is located in Appendix Section C.1. As we show there, the implied saving rate of the top 1% is consistent with the existing literature once missing forms of income with a high saving rate are taken into account.

4.1 Unveiling distributional saving

Recall from equation (4) that the underlying data for matrix $\bar{\mathbf{M}}_t$ used to unveil the financial system is available at the financial instrument level. These financial instruments can in turn be mapped into DINA wealth categories, which allows us to break up the $m = 1$ household sector by wealth distribution[19]. Although there is not a perfect one-to-one mapping between the Financial Accounts instruments and the asset classes in the DINA files, we are able to approximately map the Financial Accounts into the DINA asset classes using the mappings provided in Table A1.

The mapping allows us to separate out household sector into top 1% and bottom 99% cohorts, or divide up the bottom 99% into finer groups. In what follows, we maintain the same notation as in Section 2, but it should be understood that m now refers to the expanded categories of owners: the number of wealth-percentile household cohorts, federal government, state/local government, and rest of the world. With the augmented $\bar{\mathbf{M}}_t$ in hand, we repeat the unveiling process described earlier and allocate the n primary assets across the distribution of U.S. households using the DINA microfiles.

The new augmented matrix $\mathbf{A}_t$ contains information on the holdings of each of the eight primary assets, by each of the wealth-percentile cohorts such as top 1% etc. For example, $\mathbf{A}_t$ tells us how much of other households' debt is owned by the top 1% on the asset-side of their balance sheet. Similarly, $\mathbf{A}_t$ tells us how the rest of the primary assets are owned across the wealth distribution (i.e. federal government debt, state/local government debt, corporate sector and real estate). The mapping of primary assets to wealth distribution is only possible due to the unveiling process described in Section 2.1. For example, without unveiling, one cannot figure out how much government debt is financed by the top 1% since most government debt is held indirectly behind the veil of intermediation.

The unveiling of primary assets at the distributional level provides new insights into the financing of household, government and corporate sectors. For example, we can impute how much of federal government debt, or "safe assets" in general, is financed by the top 1% versus the next 9%, versus the bottom 90%, and how that distribution may have shifted over time.

4.2 Cross distributional household financing

The unveiling of primary assets is very useful in measuring net household debt financing across the wealth distribution. Let net debt, $ND_{it} = A^D_{it} - D_{it}$, where A^D_{it} is the holding of household debt as an asset by group i, and D_{it} is household debt borrowed by wealth cohort i itself. A^D_{it} comes from the augmented matrix $\bar{\mathbf{A}}_\mathbf{t}$ defined above. Since debt is the primary vehicle through which households

[19] Please refer to the code made available with this draft for the full procedure.

finance their spending, net debt is the key variable needed for understanding net financing patterns across the wealth distribution over time.

Figure 8 plots the net debt to national income ratio separately for top 1% and bottom 99% of the wealth distribution. Net debt is signed such that a larger positive number reflects greater holding of debt as an asset in the household portfolio. As before, we normalize levels to their respective 1982 values. While relative net debt levels are flat for top 1% and bottom 99% between 1962 and 1984, there is a steady rise in net debt holding as an asset by the top 1% till the Great Recession of 2008. At its peak, the top 1% accumulate an additional 20 percentage points of national income in net household debt in their asset portfolio.

The situation is completely the opposite for the bottom 99% who *borrow* increasingly large amounts of net debt after mid-80s, increasing their net indebtedness by 40 percentage points of national income by the time of the Great Recession. Figure 8 also provides a breakdown of net debt trend within the bottom 99% for next 9%, 51-90 percentile group, and the bottom 50% separately (dashed lines). The decline in net debt (i.e. the increase in net borrowing) is strongest and most durable for America's "middle class" - households in the 51st to 90th percentile of wealth distribution. The bottom 50% increase their net borrowing later than 51-90, and then are forced to delever the most right after the Great Recession.

Figure 8: Net Household Debt across the Wealth Distribution

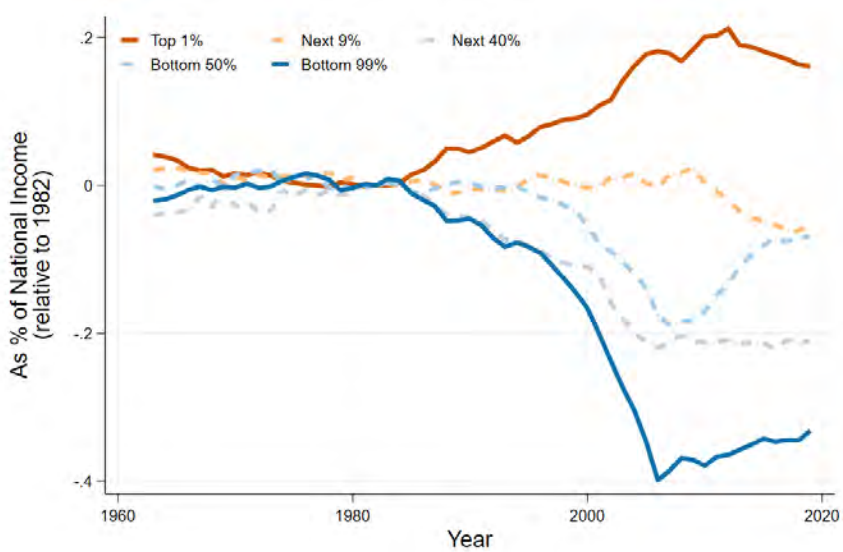

This figure shows net household debt outstanding by the U.S. household sector across the wealth distribution. Net household debt equals household debt held as an asset less household liabilities. All series are scaled by national income and expressed relative to 1982.

The key takeaway from Figure 8 is that the rising saving glut of the top 1% has been used to finance the increase in borrowing by the bottom 99%, and in particular the bottom 90%. This important finding is only possible because of the new methodology that unveils the financial sector and then allocates primary assets across the wealth distribution. Typical national accounts saving statistics aggregate the household sector and mask the dramatic shifts in saving and borrowing/lending patterns *within* the household sector.

For example, a large literature in recent decades has emphasized the importance of the "global saving glut" (e.g. Bernanke (2005)) for understanding macro trends such as the long-run decline in interest rates and the emergence of secular stagnation. Our paper uncovers an equally important saving glut within the household sector, driven by the rising split in saving rates of the very rich and the rest of the population. The macroeconomic consequences of the saving glut of the rich are likely to be very similar to the better-known global saving glut.

Figure 9 converts the stock of net debt into ten-year moving average annual flows. The annual flows reflect the net financing of household budget constraints across the wealth distribution. There has been a remarkable rise in the financing of bottom 99% household expenditures with debt since 1982. On average, the bottom 99% have financed expenditure equivalent to an additional 1.3% of national income every year between 1982 and 2019 (1.4% of national income between 1982 and 2008). To put this in perspective, the U.S. federal government's average primary deficit-to-GDP ratio was 1.1% between 1982 and 2019. The dashed line in Figure 9 shows that the top 1% have financed a majority (57.2%) of the rise in debt-financed expenditure of the bottom 99% since 1982. The shifts in aggregate demand due to debt-financed household expenditures of the bottom 99% are important for understanding the broader trends in macro-finance.

Figure 9: Flow of Debt Financing Across the Wealth Distribution

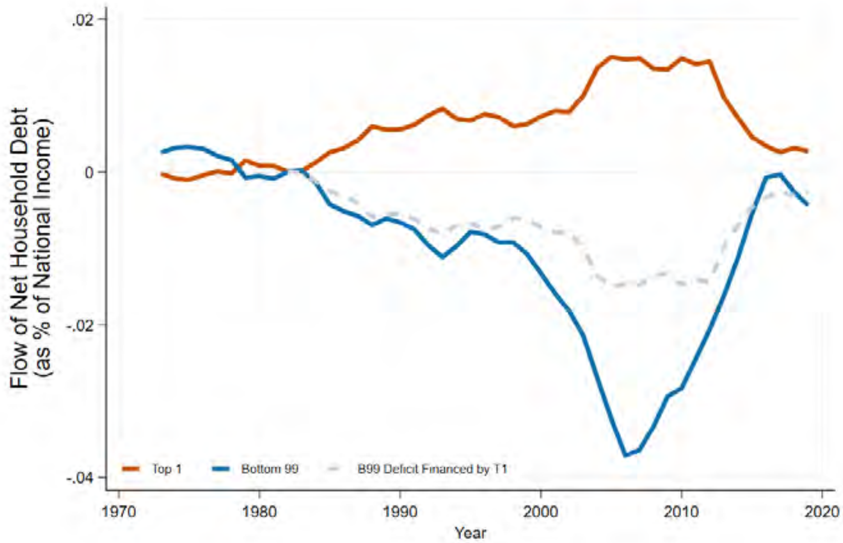

This figure shows the 10-year moving average change in net household debt for the top 1% and bottom 99% of the wealth distribution, where negative values indicate borrowing. The dashed line denotes the portion of the bottom 99%'s deficit financed by the top 1%. All series are taken as a % of National Income and expressed relative to 1982.

4.3 Saving glut of the rich and the demand for "safe assets"

Three of the unveiled primary assets include household debt, debt issued by the federal government, and debt issued by state and local governments. U.S. household debt largely includes mortgages backed by the Government-Sponsored Enterprises (GSEs), which are considered safe, just like debt issued by U.S. governments. We therefore include these three types of debt as "safe assets", and analyze the sources of demand for the rise of safe assets in the U.S. An important strand of the macro-finance literature in recent years has focused on this question.

Figure 10 plots total safe assets as a share of national income held by the rest of the world and the top 1% of U.S. households by wealth. Safe assets held by the rest of the world reflect asset accumulation due to the "global saving glut", while safe assets held by the top 1% reflect asset accumulation due to the "saving glut of the rich" identified in this paper. Prior to the early 1980s, U.S. safe assets were primarily held domestically, with the rest of the world holding a relatively small quantity of safe assets. However, starting in the 1980s, demand for U.S. safe assets from the rest of the world began to rise rapidly before stabilizing around 2014. The increase in demand for safe assets of 54.8% of national income from 1980 to 2019 reflects the well-known global saving glut.

Figure 10: Safe Asset Demand From Rest Of The World And Top 1%

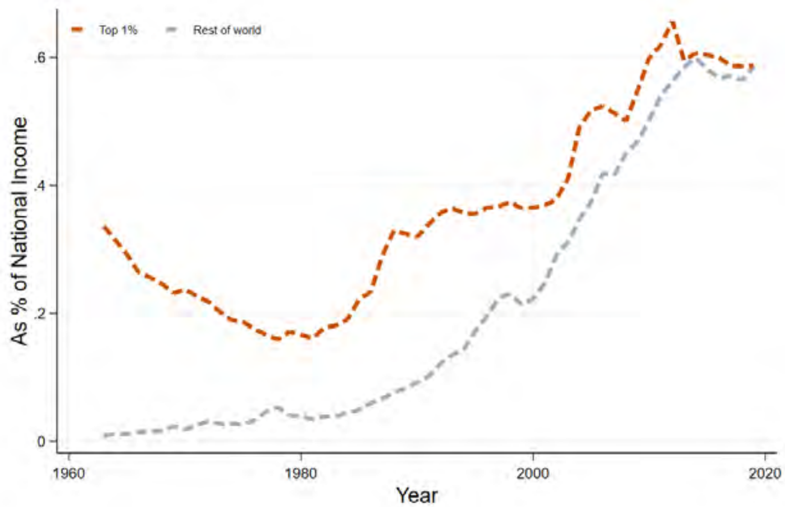

This figure shows the holdings of government and household debt by the rest of the world and by the top 1% of households, scaled by national income. These are holdings of assets after the financial sector is unveiled.

However, Figure 10 also shows a similar increase–42.1% of national income–in safe-asset demand coming from just the top 1% of U.S households. As we saw in Section 2.3, the composition of safe-asset demand is different for U.S. households, with U.S. households relatively favoring housing and state or local government debt. This likely reflects factors such as relative liquidity and tax advantages. Overall, the average increase in private saving, Θ_{it}, for the top 1% between 1980 and 2019 is 1.1% annually, while the average increase in current account deficit between 1980 and 2019 is 1.4%. The saving glut of the rich is thus of similar magnitude as the global saving glut, and as important in understanding key trends such as the persistent decline in the natural rate of interest.

4.4 What kind of expenditures did saving glut of the rich finance?

We have seen that a large share of the increase in saving coming from the top 1% went into funding expenditures of other households, or fiscal deficits. Since relatively little government expenditure is investment, these facts taken together suggest that the rise in saving of the very rich has not translated into investment, or new capital formation. In fact, while saving from the top 1% has risen significantly since the 1980s, overall investment as a share of national income has fallen in the United States.

We can use the primary assets funded by rising saving of the top 1% to estimate how much of the additional saving flow is used to fund real investment expenditure versus consumption expenditure. The unveiled primary assets map nicely into sectoral accounts that publish U.S. net investment separately by sectors such as federal government, state and local government, household (i.e. mostly real estate), and private business investment (both corporate and non-corporate). We can therefore attribute net investment to saving of the top 1% according to the particular primary assets that their saving funds. An implicit assumption in this calculation is that the marginal cash-flow is used the same way as the average cash-flow.

Figure 11 shows annual saving flow and net investment flow attributed to the top 1% of households by wealth. The sharp increase in the saving glut of the rich largely finances consumption expenditures. There is no appreciable increase in net investment done on behalf of the top 1% despite the large increase in saving by them. This should not come as a surprise given that a large share of the rising saving of the very rich was used to lend to the less-wealthy households and the government.

Figure 11: Financing of Net Investment Expenditure By Top 1% Saving

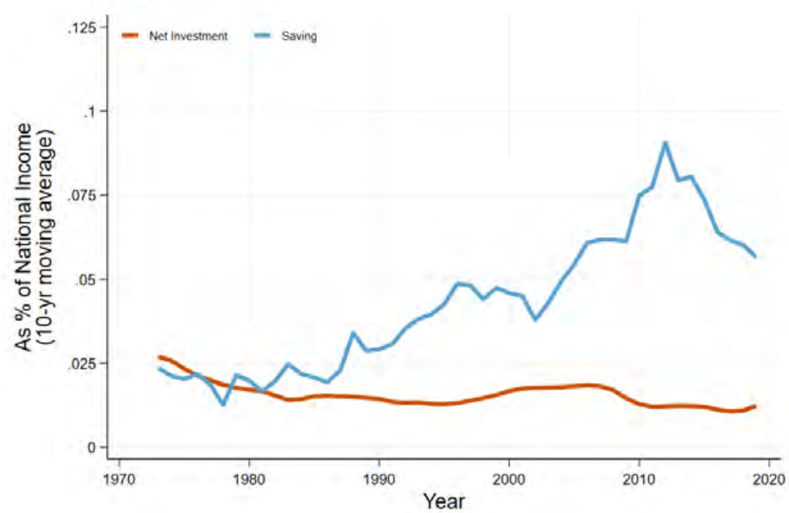

This figure plots the level of US domestic net investment attributable to saving of the top 1%, along with the total saving of the top 1%. All series are scaled by national income and expressed as a 10-year moving average.

5 External Validity Of The Saving Glut Of The Rich

This paper focused on the United States - providing a new methodology to unveil the financial sector down to primary assets, and documenting the rise of the saving glut of the rich post-1980s. The United States is naturally the most important country to perform such an analysis since it is the largest economy in the world and also its main financial hub. Nonetheless, as we discuss in this section, our results for the U.S. are also likely to hold for other major economies of the world.

5.1 International Evidence

Following up on the initial draft of this paper, Bauluz et al. (2022) unveil the financial system for China and Europe, in addition to the United States, and find broader confirmation of the key result: (i) there is a large increase in private saving for the top 10% of the wealth distribution in both China and Europe, (ii) there is a decline in saving for the rest of the distribution, and (iii) the decline in saving is cushioned by rising housing valuations. In another related paper, Acciari et al. (2024) show a significant increase in wealth inequality in Italy driven by high and rising saving rates at the top of the distribution, and declining saving rates elsewhere.

The rise of the saving glut of the rich is a global phenomenon, and one way this manifests itself is via the rise of corporate saving. Figure 12 follows the methodology in Chen et al. (2017)[20] and plots personal and corporate gross saving as a share of GDP from 1980 onward - when a consistent series is available. There is clear divergence between personal and corporate saving rates over time. Since corporate ownership is highly skewed toward the very rich globally, the rising share of corporate saving over time reflects the growing importance of saving by the rich globally. At the same time, the falling personal saving rate likely reflects increasing leverage taken on by households below the very top of the wealth distribution.

[20] We pick a reference year where all countries have saving data for the two sectors: Households and Corporations. In our sample, this is 2013. Then we regress the ratios of sector gross saving to GDP on time fixed effects, where we weigh the regressions by GDP, translated at OECD market exchange rates, and absorb country fixed effects. To benchmark the level of the each line, we calculate a mean weighted by GDP at market exchange rates over all available countries in our data in the reference year. We then add the estimated time fixed effects to extrapolate the level backwards.

Figure 12: Personal and corporate saving in OECD countries

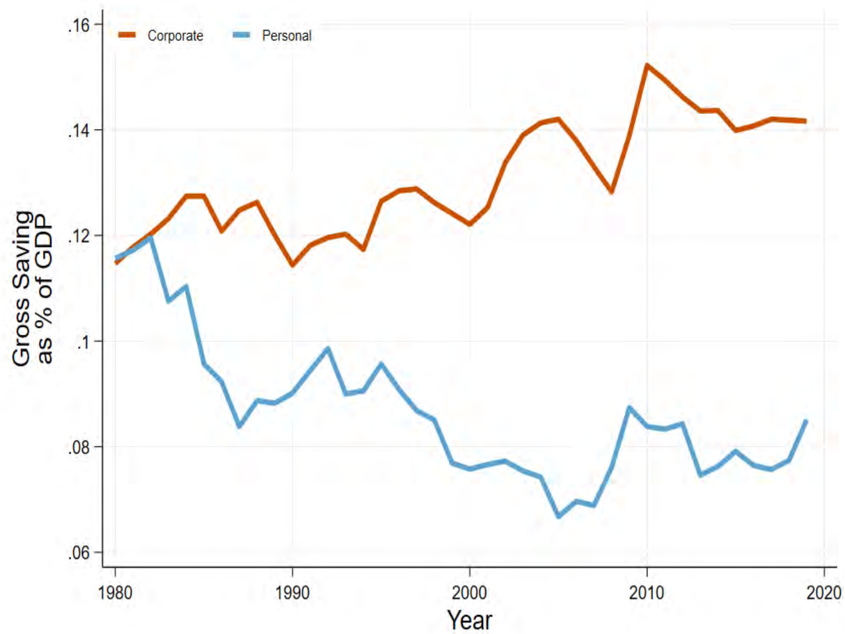

Note: This chart plots sectoral gross saving scaled by GDP across OECD countries. The series are obtained using the procedure specified in Chen et al. (2017) for an unbalanced panel.

5.2 Evidence from Norway

Norway's linked administrative records follow every household's income and balance sheet over time. Using published statistics by wealth percentile i from these data, we estimate net saving Θ_i from (7) and the associated net saving rate $s_i^N = \Theta_i / Z_i^d$. In doing so, we ensure national-accounts consistency as described in Section 3.2.

The most important component of private household saving in Norway is household saving done via corporations. On average, about 60 % of total private household saving occurs inside corporations. Since corporate equity is held almost entirely by the rich, incorporating corporate saving properly is central to understanding top-end saving behavior. An important reason for the high share of household saving done via corporations is tax law. For example, the "shareholder income tax" law that became operational from January 2006 (Alstadsaeter et al. (2019)), makes retained earnings tax-privileged. The law thus incentivized the rich to save via corporations, and correspondingly average net corporate saving went from 3.8% as a share of disposable income between 1995-2004, to 8.1% in the subsequent ten years.

Net saving rate using Iacono and Palagi (2023)

We use data from Iacono and Palagi (2023), who publish percentile-level income and wealth statistics for 2011–2018 to estimate household saving via corporations by wealth percentile. Figure D.1 in the appendix plots debt-to-income ratios, asset composition, and income and wealth shares by wealth percentile in Norway. Housing is the dominant asset class, and is typically highly leveraged, especially for the bottom quartile that has negative net wealth on average. Importantly, income and wealth are highly concentrated at the top, especially wealth. The top 1 %, 5 %, and 10 % receive 7.7 %, 16.0 %, and 23.8 % of total income, but they own 37.8 %, 54.1 %, and 63.2 % of financial wealth, respectively.

How much saving is done by each wealth percentile of households in Norway via corporations? We use equation (13) to calculate the total corporate saving done on behalf of households, S^{π}, from Norwegian national accounts. Since corporate saving is done on behalf of equity holders, and we know how financial assets are distributed across the wealth distribution, we can allocate S^{π} to wealth percentiles according to the share of financial wealth held by each percentile[21]. The solid black line in figure 13 (a) plots total corporate saving of each wealth percentile divided by its total disposable income, Z_i^d. This represents household saving rate if we only include saving that is done via corporations. The figure shows that corporate saving alone raises the saving rate to 39.3 % for the top 1 % and 20.3 % for the next 1 %. The economy-wide corporate saving rate, $\frac{S^{\pi}}{Z^d}$, is 8.1%.

The black line in figure 13 (a) shows the quantitative importance of corporate saving for understanding household saving behavior across the wealth distribution. For example, saving via corporations alone raises the saving rate of the wealthiest 1% by over 30 percentage points relative to the rest of population. The remaining saving of households is personal saving, S^p. The economy-wide personal saving rate, $\frac{S^p}{Z^d}$, is 6.9%. This is saving out of the "in-hand" income that households receive (i.e. wages, interest income and dividend income). Total private net saving of the household sector in Norway between 2011-2018 is thus, $6.9\% + 8.1\% = 15.0\%$. Household saving done via corporations is larger in magnitude than personal saving out of in-hand income. Since we do not observe personal saving at the disaggregated level, we cannot estimate total saving rate for households at the wealth-percentile level. However, the solid blue line in figure 13 (a) plots total saving rate if personal saving rate were constant across the wealth, i.e. it shifts the black line up by 6.9%. The broader takeaway from figure 13 (a) is the central role of corporate saving in raising saving rate significantly at the top-end of the wealth distribution.

Net saving rate using Fagereng et al. (2025)

Fagereng et al. (2025) use the same administrative data as Iacono and Palagi (2023) over the period

[21]This assumes that equity share within financial assets is the same across wealth distribution. It is likely that wealthier households have higher equity share. Thus our approach should be conservative in estimating saving via corporations for the richest households.

Figure 13: Net saving rate across the wealth distribution in Norway

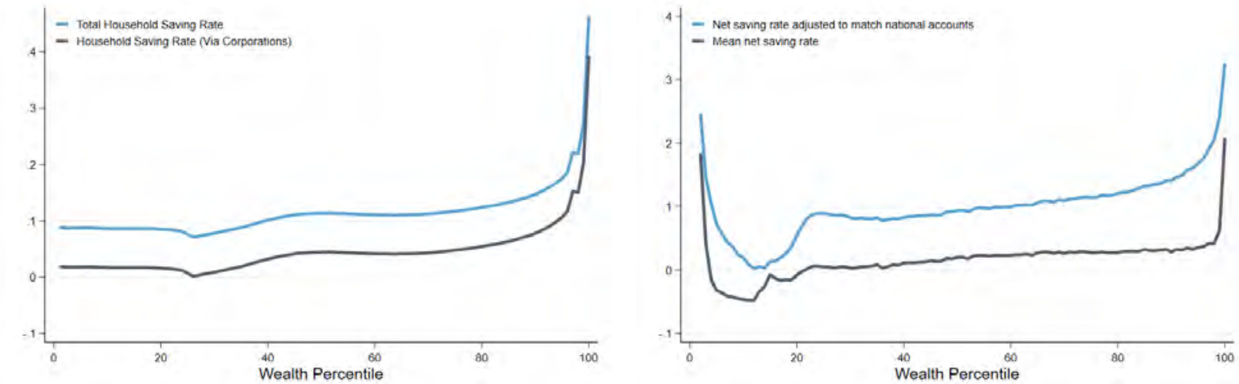

(a) Mean Net Saving Rate (Iacono and Palagi (2023))

(b) Mean Net Saving Rate (Fagereng et al. (2025))

2006–2015, and estimate net saving at the individual level using change in net wealth and removing imputed capital gains according to equation (7). They provide estimates for mean and median net saving rate, s_i^N, across the wealth distribution. The solid black line in figure 13 (b) plots their mean net saving rate against the wealth distribution from their June 2025 working paper. In line with our earlier results, the mean net saving rate is generally upward sloping over the wealth distribution[22] and rises rapidly at the top-end of the wealth distribution.

However, Fagereng et al. (2025) follow a different methodology than ours. Their focus is on building up net and gross saving rates from Norwegian micro data, but without strictly enforcing the national accounting relationships we have emphasized throughout this paper. To show that this can make a difference, we aggregate the black line in figure 13 (b) using income weights by wealth percentile from Iacono and Palagi (2023). This yields 3.7%, while the household net saving rate (inclusive of saving via corporations) is 13.1% over 2006-2015 in Norway's national accounts.[23]

Effectively, it seems that a relatively small share of the change in household net wealth (gross saving) is attributed to net saving, and instead attributed to capital gains. To explore a combination of the micro data approach of Fagereng et al. (2025) with our approach based on national accounting consistency, we uniformly scale down the capital gains component across wealth percentiles, as represented by the blue line in Figure 13(b), such that the resulting overall average net saving rate aligns with the one we measured based on the national accounts (see Appendix D.3 for details).

[22] The non-monotonicity at the bottom of the wealth distribution is likely driven by relatively high gross wealth but negative net worth households due to high leverage, as seen in figure D.1 in the appendix

[23] There is a discrepancy in Fagereng et al. (2025) between the aggregate mean net saving rate implied by Figure 8(a), i.e. 3.7%, and the mean net saving rate of 9.25% (=6,231/67,743) suggested by the summary statistics Table 1 that we are currently unable to resolve. At the moment, only the top one and bottom one percentiles in Figure 8(a) have a net saving rate above 9.25%.

This adjustment naturally elevates the net saving rates across all cohorts, with somewhat larger increases occurring among the wealthier percentiles. The adjusted saving rates between the 20th and 80th percentiles remain relatively flat, consistent with the findings reported by Fagereng et al. (2025). Of course, this adjustment is somewhat crude and more of a proof-of-concept. We believe that a more careful combination of the methodologies used in Fagereng et al. (2025) and in this paper would be a fruitful area of future research.

Taking stock, we view the evidence from Norway and the other countries mentioned in Section 5.1 as being consistent with an average net saving rate that is significantly greater at the top end of the wealth distribution. We therefore suspect that shifts in inequality happening outside the U.S. may also translate into a *saving glut of the rich*, much like it did in the U.S.

6 Concluding Remarks

This paper introduces a new *unveiling methodology* to trace household saving through the financial system and identify its ultimate use. Applying this approach to U.S. tax records from 1963 to 2019, we document the emergence of the *saving glut of the rich*—a sharp rise in saving by the top 1% since the 1980s. This glut has financed household borrowing before 2008 and government deficits thereafter, rather than productive capital investment. These findings suggest that rising inequality has fundamentally reshaped financial intermediation, fueling demand for safe assets and contributing to the secular decline in interest rates.

Our results raise important questions about the long-run relationship between inequality, financial markets, and macroeconomic stability. A key avenue for future research is understanding how the financial sector intermediates this excess saving. The expansion of financial intermediation has not necessarily translated into increased productive investment, but rather into increased leverage among middle-class households and the government. Examining the incentives and constraints that drive this allocation of saving is crucial for assessing the implications of potential imbalances in the economy.

Another important direction for future research is the global dimension of the saving glut of the rich. While our analysis focuses on the U.S., evidence suggests similar trends in other advanced economies, where saving at the top has surged while middle-class borrowing has expanded. Investigating cross-country variations in saving gluts and their implications for global capital flows, exchange rates, and financial stability can provide deeper insights into the evolving nature of financial intermediation in an era of rising inequality.

References

Acciari, Paolo, Facundo Alvaredo, and Salvatore Morelli, "The Concentration of Personal Wealth in Italy 1995–2016," *Journal of the European Economic Association*, June 2024, *22* (3), 1228–1274.

Aguiar, Mark and Mark Bils, "Has Consumption Inequality Mirrored Income Inequality?," *American Economic Review*, 2015, *105* (9), 2725–56.

Alstadsaeter, Annette, Wojciech Kopczuk, and Kjetil Telle, "Social Networks and Tax Avoidance: Evidence from a Well-Defined Norwegian Tax Reform," *International Tax and Public Finance*, 2019, *26* (6), 1291–1322.

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

Auten, Gerald and David Splinter, "Income Inequality in the United States: Using Tax Data to Measure Long-Term Trends," 2019.

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

Batty, Michael M., Elizabeth Holmquist, and Robert J. Kurtzman, "From-Whom-to-Whom Relationships in the Financial Accounts of the United States: A New Methodology and Some Early Results," FEDS Notes, Board of Governors of the Federal Reserve System March 24 2023.

Bauluz, Luis, Filip Novokmet, and Moritz Schularick, "The Anatomy of the Global Saving Glut," *ECONtribute Discussion Papers Series*, April 2022, (161).

BEA, "U.S. International Economic Accounts: Concepts and Methods," 2021.

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.

Bricker, Jesse, Alice Henriques, and Peter Hansen, "How Much Has Wealth Concentration Grown in the United States? A Re-Examination of Data from 2001-2013," 2018.

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

Caballero, Ricardo J and Arvind Krishnamurthy, "Excessive Dollar Debt: Financial Development and Underinsurance," *The Journal of Finance*, 2003, *58* (2), 867–893.

_ and _ , "Global Imbalances and Financial Fragility," *American Economic Review*, 2009, *99* (2), 584–88.

Caballero, Ricardo J., Emmanuel Farhi, and Pierre-Olivier Gourinchas, "An Equilibrium Model of "Global Imbalances" and Low Interest Rates," *American Economic Review*, March 2008, *98* (1), 358–393.

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

Chen, Peter, Loukas Karabarbounis, and Brent Neiman, "The Global Rise of Corporate Saving," *Journal of Monetary Economics*, August 2017, *89*, 1–19.

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

Fagereng, Andreas, Martin B. Holm, Benjamin Moll, and Gisle Natvik, "Saving Behavior Across the Wealth Distribution: The Importance of Capital Gains," Working Paper, BI Norwegian Business School and University of Oslo and London School of Economics jun 2025.

_ , Martin Blomhoff Holm, Benjamin Moll, and Gisle Natvik, "Saving Behavior Across the Wealth Distribution: The Importance of Capital Gains," NBER Working Paper 26588, National Bureau of Economic Research December 2019.

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

Gourinchas, Pierre-Olivier and Helene Rey, "From World Banker to World Venture Capitalist: US External Adjustment and the Exorbitant Privilege," in "G7 Current Account Imbalances: Sustainability and Adjustment," University of Chicago Press, 2007, pp. 11–66.

Greenwood, Robin, Samuel G Hanson, and Gordon Y Liao, "Asset Price Dynamics in Partially Segmented Markets," *The Review of Financial Studies*, 2018, *31* (9), 3307–3343.

Guvenen, Fatih, Greg Kaplan, Jae Song, and Justin Weidner, "Lifetime Incomes in the United States over Six Decades," Technical Report, National Bureau of Economic Research 2019.

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

Iacono, Roberto and Elisa Palagi, "A Micro Perspective on r > g," *Economica*, February 2023, *90* (358), 531–556.

Klein, Matthew C and Michael Pettis, *Trade Wars Are Class Wars: How Rising Inequality Distorts the Global Economy and Threatens International Peace*, Yale University Press, 2020.

Koijen, Ralph S. J. and Motohiro Yogo, "A Demand System Approach to Asset Pricing," *Journal of Political Economy*, August 2019, *127* (4), 1475–1515.

Kuhn, Moritz, Moritz Schularick, and Ulrike I Steins, "Income and Wealth Inequality in America," *Journal of Political Economy*, 2019.

_ , _ , **and Ulrike I. Steins**, "Income and Wealth Inequality in America, 1949–2016," *Journal of Political Economy*, September 2020, *128* (9), 3469–3519.

Kumhof, Michael, Romain Rancière, and Pablo Winant, "Inequality, Leverage, and Crises," *American Economic Review*, March 2015, *105* (3), 1217–1245.

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, Ludwig Straub, and Amir Sufi, "Indebted Demand," *The Quarterly Journal of Economics*, November 2021, *136* (4), 2243–2307.

_ , _ , **and** _ , "What Explains the Decline in R*? Rising Income Inequality Versus Demographic Shifts," *Working Paper*, 2021.

_ , _ , **and** _ , "A Goldilocks Theory of Fiscal Deficits," Working Paper w29707, National Bureau of Economic Research February 2024.

Philippon, Thomas, "Has the US Finance Industry Become Less Efficient? On the Theory and Measurement of Financial Intermediation," *American Economic Review*, April 2015, *105* (4), 1408–1438.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman, "Distributional National Accounts: Methods and Estimates for the United States," *The Quarterly Journal of Economics*, May 2018, *133* (2), 553–609.

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*, May 2016, *131* (2), 519–578.

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

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

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

UN, "System of National Accounts," 2008.

Appendix for Online Publication

A Appendix Tables and Figures

Figure A1: Average Intermediation Chain Length

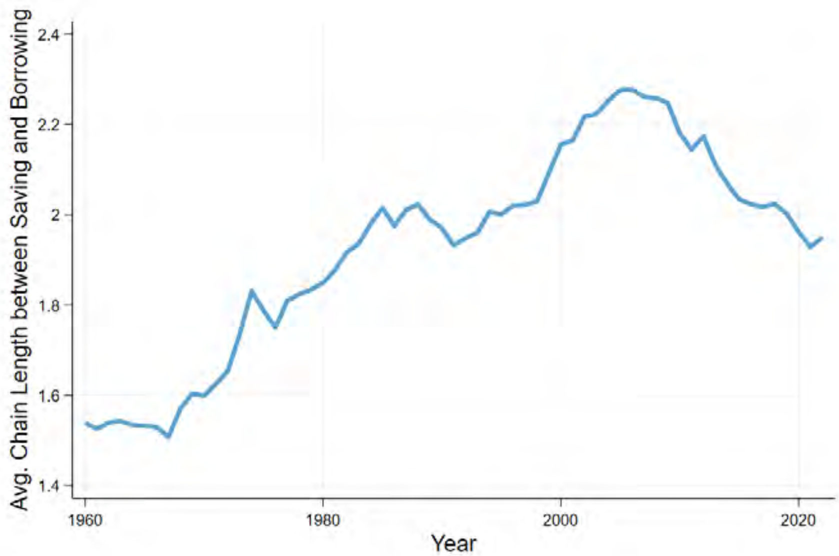

This figure plots the average number of intermediaries involved in channeling household saving to the debt it finances.

Figure A3: Real Estate and Corporate Business Q

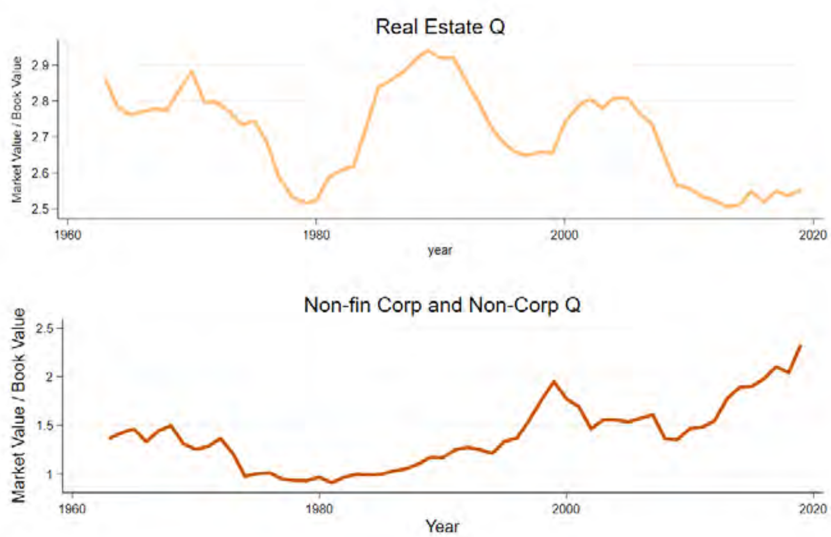

This figure plots the ratio of market value to book value for real estate and corporate business assets.

Figure A5: Personal Saving Decomposed

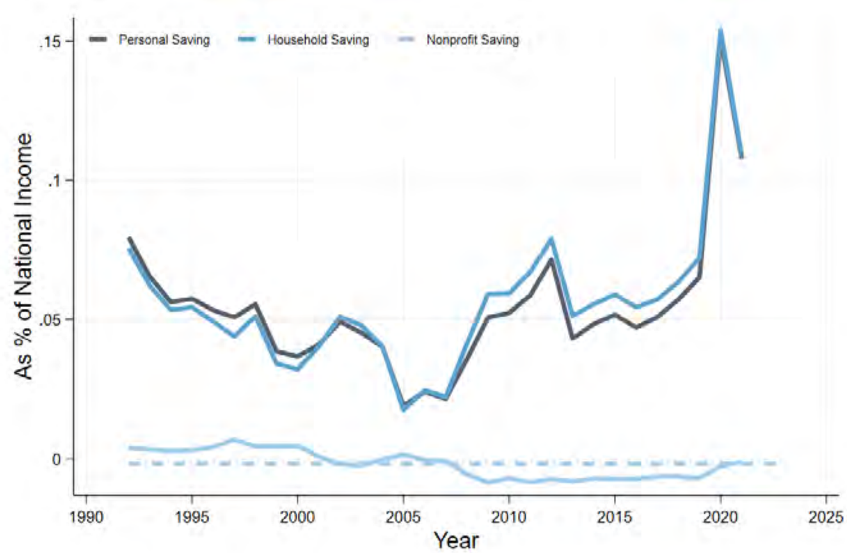

The figure illustrates the decomposition of personal saving into household and nonprofit saving. We can further decompose the national accounting identity $\Theta = S^p + S^\pi$ as $\Theta = S^{HH} + S^{NP} + S^\pi$, where HH and NP stand for households and nonprofit organizations, respectively. The dotted line denotes average nonprofit saving across time. All series are sourced from NIPA and scaled by national income.

Figure A2: The Financial Network

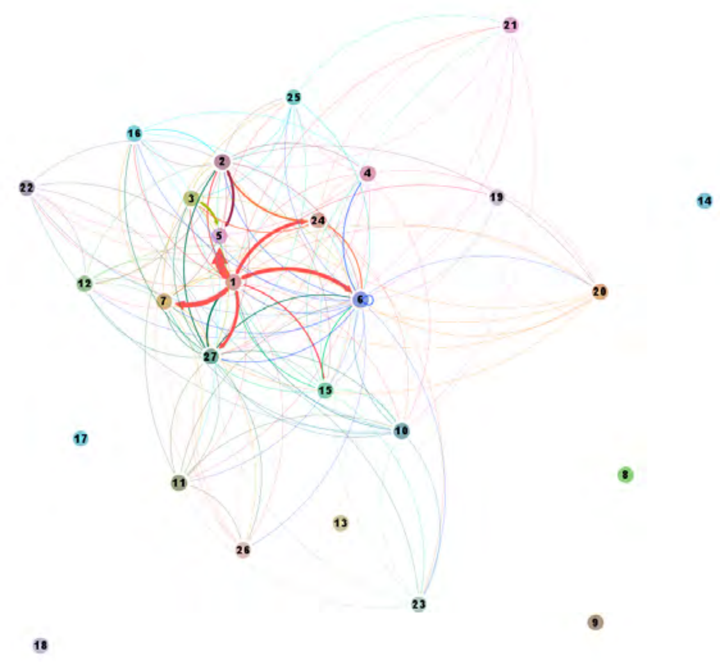

(a) 1960

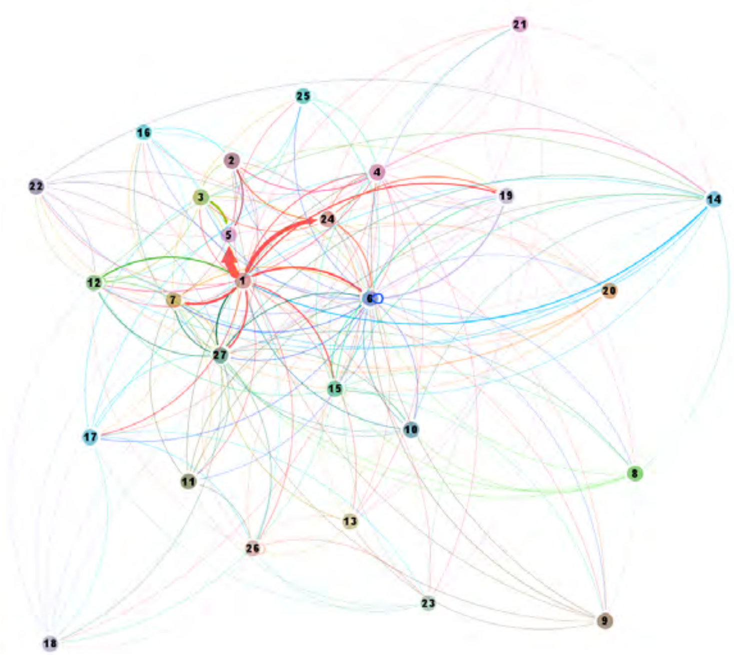

(b) 2007

This figure displays the most notable direct holdings of liabilities across the financial network in 1960 and 2007. The thickness of the edges denotes the magnitude of the holdings. Arrows point in the direction of the flow of funds. Nodes are numbered corresponding to the classification of the 23 intermediaries and 4 ultimate asset owners in Section 2.1.

Figure A4: Financial assets held by non-financial firms

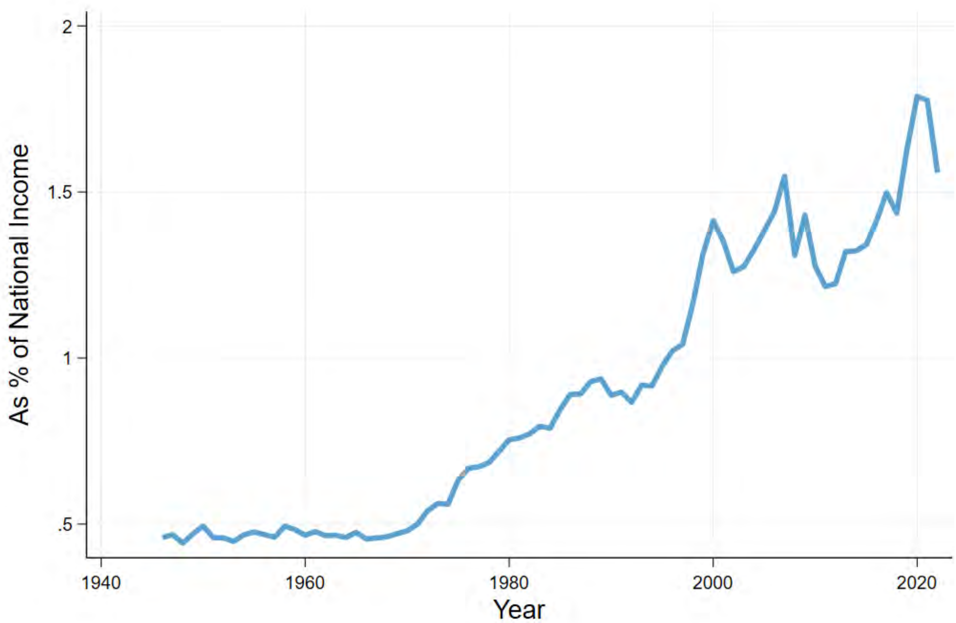

Note: The graph illustrates the evolution of financial assets held by non-financial corporate and non-corporate firms, scaled by national income. Data is sourced from the Flow of Funds dataset provided by the Federal Reserve.

Table A1: Mapping for Wealth-Based Approach

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

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

Figure A6: Comparing π^{equity} with annual S&P 500 returns

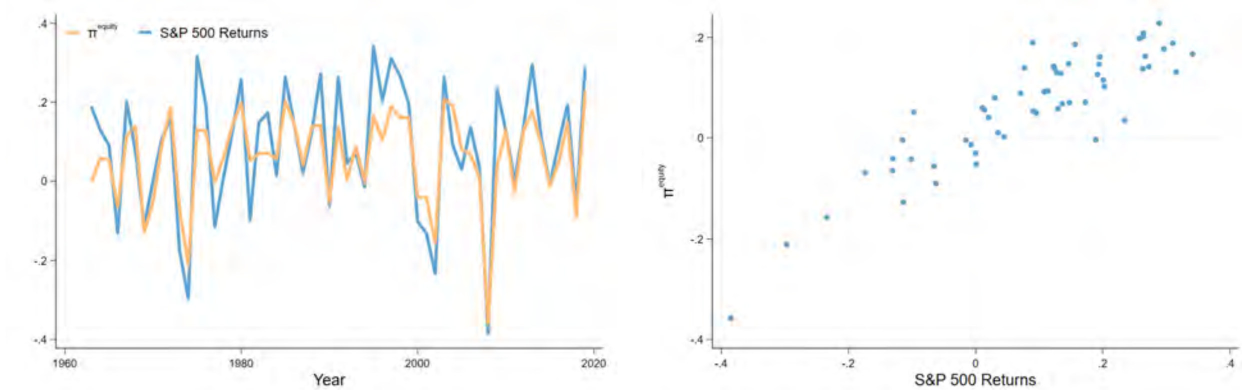

This figure compares the asset inflation rate for equities that we obtain in the wealth-based approach with annual S&P 500 returns. We find their correlation is 0.89. S&P 500 data comes from the Macrotrends website.

Figure A7: Saving Rate & Capital Gain Using Repeated Cross-Section (DINA)

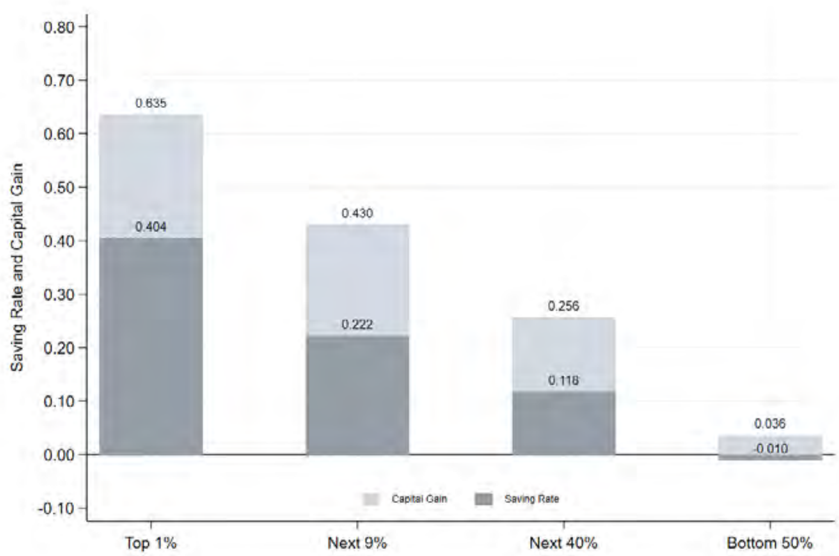

This figure depicts net saving rates across the U.S. household distribution over the period 1983-1989 if capital gains are included (from the baseline methodology).

Figure A8: Saving rate of various wealth cohorts over time

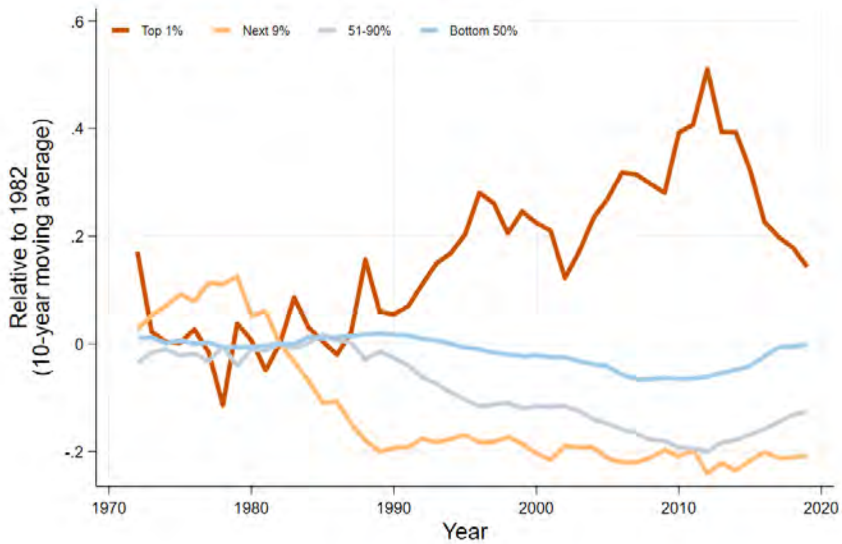

Note: This figure shows the ten-year moving average saving rate across different wealth cohorts. Saving was constructed using the wealth-based approach, and disposable income is sourced from the DINA microfiles dataset produced by Piketty et al. (2018). All series are expressed relative to 1982, when the saving rates for the top 1%, next 9%, next 40%, and bottom 50% were 0.34, 0.42, 0.17, and -0.02, respectively.

Table A2: Decomposing Saving

Period	Θ	Δ NW	Δ V	Θ^{FA}	Θ^{RE}	Δ D	ΔV^{RE}
				Top 1%			
63-82	0.026	0.047	0.021	0.025	0.002	-0.001	0.003
83-97	0.037	0.082	0.044	0.036	0.004	-0.003	0.004
98-07	0.062	0.117	0.055	0.062	0.004	-0.004	0.009
08-19	0.058	0.096	0.039	0.056	0.001	0.000	0.002
				Next 9%			
63-82	0.065	0.102	0.036	0.062	0.010	-0.006	0.013
83-97	0.040	0.094	0.055	0.037	0.010	-0.008	0.014
98-07	0.034	0.100	0.066	0.034	0.012	-0.012	0.030
08-19	0.040	0.069	0.029	0.040	0.003	-0.003	0.008
				Next 40%			
63-82	0.050	0.093	0.042	0.053	0.020	-0.022	0.030
83-97	0.032	0.092	0.060	0.043	0.018	-0.029	0.026
98-07	-0.001	0.088	0.088	0.022	0.017	-0.040	0.054
08-19	0.009	0.043	0.034	0.021	0.001	-0.013	0.013
				Bottom 50%			
63-82	-0.005	0.004	0.009	0.008	0.005	-0.018	0.007
83-97	-0.007	0.004	0.011	0.005	0.003	-0.016	0.006
98-07	-0.021	-0.001	0.020	0.005	0.009	-0.035	0.013
08-19	-0.007	0.003	0.010	0.002	-0.002	-0.007	0.002

This table decomposes the average annual saving by the top 1%, next 9% , next 40%, and bottom 50% of the wealth distribution. Saving can be decomposed according to equations $\Theta = \Delta\text{NW} - \Delta\text{V}$ and $\Theta = \Theta^{FA} + \Theta^{RE} + \Delta\text{D}$, where ΔNW signifies the change in net household wealth, ΔV represents the valuation effect, Θ^{FA} indicates saving through financial assets, Θ^{RE} denotes saving through real estate, and ΔD marks the change in household debt. V^{RE} represents the valuation effect in real estate assets. All measures are scaled by national income, and then the averages for the period are calculated.

B Mapping to UN System of National Accounts

The underbraces indicate the System of National Accounts (UN, 2008) codes assigned to the variables used in this paper. The mapping is meant to help future research in calculating distributional

Figure A9: Saving and Valuation Effects

(a) Top 1% (b) Bottom 99%

This figure plots the change in net wealth by the top 1% and bottom 99% of households, decomposed into saving and valuation effects. All series are scaled by national income and taken relative to 1982.

saving in a manner that is consistent with national accounts. We start with equation for GDP, Y,

$$\underbrace{Y}_{B1g} = \underbrace{C}_{P31S14\_15} + \underbrace{G}_{P3S13} + \underbrace{I}_{P51+P52+P53} + \underbrace{(X - M)}_{P6-P7}$$

where C is final consumption expenditure of households and NPISH[24], G refers to final consumption expenditure of general government, I refers to gross capital formation[25], and lastly X are exports and M are imports.

It is useful to subtract the consumption of fixed capital δ as it reflects accounting depreciation of capital goods, and focus on net variables.[26] Net investment, I^n, and net national income, Z, can be expressed as

$$\underbrace{I^n}_{P5n} = I - \underbrace{\delta}_{P51c/K1}$$

$$\underbrace{Z}_{B5n} = \underbrace{Y}_{B1g} - \underbrace{\delta}_{P51c} + \underbrace{U_{Row}}_{D1 \to D4S2}$$

where U_{RoW} is net income with the rest of the world, RoW.[27] Net national disposable income is,

[24]NPISH stands for "non-profit institutions serving households". These are often bundled with households since they do not provide goods and services at market prices for commercial purposes.

[25]This includes gross fixed capital formation, changes in inventories and acquisitions less disposals of valuables.

[26]The SNA code for consumption of fixed capital is P51c but it is also found under the code K1 in the OECD datasets.

[27]Net income with RoW is constituted of labor income, referred to as Compensation of employees $D1$, capital income, referred to as Property income $D4$, and taxes $D2$ less subsidies $D3$ on production and imports.

$$\underbrace{Z^d}_{B6n} = \underbrace{Z}_{B5n} - \underbrace{T_{RoW}}_{D5S2} + \underbrace{R_{RoW}}_{D61 \to D62S2 + D7S2}$$

where T_{RoW} denotes net taxes and R_{RoW} denotes net social contributions (D61), which can be decomposed as social benefits other than social transfers in kind (D62) and other current transfers (D7) with the RoW. R_{RoW} does not contain social transfers in kind (SNA code D63). If we choose to include these transfers, we refer to the resulting balance as net adjusted disposable income[28]:

$$\underbrace{Z^d_{adj}}_{B7n} = \underbrace{Z^d}_{B6n} + \underbrace{R^{kind}_{RoW}}_{D63S2}$$

Subtracting the final expenditures from net disposable income gives us aggregate net saving, S,

$$\underbrace{S}_{B8n} = \underbrace{Z^d}_{B6n} - \underbrace{C}_{P31S14\_15} - \underbrace{G}_{P3S13}$$

The aggregate net saving is used for domestic net investment, I^n, and the current account balance, F. A positive balance means that there is an excess amount of saving relative to net investment in the domestic economy.

$$\underbrace{S}_{B8n} = \underbrace{I^n}_{P5n} + \underbrace{F}_{B12}$$

$$\underbrace{F}_{B12} = \underbrace{X - M}_{P6 - P7} + W$$

where W is the balance of primary and secondary incomes with the RoW,

$$W \equiv \underbrace{U_{Row}}_{D1 \to D4S2} - \underbrace{T_{RoW}}_{D5S2} + \underbrace{R_{RoW}}_{D61 \to D62S2 + D7S2}$$

The current account balance does not strictly equate net lending / borrowing, B, in an economy, which can be written as,

$$\underbrace{B}_{B9} = \underbrace{F}_{B12} + \underbrace{R^K_{RoW}}_{D9} - \underbrace{K_{RoW}}_{NP/K2}$$

where R^K_{RoW} is net capital transfers and K_{RoW} is acquisitions less disposals of non-produced non-

[28]In the OECD National Accounts data, the RoW sector does not include social transfers in kind. Hence, at the aggregate level, net disposable income and net disposable adjusted income are the same. This data exists at the sectoral level.

financial assets.[29] The net lending / borrowing equals the balance from the financial account.

Sectoral identities SNA also provides sub-sector level breakdown of all accounting variables. There are five domestic sectors: non-financial (S11) and financial (S12) corporations, general government (S13), households (S14), non-profit institutions serving households or NPISH (S15), and one foreign sector / RoW (S2). The aggregate economy is given the code S1. As is common practice, we will bundle the household and NPISH sectors into one, and refer to it as "the household sector" below. Each important variable at the national level is also recorded at the sectoral level, and these sectoral totals add up to national aggregates. For example, disposable income at the national level equals:

$$\underbrace{Z^d}_{B6n} = \underbrace{Z^{d,p}}_{B6nS14\_15} + \underbrace{Z^{d,\pi}}_{B6nS11\_12} + \underbrace{Z^{d,g}}_{B6nS13}$$

Net aggregate saving is broken down at the sectoral level as,

$$S = S^p + S^\pi + S^g$$

where sectoral level personal, corporate and government net saving are given by,

$$\underbrace{S^p}_{B8nS14\_15} = \underbrace{Z^{d,p}}_{B6nS14\_15} + \underbrace{P_{adj}}_{D8} - \underbrace{C}_{P31S14\_15}$$

$$\underbrace{S^\pi}_{B8nS11\_12} = \underbrace{Z^{d,\pi}}_{B6nS11\_12} - \underbrace{P_{adj}}_{D8}$$

$$\underbrace{S^g}_{B8nS13} = \underbrace{Z^{d,g}}_{B6nS13} - \underbrace{G}_{P3S13}$$

where P_{adj} is the adjustment for the change in pension entitlements, and very small in practice. Thus corporate saving essentially equals corporate disposable income. Total private saving, $\Theta \equiv S^p + S^\pi$, for most countries where government does not own significant corporate assets.

C Robustness to alternative data and methodology

C.1 Income less consumption approach to measuring saving

This study measures Θ_{it}, which is the saving by group i of households in year t. By construction, $\sum_{i \in I} \Theta_{it} = \Theta_t$. The income-less-consumption approach to estimating Θ_{it} uses the following

[29] The capital account captures all the entries that do not affect measures of production, income and saving (BEA, 2021).

definition:

$$\Theta_{it} = Z_{it} - T_{ti} + R_{it} - C_{it} \tag{12}$$

where Θ_{it} is pre-tax income minus taxes plus transfers minus consumption. Starting from equation 12, we use the DINA microfiles to measure the first three terms that define Θ_{it}: Z_{it}, T_{it}, and R_{it}. When constructing these variables using DINA, we use the adult individual as the unit of observation and split income equally among spouses.[30]

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

The analysis below focuses on three main groups: the top 1% of the income distribution, the next 9%, and the bottom 90%. Unfortunately, survey datasets, typically used in the consumption literature such as the Consumer Expenditure Survey (CEX) and the Panel Study of Income Dynamics (PSID), do not measure the consumption of the highest-income households in the economy accurately.[31]

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

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

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

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

reported in the SCF as the consumption share for all goods across the income distribution.[32]

The average consumption share of the top 1% of the income distribution, as measured by Fisher et al. (2018), is 6.6% for the period from 2004 to 2016; we use this value for 2010, the midpoint of the available SCF data. How has this consumption share evolved over time? This is the most challenging variable to measure, as there is no dataset that covers the consumption of rich U.S. households over the long time period of our sample, which is 1963 to 2019.

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

The assumption of a constant consumption-to-income ratio over time is conservative, because it is likely that the consumption-to-income ratio falls with income even over long time periods (e.g., Straub (2019)).[33] If the consumption to income ratio of the top 1% has fallen over time, then the saving glut of the rich would be estimated to be even larger.[34] Figure C.1 plots the consumption share of the top 1% of the income distribution using this methodology. As it shows, the consumption share of the top 1% using this methodology has risen substantially over time, from 4 to 5% in the 1960s and 1970s, to 6 to 7% from 2010 to 2016.

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

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

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

Figure C.1: Consumption Share of the Top 1% of the Income Distribution

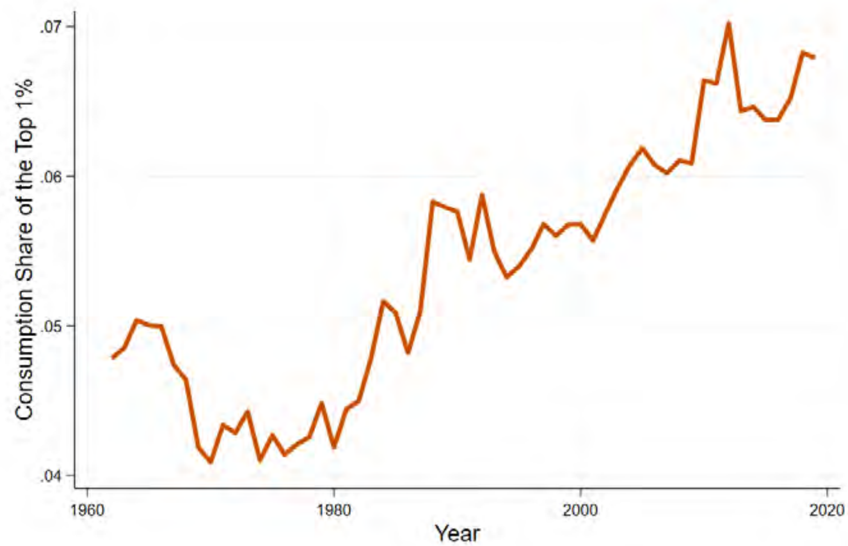

The average consumption share of the top 1% from 2004 to 2019 comes from Fisher et al. (2018). This average consumption share is used as the baseline in the year 2010. The time series is then generated using the assumption of a constant consumption to income ratio.

Figure C.2 shows average annual saving by the top 1% using the income less consumption approach with the DINA microfiles. As was the case with the wealth-based approach to measuring saving, there is a significant rise in saving of the top 1% post-1980, and a large decline in saving of the bottom 90%.

Figure C.2: Saving across the Distribution - Income less consumption approach

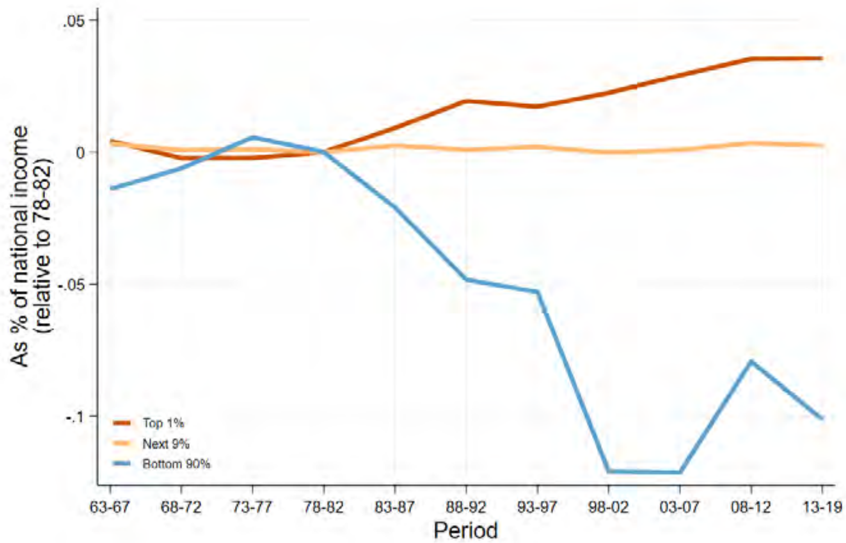

This figure illustrates the saving across the wealth distribution measured using the income less consumption approach, where income shares come from the Distributional National Accounts. Annual saving are scaled by annual national income, and five year averages are plotted. The 1978 to 1982 period is subtracted for all series.

C.2 Wealth-based saving estimate using DFA

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

Nonetheless, it is possible to compare the levels of saving from 1989 to 2016 implied by the DFA and DINA wealth shares, which is done in Table C.1. As before, annual saving is scaled by contemporaneous national income, and the averages for 1990 to 2016 are shown. For the sake of comparison, Table C.1 also shows averages using the DINA wealth shares for the same period in question.

Table C.1: Results using DFA: 1990 to 2016

Wealth Cohort	DINA	DFA
Top 1%	0.048	0.047
Next 9%	0.042	0.045
Bottom 90%	-0.001	-0.003

This table depicts measures of average annual saving by wealth cohort across 1990 to 2016 using DFA wealth shares (the period for which the DFA data are available). It also shows average annual saving using the DINA wealth shares over the same time period. Annual saving is scaled by contemporaneous national income before averaging across time.

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

The wealth-based approach for estimating saving uses data from the Distributional National Accounts (DINA) microfiles. These microfiles rely on the yearly public-use tax return files available at the National Bureau of Economic Research, along with calculations to allocate national income that is not included as part of tax returns, as outlined by Piketty et al. (2018). The DINA microfiles are not only used to calculate shares of national income across the income distribution, but also to calculate wealth shares across the wealth distribution using the capitalization technique outlined in Saez and Zucman (2016) and Smith et al. (2020).

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

However, Auten and Splinter (2019) argue that the non-taxable pension income reported on tax filings are actually rollovers from Individual Retirement Accounts, and are therefore not income but a rollover of wealth. As a result, the pension income of higher income Americans is overestimated if one uses the non-taxable part of the pension income reported on tax filings. We follow Auten and Splinter (2019) and use only the taxable pension income reported on tax filings to estimate the undistributed pension income component of national income and pension wealth. This change affects both the distribution of national income and wealth, given that pension wealth is obtained by capitalizing the income earned on pensions. As shown in Auten and Splinter (2019), this adjustment brings the estimate of pension wealth across the income distribution in the DINA closer to the distribution in the SCF.[35]

The second adjustment relates to the capitalization of fixed income on tax returns to measure fixed income wealth across the wealth distribution (e.g, Bricker et al. (2018) and Smith et al. (2020)). Translating flows of income into stocks of wealth requires an assumption of the rate of return on assets, a process detailed in Saez and Zucman (2016). Research suggests that the baseline methodology in Saez and Zucman (2016) over-states the level of fixed income asset holdings of the top 1% given the assumption of a homogeneous rate of return on fixed income assets when estimating fixed income wealth from fixed income asset cash flows (e.g., Bricker et al. (2018) and Smith et al. (2020)).[36] This manifests itself in the assumed capitalization factor one uses to multiply the fixed income asset cash flows to obtain fixed income wealth.

Given this issue, the capitalization technique used in this study assumes that the top 1% of the wealth distribution earns a return on fixed income assets that is 1.3 times (1.3X) the overall rate of return on fixed income assets earned by U.S. households. This assumption is motivated by evidence from the SCF and estate tax filings in Bricker et al. (2018) and Saez and Zucman (2020). Appendix Figure C.3 uses the SCF and shows the ratio of the top 1% of the wealth distribution's fixed income return to the fixed income return of all households from 2001 to 2019. The average is almost exactly 1.3X. The ratio is never above 1.6X in any year of the SCF.

We are confident that the results shown here are robust to issues related to the fixed income asset return of the top 1%. Appendix Table C.1 shows the main results using both the DFA and the adjusted DINA microfiles for the post 1989 period for which both data sets are available. The results are similar. Recall that the DFA is based on the SCF and therefore does not have any issue regarding capitalization factors. Furthermore, Appendix Table C.2 shows that the rise in saving by the top 1% would still be 2.6 percentage points of national income annually (as opposed to 2.7 percentage points) if a factor of 1.6X (which is the largest factor found in any year of the SCF from 2001 to 2019) were used instead of 1.3X.[37]

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

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

[37] There is an active debate in the literature on whether the interest rate on fixed income assets for the top 1% should be assumed to be even higher (e.g., Smith et al. (2020), Saez and Zucman (2020)). Smith et al. (2020) use a methodology in which they assume that the interest rate earned by the top 0.1% of the fixed income distribution is the Moody's Aaa corporate rate, and the interest rate earned by 99.0 to 99.9th percentile of the fixed income distribution is the U.S. Treasury rate. In the analysis here, these assumptions would imply a fixed income return of the top 1% of the wealth distribution that is 2.0X the interest rate earned by the U.S. household sector. This is substantially higher than the highest factor ever observed for the top 1% of the wealth distribution in the SCF from 2001 to 2019, as shown in Appendix Figure C.3. We follow the SCF because there are no issues measuring either wealth or the interest rate earned on fixed income assets for the top 1% of the wealth distribution in the SCF.

Figure C.3: Capitalization Factor of Top 1% of the Wealth Distribution from SCF

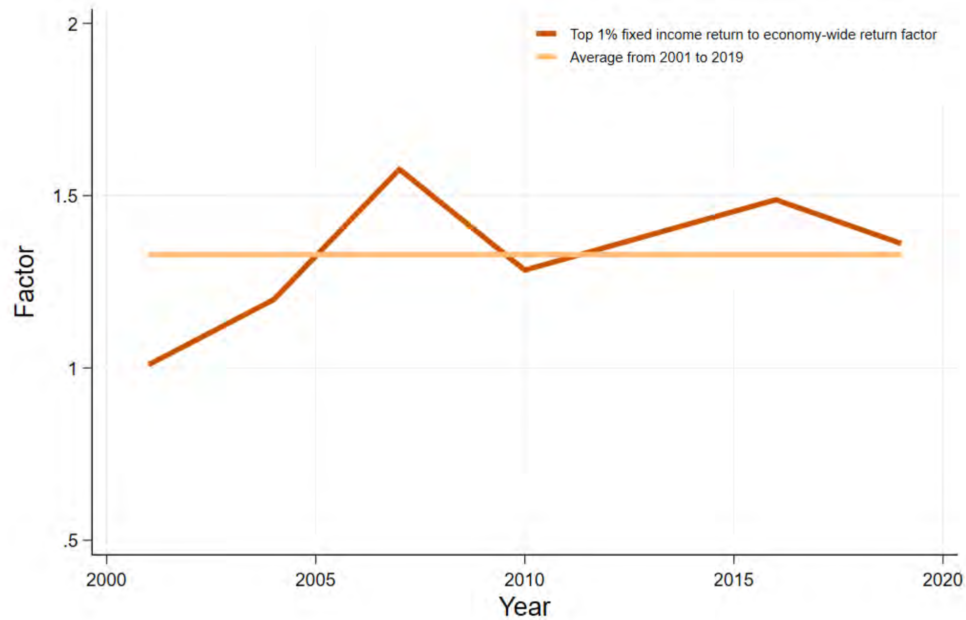

This figure plots the average fixed income asset return for the top 1% of the wealth distribution divided by the average fixed income asset return for the entire population in the SCF. The average factor is almost exactly 1.3.

Table C.2: Top 1% Saving for Different Factors of Top 1% Fixed Income Returns

Multiplicative Factor	1963-1979	2001-2016	Difference
1.0X	0.029	0.062	0.034
1.3X	0.027	0.055	0.028
1.6X	0.026	0.050	0.025

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

C.4 More details on wealth-based approach to measuring saving

This section describes the data underlying the asset inflation measures, and explores alternative methods for calculating the synthetic saving by the cohorts from wealth, which is described in Section 3.

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

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

We construct the annual net charge-off rate as net charge-offs divided by the total outstanding debt using information in Call Report data. This gives us a net charge-off rate series for mortgage debt from 1991 to 2019, and for non-mortgage consumer debt from 1983 to 2019. Net charge off on mortgage debt is not available as a separate line item prior to 1991. We therefore impute the net charge-off rate on mortgage debt from 1983 to 1990 using the non-mortgage consumer credit charge-off rate and the charge off rate on all loans issued by banks as predictors. In particular, we regress the net charge-off rate for mortgage debt between 1991 and 2019 on the net charge-off rate on non-mortgage consumer debt and the net charge-off rate on all bank loans. The R-sq of this regression is quite high at 0.75. We then use the predicted coefficient to predict the net charge-off rate on mortgage debt from 1983 to 1990.

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

Once we have the annual net charge-off rate on mortgage and non-mortgage debt, we calculate how much of debt write-down was on debt borrowed by the top 10% versus the bottom 90%. We do this using ZIP-code-level data on consumer borrowing from Equifax and merging income data from the IRS. We first sum total mortgage and total non-mortgage debt across ZIP codes in the U.S. to calculate the total dollar amount of debt written-down every year. We then allocate the total written down amount to ZIP codes based on the share of total debt default that the zip code has. We sort ZIP codes by their average income per capita (income measured by aggregate gross income)[38] and categorize zip codes into top 10% and bottom 90% by income (population weighted). Finally, we calculate the ratio of total written-down debt to total outstanding debt within each income category for both mortgage and non-mortgage debt separately.

This procedure allows us to compute the debt write-down rate WD_t^{ij} for j equal to mortgage and non-mortgage debt, and i equal to top 10% and bottom 90% from 1991 to 2019. There is no zip code level Equifax data prior to 1991. However, we can use Equifax-based estimates to impute WD_t^{ij} for years prior to 1991 by regressing WD_t^{ij} on the US-level net charge-off rate for mortgage and non-mortgage debt respectively for years 1991 to 2019. We then use the predicted coefficients and data on net charge-offs at the US level to back-fill WD_t^{ij} from 1962 to 1990. Exact details of our entire procedure can be seen in the accompanying code that is made public.

C.5 Population aging

An interesting area for future research is the interaction of population aging and the saving glut of the rich. Unfortunately, the datasets used in this study are not ideal for a comprehensive investigation of this question. The DINA microfiles, for example, only contain information placing individuals into one of two broad age groups–aged 65 or older and under 65–throughout the sample.

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

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

D Estimating distributional net saving rate in Norway

D.1 Allocating corporate saving to households in a small open economy

Households, especially the wealthy who own most of corporate wealth, save significantly via the corporate sector. We outline how sectoral national accounts data on corporate equity holding can be used to determine how much total saving is done by the domestic household sector via corporations. We use Norway as our illustrative example, but the methodology is general, and can be used with any country using standard System of National Accounts (SNA) - e.g. the OECD countries.

Norway is a small open economy with Norwegians holding foreign corporate equity and vice versa. Furthermore, the Norwegian government also saves a lot, and holds large corporate equity assets. In fact, Norway stands out in the world in terms of how much its government saves and how much assets it holds. We outline how national accounts data can be used to properly allocate domestic and foreign corporate net saving to Norwegian households.

Let S^{π} denote corporate saving attributed to Norwegian households; let $S^{\text{c,h}}$ and $S^{\text{c,f}}$ be net saving of the home and foreign corporate sectors, respectively; let $E^{\text{l,h}}$ and $E^{\text{l,f}}$ be market value of equity issued as liabilities by home and foreign corporate sectors; let $E^{\text{a,h}}$ and $E^{\text{a,f}}$ be market value of equity held as assets by home and foreign corporate sectors respectively, as claims on the other corporate sector; let $E^{\text{tot,hh}}$ be total corporate equity (home *and* foreign) held by the home household sector.

The net equity issued by the consolidated global corporate sector is,

$$E^{\text{net,global}} = (E^{\text{l,h}} - E^{\text{a,h}}) + (E^{\text{l,f}} - E^{\text{a,f}}).$$

The global corporate saving rate per dollar of net equity can be defined as, $\sigma^{\text{global}} = \frac{S^{\text{c,h}}+S^{\text{c,f}}}{E^{\text{net,global}}}$. Home and foreign corporate saving rate can analogously be defined as, $\sigma^{\text{h}} = \frac{S^{\text{c,h}}}{E^{\text{l,h}}-E^{\text{a,h}}}$, and $\sigma^{\text{f}} = \frac{S^{\text{c,f}}}{E^{\text{l,f}}-E^{\text{a,f}}}$.

If home and foreign corporates save at the same rate per dollar of equity, i.e. $\sigma^{\text{global}} \equiv \sigma^{\text{h}} = \sigma^{\text{f}}$, then the amount of corporate saving accruing to the home household sector is $S^{\pi} = E^{\text{tot,hh}} \sigma^{\text{global}}$. The proof is straightforward: Net equity, $E^{\text{net,global}}$, issued by the global corporate sector is, by definition, all held by non-corporate sectors - households, government and non-profit. Since every dollar of net equity generates σ^{global} dollars of retained earnings in saving, $E^{\text{tot,hh}} \sigma^{\text{global}}$ gives us the total corporate saving done on behalf of Norwegian households.

However, data from OECD national accounts[39] shows that the saving rate, σ^{h}, for Norway is a little over twice that of the rest of the world (σ^{f}). We can incorporate this difference in saving rate for home versus foreign corporate sectors using data from Statistics Norway (StatBank Table 10788)

[39]OECD SNA Table 710R provides corporate equity assets and liabilities at the sectoral level, and OECD SNA Table 14A provides net saving by corporations.

that decomposes household equity assets into claims on domestic ($E^{\text{home,hh}}$) and foreign corporates ($E^{\text{foreign,hh}}$), i.e. $E^{\text{tot,hh}} = E^{\text{home,hh}} + E^{\text{foreign,hh}}$. Corporate saving done on behalf of Norwegian households then equals[40],

$$S^{\pi} = E^{\text{home,hh}}\, \sigma^{\text{h}} + E^{\text{foreign,hh}}\, \sigma^{\text{f}} \tag{13}$$

D.2 Iacono and Palagi (2023) data by wealth distribution

Figure D.1

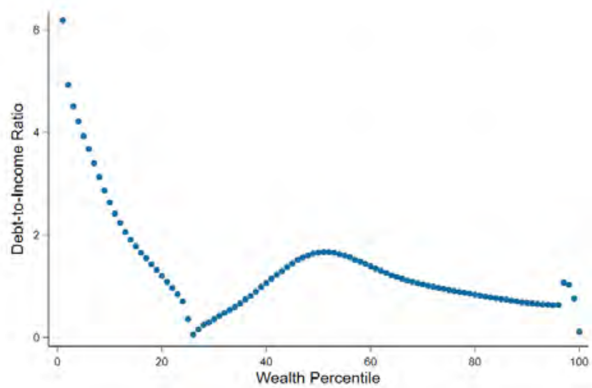

(a) Debt-to-Income Ratio

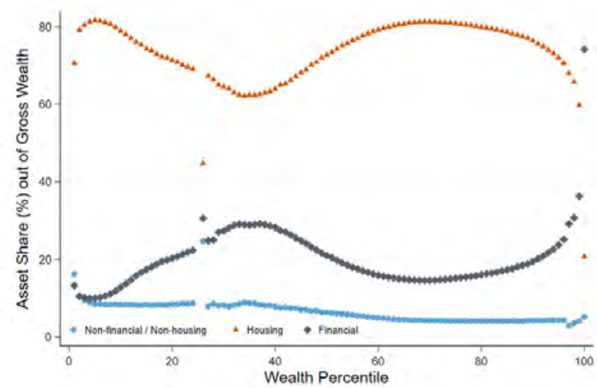

(b) Asset Decomposition Shares

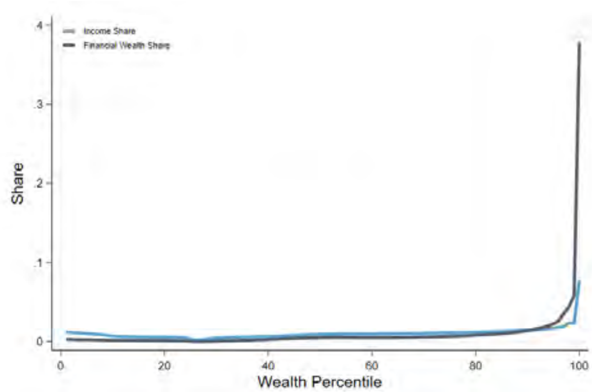

(c) Income vs. Financial Wealth Shares

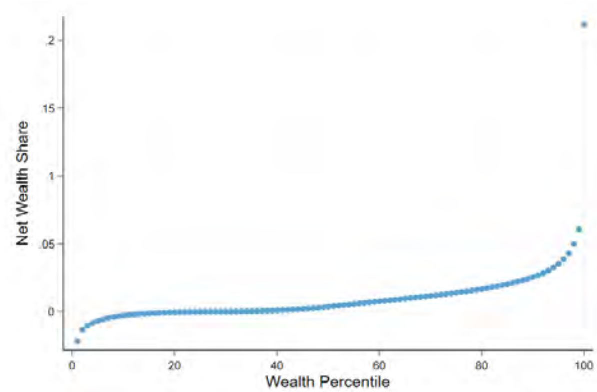

(d) Net Wealth Shares

This figure presents distributional statistics for Norway as reported by Iacono and Palagi (2023), covering 2012–2018 (with income and net wealth data extending to 2011–2018). Panel (a) reports the debt-to-income ratio by wealth percentile, Panel (b) decomposes gross wealth into non-financial/non-housing, housing, and financial asset shares, Panel (c) compares the shares of income and financial wealth, and Panel (d) displays the net wealth share; all series are calculated across the Norwegian wealth distribution.

[40] The adjustment for $\sigma^f < \sigma^h$ in equation (13) is not too important quantitatively - it lowers S^{π} by 9% relative to the estimate using $S^{\pi} = E^{\text{tot,hh}}\, \sigma^{\text{global}}$

D.3 Distributional saving rate in Fagereng et al. (2025)

Fagereng et al. (2025) (henceforth FHMN) use individual-level administrative data on income and wealth to estimate gross and net saving rates similar to those defined in section 3.2. Our methodology goes a different route, emphasizing consistency with the national accounts, as shown in equations (9a), (9b) and (9c). In this section, we highlight the differences in these approaches, and how one may be able to adjust the individual-level saving rates in FHMN to make them consistent with the national accounts.

The first row of table D.1 puts together statistics from Norwegian national accounts for the period 2006-2015 that we believe distributional saving should ideally match.[41] The starting point is the annual change in household net wealth, ΔW, that is then split between net saving and capital gains according to the equation, $\Delta W = \Theta + \pi W'$, where W' is beginning of period net wealth. We also compute a version without pension entitlements, ΔW_p, shown in column (2). FHMN refer to ΔW as "gross saving". Average gross saving is 416 billion Kroner over the sample period, which is divided into 162 billion Kroner of net saving (40.6%) and 254 billion Kroner of capital gain (59.4%). Net saving is the sum of personal household saving and corporate saving done on behalf of households according to the approach explained in D.1. Overall, as a share of total household disposable income, the gross saving rate s^G is 34.8%; the net saving rate s^N is 13.1%, and the difference is capital gains s^κ.

The second row of Table D.1 shows numbers from FHMN, using their mean net and gross saving rate estimates. Since FHMN exclude pension entitlements, we leave column (1) blank. FHMN report their net saving rate out of disposable income (Z_i^d), but their gross saving rate out of Haig-Simons income that includes capital gain. Since we prefer to compare net and gross saving rates in the same units, we convert their gross saving rate to be terms of disposable income as well. Formally, let $s_i^{G,Haig-Simons} = \frac{\Delta W_i}{Z_i^d + \tilde{\pi} W_i'}$ denote the FHMN gross saving rate out of Haig-Simons income. We can transform $s_i^{G,Haig-Simons}$ into the same units as s_i^G and s_i^N by calculating the capital gains saving rate out of disposable income, $s_i^\kappa = \frac{s_i^{G,Haig-Simons} - s_i^N}{1 - s_i^{G,Haig-Simons}}$, and then calculating the gross saving rate with respect to disposable income as, $s_i^G = s_i^\kappa + s_i^N$.

We obtain aggregate implied gross saving by aggregating the FHMN mean saving rates by wealth percentiles with the income shares from Iacono and Palagi (2023) reported in Figure D.1. Since the administrative data FHMN use excludes pension entitlements, this sum represents the total change in household net wealth excluding pension entitlement, ΔW_p. Column (3) reports the corresponding aggregate net saving implied by FHMN. The implied total net saving in FHMN (2025) is 37% of the total household net saving that we measure in the Norwegian national accounts, while gross saving exceeds 200% of what we measure in the national accounts. Columns (4) through

[41] Note that Fagereng et al. (2019) use data from 2005-2015, but Fagereng et al. (2025) use data from 2006-2015.

(8) show that the implied share of capital gains in FHMN is thus relatively high.

Table D.1: Comparing net and gross saving between FHMN and National Accounts

	Aggregates			Share of Gross Saving:		Saving Rates		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	ΔW	ΔW_p	Θ	Net Saving (%)	Capital Gain (%)	s^G (%)	s^N (%)	s^κ (%)
National Accounts	415,849	346,067	161,666	40.6	59.4	34.8	13.1	21.7
FHMN 2025 (Mean)		699,034	44,718	6.4	93.6	58.1	3.7	54.4
FHMN 2025 (Median)		378,019	59,275	15.7	84.3	31.4	4.9	26.5

This table compares saving numbers from national accounts with estimates in Fagereng et al. (2025) for the period 2006-2015. ΔW is annual change in net wealth (in millions of Norwegian Kroner, 2011 prices), ΔW_p is the same but excludes pension entitlement. Both of these series are sourced from SNA 2008 Tables 710R and 9B. Θ is net saving and is sourced from SNA 2008 Table 14A (710R, Statbank Norway Table 10788 and FRED are also used to add the household portion of corporate saving). s^G, s^N and s^κ are gross saving rate, net saving rate, and capital gain saving rate as defined in section 3.2.

While we do not have direct access to the administrative microdata, we explore a simple procedure to adjust the net saving rates so that they aggregate consistently with the national accounts. Specifically, we start with the gross saving rate by wealth percentile, denoted by s_i^G, defined with respect to disposable income as discussed above. This original gross saving rate, s_i^G, is represented by the solid blue line in panel (a) of Figure D.2. We then subtract a constant fraction, κ, of the capital gains saving rate, s_i^κ, from s_i^G across all percentiles, calibrating κ such that the disposable-income-weighted mean net saving rate, s^N, aligns precisely with the national accounts.

The adjusted net saving rate, as depicted in panel (a) of Figure D.2 alongside the original FHMN data, shifts upward relative to the original estimates and exhibits a more gradual rise at the upper end of the wealth distribution. The adjusted saving rates between the 20th and 80th percentiles remain relatively flat, consistent with findings in FHMN. In essence, the procedure increases net saving rates across all percentiles, but it disproportionately raises saving rates for cohorts with higher wealth-to-income ratios, who are naturally the wealthiest cohorts as well.

FHMN also report median gross and net saving rates across the wealth distribution and use median values in much of their analysis. We focused on their mean saving rate estimates because mean saving rates aggregate properly for comparison with national accounts. Median saving rates, strictly speaking, cannot be simply aggregated for such a comparison. With this big caveat in mind, we still believe it is useful to replicate the above analysis "as if" the median saving rates were mean saving rates, simply to get a rough comparison. The third row of Table D.1 and panel (b) of Figure D.2 illustrate this median-based analysis. We obtain similar results to the ones based on true mean saving rates.

Figure D.2: Net and Gross saving rates across the wealth distribution in Norway

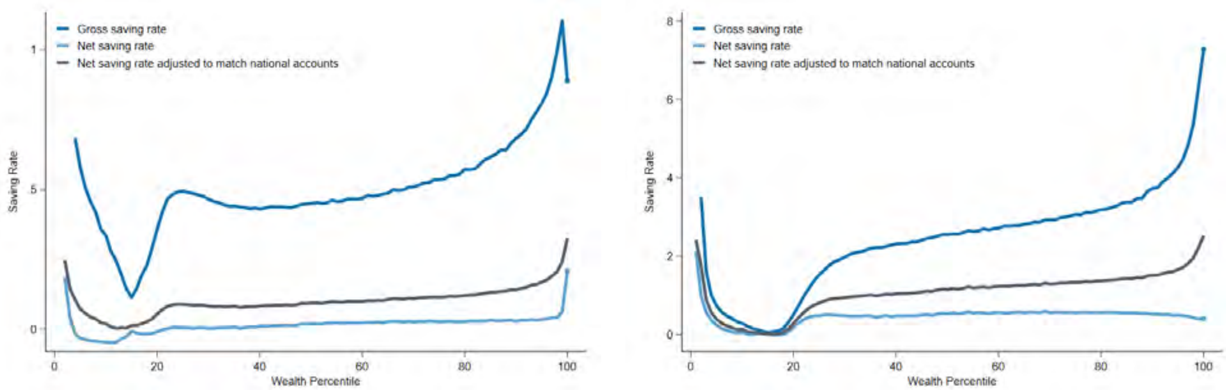

(a) Mean Saving Rates (b) Median Saving Rates

This figure plots gross and net saving rates across the Norwegian wealth distribution for the period 2006-2015, using data from Fagereng et al. (2025). The gross and net saving rates are computed out of household disposable income.