Private Innovation and Wealth Inequality[*†]

Mehran Ebrahimian[1] and Alexander Ljungqvist[2]

[1]Stockholm School of Economics and Swedish House of Finance
[2]Stockholm School of Economics, Swedish House of Finance and CEPR

January 14, 2026

Abstract

Wealth inequality has risen sharply in recent decades. We study how restricted access to private-firm innovation contributes to this rise using a structural model of wealth accumulation with scale dependence. We construct a new long-run measure of innovation value that includes both public and private firms and document a pronounced shift of value creation toward private firms, whose equity is concentrated among high-wealth individuals. Calibrated to the 1979 wealth distribution, the model shows that exclusive access to private-firm innovation accounts for between half and two-thirds of the subsequent increase in the top 1% wealth share.

JEL classification: D31, D52, E21, O31, G51
Keywords: Wealth inequality, Private markets, Innovation, Return heterogeneity

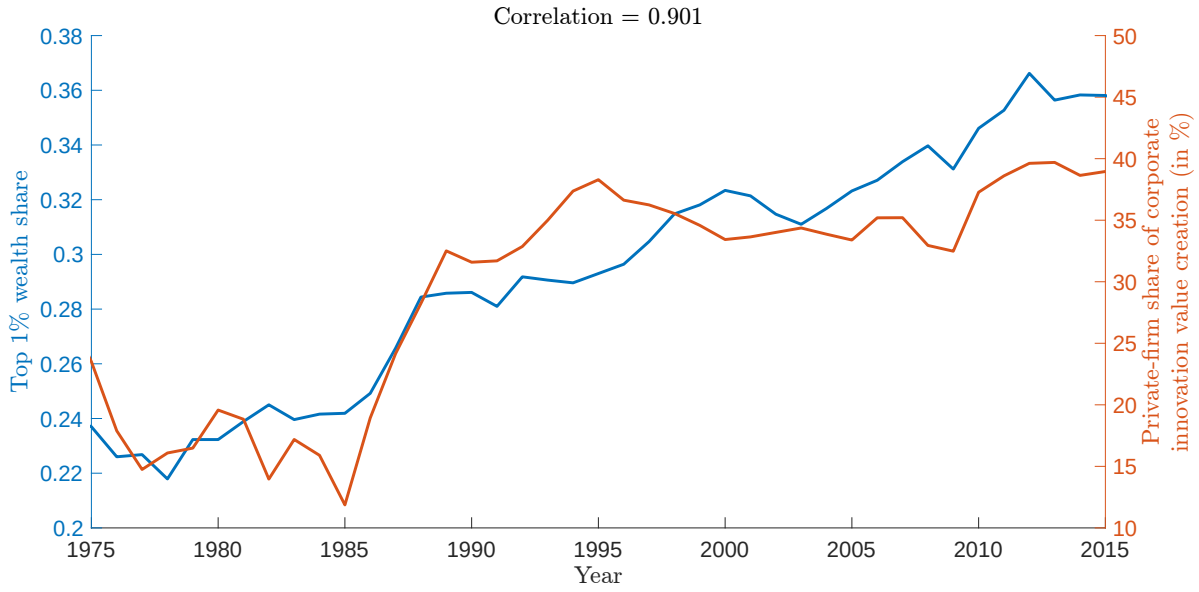
[*]Diana Tagliaferri and Mohammadreshad Esfehani provided outstanding research assistance. Ebrahimian gratefully acknowledges generous funding from Vetenskapsrådet (grant no. 2022-02555). Ljungqvist gratefully acknowledges generous funding from the Knut & Alice Wallenberg Foundation (grant no. KAW 2023.0270) and the European Research Council (grant no. 101142217).

[†]Funded by the European Union. Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

1 Introduction

Wealth inequality in the United States has risen sharply in recent decades. Between 1979 and 2015, the share of wealth held by the top 1% increased from 23.2% to 35.8%. Over the same period, an increasing share of corporate innovation—measured by the economic value of patented inventions net of R&D costs—has been generated by firms that are not listed on public stock markets. Private firms' share of net patent value more than doubled, from 16.5% in 1979 to 38.9% in 2015. As shown in Figure I, the rise in the top 1% wealth share closely tracks the rise in private firms' share of innovation value creation (the time-series correlation is 90.1%). The aim of this paper is to show that these two trends are tightly connected: the rise in top wealth inequality is largely a consequence of the shift in where corporate R&D value is created—from publicly traded firms (in which broadly anyone can invest) to private firms (in which predominantly the wealthy can invest).

Figure I: Wealth inequality and private-firm innovation



Note: The figure graphs the top 1% wealth share in the U.S. (on the left axis) and private U.S. firms' share of corporate innovation value creation (on the right axis) between 1975 and 2015. Corporate innovation value creation is the aggregate value of corporate patents (imputed from Kogan et al.'s (2017) estimates of patent values) net of R&D expenses (as measured by the NSF).

We develop a dynamic model of wealth accumulation with heterogeneous returns. Individuals earn a common return on broadly accessible assets, but only those whose wealth exceeds a threshold can invest in innovative private firms and thereby capture private-firm innovation value. We

estimate the model to match the wealth distribution in 1979 and then feed in the observed evolution of total corporate innovation value and its rising private share. This allows us to quantify how much exclusive access to private-firm innovation contributes to the subsequent increase in wealth inequality.

Our empirical starting point is to construct a long-run measure of the economic value created by corporate innovation. We follow the patent-valuation framework of Kogan et al. (2017), which infers the value of each patent granted to a U.S. public firm from the stock-market reaction around the grant date. To extend these valuations to private firms, we use an imputation model that relates patent values to observable patent characteristics in the spirit of Kline et al. (2019). Aggregating the resulting firm-level patent values across all public and private firms yields an annual series of total patent value back to 1975. Subtracting contemporaneous R&D spending gives a net measure of corporate innovation value creation. When scaled by aggregate household wealth, this series implies that net corporate innovation value creation has contributed about 2% to annual wealth growth on average since 1975.

A key empirical fact underlying our analysis is that an increasingly large share of corporate innovation value is generated by private firms. Decomposing total innovation value by firm type, we find that private firms accounted for roughly 15% of net patent value between 1975 and 1985, but their share rose sharply from the mid-1980s onward and stabilized at around 35%–40% in 1995–2015. This shift is not driven by a growing number of private-firm patents. Rather, patents produced by private firms have become substantially more valuable relative to those of public firms, and private firms now generate this value at lower R&D cost. In effect, private firms increasingly account for the most net-valuable innovations in the economy. Because both ownership of and investment in innovative private businesses are highly concentrated among wealthy individuals (Cagetti and De Nardi, 2006; Bach, Calvet, and Sodini, 2020), the rising private share of innovation value channels a growing fraction of aggregate value creation to the top of the wealth distribution.

To quantify the implications of rising private-firm innovation value, we embed a single source of return heterogeneity via scale dependence into a dynamic wealth-accumulation model for the upper tail of the distribution (Gabaix et al., 2016; Hubmer, Krusell, and Smith, 2021). In the model, individuals who cross a wealth threshold receive an extra return component that reflects exposure to

private-firm innovation value. Idiosyncratic shocks to wealth generate mobility around the wealth threshold, so that access to this additional component changes endogenously over time. Our model allows us to isolate the role of exclusive access to private-firm innovation in shaping wealth dynamics. It yields a law of motion in which the upper tail responds not only to aggregate returns and economic growth—the standard "r minus g" channel—but also to the evolving magnitude and exclusivity of innovation value created by private firms.

We calibrate the model to match the U.S. wealth distribution in 1979, the year when top wealth concentration reached its historical low. We infer the common return from aggregate post-tax capital-income data. Guided by evidence from the Survey of Consumer Finances that private-business holdings are highly concentrated, our baseline specification assumes that only individuals in the top 1% have access to the additional return component associated with private-firm innovation. The time-varying wedge in wealth accumulation between high-wealth individuals and the rest is disciplined directly by our series on total innovation value (relative to aggregate wealth) and its private-firm share: larger innovation-driven wealth gains, and a larger private share, imply a larger wedge. The remaining free parameter—the subjective discount rate governing consumption behavior—is chosen so that the stationary wealth distribution implied by the model matches the observed top 1% wealth share in 1979.

We then feed in the observed evolution of innovation value, its private-firm share, aggregate capital income and taxation, and economic growth from 1979 onward to simulate the wealth distribution through 2015. This exercise quantifies how much exclusive access to the rising private share of innovation value contributes to the observed increase in the top 1% wealth share.

Our baseline quantitative results show that our model with exclusive access to private-firm innovation for high-wealth individuals explains a large share of the observed rise in top wealth inequality. In the data, the top 1% wealth share increases from about 23% in 1979 to 36% in 2015, a rise of roughly 13 percentage points. Our model generates nearly 9 percentage points of this increase, accounting for about two-thirds of the empirical trend. By contrast, a standard wealth-accumulation model without exclusive access—in which all individuals earn the same average post-tax return—produces an increase of less than 1 percentage point over the same period. This stark divergence highlights the central role played by private-firm innovation value, and by exclusive access to it, in

driving the rise in top wealth concentration.

To assess the importance of the rising private share of innovation value, we run a counterfactual in which this share is held fixed at its initial level of 16.5% in 1979. In this scenario, high-wealth individuals continue to receive, every year, the same fraction of innovation value as in 1979, even as total innovation value evolves over time. As a result, the wealth-accumulation wedge between high-wealth individuals and the rest increases only modestly. The counterfactual produces a rise of just 4 percentage points in the top 1% wealth share, compared with 9 percentage points in the baseline. The gap between these simulations shows that the increasing private share is a powerful amplification mechanism: without it, the model yields only a muted rise in top wealth inequality.

The analysis so far highlights the central role of exclusive access to private-firm innovation in shaping wealth dynamics. We next examine how tax policy interacts with this mechanism. We find that policies affecting the taxation of capital gains and R&D subsidies have first-order effects on top wealth concentration, whereas taxes on interest income and dividends play a more limited role.

We conduct several robustness analyses to verify that our main result is not sensitive to alternative calibrations of idiosyncratic risk, access thresholds, or within-top inequality targets. For example, re-estimating the model to match inequality within the top 10% yields a nearly identical rise in the top 1% wealth share. Across all robustness specifications, our model continues to explain at least half of the observed increase in the top 1% wealth share since 1979.

Related Literature. A large body of work studies the pronounced rise in U.S. wealth inequality in recent decades (Saez and Zucman, 2016; Piketty, Saez, and Zucman, 2018). A central explanation in quantitative macroeconomic models is heterogeneous returns to wealth. Benhabib, Bisin, and Zhu (2011), Benhabib and Bisin (2018), and Benhabib, Bisin, and Luo (2019) emphasize the role of saving behavior and return heterogeneity in shaping top wealth shares and intergenerational mobility. Hubmer, Krusell, and Smith (2021) combine several mechanisms—earnings heterogeneity, income-tax progressivity, discount-factor heterogeneity, and return heterogeneity—to account for the evolution of top wealth concentration. Gabaix et al. (2016) highlight how return heterogeneity, whether arising from exogenous type differences or scale dependence, affects the dynamics of in-

equality.[1] Empirically, recent papers using micro data document that returns vary systematically across individuals and that this heterogeneity is sizable and persistent (Bach, Calvet, and Sodini, 2020; Fagereng et al., 2020). Using administrative data, Ebrahimian and Sodini (2024) estimate preference heterogeneity and show first-order implications for returns on wealth and intergenerational persistence.

We contribute to this literature by identifying and quantifying a specific, previously unmeasured source of scale dependence: exclusive access for high-wealth individuals to the value created by private-firm innovation. Our analysis is complementary to Gomez and Gouin-Bonenfant (2025), who show that lower required returns raise top wealth inequality by accelerating entrepreneurial capital accumulation; we focus instead on how unequal access to private-firm innovation reallocates aggregate innovation rents toward the top of the wealth distribution. While existing work introduces return heterogeneity as a structural assumption or estimates it from micro returns, no study links this heterogeneity to differential access to a particular set of investment opportunities—namely, the innovation value generated by private firms. Using our aggregate series on innovation value and its decomposition by firm type, we show that this access-driven form of scale dependence has substantial explanatory power for the rise in top wealth inequality. Our focus differs from research on household portfolio heterogeneity and valuation channels (Greenwald et al., 2021; Catherine et al., 2023; Catherine, Miller, and Sarin, 2025): those papers study how time-varying asset yields affect valuations, whereas we examine heterogeneity in access to a source of value creation itself—innovation undertaken by private firms.

A narrower line of empirical research has examined individuals' investment in private markets and its implications for inequality. Canipek (2024) studies investment behavior around the "accredited-investor" eligibility cutoff,[2] while Gocmen, Martínez-Toledano, and Mittal (2025) analyze how capital-gains taxation affects high-net-worth individuals' investment in early-stage companies and their realized returns. Our analysis differs in two important ways. First, we are not concerned with a particular regulatory threshold or a specific segment of the private market. Instead, we aggregate the full value of private-firm innovation in the U.S. economy and quantify the extent to which

[1] Seminal early contributions such as Quadrini (1999) and Cagetti and De Nardi (2006) feature return heterogeneity through occupational choice and entrepreneurship.

[2] Lindsey and Stein (2019) analyze angel investment under the accredited-investor definition, but focus on firm outcomes rather than investors' wealth accumulation.

access to this value creation is concentrated among wealthy individuals. Second, we embed this access mechanism into a structural, economy-wide model of wealth accumulation and show that it has first-order implications for the rise in top wealth concentration.

In summary, we assemble a new long-run measure of corporate innovation value that, for the first time, covers both public and private firms. We document a sharp rise in the share of innovation value generated by private firms and show that this shift is driven by the increasing economic value of their patents. We build a structural model in which access to this private-firm innovation value is limited to high-wealth individuals, calibrate it to the U.S. wealth distribution in 1979, and show that this access mechanism can account for a substantial share of the subsequent rise in top wealth inequality. Our framework highlights a new channel through which the organization of corporate innovation—particularly its movement outside public markets—feeds directly into the concentration of wealth at the top.

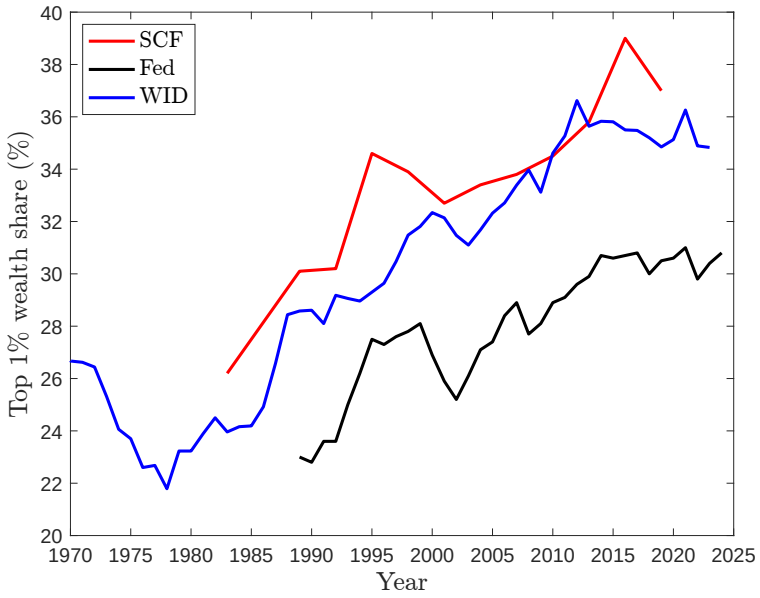
The remainder of the paper is organized as follows. Section 2 introduces the data and documents the evolution of corporate innovation value and its private share. Section 3 presents the model, and Section 4 describes its calibration. Section 5 reports the baseline results and counterfactuals, and Section 6 provides robustness checks. Section 7 concludes.

2 Data and Facts

2.1 The Rise in Wealth Inequality

Wealth inequality in the United States has increased markedly over the past several decades. Figure II plots the top 1% wealth share using three widely used sources: the Survey of Consumer Finances (SCF), the Federal Reserve's Distributional Financial Accounts (DFA), and the World Inequality Database (WID). The SCF provides survey-based estimates of household wealth. The DFA combines SCF microdata with national accounts to incorporate components of wealth that are poorly captured in surveys, such as defined-benefit pensions, annuities, and insurance reserves (Batty et al., 2021). The WID series derives wealth estimates from administrative income tax data using a structural mapping between capital income and asset holdings. Because it offers consistent coverage over a longer horizon, our analysis primarily relies on the WID data.

Figure II: Wealth inequality



Note: The figure plots the top 1% wealth share using data from three sources: the Survey of Consumer Finances (SCF), the Federal Reserve DFA Survey, and the World Inequality Database (WID).

Despite differences in levels, all three sources show a pronounced rise in top wealth concentration. In the WID series, for example, the top 1% share increased from about 23% in 1979 to roughly 35% in 2023.

2.2 Private Firms' Patenting

2.2.1 Patent Values and R&D Spending

To estimate our model, we require data on patent values and corporate R&D spending. We follow Kogan et al. (2017) to estimate the value of patents issued to stock market listed firms and Kline et al. (2019) to impute the value of patents issued to private firms. Our sample period begins in 1975 (when accounting standard changes allow us to estimate public- and private-firm R&D spending) and ends in 2015 (the last year for which a key variable used in the imputation model is available).

Kogan et al. (2017) estimate the market value of a public firm's patent from stock price movements around the announcement of the patent's grant to the firm. Estimating their model requires identifying patents granted to domestic stock market listed firms, rather than to domestic unlisted firms, individuals, universities, hospitals, non-profits, the government, or foreign corporations. Our

starting point is the linking table provided by Kogan et al. (2017), which links 1.9 million patents granted between 1926 and 2010 to public (not necessarily domestic) firms in the CRSP database. We refine this linking table in two principal steps. First, we update the CRSP mapping through 2015. Second, we use data on corporate trees extracted from the National Establishment Time Series (NETS) database to identify patents granted to subsidiaries of public firms starting in 1989.[3] We then reestimate Kogan et al.'s model using our updated public-firm links.

To estimate the value of patents granted to private firms, we first create a linking table that classifies patents assigned to corporations as belonging to private firms if they are not assigned to public firms according to our updated public-firm linking table. We then estimate an imputation model following Kline et al. (2019), who extrapolate Kogan et al.'s (2017) estimates based on patent characteristics using a Poisson regression model. In our application, we regress our estimates of public-firm Kogan et al. patent values on the following six characteristics: a patent's 5-year forward citations, the number of independent claims it has been granted, its generality and originality scores (Trajtenberg, Rebecca, and Jaffe, 1997; Hall, Jaffe, and Trajtenberg, 2001), whether it is a breakthrough invention (Kelly et al., 2021), and whether it is a strategic patent (Veihl, 2022). We estimate this valuation model recursively over 10-year intervals, absorbing fixed effects at the grant-year level and at the Cooperative Patent Classification (CPC) main group level.[4] While the effect sizes vary over time, we find that patent values generally increase in citations, the number of claims, generality, and originality, and that breakthrough patents and strategic patents are more valuable.

Figure III, Panel (a) plots the annual aggregate imputed value of all U.S. patents granted to U.S. firms over the period from 1975 to 2015, in constant 1982 dollars. Through the mid-1990s, U.S. firms produced patented inventions valued at between $100 billion and $350 billion per year. In the second half of the 1990s, corporate patent values increased sharply, peaking at $2.7 trillion in 2000. In the 15 years since, corporate patent values have averaged $1.1 trillion per year. The graphs for public and private U.S. firms follow a similar pattern.

Figure III, Panel (b) plots aggregate corporate R&D spending by U.S. firms using data made avail-

[3]This step captures patents assigned to subsidiaries whose names bear little relation to the name of their listed corporate parent. This helps correct erroneous links in the Kogan et al. (2017), which is primarily based on a fuzzy name match.

[4]We include pre-sample data starting in 1966 so that each year's imputation model is based on 10 years of recursively estimated coefficients.

able by the National Science Foundation (NSF) and the National Center for Science and Engineering Statistics, in constant 1982 dollars.[5] Corporate R&D spending has increased at a compound annual growth rate of 3% between 1975 and 2015, but with considerable variation: in the 1970s and 1980s, corporate R&D grew by 4.5% per annum, slowing to 3.9% in the 1990s and 1.4% in the 2000s, and finally increasing to 4.8% in the first half of the 2010s.

In 1972, the Accounting Principles Board and the Securities and Exchange Commission made the disclosure of R&D expenses mandatory for public firms, and in 1975, the Financial Accounting Standards Board required public firms to expense all R&D expenditures in the year incurred, rather than to capitalize some or all of them (Nix and Nix, 1992). These accounting standard changes allow us to estimate the aggregate annual R&D expenditures of public U.S. firms from 1975 using firm-level data available in the Compustat database. Subtracting these from the NSF's estimates for total corporate R&D spending gives us an estimate of the aggregate annual R&D spending by private U.S. firms. Figure III, Panel (b) plots these two time series. While R&D spending by public firms has increased at a compound average rate of 4.1% per year between 1975 and 2015, R&D spending by private firms has actually fallen in real terms, at a compound average rate of −2.6% per year. (The 2000s is the only decade in which private firms grew their R&D spending faster than public firms, by 5.3% vs. 0.7%.) In relative terms, then, private firms account for a shrinking share of corporate R&D spending in the U.S.

2.2.2 Private-firm Innovation

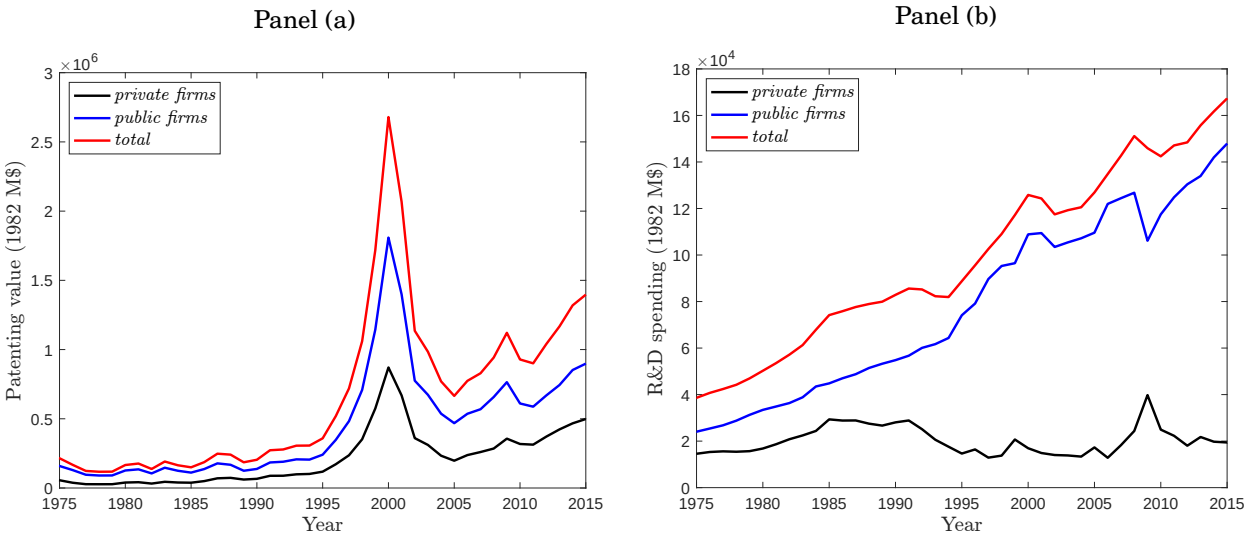
The focus of our analysis is on the effects of exclusive access for the wealthy to the fruits of innovation created by private firms. To measure time variation in the private share of corporate innovation value creation, we define e_t as the ratio of the aggregate imputed value of patents granted to private U.S. firms (net of aggregate R&D spending by private U.S. firms) in year t and the aggregate imputed value of patents granted to all U.S. firms (net of aggregate corporate R&D spending) in year t.[6] We refer to e_t as the private innovation share.

Figure IV, Panel (a) plots e_t over the period 1975-2015. The private innovation share halved between 1975 and 1985, from 23.5% to 11.9%, then more than tripled over the next 10 years, to

[5] The latest release can be found here.

[6] Our results are robust to using actual estimates for public firms and imputed estimates for private firms.

Figure III: Patent values and R&D spending

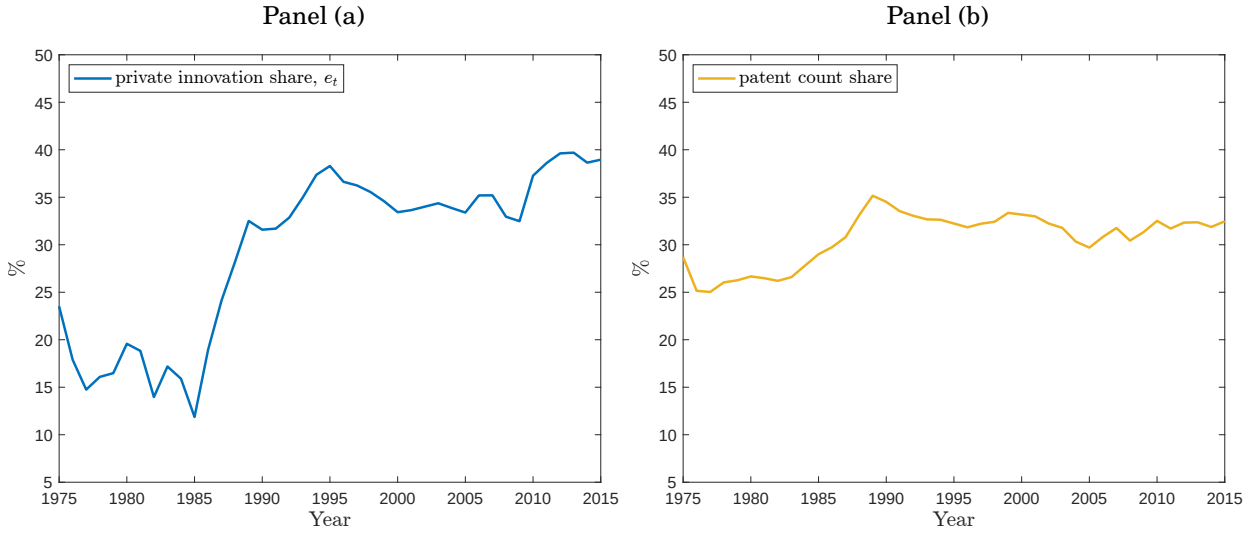


Note: Panel (a) plots the aggregate imputed value of all U.S. patents granted to U.S. firms as well as the imputed value of patents granted to public and private U.S. firms, respectively, in constant 1982 dollars. Panel (b) plots aggregate corporate R&D spending by U.S. firms as well as the R&D spending of public and private U.S. firms, respectively, in constant 1982 dollars. Data on private firms' R&D spending are not available before 1975.

38.3% in 1995, and has since been fairly stable at around 32% to 40%. By implication, at least a third of the total value created by corporate R&D programs in the U.S. since 1995 has accrued exclusively to those individuals who have access to, and can invest in, private firms.

A priori, the private innovation share could increase either because private firms generate an increasing share of patented inventions or because private firms increasingly generate inventions with more valuable characteristics compared to public firms (or a combination of both). Figure IV, Panel (b) shows that private firms' share of the annual number of corporate patents has been remarkably stable since the 1970s, at around 30%. Thus, the dynamics of e_t appears to be entirely driven by changes in value creation.

Figure IV: The private innovation share and the private patent-count share



Note: Panel (a) plots the time series of the private innovation share, e_t, defined as the share of the value of private U.S. firms' patents in the aggregate value of U.S. corporate patents, each net of R&D spending. Data on private firms' R&D expenditures are not available before 1975. Panel (b) plots the time series of the annual count share of corporate patents granted to private U.S. firms.

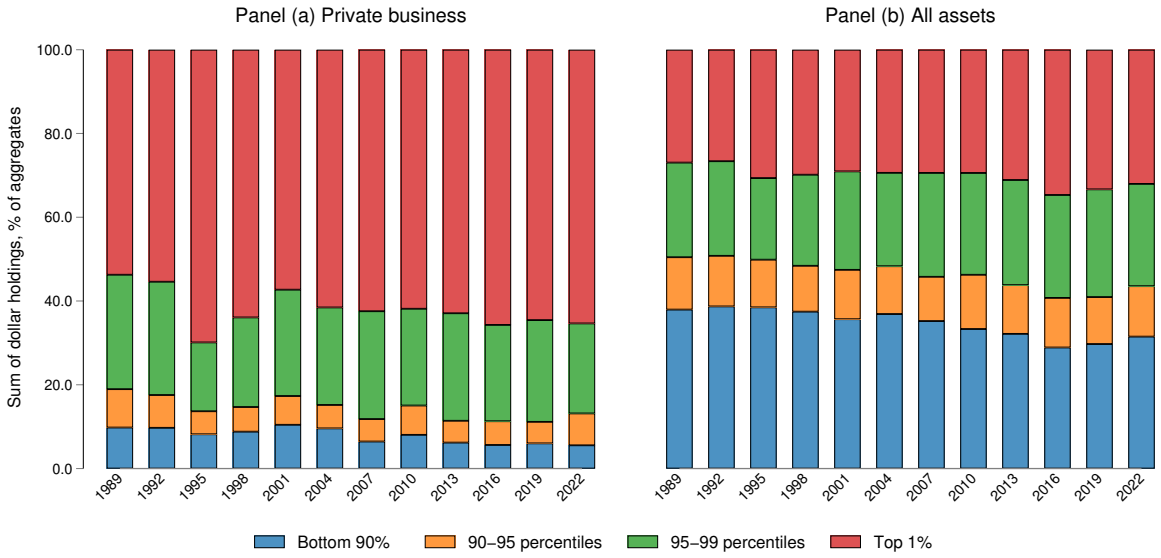
2.3 The Concentration of Private-Business Equity

Ownership of, and investment in, private businesses is extremely concentrated in the United States. Figure V plots the shares of private-business equity (Panel a) and total assets (Panel b) held by four groups: the top 1%, individuals in percentiles 95–99, individuals in percentiles 90–94, and the bottom 90%. The data come from the SCF for 1989–2022 and cover both active and passive private-business holdings. Active holdings include a wide range of privately held firms, from sole proprietorships and professional practices to founders' equity stakes in their start-ups.

The figure shows that the top 1% consistently hold the majority of private-business equity in every SCF wave, and this share has remained remarkably stable over time. The top 5% collectively hold more than 80% of private-business equity. By contrast, the bottom 90% hold less than 10%, a small and declining share that presumably largely reflects self-employment businesses rather than innovative private firms.

Importantly, this concentration is not simply a reflection of the overall concentration of wealth. As shown in Panel (b) of Figure V, total assets are far less concentrated. Consistent with Figure II, the top 1% hold roughly 25–35% of aggregate assets—only about half of their share of private-business

Figure V: Private-business equity, by net-wealth percentile



Note: The figure plots the shares of private-business equity (in Panel a) and total assets (in Panel b) that are held by four groups of individuals: the top %1, individuals in percentiles 95-99, individuals in percentiles 90-94, and the bottom 90%. The data come from all available waves of the Survey of Consumer Finance.

equity. Thus, high-wealth individuals hold a disproportionately large fraction of their portfolios in private-business equity.[7]

The high concentration of private-business equity implies that any rise in the economic value generated by private firms will necessarily channel a growing share of aggregate value creation to the top of the distribution. Section 3 develops a structural model to quantify this mechanism.

2.4 Other Data

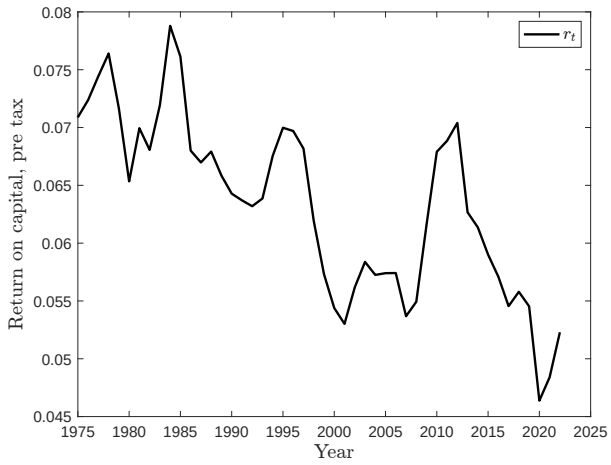
2.4.1 Aggregate Return on Capital

We construct the average annual pre-tax return on capital, r_t, following the methodology of Piketty and Zucman (2014). Specifically, we measure r_t using the relationship $r_t = \theta_t/\beta_t$, where θ_t represents the capital share of national income and β_t denotes the wealth-to-income ratio. To estimate θ_t and β_t, we use data from the Integrated Macroeconomic Accounts (IMA) maintained by the Federal Reserve and the National Income and Product Accounts (NIPA) maintained by the Bureau of Economic Analysis. The capital share of national income, θ_t, is computed as total capital income (consisting of

[7] Table B.1 reports the corresponding dollar values across wealth groups and SCF waves.

corporate profits, housing capital income, self-employment capital income, net foreign capital income, and net government interest payments) divided by national income. The wealth-to-income ratio, β_t, is obtained by dividing net private wealth (consisting of financial and non-financial assets, net of financial liabilities) by national income.[8] Figure VI graphs the time series of r_t over the period from 1975 to 2020.

Figure VI: Average pre-tax return on capital



Note: The figure plots the average pre-tax return on capital, r_t, for the period from 1975 to 2020.

2.4.2 Capital Income Taxes

Our measure of personal taxes on capital income departs from the literature by combining a focus on *marginal* tax rates (those applying to the top of the income distribution) with consideration of the different components of capital income: interest income, dividends, and capital gains. As we will show, distinguishing between these components turns out to be important for two reasons. First, long-term capital gains and the taxes due on them are the most relevant to understanding the evolution of wealth for the top 1%. Second, marginal tax rates on interest income, dividends, and capital gains differ greatly from each other and vary considerably over time.[9]

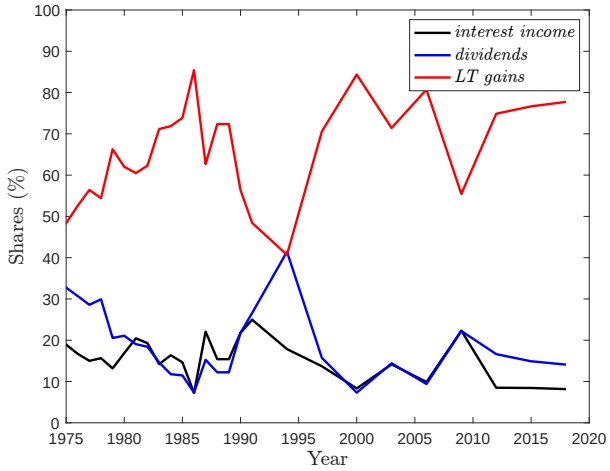
Figure VII graphs the shares of their total capital income that the top 1% receive from interest,

[8] See Section A.1 in the Appendix for further details.

[9] Moreover, the U.S. applies different tax rates to capital gains on short-term investments (those held for less than a year) and long-term investments (those held for more than a year). There are at least three reasons why patent-related capital gains would usually be taxed at the long-term capital gains rate in our setting: investments in R&D typically take a long time to bear fruit; once it has, obtaining a patent usually takes at least a year (and often longer); and private firms are illiquid, meaning an investor cannot sell their shares in private firms at will, which tends to increase holding periods.

dividends, and long-term capital gains, respectively.[10] Long-term capital gains are consistently and by far the largest source of capital income in every year between 1975 and 2020, while interest income is almost always the smallest source (with dividends in between). To illustrate, in 1979, the top 1% received as much as 66% of their total capital income from long-term capital gains and only 13% from interest income.

Figure VII: Capital income shares



Note: The figure graphs the shares of various capital income components in total capital income. We use these shares as weights to compute the overall marginal tax rate on capital income. Before 1988, the data source is WID, with minor adjustments to long-term capital gains; from 1988 onward, capital income shares are calculated using the SCF dataset.

Figure VIII, Panel (a) graphs the marginal tax rates on each component of capital income. To estimate these tax rates, we follow Auerbach and Hassett (2015) and others and use the TAXSIM model maintained by the National Bureau of Economic Research (NBER).[11,12] Except in 1988, when all components of capital income are briefly taxed at the same rate, taxes on long-term capital gains are on the order of 15-30 percentage points lower than taxes on interest income; before the Jobs and Growth Tax Reconciliation Act of 2003, they are also lower than taxes on dividends. Collectively, Figure VII and Figure VIII show that using the marginal tax on interest income to compute the

[10] For details on how we construct each of the series discussed in this section, see Section A.2 in the Appendix.

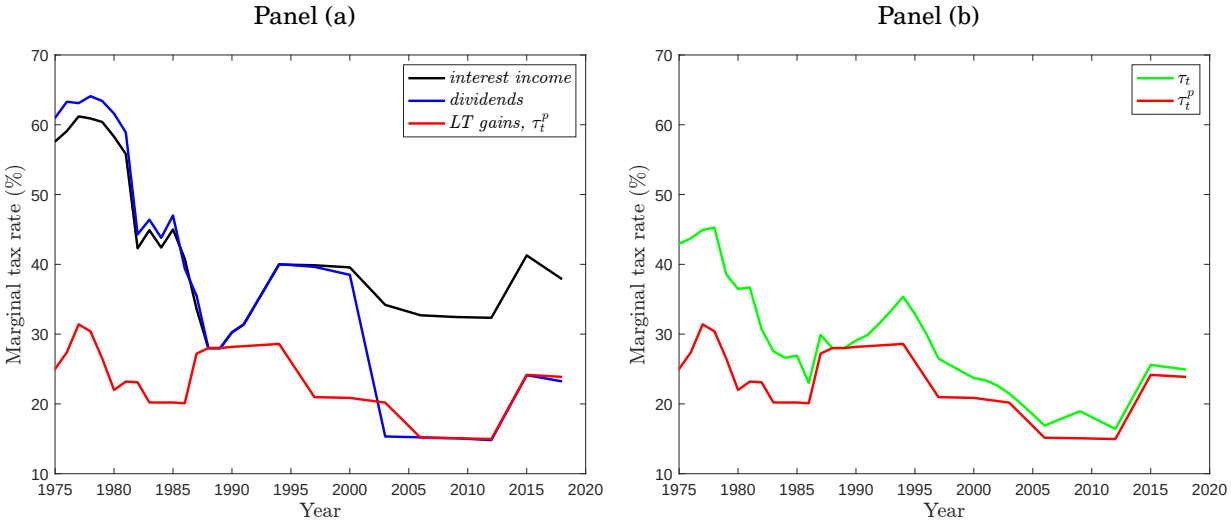
[11] For details on TAXSIM, see Feenberg and Coutts (1993).

[12] TAXSIM simulates an individual's marginal tax rate on each component of capital income when fed the individual's data on labor income (including pension and business income), capital income (i.e., interest income, dividends, and short-term and long-term capital gains), and deductions. A practical consideration arises because the relevant input data are available in the Survey of Consumer Finances (SCF) only from 1988 onward. For the period from 1975 to 1987, we use Feenberg and Coutts' (1993) estimates of marginal tax rates based on tax returns filed with the Internal Revenue Service. Because Feenberg and Coutts focus on the top 1% of the *income* rather than the *wealth* distribution, our TAXSIM simulations after 1988 use SCF data for the top 1% by income. Figure A.3 in the Appendix shows that at least after 1988, marginal capital income tax rates are very similar in the top 1% of the income and wealth distributions.

after-tax return of the top 1% (as is done in, e.g., Auerbach and Hassett, 2015) is misleading for two reasons: the top 1% derive most of their capital income from long-term capital gains, and long-term capital gains are taxed at a much lower rate than is interest income. Moreover, the dynamics of marginal tax rates on interest and on long-term capital gains are very different: while the former have generally fallen over time, the latter have been much more stable.

Our model calibration in Section 4 uses two tax rates: the marginal tax rate on long-term capital gains, denoted τ_t^p, and the weighted average marginal tax rate on capital income, denoted τ_t. The latter uses as weights the capital income shares of the top 1%, plotted in Figure VII. Figure VIII, Panel (b) graphs the two tax-rate series.

Figure VIII: The top 1%'s marginal tax rates on capital income



Note: Panel (a) plots the top 1%'s marginal tax rates on the three main components of capital income: interest income, dividends, and long-term capital gains. (We ignore short-term capital gains, which account for no more than a few percent of capital income, because the required data are not available before 1988.) Panel (b) plots the weighted-average marginal tax rate on capital income, τ_t (using the capital income shares shown in Figure VII as weights), alongside the marginal tax rate on long-term capital gains, τ_t^p, from Panel (a). For the period before 1988, estimates of marginal capital-income tax rates are obtained from Feenberg and Coutts (1993). From 1988 onward, estimates are obtained from TAXSIM using data from the SCF as inputs.

3 Model

We build on the models for the dynamics of inequality presented by Gabaix et al. (2016) to study the evolution of top wealth inequality. We include return heterogeneity via scale dependence by considering exclusive access for wealthier individuals to the innovation value creation by private

firms.

Denote by w_{it} individual i's wealth at time t. Every individual receives a common return on her wealth in period t, which we call the "market return" and denote by r_t^m. We can think of r_t^m as the return generated by investments in companies that are listed on the stock market. Individuals whose wealth w_{it} is sufficiently high (to be specified shortly) have exclusive access to an additional investment technology that increases their wealth at a rate Δ_t (on top of r_t^m). We can think of this technology as investments in innovative private firms. Individuals also experience idiosyncratic shocks to the return on their wealth, denoted by σdZ_{it}, where σ is the standard deviation of the shock and dZ_{it} is *iid* across individuals and time. The rate of return on wealth thus varies across individuals because *some* individuals have exclusive access to Δ_t and because individuals experience different shocks dZ_{it}; it also varies over time because both r_t^m and Δ_t vary over time. We consider all these terms to represent *after-tax* returns.

In the baseline model presented in this section, we consider a form of scale dependence in which an individual has access to Δ_t if and only if her wealth w_{it} exceeds $\bar{w}_t$, which we call the "access cutoff." In the generic representation, this threshold may vary over time. In our simulations in Section 5, it is set flexibly, either at the 99th wealth percentile at each point in time (corresponding to only the top 1% having access to Δ_t) or at a fixed value in absolute terms over time (reflecting, for example, regulatory criteria defining "accredited investors" permitted to invest in private companies). Later in Section 6, we also explore cases with multiple cutoffs, modeling individuals with heterogeneous (and potentially partial) access to Δ_t.

All individuals earn the same deterministic labor income y_t, which grows at the rate g_t over time. In our setup, labor income serves only as the stabilizer of the wealth distribution.[13] At each point in time, an individual consumes at a rate proportionate to her wealth: $c_{it} = \rho w_{it}$, where ρ represents the time discount rate in individuals' preferences. This consumption policy can be micro-founded by considering a unit elasticity of intertemporal substitution (EIS),[14] when ignoring human capital in total wealth, as is reasonable for individuals at the top of the wealth distribution.

[13] A vast literature shows that variations in labor income do not help explain wealth inequality in the data, especially regarding the concentration of wealth at the very top. See Benhabib and Bisin (2018) for a review.

[14] See Wachter (2013) for the analytical derivation.

The individual's wealth accumulation follows

$$dw_{it} = y_t dt - \rho w_{it} dt + (r_t^m + \mathbf{1}\{w_{it} \geq \bar{w}_t\}\Delta_t) w_{it} dt + \sigma w_{it} dZ_{it}, \tag{1}$$

where the indicator function $\mathbf{1}\{w_{it} \geq \bar{w}_t\}$ represents the criterion for exclusive access to Δ_t. By definition, Δ_t is the difference in the wealth accumulation rate between those whose wealth is above and below the access cutoff $\bar{w}_t$. We refer to Δ_t as the "wealth-accumulation wedge."

Define x_{it} as the *log* of wealth w_{it}, scaled by the accumulated economic growth g_t over time. We work with x_{it} when simulating the model to generate the wealth distribution. The stochastic law of motion for x_{it} follows

$$dx_{it} = (ye^{-x_{it}} - \rho + r_t^m - g_t + \mathbf{1}\{x_{it} \geq \bar{x}_t\}\Delta_t - \tfrac{1}{2}\sigma^2)dt + \sigma dZ_{it}, \tag{2}$$

where $\bar{x}_t$ is the log scaled transform of the wealth access cutoff $\bar{w}_t$. This law of motion results in a *stationary* distribution of x on the real line, so we do not need to move our grid for x over time when conducting simulations of wealth accumulation.

Ultimately, the wealth distribution is governed, firstly, by the base rate of return on wealth net of economic growth, that is, the well-known "r minus g" term, $r_t^m - g_t$. Secondly, the wealth distribution depends on the wealth-accumulation wedge, Δ_t, and the access metric, here determined by $\bar{x}_t$. Finally, the time discount rate ρ and the idiosyncratic volatility σ affect the distribution of wealth, including the pace at which changes in $r_t^m - g_t$, Δ_t, and $\bar{x}_t$ translate into a change in the wealth distribution over time. In the following section, we explain how we calibrate and estimate these model elements using real-world data.

4 Calibration and Estimation

Scale dependence in our model of wealth accumulation takes the form of exclusive access for the wealthy to the fruits of innovation created by private firms. Our goal is to explain the increase in wealth inequality in recent decades, with a particular focus on the top 1% share of wealth. As a natural first step, we assume that the access cutoff, $\bar{w}_t$, is set such that only individuals in the top 1% of the wealth distribution have access to private innovation value creation. Under this assumption,

the trend in the wealth-accumulation wedge, Δ_t, between individuals in the top 1% and the rest is governed by three factors: the dynamics of corporate innovation value creation (shown in Figure III), the dynamics of the private innovation share (e_t in Figure IV), and (because Δ_t is expressed in after-tax terms) the dynamics of the marginal tax rate on long-term capital gains (τ_t^p in Figure VIII).

Formally, denote by V_t corporate innovation value creation, measured as the aggregate value of corporate patents in year t net of R&D expenses incurred in year t (that is, the year-by-year difference between the two red lines in Panels (a) and (b) of Figure III). Similarly, denote by V_t^l and V_t^u innovation value creation by stock market listed (public) and unlisted (private) firms in year t, respectively. By definition, V_t^l and V_t^u are related to V_t by:

$$V_t^l = (1 - e_t) \cdot V_t,$$
$$V_t^u = e_t \cdot V_t.$$

In what follows, we explain how we incorporate data on the time-varying private innovation share e_t and corporate innovation value creation V_t into the simulation of wealth dynamics.

The top 1% have access to both public and private firms and thus can benefit from both public and private innovation value creation, i.e., they share in $V = V^l + V^u$. Our maintained assumption for now is that the other 99% can only invest in public firms and thus share only in V^l. We assume homogeneity in splitting public-firm value creation, such that all individuals in the economy share in V^l in proportion to their wealth.[15] We similarly assume homogeneity within the top 1%, such that individuals in the top 1% share in V^u in proportion to their wealth. As a result, the wealth accumulation rate is the same for all individuals within (but not across) the two groups, which we specify next.

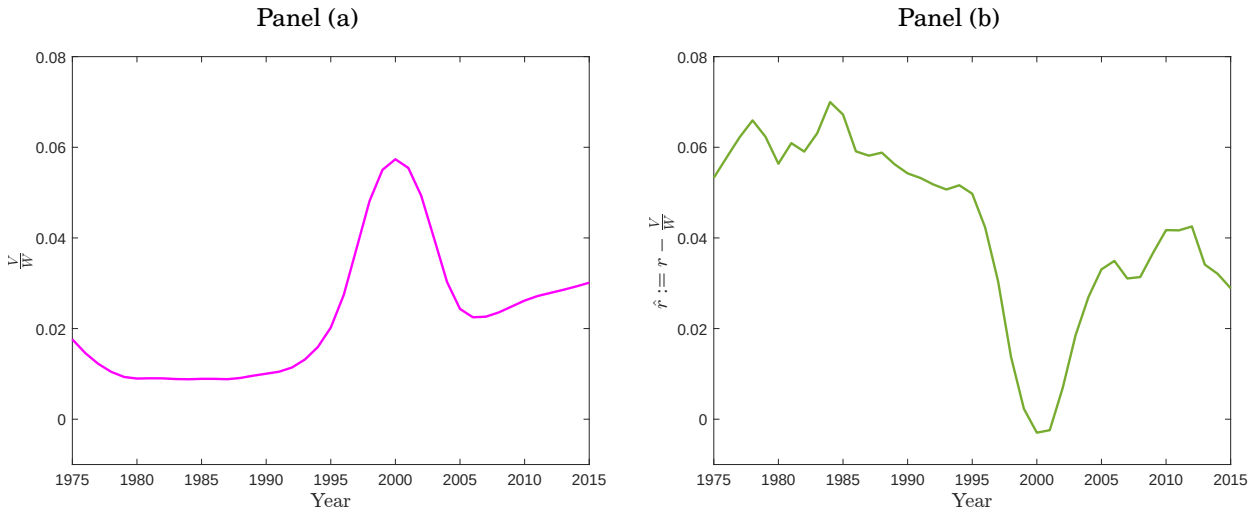
The pre-tax rate of return on wealth in the economy, r, consists of two components: one that is related to corporate innovation value creation, and another (denoted $\hat{r}_t$) that captures all other sources of capital income. Recall from Section 2.2.1 that Kogan et al. (2017) measure the value of a patent based on the market-adjusted change in a firm's stock market value when the patent office

[15] In other words, we are assuming that the top 1% do *not* substitute investments in private firms for investments in public firms. This assumption is conservative to the extent that substitution would amplify the wealth-accumulation wedge Δ_t, leading to a larger rise in the top 1% wealth share.

approves the firm's patent application. Ignoring returns from other sources of capital income ($\hat{r}_t$), the value of the representative investor's wealth in year t increases by their share in the aggregate Kogan et al. value of corporate patents granted that year. At the same time, their wealth decreases by their share in that year's corporate R&D expenses (which fund the development of future patentable inventions). Since V_t is defined as the difference between the value of patents granted in year t and R&D expenses incurred in year t, we have that V_t/W_t captures the percent change in aggregate wealth due to corporate innovation value creation. We refer to V_t/W_t as the "innovation wealth change."

By definition, the pre-tax rate of return on wealth in the economy, r_t, equals $\frac{V_t}{W_t} + \hat{r}_t$. Figure IX plots its two components, the innovation wealth change, $\frac{V}{W}$, in Panel (a) and the wealth change from all other sources of capital income, $\hat{r}_t = r_t - \frac{V_t}{W_t}$, in Panel (b).

Figure IX: Disaggregating the aggregate pre-tax return on wealth



Note: Panel (a) plots V_t/W_t, the percent change in aggregate wealth due to corporate innovation value creation. V_t is the aggregate imputed value of all U.S. patents granted to U.S. firms net of R&D spending, as plotted in Figure III. W_t is net aggregate private wealth (i.e., financial and non-financial assets net of financial liabilities). We use a 10-year moving average to smooth the data. Panel (b) plots $\hat{r}_t = r_t - V_t/W_t$, the aggregate return on wealth from all sources of capital income other than corporate innovation value creation.

Disaggregating r into $\hat{r}$ and V/W, we can write the post-tax return on wealth for the top 1% and

for the other 99% as

$$\textbf{Top 1\%:} \quad r_t^m + \Delta_t = (1-\hat{\tau}_t)\hat{r}_t + (1-\tau_t^p)\left(\frac{V_t^l}{W_t} + \frac{V_t^u}{W_t^{top}}\right),$$

$$\textbf{Bottom 99\%:} \quad r_t^m = (1-\hat{\tau}_t)\hat{r}_t + (1-\tau_t^p)\frac{V_t^l}{W_t},$$

where W_t is aggregate wealth and W_t^{top} is the wealth of individuals in the top 1%. Recall that we assume that wealth changes due to corporate innovation value creation constitute long-term capital gains and so are taxed at the marginal long-term capital gains tax rate τ^p. Returns from all other sources of capital income, $\hat{r}$, are taxed at the weighted-average tax rate $\hat{\tau}$, which differs from τ introduced in Section 2.4.2 only because the weights need adjusting for the fact that we decompose total capital income into two components: $\hat{r}$ and V/W.[16]

Using the definitions of V^u and V^l, we can rewrite the post-tax return on wealth for the top 1% and for the other 99% as a function of the private innovation share (e_t) and the innovation wealth change (V_t/W_t):

$$\textbf{Top 1\%:} \quad r_t^m + \Delta_t = r_t^{ave} + (1-\tau_t^p)\, e_t \frac{V_t}{W_t}\left(\frac{1}{s_t^{top}} - 1\right), \tag{3}$$

$$\textbf{Bottom 99\%:} \quad r_t^m = r_t^{ave} - (1-\tau_t^p)\, e_t \frac{V_t}{W_t}, \tag{4}$$

where

$$r_t^{ave} = (1-\hat{\tau}_t)\hat{r}_t + (1-\tau_t^p)\frac{V_t}{W_t} = (1-\tau_t)r_t \tag{5}$$

is the average after-tax return on aggregate wealth in the economy, and $s_t^{top} = W_t^{top}/W_t$ is the wealth share of the top 1%.

Owing to their exclusive access to private innovation value creation, the top 1% will accumulate wealth at a higher rate given by the wealth-accumulation wedge, Δ_t:

$$\Delta_t = (1-\tau_t^p)\, e_t \frac{V_t}{W_t}\frac{1}{s_t^{top}}. \tag{6}$$

[16] If $\tau^p = \hat{\tau} = \tau$, we can verify that $r_t^m(W_t - W_t^{top}) + (r_t^m + \Delta_t)W_t^{top} = (1-\tau_t)r_t W_t$, i.e., the average post-tax return on aggregate wealth in the economy is $(1-\tau_t)r_t$, as expected.

In our model, Δ_t governs the dynamics of wealth inequality, and in particular the top 1% wealth share. A equation (6) makes clear, Δ_t has four determinants: it is decreasing in the tax on long-term capital gains (τ_t^p), increasing in the private innovation share (e_t) and in the innovation wealth change ($\frac{V_t}{W_t}$), and decreasing in the top 1% wealth share (s_t^{top}).[17]

Figure X, Panel (a) plots the term $(1-\tau_t^p)\, e_t\, \frac{V_t}{W_t}$, the after-tax percent change in aggregate wealth that is due to private innovation value creation—the private innovation wealth change for short. It is this term that is the source of the wealth-accumulation wedge Δ_t, as derived in equation (6). The figure shows that the private innovation wealth change has risen since the 1990s, to values of around 1 percentage point a year.

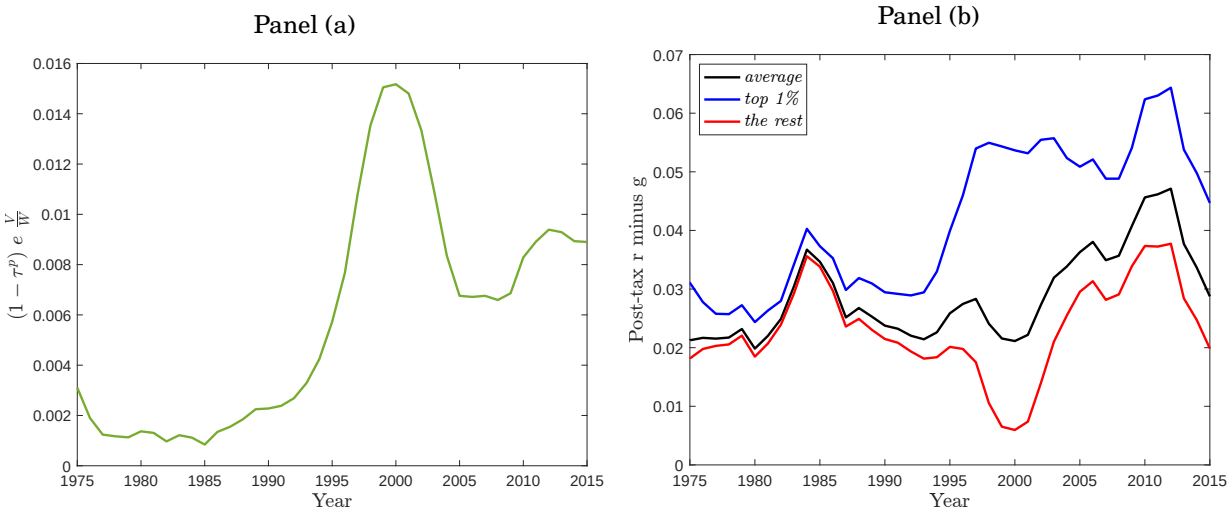
Figure X, Panel (b) uses the rates of returns on wealth derived in equations (3) to (5) to plot the "r minus g" term that is at the heart of existing models of wealth inequality dynamics, alongside the time series of our measures of the returns earned by the top 1%, $r_t^m + \Delta_t$, and by the other 99%, r_t^m (net of economic growth g_t). Over the period from 1975 to 2015, "r minus g" ranges between 2.1% and 4.8%, with an upward trend over time. The series for the top 1% is consistently above the "r minus g" line while the series for the bottom 99% is consistently lower. The returns of the top 1% have trended upward, from 2.8% in 1979 to 4.5% in 2015. We see no comparable upward trend for the bottom 99% over the post-1979 period. The divergence reflects the fact that the top 1% have increasingly earned "exclusive returns" as both the value of corporate innovation and the share of corporate innovation that is generated by private firms have increased over time.

Calibrating σ and estimating ρ. Since our goal is to explain the *rise* in wealth inequality in the past few decades, a natural initial steady-state is 1979, the year in which wealth inequality began to rise (see Figure II). We estimate the parameter of the time discount rate ρ by matching the top 1% wealth share in the steady state of the model with data in 1979.

We obtain the steady-state distribution of wealth in the model as follows. First, we construct r^m and Δ as of 1979 from equations (3) to (5), using data on the (pre-tax) return on capital r_t, the innovation wealth change $\frac{V_t}{W_t}$, the private innovation share e_t, the tax rates τ_t and τ_t^p, and the

[17] The last comparative static follows because at higher s_t^{top}, the private innovation wealth change, $e_t \cdot \frac{V_t}{W_t}$, is shared among individuals with greater wealth, which reduces the rate of wealth accumulation for each of them.

Figure X: Heterogeneous returns on wealth



Note: Panel (a) plots the after-tax return on wealth that is derived from exclusive access to private innovation value creation, $(1-\tau^p)\ e\ \frac{V}{W}$, which is the source of the wealth-accumulation wedge Δ_t between the top 1% and the rest. Panel (b) plots the estimated rates of return for the top 1% of the wealth distribution and for the other 99%, as well as the average rate of return in the economy net of the economic growth rate g_t. Returns are estimated from equations (3) to (5) using data on s_t^{top}, r_t, $\frac{V_t}{W_t}$, e_t, τ_t, and τ_t^p.

top 1% wealth share s_t^{top}. We then simulate the dynamics of log-scaled wealth x_{it} in equation (2) forward until converging to a steady-state distribution for x. Our starting point in this simulation is the steady-state distribution of wealth in the basic model without exclusive access ($e = 0$), for which Gabaix et al. (2016) provide a unique, analytically derived steady-state distribution. At each point in the dynamic simulation, we then obtain the top 1% wealth cutoff $\bar{x}_t$ for exclusive access (see equation 2) using the *lagged* simulated distribution. We stop this dynamic forward simulation when the wealth distribution stabilizes, which gives the *steady-state* distribution with exclusive access.

We find an estimate of ρ by iterating over alternative values until the resulting top 1% wealth share in the steady-state distribution matches the data in 1979. In the model with exclusive access to private innovation value creation, the estimated value of ρ is 0.052. This value corresponds to a time discount factor of $1/(1+0.052) \approx 0.95$, which is a standard calibration in the macroeconomics literature. In the basic model without exclusive access ($e = 0$), the estimate of ρ is slightly lower, at around 0.049.[18] Using these two estimates of ρ and starting from the two steady-state distributions, we simulate the *dynamics* of the wealth distributions in the basic model without exclusive access

[18]This is the case because the basic model generates a lower top 1% wealth share when fixing ρ, so in order to match the data, the estimation converges to a lower value of ρ to raise the simulated top 1% share back to the 1979 target level in the data.

and in our model with exclusive access to private innovation value creation.

Measuring the idiosyncratic return volatility, σ, in the data is challenging, as it requires micro data on individuals' balance sheets. Bach, Calvet, and Sodini (2020) report values as high as 25% for wealthy Swedish households, while Fagereng et al. (2020) find values in the range of 10-20% for Norwegians. Gabaix et al. (2016) use a calibration of 30%, which they note is at the upper end of values estimated (or calibrated) in the U.S. context. We conservatively calibrate σ to a high value of 30% and consider robustness to alternative calibrations in Section 6.[19]

5 Results

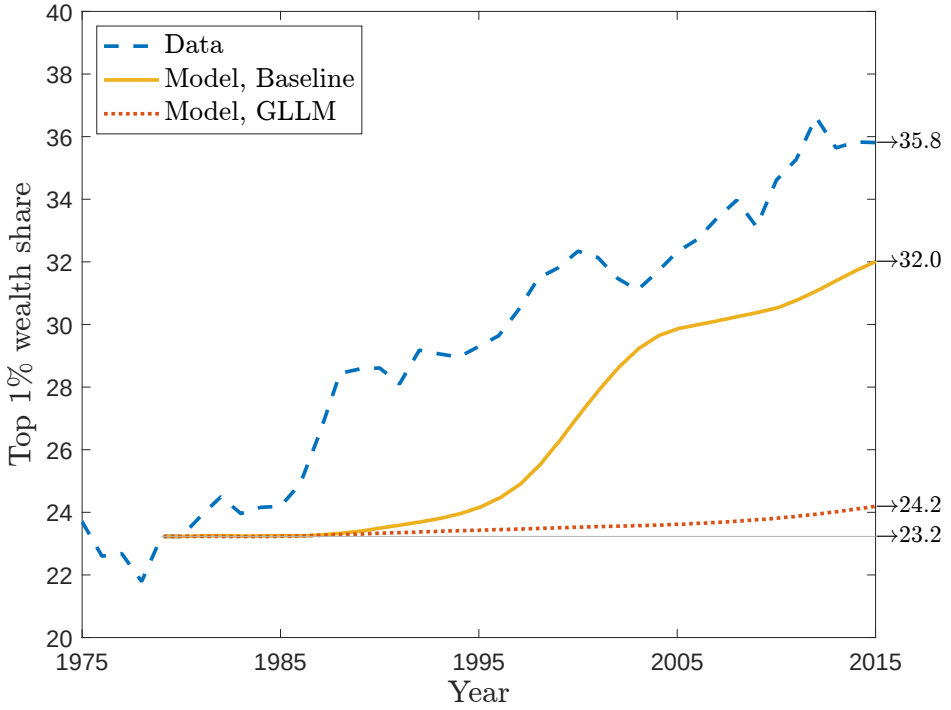
5.1 Baseline Findings

In this section, we present our main results for the dynamics of wealth inequality when the top 1% have exclusive access to private innovation value creation. Figure XI shows the dynamics of the top 1% wealth share in the data, alongside the simulated dynamics from our model. For comparison, the figure also includes the simulation results from the basic model of Gabaix et al. (2016) without exclusive access (GLLM). Both models are calibrated to match the observed top 1% wealth share in 1979 in their respective steady states, as described in Section 4. From this initial steady state, the dynamic simulation in our model is done by feeding the average wealth accumulation rate r_t^m and the wealth-accumulation wedge Δ_t at each point in time into model equation (2) in Section 3. We construct r_t^m and Δ_t from equations (3) to (5) in Section 4, in which we use data values of r_t^{ave}, τ_t^p, V_t/W_t, and e_t. We use the simulated distribution of wealth with a lag to set the top 1% wealth share, s_t^{top}. The (log-scaled) wealth access cutoff $\bar{x}_t$ in equation (2) is also obtained from the lagged simulated distribution at each time. The GLLM simulation corresponds to a special case of our model with $e_t = 0$.

We find that exclusive access to private innovation value creation can quantitatively account for a large part of the increase in wealth inequality. Figure XI shows that, in the data, the top 1% wealth share rises by 12.6 percentage points (from 23.2% to 35.8%) between 1979 and 2015. The model with exclusive access predicts an increase of 8.8 percentage points (from 23.2% to 32.0%) over the same

[19] As we will show, a high value for σ is conservative because in our framework, a higher value of σ weakens the impact of exclusive access to private innovation value creation on wealth inequality.

Figure XI: Dynamics of wealth inequality



Note: The figure shows the evolution of the top 1% wealth share in the data and in two simulation models: our baseline, featuring exclusive access to private innovation value creation; and another without exclusive access, as presented in Gabaix et al. (2016). The two models are separately estimated to match the top 1% wealth share in 1979 by their respective steady states; they are calibrated with data for the wealth accumulation rate from 1979 to 2015.

period. In other words, the widening wealth-accumulation wedge (between the top 1% and the rest) shown in Figure X, Panel (a) can explain as much as 70% of the observed increase in the top wealth share.[20]

Importantly, Figure XI shows that the basic model without exclusive access (GLLM) fails to generate a rise in wealth inequality. The predicted increase is just 1 percentage point (from 23.2% to 24.2%). The basic model only relies on the movements of the average post-tax return on capital net of economic growth, "r minus g" in the data. However, from the late 1970s to the mid-2010s, "r minus g" does not increase by much (see Figure X, Panel b). First, despite the relatively stable level of the pre-tax return from the 1970s until the 1990s, this return generally falls after 2000 (see Figure VI). Second, our calculation of the capital tax for the wealthy takes into account the significant share of capital gains and dividends besides interest income as part of the capital income, and, unlike the

[20] Recall that both the increase in aggregate wealth due to corporate innovation, V_t/W_t, and the increase in the private innovation share, e_t, contribute to the wealth-accumulation wedge; see equation (6). In Section 5.2, we will show that each plays a nontrivial role in explaining the rise in the top wealth share.

interest income tax, capital-gains and dividend tax rates do *not* fall much during the period 1970s to 2010s (see Figure VIII). In the end, "r minus g" does not increase significantly during this period on a *post-tax* basis. Thus, the resulting increase in the top wealth share in the GLLM simulation is minimal.

The divergence between the two models underscores the central role of exclusive access to private innovation value creation in shaping the wealth distribution. While a model with uniform returns fails to reproduce the empirical trajectory of inequality, incorporating exclusive access to private innovation value creation substantially improves the model's explanatory power. These results highlight the importance of heterogeneity in access across different segments of the wealth distribution, not just a general rise in return on capital in the economy.

5.2 Counterfactual Analyses

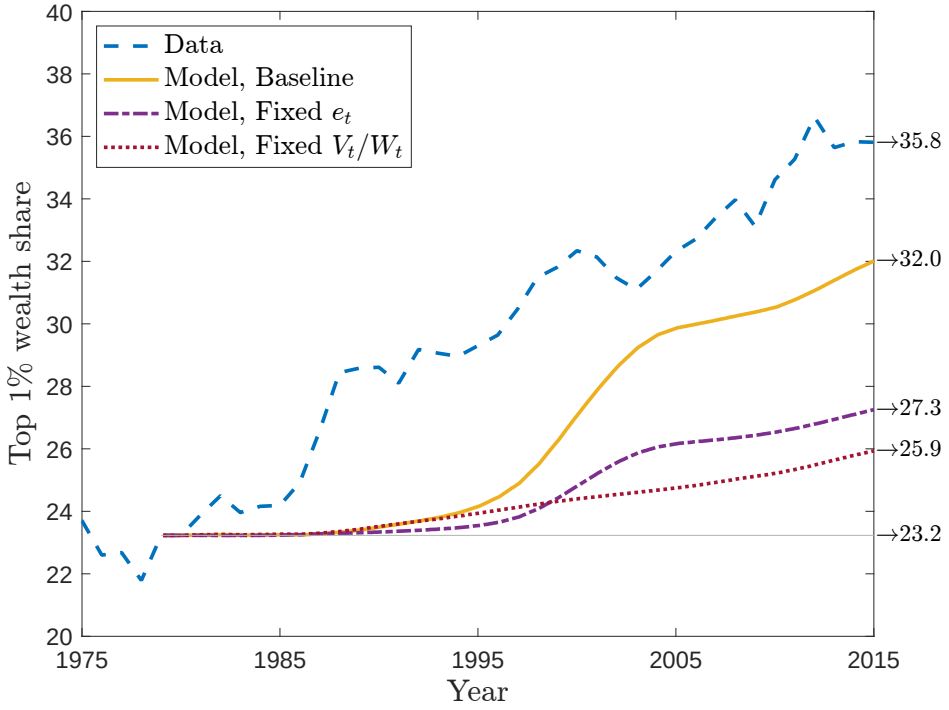
5.2.1 Private Innovation Wealth Change

Given exclusive access, the wealth-accumulation wedge between the top 1% and the rest, Δ_t, depends on two factors: 1) the percent change in aggregate wealth due to corporate innovation value creation, V_t/W_t, and 2) private firms' share of corporate innovation value creation, e_t. In this section, we first show that increases in both V/W and e contribute to the simulated rise in wealth inequality documented in Section 5.1 and then quantify their relative contribution. To do so, we run counterfactual analyses that separately change V/W and e from the observed values in the data.

Figure XII graphs the simulated dynamics of the top 1% wealth share under two scenarios: one that holds the wealth change due to innovation, V_t/W_t, fixed at its initial 1979 value and lets the private innovation share e_t vary as per the data, and another that holds e_t fixed at its initial 1979 value and lets V_t/W_t vary as per the data. The figure shows that both V_t/W_t and e_t have a sizable impact on the top 1% wealth share. Specifically, with e_t held constant, the upward trend in V_t/W_t seen in Figure X would on its own raise the top 1% wealth share by 4.1 percentage points (from 23.2% to 27.3%), while the increase in e_t seen in Figure IV would on its own raise the top 1% wealth share by 2.7 percentage points (from 23.2% to 25.9%).

The counterfactuals shown in Figure XII imply that the combined effect of increases in V_t/W_t and in e_t is greater than each increase's individual effect. In isolation, changes in V_t/W_t and e_t give

Figure XII: Counterfactual analyses of the wealth-accumulation wedge



Note: The figure shows the evolution of the top 1% wealth share in the data, in our baseline model, and in counterfactual analyses of the wealth-accumulation wedge. Each model is estimated to match its steady-state top 1% wealth share with data in 1979. The counterfactual analyses hold either the private innovation share (e_t) or the innovation wealth change (V_t/W_t) fixed at its initial 1979 value.

rise to a $4.1 + 2.7 = 6.8$ percentage-point increase in the top 1% wealth share. This is smaller than the 8.8 percentage-point increase in our baseline simulation, which combines changes in V_t/W_t and e_t. The difference of 2 percentage points reflects complementarity in the effects of V_t/W_t and e_t: a given increase in corporate innovation value creation V_t has a larger effect on wealth inequality the higher the private-firm innovation share e_t; and a given increase in e_t has a larger effect on wealth inequality the more valuable is corporate innovation. The empirical fact that both V_t/W_t and e_t have been trending up over the 1979-2015 period means that both trends reinforce each other, resulting in a larger increase in the top 1% wealth share.

5.2.2 The Exclusive Access Criterion

Our baseline simulation assumes that access to private innovation wealth creation, $e_t V_t/W_t$, is restricted to the top 1% in every year. This assumption can be motivated based on the empirical fact that the concentration of private-business equity is remarkably constant over time (see Figure V). In

other words, in the data, it is not the case that individuals in the lower wealth percentiles increase their share of private-business equity over the sample period. In this section, we consider alternative access criteria that expand, over time, the set of individuals with access to private innovation wealth creation. As a result, wealth should become less concentrated at the top.

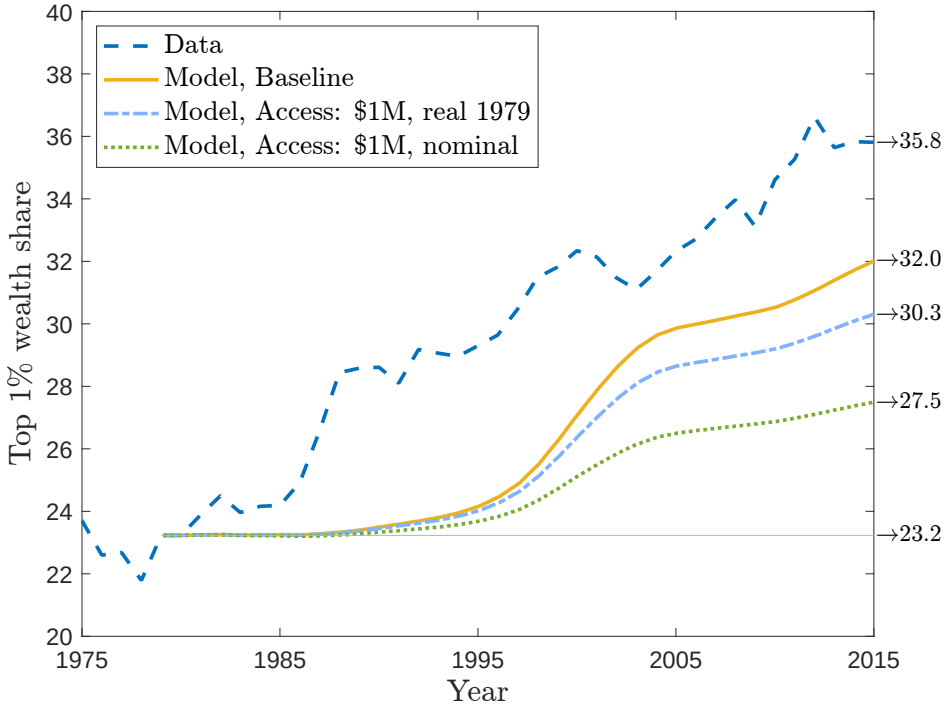
We consider two counterfactual analyses. The first captures the fact that the SEC's "accredited-investor" rules are based on a *nominal* threshold (here: $1 million in net worth not including the primary residence). This threshold roughly corresponds to the minimum wealth of the top 1% in 1979 (same as in our baseline model), but over time, asset inflation will naturally expand the set of accredited investors beyond the top 1%.

The second counterfactual analysis considers the possibility that investment in private firms requires a fixed cost (perhaps a search cost) whose size increases with inflation but not with wealth. We implement this counterfactual by restricting access to private innovation value creation to those individuals whose wealth in a given year exceeds the minimum wealth in 1979's top 1% cohort (about $1 million) in *real* terms.

Figure XIII reports the two counterfactual simulations and compares them to the data and to the baseline model. As expected, wealth becomes less concentrated at the top when a broader set of individuals (beyond the top 1%) gain access to private innovation value creation over time. In particular, the top 1% wealth share increases from 23.2% in 1979 to 27.5% in 2015 under the accredited-investor scenario of a nominal wealth cutoff—less than the increase to 32.0% in the baseline. This implies that the SEC's accredited-investor rules are *not* the binding constraint on access to private innovation value creation and so cannot explain the rise in wealth inequality.

Restricting access based on a real wealth cutoff of $1 million in 1979 dollars comes closer to the baseline increase: the top 1% wealth share increases from 23.2% in 1979 to 30.3% in 2015 in this counterfactual. The shortfall compared to the baseline increase to 32.0% suggests that frictions that keep access tied to the *relative* wealth in the cross-section—rather than absolute wealth levels—play a central role in explaining the rise in wealth inequality over time.

Figure XIII: Counterfactual analyses of the access criteria



Note: The figure shows the evolution of the top 1% wealth share in the data, in our baseline model, and in counterfactual analyses of the access criteria. Each model is estimated to match its steady-state top 1% wealth share with data in 1979. In the baseline model, access to private innovation value creation is restricted to the top 1% in a given year. The counterfactual analyses instead restrict access to those individuals whose wealth in a given year exceeds the minimum wealth in 1979's top 1% cohort (about $1 million), either in nominal or in real terms.

5.3 Policy Experiment: Taxes and Subsidies

We study four tax-policy counterfactuals to assess how taxation interacts with innovation-driven wealth accumulation. Figure VIII in Section 2.4.2 summarizes the evolution of U.S. taxes on interest income, dividends, and long-term capital gains since 1979.

Table I: Counterfactual tax policies

	Baseline	fixed τ^p at 1979 value	fixed τ^i & τ^d at 1979 value	$\downarrow \tau^p$ to 0 for top 1%, by 2015	remove 10% R&D subsidy
	(1)	(2)	(3)	(4)	(5)
Top 1% wealth share (%), 1979	23.2	23.2	23.2	23.2	23.2
Top 1% wealth share (%), 2015	32.0	30.8	30.7	34.6	31.0
Change (percentage points)	8.8	7.6	7.5	11.4	7.8

Note: The table reports the predicted rise in our baseline simulation and in counterfactual scenarios of the tax rate on alternative sources of capital income (τ^i_t, τ^d_t, and τ^p_t) and the subsidy on R&D expenses in the data.

Capital gains taxation. Capital gains taxes declined substantially over the period, with the statutory rate τ_t^p falling from 26.5% in 1979 to a low of 15% in 2006–2012. Holding τ_t^p fixed at its 1979 level reduces the increase in the top 1% wealth share from 8.8 percentage points in the baseline to 7.6 percentage points (Table I, column (2)), indicating a meaningful contribution of capital gains tax reductions to rising wealth inequality.

Taxes on interest income and dividends. In a second counterfactual, we hold the tax rates on interest income and dividends, τ_t^i and τ_t^d, fixed at their 1979 levels. Despite much larger declines in these rates in the data, this counterfactual yields a reduction in top wealth accumulation no greater than in the firt counterfactual: the top 1% share rises by 7.5 percentage points (Table I, column (3)). The reason is that capital gains taxes directly governs the wealth-accumulation wedge at the top, and long-term capital gains account for an increasing share of capital income (Figure VII).

Zero effective capital gains taxation at the top. We next consider a polar counterfactual motivated by the "buy–borrow–die" mechanism, under which the effective capital gains tax rate for the top 1% declines linearly to zero by 2015 (Fox and Liscow, 2025). Under this scenario, the top 1% wealth share rises by 11.4 percentage points, reaching 34.6% in 2015—close to the observed value of 35.8% (Table I, column (4)). This experiment shows that exclusive access to private innovation value combined with very low effective capital gains taxation can account for nearly the entire rise in top wealth concentration.

R&D tax credits. Finally, we evaluate the distributional effects of R&D tax credits, which average about 10% since the introduction of the federal research credit in 1981 (Hall and Van Reenen, 2000). Eliminating the subsidy and assuming a one-for-one reduction in innovation value creation lowers the increase in the top 1% wealth share to 7.8 percentage points (Table I, column (5)). In the model, R&D tax incentives therefore contribute to rising wealth inequality by amplifying innovation-driven wealth accumulation at the top.

6 Robustness

We consider three key robustness tests to the baseline analysis presented in previous sections and assess their implications for explaining the rise in wealth inequality. By way of preview, we find that, across the alternative specifications we consider, exclusive access to private innovation value creation can explain at least *half* of the observed rise in top wealth shares from 1979 to 2015.

6.1 The Idiosyncratic Return on Wealth

The idiosyncratic return on wealth, controlled by the parameter σ in equation (1), affects the extent to which Δ_t, the wealth-accumulation wedge between those at the top of the wealth distribution and the rest, leads to rising wealth inequality over time. Larger idiosyncratic shocks to wealth σ generate more mobility into and out of the top 1%. Thus, when σ is larger, the private innovation value creation to which individuals in the top 1% have access will leak to lower percentiles over time, resulting in a *smaller* predicted rise in the top 1% wealth share.[21]

In this section, we consider the robustness of our findings to alternative choices of σ. Recall that our baseline calibration sets $\sigma = 0.3$, which, as discussed in Section 4, is at the upper end of estimates reported in the literature and so is conservative. Here, we consider two alternative choices: $\sigma = 0.2$ (which corresponds to empirical estimates for Swedish and Norwegian households) and an extreme case of $\sigma = 0.4$.[22]

Table II summarizes our findings. The table reports the top 1% wealth shares in 1979 and in 2015 and the change between 1979 and 2015. Column (1) shows the realizations from the data while column (2) shows our baseline estimates. Columns (3) and (4) report results for $\sigma = 0.2$ and $\sigma = 0.4$, respectively. Our results are sensitive to the choice of σ in the expected direction. The lower $\sigma = 0.2$ calibration generates a larger rise in the top 1% wealth share of 10.4 percentage points between 1979 and 2015 (compared to the 8.8 percentage-point baseline estimate). This confirms that our baseline estimates are conservative.

[21] In the basic model without exclusive access (GLLM), the change in the predicted rise goes in the opposite direction. In the GLLM model, idiosyncratic shocks to wealth σ and a rising "r minus g" are the only source of wealth inequality. Greater dispersion in returns as σ increases then leads to a *larger* rise in the top wealth share over time.

[22] Given a new σ, we re-estimate the parameter ρ of our model with exclusive access to private innovation value creation by re-matching the steady state of the model in 1979. We then re-simulate the economy over time.

Even with an extremely (and arguably implausibly) high value of $\sigma = 0.4$, our model of exclusive access to private innovation value creation predicts a substantial rise in the top 1% wealth share of 7.2 percentage points—only marginally smaller than the baseline estimate of 8.8 percentage points and still large in relation to the realized rise of 12.6 percentage points in the data.

Table II: Robustness tests: σ and $\bar{w}_t$

	Data	Model				
		Baseline	$\sigma = 0.2$	$\sigma = 0.4$	Access for top 5%	Access for top 10%
	(1)	(2)	(3)	(4)	(5)	(6)
Top 1% wealth share (%), 1979	23.2	23.2	23.2	23.2	23.2	23.2
Top 1% wealth share (%), 2015	35.8	32.0	33.6	30.4	30.2	29.6
Change (percentage points)	12.6	8.8	10.4	7.2	7.0	6.4

Note: The table reports the realized rise from 1979 to 2015 in the top wealth share in the data (column 1) along with the predicted rise in our baseline simulation (column 2) and in alternative specifications that vary the parameters for idiosyncratic shocks to wealth, σ, and for who has access to private innovation value creation, $\bar{w}_t$. In the baseline model, σ is calibrated to 0.3, and access is restricted to individuals in the top 1%.

6.2 Broader Access to Private Innovation Value Creation

Our baseline model sets $\bar{w}_t$ such that access to private innovation value creation is restricted to individuals in the top 1% of the wealth distribution. Data on who invests in the private firms that create innovation value are not available for the U.S., but we can view the SCF data on private-business equity in Figure V as a lower bound. The reason is that private-business equity among the less well off more likely represents self-employment (e.g., plumbers or medical practices) than investments in risky innovative private firms (investment in which is restricted by the SEC's accredited-investor rules). Naturally, broadening the access criterion beyond the top 1% dilutes the top 1%'s wealth-accumulation wedge, Δ_t, which in turn will reduce the predicted rise in the top 1% wealth share. The empirical question is, by how much?

In this section, we consider the robustness of our findings to alternative access assumptions. Specifically, we broaden access to private innovation value creation to individuals in either the top 5% or the top 10%, on the following terms (chosen to match the patterns in Figure V): we assume that individuals in the top 1% have access to all private innovation value creation $e_t V_t$ (as in the baseline), that individuals in percentiles 95-99 have access to half of $e_t V_t$, and that individuals in percentiles

90-94 have access to a quarter of $e_t V_t$. We calculate each individual's wealth accumulation rate by distributing aggregate $e_t V_t$ in proportion to the product of the individual's wealth and the calibrated access weight of her wealth percentile. We then re-estimate the model using the steady-state top 1% wealth share outcome in 1979.

Table II, columns (5) and (6) report the results. Whether access is broadened to the top 5% or to the top 10%, our model predicts nearly the same rise in the top 1% wealth share from 1979 to 2015, of 7.0 and 6.4 percentage points, respectively. Consistent with the expected dilution, these estimates are smaller than the baseline estimate of 8.8 percentage points. Still, even when we relax our baseline access criterion, the model can explain at least half of the observed 12.6 percentage-point rise in the top 1% wealth share.

6.3 Within-Top Wealth Inequality

Our model is designed to explain the rise in the concentration of wealth in the right tail of the distribution. To do so, the model focuses on return heterogeneity, which Benhabib and Bisin (2018), Benhabib, Bisin, and Luo (2019), and Hubmer, Krusell, and Smith (2021) show is crucial to understanding the rise in the top wealth share. To keep the analysis tractable, we abstract from earnings heterogeneity and progressive income taxation, which are required to generate realistic outcomes for the rest of the wealth distribution (which is not our focus).

In this section, we narrow our focus to the right tail (for which our assumptions are considered the most appropriate) and model *within-top* wealth inequality, which we measure as the ratio of the wealth of the top 1% to the wealth of the top 10%. We proceed as follows. We first re-estimate the time discount rate in individuals' preferences, ρ, by matching the ratio of the wealth of the top 1% to the wealth of the top 10% in 1979 and then re-simulating wealth accumulation over time. The new estimate of ρ is 0.067, which is larger than the value 0.052 in the baseline.

Table III reports the results for within-top wealth inequality. In the data, the top 1% increase their share of wealth in the top 10% by 13.4 percentage points between 1979 and 2015, from 35.5% to 48.9%. Our baseline model, matched to the top 1% wealth share in 1979, predicts an increase of 10.5 percentage points. The alternative specification, matched to the ratio of the wealth shares of the top 1% and the top 10% in 1979, predicts a 12.8 percentage-point increase. In other words, our

model of exclusive access to private innovation value creation can explain almost the entire increase in the within-top inequality between 1979 and 2015.

Table III: Robustness test: Within-top inequality

	Data	Model	
		Baseline	Matching 1979's wealth share ratio
	(1)	(2)	(3)
Top 1% to top 10% wealth share ratio, 1979 (%)	35.5	45.1	35.5
Top 1% to top 10% wealth share ratio, 2015 (%)	48.9	55.6	48.3
Change (percentage points)	13.4	10.5	12.8

Note: The table shows robustness results for the predicted rise in the ratio of the top 1% and the top 10% wealth share. Column (1) shows the values in the data. Column (2) shows the baseline results, in which the model is estimated by matching the top 1% wealth share in 1979. Column (3) shows results for a model that matches the ratio of the top %1 and top 10% wealth shares in 1979.

7 Conclusion

Our analysis reveals a previously unmeasured mechanism linking the organization of corporate innovation to the distribution of wealth. By assembling a long-run series of innovation value that includes both public and private firms, and by documenting the sharp rise in the private-firm share of that value, we show that access to private-firm innovation is both highly concentrated and quantitatively important for the evolution of top wealth shares. Embedding this access channel into a structural model of wealth accumulation allows us to account for much of the observed rise in top wealth concentration since 1979. These findings suggest that changes in where innovation occurs—and who can invest in it—have first-order implications for inequality.

More broadly, our results highlight the importance of understanding how shifts in the frontier of corporate activity interact with household balance sheets. The growing role of private markets in generating innovation value raises questions about the allocative, distributional, and policy consequences of innovation occurring outside public capital markets. Future research could explore how access to private-firm equity varies across countries, how innovations financed through private markets diffuse into the broader economy, and how policy interventions could influence the distributional effects of innovation. By tracing return heterogeneity to concrete investment opportunities,

our framework opens the door to a richer understanding of how modern innovation ecosystems shape long-run wealth inequality.

References

Auerbach, Alan J, and Kevin Hassett. 2015. "Capital taxation in the twenty-first century." *American Economic Review*, 105(5): 38–42.

Bach, Laurent, Laurent E. Calvet, and Paolo Sodini. 2020. "Rich pickings? Risk, return, and skill in household wealth." *American Economic Review*, 110(9): 2703–47.

Batty, Michael, Jesse Bricker, Joseph Briggs, Sarah Friedman, Danielle Nemschoff, Eric Nielsen, Kamila Sommer, and Alice Henriques Volz. 2021. "The Distributional Financial Accounts of the United States." *Measuring Distribution and Mobility of Income and Wealth*, 641–677. University of Chicago Press.

Benhabib, Jess, Alberto Bisin, and Mi Luo. 2019. "Wealth Distribution and Social Mobility in the US: A Quantitative Approach." *American Economic Review*, 109(5): 1623–47.

Benhabib, Jess, Alberto Bisin, and Shenghao Zhu. 2011. "The distribution of wealth and fiscal policy in economies with finitely lived agents." *Econometrica*, 79(1): 123–157.

Benhabib, Jess, and Alberto Bisin. 2018. "Skewed wealth distributions: Theory and empirics." *Journal of Economic Literature*, 56(4): 1261–1291.

Cagetti, Marco, and Mariacristina De Nardi. 2006. "Entrepreneurship, frictions, and wealth." *Journal of political Economy*, 114(5): 835–870.

Canipek, Aras. 2024. "The Accredited Investor Definition, Private Investments, and Wealth Inequality in the US." *Private Investments, and Wealth Inequality in the US (November 06, 2024)*.

Catherine, Sylvain, Max Miller, and Natasha Sarin. 2025. "Social security and trends in wealth inequality." *Journal of Finance*, 80(3): 1497–1531.

Catherine, Sylvain, Max Miller, James D Paron, and Natasha Sarin. 2023. "Interest-rate risk and household portfolios." *Jacobs Levy Equity Management Center for Quantitative Financial Research Paper*.

Ebrahimian, Mehran, and Paolo Sodini. 2024. "More than Money: The Role of Preferences on Wealth Mobility." *Available at SSRN 5577751*.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri. 2020. "Heterogeneity and persistence in returns to wealth." *Econometrica*, 88(1): 115–170.

Feenberg, Daniel, and Elisabeth Coutts. 1993. "An introduction to the TAXSIM model." *Journal of Policy Analysis and management*, 12(1): 189–194.

Fox, Edward, and Zachary Liscow. 2025. "The Role of Unrealized Gains and Borrowing in the Taxation of the Rich." *Law & Economics Working Papers*, , (286).

Gabaix, Xavier, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. 2016. "The dynamics of inequality." *Econometrica*, 84(6): 2071–2111.

Gocmen, Ararat, Clara Martínez-Toledano, and Vrinda Mittal. 2025. "Private Capital Markets and Inequality." *Available at SSRN 5166981*.

Gomez, Matthieu, and Emilien Gouin-Bonenfant. 2025. "Wealth Inequality in a Low Rate Environment." Unpublished Working Paper, Columbia University.

Greenwald, Daniel L, Matteo Leombroni, Hanno Lustig, and Stijn Van Nieuwerburgh. 2021. "Financial and total wealth inequality with declining interest rates." National Bureau of Economic Research.

Hall, Bronwyn H, Adam B Jaffe, and Manuel Trajtenberg. 2001. "The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools." National Bureau of Economic Research Working Paper 8498.

Hall, Bronwyn H., and John Van Reenen. 2000. "How Effective Are Fiscal Incentives for R&D? A Review of the Evidence." *Research Policy*, 29(4–5): 449–469.

Hubmer, Joachim, Per Krusell, and Anthony A Smith. 2021. "Sources of US wealth inequality: Past, present, and future." *NBER Macroeconomics Annual*, 35(1): 391–455.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy. 2021. "Measuring Technological Innovation over the Long Run." *American Economic Review: Insights*, 3(3): 303–20.

Kline, Patrick, Neviana Petkova, Heidi Williams, and Owen Zidar. 2019. "Who Profits from Patents? Rent-Sharing at Innovative Firms." *Quarterly Journal of Economics*, 134(3): 1343–1404.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. 2017. "Technological Innovation, Resource Allocation, and Growth." *Quarterly Journal of Economics*, 132(2): 665–712.

Lindsey, Laura Anne, and Luke CD Stein. 2019. "Angels, entrepreneurship, and employment dynamics: Evidence from investor accreditation rules."

Nix, Paul E., and David E. Nix. 1992. "A historical review of the accounting treatment of research and development costs." *Accounting Historians Journal*, 19(1): 51–78.

Piketty, Thomas, and Gabriel Zucman. 2014. "Capital is back: Wealth-income ratios in rich countries 1700–2010." *Quarterly Journal of economics*, 129(3): 1255–1310.

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

Quadrini, Vincenzo. 1999. "The importance of entrepreneurship for wealth concentration and mobility." *Review of Income and Wealth*, 45(1): 1–19.

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

Trajtenberg, Manuel, Henderson Rebecca, and Adam Jaffe. 1997. "University Versus Corporate Patents: A Window On The Basicness Of Invention." *Economics of Innovation and New Technology*, 5(1): 19–50.

Veihl, Maria. 2022. "Strategic patenting: The dark side of patents." Unpublished Working Paper.

Wachter, Jessica A. 2013. "Can time-varying risk of rare disasters explain aggregate stock market volatility?" *Journal of Finance*, 68(3): 987–1035.

Appendix

A Data

A.1 Aggregate Return on Capital

This section outlines the methodology used to generate the return series presented in Figure VI. The process involves updating and replicating the return series originally provided by Piketty and Zucman (2014) using contemporary data sources. The original series, which was used in Gabaix et al. (2016) for the return on wealth, spans from 1946 to 2010 and is based on a pre-tax rate of return.

The average rate of return on domestic capital is defined as $r_t = \theta_t/\beta_t$, where θ_t represents the capital share and β_t is the wealth-national income ratio. To construct the return series following Piketty and Zucman (2014), we decompose θ_t and β_t into their detailed components and update them using data from the Federal Reserve's Integrated Macroeconomic Accounts (IMA) and the National Income and Product Accounts (NIPA) maintained by the Bureau of Economic Analysis.

Capital share (θ_t): The capital share (θ_t) is defined as total capital income divided by total income. Total income is directly taken from the IMA as nominal national income. Constructing total capital income, however, is more complex. It consists of several components:

- **Corporate capital income**: This includes profits from corporations, which is derived from NIPA tables.
- **Housing capital income**: This is the net product of the housing sector, which is derived from NIPA tables.
- **Self-employment capital income**: This component captures the capital income of self-employed individuals, which is estimated using data on mixed income drived from NIPA and IMA.
- **Net foreign capital income**: This represents income from foreign investments, adjusted for outflows and inflows, and is sourced from NIPA Table 4.1.
- **Net government interest payments**: This includes interest payments made by the government, net of receipts, and is updated using NIPA Table 3.1.

Each of these components is further divided into detailed sub-components, which are updated using the latest available data from NIPA and IMA. Piketty and Zucman (2014) provide the formula for deriving most capital income components using IMA and NIPA as sources. However, in some cases, data from one source has been discontinued. For example, certain series, such as NIPA Table 7.12, are no longer available. We update the corresponding variable using only the remaining data source in such instances.

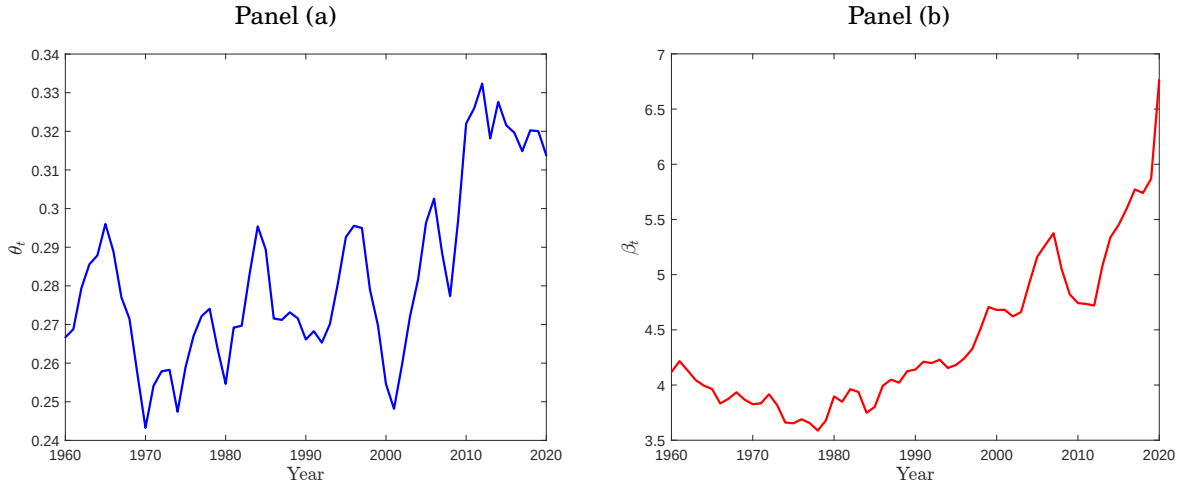
Although the differences are negligible when the series can be derived from both NIPA and IMA, we prioritize the formula based on NIPA data. The IMA-based formula is used only when deriving the series from NIPA is no longer possible. The replicated series for θ_t is presented in Figure A.1, Panel (a).

Wealth-national income ratio (β_t): The wealth-national income ratio (β_t) is defined as net private wealth divided by total income. As with θ_t, total income is taken directly from IMA as nominal national income. Constructing net private wealth involves a more detailed process. Net private wealth is calculated as the sum of financial and non-financial assets, net of financial liabilities. Each of these components is further broken down into detailed sub-components:

- **Financial assets**: This includes assets such as stocks, bonds, and other financial instruments.
- **Non-financial assets**: This includes real estate, durable goods, and other tangible assets.
- **Financial liabilities**: This includes debts such as mortgages, loans, and other liabilities.

Each of these components is updated using the most recent data available from IMA. The replicated series for β_t is presented in Figure A.1, Panel (b).

Figure A.1: Capital share of income and wealth-national income ratio



Note: This graph illustrates the capital share of income (θ_t) (in Panel a) and the wealth-to-national-income ratio (β_t) (in Panel b). The return series is derived from the ratio of these two variables.

A.2 Capital Income Tax

This appendix explains the methodology used to generate the tax-rate series presented in Figure VIII. We construct two time series for tax rates spanning the longest possible periods: one for the marginal tax rate on a composite capital income category, including interest income, dividends, and long-term capital gains, and another for the tax rate explicitly applied to long-term capital gains.

To construct these time series, we rely on the following two series:

- Marginal tax rates on interest income, dividends, and long-term capital gains, presented in Figure A.2, Panel (a).
- Income shares of these components, presented in Figure A.2, Panel (b).

The marginal tax rate on capital income is computed as a weighted average of the marginal tax rates of its components, using the income shares as weights.

Post-1988 data. For years after 1988, we use the SCF dataset and the NBER TAXSIM model to compute individual-level marginal tax rates. Since we require aggregate tax rates for the top 1% of income earners, we sort individuals in each year by their adjusted gross income (AGI) and select the top 1%. To compute the aggregate marginal tax rate for a specific income type, we calculate a weighted average of the marginal tax rates for the top 1% of individuals in each year, using the corresponding income shares derived from SCF as weights. Because SCF data is available only every three years, this procedure provides marginal tax rates for the years 1988 to 2019 at three-year intervals.

Pre-1988 marginal tax rates. Although the SCF has been conducted since 1962, data prior to 1988 is not publicly available. Therefore, we look for alternative sources for marginal tax rates before 1988. We find a report on the NBER TAXSIM website that presents marginal tax rates for the top 1% of income earners across various income types, including interest income, dividends, and long-term capital gains. These series are based on TAXSIM simulations using microdata from the Internal Revenue Service (IRS). The reported marginal tax rate series covers the years 1960 to 2000, making them a reasonable source for pre-1988 data. While using a consistent data source for both pre-and post-1988 periods would be ideal, this is not feasible.

Pre-1988 income shares. The required income shares are obtained from the World Inequality Database (WID), which provides historical data on fiscal capital income components, including:

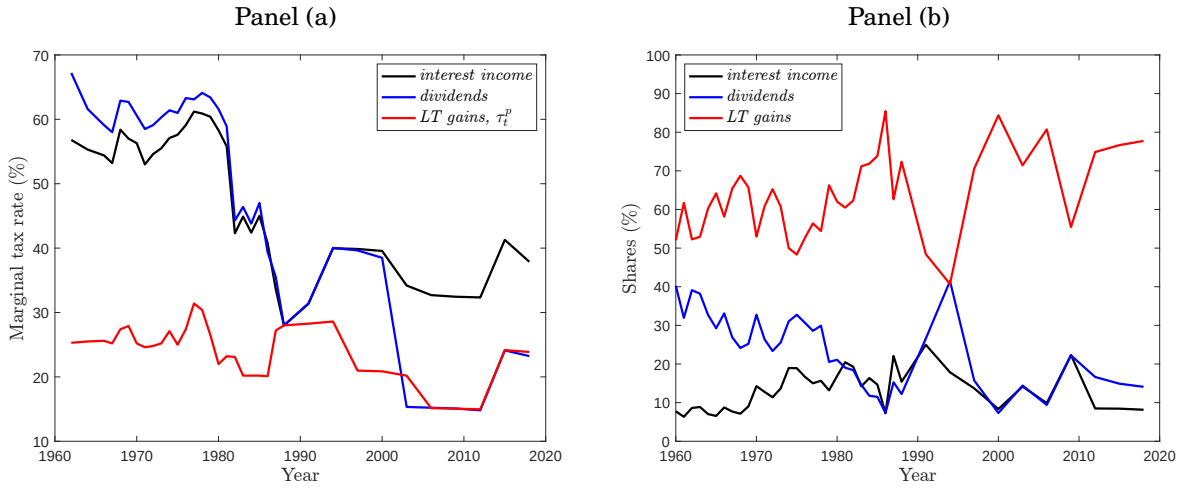
- Rents
- Interest
- Dividends
- Capital gains
- The capital component of mixed income

Adjustments for long-term capital gains. Since WID does not distinguish between long-term and short-term capital gains, we estimate long-term capital gains shares using historical data from the U.S. Department of the Treasury. We assume a stable 95% share for long-term capital gains in total capital gains, based on available data from 1977 to 1988.

Marginal tax rate for top 1% wealth holders. Following Gabaix et al. (2016), we construct tax series for the top 1% of income earners. Ideally, we would focus on the top 1% of wealth holders; however, data limitations make this infeasible. In particular, estimating marginal tax rates requires detailed microdata, which are available from the Survey of Consumer Finances (SCF) only from 1988 onward. For earlier periods, such data are not available. For the post-1988 period, we can compare marginal tax rates for the top 1% of wealth holders and the top 1% of income earners. To perform this comparison, we rank households in the SCF by net worth, select the top 1%, and compute marginal tax rates using the same methodology applied to the top 1% of income earners. Figure A.3 plots the

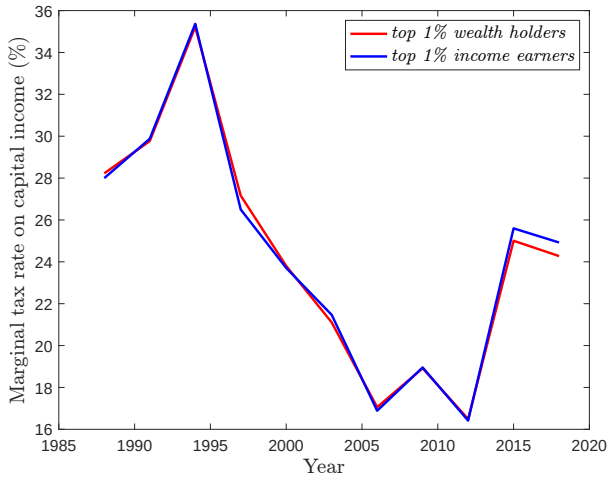
two series. The difference in tax rates between the two groups (the top 1% wealth holders and the top 1% income earners) is minimal.

Figure A.2: Marginal tax rates and capital income shares used in calculating the marginal tax rate on capital income



Note: Panel (a) shows the marginal tax rates for the different components of capital income used in our replication of the marginal tax rate on total capital income. For the period before 1988, the data source is a report from the TAXSIM website, while for 1988 onward, the rates are based on our calculations using SCF data. Panel (b) illustrates the shares of various capital income components in total capital income for top 1% income earners, which are used as weights to compute the overall marginal tax rate on capital income. Before 1988, the data source is WID, with minor adjustments to long-term capital gains. From 1988 onward, the shares are calculated using the SCF dataset.

Figure A.3: Comparison of marginal tax rates on capital income for the top 1% wealth and income



Note: For the red line, we restrict the SCF sample to the top 1% of wealth holders and calculate their marginal tax rates. Similarly, for the blue line, we limit the sample to the top 1% of income earners and perform the same calculation.

B Supplementary Tables and Figures

Table B.1: Private-business equity across the wealth distribution

	bottom 90%		90-94		95-99		top 1%	
SCF wave	prv-bus	assets	prv-bus	assets	prv-bus	assets	prv-bus	assets
1989	0.8 (10%)	17.5 (38%)	0.8 (9%)	5.8 (12%)	2.4 (27%)	10.4 (23%)	4.6 (54%)	12.5 (27%)
1992	0.8 (10%)	16.8 (39%)	0.6 (8%)	5.3 (12%)	2.1 (27%)	9.8 (23%)	4.3 (55%)	11.6 (27%)
1995	0.7 (8%)	18.4 (38%)	0.5 (5%)	5.4 (11%)	1.4 (16%)	9.3 (19%)	5.8 (70%)	14.7 (31%)
1998	0.9 (9%)	23.2 (37%)	0.6 (6%)	6.8 (11%)	2.2 (21%)	13.4 (22%)	6.7 (64%)	18.5 (30%)
2001	1.4 (10%)	28.7 (36%)	0.9 (7%)	9.5 (12%)	3.4 (25%)	18.9 (24%)	7.6 (57%)	23.4 (29%)
2004	1.5 (10%)	34.3 (37%)	0.9 (6%)	10.6 (11%)	3.6 (23%)	20.7 (22%)	9.5 (62%)	27.4 (29%)
2007	1.4 (6%)	38.3 (35%)	1.1 (5%)	11.5 (11%)	5.4 (26%)	26.9 (25%)	13.2 (62%)	32.0 (29%)
2010	1.3 (8%)	31.7 (33%)	1.2 (7%)	12.3 (13%)	3.8 (23%)	23.1 (24%)	10.3 (62%)	28.0 (29%)
2013	1.1 (6%)	31.1 (32%)	0.9 (5%)	11.3 (12%)	4.4 (26%)	24.2 (25%)	10.8 (63%)	30.1 (31%)
2016	1.3 (6%)	35.3 (29%)	1.4 (6%)	14.4 (12%)	5.5 (23%)	30.0 (25%)	15.7 (66%)	42.3 (35%)
2019	1.5 (6%)	37.9 (30%)	1.3 (5%)	14.3 (11%)	6.0 (24%)	32.8 (26%)	16.0 (65%)	42.5 (33%)
2022	1.7 (6%)	49.1 (31%)	2.3 (8%)	18.8 (12%)	6.6 (21%)	38.0 (24%)	20.1 (65%)	49.9 (32%)

Note: The table shows the aggregate holdings of private-business equity (variable 'bus' in SCF) and total assets (variable 'asset' in SCF) in trillions of 2022 dollars, for four groups of individuals: the top %1, individuals in percentiles 95-99, individuals in percentiles 90-95, and the bottom 90%. The numbers in parentheses show the holdings in percent of the sum across all individuals, separately for private-equity and total-asset values. The data come from the Survey of Consumer Finance.