Do the Wealthy Earn Higher Returns? Evidence from U.S. Survey Data

Christophe Van Langenhove*

September 2026

Abstract

Estimates of the return gradient across the wealth distribution differ by 36.5 points in U.S. survey data. I recompute it on common PSID and SCF data under 880,317 specifications that combine four return types with thirteen measurement choices. The return gradient ranges from -100.0 to $+42.2$ points. At the baseline, the median return gradient on the gross return on assets is only $+2.8$ points in the PSID and 0.0 in the SCF. The other three return types diverge from it because debt service and leverage fall with wealth.

Keywords: return on wealth, return heterogeneity, wealth inequality, PSID, SCF.

JEL classification: C18, D31, G11

*Vrije Universiteit Brussel (VUB) (email: christophe.van.langenhove@vub.be; ORCID: 0000-0002-9761-2155). I acknowledge the financial support from the Research Foundation Flanders (FWO, project number: 11I5324N). I thank Arthur Apostel and Marthe Mareels for valuable feedback. The Online Supplement and most recent version of this working paper are available at christophevanlangenhove.github.io/.

1 Introduction

A leading explanation for the concentration of wealth at the top is that richer households earn higher returns on their wealth (Benhabib, Bisin and Zhu, 2011; Gabaix et al., 2016; Benhabib, Bisin and Luo, 2019). A higher return compounds faster, so the wealth of the rich diverges from the rest. The strength of that return heterogeneity channel depends on the return gradient, the return of the top wealth decile minus the return of a lower reference wealth decile, in percentage points.

Existing estimates of return heterogeneity from the Panel Study of Income Dynamics (PSID) and the Survey of Consumer Finances (SCF) differ by 36.5 points. Among the PSID studies, Daminato and Pistaferri (2024) report a headline return gradient of +30.0 points, Gaillard and Wangner (2023) +9.6, and Sundram (2025) +4.0. Snudden (2024) instead finds a return that falls with wealth, a return gradient of −6.5 points. Among the SCF studies, Xavier (2021) reports +4.7 points and Hayati (2026) +1.7 after 1980.

The spread in U.S. survey data has two sources. First, the return type: whether the return subtracts the cost of debt (gross versus net return), and whether it divides by gross assets or by wealth (return on assets versus equity). Those two decisions give four return types. Second, thirteen measurement choices: the debts in the cost of debt and their interest rate, vehicle loans in the cost of debt, the rent imputed to owner-occupied housing, the business-income share in the return on assets, vehicles in the asset base, the wave point of the stock, the reference wealth decile, the wealth floor, the outlier trim on each tail, the horizon, the wave point of the ranking, and the mean or the median as each wealth decile's return.

Contributions I make four contributions. First, I recompute the return gradient for each return type on common PSID and SCF data, under 880,317 specifications, one return gradient each. On U.S. survey data the return gradient ranges from −100.0 to +42.2 points, wider than the spread of the existing headlines. Across the specifications the median is +2.5 points and the mean +1.4. Only 2.4 percent of specifications give a return gradient above 10 points.

Second, the return type determines the range of the return gradient. In both the PSID and the SCF, the net return on equity gives the widest range of the four return types, and in the SCF the gross return on assets gives the narrowest. In the SCF the net return on equity runs from −30.1 to +42.2 points, and the gross return on assets from −0.1 to +6.0. In the PSID the net return on equity runs from −99.9 to +26.1 points, and the gross return on assets from −97.4 to +21.0.

Third, within each return type the reference wealth decile moves the return gradient most on the median. In the PSID a lower reference wealth decile steepens the return gradient on the median by up to 5.3 points on the gross return on assets and by up to 13.2 on the gross return on equity. On the mean the largest single effect is the outlier trim's, 11.3 points on the PSID's gross return on equity. The steepest specifications combine a low reference wealth decile, an end-of-wave ranking, a single wave, and no imputed rent.

Fourth, at the baseline the median return gradient on the gross return on assets is only $+2.8$ points in the PSID and 0.0 in the SCF. The other three return types diverge from it in both directions, because the middle wealth deciles hold the household debt. In the PSID the median return gradient is $+6.1$ on the net return on assets, -5.3 on the gross return on equity, and $+3.4$ on the net return on equity. In the SCF the three are $+2.0$, -4.9, and $+0.4$, respectively.

Related literature I contribute to four strands of the literature. First, seven U.S. studies estimate returns across the wealth distribution from the PSID and the SCF (Cao and Luo, 2017; Xavier, 2021; Gaillard and Wangner, 2023; Daminato and Pistaferri, 2024; Snudden, 2024; Hayati, 2026; Sundram, 2025). I place the six numeric headlines and the 880,317 specifications on one specification curve (Leamer, 1983; Sala-i-Martin, 1997; Simonsohn, Simmons and Nelson, 2020).

Second, the U.S. studies that observe the very top find a return premium. Balloch and Richers (2023) observe the security-level portfolios of wealthy investors and find returns rising with wealth, more steeply still after risk adjustment. Smith, Zidar and Zwick (2023) use tax records and find returns rising with wealth, especially on fixed income. Wolff (2025) builds SCF housing returns from imputed rent and accrued capital gains and finds them rising with wealth.

Third, administrative registers in Norway and Sweden show a return premium that neither the PSID nor the SCF can detect (Fagereng et al., 2020; Bach, Calvet and Sodini, 2020; Hubmer et al., 2026). Fagereng et al. (2020) find returns heterogeneous within narrow asset classes. Bach, Calvet and Sodini (2020) find the premium determined primarily by systematic risk, with no evidence that wealthy households select better-performing assets within a class.

Fourth, a class of heterogeneous-agent models builds the long upper tail of wealth from return heterogeneity (Benhabib, Bisin and Zhu, 2011; Gabaix et al., 2016; Aoki and Nirei, 2017; Benhabib and Bisin, 2018; Benhabib, Bisin and Luo, 2019). Hubmer, Krusell and Smith (2021) find return differences across and within wealth groups important for matching U.S. wealth concentration and its dynamics.

Roadmap Section 2 sets out the return gradient, the four return types, and the PSID and the SCF. Section 3 details the thirteen measurement choices, eight common to the PSID and the SCF and five specific to the PSID. Section 4 replicates each existing study at its own configuration, then varies every measurement choice and shows which ones move the return gradient. Section 5 fixes a baseline and reports the return gradient for each return type. Section 6 concludes.

2 Return concepts and data

2.1 Return gradient and return types

The return gradient I summarize return heterogeneity as a return gradient:

$$g = \bar{r}_{10} - \bar{r}_m, \tag{1}$$

where $\bar{r}_{10}$ is the mean or median return of households in the top wealth decile and $\bar{r}_m$ the same for a lower wealth decile m. I report the return gradient in percentage points. The reference wealth decile m is a measurement choice.

Four return types Four return types follow from two decisions: whether to subtract the cost of debt from the return, and whether to divide by gross assets or wealth.

The gross return on assets divides household i's capital income and capital gains by its gross assets:

$$r^g_{A,i} = \frac{Y_i + G_i}{A_i}, \tag{2}$$

where Y_i is capital income, G_i capital gains, and A_i gross assets. Debt enters neither side. Each household quantity refers to one survey wave t, and I drop t wherever a single wave enters.

The net return on assets keeps the gross-asset denominator but charges the cost of debt C_i in the numerator:

$$r^n_{A,i} = \frac{Y_i + G_i - C_i}{A_i} = r^g_{A,i} - \frac{C_i}{A_i}. \tag{3}$$

The subtracted term $s_i = C_i / A_i$ is debt service per dollar of assets.

The gross return on equity divides capital income and capital gains by wealth $E_i = A_i - D_i$:

$$r^g_{E,i} = \frac{Y_i + G_i}{E_i} = (1 + \ell_i)\, r^g_{A,i}, \tag{4}$$

where D_i is debt and $\ell_i = D_i / E_i$ the debt-to-equity ratio. Dividing by wealth rather than gross assets scales the gross return on assets by the leverage factor $1 + \ell_i$, and no cost of debt is charged. No existing study uses this return type.

The net return on equity charges the cost of debt and divides by wealth:

$$r^n_{E,i} = \frac{Y_i + G_i - C_i}{E_i} = (1 + \ell_i)\, r^n_{A,i}. \tag{5}$$

The four coincide only for a household with no debt. Where no single household is meant, I write them r^g_A, r^n_A, r^g_E and r^n_E. Throughout, the asset-based returns are the gross and net returns on assets, and the equity-based returns the gross and net returns on equity. The net-of-debt returns are the net return on assets and the net return on equity, and the gross returns the other two. Snudden (2025) draws the asset-based versus equity-based distinction on U.S. panel data and measures returns on assets, setting aside leverage and debt service.

The wealth-concentration models hold no household debt, so their state variable is gross assets. The return they require is therefore the gross return on assets.

Cao and Luo (2017) and Sundram (2025) use the gross return on assets, and Gaillard and Wangner (2023) and Hayati (2026) the net return on assets. Xavier (2021) uses the net return on equity. Snudden (2024) reports the net return on equity alongside the gross return on assets, and Daminato and Pistaferri (2024) alongside the gross and net returns on assets.

2.2 Two U.S. surveys

The PSID is a biennial household panel. It observes each household's capital income (interest, dividends, rents, business, farm, and trust income), assets, debts, and primary-residence mortgage interest directly. Capital gains and the cost of debt must be imputed. The sample comprises 12 biennial waves, 2001 to 2023, and covers heads aged 20 to 80. The PSID contributes 27,482 household-year returns from 7,174 households.

The SCF is a triennial repeated cross-section.[1] It oversamples wealthy households by design and covers the top of the wealth distribution that the PSID misses (Pfeffer et al., 2016; Van Langenhove, 2026*b*). The sample comprises 12 triennial cross-sections, 1989 to 2022, and covers heads aged 20 to 80. The SCF holds 223,068 observations. The SCF requires the same two imputations.

Returns from both the PSID and the SCF are before tax and in real terms throughout. The inflation adjustment is common to both wealth deciles and nearly cancels in their difference, so the real and nominal return gradients are close.

3 Measurement choices

Neither the PSID nor the SCF observes the household return directly, so I build it from the components of equations (2) to (5). Capital gains and the cost of debt are unobserved and must be imputed. Within each survey, every specification uses one capital-gains construction. Thirteen measurement choices remain: eight common to both the PSID and the SCF, five specific to the PSID.

Table 1 lists the thirteen measurement choices, the values each can take, and the value most existing studies use. Twelve measurement choices have such a common value. For the outlier trim, the existing studies disagree. Table 2 records the configuration of the seven existing studies, every one that measures household returns across the wealth distribution from the PSID or the SCF.

3.1 Capital gains

Accrued capital gains G_i are unobserved in both surveys, so every study imputes them, per asset in the PSID and synthetically in the SCF.

The PSID panel reports the wave-to-wave asset values and the net investment flows, and accrued capital gains sums each class's value change net of its own investment, per year of the two-year wave:

$$G_i = \tfrac{1}{2} \sum_a (\Delta A_{ia} - N_{ia}) \,, \tag{6}$$

[1] I use the standard public SCF. Hayati (2026) instead draws on the SCF+ extension of Kuhn, Schularick and Steins (2020).

Table 1: The thirteen measurement choices and their values

Measurement choice	Values	Common value
Common to both the PSID and the SCF		
1. Cost of debt: which debts	Mortgage debt only, or all debt (mortgage and consumer)	All debt
2. Cost of debt: rate	Flat benchmark, or a wealth-varying credit spread	Flat benchmark†
3. Imputed rent	Omitted, or a 3%, 4%, or 5% rental yield on owner-occupied housing	5%
4. Summary statistic	Mean or median	Mean
5. Reference wealth decile	Third, fourth, fifth, or sixth (21st to 60th wealth percentile)	Fifth
6. Wealth floor	None, \$100, \$500, \$1,000, \$5,000, or \$10,000 minimum wealth	\$1,000
7. Outlier trim per tail	None, 0.5%, or 3%	No clear default
8. Vehicle scope	Exclude vehicles, or include them in the asset base	Include‡
Specific to the PSID		
9. Horizon	Single wave (annual) or long run, both on a one-year basis	Single wave
10. Denominator timing	Midpoint, start, or start plus half flows	Midpoint
11. Ranking timing	Wealth deciles ranked on start-of-wave, midpoint, or end-of-wave wealth	Start‡
12. Vehicle cost of debt	Exclude, or charge the imputed vehicle-loan interest (net-of-debt returns only)	Exclude
13. Business income	Full, the asset part only, or owners only	Full

Notes: Each row is a measurement choice and the values it can take, numbered as in the text. The common measurement choices take the same values in the PSID and the SCF, and the PSID-specific ones apply to the PSID alone. The final column is the value most existing studies use, or the central construction for a PSID-specific measurement choice. The outlier trim has no common value and is set in Section 5.1. †Every existing study that imputes a rate uses a flat one, and the baseline charges a wealth-varying credit spread instead (Section 5.1). ‡The baseline departs from the final column on the ranking timing and the vehicle scope (Section 5.1). The two cost-of-debt measurement choices and the vehicle cost of debt enter only the two net-of-debt return types. Crossing the four return types with these values gives 880,317 specifications, 868,817 in the PSID and 11,500 in the SCF (Appendix A).

Table 2: The measurement configuration of each study

Measurement choice	Cao and Luo (2017)	Gaillard and Wangner (2023)	Daminato and Pistaferri (2024)	Snudden (2024)	Sundram (2025)	Xavier (2021)	Hayati (2026)
Survey, period	PSID 1984–99	PSID 1998–2019	PSID 1999–2019	PSID 1999–2019	PSID 1999–2019	SCF 1989–2019	SCF+ 1949–2022
Return type	Gross r_A^g	Net on assets r_A^n	All but r_E^g	r_A^g and r_E^n	Gross r_A^g	Net on equity r_E^n	Net on assets r_A^n
1. Cost of debt: which debts	None	All debt	All debt	Mortgage only	None	All debt	All debt
2. Cost of debt: rate	None	Reported + flat	Unstated	Reported	None	Reported	Reported + flat
3. Imputed rent	Omitted	5%	Omitted	5%	Actual rent	Rent-to-value ratio	Market yield
4. Summary statistic	n.a.	Mean	Mean	Mean	Mean	Mean	Mean
6. Wealth floor	\$1,000	\$1,000	None stated	\$500	None stated	None stated	None stated
7. Outlier trim per tail	Bound $[-100, 300]\%$	0.5%	None stated	Drop 5/tail, stud. 99%	0.5%	None stated	3%
8. Vehicle scope	Exclude	Include	Include	Include	Include	Include	Include
9. Horizon	5-year, ann.	Single wave	Single wave	Single wave	Single wave	3-year, ann.	Triennial, long-run
10. Denominator timing	Beginning	Midpoint $+\frac{1}{2}$ flow	Beginning $+\frac{1}{2}$ flow	Beginning (hybrid)	Beginning (hybrid)	Beginning	Beginning $+\frac{1}{2}$ flow
11. Ranking timing	Unstated	Unstated	Unstated	Beginning	Beginning	n.a.	n.a.

Notes: Each column is one study and each row a measurement choice, numbered as in Table 1. PSID studies use realized panel returns and SCF studies synthetic returns. r_A^g, r_A^n, r_E^g and r_E^n are the gross return on assets, the net return on assets, the gross return on equity and the net return on equity. "All but r_E^g" marks a study reporting every type except the gross return on equity. The core-asset base of Cao and Luo (2017) nets non-mortgage debts and drops owner-occupied housing (Online Supplement SA.D). "$+\frac{1}{2}$ flow" adds half of net flows to the stated base. SCF+ is the extended SCF of Kuhn, Schularick and Steins (2020). "Stud. 99%" is a studentized 99 percent outlier rule. The beginning-of-period bases of Snudden (2024) and Sundram (2025) are lagged for the asset classes with reported flows and wave-averaged for the rest. In row 2, "Reported" charges the interest rates households report, "flat" an imputed benchmark rate on the debts without one, "Unstated" a debt payment with no stated rate, and "None" no cost of debt. In row 11, the ranking timing is the wave point of the wealth that defines the wealth groups, and "n.a." marks the SCF studies, whose single cross-section has none. Cao and Luo (2017) report no return by wealth group, so row 4 is n.a. for them. Each study's reference wealth decile, measurement choice 5, is stated in its section of the Online Supplement (Online Supplement SA.B). The headline return gradients of Table 3 are not comparable across these configurations.

where A_{ia} is the market value of the household's holding of asset class a, ΔA_{ia} its change over the wave, and N_{ia} the household's net investment in the class. Gaillard and Wangner (2023), Snudden (2024) and Sundram (2025) use this construction, and Daminato and Pistaferri (2024) describe their capital gains as the change in reported holdings.[2] The PSID reports no purchase flows for the owner-occupied home, so I treat a home that first appears between two waves as an acquisition, with no capital gain in that wave (Appendix B).

The SCF observes each household once, so accrued capital gains cannot be computed from the change in its asset values. I build them synthetically instead, assigning each asset class its market return for the wave:

$$G_i = \sum_a A_{ia}\, \gamma_{at}, \tag{7}$$

where γ_{at} is the capital-gain rate of asset class a (listed equity, housing, fixed income, business, pensions) in wave t, from Jordà et al. (2019) and the Federal Reserve Economic Data (FRED). The synthetic construction captures accrued capital gains, but assigns a common γ_{at} within each class. The construction therefore cannot detect a within-class premium, the higher return on the same asset that one household earns over another. Both SCF studies use it (Xavier, 2021; Hayati, 2026).[3]

3.2 Cost of debt

The PSID observes the interest on a household's primary-residence mortgage, and the SCF observes no interest at all. A study that charges a cost of debt C_i therefore imputes the rest. The cost of debt sums interest across a household's debts:

$$C_i = \sum_d r_{id}\, D_{id}, \tag{8}$$

where D_{id} is the household's balance of debt type d (mortgage, consumer) and r_{id} the interest rate it pays on that debt. I impute every interest payment the PSID does not report, and all SCF interest, at a FRED rate, with medical debt and loans from relatives at no interest (Online Supplement SA.E). Imputing the cost of debt involves two measurement choices: which debts to charge, and whether the interest rate varies with wealth.

[2] The residual construction, which recovers the capital gains from the change in wealth less total net investment, attributes any error in measured saving to the capital gains, most of all for the wealthy. No specification uses it.

[3] The SCF's reported realized capital gains omit unrealized capital gains, the bulk of the return at the top, and no study uses the reported ones.

First, which debts to charge interest on (*measurement choice 1*) has two options: all debt, mortgage and consumer, as Daminato and Pistaferri (2024) do, or mortgage debt only, as Snudden (2024) does. Second, whether the imputed rate r_{id} varies with wealth (*measurement choice 2*) also has two options: one flat benchmark rate for every household, or a rate that falls with wealth. Every existing study that imputes a rate uses a flat one. The wealth-varying rate adds a credit spread to the benchmark, positive in the lowest wealth deciles and negative in the highest, so the poorest households pay the highest rate on their imputed debt.[4]

3.3 Other common measurement choices

Five other measurement choices common to both the PSID and the SCF are the imputed rent, the summary statistic, the reference wealth decile, the wealth floor, and the outlier trim.

First, the imputed rent is a rental yield ρ times the value of an owner-occupied home (*measurement choice 3*). It enters capital income in the realized PSID return and the synthetic housing return in the SCF. The options are none, 3, 4, and 5 percent. The literature sets ρ to zero (Cao and Luo, 2017; Daminato and Pistaferri, 2024) or 5 percent (Gaillard and Wangner, 2023; Snudden, 2024). Three studies use no fixed yield: Sundram (2025) counts only the rent a household actually receives, Hayati (2026) applies each year's market rental yield, and Xavier (2021) the rent-to-value ratio of the SCF's other real estate.

Second, each wealth decile's return is the mean or the median of its households' returns (*measurement choice 4*).[5] Every study with a return gradient uses a mean, though the two SCF studies use dollar-weighted bin aggregates rather than household means (Xavier, 2021; Gaillard and Wangner, 2023; Daminato and Pistaferri, 2024; Snudden, 2024; Hayati, 2026; Sundram, 2025).

Third, the reference wealth decile is the lower wealth decile m that the top is measured against in the return gradient (1) (*measurement choice 5*). The literature places it at the 20th to 40th wealth percentile (Xavier, 2021), at the 40th to 50th (Gaillard and Wangner, 2023), at the bottom quintile (Sundram, 2025), against the bottom 90 percent (Hayati, 2026), or at median wealth (Daminato and Pistaferri, 2024). I vary it across the third to sixth wealth decile.

Fourth, the wealth floor drops households below a minimum wealth: no floor, \$100, \$500, \$1,000, \$5,000, or \$10,000 (*measurement choice 6*). Even with no floor, every PSID specification

[4] The spreads come from the Fannie Mae loan-level price-adjustment matrices for mortgage debt and from credit-score-based rate schedules for consumer debt, with each wealth decile assigned its own rate. Both are parameters of the construction rather than estimates. The consumer spread ranges from $+6.73$ points at the lowest wealth decile to -5.08 at the highest, a gap of 11.81 points, and the mortgage spread from $+1.75$ to -0.25, a gap of 2.00.

[5] Van Langenhove (2026*a*) formalizes the same household-level median for saving rates across the wealth distribution.

keeps only households with positive wealth and every SCF specification only households with non-negative wealth. These options cover every floor the existing studies impose (Table 2).

Fifth, the outlier trim drops the most extreme returns in each tail: none, 0.5 percent, or 3 percent (*measurement choice 7*). Gaillard and Wangner (2023) and Sundram (2025) trim 0.5 percent per tail and Hayati (2026) trims 3 percent. Snudden (2024) instead drops the five most extreme observations in each tail, and Cao and Luo (2017) bound returns to $[-100, 300]$ percent. Daminato and Pistaferri (2024) and Xavier (2021) do not report any trimming. No outlier trim is common to most studies, so the baseline sets it by reasoning rather than from the literature (Section 5.1).

3.4 Vehicle measurement choices

Two measurement choices concern vehicles, which both the PSID and the SCF record and most existing studies count as an asset (Table 2).

First, the vehicle scope excludes vehicles from the asset base or includes them (*measurement choice 8*). The PSID records vehicle wealth net of any loan, and the SCF records the value and the loan separately (Online Supplement SA.F). Every existing study but Cao and Luo (2017) includes vehicles. Including them raises the denominator, wealth and gross assets alike, and adds no capital income, since a vehicle earns none.

Second, the vehicle cost of debt omits the interest on vehicle loans or charges it (*measurement choice 12*). The PSID reports only the annual loan payment, from which I impute the interest at an external auto-loan rate (Online Supplement SA.F). The measurement choice is PSID-specific because the SCF folds vehicle loans into its installment debt, so their interest is charged in every SCF specification that charges consumer debt, whether or not vehicles enter the asset base.

3.5 PSID-specific measurement choices

Four PSID-specific measurement choices remain: the return horizon, the denominator timing, and the ranking timing, which arise from the panel structure, and the attribution of business income. The fifth, the vehicle cost of debt, stays with the vehicle scope in Section 3.4.

First, the PSID's wave return is an annual return. The interview is biennial, the income covers the calendar year before it, and the capital gains are averaged per year over the wave (Appendix B). The horizon takes that return over a single wave or compounds a household's

wave returns into one long-run average over all its waves (*measurement choice 9*).[6] The long-run average removes the transitory variation in a single wave. Both horizons are on a one-year basis, like the SCF return. Gaillard and Wangner (2023) and Daminato and Pistaferri (2024) use a single wave, Cao and Luo (2017) a five-year average.

Second, a wave's return divides income and capital gains by a stock. The denominator timing is which point of the wave that stock is measured at (*measurement choice 10*). The options are its value at the wave's start, the midpoint of its start and end, or the start plus half the wave's net flows.[7] The half-flow denominator is the Dietz (1968) convention of Fagereng et al. (2020). Gaillard and Wangner (2023) apply it to the PSID on a gross-asset base and Daminato and Pistaferri (2024) on a net-wealth base. Cao and Luo (2017) divide by the start, and Snudden (2024) and Sundram (2025) by the start for the asset classes with reported flows and by the wave average for the rest. All three options require a household's stock tracked across waves, which only the PSID panel provides.[8]

Third, the ranking timing is the point of the wave at which the wealth that sorts households into wealth deciles is measured (*measurement choice 11*): the wave's start, its midpoint, or its end. Ranking at the end counts the wave's capital gains in the wealth that ranks, which steepens the return gradient. Snudden (2024) and Sundram (2025) rank at the start, Snudden (2024) for that reason. Gaillard and Wangner (2023) and Daminato and Pistaferri (2024) do not state their timing (Online Supplement SA.B.4).

Fourth, the attribution of business income is the part of unincorporated business income that counts as a return (*measurement choice 13*). The PSID reports that income in two parts, the owner's labor income and the asset part, a return on the business equity. The SCF reports it as one amount, so the measurement choice is PSID-specific. The options are all of it, the asset part alone, or all of it for business owners only. Business owners only is the closest reading of the asset part Sundram (2025) uses. The baseline counts all of it, as the literature does. The options differ most in the lowest wealth deciles. There, counting all of it turns a self-employed household's labor income into a large return on little business equity (Online

[6] The long-run return is the geometric mean across the household's T_i waves, $\left(\prod_t (1+r_{it})\right)^{1/T_i} - 1$, which stays on a one-year basis.

[7] On a gross-asset base the three denominators are $A_{i,t-1}$, the geometric midpoint $\sqrt{A_{i,t-1}\,A_{it}}$, and $A_{i,t-1} + \frac{1}{2}N_{it}$. Here $A_{i,t-1}$ and A_{it} are gross assets at the wave's start and end, and N_{it} is the household's net investment over the wave summed across asset classes. On a wealth base E replaces A. The midpoint is arithmetic where the two stocks differ in sign or either is zero.

[8] The standard SCF observes a household's wealth only once, so this measurement choice and the ranking timing below do not arise there. The exception is Hayati (2026), whose SCF+ pseudo-panel permits a beginning-plus-half-flow denominator, and I recompute that study on the standard SCF.

Supplement SA.B.7). Farm income is not split in the PSID and counts as a return under every option.

4 Return gradient estimates and their drivers

Crossing the thirteen measurement choices with the four return types, in the PSID and in the SCF, gives 880,317 specifications, one return gradient each (Table 1, Appendix A). Together they form the specification curve of Simonsohn, Simmons and Nelson (2020).

4.1 Replicating the existing studies

I replicate the seven existing studies twice, once at each study's own measurement configuration and once on a common comparison.

First, I recompute each study on my PSID or SCF data at its own configuration, comparing its own top and reference groups (Table 3). The headline estimates differ by 36.5 points and the replications by 13.3. 3 of the 6 numeric headlines fall inside their replication's 95 percent interval. The steepest, Daminato and Pistaferri (2024), does not. Their paper leaves the timing of the wealth ranking unstated, and their +30.0-point headline reproduces only when households are ranked on end-of-wave wealth.[9] The Gaillard and Wangner (2023) headline falls just above its interval. Their top group, the PSID's top wealth percentile, holds 25 household-years per wave (Online Supplement SA.B.3).

Second, I recompute each study's return type on one common comparison: the mean return of the top wealth decile against the fifth, with a $10,000 wealth floor and a 0.5 percent outlier trim. Those estimates differ by 11.5 points, and no headline falls inside its common-comparison interval (Online Supplement SA.C). The return type and the horizon set the sign and the size of the baseline return gradient (Online Supplement SA.G).

4.2 The distribution of return gradients

Three findings follow from the specification curve.

First, across the 880,317 specifications the return gradient ranges from −100.0 to +42.2 points and contains every existing headline. The median is +2.5 points and the mean +1.4. Only 2.4 percent of specifications exceed 10 points, and 28 percent are negative. Both steep estimates,

[9] Ranked on end-of-wave wealth and lightly trimmed, the replication gives +29.6. Ranked on end-of-wave wealth and left untrimmed, as their rule leaves it, it gives −132.3. Table 3 ranks at the start of the wave, as every study that states a timing does. The end-of-wave reading also needs an outlier trim, and their paper applies none.

Table 3: Each study replicated on my data at its own measurement configuration

Study	Return type	Groups compared	Reported	Replicated	
			Return gradient	Return gradient	Interval
Cao and Luo (2017)	r_A^g	–	–	+0.9	(−1.5, +3.2)
Xavier (2021)	r_E^n	Top 1% vs P20–40	+4.7	+4.9	(+4.0, +5.7)
Gaillard and Wangner (2023)	r_A^n	Top 1% vs P40–50	+9.6	+4.7	(+1.2, +8.3)
Daminato and Pistaferri (2024)	r_E^n	Top 1% vs median	+30.0	+6.6	(+2.4, +10.9)
Snudden (2024)	r_E^n	Top 10% vs P50–60	−6.5	−6.7	(−8.9, −4.6)
Hayati (2026)	r_A^n	Top 10% vs bottom 90%	+1.7†	+4.2	(+4.0, +4.3)
Sundram (2025)	r_A^g	Top 20% vs bottom 20%	+4.0	+3.5	(+2.7, +4.2)

Notes: Each row is a study's reported return gradient and its recomputation on my PSID or SCF data at the study's own return type, age window, sample period, weighting, and top and reference groups, in percentage points. r_A^g, r_A^n and r_E^n are the gross return on assets, the net return on assets and the net return on equity. P20–40 is the 20th to 40th wealth percentile. The interval is a 95 percent bootstrap interval over whole households, within survey year in the SCF. Where a study's own imputed rent or interest rate cannot be built on the public files, the row takes this paper's 5 percent yield and wealth-varying rate. Cao and Luo (2017) report no return gradient, so their row compares the top wealth decile against the fifth at their configuration. The Snudden (2024) headline is read from his Figure 4 and the Sundram (2025) headline is the difference of his two stated quintile returns. The Sundram (2025) row counts vehicles and is nominal, as his paper is. The other rows exclude vehicles and are deflated (Appendix B). †Hayati (2026) reports +1.7 on the SCF+ extension after 1980, which the standard SCF from 1989 tests.

Gaillard and Wangner (2023) and Daminato and Pistaferri (2024), lie in the highest few percent of all specifications (Figure 1).

Second, the return types on gross assets vary less across specifications than those on wealth (Table 4). In the SCF the gross return on assets runs from −0.1 to +6.0 points and the net return on equity from −30.1 to +42.2. In the PSID the gross return on assets runs from −97.4 to +21.0 points and the net return on equity from −99.9 to +26.1.

Third, the PSID and the SCF produce nearly the same central distribution, with means of +1.3 and +2.2 points. Steep specifications are rare in both: 21,321 of the 868,817 PSID specifications and 55 of the 11,500 SCF specifications exceed 10 points (Table 4). The lowest and highest return gradients are −100.0 and +26.1 points in the PSID and −64.2 and +42.2 in the SCF. No PSID specification reaches the steepest PSID headline of +30.0 points.

4.3 The drivers of the return gradient

The reference wealth decile and the imputed rent move the return gradient most on the median. On the mean the outlier trim gives the largest single effect.

Figure 1: The specification curve of the return-wealth gradient

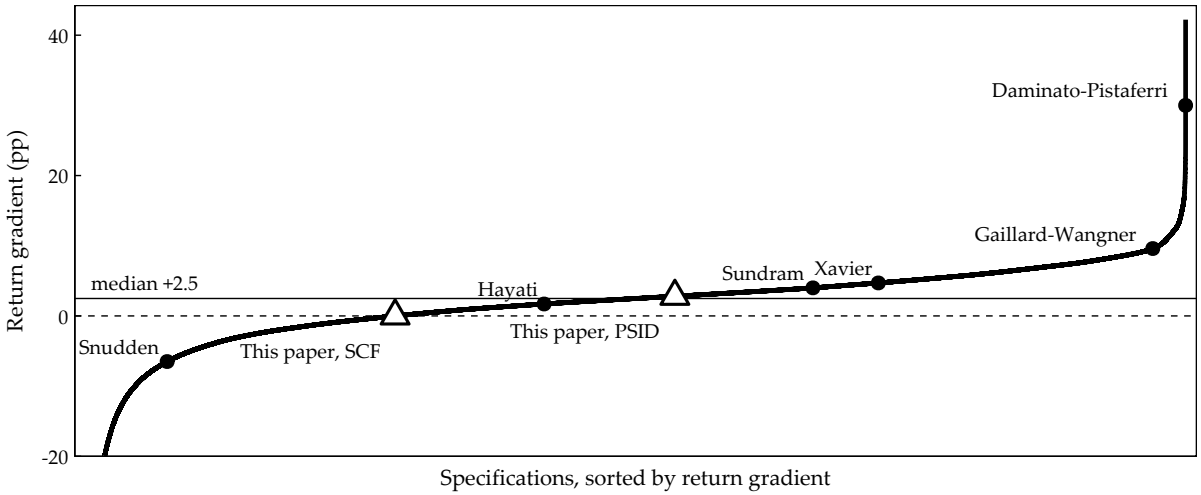

Notes: All 880,317 specifications, sorted from the most negative return gradient to the most positive. Filled points are the six existing studies with a numeric headline, at the rank of their value. Open triangles are this paper's baseline on the gross return on assets, on the median, in each survey. The dashed line marks a zero return gradient and the solid line the median specification. The lower tail continues to −100.0 points, below the plotted range.

Decomposition strategy Within a return type and survey, I measure each measurement choice by how far its options move the return gradient. Let $\tilde{g}_c(o)$ be the median of the return gradient across the specifications that take option o of measurement choice c. The two vehicle measurement choices and the business-income attribution are held at their baseline values, and nothing else is fixed. The effect of measurement choice c is the segment between the smallest and largest of these medians:

$$\left[\min_o \tilde{g}_c(o),\ \max_o \tilde{g}_c(o) \right]. \tag{9}$$

The two vehicle measurement choices and the business-income attribution are measured on the specifications where each of them varies (Online Supplement SA.H.2).

I report the effects separately for the median-based and the mean-based specifications, since a mean wealth-decile return includes the extreme household returns at low wealth and a median does not (Figures 2 and 3, on one shared axis). A variance decomposition and a Shapley decomposition rank the measurement choices the same way (Online Supplement SA.H.2).

Results I obtain three results. First, on the median the reference wealth decile and the imputed rent move the return gradient most (Figure 2). A lower reference wealth decile steepens the return gradient on the net return on equity by up to 7.3 points in the PSID and 3.9 in the SCF. Imputed rent flattens the return gradient on the net return on equity by up to 6.3 points

Figure 2: How each measurement choice moves the median-based return gradient

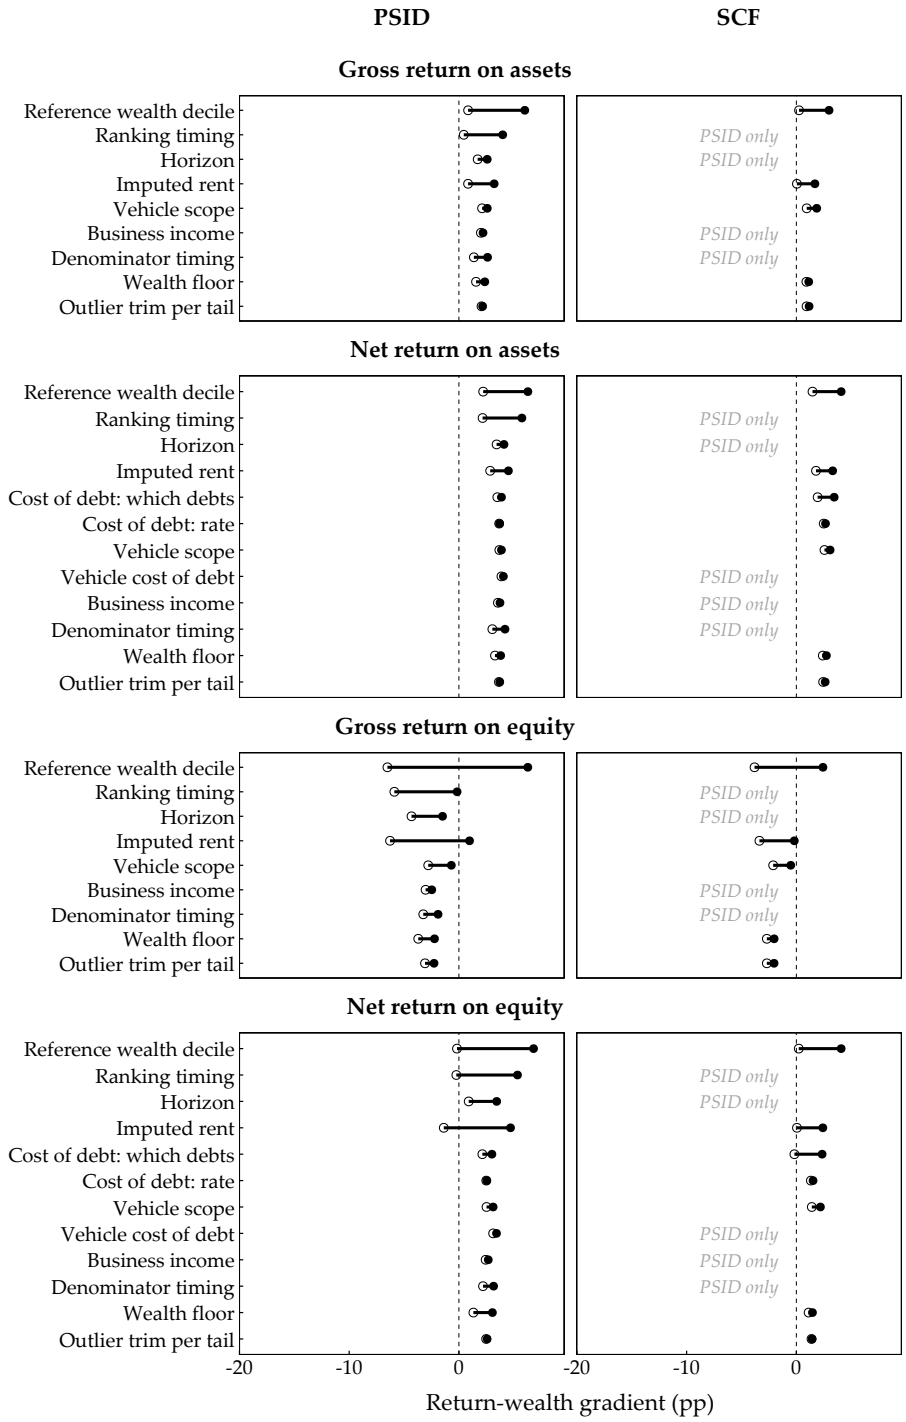

Notes: Specifications that take each wealth decile's return as the median of its households' returns. Columns are the PSID and the SCF, rows the four return types, on the return gradient axis shared with Figure 3. Each segment runs from the smallest of a measurement choice's option medians (open circle) to the largest (filled circle), as in equation (9). The dashed line marks a zero return gradient. The segments hold the two vehicle measurement choices at baseline, as the effects reported in the text do. They let the business-income attribution vary, and the effects in the text hold it at baseline, so a PSID segment and its effect can differ a little. The horizon, the denominator timing, the ranking timing, the vehicle cost of debt and the business-income attribution apply to the PSID only, and the cost of debt does not enter the two gross returns.

Figure 3: How each measurement choice moves the mean-based return gradient

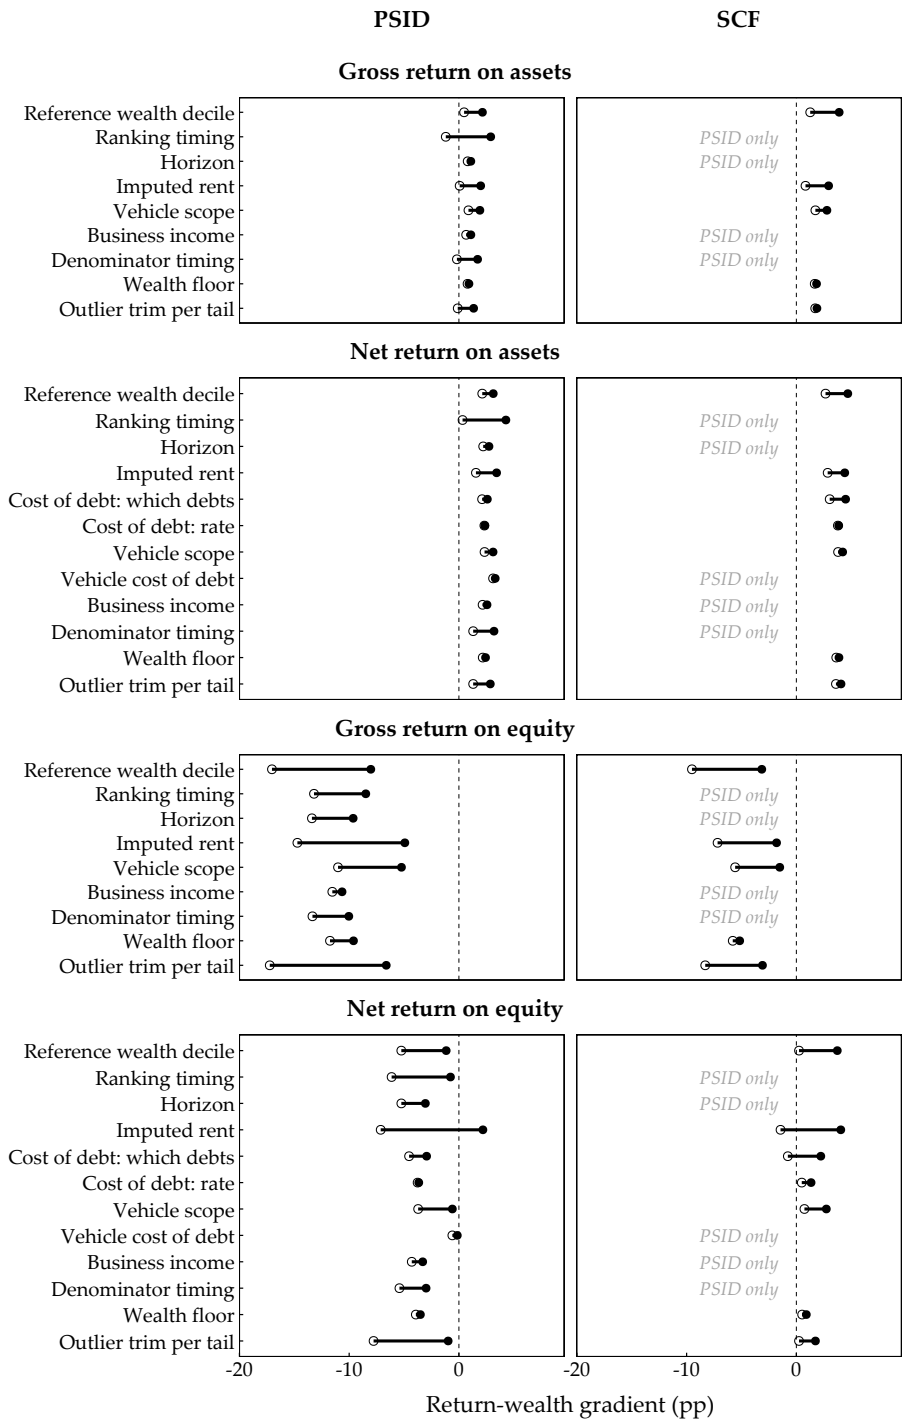

Notes: Specifications that take each wealth decile's return as the mean of its households' returns. Columns, rows, segments, axis and dashed line are as in Figure 2.

Table 4: Distribution of the return gradient across specifications, by return type and survey

Return type	Mean	Median	SD	Range
PSID				
Gross return on assets (r_A^g)	+1.9	+1.8	4.6	[−97.4, +21.0]
Net return on assets (r_A^n)	+3.4	+3.4	4.2	[−94.9, +22.8]
Gross return on equity (r_E^g)	−5.8	−4.3	12.0	[−100.0, +18.4]
Net return on equity (r_E^n)	+0.3	+1.2	8.6	[−99.9, +26.1]
SCF				
Gross return on assets (r_A^g)	+2.1	+1.9	1.4	[−0.1, +6.0]
Net return on assets (r_A^n)	+3.6	+3.5	1.5	[+0.5, +8.0]
Gross return on equity (r_E^g)	−2.7	−1.9	6.4	[−64.2, +6.4]
Net return on equity (r_E^n)	+2.1	+1.9	3.4	[−30.1, +42.2]

Notes: Each row summarizes the return gradient across every specification of one return type in one survey, including the specifications without an outlier trim. The mean, median, standard deviation, minimum and maximum are in percentage points. The data are fixed within each row, so the spread is measurement.

in the PSID and 2.4 in the SCF. Imputed rent raises the return of the housing-heavy middle wealth deciles toward the top wealth decile's.

Second, the horizon and the ranking timing steepen the PSID return gradient through transitory capital gains. On the median, a single wave rather than a long-run average raises the return gradient by up to 2.9 points on the gross return on equity, and ranking at the wave's end rather than its start by up to 5.7 points on the net return on equity. A household with a large capital gain in a wave has a high return and a high wealth rank at once, so it enters the top wealth decile with the gain still in its return.

Third, the outlier trim and the wealth floor matter on the mean and not on the median. On the median the outlier trim moves the return gradient by at most 0.9 points and the wealth floor by at most 1.7. On the mean the outlier trim gives the largest single effect, 11.3 points on the PSID's gross return on equity, because a mean includes the extreme returns of households with near-zero wealth (Figure 3). Daminato and Pistaferri (2024) state neither a wealth floor nor an outlier trim. The reference wealth decile, the ranking timing, the horizon and the imputed rent separate the steepest specifications from the flattest (Online Supplement SA.I).

5 A baseline return gradient

The baseline fixes each measurement choice at a defensible value and reports the resulting return gradient for each of the four return types.

5.1 Baseline measurement choices

The baseline takes each measurement choice at its common value in Table 1, with four exceptions. First, the outlier trim has no clear default and is set to 0.5 percent per tail. Across its three options the outlier trim moves the return gradient on the median by at most 0.9 points. Second, the interest rate on debt falls with wealth, where every existing study imputes a flat rate. A wealth-varying rate steepens the two net-of-debt return gradients. Third, households are ranked on end-of-wave wealth, where the studies that state a timing rank at the start. End-of-wave ranking steepens the return gradient (Section 3.5). Fourth, vehicles are excluded from the asset base, where most existing studies count them. A vehicle earns no capital income, so counting it only enlarges the denominator.

The baseline does not fix the summary statistic and reports each wealth decile's return as both a mean and a median.

5.2 Baseline results

The baseline gives two results (Figures 4 and 5). First, the return gradient on the gross return on assets is small. On the median it is $+2.8$ in the PSID and 0.0 in the SCF, and on the mean $+3.1$ and $+0.8$ (Figure 4). Second, the other three return types diverge from it in both directions. The net return on assets, which charges the cost of debt, has a median return gradient of $+6.1$ in the PSID and $+2.0$ in the SCF. The gross return on equity, which divides by wealth, has a negative return gradient, -5.3 in the PSID and -4.9 in the SCF (Figure 5). The net return on equity, which applies both, lies between the two at $+3.4$ and $+0.4$. On the mean the two equity-based return gradients in the PSID lie below their medians, at -19.9 and -5.4, and the means on the gross and net returns on assets stay close to their medians, at $+3.1$ and $+5.5$.

A return premium appears at the very top of the SCF. Measured to the top wealth percentile instead of the top wealth decile, the SCF's mean return gradient on the gross return on assets is $+1.9$, different from zero at the 5 percent level. The PSID's top wealth percentile holds only 25 household-years per wave, against 649 households in the SCF, too few to measure a premium (Online Supplement SA.J). Neither survey observes the billionaires above that (Vermeulen, 2018), where Balloch and Richers (2023) find returns rising with wealth through private equity and other alternatives.

The small median return gradient on the gross return on assets stays below the steep existing headlines across the defensible measurement choices and nine further robustness checks (Ap-

Figure 4: Return by wealth decile: the asset-based returns

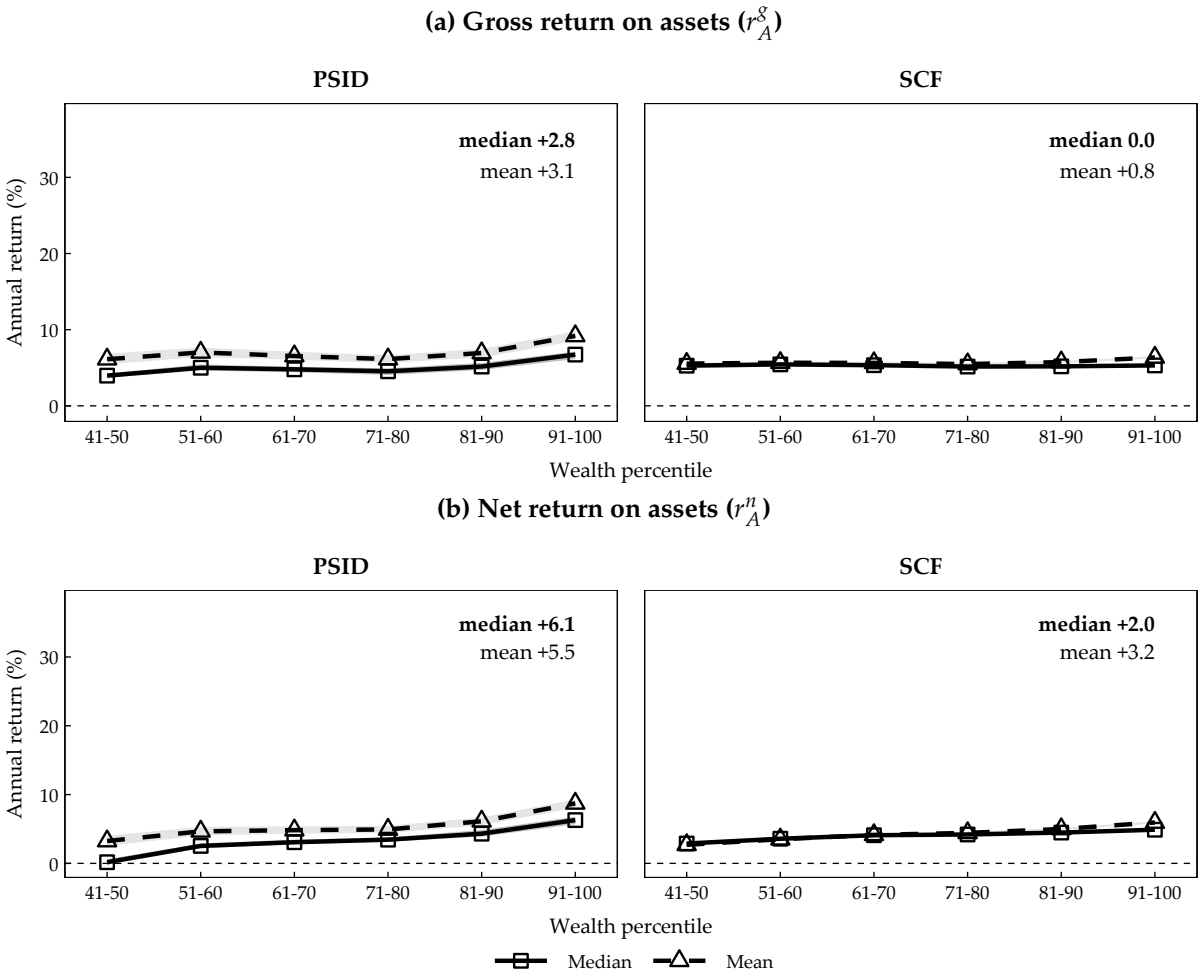

Notes: Each subfigure plots the annual return by wealth decile in the PSID and the SCF for the return type in its subcaption. The solid line with open squares is the median and the dashed line with open triangles the mean, each with a 95 percent bootstrap band over whole households, a drawn SCF household carrying its five imputed copies. PSID returns are realized wave returns and SCF returns synthetic returns, on heads aged 20 to 80, with a 5 percent imputed rent, a $1,000 wealth floor and a 0.5 percent outlier trim, ranked on end-of-wave wealth and with vehicles excluded. Each panel starts at the reference wealth decile, the fifth. The return gradient at the top right is the top wealth decile's return less the fifth's, for the median and the mean. All four return-type subfigures, here and in Figure 5, share one vertical scale. The levels are not comparable across surveys, since realized PSID and synthetic SCF returns are different objects over different periods.

pendix A.2), in each of four multi-wave PSID subperiods (Appendix A.3), net of a quadratic in age and year fixed effects (Appendix A.4), and after risk adjustment (Online Supplement SA.K). The baseline medians also reject the steep existing headlines. At the 5 percent level they differ from +9.6 and +30.0, which are net returns, and from the +4.0 of Sundram (2025), a gross return on assets like the baseline (Online Supplement SA.B.7). A household's average return over its odd waves and its average return over its even waves correlate at −0.04. The wealth-decile means from the same two halves correlate at 0.94, and the return gradient is +2.7 in one half and +3.5 in the other.

Figure 5: Return by wealth decile: the equity-based returns

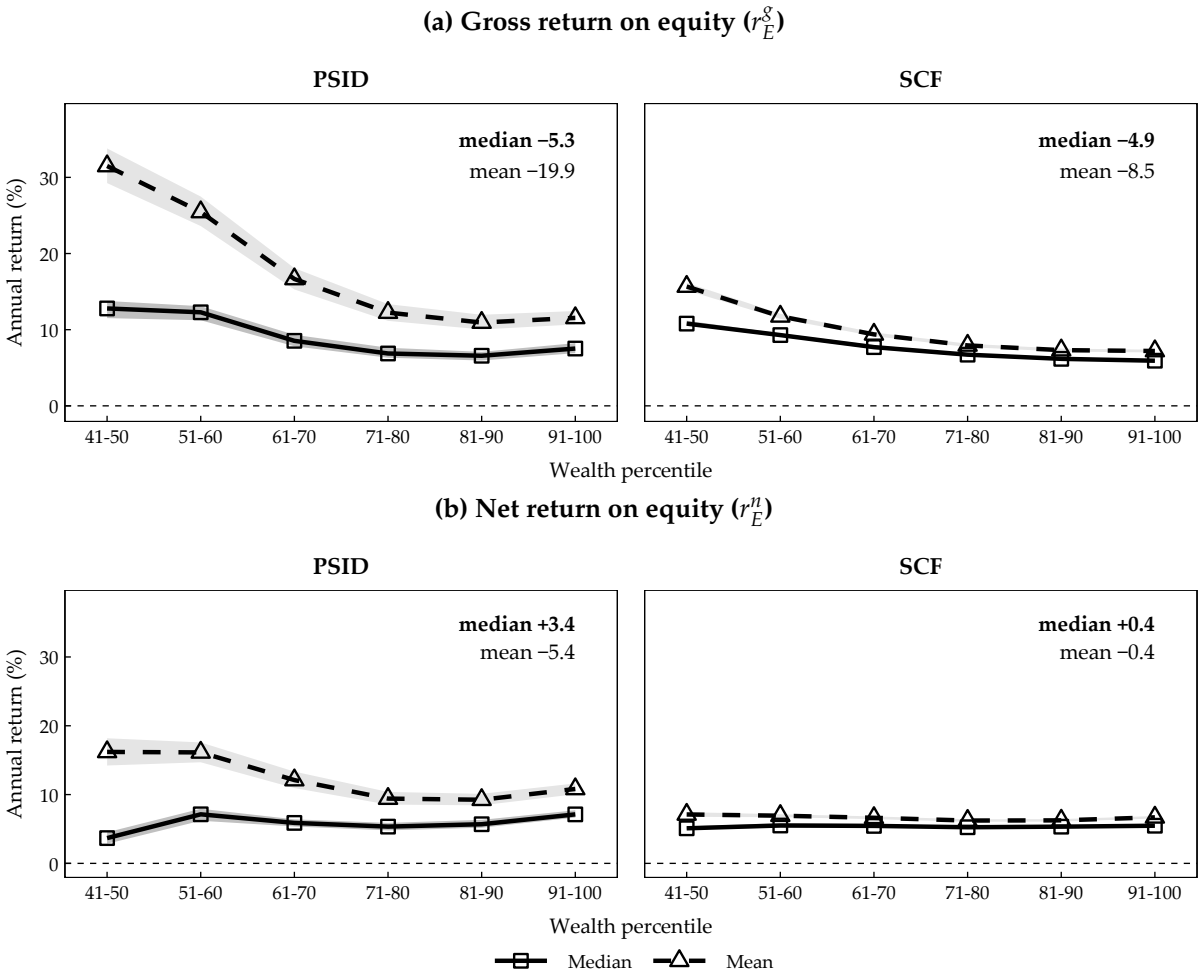

Notes: Each panel plots the annual return by wealth decile for one survey and the return type in its subcaption, the two returns that divide by wealth. Lines, bands, sample, vertical scale and the return gradient at the top right are as in Figure 4.

5.3 Divergence across the four return types

Two wedges separate the gross return on assets from the other three return types: debt service and leverage. Both come from household debt. Debt service per dollar of assets is $s = C/A$ and leverage is $\ell = D/E$ (Section 2.1). Equations (3) to (5) give the other three return types from the gross return on assets:

$$r_A^n = r_A^g - s, \qquad r_E^g = (1+\ell)\, r_A^g, \qquad r_E^n = (1+\ell)\, r_A^n.$$

Subtracting debt service turns a gross return into a net return. Multiplying by $1+\ell$ turns an asset return into an equity return.

Figure 6: Leverage and the cost of debt across the wealth distribution

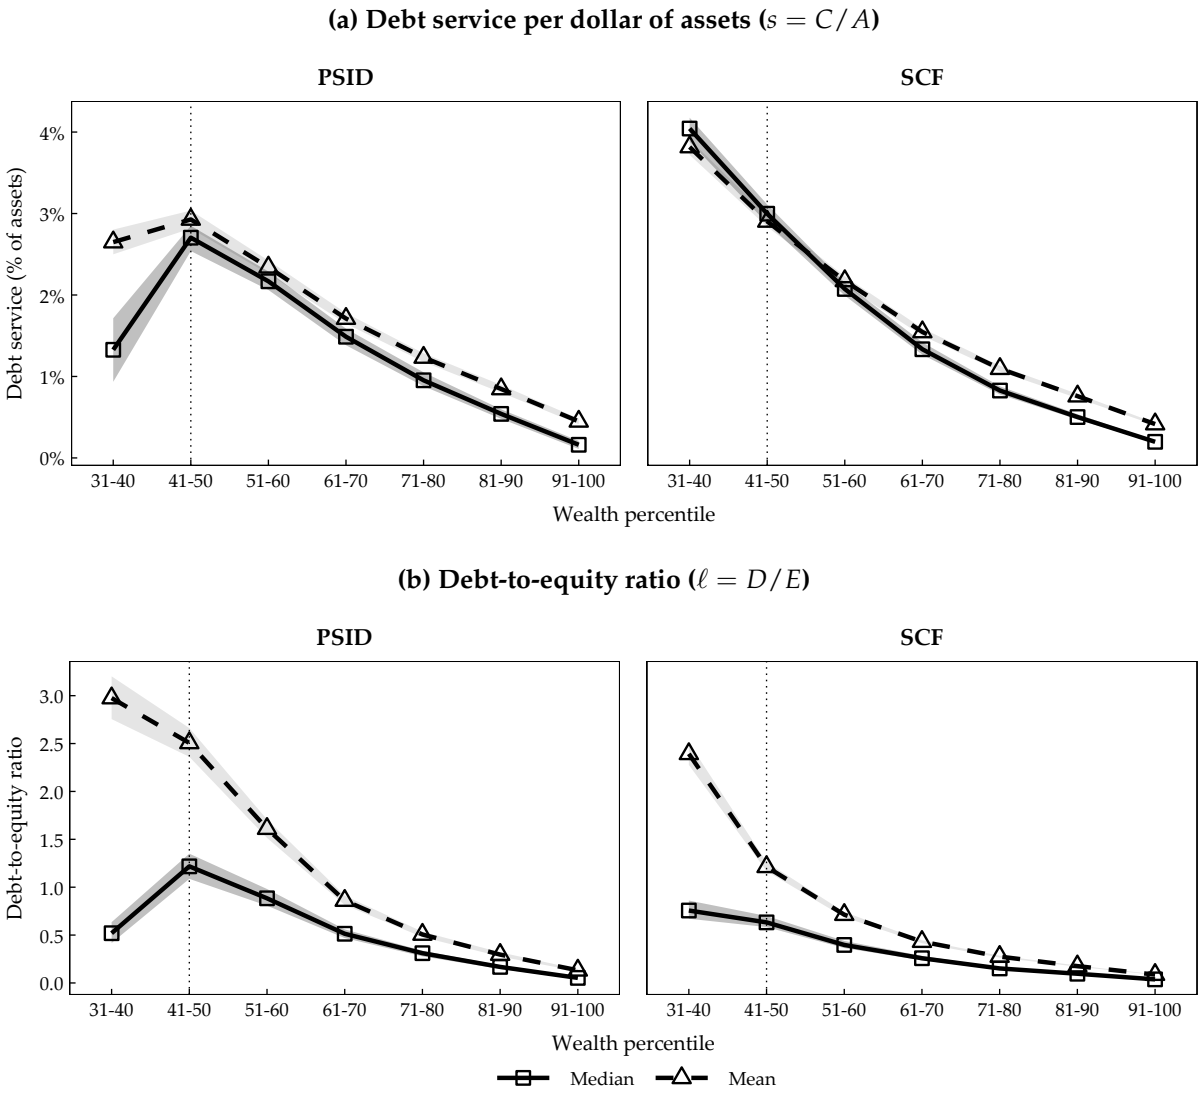

Notes: Each panel plots one wedge by wealth decile for one survey, the median (solid) and the mean (dashed), each with a 95 percent bootstrap band over whole households. Panel (a) is debt service per dollar of gross assets, $s = C/A$. Panel (b) is the debt-to-equity ratio, $\ell = D/E$. The PSID charges each household's reported primary-mortgage interest and the wealth-varying rate on its other debt, and the SCF the wealth-varying rate on every debt, as in the net panels of Figures 4 and 5. Sample, wealth floor and outlier trim are as in Figure 4. The dotted line marks the reference wealth decile, the fifth.

Both wedges are largest in the middle wealth deciles and vanish at the top (Figure 6). The three other return gradients follow from this pattern. In the PSID, the median household at the reference wealth decile pays 2.7% of its assets in debt service and the median household at the top wealth decile 0.2%. Subtracting debt service therefore lowers the middle's return more than the top's, and the net return on assets has a higher return gradient than the gross return on assets. In addition, the reference wealth decile's debt-to-equity ratio is 1.22 and the top's 0.05. Dividing by wealth therefore doubles the middle's return and leaves the top's unchanged, and the gross return on equity has a negative return gradient. The net return on equity applies both wedges, and its return gradient lies between the two. The SCF profile is the same.

6 Conclusion

The six existing headline estimates of the return gradient in U.S. survey data differ by 36.5 points. Replicated on the same PSID and SCF data, each at its own measurement configuration, they differ by 13.3 points. The return type determines the range of the return gradient. Within a return type, the reference wealth decile and the imputed rent move the return gradient most on the median. On the mean the largest single effect is the outlier trim's.

At the baseline the median return gradient on the gross return on assets is $+2.8$ in the PSID and 0.0 in the SCF. The other three return types diverge from it in both directions. The cost of debt raises the median return gradient on the net return on assets to $+6.1$ in the PSID. Leverage makes the gross return on equity negative, at -5.3 in the PSID and -4.9 in the SCF. The net return on equity applies both wedges and lies between the two at $+3.4$ in the PSID.

For the heterogeneous-agent models that build the long upper tail of wealth from a return rising with wealth, the relevant return is the gross return on assets, and its return gradient is small.

Neither survey can find a within-class premium at the very top. The SCF's mean return gradient of $+1.9$ on the gross return on assets at the top wealth percentile is portfolio composition by construction, since its synthetic return gives every household in an asset class the same return, and the PSID holds too few households there to measure one. The U.S. studies with account-level or tax data find returns rising with wealth (Balloch and Richers, 2023; Smith, Zidar and Zwick, 2023).

References

Aoki, Shuhei, and Makoto Nirei. 2017. "Zipf's Law, Pareto's Law, and the Evolution of Top Incomes in the United States." *American Economic Journal: Macroeconomics*, 9(3): 36–71.

Bach, Laurent, Laurent E. Calvet, and Paolo Sodini. 2020. "Rich Pickings? Risk, Return, and Skill in Household Wealth." *American Economic Review*, 110(9): 2703–2747.

Balloch, Cynthia, and Julian Richers. 2023. "Asset Allocation and Returns in the Portfolios of the Wealthy." Financial Markets Group Discussion Paper 885, London School of Economics.

Benhabib, Jess, Alberto Bisin, and Mi Luo. 2019. "Wealth Distribution and Social Mobility in the US: A Quantitative Approach." *American Economic Review*, 109(5): 1623–1647.

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.

Cao, Dan, and Wenlan Luo. 2017. "Persistent Heterogeneous Returns and Top End Wealth Inequality." *Review of Economic Dynamics*, 26: 301–326.

Daminato, Claudio, and Luigi Pistaferri. 2024. "Returns Heterogeneity and Consumption Inequality over the Life Cycle." NBER Working Paper 32490.

Dietz, Peter O. 1968. "Components of a Measurement Model: Rate of Return, Risk, and Timing." *Journal of Finance*, 23(2): 267–275.

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

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

Gaillard, Alexandre, and Philipp Wangner. 2023. "Wealth, Returns, and Taxation: A Tale of Two Dependencies." SSRN Working Paper 3966130.

Hayati, Mojtaba. 2026. "Scale-Dependent Returns and the Interest Rate." Working Paper, University of Zurich.

Hubmer, Joachim, Elin Halvorsen, Sergio Salgado, and Serdar Ozkan. 2026. "Why Are the

Wealthiest So Wealthy? New Longitudinal Empirical Evidence and Implications for Theories of Wealth Inequality." *Econometrica*. Forthcoming.

Hubmer, Joachim, Per Krusell, and Anthony A. Smith. 2021. "Sources of US Wealth Inequality: Past, Present, and Future." *NBER Macroeconomics Annual*, 35: 391–455.

Jordà, Òscar, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick, and Alan M. Taylor. 2019. "The Rate of Return on Everything, 1870–2015." *Quarterly Journal of Economics*, 134(3): 1225–1298.

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

Leamer, Edward E. 1983. "Let's Take the Con Out of Econometrics." *American Economic Review*, 73(1): 31–43.

Pfeffer, Fabian T., Robert F. Schoeni, Arthur Kennickell, and Patricia Andreski. 2016. "Measuring Wealth and Wealth Inequality: Comparing Two U.S. Surveys." *Journal of Economic and Social Measurement*, 41(2): 103–120.

Sala-i-Martin, Xavier X. 1997. "I Just Ran Two Million Regressions." *American Economic Review*, 87(2): 178–183.

Simonsohn, Uri, Joseph P. Simmons, and Leif D. Nelson. 2020. "Specification Curve Analysis." *Nature Human Behaviour*, 4(11): 1208–1214.

Smith, Matthew, Owen Zidar, and Eric Zwick. 2023. "Top Wealth in America: New Estimates under Heterogeneous Returns." *Quarterly Journal of Economics*, 138(1): 515–573.

Snudden, Stephen. 2024. "Leverage and Rate of Return Heterogeneity among U.S. Households." LCERPA Working Paper 2024-8.

Snudden, Stephen. 2025. "Idiosyncratic asset return and wage risk of US households." *Economic Inquiry*, 63(2): 636–657.

Sundram, Jacob. 2025. "Monetary Policy, Heterogeneous Returns, and Top Inequality." Job Market Paper, University of Copenhagen.

Van Langenhove, Christophe. 2026*a*. "Saving Rate Heterogeneity across the Wealth Distribution in the United States."

Van Langenhove, Christophe. 2026*b*. "Wealth Mobility in the United States: Empirical Evi-

dence from the PSID." UGent, Department of Economics Working Paper WP25/1104.

Vermeulen, Philip. 2018. "How Fat Is the Top Tail of the Wealth Distribution?" *Review of Income and Wealth*, 64(2): 357–387.

Wolff, Edward N. 2025. "Heterogenous Rates of Return on Homes: Do the Rich Do Better?" *Review of Income and Wealth*, 71(1): e12704.

Xavier, Inês. 2021. "Wealth Inequality in the US: The Role of Heterogeneous Returns." Federal Reserve Board Working Paper. SSRN 3915439.

A The specifications, inference, and robustness

A.1 The set of specifications

The four return types and the values of the thirteen measurement choices (Table 1) combine into 882,432 specifications. I drop the 2,115 specifications with a return gradient beyond 100 points in absolute value. Nearly all of them take the mean without an outlier trim and rank on start-of-wave or midpoint wealth. That leaves 880,317 specifications, 868,817 in the PSID and 11,500 in the SCF. The two exclusions of Appendix B apply under every specification, the option without an outlier trim included.

The values are those the seven existing studies use, intermediate values, and the wealth-varying interest rate of Section 3. Variance shares measure the spread across these values (Online Supplement SA.H.2).

A.2 Robustness of the return gradient on the gross return on assets

Nine robustness checks keep the median return gradient on the gross return on assets below the steep existing headlines (Table 5). First, the common period 2001 to 2022 gives a median of $+2.8$ in the PSID and -1.6 in the SCF. Second, the rise of a line fitted from the fifth wealth decile to the top gives $+2.0$ and -0.2. Third, a \$100,000 wealth floor gives $+7.2$ and $+0.5$. Fourth, an imputed-rent yield of 3 percent gives $+4.0$ and $+0.4$, and of 7 percent $+1.5$ and -1.3. Fifth, the SCF without its survey weights gives $+0.5$. Sixth, a pension equity share of 50 percent gives -0.2 in the SCF and of 70 percent 0.0. Seventh, ranking on gross assets gives $+0.2$ and 0.0. Eighth, nominal returns give $+3.1$ and -0.9. Ninth, the 1,816 PSID households in at least six return waves give $+3.2$.

A.3 The return gradient across subperiods

Across four subperiods the PSID median return gradient on the gross return on assets is between $+0.9$ and $+3.8$ points (Table 5). The SCF median is between -3.8 and $+0.2$. The SCF return gradient moves more across subperiods, because its synthetic return varies only with the wave's asset-class returns. No subperiod in either survey shows a steep positive return gradient.

Table 5: Robustness of the return gradient on the gross return on assets

	PSID		SCF	
Check	Median	Mean	Median	Mean
Baseline	+2.8	+3.1	0.0	+0.8
Common period, SCF 2001 to 2022	+2.8	+3.1	−1.6	−0.7
Fitted slope	+2.0	+2.1	−0.2	+0.6
Wealth floor \$10,000	+3.0	+3.9	−0.1	+0.8
Wealth floor \$100,000	+7.2	+10.6	+0.5	+1.5
Imputed rent, 3 percent yield	+4.0	+4.0	+0.4	+1.7
Imputed rent, 7 percent yield	+1.5	+2.0	−1.3	−0.1
Unweighted SCF	–	–	+0.5	+1.7
Pension equity share 50 percent	–	–	−0.2	+0.7
Pension equity share 70 percent	–	–	0.0	+0.9
Ranking on gross assets	+0.2	0.0	0.0	+0.7
Nominal returns	+3.1	+3.1	−0.9	+0.8
Households in at least four waves	+3.0	+3.8	–	–
Households in at least six waves	+3.2	+4.3	–	–
Subperiod 2001–2005	+3.8	+5.0	−3.8	−4.9
Subperiod 2007–2011	+2.6	+3.1	+0.2	+1.7
Subperiod 2013–2017	+3.1	+2.9	−2.4	+0.8
Subperiod 2019–2023	+0.9	+1.8	−0.4	−0.9

Notes: Each entry is the median or the mean return gradient on the gross return on assets at the baseline of Section 5.1, on heads aged 20 to 80, with one check applied, in percentage points. A dash marks a check that does not apply to the survey. The baseline row is the estimate of Section 5.2. The fitted-slope row is the rise from the fifth wealth decile to the top of a line fitted through the wealth-decile returns. The unweighted-SCF row drops the survey weight within wealth deciles formed on the weighted ranking. The PSID is unweighted throughout. The panel rows keep the PSID households observed in at least four or six return waves. The subperiod rows compute the return gradient within the subperiod, from three PSID waves and two SCF waves each.

A.4 Age and period adjustment

The baseline return gradient is an unconditional difference between wealth-decile returns, and older households hold more wealth and earn different returns. To separate age from wealth, I regress the household gross return on assets on wealth-decile dummies, a quadratic in age, and year fixed effects, and read the return gradient off the top-wealth-decile coefficient against the fifth. The adjusted mean return gradient is +4.6 in the PSID, against the unconditional +3.1, and the SCF moves from +0.8 to +0.9. Age does not account for the return gradient in either survey.

B Constructing the household return

Table 6 lists which input of the return each survey observes and which it imputes. The PSID observes its capital income and net investment and imputes its capital gains and most of its cost of debt. The SCF imputes every flow. Appendix B.1 builds the PSID return from those inputs and Appendix B.2 the SCF return.

B.1 The PSID realized return

Four timing conventions apply. First, the wave return is annual: the capital income covers the calendar year before the interview, the capital gains enter as half the two-year value change, and the interest paid is annual. Second, total debt is at the wave midpoint under every denominator timing. Third, on the long-run horizon a household takes the wealth decile it holds in most of its waves. Fourth, the return is real, deflated by the average annual change in the consumer price index over the wave:

$$r_{it} = \frac{1 + r_{it}^{\text{nom}}}{\left(\text{CPI}_t / \text{CPI}_{t-1}\right)^{1/2}} - 1.$$

where CPI_t is the consumer price index in the year of wave t and consecutive waves are two years apart.

A first home purchase is not a capital gain. The PSID reports no purchase flow for the owner-occupied home, so a home that first appears between two waves earns no capital gain in that wave. Bracketed asset values enter at their bracket midpoints. A switch between a bracket and a reported value therefore creates a spurious capital gain, and the outlier trim and the wealth floor remove the largest ones.

Through 2011 the PSID records other real estate net of its mortgage, so I recover the value and the debt with a loan-to-value model estimated on the SCF (Online Supplement SA.L.4). Interest on every debt but the primary mortgage is the balance times a FRED rate: the mortgage rate plus 0.625 points on investment-property debt, the consumer-credit rate on consumer debt, and the new-car finance rate on vehicle loans (Online Supplement SA.E). Under the wealth-varying rate of measurement choice 2, a credit spread by wealth decile is added to each rate (Section 3.2).

Table 6: What the PSID and the SCF observe in the return identity

	PSID	SCF
Capital income Y	Observed	Imputed
Capital gains G	Imputed	Imputed
Net investment N	Observed	Not observed
Cost of debt C	Mortgage interest observed	Imputed
Asset and debt stocks	Observed	Observed
Waves	12, 2001 to 2023	12, 1989 to 2022
Head age	20 to 80	20 to 80
Households	7,174	–
Observations	27,482 household-years	223,068

Notes: "Observed" is a recorded survey item and "imputed" a value built from other inputs. The PSID observes the interest on the primary mortgage and imputes the rest. The SCF records capital income, but the synthetic return assigns each asset class its market yield. Sample sizes are for the gross return on assets at the baseline wealth floor and outlier trim. A dash marks the SCF, whose observations are household records across five imputation replicates, so it has no household count. "Not observed" is an input the synthetic return does not use.

Every specification drops a household-year on two conditions. First, its net return on equity at the baseline is -100 percent or below. Second, its capital income is below $-\$800{,}000$. Only a business or farm loss reaches that value.

B.2 The SCF synthetic return

Each of five asset classes earns the wave's market return, from Jordà et al. (2019) and FRED. Equity and business earn the listed-equity capital gains and dividend yield. Housing earns the housing capital gains, a 5 percent imputed rent on the owner-occupied home, and the market rental yield on other residential property. Fixed income earns the bill rate, and pensions a sixty-forty mix of the equity and bill returns. The 2022 wave takes each series at its most recent observed year, 2021 for the two capital gains and 2020 for the rent yield, the dividend yield and the bill rate.