Wealth Proxies in the PSID:
A Wealth Series from 1969

Christophe Van Langenhove*

September 2026

Abstract

The Panel Study of Income Dynamics (PSID) records wealth from 1984. Pfeffer and Killewald (2018) and Fisher and Johnson (2023) predict the earlier waves from one or two of the sixteen characteristics the PSID records from 1969. They compare their wealth proxies against actual wealth on an intergenerational rank-rank slope alone. I build a gradient-boosting (GB) wealth proxy on all sixteen characteristics. Over the 16 waves recording wealth it outperforms both wealth proxies on the wealth rank distribution, intergenerational wealth mobility, wealth inequality and wealth accumulation trajectories, cutting the wealth rank error by 27 to 34% and more than halving the Gini distance. The wider set of input variables, not the estimator, explains the outperformance. The GB-proxy covers every pre-1984 household-year, giving PSID research a 1969 to 2023 wealth series with three biases of known size: it overstates intragenerational wealth persistence, understates wealth inequality and overstates wealth accumulation.

JEL classification: C81, D31, J62

Keywords: PSID, wealth measurement, imputation, wealth mobility, gradient boosting

*Vrije Universiteit Brussel (VUB) (email: christophe.van.langenhove@vub.be). I acknowledge the financial support from the Research Foundation Flanders (FWO, project number: 11I5324N). I am grateful to Arthur Apostel, Paula Gobbi, Freddy Heylen, Selma Malmberg, Marthe Mareels, Ive Marx, Yasin Kürşat Önder, Gert Peersman, Alberto Russo and Dirk Van de gaer for valuable comments. I also thank participants at the UGent internal seminar, ESPANet Early Career Researchers Day, 20th PSE Doctorissimes Conference, Inequalities & Opportunities Conference, 11th ECINEQ Conference, ISI Wealth Conference 2025, and World Inequality Conference for discussion and feedback. This paper previously circulated as part of "Wealth Mobility in the United States: Empirical Evidence from the PSID".

1 Introduction

The Panel Study of Income Dynamics (PSID) runs from 1968 to the present and records household wealth only from 1984. Two published wealth proxies predict wealth in the pre-1984 waves. Pfeffer and Killewald (2018) (PK) rank households by the value of the main residence and set renters' value to zero. Fisher and Johnson (2023) (FJ) impute home equity and non-home-equity wealth from the housing value and from interest and dividend income. I call their two series the PK-proxy and the FJ-proxy.

Both wealth proxies have two shortcomings. First, they do not use all the socio-economic information the PSID holds. The PSID records sixteen household characteristics in every wave since 1969. The PK-proxy takes one input variable and the FJ-proxy two. Second, PK and FJ compare their wealth proxies against actual wealth on one performance metric, the parent-child rank-rank slope. Neither paper compares its wealth proxy against actual wealth on the four other dimensions: the overall wealth rank distribution, intragenerational wealth mobility, wealth inequality and wealth accumulation trajectories.

I construct the GB-proxy with gradient boosting from all sixteen characteristics the PSID has recorded since 1969. I compare it against actual wealth in the 16 waves that record wealth: 1984, 1989, 1994, and every second year from 1999 to 2023. The comparison covers five dimensions: the overall wealth rank distribution, intragenerational wealth mobility, intergenerational wealth mobility, wealth inequality and wealth accumulation trajectories. I obtain three results.

First, the GB-proxy outperforms the PK-proxy and the FJ-proxy on four of the five dimensions, all but intragenerational wealth mobility. It cuts the wealth rank error from 21 and 19 units to 14 and the Gini distance from 0.16 and 0.17 to 0.06. It outperforms both in every quintile of the wealth distribution. Its error is largest among the poorest households: it misplaces 17% of the poorest 20% by more than 25 rank units, against 3% of the richest 20%. On households held out of its training sample, its wealth rank error is 32 units among renters against 15 among homeowners.

Second, the outperformance comes mostly from the wider set of socio-economic characteristics the GB-proxy takes as input variables and not from its estimator. The set of input variables accounts for 84% of the improvement on wealth inequality, 97% on wealth accumulation trajectories, 93% on the overall wealth rank distribution and 94% on intergenerational wealth mobility. On intragenerational wealth mobility the estimator, not the set of input variables, creates the overstatement of wealth persistence.

Third, the GB-proxy gives future research a PSID wealth series over the whole period 1969 to 2023, with three biases of known size. It overstates intragenerational wealth persistence: over ten years its rank-rank slope is 0.76 against 0.70 in actual wealth. It understates wealth inequality: its Gini lies below actual wealth's in every wave that records actual wealth, by 0.06 on average against an actual Gini of 0.78. It overstates wealth accumulation: within a birth cohort, median wealth rises 10.4 times between ages 30 to 34 and ages 60 to 64, against 8.5 times in actual wealth.

Literature I contribute to three strands. First, Hurst, Luoh and Stafford (1998) introduce the PSID's 1984 to 1994 wealth supplements. Pfeffer et al. (2016) compare the PSID's wealth measure against the Survey of Consumer Finances. Cooper, Dynan and Rhodenhiser (2019) examine the PSID's treatment of retirement assets. Si et al. (2023) test a multiple imputation of missing wealth components against the imputation the PSID itself applies. I compare the wealth proxies built for the 1969 to 1983 waves, which record no wealth, against the wealth the PSID records from 1984.

Second, supervised learning turns input variables into estimates of the income distribution, wealth inequality and intergenerational mobility. Bloise, Brunori and Piraino (2021) impute parental income in the PSID with a supervised learner. Salas-Rojo and Rodríguez (2022) use a random forest to group households by the inheritance they received. Cornwall and Gindelsky (2025) nowcast the income distribution from national accounts aggregates with an elastic net. Each selects its model on out-of-sample prediction error, yet a low prediction error is not enough: the GB-proxy has the lowest wealth rank error of the three wealth proxies and still understates wealth inequality and overstates wealth accumulation.

Third, housing values serve as a wealth proxy outside the PSID. Pfeffer and Killewald (2018) propose them for non-survey sources, naming the 1940 Census and Zillow. Kalsi and Ward (2025) rank households in the 1850 to 1940 Censuses on real estate and home values. Binder, Risch and Voorheis (2025) build housing wealth from property tax assessments and deed records. Each ranks households on one asset and assigns renters none. Surveys modelled on the Demographic and Health Surveys record neither income nor expenditure. Studies of them rank households on an asset index, which weights indicators for owning a radio, a television or a car and for piped water or a flush toilet. The weights come from principal components (Filmer and Pritchett, 2001) or factor analysis (Sahn and Stifel, 2003).

The result extends beyond the PSID. Rebuilt in the Survey of Consumer Finances, the PK-proxy leaves a Gini distance of 0.17 against 0.16 in the PSID. Wherever a dataset records labour income, capital income, the mortgage and the rent beside the housing value, a wealth proxy that takes all of them as input variables should outperform one built on the housing value alone.

Roadmap Section 2 sets out the PSID sample, defines actual wealth and builds the three wealth proxies. Section 3 compares the three wealth proxies against actual wealth on all five dimensions. Section 4 decomposes the outperformance between the set of input variables and the estimator, and explains why the GB-proxy's error is largest among the poorest households and renters. Section 5 concludes.

2 PSID sample and three wealth proxies

All three wealth proxies run on one PSID sample, so a comparison between them moves the wealth proxy and nothing else. On it I rebuild the two published wealth proxies, Pfeffer and Killewald (2018) (PK) and Fisher and Johnson (2023) (FJ), and fit the GB-proxy.

2.1 PSID samples

The PSID records actual wealth in 16 waves: 1984, 1989, 1994, and every second year from 1999 to 2023. These waves form the 1984 to 2023 sample, which holds 67,665 household-years descending from 1,923 families interviewed in 1968. Every comparison against actual wealth uses this sample. I fit the FJ-proxy on it as well. I fit the GB-proxy on a wider set: every household-year in these waves with recorded wealth. Actual wealth is the PSID's own wealth variable:

$$w_{i,t} = a_{i,t} - d_{i,t} \tag{1}$$

where $a_{i,t}$ is household i's real and financial assets at t and $d_{i,t}$ its debts. Both exclude vehicles.

The other waves record no wealth: 1969 to 1983, and the annual waves from 1985 to 1997 other than 1989 and 1994. A wealth proxy predicts wealth in those waves from input variables the PSID records in every wave from 1969. The GB-proxy therefore covers 173,036 household-years over 1969 to 2023, 48,859 of them before 1984.

Every series uses the PSID's Survey Research Center (SRC) sample, drawn in 1968 to be nationally representative, and excludes the low-income oversample of that year and the immigrant samples added from 1997. On the oversample the GB-proxy does not rank best, and on the oversample pooled with this sample it does (Supplement SA.B). The PSID interviews the same families every year to 1997 and every second year since. Children who leave home and form their own households stay in the panel. No series is survey-weighted. I report the survey-weighted variants in Appendix B and Supplement SA.B.

2.2 Published proxies: the PK-proxy and FJ-proxy

Pfeffer and Killewald (2018) predict wealth as the housing value multiplied by a factor:

$$\hat{w}_{i,t}^{\text{PK}} = \begin{cases} f_h \, h_{i,t} & \text{if } h_{i,t} > 0 \\ f_r \, r_{i,t} & \text{if } h_{i,t} = 0 \end{cases} \tag{2}$$

where $h_{i,t}$ is the gross value of household i's main residence in wave t, $r_{i,t}$ is its rent, and $\hat{w}_{i,t}^{\text{PK}}$ is the PK-proxy's prediction.[1] A hat denotes a prediction and a superscript the series. The factors f_h and f_r convert the housing value and the rent into a wealth level. The PSID records both from 1969.[2] PK set f_h to one and f_r to zero: a homeowner's prediction is the housing value itself and a renter's is zero. PK do not state how renters rank among themselves. I draw their order at random and report two alternatives as robustness: all renters tied, and renters ordered by rent (Appendix A.1).

Fisher and Johnson (2023) report their two input variables, the housing value and interest and dividend income, but not the functional form. I fit a linear regression of actual wealth on the housing value and capital income, through the origin:

$$w_{i,t} = \beta_h \, h_{i,t} + \beta_y \, y_{i,t} + \varepsilon_{i,t} \tag{3}$$

where $y_{i,t}$ is capital income, measured as family income less labour income, other members' taxable income and transfer income. β_h and β_y are the coefficients on the housing value and on capital income, and $\varepsilon_{i,t}$ is the residual (Supplement SA.C). On the 1984 to 2023 sample the fitted $\hat{\beta}_h$ is 2.15 and $\hat{\beta}_y$ is 5.9. The FJ-proxy is the fitted value $\hat{w}_{i,t}^{\text{FJ}} = \hat{\beta}_h \, h_{i,t} + \hat{\beta}_y \, y_{i,t}$.

[1] The PSID records the housing value in round numbers. In 23% of consecutive annual pairs a homeowner's housing value does not change at all. Even if every repeat were a value repeated from an earlier year without revaluing the house, the level error would be 2% of housing wealth (Supplement SA.F).

[2] In 1988 and 1989 the PSID records whether a household rents but not the amount. For those two waves I interpolate a renter's rent linearly between its nearest reported rents in earlier and later waves. A renter with a reported rent on one side only takes that rent.

The PK-proxy therefore takes one of the sixteen characteristics the PSID records in every wave and the FJ-proxy two. PK and FJ each compare theirs against actual wealth on one performance metric, the parent-child rank-rank slope.

2.3 The gradient-boosting (GB) wealth proxy

I build the GB-proxy with a gradient-boosting machine learning model (Friedman, 2001). It takes as input variables all sixteen characteristics the PSID records in every wave since 1969 (Table 1). I estimate three models: one for homeowners, one for renters and one for the 5% of households in a wave that report neither a housing value nor rent. The homeowner model adds the housing value and the mortgage ratio. The renter model adds the rent. I state the reason for each in Appendix A.

The thirteen shared input variables fall into three groups (Table 1). Seven describe the household and its income: age, household size, family structure, labour income, capital income, whether it holds a business and the years a member could not work. Three measure need: food-stamp receipt, income over the needs standard and labour income against a year at the minimum wage. Three measure wealth accumulation and surplus: the waves spent owning a home, rooms per person and money given to others. Dollar input variables enter in constant dollars and the rest as ratios, shares or indicators. For every input variable, at most a tenth of the pre-1984 household-years lie below the fifth percentile of its 1984 to 2023 values.

Each model predicts a household's wealth relative to an aggregate of its wave rather than in dollars. For a homeowner, $\tau_{i,t}$ is the household's wealth over its housing value, divided by the same ratio for the country as a whole. For every other household, $\tau_{i,t}$ is its wealth divided by mean wealth per household:

$$\tau_{i,t} = \begin{cases} \dfrac{w_{i,t}/h_{i,t}}{F(t)} & \text{if } h_{i,t} > 0 \\ \dfrac{w_{i,t}}{\bar{w}(t)} & \text{otherwise} \end{cases} \tag{4}$$

where $w_{i,t}$ is actual wealth, $h_{i,t}$ the housing value, $F(t)$ the ratio of aggregate household and nonprofit wealth to the aggregate value of residential real estate in wave t, and $\bar{w}(t)$ aggregate household and nonprofit wealth divided by the number of households. The aggregates come from the Financial Accounts of the United States and the household count from the Census. The GB-proxy is the prediction with the scaling undone: $\hat{w}^{\mathrm{GB}}_{i,t} = \hat{\tau}_{i,t}\, F(t)\, h_{i,t}$ for a homeowner and $\hat{w}^{\mathrm{GB}}_{i,t} = \hat{\tau}_{i,t}\, \bar{w}(t)$ for everyone else.

Table 1: The input variables, and which wealth proxy takes each

	Used by		
Input variable	GB-proxy	PK-proxy	FJ-proxy
Demographics and income			
Age	•		
Household size	•		
Family structure	•		
Labour income	•		
Capital income	•		•
Business held	•		
Years unable to work	•		
Bottom of the distribution			
Food stamps received	•		
Income to needs	•		
Minimum wage ratio	•		
Wealth accumulation and surplus			
Waves spent owning	•		
Rooms per person	•		
Money given away	•		
Homeowners only			
Housing value	•	•	•
Mortgage ratio	•		
Renters only			
Rent	•		
Input variables in total	16	1	2

Notes: The PSID records every input variable from its first wave in 1969, so none limits how far back a wealth proxy reaches. Demographics and income, Bottom of the distribution and Wealth accumulation and surplus enter every model. Homeowners' model adds the housing value and the mortgage ratio to those shared groups, while renters' model adds the rent. Households that report neither a housing value nor rent get a third model, fitted on the shared groups by themselves. The GB-proxy's total counts each input variable once, whichever models it enters. The PK-proxy is Pfeffer and Killewald's (2018). The FJ-proxy follows Fisher and Johnson (2023), fitted here on the housing value and the capital-income residual.

Why scale by a wave aggregate? A gradient-boosting model predicts $\tau_{i,t}$ as an average over households of the 1984 to 2023 sample with similar input variables, so its predictions stay within that sample's range. In the Financial Accounts, household wealth relative to residential real estate lies between 2.52 and 3.54 over 1984 to 2023 and between 2.97 and 3.87 before 1984. A model fitted on wealth over the housing value alone would give a pre-1984 household a ratio from the 1984 to 2023 range, whatever the aggregate ratio in its wave. Dividing by $F(t)$ removes the movement between waves from $\tau_{i,t}$, so the model only places a household within its wave. Both aggregates exist from 1969, and multiplying by each wave's aggregate restores the level.

The GB-proxy is the model's point prediction, one value for each combination of input variables. Two households with the same input variables get the same prediction, whatever their actual wealth. The predictions are therefore less dispersed than the wealth behind them. Before fitting, I drop each model's training household-years with $\tau_{i,t}$ outside its 2.5th to 99.5th percentiles. Dropping that highest half percent compresses the top of the predictions further. A robustness series adds to each point prediction a random draw of the part the model leaves unexplained, which restores the missing dispersion (Supplement SA.H). Every result below uses the point prediction (Appendix A). On wealth inequality the robustness series comes closer to actual wealth than the point prediction does (Supplement SA.C).

Which input variables contribute most to the predictions? A Shapley contribution attributes each household's prediction to its input variables. For homeowners the mortgage ratio accounts for 27% of the summed absolute contributions, and with capital income and income to needs the three leading input variables account for 53%. For renters income to needs accounts for 29%, and with capital income and age the three leading input variables account for 67%. I state each model's tuning in Appendix A. I report every input variable's share, the dependence of the prediction on each leading input variable and the learning curves in Supplement SA.H.

3 Three wealth proxies against actual wealth

Each of the five dimensions has its own performance metric, computed on the 1984 to 2023 sample as the gap between a wealth proxy's estimate and actual wealth's. The GB-proxy outperforms the PK-proxy and the FJ-proxy on four of them: the overall wealth rank distribution, intergenerational wealth mobility, wealth inequality and wealth accumulation trajectories. On the fifth, intragenerational wealth mobility, it overstates wealth persistence.

3.1 The overall wealth rank distribution

Performance metric The performance metric $\mathcal{M}$ for the overall wealth rank distribution compares a household's rank under a wealth proxy with its rank under actual wealth within each wave, and averages across waves:

$$\mathcal{M} = \frac{1}{T}\sum_{t=1}^{T} g\left(\frac{1}{N_t}\sum_{i=1}^{N_t} m(\kappa_{i,t}, \hat{\kappa}_{i,t})\right) \tag{5}$$

where $\kappa_{i,t}$ is household i's percentile rank under actual wealth in wave t, $\hat{\kappa}_{i,t}$ is its rank under the wealth proxy, N_t is the number of households in that wave and T the number of waves. The function m takes four forms: a squared error, an absolute error, and an indicator that the two ranks differ by more than 25 or by more than 50 rank units. The function g is the square

Table 2: Performance on the overall wealth rank distribution

	Rank units		% of households	
Wealth proxy	Rank error	Absolute error	Beyond 25 ranks	Beyond 50 ranks
GB-proxy	14	10	7	0.8
PK-proxy	21	15	19	3
FJ-proxy	19	14	15	2

Notes: Every column is a distance from the ranking under actual wealth, so a lower figure is a closer wealth proxy. Every performance metric is Equation 5 under a different loss, computed within a wave and averaged with equal weight over the 16 waves of the 1984 to 2023 sample, on the household-years where actual wealth, the three wealth proxies and the rent-ordered variant of Appendix A.1 are all present. Wealth rank error is the root mean squared percentile-rank error and absolute error the mean absolute error. The last two columns are the percentage of households a wealth proxy places more than 25 and more than 50 rank units from their rank under actual wealth. PK's renters are ranked at random, one of three orderings the published construction leaves open. Appendix A.1 scores all three.

root under a squared error and leaves the average unchanged under the other three. The four performance metrics are therefore the root mean squared wealth rank error (RMSE), the mean absolute wealth rank error, and the shares of households placed more than 25 and more than 50 rank units from their actual rank.

Results The GB-proxy outperforms the PK-proxy and the FJ-proxy on all four performance metrics, in each of the 16 waves of the 1984 to 2023 sample (Table 2). Its RMSE is 14 rank units, 34% below the PK-proxy's and 27% below the FJ-proxy's. It places 7% of households more than 25 rank units from their actual rank, against 19% under the PK-proxy and 15% under the FJ-proxy. For the FJ-proxy and the GB-proxy these figures are in sample, computed on the households in the fit. On households held out of its training sample, the GB-proxy's error is 13.90 rank units against 18.29 for the FJ-proxy (Appendix A).

The GB-proxy outperforms the PK-proxy and the FJ-proxy in every quintile of actual wealth on both performance metrics (Figure 1). The three wealth proxies differ in where their errors fall. The GB-proxy's and the FJ-proxy's errors are largest among the poorest households. The GB-proxy places 17% of the poorest wealth quintile more than 25 rank units from their actual rank, against 3% of the richest. A rank unit spans fewer dollars at the bottom, so the same dollar error moves a poor household further. The PK-proxy's error is spread evenly: it misplaces 18% of the richest wealth quintile and 20% of the poorest. In the poorest wealth quintile the PK-proxy with renters tied misplaces as many households as the GB-proxy, within sampling error (Appendix A.1, Appendix B).

Figure 1: Performance on the overall wealth rank distribution, by quintile of actual wealth

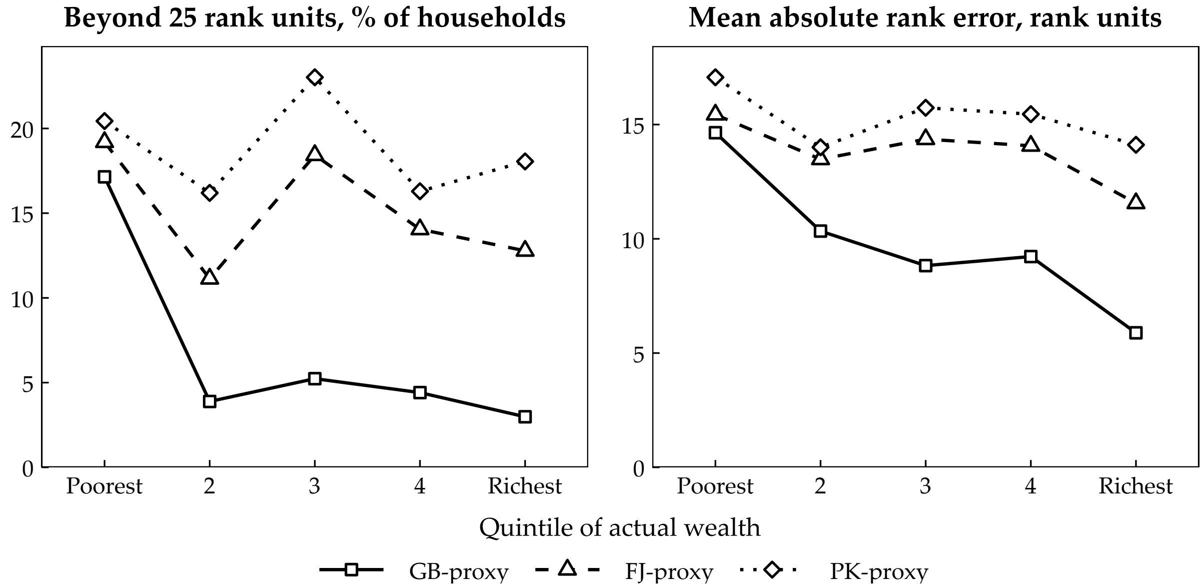

Notes: The panels split households into quintiles of actual wealth within each wave of the 1984 to 2023 sample and pool them over the 16 waves. The left panel is the percentage of households a wealth proxy places more than 25 rank units from their rank under actual wealth. The right panel is the mean absolute wealth rank error, in rank units. An open square on a solid line marks the GB-proxy, an open triangle on a dashed line the FJ-proxy and an open diamond on a dotted line the PK-proxy.

The GB-proxy's error is smallest for homeowners and largest for renters: on households held out of its training sample, its wealth rank error is 15 units for homeowners, 32 for renters and 29 for households that report neither a housing value nor rent. Each group's wealth rank error uses ranks within the group (Appendix A). Homeowners are 64% of the 1984 to 2023 sample, renters 30% and households reporting neither 6%. In the pre-1984 waves the shares are 64%, 31% and 5%.

Robustness These findings are robust to four of five tests. First, dropping any one of the 16 waves leaves the GB-proxy's wealth rank error between 13.6 and 13.9 units, the lowest each time. Second, the PSID's longitudinal weight and a reweighting for attrition each move that error by less than one unit (Appendix B). Third, resampling the 1968 families places the GB-proxy's error between 13 and 14 units, below the 21 and 19 of the PK-proxy and the FJ-proxy. Fourth, a plain rank correlation with actual wealth gives the same ordering as the four performance metrics (Appendix B). Fifth, on the low-income oversample alone the PK-proxy with renters tied has the lowest wealth rank error, 22 units against the GB-proxy's 25. On both samples pooled the GB-proxy's is lowest, 17 against 20 (Supplement SA.B).

3.2 Intragenerational wealth mobility

Performance metric I measure intragenerational wealth mobility by the rank-rank slope of a household on itself. The unit is the household. I regress its within-cohort wealth rank ten years after an initial wave on its rank at that initial wave, with a fixed effect for the initial wave:

$$\kappa_{i,t+10} = \alpha_t + \beta\, \kappa_{i,t} + \varepsilon_{i,t} \tag{6}$$

where $\kappa_{i,t}$ is household i's wealth rank within its birth cohort in wave t, α_t absorbs the initial wave, β is ten-year wealth persistence and $\varepsilon_{i,t}$ is the residual. On individuals followed from their early thirties to their late fifties in the PSID, Shiro et al. (2022) report a slope of 0.59. A wealth proxy's gap is its β less actual wealth's on the same 14,822 pairs (Table 3). The initial waves start in 1999, the reference person's age at the initial wave is 20 to 64, and the standard errors come from 500 resamples of the 1968 family. A pair needs the same reference person and partner at both dates. The sample therefore holds the 52% of households the PSID observes again ten years on with the same reference person and partner.[3]

Results The GB-proxy overstates ten-year wealth persistence, and the overstatement comes from the poorest households and homeowners (Table 3, Figure 2). I obtain two findings. First, the GB-proxy's slope is 0.76 against 0.70 in actual wealth, a gap of 0.07 and the largest of the three wealth proxies. Under the GB-proxy 62% of the poorest wealth quintile remain in it ten years on, against 51% under actual wealth. Among households that own at both dates its slope is 0.80 against 0.68, and among households that rent at both dates 0.48 against 0.47.

Second, the FJ-proxy's closer aggregate slope comes from errors that offset between homeowners and renters, and the PK-proxy's depends on the renter ordering. The FJ-proxy's slope is 0.83 against 0.68 among households that own at both dates and 0.34 against 0.47 among households that rent at both dates. The PK-proxy's slope is 0.66 with renters drawn at random and 0.71 with renters tied, against 0.70 in actual wealth, and 0.85 among households that own at both dates (Appendix A.1).

Robustness The GB-proxy overstates ten-year wealth persistence under every one of six changes to the sample. The overstatement survives 1984 and 1989 added as starting waves, on each household's first ten-year pair alone, at ages 30 to 54, and under the PSID's longitudinal weight. It also survives when I follow a household through its reference person alone,

[3] Of the households that start in the poorest quintile of actual wealth, 47% form a ten-year pair, against 58% of those that start in the richest.

Table 3: Performance on intragenerational wealth mobility

	Rank-rank slope		Ten years						Thirty years
Series	Two years	Ten years	Gap to actual (se)	Hybrid slope	Stay top wealth quintile, %	Stay bottom wealth quintile, %	Mean move, rank units	Spread ratio	Ages 30-34 to 60-64
Actual wealth	0.83	0.70	0.00 (0.00)	0.70	66	51	15.7	1.00	0.59
GB-proxy	0.87	0.76	0.07 (0.01)	0.71	67	62	14.0	0.89	0.66
FJ-proxy	0.83	0.70	0.01 (0.01)	0.63	65	57	15.0	0.97	0.58
PK-proxy, renters at random	0.78	0.66	-0.03 (0.01)	0.58	66	52	15.4	0.99	0.53
PK-proxy, renters tied	0.86	0.71	0.01 (0.01)	0.62	67	65	14.1	0.91	0.55
GB-proxy, draws (calibration check)	0.85	0.71	—	0.55	—	—	15.5	—	—

Notes: Each row pairs a household with itself two or ten years later on the 1984 to 2023 sample, with initial waves from 1999 and the reference person aged 20 to 64 at the initial wave, ranked within five-year birth cohort and wave. The 14,822 ten-year pairs and 32,389 two-year pairs are 52% and 75% of the households observed at the initial wave. Slopes are coefficients of the end rank on the initial rank within initial wave, with both ranks from the series in the row. The hybrid column takes the end rank from actual wealth instead, as a study of the pre-1984 waves would. The gap is the series' slope less actual wealth's, with a standard error from 500 resamples of the 1968 family. Within a wealth quintile, the spread is the ninetieth less the tenth percentile of the rank change. The spread ratio divides the series' spread, averaged over the five quintiles, by actual wealth's. The PK-proxy's slope depends on how renters are ordered, by a fresh random draw in each wave in its first row and by one tied value in its second, which bracket it. The draws row follows a wealth persistence process fitted to actual wealth, so it tests that fit and not the series. The last column ranks a household within its birth cohort on its median wealth at ages 30 to 34 and again at 60 to 64, over 504 households observed at both ages.

Figure 2: Ten-year rank change by starting wealth quintile, each series against actual wealth

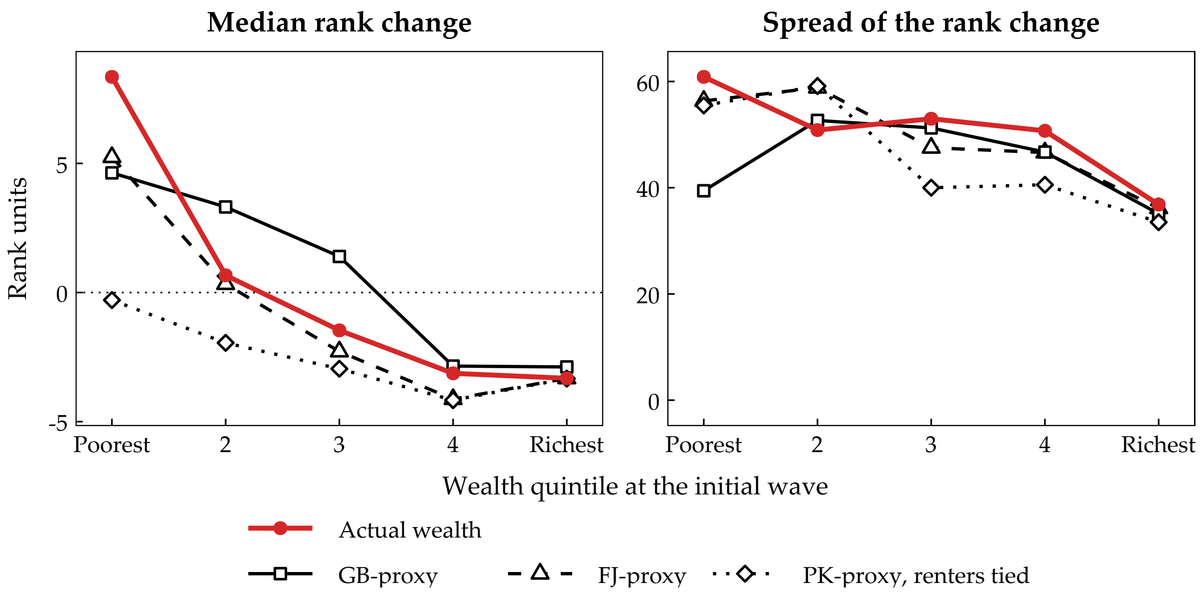

Notes: Each panel pairs a household with itself ten years later on the 1984 to 2023 sample, initial waves from 1999, and groups the pairs by the household's wealth quintile at the initial wave under that series. The left panel shows the median change in within-cohort rank, in rank units, with a dotted line at no change, and the right its spread, the ninetieth percentile less the tenth. In the left panel the GB-proxy's median rise out of the poorest wealth quintile falls short of actual wealth's, which is the overstatement of wealth persistence that Section 3.2 reports. Actual wealth is the red filled circles on a solid line and the GB-proxy the open squares on a solid line. The FJ-proxy is the open triangles on a dashed line and the PK-proxy with renters tied the open diamonds on a dotted line. I draw the tied ordering because a random one redraws every renter's rank in each wave, which is not movement.

which keeps the households whose partner changes over the decade. Dropping each initial wave in turn keeps the GB-proxy's slope between 0.76 and 0.77, against 0.69 to 0.70 in actual wealth (Appendix B, Supplement SA.D).

3.3 Intergenerational wealth mobility

Performance metric I measure intergenerational wealth mobility by the rank-rank slope, the child's within-cohort wealth rank regressed on the parent's, in a parent-child and a grandparent-grandchild comparison. Both rank individuals, because a parent and an adult child live in different households. An individual's wealth is a share of their household's: the whole for a single reference person, half for a married one and their share of household labour income for a cohabiting one (Appendix A).

Parent-child comparison The parent-child comparison uses the 1984 to 2023 sample, where both generations have actual wealth. I regress the child's rank on the parent's rank at the same lifecycle stage, each a five-year age band. A wealth proxy's gap is its slope less actual wealth's, estimated on the same parent-child pairs in one stacked regression:

$$\kappa^c_{i,s} = \alpha_s + \beta_s\, \kappa^p_{i,s} + \gamma_s\, \kappa^p_{i,s} D_i + \delta_s\, D_i + \varepsilon_{i,s} \tag{7}$$

where $\kappa^p_{i,s}$ and $\kappa^c_{i,s}$ are the parent and child within-cohort wealth ranks at lifecycle stage s, α_s is its intercept, β_s the parent-child rank-rank slope under actual wealth, D_i marks the wealth proxy series, γ_s is the gap, δ_s the wealth proxy's own level shift and $\varepsilon_{i,s}$ the residual. Standard errors cluster on the 1968 family. Table 4 reports two panels. Panel A ranks both generations under the wealth proxy. Panel B ranks the parent under the wealth proxy and the child under actual wealth, as a study of the pre-1984 waves must. Under actual wealth the slope rises from 0.34 at ages 30 to 34 to 0.49 at ages 60 to 64 (Table 4).[4]

Grandparent-grandchild comparison The grandparent-grandchild comparison uses the pre-1984 waves. The PSID records actual wealth for the grandchild and the parent. Each individual's rank is the mean of their within-cohort wealth ranks over the waves that observe them aged 25 to 75. Averaging over waves reduces the attenuation a single noisy wave causes (Solon, 1992). The grandparent has actual wealth from 1984 only, and a wealth proxy from 1969. I regress the grandchild's mean rank on the grandparent's:

$$\bar{\kappa}^{\text{gc}}_i = \alpha + \beta\, \bar{\kappa}^{\text{gp}}_i + \varepsilon_i \tag{8}$$

where $\bar{\kappa}^{\text{gc}}_i$ and $\bar{\kappa}^{\text{gp}}_i$ are the mean ranks of the grandchild and of the grandparent in pair i, the bar denoting the average over the waves that observe them, α is the intercept and ε_i the residual. A wealth proxy's gap is the change in β when it replaces actual wealth in the grandparent's rank, estimated as in Equation 7 on the pairs where the grandparent has both ranks.

Results The GB-proxy's parent-child slope lies closest to actual wealth's (Table 4). I obtain three findings. First, the GB-proxy's gap, its slope less actual wealth's, is jointly indistinguishable from zero across the seven lifecycle stages. A joint test gives a p-value of 0.62 (Panel B). Second, the PK-proxy and the FJ-proxy place the slope below actual wealth's at all seven lifecycle stages, and so understate wealth persistence. A joint test rejects zero for the PK-proxy's gaps, at 0.01, and not for the FJ-proxy's, at 0.06 (Supplement SA.D).[5] Third, in Panel A the wealth proxy replaces actual wealth for the child as well as the parent. There the GB-proxy

[4] On every parent-child pair with actual wealth it rises from 0.34 to 0.48 (Van Langenhove, 2026*a*).

[5] Neither published comparison isolates the wealth proxy's own error. Pfeffer and Killewald (2018) take a rank slope of 0.37 in housing values against 0.39 in wealth as evidence that housing values serve as a wealth proxy. That closeness averages gaps of one sign at every lifecycle stage. Fisher and Johnson (2023) report 0.30 on 395 offspring aged 30 to 32 and 0.21 on the 2,853 pairs formed from the waves that record wealth alone.

Table 4: Parent-child wealth mobility under each wealth proxy

Age band	Pairs	Actual wealth	Gap against the actual slope: GB-proxy	PK-proxy	FJ-proxy
Panel A. Both generations ranked under the wealth proxy					
30–34	3,346	0.34	0.08*** (0.02)	−0.05* (0.03)	0.01 (0.02)
35–39	2,633	0.40	0.07*** (0.02)	−0.05* (0.03)	−0.03 (0.03)
40–44	2,606	0.39	0.03* (0.02)	−0.06** (0.03)	−0.04* (0.03)
45–49	1,539	0.40	0.09*** (0.02)	0.04 (0.03)	0.04 (0.03)
50–54	1,024	0.46	0.03 (0.03)	0.00 (0.04)	−0.01 (0.04)
55–59	1,333	0.45	0.02 (0.03)	−0.11*** (0.04)	−0.08** (0.03)
60–64	1,009	0.49	0.04 (0.03)	−0.08* (0.04)	−0.07* (0.04)
Panel B. Parent ranked under the wealth proxy, child under actual wealth					
30–34	3,346	0.34	0.00 (0.01)	−0.04** (0.02)	−0.02 (0.02)
35–39	2,633	0.40	0.01 (0.01)	−0.03 (0.02)	−0.01 (0.02)
40–44	2,606	0.39	0.02 (0.02)	−0.04* (0.02)	−0.02 (0.02)
45–49	1,539	0.40	0.03* (0.02)	−0.02 (0.02)	−0.04 (0.02)
50–54	1,024	0.46	−0.01 (0.02)	−0.06* (0.03)	−0.07* (0.04)
55–59	1,333	0.45	0.00 (0.02)	−0.07** (0.03)	−0.04 (0.02)
60–64	1,009	0.49	0.01 (0.02)	−0.07** (0.03)	−0.06** (0.02)

Notes: The first column is the slope of a child's within-cohort wealth rank on the parent's at the same lifecycle stage, under actual wealth, on the pairs counted in the second column. Each individual's wealth is their share of their household's, the same share under every series (Appendix A). The remaining columns are each wealth proxy's gap, from one stacked regression of both rank series with the parent rank interacted with a wealth proxy indicator. Standard errors, in parentheses, are clustered on the 1968 family. Parents are linked through the Family Identification Mapping System. A child with two linked parents contributes one row per parent. Stars mark gaps distinguishable from zero at the ten, five and one percent levels. Panel A substitutes the wealth proxy for both generations. Panel B substitutes it for the parent alone and takes the child on actual wealth, the design a study of the pre-1984 waves runs, on the seven lifecycle stages every series reaches. Every lifecycle-stage count above covers these seven lifecycle stages, every lifecycle stage actual wealth reaches.

overstates the slope at every lifecycle stage, and the PK-proxy's and the FJ-proxy's gaps change sign across lifecycle stages.

The GB-proxy's pooled gap is near zero because its errors on the two halves of the parent distribution offset. Below the median parent rank its gap is -0.05, and above the median 0.07. Both are distinguishable from zero (Table 5). The PK-proxy and the FJ-proxy understate wealth persistence below the median and are indistinguishable from actual wealth above it. On the wealth quintile transition matrix the GB-proxy keeps 45% of the richest wealth quintile's children in the richest wealth quintile against 42% under actual wealth (Supplement SA.D).

In the grandparent-grandchild comparison a wealth proxy adds pairs, and the GB-proxy reproduces the slope on the pairs with actual wealth (Table 6). I obtain two findings. First, a wealth proxy adds the pairs whose grandparent the PSID observes only before 1984. Actual wealth alone gives about 1,000 grandparent-grandchild pairs, and each of the three wealth proxies about 1,300. Second, on the pairs where the grandparent has an actual rank, the GB-proxy

Table 5: The parent-child slope below and above the median parent rank, under each series

	Rank-rank slope		Children who stay, %	
	Below the median	Above the median	Poorest wealth quintile	Richest wealth quintile
Actual wealth	0.42 (0.03)	0.37 (0.04)	34	42
GB-proxy	0.38 (0.03)	0.44 (0.04)	34	45
PK-proxy	0.36 (0.04)	0.35 (0.04)	32	41
FJ-proxy	0.37 (0.03)	0.38 (0.04)	33	41

Notes: The slope is the child's within-cohort wealth rank on the parent's, under the parent-only design of Table 4, Panel B: the parent ranked under the series named in the row and the child under actual wealth, the same parent-child pairs in every row. One piecewise-linear regression per series places a knot at rank 50 on the parent's rank under that series, pooled over the 7 lifecycle stages with a lifecycle-stage fixed effect. Standard errors, in parentheses, are clustered on the 1968 family. The last two columns are the share of the poorest wealth quintile's children whose own rank lies in the poorest wealth quintile, and the same for the richest, the parent's wealth quintile under the series named in the row and the child's under actual wealth. The PK-proxy is Pfeffer and Killewald's (2018).

and the FJ-proxy reproduce the slope: their gaps are 0.01 and -0.01, both indistinguishable from zero. The PK-proxy lowers it: its gap is -0.04, and the difference between its gap and the GB-proxy's is distinguishable from zero.

Robustness These findings are robust to five tests. First, a Holm correction for testing the seven gaps at once leaves no individual gap in Panel B distinguishable from zero (Holm, 1979). Clustering on the parent instead of the 1968 family leaves the significance of every gap unchanged. Second, a bootstrap over the 1968 families places the GB-proxy's parent-child gap between -0.02 and 0.02, on both sides of zero. Third, restricting the sample to individuals observed at 3 or more lifecycle stages leaves the GB-proxy's pooled gap unchanged at 0.01. Fourth, the GB-proxy's gap stays within 0.10 of zero, the equivalence band, under the PSID's longitudinal weight. Fifth, dropping each birth cohort of children in turn leaves it inside the equivalence band as well (Appendix B, Supplement SA.D).

3.4 Wealth inequality

Performance metric I measure wealth inequality by the Gini coefficient, wave by wave. The performance metric $\mathcal{D}_p$ is a wealth proxy's mean absolute distance from the actual Gini:

$$\mathcal{D}_p = \frac{1}{T}\sum_{t=1}^{T} \left| G_p(t) - G_a(t) \right| \tag{9}$$

where $G_p(t)$ is the Gini of wealth proxy p in wave t, $G_a(t)$ is the Gini of actual wealth in the same wave and T is the number of waves. I set negative wealth to zero before computing any Gini, under actual wealth and under every wealth proxy alike. I also report the wealth shares of the top 10%, the middle 50 to 90% and the bottom 50%, averaged over the same waves.

Table 6: **Grandparent-grandchild wealth mobility under each wealth proxy**

Grandparent series	Pairs	Slope [95% interval]	Gap against actual wealth (SE)
Actual wealth	1,007	0.24 [0.16, 0.31]	
GB-proxy	1,307	0.27 [0.20, 0.33]	0.01 (0.02)
PK-proxy	1,306	0.22 [0.14, 0.30]	-0.04 (0.03)
FJ-proxy	1,307	0.25 [0.17, 0.32]	-0.01 (0.02)

Notes: Each individual's rank is the mean of their within-cohort ranks over the waves in which they are aged 25 to 75. The grandchild and the parent are on actual wealth and the grandparent on the series named, on every wave the series reaches. Each individual's wealth is their share of their household's, the same share under every series (Appendix A). Slope is the grandchild's rank on the grandparent's, with no other regressor. Intervals are 95 percent normal intervals from standard errors clustered on the grandchild's 1968 family. The gap is the change in the slope when the series replaces actual wealth on the pairs where the grandparent has both, with its clustered standard error in parentheses. With the parent's rank added to both regressions, the gap is -0.02 (0.02) for the GB-proxy, -0.07 (0.03) for the PK-proxy and -0.05 (0.03) for the FJ-proxy. A design that ranks all three at one lifecycle stage reaches at most 9 pairs on actual wealth, and its estimates are in the replication files only.

Results Every wealth proxy understates wealth inequality, and the GB-proxy understates it least. I obtain two findings. First, on the 1984 to 2023 sample the GB-proxy lies closest to the actual Gini in every wave: its Gini distance is 0.06, against 0.16 for the PK-proxy and 0.17 for the FJ-proxy (Figure 3, panel b). Second, the GB-proxy understates wealth inequality before 1984 as well. The PSID records no actual wealth there, so the comparison is against the implied actual Gini: the harmonized Survey of Consumer Finances of Kuhn, Schularick and Steins (2020), the SCF+, less the coverage gap (Figure 3, panel a, Supplement SA.E). The GB-proxy's Gini lies below the implied actual Gini in every SCF+ window, by up to 0.13, against up to 0.09 on the 1984 to 2023 sample (Supplement SA.E).

The understatement of wealth inequality comes from the top of the wealth distribution. First, averaged over the 1984 to 2023 sample, the GB-proxy misses 15% of the top 10% wealth share, the PK-proxy 36% and the FJ-proxy 36% (Figure 3, panel d). Second, the wealth proxies place most of the missing wealth in the middle 50 to 90%. The middle's share is 0.35 under actual wealth, 0.41 under the GB-proxy and 0.52 and 0.51 under the PK-proxy and the FJ-proxy (Figure 3, panel f). The bottom 50% holds 0.01 under actual wealth and no more than 0.08 under any wealth proxy (Figure 3, panel h). Rebuilt in the Survey of Consumer Finances, which reaches the top of the wealth distribution, the PK-proxy leaves a Gini distance of 0.17, against 0.16 in the PSID (Supplement SA.E).

Robustness These findings are robust to six tests. First, four further wealth inequality measures give the same ordering as the Gini (Supplement SA.C). Second, subtracting employer retirement balances widens the PK-proxy's and the FJ-proxy's distances, moves the GB-proxy's by less than a hundredth and leaves the ordering unchanged (Appendix B). Third, dropping

Figure 3: Wealth inequality under each series, 1969 to 1984 and 1984 to 2023

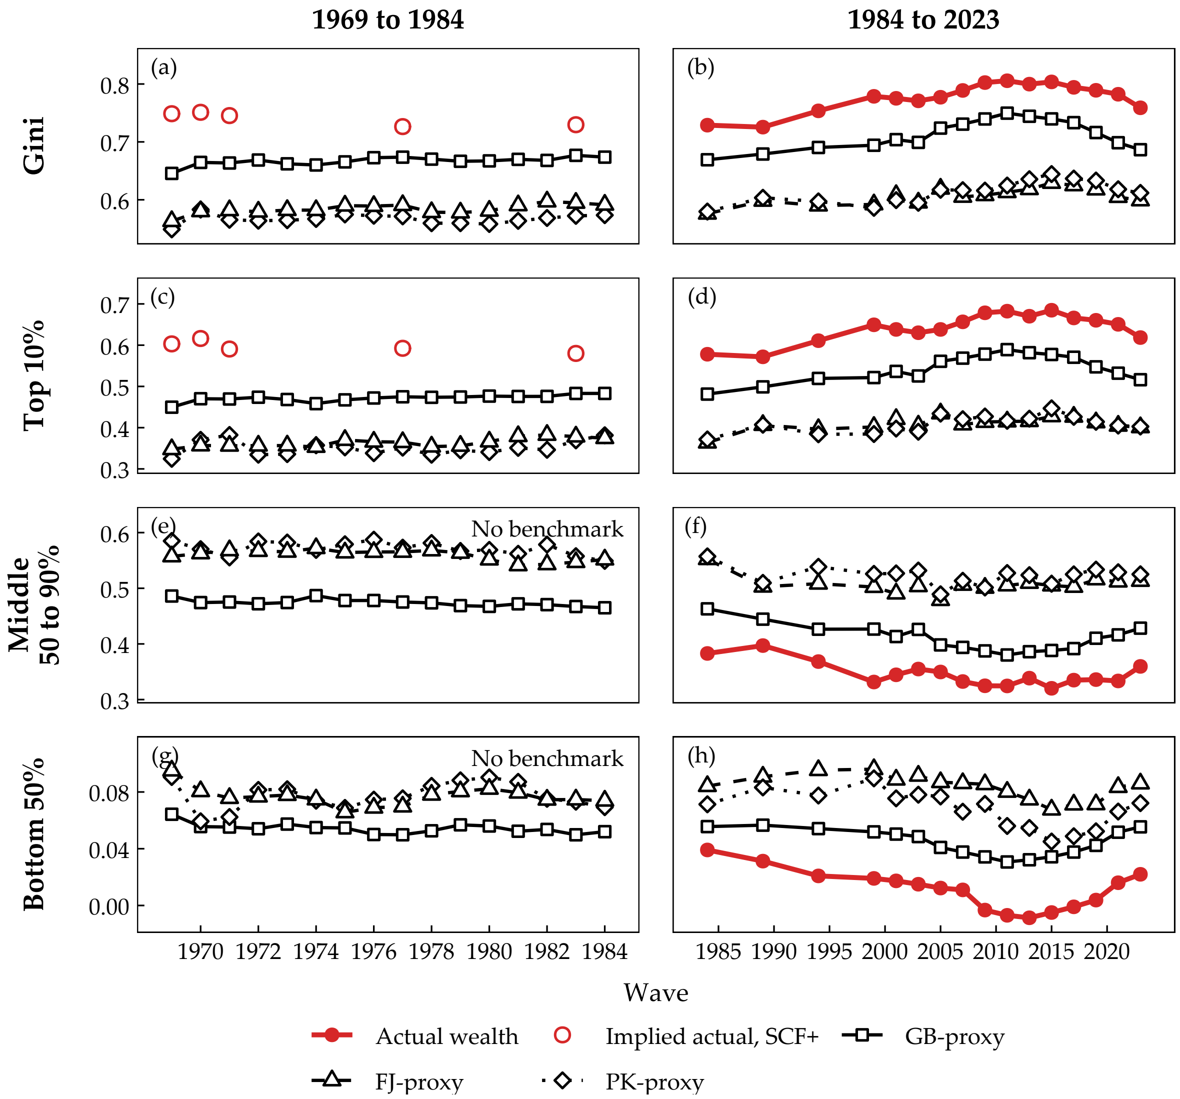

Notes: Actual wealth in the PSID is the red solid line with filled circles. The red open circles are the implied actual value from the SCF+ at the five waves the SCF+ reaches. That value is the SCF+ Gini, or top 10% wealth share, net of the coverage gap between the two surveys on that performance metric (Supplement SA.E). I measure that gap for the Gini and the top 10% share alone, so panels (e) and (g) show no benchmark. The black lines are the three wealth proxies, an open square on a solid line the GB-proxy, an open triangle on a dashed line the FJ-proxy and an open diamond on a dotted line the PK-proxy. The rows are the Gini of wealth, with negative wealth set to zero, and the shares of wealth held by the top 10%, the middle 50 to 90% and the bottom 50%, which sum to one. The columns are two periods on two samples. The left column, 1969 to 1984, is each series on the households it reaches. The right column is the 1984 to 2023 sample, the household-years with actual wealth and all three wealth proxies. The wave 1984 appears in both columns, once on each sample. The panel letters run row by row.

each wave in turn leaves the GB-proxy's Gini distance between 0.06 and 0.07. Fourth, a bootstrap over the 1968 families places it between 0.06 and 0.07, below the PK-proxy's and the FJ-proxy's point estimates (Supplement SA.D). Fifth, a Gini that keeps negative wealth instead of setting it to zero gives the same ordering. Sixth, restricting the sample to reference persons aged 25 to 64 gives the same ordering, with the PK-proxy and the FJ-proxy so close that two decimals do not separate them (Supplement SA.C).

3.5 Wealth accumulation trajectories

Performance metric I measure wealth accumulation as the ratio of median wealth at ages 60 to 64 to median wealth at ages 30 to 34, within a birth cohort:

$$\mathcal{R}_p = \frac{\mu_p^{\text{old}}}{\mu_p^{\text{young}}} \tag{10}$$

where $\mathcal{R}_p$ is that ratio for series p, and μ_p^{young} and μ_p^{old} are the medians of series p at the two ages. I compute the ratio for each of the three birth cohorts the 1984 to 2023 sample observes at both ages and report the median birth cohort.

Results The GB-proxy comes closest to actual wealth accumulation but overstates it. I obtain two findings. First, the GB-proxy overstates wealth accumulation by 22% on the median birth cohort. The PK-proxy understates it by a factor of 2.8 to 3.5 across the three birth cohorts. Median wealth rises 8.5 times between the two ages under actual wealth, 10.4 times under the GB-proxy, 2.4 under the PK-proxy and 2.8 under the FJ-proxy. Second, the overstatement comes from households observed at one age only, since households enter and leave the PSID. On the 504 households observed at both ages, median wealth rises 7 times under both actual wealth and the GB-proxy, against 2 under the PK-proxy. The GB-proxy therefore reproduces the wealth accumulation trajectory of households the PSID follows through working life.

Within the top 10% and the middle 50 to 90% the GB-proxy's rise comes nearest actual wealth's, and the PK-proxy's and the FJ-proxy's do not (Figure 4).[6] I split households into the bottom 50%, the middle 50 to 90% and the top 10% of each series' own ranking within birth cohort and age band. In the top 10%, median wealth rises 6.2 times under actual wealth and 5.3 times under the GB-proxy, and in the middle 50 to 90% 7.5 and 6.6 times. Within each group the GB-proxy therefore rises less than actual wealth, and within birth cohort across all households more: two biases of opposite sign.[7]

Why do the PK-proxy and the FJ-proxy understate wealth accumulation? Their input variable, the main residence, grows less than the rest of the balance sheet. Its gross value is 114% of the median household's wealth at ages 30 to 34 and 53% at ages 60 to 64. The first figure exceeds

[6] Figure 4 pools birth cohorts, so its age bands are different birth cohorts observed at one date. Its ratios therefore differ from the within-cohort ones (Supplement SA.G).

[7] The bottom 50% has no ratio, because its median at ages 30 to 34 is at or below zero under actual wealth, the PK-proxy and the FJ-proxy. At ages 60 to 64 its median is $85,000 under actual wealth, against $141,000 under the GB-proxy, $88,000 under the PK-proxy and $252,000 under the FJ-proxy.

Figure 4: Median wealth by age band, for every household and within three wealth groups

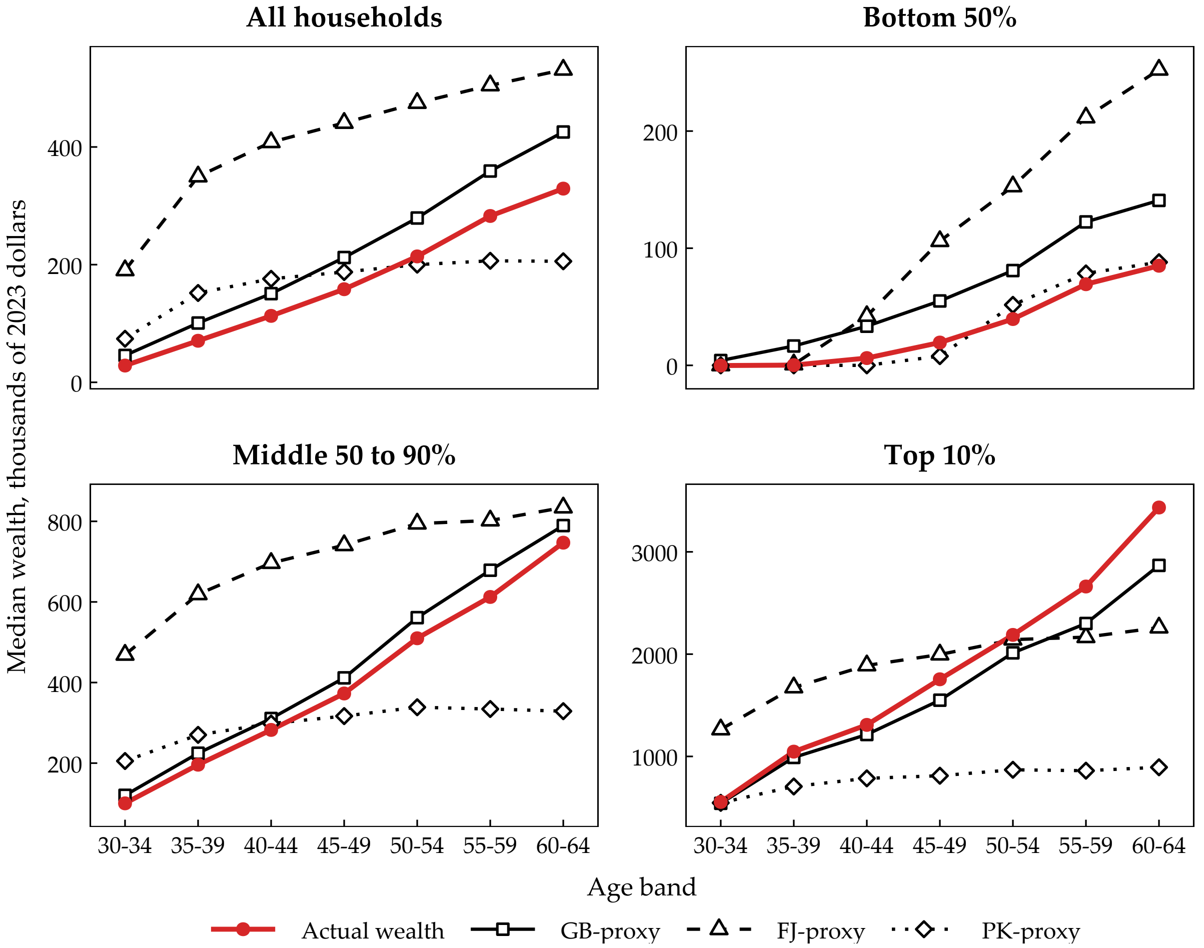

Notes: Each point is the median wealth in a five-year age band, in thousands of 2023 dollars, on the 1984 to 2023 sample and on the same households under each series. The four panels are every household, the bottom 50%, the middle 50 to 90% and the top 10%, each on its own vertical scale. Each series ranks households within birth cohort and age band, so a household in the GB-proxy's top 10% need not be in actual wealth's. In the top 10% and the middle 50 to 90% the GB-proxy's line rises less steeply than actual wealth's. The all-household panel pools birth cohorts and so omits the within-cohort overstatement of Section 3.5 (Supplement SA.G). Actual wealth is the red filled circles on a solid line, the GB-proxy the open squares on a solid line, the FJ-proxy the open triangles on a dashed line and the PK-proxy the open diamonds on a dotted line.

one hundred because of the mortgage. A wealth proxy built on the housing value therefore misses the rise in the rest of the balance sheet.

Robustness Two of seven tests reverse which wealth proxy comes closest to actual wealth accumulation: the PSID's longitudinal weight and a wealth concept without retirement balances. The other five, the personal consumption expenditure deflator, wealth per household member, each wave dropped in turn, each of the three birth cohorts and the 35 to 39 and 55 to 59 age bands, leave the GB-proxy closest (Table 9, Table 10, Supplement SA.C). On the same rows, the PSID's longitudinal weight raises the PK-proxy's ratio from 2.4 to 9.1 times. It moves actual wealth's only from 8.5 to 9.6 and the GB-proxy's from 10.4 to 10.8, so the PK-proxy then lies nearer actual wealth (Table 9, Appendix B, Supplement SA.B). A wealth concept without retirement balances puts actual wealth's rise nearer the PK-proxy's than the GB-proxy's. Employer balances alone do not.

4 The sources of the GB-proxy's performance

The GB-proxy's outperformance comes from its wider set of input variables and not from its estimator. Its own error concentrates among the poorest households and renters (Section 3).

4.1 The outperformance

The GB-proxy differs from a housing wealth proxy in two ways. It takes 15 or 14 input variables where a housing wealth proxy takes one. It fits them with a gradient-boosting model rather than a linear regression. Four models, each built and tuned as the GB-proxy is, separate what the input variables add from what the estimator adds. Model LA is a linear regression on one input variable, the housing value for homeowners and the rent for renters. Model LF is a linear regression on all 15 or 14 input variables. Models BA and BF are gradient-boosting models on the same two sets. Model BF therefore has the GB-proxy's input variables and estimator. I fit all four models out of fold. Each household falls in one of five folds, and its prediction comes from a model fitted on the other four.

A model's distance $\mathcal{D}$ is its gap from actual wealth on the performance metric of each dimension. The GB-proxy's improvement over model LA is $\mathcal{D}_{\text{LA}} - \mathcal{D}_{\text{BF}}$. The Shapley decomposition gives each source a value ϕ, the average of what the source adds when the other is absent and what it adds when the other is present:

$$\begin{aligned} \phi_{\text{input variables}} &= \tfrac{1}{2}\left[(\mathcal{D}_{\text{LA}} - \mathcal{D}_{\text{LF}}) + (\mathcal{D}_{\text{BA}} - \mathcal{D}_{\text{BF}})\right], \\ \phi_{\text{estimator}} &= \tfrac{1}{2}\left[(\mathcal{D}_{\text{LA}} - \mathcal{D}_{\text{BA}}) + (\mathcal{D}_{\text{LF}} - \mathcal{D}_{\text{BF}})\right]. \end{aligned} \tag{11}$$

The two values sum to the improvement, which is what makes them shares of it. Table 7 reports each term as a share of the improvement.

Most of the GB-proxy's improvement comes from its set of input variables and not from its estimator (Table 7), with one exception. First, on the four dimensions where it outperforms, the set of input variables accounts for 84 to 97% of the improvement. The housing value or rent alone leaves a wealth rank error of 20 units. All the input variables in a linear regression bring it to 16 and the GB-proxy, refitted on five folds, to 15. The Gini distance falls the same way,

Table 7: The GB-proxy's distance from actual wealth on each dimension, split between its set of input variables and its estimator

	One input variable		All input variables			Share of the improvement, %	
Dimension, and its distance	Linear	Boosted	Linear	Boosted	Improvement	Input variables	Estimator
Overall wealth rank distribution (rank units)	20	20	16	15	5	93	7
Intragenerational wealth mobility (ten-year slope)	0.04	0.03	0.02	0.06	-0.02	17	83
Intergenerational wealth mobility (parent-child slope)	0.04	0.04	0.02	0.01	0.03	94	6
Wealth inequality (Gini)	0.16	0.16	0.09	0.06	0.10	84	16
Wealth accumulation trajectories (growth ratio)	5	5	2	1	4	97	3

Notes: Each cell is a model's distance from actual wealth on the dimension's performance metric, named in the row and stated in the unit in brackets: the wealth rank error in rank units (Equation 5), the absolute gap on the ten-year rank-rank slope (Equation 6), the mean absolute gap over the five lifecycle stages of Panel B (Equation 7), the Gini distance (Equation 9) and the absolute gap on the cohort growth ratio (Equation 10). The columns cross one input variable against all of them, and a linear fit against a boosted one. The four models differ in nothing else: each takes the GB-proxy's construction and tuning, fitted out of fold on five folds blocked by household. The single input variable is the housing value for a homeowner and the rent for a renter. The improvement is the first column less the fourth, and a negative value means the GB-proxy's construction lies further from actual wealth than the linear model on one input variable. The last pair is Equation 11, and the two shares of the improvement sum to 100. Every figure is computed on the 66,297 household-years of the 1984 to 2023 sample that all four models score.

from 0.16 to 0.09 to 0.06. Second, the estimator moves the GB-proxy away from actual wealth on intragenerational wealth mobility. Model BF's ten-year slope gap is 0.06, against 0.04 on one input variable and 0.02 on all. The estimator accounts for 83% of the gap's growth from 0.04 on one input variable to 0.06. A linear regression on all of them would therefore lie closer there and further on the other four dimensions.

4.2 Performance across the wealth distribution

The GB-proxy's wealth rank error is largest in the poorest wealth quintile and among renters because a rank unit spans the fewest dollars at the bottom and renters are poorer. I obtain four findings, out of fold on the 71,952 household-years that record actual wealth, with every household ranked against its whole wave (Table 8).

First, a rank unit spans $5 thousand in the poorest wealth quintile and $158 thousand in the richest. The poorest wealth quintile's wealth rank error is therefore 16 units against 8, although its median dollar error is $25 thousand against $423 thousand. Second, within the poorest wealth quintile the error falls on the 84% of households with negative wealth, at 17 units against 10. Adding the PSID's non-mortgage debt as an input variable lowers theirs to

Table 8: The GB-proxy's wealth rank error by quintile of actual wealth and by household group

	All households			Wealth rank error by household group		
Wealth quintile	Wealth rank error	Dollar error	Dollars per rank	Homeowners	Renters	Neither
1, poorest	16	25	5	33	12	14
2	12	11	2	8	13	13
3	10	50	6	8	18	15
4	10	132	16	9	25	22
5, richest	8	423	158	7	29	25
All	11	62	—	9	14	15

Notes: The GB-proxy's construction, out of fold on five folds blocked by household, every household ranked against its whole wave. Wealth rank error is the mean absolute error in percentile-rank units. Dollar error is the median absolute error and dollars per rank the span of the wealth quintile's central sixteen rank units over sixteen, both in thousands of 2023 dollars, the second averaged over waves. Quintiles are of actual wealth within wave. Neither is the households reporting neither a housing value nor rent. The 71,952 household-years of the 16 waves that record actual wealth.

13. Third, renters' error exceeds homeowners', 14 units against 9, only because 47% of renters are in the poorest wealth quintile against 6% of homeowners. Fourth, the housing value and the mortgage ratio, which only the homeowner model takes, account for 43% of that gap. Without them the homeowner error rises by 2 units.

The renter-homeowner gap in wealth rank error comes from renters' place in the wealth distribution, not from the renter model. I split the gap into a composition part and a part within wealth quintiles. With π_q^g the share of group g in wealth quintile q and e_q^g its wealth rank error there, summing over the five wealth quintiles:

$$e^R - e^H = \sum_q \left(\pi_q^R - \pi_q^H\right) e_q^H + \sum_q \pi_q^R \left(e_q^R - e_q^H\right), \tag{12}$$

where e^R and e^H are renters' and homeowners' overall wealth rank errors. The first term is the composition part, the gap renters would have with homeowners' error in each wealth quintile. The second is the gap that remains within wealth quintiles. With homeowners' error in each wealth quintile, renters' error would exceed homeowners' by 10 units, the first term of Equation 12. The observed gap is 5 units because the second term is negative. Renters rank better than homeowners in the poorest wealth quintile, 12 units against 33, and 47% of renters are there. In the other four wealth quintiles they rank worse.

5 Conclusion

The GB-proxy takes all sixteen characteristics the PSID records as input variables. It outperforms the PK-proxy and the FJ-proxy on four of the five dimensions: the overall wealth rank distribution, intergenerational wealth mobility, wealth inequality and wealth accumulation trajectories. It cuts the wealth rank error from 21 and 19 units to 14, and the Gini distance from 0.16 and 0.17 to 0.06. Within a birth cohort, median wealth rises 8.5 times in actual wealth and 10.4 times under the GB-proxy, where the PK-proxy and the FJ-proxy give 2.4 and 2.8. On intergenerational wealth mobility the parent-child slope is indistinguishable from actual wealth's at all seven lifecycle stages.

The GB-proxy leads in every wealth quintile. The wider set of input variables accounts for 84 to 97% of its outperformance on those four dimensions, and the estimator for the rest. The GB-proxy overstates intragenerational wealth persistence, at a ten-year rank-rank slope of 0.76 against 0.70 in actual wealth. It understates the Gini by 0.06 on average and overstates the within-cohort rise. A study built on the series therefore knows on which side of the actual figure its own estimate lies.

The GB-proxy's own error is largest in the poorest wealth quintile, where a rank unit spans $5 thousand against $158 thousand in the richest. Renters' error exceeds homeowners', 14 rank units against 9. The renter model is not the cause: 47% of renters are in the poorest wealth quintile, against 6% of homeowners.

Rebuilt in the Survey of Consumer Finances, which reaches the top of the wealth distribution, the PK-proxy leaves a Gini distance of 0.17, against 0.16 in the PSID. Pfeffer and Killewald (2018) propose theirs for the 1940 Census and Zillow, sources that record the housing value and no other asset. Wherever a dataset records income, household structure and need beside the housing value, a wealth proxy fitted on all of them will likely rank households better.

References

Binder, Ariel J., Max Risch, and John Voorheis. 2025. "Housing Capital and Intergenerational Mobility in the United States." U.S. Census Bureau, Center for Economic Studies Working Paper 25-55.

Bloise, Francesco, Paolo Brunori, and Patrizio Piraino. 2021. "Estimating Intergenerational Income Mobility on Sub-optimal Data: A Machine Learning Approach." *The Journal of Economic Inequality*, 19(4): 643–665.

Boserup, S., W. Kopczuk, and C. Kreiner. 2018. "Born with a Silver Spoon? Danish Evidence on Wealth Inequality in Childhood." *The Economic Journal*, 128(612): F514–F544.

Cooper, Daniel, Karen Dynan, and Hannah Rhodenhiser. 2019. "Measuring Household

Wealth in the Panel Study of Income Dynamics: The Role of Retirement Assets." Federal Reserve Bank of Boston Research Department Working Paper 19-6.

Cornwall, Gary, and Marina Gindelsky. 2025. "Nowcasting Distributional National Accounts for the United States: A Machine Learning Approach." *AEA Papers and Proceedings*, 115: 79–84.

Filmer, Deon, and Lant H. Pritchett. 2001. "Estimating Wealth Effects without Expenditure Data—or Tears: An Application to Educational Enrollments in States of India." *Demography*, 38(1): 115–132.

Fisher, Jonathan, and David S. Johnson. 2023. "Intergenerational Mobility using Income, Consumption, and Wealth." Washington Center for Equitable Growth Working Paper.

Friedman, Jerome H. 2001. "Greedy Function Approximation: A Gradient Boosting Machine." *The Annals of Statistics*, 29(5): 1189–1232.

Holm, Sture. 1979. "A Simple Sequentially Rejective Multiple Test Procedure." *Scandinavian Journal of Statistics*, 6(2): 65–70.

Hurst, Erik, Ming Ching Luoh, and Frank P. Stafford. 1998. "The Wealth Dynamics of American Families, 1984–94." *Brookings Papers on Economic Activity*, 1998(1): 267–337.

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

Kalsi, Priti, and Zachary Ward. 2025. "The Gilded Age and Beyond: The Persistence of Elite Wealth in American History." National Bureau of Economic Research Working Paper 33355.

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.

Pfeffer, Fabian T., and Alexandra Killewald. 2018. "Generations of Advantage. Multigenerational Correlations in Family Wealth." *Social Forces*, 96(4): 1411–1442.

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.

Sahn, David E., and David Stifel. 2003. "Exploring Alternative Measures of Welfare in the Absence of Expenditure Data." *Review of Income and Wealth*, 49(4): 463–489.

Salas-Rojo, Pedro, and Juan Gabriel Rodríguez. 2022. "Inheritances and Wealth Inequality: A Machine Learning Approach." *The Journal of Economic Inequality*, 20(1): 27–51.

Shiro, Ariel Gelrud, Christopher Pulliam, John Sabelhaus, and Ember Smith. 2022. "Stuck on the Ladder: Intragenerational Wealth Mobility in the United States." Brookings Institution Report.

Si, Yajuan, Steve Heeringa, David Johnson, Roderick J. A. Little, Wenshuo Liu, Fabian Pfeffer, and Trivellore Raghunathan. 2023. "Multiple Imputation with Massive Data: An Application to the Panel Study of Income Dynamics." *Journal of Survey Statistics and Methodology*, 11(1): 260–283.

Solon, Gary. 1992. "Intergenerational Income Mobility in the United States." *American Economic Review*, 82(3): 393–408.

Van Langenhove, Christophe. 2026*a*. "Wealth Mobility in the United States: Empirical Evidence from the PSID." UGent, Department of Economics Working Paper WP25/1104.

Van Langenhove, Christophe. 2026*b*. "Who Reaches the Top? Intergenerational Wealth Transfers, Wealth Inequality and Wealth Mobility in the United States."

A Wealth proxy construction

This appendix sets out the PK-proxy's renter ordering, the GB-proxy's estimator, the candidate input variables it rejects, its out-of-sample performance, and the split of household wealth between individuals.

A.1 The PK-proxy's renter ordering

Because Equation 2 gives every renter zero wealth, the PK-proxy needs a rule for how renters rank among themselves. This paper puts renters in random order. Two others exist: all renters tied at one rank, as in Binder, Risch and Voorheis (2025), or renters ordered by their rent capitalized at the rental yield of Jordà et al. (2019). Pfeffer and Killewald (2018) set renters' housing value to zero and state no such rule. The rule changes no ordering in this paper. The PK-proxy's wealth rank error is 21 rank units with renters in random order, 20 with renters tied and 20 with renters ordered by rent, against 14 for the GB-proxy. The rule does move the PK-proxy's ten-year rank-rank slope: 0.66 with renters in random order and 0.71 with renters tied, against 0.70 in actual wealth (Section 3.2).

A.2 The GB-proxy's estimator

The GB-proxy is three gradient-boosting models (Friedman, 2001), one per group of households, because the housing value exists only for homeowners and the rent only for renters. The homeowner model adds the housing value and the mortgage ratio to the thirteen shared characteristics, while the renter model adds the rent and the model for households reporting neither has the thirteen shared characteristics by themselves. All three fit on the 16 waves that record actual wealth and predict all 42 waves from 1969.

Hyperparameters The homeowner model runs 140 rounds at depth 9 with a learning rate of 4.5 percent, the other two 90 rounds at depth 6 and 6.0 percent. A search over 336 settings per group on five household-blocked folds lowers held-out target error by at most 1%. In every group that is less than one standard error across folds, the bar for a change I set before searching. At the chosen number of rounds, held-out wealth rank error lies 0.0%, 0.2% and 1.1% above the minimum of the learning curve in the three groups. Moving depth and learning rate one step in either direction spans 0.48 rank units of wealth rank error (Supplement SA.H).

Trend Only one group's target trends over the 1984 to 2023 sample by more than 10% of its mean with a p-value below five percent, the two bars I set beforehand. That group is the households reporting neither a housing value nor rent. Extended back to 1969, a linear trend moves their target by -100% of its mean ($p = 0.01$), so their pre-1984 predictions drift.

Two forms Every performance metric uses the point series. A conditional mean is less dispersed than the wealth it predicts, so the point series is a lower bound on actual wealth inequality in every wave (Supplement SA.C). The 50 draws per household-year, calibrated to wealth persistence under actual wealth, put that dispersion back: their Gini distance is 0.01 against 0.06 for the point series (Supplement SA.C). Before 1984 the draws use the 1984 to 2023 fit alone, since the PSID records actual wealth only from 1984.

A.3 Candidate input variables

An input variable enters the GB-proxy only if it improves the backward split and reaches the pre-1984 waves. Education, industry and the type of home improve the backward split but reach only 39% of the pre-1984 household-years. The PSID records education only from 1975, and industry and the type of home reach 71% and 80% of those household-years (Supplement SA.A). National aggregates reach every wave but take one value per wave. Real disposable income per head is lower in every pre-1984 wave than in any 1984 to 2023 wave. Every pre-1984 household therefore falls on the same side of every split on it, and takes the same contribution.

A.4 Out-of-sample performance

The errors in the main text are in sample, since the GB-proxy and the FJ-proxy fit on the 1984 to 2023 household-years I measure their errors on. On households held out of the fit the GB-proxy's wealth rank error is 13.89 rank units over 65,635 held-out household-years, and 13.90 against the FJ-proxy's 18.29 on the rows the FJ-proxy also covers. Fitted on the waves from 1999 alone, a model predicts 1984, 1989 and 1994 at 13.34 rank units, against 13.60 when fitted on the waves from 2009, both windows fixed in advance. The held-out errors by household group in Section 3.1 rank homeowners among homeowners and renters among renters. Those errors are not comparable with the pooled error in Table 2, which ranks every household against every other.

A.5 The individual split

The intergenerational wealth mobility comparison runs on individuals and splits each household's wealth between reference person and partner by marital status. A single reference person keeps the whole value, a married one half, and a cohabiting one their share of household labour income. The equal split for married couples follows Van Langenhove (2026*b*) on this panel and the Danish registers, which record half of each jointly owned asset to each spouse (Boserup, Kopczuk and Kreiner, 2018). The same rule splits actual wealth and the three wealth proxies. The cohabiting rule is my own and covers 9% of individual-years. An equal split for cohabiting couples moves the GB-proxy's parent-child gap by less than a hundredth from 0.00. Both rules' series appear in Supplement SA.C.

B Robustness of the results

Tables 9 and 10 hold the three headline performance metrics under every perturbation of the sample and the wealth concept, and the wealth accumulation ratio within each birth cohort. The paragraphs below say what each test is for and add what the tables do not print. Two perturbations reverse the ordering on wealth accumulation: the PSID's longitudinal weight and a wealth concept without retirement balances from 1999.

The overall wealth rank distribution The GB-proxy ranks first under the PSID's longitudinal weight and under a reweighting for attrition. The same holds in a panel balanced over the whole 1984 to 2023 sample, which keeps 127 households, and in one balanced over half the waves, which keeps 3,048 (Table 9). Its Spearman rank correlation with actual wealth is 0.89, against 0.78 for the FJ-proxy and 0.74 for the PK-proxy. The PK-proxy comes closest in the poorest wealth quintile with renters tied. There a bootstrap on the 1968 families puts the share of households the GB-proxy misplaces by more than 25 rank units between 15% and 19%, and the difference from the PK-proxy's share between -1 and 2 percentage points, an interval that includes zero.

Intragenerational wealth mobility The GB-proxy overstates ten-year wealth persistence in every variant of the design. The gap is 0.07 once the 1984 and 1989 initial waves enter, 0.09 on each household's first pair alone and 0.06 at ages 30 to 54, against 0.07 in the main text. Under the PSID's longitudinal weight it is 0.06, and with households ranked within wave instead of within birth cohort 0.04 (Supplement SA.D). Pairing a household on its reference person alone admits 60% of the households at risk, against 52% under the same-partner requirement, and gives 0.07. The overstatement sits below the median: below the median initial rank the gap is 0.29 and above it -0.14.

Intergenerational wealth mobility The joint test of the seven parent-only gaps gives a p-value of 0.62 for the GB-proxy and 0.06 for the FJ-proxy. Tested separately, 7 of the 28 gaps, one per series and lifecycle stage, have a p-value below five percent. None does after a Holm adjustment for the number of tests (Holm, 1979). Under the PSID's longitudinal weight the GB-proxy's gap is 0.01, with no lifecycle-stage gap above 0.03 in absolute value. Dropping each of the eleven birth cohorts of children in turn keeps it between 0.00 and 0.01. Its grandparent-grandchild gap is 0.00 at ages 25 to 64 and 0.01 on individuals observed in three or more waves, against 0.01 in Section 3.3, with a bootstrap interval from -0.04 to 0.05 (Supplement SA.D).

Wealth inequality Every measurement choice leaves the Gini ordering of the three wealth proxies unchanged, with the GB-proxy's distance the lowest. Reweighting for attrition leaves the GB-proxy's distance at 0.07 against 0.07 unweighted on those rows. Subtracting employer retirement balances from 1999 takes the PK-proxy's distance from 0.16 to 0.17, the FJ-proxy's from 0.17 to 0.18, and moves the GB-proxy's by less than a hundredth, to 0.07. Four further performance metrics give the same ordering as the Gini. So do a Gini that keeps negative wealth and the 25 to 64 age window the SCF+ comparison uses (Supplement SA.C).

The FJ-proxy ranks households better than the PK-proxy and measures wealth inequality worse, at a Gini distance of 0.17 against 0.16. A bootstrap over the 1968 families puts that difference between 0.00 and 0.01. The PK-proxy's distance is the wealth proxy's and not the PSID's: in the Survey of Consumer Finances it is 0.17 against 0.16 here, across the five years both surveys cover. The SCF places the actual Gini at 0.83 against 0.78 in the PSID. That difference, the coverage gap, is why no Gini level in this appendix is a population estimate (Supplement SA.E).

Wealth accumulation trajectories On the household-years that have a longitudinal weight, weighting moves actual wealth's ratio from 8.5 to 9.6 times, the GB-proxy's from 10.4 to 10.8 and the PK-proxy's from 2.4 to 9.1. The 1984 weight takes 70 distinct values across 3,486 households. Dropping that wave gives 11.1, 11.7 and 9.7 on the remaining two birth cohorts. The personal consumption expenditure deflator puts actual wealth's ratio at 9.7 against 8.5 under the consumer price index. Per household member, actual wealth's ratio is 13.3. Dropping each wave in turn keeps the GB-proxy's ratio between 8.6 and 10.8 times and the PK-proxy's between 2.2 and 3.0. On the 35 to 39 and 55 to 59 age bands the ratios are 3.1, 3.2 and 1.5 over five birth cohorts. Each of the three birth cohorts puts the GB-proxy closest to actual wealth (Table 10).

Table 9: The three performance metrics under every perturbation

	Wealth rank error		Gini distance		Cohort growth		
	GB-proxy	PK-proxy	GB-proxy	PK-proxy	Actual	GB-proxy	PK-proxy
Manuscript	13.8	20.9	0.06	0.16	8.5	10.4	2.4
Survey weighted	13.4	20.7	0.07	0.15	9.6	10.8	9.1
the same rows unweighted	13.7	20.9	0.06	0.16	8.5	10.4	2.4
the 1984 weight flattened	13.4	20.6	0.07	0.15	9.6	10.8	9.1
1984 dropped, weighted	13.5	20.7	0.07	0.15	11.1	11.7	9.7
the same rows unweighted	13.8	21.0	0.07	0.16	9.3	10.5	2.8
1984 and 1989 dropped, weighted	13.6	20.7	0.07	0.16	12.7	10.8	10.2
the same rows unweighted	13.9	21.0	0.07	0.17	10.1	10.7	3.2
Attrition weighted	13.3	21.1	0.07	0.17	8.0	8.3	2.1
the same rows unweighted	13.2	21.1	0.07	0.18	8.0	8.0	2.0
Present in every wave	15.0	24.5	0.09	0.20	—	—	—
Present in at least half	12.6	20.7	0.07	0.19	6.7	7.8	2.0
Reference person aged 25 to 64	14.7	21.5	0.07	0.17	8.5	10.4	2.4
Employer retirement balances subtracted from 1999	14.1	21.2	0.07	0.17	7.0	10.4	2.4
without retirement balances from 1999	14.5	21.3	0.08	0.18	5.4	10.4	2.4
the same rows, the PSID's own variable	13.8	20.9	0.06	0.16	8.5	10.4	2.4
Waves before 1999	12.9	20.6	0.06	0.14	—	—	—
Waves from 1999	14.0	21.0	0.07	0.17	—	—	—
PCE deflator	13.8	20.9	0.06	0.16	9.7	11.8	2.8
Per household member	13.8	20.9	0.06	0.16	13.3	15.2	4.9
Ages 35 to 39 against 55 to 59	13.8	20.9	0.06	0.16	3.1	3.2	1.5
Leave one wave out, range	[13.6, 13.9]	[20.8, 21.1]	[0.06, 0.07]	[0.16, 0.17]	[7.6, 9.4]	[8.6, 10.8]	[2.2, 3.0]

Notes: Wealth rank error is the root mean squared percentile-rank error against actual wealth, averaged over waves with equal weight. The Gini distance is the mean absolute deviation of the wealth proxy's Gini from the actual Gini across waves. Cohort growth is the median ratio of median real wealth at ages 60 to 64 to that at ages 30 to 34 within a five-year birth cohort, or of ages 55 to 59 to 35 to 39 in the row that names them. The Actual column is that ratio under actual wealth. The PK-proxy is Pfeffer and Killewald's (2018). Each weighted row is followed by the same performance metric on the same rows unweighted, and each concept row by the PSID's own variable on the same rows. Flattening the 1984 weight gives every 1984 household that wave's mean weight. A dash means the variant leaves no birth cohort observed at both ends of the working life.

Table 10: Wealth accumulation within each birth cohort, under each series

Birth cohort	Actual wealth	Gradient boosting	PK, housing	FJ, capital income
1950 to 1954	7	7	2	2
1955 to 1959	9	10	2	3
1960 to 1964	10	11	4	3

Notes: Each cell is median real wealth at ages 60 to 64 over median real wealth at ages 30 to 34 within the birth cohort, on the 1984 to 2023 sample, under the series named in the column. A cohort enters when every series has at least 100 household-years in both age bands. The PK-proxy is Pfeffer and Killewald's (2018).