Online Supplement

Wealth Proxies in the PSID: A Wealth Series from 1969

Christophe Van Langenhove

September 2026

Contents

SA.A The information set

Adding input variables takes a model of the housing value alone from 21 rank units of out-of-fold error to 16 for a linear model on all ten, and the gradient-boosting estimator on the same ten takes it only to 15 (Table SA.1). Section 4.1 makes the same split on each dimension, on the GB-proxy's own construction.

SA.A.1 The comparison and its rule

I count a set of input variables as an improvement only when its wealth rank error falls by more than the 0.07 rank units the baseline moves across its own seeds. At least two of the Gini distance, the within-cohort growth ratio and the parent-child rank-rank gap must improve as well, and the backward wave split must not worsen. The comparison covers 45 sets at three seeds each, out of fold, on the two-model construction the GB-proxy replaced, with its ten input variables and separate models for homeowners and renters.

A candidate must reach 1969 and must vary within a wave, and both constraints act before any fitting. A wealth proxy scores only rows complete on every input variable, so a variable missing before 1984 empties the waves it exists to fill. A variable the PSID records once per wave gives a tree nothing to split on within that wave.

SA.A.2 The coverage constraint

The coverage constraint disqualifies 16 of the 45 sets whatever they score, education, industry and the type of home among them. Those three improve the backward split, but together they reach 39% of the pre-1984 household-years against 100% for the baseline ten, because the PSID records education only from 1975 and reaches 71% and 80% of those household-years with industry and the type of home. Anchoring a household to its own first observed wealth fails the same constraint, because only 61% of pre-1984 household-years belong to a household the PSID still records in 1984.

SA.A.3 Aggregates and the training range

An aggregate indexed by wave helps before 1984 only where its 1969 to 1983 values fall inside its 1984 to 2023 range, because a tree returns its boundary value outside its training range. The aggregate ratio of household wealth to household real estate runs from 2.52 to 3.54 over the 1984 to 2023 sample and from 2.97 to 3.87 before 1984, so 67% of the 15 pre-1984 waves lie inside its training range. Real disposable income per head lies inside in 0% of those waves,

Table SA.1: What the GB-proxy gains from more input variables

Set	What it adds	Added	Coverage	Wealth rank error	Gini distance	Backward split
A0	baseline	0	100%	14.96 =	0.07 =	15.03 =
A1	household characteristics	10	98%	14.72 −	0.06 −	15.11 =
A2	raw-file characteristics	3	100%	14.96 =	0.07 =	15.00 =
A3	national aggregates	7	100%	14.90 =	0.07 +	15.00 =
A4	geography	2	100%	14.94 =	0.07 +	14.97 =
A5	tenure length	1	100%	14.94 =	0.07 +	14.92 =
A6	everything at full coverage	23	98%	14.63 −	0.06 −	15.11 =
A7	short-coverage additions	3	39%	15.40 +	0.07 −	14.90 =
A8	everything, coverage ignored	26	38%	15.14 +	0.06 −	14.63 −
A9	every candidate at exactly full coverage	18	100%	14.72 −	0.06 −	15.01 =
C1	linear, same information set	0	100%	16.08 +	0.09 +	15.51 +
C2	linear, extended information set	10	98%	15.91 +	0.09 +	15.35 +
C3	housing value alone	0	100%	20.63 +	0.15 +	21.79 +
C4	labour income and scale	0	100%	19.55 +	0.13 +	20.45 +
EP	earner share and head's labour income	2	84%	15.35 +	0.06 −	14.90 =
ES	the reference person's earner share	1	84%	15.38 +	0.07 −	15.06 =
L1	robust loss, deployed input variables	0	100%	14.74 −	0.05 −	15.09 =
N1	calendar year, negative control	1	100%	14.89 =	0.07 +	14.97 =
PA	head's labour income and race	2	100%	14.87 −	0.07 =	14.86 =
PB	head's labour income and lump-sum receipts	2	100%	14.86 −	0.07 =	14.78 =
PC	race and lump-sum receipts	2	100%	14.89 =	0.07 =	14.93 =
PD	labincome rp, race and lump-sum receipts	3	100%	14.85 −	0.07 −	14.77 =
VA	plus age of the youngest child	1	100%	14.91 =	0.07 +	14.84 =
VB	plus the wave factor	1	100%	14.88 −	0.07 +	14.81 =
VC	plus financial to real-estate wealth, national	1	100%	14.89 =	0.07 +	14.74 =
VD	plus housing share of wealth, national	1	100%	14.87 −	0.07 +	14.84 =
VE	plus household leverage, national	1	100%	14.89 =	0.07 +	14.82 =
VF	plus children	1	100%	14.91 =	0.07 +	14.81 =
VG	plus composition change	1	100%	14.91 =	0.07 +	14.82 =
VH	plus family income	1	100%	14.83 −	0.07 −	14.97 =
VI	plus household type	1	100%	14.92 =	0.07 +	14.79 =
VJ	plus head's labour income	1	100%	14.93 =	0.07 =	14.87 =
VK	plus partner's labour income	1	100%	14.97 =	0.07 =	14.98 =
VL	plus long interest rate	1	100%	14.84 −	0.07 +	14.84 =
VM	plus adults	1	100%	14.91 =	0.07 +	14.85 =
VN	plus ownership rate, national	1	100%	14.89 =	0.07 +	14.97 =
VO	plus race	1	100%	14.93 =	0.07 =	14.96 =
VP	plus head's sex	1	100%	14.93 =	0.07 +	14.86 =
VQ	plus waves in the home	1	100%	14.89 =	0.07 +	14.72 =
VR	plus unemployment rate	1	100%	14.90 =	0.07 +	14.79 =
VS	plus food stamps	1	100%	14.88 −	0.07 +	14.82 =
VT	plus lump-sum receipts	1	100%	14.92 =	0.07 =	14.93 =
VU	plus state	1	100%	14.88 −	0.07 +	14.78 =
VV	plus state income per person, relative	1	100%	14.91 =	0.07 +	14.82 =
VW	plus rooms	1	98%	14.90 =	0.07 =	14.82 =

Notes: Every figure is out of fold on folds blocked by household, averaged over three seeds. Coverage is the share of pre-1984 household-years with every variable the set requires, and a set below the baseline's coverage is disqualified whatever it scores. A minus sign marks a movement better than the baseline by more than its across-seed range, a plus sign one worse by more than that range, and an equals sign one inside it. The backward split fits the waves from 1999 and predicts 1984, 1989 and 1994.

and a tree given it returns one number for every pre-1984 wave. I include it and the calendar year in the comparison to show that limit.

Table SA.2: The construction variants: three performance metrics, the GB-proxy's input variables, mean over seeds

Tag	Construction	Gini distance	Within-cohort growth	Parent-only gap
v2r	the construction version 4a replaced, refit	0.04	8.4	0.01
v4a	GB-proxy: trailing window, real dollars, no equity input variable	0.04	10.4	0.01
v4b	trailing window alone	0.04	8.5	0.01
v4c	real dollars alone	0.04	9.4	0.01
v4f	equity input variable dropped alone	0.04	8.5	0.01
v4d	trailing window, ranked dollars, no equity input variable	0.04	10.4	0.01
v4e	the GB-proxy's construction less the four wealth accumulation input variables	0.04	10.4	0.01
v5a	trailing window, relative dollars, no equity input variable	0.04	10.4	0.01
v5b	as v5a, with wealth per household as the homeowner carrier	0.05	10.4	0.01
v5c	relative dollars alone	0.04	10.1	0.01
v5d	trailing window, national-aggregate dollars, no equity input variable	0.04	9.8	0.01

Notes: Every variant is refit on the consistent wealth concept and residual, with three seeds and shared folds, trim and tuning. On these rows the actual within-cohort growth ratio is 6.7 and the refitted two-model baseline's 9.5, and the parent-only gap is the mean over the five common lifecycle stages. The Construction column says what each tag changes.

SA.A.4 The wealth accumulation input variables

Dropping the four wealth accumulation input variables from the GB-proxy's own construction leaves the parent-only gap at 0.01 against 0.01 with them, the within-cohort growth ratio at 10 times against 10, and the Gini distance at 0.04 against 0.04 (Table SA.2). At the precision the paper prints, the four move none of the three. The table gives all eleven variants of that construction.

SA.A.5 The prediction outside the sample

A 1983 prediction scored against the same households' actual wealth in 1984 misses by 10.8 rank units, against 15.8 for the PK-proxy, and most of its excess over the 1984 to 2023 sample's own 9.7 is the year between the prediction and the measurement. The test exists because the PSID records wealth from 1984, so the 3,017 households active in 1983 and observed with actual wealth in 1984 have a prediction from outside the 1984 to 2023 sample and a measure-

ment inside it. I score the 1983 prediction on the mean absolute wealth rank error within those households, unweighted, and I refit the construction on the 15 waves without 1984, because I fitted the GB-proxy's own model on all 16 wealth waves.

The refit records 10.8 rank units against 10.6 for the fitted model, a contamination of 0.2 units, and misplaces 9% of households by more than a quarter of the distribution. Its excess over the 1984 to 2023 sample is 1.1 rank units. The same model predicts 1984 on the same households at 9.5, so the one-year gap costs 1.3 rank units and a 1983 figure scored against 1984 is an upper bound on the prediction error.

SA.B The sample

None of the 16 waves of the 1984 to 2023 sample produces the GB-proxy's lead on its own. Dropping each in turn leaves it with the lowest wealth rank error in all 16 of the samples that remain, between 13.6 and 13.9 units against 20.8 to 21.1 for the PK-proxy (Table 9). Its Gini distance stays between 0.06 and 0.07, and the within-cohort growth ratio moves most, from 8.6 to 10.8.

SA.B.1 The survey weight

The PSID's longitudinal weight leaves the ordering of the wealth proxies on the overall wealth rank distribution and on wealth inequality where it is, and moves the PK-proxy's within-cohort growth ratio further than any other figure, from 2.4 to 9.1 times (Table 9). The main text applies no weights, because its Survey Research Center sample excludes the low-income oversample and the immigrant samples added from 1997, so an unweighted Gini here is not an estimate of the population Gini. I compare every weighted figure with the unweighted one on the same rows, the 67,227 of the 67,665 matched household-years, or 99%, with a positive weight. The GB-proxy's Gini distance stays at 0.07 and the FJ-proxy's moves from 0.17 to 0.17, and the ordering holds.

The rise in the PK-proxy's ratio survives every treatment of the 1984 weight, which takes few distinct values, so the weight itself produces it. It takes only 70 distinct values across 3,486 households, in the PSID's own file as well, against 3,862 in 1994, whose weight comes from the family file with 7,120 distinct values. Giving every 1984 household that wave's mean weight leaves the ratio at 9.1. Dropping 1984 gives 9.7 over two birth cohorts, and dropping 1989 as well gives 10 over one, because a birth cohort reaches ages 30 to 34 in one wave alone.

SA.B.2 Balanced panels

A panel balanced over the whole 1984 to 2023 sample keeps too few households to reach a birth cohort at both ends of the working life, 127 households and 2032 household-years. On it the GB-proxy's Gini distance is 0.09 against 0.20 for the PK-proxy (Table 9). Requiring presence in at least 8 waves keeps 3,048 households and 33,625 household-years, on which the GB-proxy's wealth rank error falls to 13 units and its Gini distance is 0.07.

SA.B.3 Attrition

Households that leave the panel are not a random draw from those that stay (Fitzgerald, Gottschalk and Moffitt, 1998). Reweighting for that leaves the GB-proxy's Gini distance at 0.07 against 0.07 unweighted, on the same 46,260 rows (Table 9). Each household takes one over its predicted chance of reappearing in the following wave, from a wave-by-wave model on actual wealth, the housing value and the wave. I trim the weights at the first and ninety-ninth percentiles so that no household dominates a wave. The PK-proxy's distance moves from 0.18 to 0.17.

SA.B.4 The low-income oversample

On the low-income oversample alone the GB-proxy ranks households worse than the PK-proxy with renters tied (Table SA.3). Every figure in this paper excludes that oversample, the Survey of Economic Opportunity sample of low-income households the PSID drew in 1968, and the GB-proxy's three models never saw those households. I rebuild the processing tier with the oversample in it and apply the models without refitting, which reaches 40,075 household-years of 11,247 oversample households over the 16 waves that record actual wealth. On the 71,952 household-years of the Survey Research Center (SRC) sample the two object trees share, the models return the GB-proxy series exactly.

On the oversample alone the GB-proxy records 25 rank units, against 22 for the PK-proxy with renters tied and 25 for the FJ-proxy, and a Gini distance of 0.07 against 0.06 and 0.11 (Table SA.3). Ranked together with the SRC sample, which is 64% of the pool, it records 17 against 20 and 20. The oversample is poorer, with 24% of its household-years holding negative wealth against 17% of the SRC sample's, and the GB-proxy's errors are largest among the poorest (Section 3.1). The certified verdicts rest on the SRC sample for that reason (Supplement SA.E.7). The wealth mobility dimensions are not scored here, because the rebuild does not produce the parent-child linkage.

Table SA.3: The wealth proxies on the low-income oversample

	Household-years	Wealth rank error	Beyond 25 units, %	Gini distance
The low-income oversample				
GB-proxy	40,075	25	28	0.07
PK-proxy, renters tied	40,075	22	22	0.06
FJ-proxy	40,075	25	24	0.11
The SRC sample				
GB-proxy	71,952	14	7	0.06
PK-proxy, renters tied	71,952	20	17	0.15
FJ-proxy	71,951	19	16	0.17
Both together				
GB-proxy	112,027	17	13	0.06
PK-proxy, renters tied	112,027	20	14	0.13
FJ-proxy	112,026	20	17	0.15

Notes: The low-income oversample is the Survey of Economic Opportunity sample the PSID drew in 1968, which every other figure in this paper excludes. The GB-proxy's three models are fitted on the SRC sample and applied here without refitting, and each block ranks its own households against each other. Every figure is unweighted. The PK-proxy takes renters tied at zero, and the FJ-proxy the PSID's own capital-income residual, because the consistent residual of Appendix A is built for the SRC sample alone.

SA.C Concept, measure and scale

SA.C.1 The wealth concept

Neither treatment of the retirement balances reorders the wealth proxies on wealth inequality, and subtracting the employer balances from 1999 moves the GB-proxy's Gini distance by less than a hundredth, to 0.07, and the PK-proxy's from 0.16 to 0.17 on the same rows (Table 9). Removing the whole retirement item instead, individual balances included, moves each distance by a further hundredth, to 0.08 and 0.18. I keep the PSID's own variable, which holds retirement balances in every wave, and Section 2 discloses the change of question at 1999 that I do not adjust for.

SA.C.2 The capital-income input variable

With an intercept, or on the PSID's asset-income item, the FJ-proxy lies closer to the actual Gini than the PK-proxy does, so its place between the other two depends on how I fit it. I draw the contrast between ranking and scale from the GB-proxy and the PK-proxy. Fitted with an intercept of $-91,000 in Equation 3, the FJ-proxy records 19 rank units and a Gini distance of 0.10, inside the PK-proxy's on both. Refitted on the PSID's asset-income item over the 66,919 household-years that item reaches, it records 18 rank units against the PK-proxy's 21, and a Gini distance of 0.14 against 0.17 on the residual and the PK-proxy's 0.16.

SA.C.3 The break at 1999

The 1999 change in how the PSID asks about retirement balances leaves no step in actual wealth and no change in the ordering of the wealth proxies. The balances the later waves itemize are 8% of wealth on average, and the Gini moves 0.03 from 1994 to 1999, against 0.03 over the preceding five years. Split at 1999, the GB-proxy's Gini distance is 0.06 before and 0.07 after, against 0.14 and 0.17 for the PK-proxy.

SA.C.4 Measures beside the Gini

The top one, five and ten percent wealth shares and the ratio of the ninetieth percentile to the median give the Gini's ordering, on the same rows and with the same bottom-coding. On an actual top one percent share of 0.25, the GB-proxy misses by 0.08 and the PK-proxy by 0.15, and the other three measures place the two the same way round.

Table SA.4: The GB-proxy's Gini under the point form and the draw form, wave by wave

Wave	Actual	Point form	Draw form	Draw form against actual
1984	0.73	0.67	0.75	above
1989	0.73	0.68	0.75	above
1994	0.75	0.69	0.77	above
1999	0.78	0.69	0.77	below
2001	0.78	0.70	0.78	above
2003	0.77	0.70	0.77	above
2005	0.78	0.72	0.79	above
2007	0.79	0.73	0.79	above
2009	0.80	0.74	0.81	above
2011	0.81	0.75	0.82	above
2013	0.80	0.74	0.81	above
2015	0.80	0.74	0.81	above
2017	0.79	0.73	0.80	above
2019	0.79	0.72	0.79	below
2021	0.78	0.70	0.77	below
2023	0.76	0.69	0.76	below

Notes: The Gini of wealth, bottom-coded at zero, on the rows of the 1984 to 2023 sample the implicates reach. The draw form is the mean over implicates of each implicate's Gini.

Keeping the negative tail in the Gini, or restricting the rows to the SCF+ comparison's age window, leaves the GB-proxy closest to actual wealth. A Gini that keeps the negative tail, which 16% of actual household-years hold, puts actual wealth at 0.80 against 0.78 bottom-coded, and the distances at 0.09 for the GB-proxy and 0.19 and 0.19 for the PK-proxy and the FJ-proxy. On the 51,795 household-years with a reference person aged 25 to 64, the window of that comparison, the distances are 0.07, 0.17 and 0.17 (Supplement SA.E). I do not use the ratio of the seventy-fifth percentile to the twenty-fifth, because 33% of a housing wealth proxy's household-years lie within a dollar of zero.

SA.C.5 The draw form

The GB-proxy's draw form, the point prediction plus a draw of the part the model leaves unexplained, lies closer to actual wealth inequality than the point form every performance metric above uses. Its Gini distance is 0.01 against 0.06, and its distance from the top 10% wealth share 0.01 against 0.09 (Table SA.4). The point form lies below actual wealth in all 16 waves, so a wealth inequality figure computed from it is a lower bound. The draw form lies above in twelve waves and below in four. The draws sample a residual distribution fitted to actual wealth, which the PSID records only from 1984, so their closeness after 1984 partly reflects that fit.

SA.C.6 Which way a construction misses

A construction that capitalizes income flows stretches the distribution, and one that scales a reported asset value compresses it. Fagereng et al. (2016) capitalize asset-class income at mean returns in Norwegian administrative data, and their implied Gini lies above the actual one in all twenty years, by 0.03 on average. Smith, Zidar and Zwick (2023) put the top interest rate at around 3.5 times the equal-returns rate, which lowers the capitalized top one percent share from 38.9 to 33.7 percent. The GB-proxy misses that share by 0.08 of total wealth, more than the correction.

SA.C.7 The deflator and household size

The personal consumption expenditure deflator puts the within-cohort growth ratio at 9.7 for actual wealth, 11.8 for the GB-proxy and 2.8 for the PK-proxy, against 8.5, 10.4 and 2.4 under the consumer price index. Wealth per household member gives 13.3, 15 and 5, because household size falls across the working life as children leave, and under both the GB-proxy stays closest to actual wealth.

SA.C.8 Ranking against scale in wealth accumulation

To separate the ranking a wealth proxy gives from the scale it gives, I map each series onto the distribution of actual wealth within its wave:

$$\tilde{w}_{i,t} = \Phi_{a,t}^{-1}\left(\Phi_{p,t}(\hat{w}_{i,t})\right) \tag{SA.1}$$

where $\tilde{w}_{i,t}$ is the mapped series, $\hat{w}_{i,t}$ is the wealth proxy's prediction of Section 2, $\Phi_{p,t}$ is the distribution of wealth proxy p in wave t and $\Phi_{a,t}$ is the distribution of actual wealth. The mapped value keeps the wealth proxy's ranking and takes actual wealth's units. On its ranking alone the GB-proxy gives 141% of the actual rise over the working life, the PK-proxy 28% and the FJ-proxy 40%, so the GB-proxy's ranking overstates the rise and its scale corrects part of that.

SA.C.9 The individual split

How a couple's wealth is split between its members changes 2% of individual-years and no parent-only gap at the precision the paper prints. The manuscript gives a married reference person half of household wealth and a cohabiting one their share of household labour income (Appendix A). Splitting every couple equally instead leaves the gap, averaged over the seven

lifecycle stages, at 0.00 for the GB-proxy against 0.00 and -0.06 for the PK-proxy against -0.06, and no lifecycle-stage gap of the GB-proxy lies further than 0.03 from zero, inside the equivalence band of 0.10.

SA.D Inference and multiplicity

Tested one at a time, 7 of the 28 parent-only gaps, one per series and lifecycle stage, have a p-value below five percent (Table SA.5). None does after a Holm or a Benjamini-Hochberg adjustment for the number of tests (Holm, 1979; Benjamini and Hochberg, 1995). Under the design that replaces both generations, 7 of 28 gaps have a p-value below five percent, 3 after the Holm adjustment and 3 after Benjamini-Hochberg.

SA.D.1 The joint test

The seven lifecycle-stage samples nest inside one another, so their gaps covary and a sum of squared t-ratios overstates the joint statistic. I therefore test the seven gaps jointly in the inverse of their covariance:

$$W = \hat{\gamma}' \, \hat{V}^{-1} \, \hat{\gamma} \tag{SA.2}$$

where W is the joint statistic, $\hat{\gamma}$ stacks the seven estimated gaps γ_s of Equation 7, and $\hat{V}$ is their covariance from one stacked regression across the lifecycle stages, clustered on the 1968 family. The correction leaves the GB-proxy's Wald statistic at 5 and lowers the FJ-proxy's from 18 to 14, whose p-value of 0.06 no longer rejects at five percent. A covariance from 500 resamples of the 1,078 1968 families changes no joint verdict. It moves the GB-proxy's p-value most, from 0.62 to 0.60.

SA.D.2 The gap under a moved sample

Neither a balanced set of individuals nor an attrition weight moves the GB-proxy's gap from 0.01, and both move the PK-proxy's by at most a hundredth from -0.05. Restricting the sample to the 8,446 individuals observed at 3 or more lifecycle stages leaves the GB-proxy's gap at 0.01

Table SA.5: The joint test and the multiplicity adjustment under the parent-only design

	Naive		Clustered on the 1968 family		Bootstrap	Lifecycle stages	Holm
	Wald	p	Wald	p	p	in band	survivors
GB-proxy	5	0.62	5	0.62	0.60	7 of 7	0 of 7
PK-proxy	25	0.00	17	0.01	0.02	4 of 7	0 of 7
Housing, rent-ordered	24	0.00	17	0.02	0.02	4 of 7	0 of 7
FJ-proxy	18	0.01	14	0.06	0.07	5 of 7	0 of 7

Notes: The naive statistic sums the squared t-ratios of the lifecycle-stage gaps as if independent, and the clustered statistic is the quadratic form of the same gaps in the inverse of their covariance, clustered on the 1968 family. The bootstrap p uses the covariance of the gap vector over replicates that resample 1968 families. Lifecycle stages in band counts the lifecycle stages whose 95 percent interval lies inside the equivalence band of 0.10, and Holm survivors the gaps significant at five percent after a Holm adjustment within the wealth proxy's lifecycle stages.

and puts the PK-proxy's at -0.04. Weighting for attrition instead, from a model of who is still observed at the following lifecycle stage, gives 0.01 and -0.05.

SA.D.3 The gap and the prediction error

The GB-proxy's prediction error predicts the child's rank less than the PK-proxy's does, at a coefficient of 0.18 against 0.32. I regress the child's rank on the parent's predicted rank and on the parent's error rank, the actual rank less the predicted one:

$$\kappa_i^c = \alpha + \beta_{\text{pred}}\,\hat{\kappa}_i^p + \beta_{\text{error}}\left(\kappa_i^p - \hat{\kappa}_i^p\right) + \varepsilon_i \tag{SA.3}$$

where $\hat{\kappa}_i^p$ is the parent's predicted rank, the bracket is the parent's error rank and α is the intercept. The GB-proxy's error takes a coefficient of 0.18 against 0.46 on the prediction, and the PK-proxy's error 0.32 against 0.46.

The GB-proxy's near-zero gap survives the two checks that could have produced it from the fit alone. With the parent's thirteen input variables added to the child regression, the gap averages 0.04 and stays inside the equivalence band at every lifecycle stage, and with the parent's rank taken from household-blocked out-of-fold predictions the mean gap is 0.01. I set the equivalence band on the 1984 to 2023 sample, where the wealth rank error is lower than on the pre-1984 mix (Section 3.1), so the gaps may leave it in the waves the series serves.

SA.D.4 The ten-year gap under moved samples

The GB-proxy overstates ten-year wealth persistence under every moved sample, by no less than 0.04. Dropping each initial wave in turn keeps its ten-year slope between 0.76 and 0.77. The gap of Section 3.2 is 0.07 once the 1984 and 1989 initial waves enter, 0.09 on each household's first pair alone and 0.06 for reference persons aged 30 to 54 at the initial wave. Under the PSID's longitudinal weight it is 0.06, and with households ranked within wave instead of within birth cohort 0.04.

SA.D.5 The gap across the parent distribution

The GB-proxy's gap of Section 3.3 is -0.05 below the median parent rank and 0.07 above it, each distinguishable from zero at five percent. I estimate one piecewise-linear regression per wealth proxy, with a knot at rank 50 of the parent's rank under that wealth proxy and a fixed effect for the lifecycle stage, on the same 13,490 pairs over seven lifecycle stages for every arm.

Table SA.6: Wealth quintile transition matrices, parent to child, under each series

Parent's wealth quintile	Child's wealth quintile under actual wealth, %				
	1 (poorest)	2	3	4	5 (richest)
Actual wealth					
1	34	29	18	12	7
2	24	25	23	16	12
3	16	19	23	23	19
4	13	13	20	26	28
5	10	9	15	23	42
GB-proxy					
1	34	30	19	11	7
2	24	25	23	17	11
3	18	20	23	20	19
4	14	12	20	28	27
5	11	8	14	22	45
PK-proxy					
1	32	28	19	13	9
2	23	28	22	16	11
3	18	17	24	22	19
4	13	13	20	26	28
5	13	9	14	23	41
FJ-proxy					
1	33	28	18	12	8
2	23	27	23	16	12
3	18	18	23	22	19
4	14	13	20	25	28
5	12	9	14	24	41

Notes: Each row is the distribution of the children of parents in that wealth quintile over the child's own wealth quintile, in percent, summing to one hundred. The parent's wealth quintile is under the series named in the block and the child's under actual wealth, on the parent-only design of Table 4, Panel B, pooled over the 7 lifecycle stages on the same 13,490 lifecycle-stage pairs in every block. The PK-proxy is Pfeffer and Killewald's (2018).

The PK-proxy's gaps are -0.06 and -0.02 and the FJ-proxy's -0.06 and 0.00, so both understate wealth persistence below the median only.

Splitting on the parent's actual rank instead would put the GB-proxy's gap at -0.08 below the median and -0.02 above, because that split selects on the prediction error, which is largest at the bottom (Section 3.1). The knot sits on each wealth proxy's own rank for that reason.

The transition matrix shows the same pattern in shares, with the GB-proxy keeping 34% of the poorest wealth quintile's children in the poorest wealth quintile and 45% of the richest wealth quintile's children in the richest (Table SA.6). Actual wealth keeps 34% and 42%, and the PK-proxy 32% and 41%.

SA.D.6 The three-generation gap under moved samples

The grandparent-grandchild gap of Section 3.3 is 0.01 for the GB-proxy. Narrowing the age band from 25 to 75 to 25 to 64 gives 0.00, and requiring three waves inside that band gives 0.01.

A bootstrap over the 1968 families places the gap between -0.04 and 0.05, and an equivalence test puts it inside the equivalence band of 0.10.

SA.D.7 The bootstrap

Resampling the 1,078 1968 families with replacement, 500 times, places the GB-proxy's gap between -0.02 and 0.02 at the lifecycle stage with the most pairs, and the PK-proxy's between -0.07 and 0.00. I resample 1968 families because a household that recomposes changes identifier and would otherwise enter as several independent draws.

A second resample, of the 1,923 1968 families in the 1984 to 2023 sample, 500 times, places the GB-proxy's wealth rank error between 13 and 14 rank units and its Gini distance between 0.06 and 0.07. Its outperformance on wealth rank error runs from 33% to 36% against the PK-proxy and from 26% to 29% against the FJ-proxy. Clustering on the parent instead of the 1968 family leaves every significance verdict unchanged, and Table SA.7 gives the standard errors under both units.

Table SA.7: Parent-only gaps under two clustering units

Wealth proxy	Age band	Gap	SE, parent id	SE, 1968 family	Ratio
GB-proxy	30–34	0.00	0.01	0.01	1.02
	35–39	0.01	0.01	0.01	1.02
	40–44	0.02	0.02	0.02	1.00
	45–49	0.03	0.02	0.02	0.99
	50–54	-0.01	0.02	0.02	1.00
	55–59	0.00	0.02	0.02	1.01
	60–64	0.01	0.02	0.02	1.00
PK-proxy	30–34	-0.04	0.02	0.02	1.05
	35–39	-0.03	0.02	0.02	0.98
	40–44	-0.04	0.02	0.02	1.00
	45–49	-0.02	0.03	0.02	1.00
	50–54	-0.06	0.03	0.03	1.00
	55–59	-0.07	0.03	0.03	1.01
	60–64	-0.07	0.03	0.03	0.99
Housing, rent-ordered	30–34	-0.04	0.02	0.02	1.04
	35–39	-0.02	0.02	0.02	0.98
	40–44	-0.04	0.02	0.02	1.00
	45–49	-0.02	0.02	0.02	1.00
	50–54	-0.06	0.03	0.03	1.00
	55–59	-0.07	0.03	0.03	1.01
	60–64	-0.07	0.03	0.03	0.99
FJ-proxy	30–34	-0.02	0.01	0.02	1.02
	35–39	-0.01	0.02	0.02	0.98
	40–44	-0.02	0.02	0.02	1.00
	45–49	-0.04	0.02	0.02	1.00
	50–54	-0.07	0.04	0.04	1.00
	55–59	-0.04	0.02	0.02	1.00
	60–64	-0.06	0.02	0.02	1.00

Notes: The same stacked regression on the same pairs, with the covariance clustered on the parent's individual id and on the child's 1968 family.

SA.E An external benchmark

Rebuilt in the Survey of Consumer Finances, which draws a list sample from tax records to reach the top, the PK-proxy leaves a Gini distance of 0.17 across the five years from 1989 to 2019 that both surveys cover, against 0.16 in the PSID. The distance is the wealth proxy's own, since it survives a survey that reaches the top. Panels without such a sample miss the top, and Juster and Kuester (1991) find that two of them understate wealth and wealth inequality against the 1983 SCF.

SA.E.1 The two constructions

I rebuild the PK-proxy in the SCF under both of its renter orderings, from the gross value of the primary residence the SCF records in the same interview as wealth. Renters take a positive value negligible against any homeowner's residence, in random order in the first construction and in the order of their rent in the second. I compute every performance metric inside each of the SCF's five implicates and average across them, as its documentation asks. The GB-proxy and the FJ-proxy have no SCF rebuild, because several of their thirteen input variables are household histories that a cross-section does not record.

SA.E.2 The level and the distance

The PK-proxy understates wealth inequality in the SCF as it does in the PSID, in five of the five SCF years, and the SCF's better reach into the top shows in the level of the Gini while the distance survives it. Wealth has a Gini of 0.83 in the SCF against 0.78 in the PSID, a coverage gap of 0.06. In the SCF 16% of households hold negative wealth and in the PSID 16%, and since every Gini here bottom-codes wealth at zero, that part of the level cancels from a distance. The PK-proxy ranks households better in the SCF, at 17 rank units against 21 in the PSID, because the SCF reaches further up the top, where the housing value and wealth move together most closely.

SA.E.3 The SCF+ before 1984

The SCF+ is the external distribution I score the pre-1984 series against, in three windows, 1971, 1977 and 1983. It harmonizes the historical Surveys of Consumer Finances with the modern ones (Kuhn, Schularick and Steins, 2020), records wealth in 1969, 1970, 1971, 1977 and 1983, and pools 1969 to 1971 into the 1971 window. Both sides use wealth less vehicles, on the survey weight, for heads aged 25 to 64. My import reproduces the Gini of the SCF summary extract in

four years within 0.00 and the authors' own window Ginis for 1950 to 2016 within 0.01. I apply the manuscript's deflator to nominal values, because the file's price index runs 9% above the CPI-U in 1969 to 1977.

SA.E.4 The coverage gap and the implied actual Gini

The SCF+ Gini runs 0.05 above the PSID's on the same concept, weight and age window, and that mean is what turns an SCF+ Gini into the implied actual Gini before 1984. Over the six overlap pairs the gap runs from 0.03 to 0.06, and the same construction gives the implied top 10% wealth share, whose coverage gap runs from 0.06 to 0.11, 0.08 on average. The tolerance is the 0.05 to 0.09 by which actual wealth's Gini lies above the GB-proxy's across the 16 waves of the 1984 to 2023 sample, widened by each window's replicate interval.

SA.E.5 The result

The GB-proxy is a looser lower bound on wealth inequality before 1984 than after it in the 1971 window alone, and the trend test passes (Table SA.8). It lies below the implied actual Gini in every window, by 0.13 in 1971 against a widened band whose upper edge is 0.11, and inside the band in 1977 and 1983. The distance narrows by 0.03 between 1971 and 1977, where the 1984 to 2023 sample admits 0.04. Across the variants of Table SA.9 the series is looser in none to two windows. I fixed this rule after a preliminary calculation had already produced an approximate result that failed the trend test.

The point series understates the rise in the median between the 1971 and 1983 windows, so its rise is a floor on the actual one. The SCF+ median for the same heads rises from $78,665 to $103,094, while the point series' median rises from $137,122 to $138,856. The draws come within 0.04 of the implied actual Gini, while the PK-proxy with renters tied and the FJ-proxy lie at least 0.19 and 0.17 below it.

SA.E.6 Selecting the construction

I chose the GB-proxy's construction on the 1977 and 1983 windows and held the 1971 window out, because on the 1984 to 2023 sample the candidates lie within a fraction of a rank unit of each other. I count a construction as holding in a window when its distance there, less its own mean distance on the 1984 to 2023 sample, falls inside the half-width of 0.02 with the distance still above zero. The GB-proxy's construction holds in 1983, at 0.01, and misses in 1977, at 0.03, and its level error over the two windows is 0.01 against a tolerance of 0.10. The level match

Table SA.8: The 1969 to 1983 series against the SCF+, by the SCF+'s three-year windows

Window	SCF+	5th–95th	Implied actual	GB-proxy	Draws	PK-proxy, renters tied	FJ-proxy	*D* (GB-proxy)	*D* on the 1984 actual
Panel A. Gini of wealth, bottom-coded									
1969–1971	0.80	0.77 to 0.82	0.75	0.62	0.72	0.51	0.52	0.13	
1977	0.77	0.75 to 0.78	0.73	0.63	0.72	0.51	0.53	0.10	
1983	0.78	0.76 to 0.79	0.73	0.65	0.73	0.54	0.56	0.08	0.07
Panel B. Top-decile share of wealth, bottom-coded									
1969–1971	0.69			0.43	0.56	0.30	0.32		
1977	0.67			0.43	0.55	0.30	0.32		
1983	0.66			0.46	0.58	0.33	0.35		
Panel C. Median wealth, 2023 dollars									
1969–1971	78,665			137,122	103,204	112,901	253,235		
1977	107,980			146,984	113,218	130,694	301,685		
1983	103,094			138,856	104,189	131,571	292,007		

Notes: SCF+ wealth less vehicles, with individual retirement accounts and employer accounts kept, on the survey weight, 65 rows without one dropped, heads aged 25 to 64 on both sides. The SCF+ unit is its three-year window (1969, 1970 and 1971 pooled). The PSID side pools the same calendar years with each wave's weights scaled to one total. The 5th to 95th percentile is the spread of the Gini over the authors' 999 replicate weights on the first implicate. Every other performance metric is computed inside each of five implicates and averaged. The implied actual Gini is the SCF+ Gini less the coverage gap of 0.05, the mean over six overlap pairs of the SCF+ Gini minus the PSID actual Gini on the same concept, weight and age window. *D* is the implied actual Gini minus the GB-proxy's Gini. The level test compares *D* with the 1984 to 2023 sample's distance range [0.05, 0.09], widened by the half-width of the window's interval. SCF+ values are returned to nominal by the file's own price index and deflated to 2023 dollars by the CPI-U. In 1983 the high-income sample holds 1.9 percent of the weight and raises the SCF+ Gini by 0.07, from 0.70. The last column scores the 1983 window against the PSID's 1984 actual Gini of 0.73 with no coverage-gap adjustment. The draws are the fifty implicates of the same model, averaged. Their outcome under the same rule is C-low. Their dispersion is calibrated on the target they are compared against, so they are not an independent check.

Table SA.9: The SCF+ comparison under alternative choices

Variant	*G*	*D*, 1969–1971	*D*, 1977	*D*, 1983	Windows failing	Outcome
headline	0.05	0.13	0.10	0.08	1	C-high
G from 1984 alone	0.05	0.13	0.10	0.07	1	C-high
G from 1989 alone	0.06	0.12	0.09	0.07	1	C-high
G from 1994 alone	0.04	0.14	0.11	0.09	2	C-high
G from 2001 alone	0.04	0.14	0.11	0.08	1	C-high
G from 2007 alone	0.03	0.15	0.12	0.09	2	C-high
G from 2013 alone	0.06	0.12	0.09	0.07	1	C-high
all ages	0.06	0.10	0.05	0.06	0	B
SCF+ wealth with vehicles	0.03	0.14	0.10	0.08	1	C-high
SCF+ without pensions (1983 on)	0.06	0.12	0.08	0.06	1	C-high
SCF+ on wgt, PSID weighted	0.05	0.10	0.07	0.08	0	B
SCF+ on wgtI95W95, PSID unweighted	0.05	0.12	0.08	0.07	1	C-high
1983 window against the observed 1984 actual	0.00			0.07	0	B
all households (unmatched actual)	0.04	0.14	0.11	0.09	1	C-high
G with the concept delta added	0.06	0.12	0.09	0.06	1	C-high
G from all 13 waves within one year of a window	0.05	0.13	0.10	0.07	1	C-high
single years 1969, 1970, 1971 (no interval)	0.05	0.13	0.10	0.08	2	C-high

Notes: Each row recomputes the coverage gap, the band where the variant changes the PSID side, and the three *D* values for the GB-proxy's point series under one changed choice, everything else as in the headline. Outcome A means the level and trend tests both pass. B means the level passes and the trend fails. C-high means *D* lies above the band in some window, and C-low that it lies below. The unweighted row uses the SCF+'s unadjusted weight, one value in the historical years named, and drops the PSID weight. The single-year row scores 1969, 1970 and 1971 separately, without an interval. Deflating by the file's own price index throughout, instead of returning to nominal first, puts the SCF+ medians at 72,241 and 98,067 dollars of 2023 in the 1971 and 1983 windows, against 78,665 and 103,094 under the manuscript's deflator. The Gini rows do not depend on the deflator.

Table SA.10: What the pre-1984 series supports, by performance metric

Performance metric	Verdict	What the evidence is
Where a household stands in its wave	bounded	14 rank units, the lowest of the five series in every wave and in four of five quintiles of actual wealth. The SCF+ disagrees on the share of non-homeowners above the median owner in three of three windows
The parent-child rank-rank gap	certified	inside the 0.10 band at seven of seven lifecycle stages, joint p 0.62
The grandparent-grandchild slope	certified	moves the slope by 0.01, and by -0.03 on the waves the actual rank can also use
How concentrated wealth is	bounded below	0.06 from the actual Gini, below it in every wave of the 1984 to 2023 sample and in every SCF+ window before 1984
How much wealth accumulates over a working life	bounded above	10.4 times against an actual 8.5 times, from two biases of opposite sign, in the same direction in every birth cohort the profile reaches
How many dollars a household holds	bounded	the SCF+ median moves 31% between the 1971 and 1983 windows where the series' median moves 1%

Notes: Certified means the series passes a test fixed in advance on the 1984 to 2023 sample and no external check stands against it. Bounded means the sign of the bias is known and steady while its size is not certified, below for a series under the benchmark on every comparison and above for one over it. Not supported means the series and the external benchmark disagree in sign. Each magnitude is the macro the text quotes, and the rank-rank row proxies the parent and takes the child from actual wealth, the design a study of the pre-1984 waves runs.

costs the wealth accumulation overshoot of Section 3.5, 10.4 times within a birth cohort against 8.5 in actual wealth.

SA.E.7 Certified and bounded claims

Table SA.10 certifies the GB-proxy's parent-child and grandparent-grandchild slopes for the pre-1984 waves and bounds its wealth rank error, its Gini, its wealth accumulation and its dollar level.[1] I call a row certified when it passes a test fixed in advance on the 1984 to 2023 sample and fails no external check, and bounded when the evidence fixes the direction of the series' miss. The parent-child gap lies inside the equivalence band of 0.10 at seven of seven lifecycle stages, a band twice the smallest gap those lifecycle stages could detect, and the grandparent-grandchild gap is 0.01.

Every verdict is a verdict about the Survey Research Center sample, the 1984 to 2023 sample and the households with a housing value or rent, and each is looser beyond them. On the low-income oversample the series does not rank best (Supplement SA.B). Households that report

[1] Van Langenhove (2026) uses the series on those terms. It measures intergenerational and intragenerational wealth mobility in the PSID over 1969 to 2023 on the GB-proxy's ranks before 1984 and on actual wealth from then, and states the GB-proxy's overstatement of wealth persistence for each analysis.

neither a housing value nor rent it ranks at 29 rank units against 15 for homeowners. The 1984 to 2023 sample is where its distances are the tightest they are anywhere. It also puts too few non-homeowners above the median homeowner, in all three SCF+ windows and by up to 8 percentage points, so a study should rank homeowners against homeowners and non-homeowners against non-homeowners.

SA.F Rounding in the reported housing value

The GB-proxy's largest input variable is a self-reported housing value that the PSID records on a grid of round numbers. Across the 29 annual waves a wave holds between 186 and 312 distinct values among several thousand homeowners, and 75% to 98% of them are multiples of a thousand dollars. Multiples of ten thousand cover 17% to 41% of values before 1984 and 17% to 53% across the annual waves, because the unit a household rounds to grows with the price of its house. On the 1984 to 2023 sample 42% to 70% of positive values are multiples of ten thousand dollars, and a wave holds 249 to 503 distinct values.

Every comparison in this paper is against a measure built on the same self-report, because actual wealth uses it for the residence, and the report has an error of its own. Homeowners overstate the value by 5.1% on average against sale prices in the American Housing Survey, and only the length of tenure predicts the error (Kiel and Zabel, 1999). In the PSID the gap between the reported value and a market estimate is largest for homeowners in negative equity (Choi and Painter, 2018). The bracket imputation of the main residence from 2005 likewise reaches the housing input variable and actual wealth in the same waves.

The tests below cover 53,389 pairs over the 28 annual pairs from 1969 to 1997. The rows are the Survey Research Center sample's homeowner-occupiers who report a positive value in both waves of a pair, with the raw report cleared of missing-value codes only, movers excluded and no weights.

SA.F.1 Household or PSID

The PSID's accuracy flag marks each value it assigns, from 1969 to 1993, which covers 24 of the 28 pairs on both legs. Of 56,445 homeowner-years through 1993, 5% have an assigned value, and among pairs that report the same number twice only 6% have either value assigned. Dropping the assigned values raises the repeat rate from 22.1% to 22.3%, so the repeats come from the households' own reports.

SA.F.2 Why one value in four repeats

23% of pairs report the same value twice, although a median homeowner's value moves 9% in a year, 6% to 16% across the annual pairs. Three tests attribute the repeats to rounding, and a household repeating an earlier year's value without revaluing its house would fail each of them. For a first-wave value that is not a multiple of a hundred dollars, the repeat rate falls to

2%. Across the 28 pairs, regressing the annual repeat rate on national real-estate growth gives a coefficient of -0.60 with a standard error of 0.11, because a price that moves faster crosses a grid boundary more often.

In a simulation where each household revalues by the wave's growth and reports to its grid unit, simulated repeats are 65% to 129% of observed across 6 combinations of three grid models and two growth rates. 2 of the combinations fall below the 80% line fixed in advance.

SA.F.3 What it costs this paper

Housing values that a household repeats from an earlier year without revaluing its house cost this paper at most 2% of housing wealth, under the 3% threshold fixed in advance. 19% of homeowner-years repeat a value reported in an earlier year, and the bound treats every one as unrevalued and prices it at the national appreciation accumulated since the value was first given. The pipeline's handling of sentinels, negative values and outliers moves the one-year repeat rate by 0.07 percentage points in total, of which the outlier smoothing accounts for 0.05.

SA.G The input variables, wave by wave

SA.G.1 Coverage before 1984

I can build every wealth proxy in each of the 15 waves between 1969 and the first wave that records wealth, but they reach different shares of a wave. The capital-income residual the FJ-proxy takes as an input variable has a value for every household. The housing wealth proxies reach 95% of the households interviewed in the median wave, because they cannot place a household that reports neither a housing value nor rent. The GB-proxy reaches 100% in every wave, because those households have a model of their own.

Over 1969 to 1983 the GB-proxy's real median falls from $103,052 to $98,522, and its Gini rises from 0.65 to 0.68 and its top 10% wealth share from 0.45 to 0.48. The draws put the Gini at 0.73 and 0.75. In the 1971 window the point series lies further below the implied actual Gini than it ever lies below actual wealth after 1984, so before 1984 it is a looser lower bound on wealth inequality (Supplement SA.E).

SA.G.2 Where the input variables come from

The rent field and the asset-income item change form across the 43 waves, and Tables SA.11 and SA.12 give the code, the unit and the form of each input variable and of the weight in every wave. In 1988 and 1989 the rent field is a homeowner indicator, so the pipeline interpolates the rent from the neighbouring waves before it corrects outliers. The PSID brackets the asset-income item up to 1975 and asks for an amount with a reporting period from 1994 to 2003.

SA.G.3 The top-code

The PSID caps the housing value at $99,999 in the six waves from 1969 to 1974, where up to 0.6% of positive homeowner-reported values lie at the cap, and in zero of the waves this paper compares. From 1975 the field takes six digits, and the import maps its ceiling of $999,999 to missing. 18 homeowner-reports lie at that ceiling, in 7 of its 36 waves and 2 of the 16 this paper compares, at most 0.08% of a wave's homeowners. The housing value's own top tail is cut in the earliest waves, and no performance metric on the rescaled series can detect it.

SA.G.4 Assignment inside actual wealth

The PSID assigns a value inside actual wealth for between 8% and 14% of the matched household-years in the eight waves its accuracy flag covers, and the item built from the hous-

Table SA.11: Where the four raw input variables come from, wave by wave: 1968 to 1992

Wave	Housing value	Rent	Asset income	Weight
1968	none	none	none	none
1969	V449	V455 (annual dollars)	V703 (bracket, nine intervals)	V1014 (pipeline import)
1970	V1122	V1128 (annual dollars)	V1203 (bracket, nine intervals)	V1609 (pipeline import)
1971	V1823	V1829 (annual dollars)	V1904 (bracket, nine intervals)	V2321 (pipeline import)
1972	V2423	V2429 (annual dollars)	V2505 (bracket, nine intervals)	V2968 (pipeline import)
1973	V3021	V3023 (annual dollars)	V3058 (bracket, nine intervals)	V3301 (pipeline import)
1974	V3417	V3419 (annual dollars)	V3470 (bracket, nine intervals)	V3721 (pipeline import)
1975	V3817	V3819 (annual dollars)	V3870 (bracket, nine intervals)	V4224 (pipeline import)
1976	V4318	V4326 (annual dollars)	V4384 (annual dollars, rent, dividends and interest)	V5099 (pipeline import)
1977	V5217	V5225 (annual dollars)	V5295 (annual dollars, rent, dividends and interest)	V5665 (pipeline import)
1978	V5717	V5723 (annual dollars)	V5794 (annual dollars, rent, dividends and interest)	V6212 (pipeline import)
1979	V6319	V6326 (annual dollars)	V6404 (annual dollars, rent, dividends and interest)	V6805 (pipeline import)
1980	V6917	V6925 (annual dollars)	V6994 (annual dollars, rent, dividends and interest)	V7451 (pipeline import)
1981	V7517	V7525 (annual dollars)	V7586 (annual dollars, rent, dividends and interest)	V8103 (pipeline import)
1982	V8217	V8221 (annual dollars)	V8279 (annual dollars, rent, dividends and interest)	V8727 (pipeline import)
1983	V8817	V8825 (annual dollars)	V8887 (annual dollars, rent, dividends and interest)	V9433 (pipeline import)
1984	V10018	V10026 (annual dollars)	V10272 (annual dollars, dividends and interest)	V11079 (pipeline import)
1985	V11125	V11133 (annual dollars)	V11414 (annual dollars, dividends and interest)	V12446 (pipeline import)
1986	V12524	V12532 (annual dollars)	V12813 (annual dollars, dividends and interest)	V13687 (pipeline import)
1987	V13724	V13732 (annual dollars)	V13915 (annual dollars, dividends and interest)	V14737 (pipeline import)
1988	V14824	V14824 (homeowner indicator, from the housing-value field)	V14930 (annual dollars, dividends and interest)	V16208 (pipeline import)
1989	V16324	V16324 (homeowner indicator, from the housing-value field)	V16430 (annual dollars, dividends and interest)	V17612 (pipeline import)
1990	V17724	V17733 (annual dollars)	V17846 (annual dollars, dividends and interest)	V18943 (pipeline import)
1991	V19024	V19033 (annual dollars)	V19146 (annual dollars, dividends and interest)	V20243 (pipeline import)
1992	V20324	V20333 (annual dollars)	V20446 (annual dollars, dividends and interest)	V21547 (pipeline import)

Table SA.12: Where the four raw input variables come from, wave by wave: 1993 to 2023

Wave	Housing value	Rent	Asset income	Weight
1993	V21610	V21620, V21621 (amount with a period, annualised)	V21829, V21845, V21861, V21877 (annual dollars per component)	V23361 (family file)
1994	ER2033	ER2049, ER2050 (amount with a period, annualised)	ER3203, ER3218, ER3233, ER3248 (amount and period per component)	ER4160 (family file)
1995	ER5032	ER5048, ER5049 (amount with a period, annualised)	ER6203, ER6218, ER6233, ER6248 (amount and period per component)	ER7000 (family file)
1996	ER7032	ER7121, ER7122 (amount with a period, annualised)	ER8320, ER8335, ER8350, ER8365 (amount and period per component)	ER9251 (family file)
1997	ER10036	ER10060, ER10061 (amount with a period, annualised)	ER11213, ER11228, ER11243, ER11258 (amount and period per component)	ER12084 (pipeline import)
1999	ER13041	ER13065, ER13066 (amount with a period, annualised)	ER14479, ER14494, ER14509, ER14524 (amount and period per component)	ER16518 (pipeline import)
2001	ER17044	ER17074, ER17075 (amount with a period, annualised)	ER18634, ER18650, ER18666, ER18682 (amount and period per component)	ER20394 (pipeline import)
2003	ER21043	ER21072, ER21073 (amount with a period, annualised)	ER22003, ER22020, ER22037, ER22054 (amount and period per component)	ER24179 (pipeline import)
2005	ER25029	ER25063, ER25064 (amount with a period, annualised)	ER27932, ER27934, ER27936, ER27938 (annual dollars per component)	ER28078 (pipeline import)
2007	ER36029	ER36065, ER36066 (amount with a period, annualised)	ER40922, ER40924, ER40926, ER40928 (annual dollars per component)	ER41069 (pipeline import)
2009	ER42030	ER42080, ER42081 (amount with a period, annualised)	ER46830, ER46832, ER46834, ER46836 (annual dollars per component)	ER47012 (pipeline import)
2011	ER47330	ER47387, ER47388 (amount with a period, annualised)	ER52238, ER52240, ER52242, ER52244 (annual dollars per component)	ER52436 (pipeline import)
2013	ER53030	ER53087, ER53088 (amount with a period, annualised)	ER58039, ER58041, ER58043, ER58045 (annual dollars per component)	ER58257 (pipeline import)
2015	ER60031	ER60088, ER60089 (amount with a period, annualised)	ER65217, ER65219, ER65221, ER65223 (annual dollars per component)	ER65492 (pipeline import)
2017	ER66031	ER66090, ER66091 (amount with a period, annualised)	ER71294, ER71296, ER71298, ER71300 (annual dollars per component)	ER71570 (pipeline import)
2019	ER72031	ER72090, ER72091 (amount with a period, annualised)	ER77316, ER77318, ER77320, ER77322 (annual dollars per component)	ER77631 (pipeline import)
2021	ER78032	ER78091, ER78092 (amount with a period, annualised)	ER81643, ER81645, ER81647, ER81649 (annual dollars per component)	ER81958 (pipeline import)
2023	ER82033	ER82092, ER82093 (amount with a period, annualised)	ER85497, ER85499, ER85501, ER85503 (annual dollars per component)	ER85812 (pipeline import)

Notes: Variable codes are those of the PSID cross-year index, and the housing value and the rent are the codes the pipeline's crosswalk maps. Where the rent code equals the housing-value code, the field is a homeowner indicator and the rent is interpolated from the neighbouring waves. Asset income is the reference person's item, with the partner's taken alongside, and the weight is the core longitudinal weight in every wave it exists, from the source the column names.

Table SA.13: The part of actual wealth the PSID assigned, wave by wave

Wave	Rows	Including home equity	Excluding home equity	Vehicles	Home equity	Annuities and IRAs	Stocks	Other assets
2009	4,593	14	13	6	1.0	3	2	2
2011	4,717	8	7	2	0.8	2	1	1
2013	4,668	8	8	3	0.6	1	1	1
2015	4,580	8	8	3	0.4	1	1	1
2017	4,602	9	8	3	0.5	2	1	1.0
2019	4,536	8	8	3	0.6	2	1	1.0
2021	4,408	10	10	4	0.7	2	1	0.9
2023	4,376	12	12	6	0.8	3	2	1

Notes: Percentage of household-years in each wave with the PSID's accuracy flag for an assigned value, on the rows of the 1984 to 2023 sample, unweighted. The first two columns are the PSID's constructed wealth summaries, and either records a household-year in which any component was assigned. The five component columns are the items whose flag exists in all eight waves. They are not the whole of wealth and do not sum to the first column, and the cash, business, other real estate and debt items change their flag structure inside the period. Across all family records the first column runs 13.6 to 20.6 percent. The flags reach the family file from 2009, and for 1984 to 2007 the PSID releases them in separate wealth supplement files this paper does not use.

ing value is the one it completes least often (Table SA.13). Actual wealth is the PSID's own construction, and from 2009 the flag in the family file marks each component a household would not report. The share runs against 14% to 21% of all family records, and the assignments concentrate in vehicles, at most 6% of household-years, and reach home equity in at most 1.0%.

The first certification verdict is certified on both sets, the rows the PSID assigned nothing in and the full set. I recompute its rule on the 32,971 of the 36,480 covered rows that hold no assignment, and the GB-proxy still has the lowest wealth rank error and the lowest share misplaced by more than a quarter of the distribution in all eight waves. Dropping the assigned rows raises its wealth rank error by 0.16 units and moves no other series by more than 0.04, against a lead of 5 units. The flag does not reach the eight waves from 1984 to 2007, so these shares describe the later half of the 1984 to 2023 sample.

SA.G.5 The break in the capital-income residual

The capital-income residual breaks where the PSID changed its income questions, and both fitted wealth proxies pool across the break. The pipeline's residual subtracts a labour-income total that adds other members' earnings only in 1993 and from 2005, so on the households present in both waves its mean moves by -30% between 1992 and 1993 and by -5% between 2003 and 2005. The residual this paper uses subtracts every member's taxable income in every wave, and the same two movements are -1% and 19%. One step survives every definition, a rise in the mean of 90% between 1993 and 1994, when the PSID redesigned its income ques-

Table SA.14: What the 1984 to 2023 sample keeps, wave by wave

Wave	Rows with wealth	Sample rows	Retained (%)	Homeowners (%)
1984	3,486	3,303	95	66
1989	3,772	3,544	94	70
1994	4,290	4,074	95	68
1999	4,071	3,825	94	70
2001	4,224	3,988	94	69
2003	4,246	3,985	94	70
2005	4,369	4,084	93	70
2007	4,655	4,382	94	68
2009	4,912	4,593	94	67
2011	5,049	4,717	93	65
2013	4,993	4,668	93	64
2015	4,890	4,580	94	63
2017	4,901	4,602	94	64
2019	4,790	4,536	95	64
2021	4,679	4,408	94	67
2023	4,625	4,376	95	68

Notes: Household-years with actual wealth, and those also with the three wealth proxies and the rent-ordered variant of Appendix A.1, the 1984 to 2023 sample. The 2011 wave retains 93 percent of its rows against 93 to 95 elsewhere, with homeowners 65 percent of what survives there.

tions, so a study of the pre-1984 waves inherits a coefficient fitted partly on a differently defined variable.

SA.G.6 Retention and the pooled profile

The wave with the weakest retention, 2011, keeps 93% of the household-years with actual wealth against 93 to 95% elsewhere (Table SA.14). Dropping it leaves the GB-proxy's outperformance on wealth rank error at 32% against 32% and the housing wealth proxies' mean distance from the actual Gini at 0.16 against 0.16.

A profile pooled across waves is an age gradient across birth cohorts, and under it the GB-proxy understates the rise it overstates within a birth cohort. Pooling the 1984 to 2023 sample and taking medians by age band gives a rise of 11.6 times for actual wealth against 2.8 for the PK-proxy and 9.3 for the GB-proxy. The mean survey year is 2007.4 in the youngest age band and 2009.6 in the oldest, so the two bands are different birth cohorts seen at one date. The PK-proxy understates the rise by a factor of 4.2 pooled against 2.8 to 3.5 within birth cohort. Dropping 2011 moves the pooled ratio from 11.6 to 11.0 times for actual wealth and from 2.8 to 2.7 times for the PK-proxy.

SA.H The estimator

Every figure in this section comes from the GB-proxy's own fitted models, refitted by the production code and checked to reproduce the series exactly in every wave. The GB-proxy is the construction of Section 2, one gradient-boosting model per group of households and set of waves, fitted on the 16 waves that record actual wealth, with dollar input variables in constant dollars, an ownership history over the four preceding waves and no home-equity input variable.

SA.H.1 The fit

Boosting builds each prediction as a sum of corrections, starting from the training mean of its group's target:

$$\hat{\mathcal{F}}^{(0)} = \frac{1}{N} \sum_{i,t} \tau_{i,t}, \tag{SA.4}$$

where $\hat{\mathcal{F}}^{(0)}$ is the initial prediction, $\tau_{i,t}$ is the modelling target of Equation 4 for household i in wave t, and N is the number of training rows in its group, one per household and wave. At each round k a weak learner $\zeta^{(k)}$ fits the residuals of the current prediction, which updates as:

$$\hat{\mathcal{F}}^{(k)}(\mathbf{x}) = \hat{\mathcal{F}}^{(k-1)}(\mathbf{x}) + \lambda\, \zeta^{(k)}(\mathbf{x}). \tag{SA.5}$$

The learning rate λ scales each correction, $\mathbf{x}$ is a household's vector of input variables, and the loss is mean squared error. The number of rounds, the tree depth and the learning rate appear in Table SA.15, and a grid on this construction finds no setting that improves on them by more than one fold standard error (Appendix A). Each tree takes 0.8 of the rows and 0.8 of the input variables, I trim each group's target at its 2.5 and 99.5 percentiles on the training rows alone, and every out-of-fold figure below comes from five folds blocked by household.

SA.H.2 Construction choices

The fold rule, the trim and the loss come from the construction that preceded this one, which had two models and ten input variables. Folds assigned to household-years let a model fit a household it then predicts. They gave a wealth rank error of 14.77 units, against 14.92 for folds blocked by household and 15 for folds blocked by 1968 family, so every figure here blocks on households. Removing the trim gave the worst wealth rank error of four rules, 15.42 against 14.92, and the closest Gini distance, 0.05 against 0.07. A pseudo-Huber loss improved both, to

Table SA.15: The tuning of the fitted models

	Homeowners	Renters	Neither
Boosting rounds	140	90	90
Tree depth	9	6	6
Learning rate, percent	4.5	6.0	6.0
Row subsample per tree	0.8	0.8	0.8
Column subsample per tree	0.8	0.8	0.8
Loss	squared error	squared error	squared error
Starting value	training mean	training mean	training mean
Trim, percentiles of the target	2.5 to 99.5	2.5 to 99.5	2.5 to 99.5
Folds, blocked by household	5	5	5
Scale-model depth	4	4	4
Input variables	15	14	13
Fitted rows	43,976	21,658	4,156

Notes: Every value is read from the fitted booster and asserted equal to the call that fitted it. The rounds, depth and learning rate were chosen once, by a grid on the preceding construction, and inherited unchanged. The modelling target is wealth over the housing value for homeowners and wealth over the wave's mean wealth for the other two groups, the homeowner ratio divided by the wave's aggregate ratio. The starting value is the mean of that target on the fitted rows, and the trim drops rows outside the stated percentiles of it before fitting.

14.68 rank units and a Gini distance of 0.05, but I keep squared error because on the backward split, fitted from 1999 and scored on 1984 to 1994, it recorded 15.09 against 15.03.

SA.H.3 The leading input variables

The mortgage ratio is what the homeowner model uses most, at 27% of its Shapley contributions and 53% for the top three input variables, and the two other ways of ranking the input variables agree with it (Table SA.16, Figure SA.1). For renters it is income to needs, at 29% and 67% for the top three, and for households reporting neither a housing value nor rent it is capital income, at 29% and 63%, from the exact TreeSHAP decomposition on each model's fitting rows. Split gain ranks the input variables the same way, at rank correlations of 0.91, 0.78 and 0.77. Out of fold, permuting the mortgage ratio raises the homeowner model's wealth rank error by 7 rank units from 15, more than any other, and the permutation drops rank the input variables as the Shapley shares do, at correlations of 0.96, 0.92 and 0.83.

SA.H.4 The shape of the response

Across the middle 95 percent of its leading input variable, the homeowner prediction moves 0.75 standard deviations of its target, the renter prediction 0.95 and the third group's 0.79 (Figure SA.2). The figure draws each model's partial dependence on its five leading input variables, so the shape of each response is visible as well as its range.

Table SA.16: What the fitted models use, by household group

	Homeowners			Renters			Neither		
Input variable	Shapley	Gain	Permutation	Shapley	Gain	Permutation	Shapley	Gain	Permutation
Mortgage ratio	27	19	7	—	—	—	—	—	—
Capital income	15	19	2	22	16	2	29	37	3
Income to needs	12	13	2	29	28	2	18	16	2
Housing value	11	13	1	—	—	—	—	—	—
Age	9	10	1.0	16	14	1.0	16	19	2
Business held	9	5	1.0	5	3	0.5	5	4	0.2
Rooms per person	5	4	0.3	2	3	0.2	6	6	0.6
Minimum wage ratio	4	5	0.4	4	8	0.3	3	3	0.7
Family structure	3	1.0	0.2	0.6	0.9	0.0	6	2	0.5
Labour income	2	5	0.4	4	7	0.3	6	6	2
Money given away	1	3	0.1	1	3	0.0	1	2	0.0
Years unable to work	1	1	0.1	2	0.9	0.0	2	2	0.0
Household size	1.0	2	0.0	1	2	0.1	2	3	0.0
Food stamps received	0.5	0.3	0.0	4	0.7	0.5	2	0.3	0.5
Waves spent owning	0.5	0.4	0.0	3	2	0.1	4	2	0.0
Rent	—	—	—	8	12	0.3	—	—	—

Notes: One row per input variable, in the homeowner model's order. Shapley is the input variable's share of the mean absolute Shapley contribution over the fitted rows, in percent. Gain is its share of the model's split gain, in percent. Permutation is the wealth rank error a random permutation of the input variable adds on one household-blocked held-out fold, in rank units, the mean of five permutations. A dash marks an input variable the group's model does not take. The models are the GB-proxy's construction of Section 2, one per household group and set of waves, for the set serving the most waves.

SA.H.5 Rounds and tuning

Held-out wealth rank error reaches its minimum at round 100 for homeowners against the 140 chosen, a cost of 0.0% (Figure SA.3). For renters the minimum is at round 120 against 90, a cost of 0.2%, and for the third group at round 30 against 90, a cost of 1.1%. Moving each model's tree depth and learning rate one step in each direction gives nine settings whose wealth rank error spans 0.48 units and whose Gini distance spans 0.02 (Table SA.17). I retune nothing on this evidence, since the grid of Appendix A already finds no setting better by more than one fold standard error.

SA.H.6 The conditional distribution

The point series is a conditional mean, less dispersed than wealth, so the GB-proxy also comes as 50 draws per household-year. A draw adds to the location model's prediction a standardized residual from its empirical distribution, scaled by a second model fitted at depth four to the log squared out-of-fold residual, and 20 household bootstraps put model uncertainty at 32% of the draw variance.

The draws reproduce actual wealth's rank autocorrelations to within a hundredth and its out-of-fold Gini, 0.79 against 0.79, where the point series gives 0.72. I calibrate a household's posi-

Figure SA.1: What each fitted model uses: Shapley share and split-gain share, by input variable

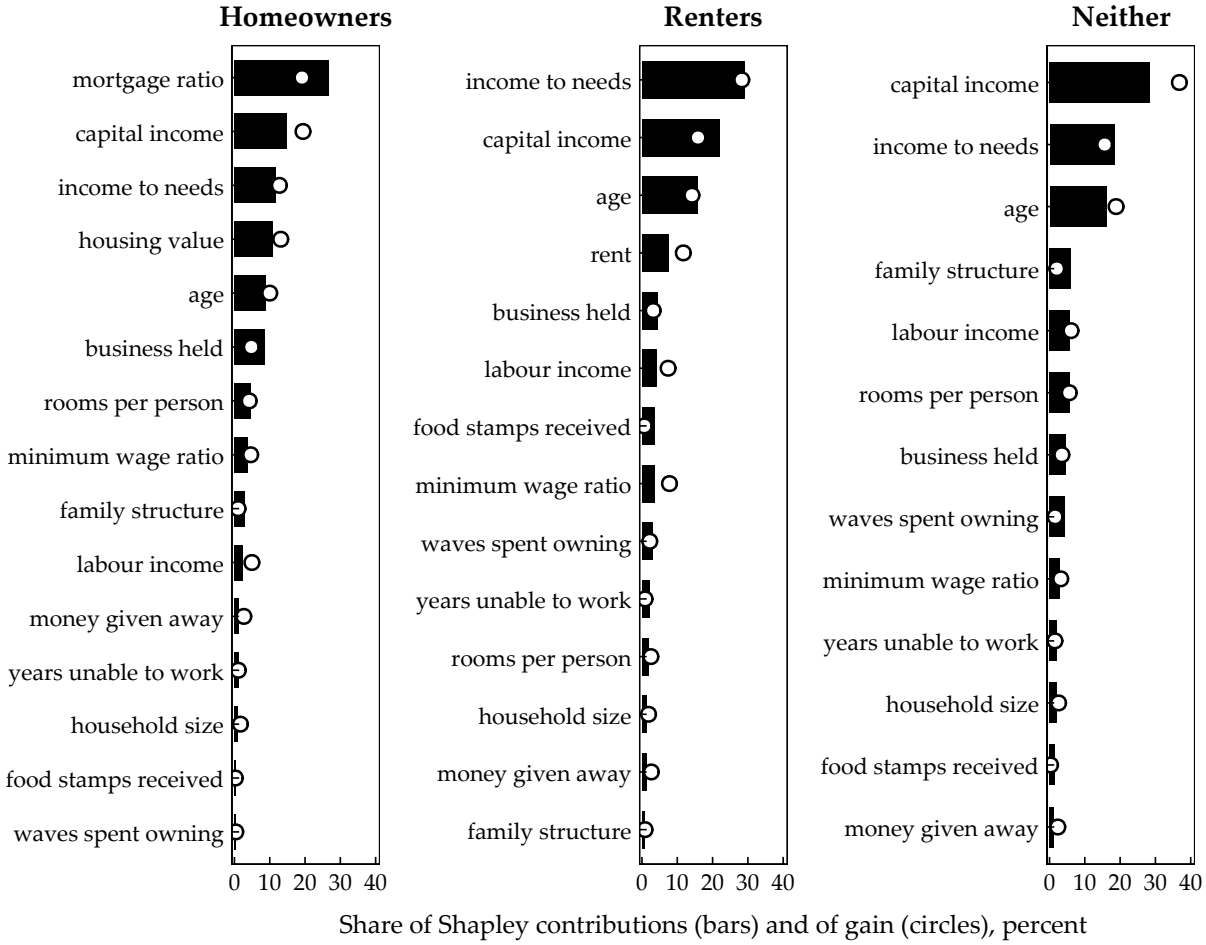

Notes: Each panel is one household group and each bar one input variable, ordered within the panel. A bar is the input variable's share of the mean absolute Shapley contribution over the group's fitting rows, in percent. An open circle is the input variable's split-gain share in the same model. The models are the GB-proxy's construction, for the set of waves serving the most waves.

tion in its own conditional distribution as a mixture of a slow and a fast process, with weight 0.90 on the slow one and one-year correlations of 0.98 and 0.95, to match wealth persistence under actual wealth. The within-wave rank autocorrelations at two and five years are then 0.88 and 0.81, against 0.88 and 0.80 for actual wealth, and a household moves 9.6 rank units between consecutive wealth waves in a draw against 9.2 in actual wealth.

Figure SA.2: Partial dependence of each fitted model on its five leading input variables

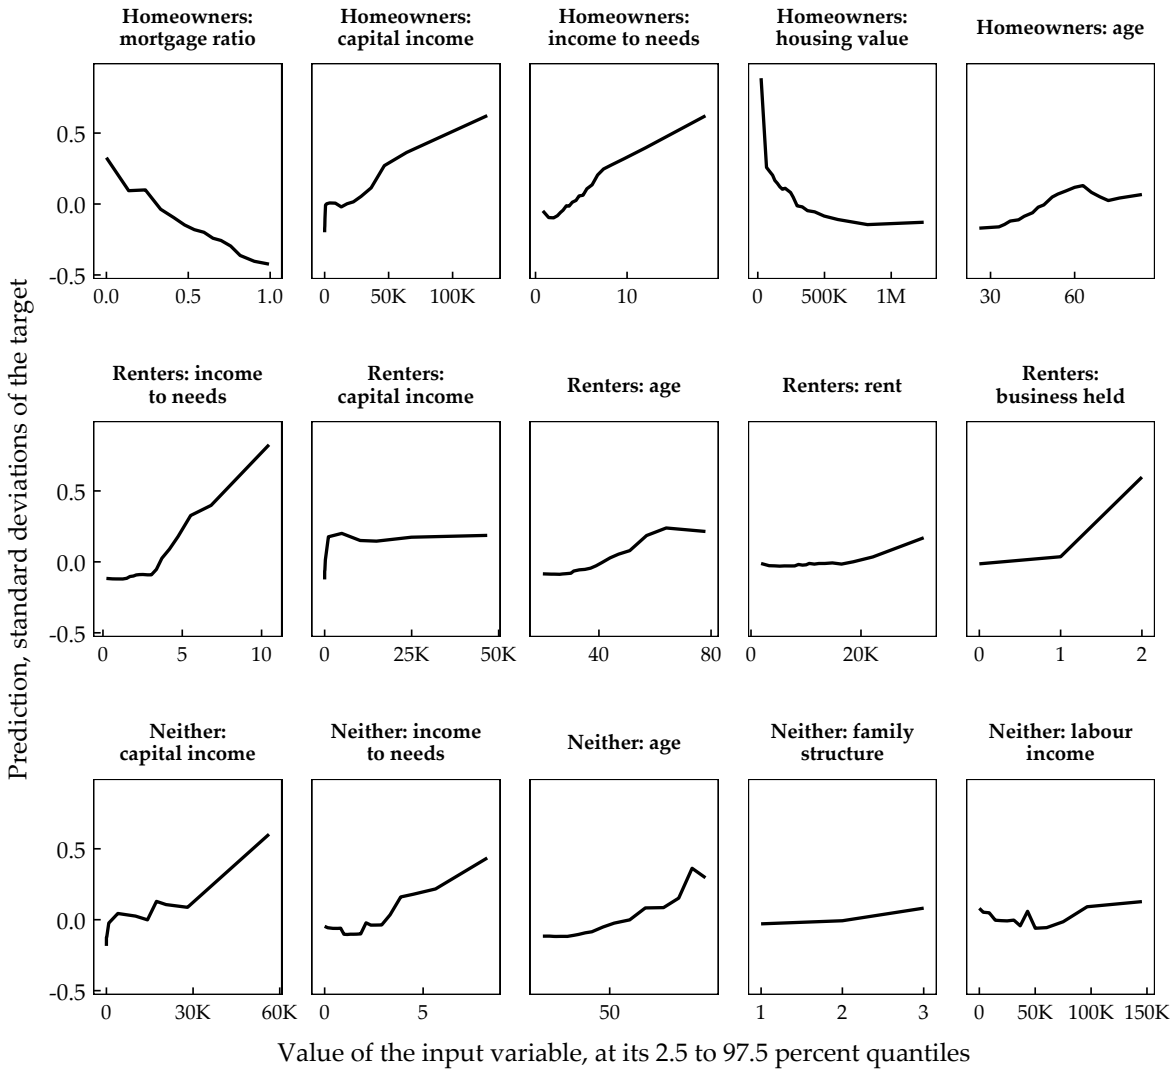

Notes: Each panel is one input variable of one household group, the five input variables with the largest Shapley share in each group. The line is the mean prediction over 2000 rows drawn from the group's fitting rows when the input variable takes the value on the horizontal axis, less the mean prediction over the same rows unchanged. The unit is a standard deviation of the group's modelling target. The grid is twenty quantiles between the input variable's 2.5 and 97.5 percent points. An input variable with ten or fewer distinct values uses those values.

Table SA.17: The chosen tuning against its neighbours, out of fold

Tree depth	Learning rate	Wealth rank error	Gini distance
chosen -3	chosen × 0.5	15	0.08
chosen -3	chosen	15	0.07
chosen -3	chosen × 2	15	0.07
chosen	chosen × 0.5	15	0.07
chosen	**chosen**	**15**	**0.06**
chosen	chosen × 2	15	0.06
chosen +3	chosen × 0.5	15	0.07
chosen +3	chosen	15	0.06
chosen +3	chosen × 2	15	0.06

Notes: Each row is one tuning of all three models, out of fold on the five household-blocked folds, at the chosen number of rounds. The depth and the learning rate move by the same step in every model, so a row is a tuning decision a user of this construction could take. The chosen row is in bold, and no series in this paper is refitted on this evidence.

Figure SA.3: The learning curve of each fitted model on one held-out fold

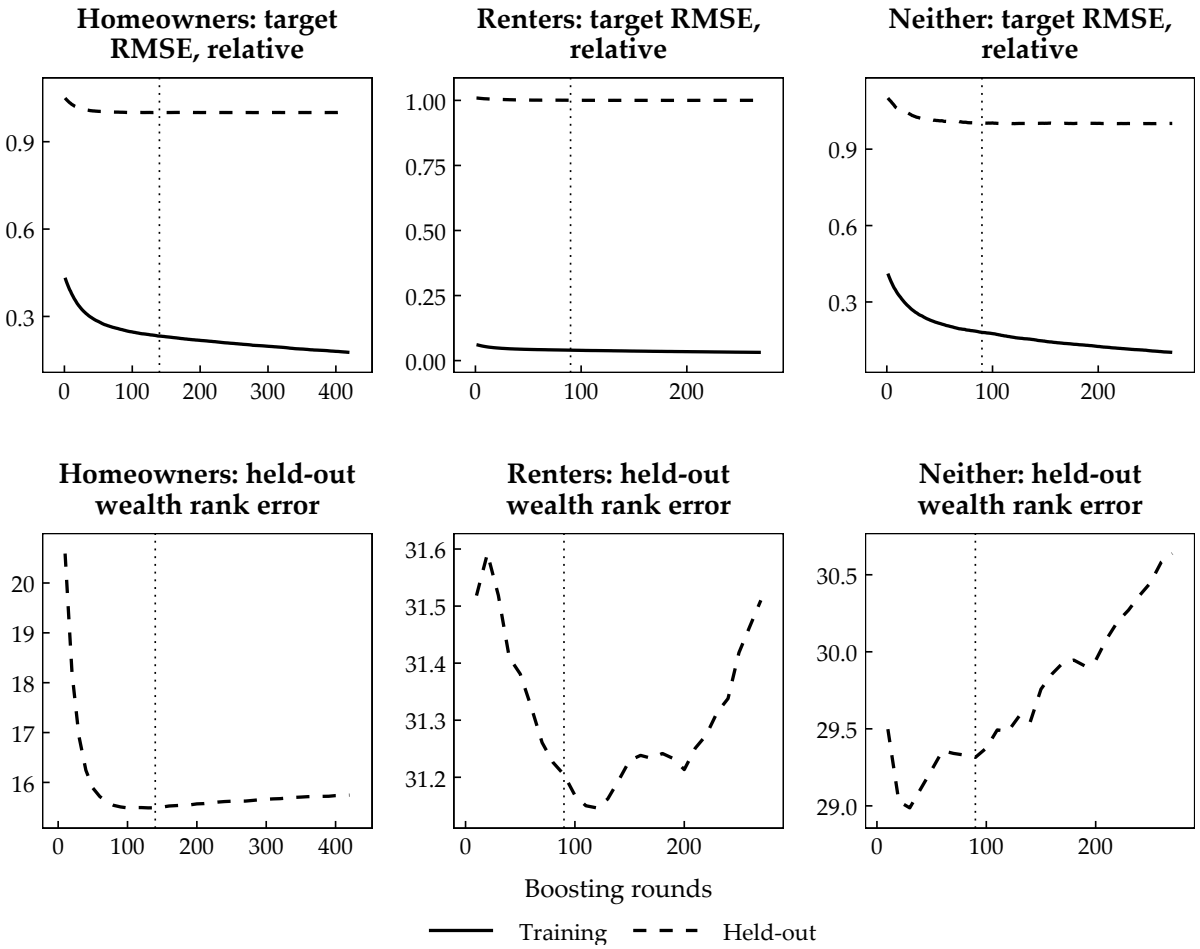

Notes: Each panel is one household group and one measure, on its own vertical scale. The target RMSE panels show the root mean squared error on the modelling target, on the training rows (solid) and on the held-out fold (dashed), relative to the held-out minimum. The wealth rank error panels show the held-out wealth rank error every ten rounds. The dotted line marks the chosen number of rounds, and the held-out fold is one of five blocked by household.

References

Benjamini, Yoav, and Yosef Hochberg. 1995. "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing." *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1): 289–300.

Choi, Jung Hyun, and Gary Painter. 2018. "Self-Reported vs. Market Estimated House Values: Are Homeowners Misinformed or Are They Purposely Misreporting?" *Real Estate Economics*, 46(2): 487–520.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri. 2016. "Heterogeneity in Returns to Wealth and the Measurement of Wealth Inequality." *American Economic Review*, 106(5): 651–655.

Fitzgerald, John, Peter Gottschalk, and Robert Moffitt. 1998. "An Analysis of Sample Attrition in Panel Data: The Michigan Panel Study of Income Dynamics." *The Journal of Human Resources*, 33(2): 251–299.

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

Juster, F. Thomas, and Kathleen A. Kuester. 1991. "Differences in the Measurement of Wealth, Wealth Inequality and Wealth Composition Obtained from Alternative U.S. Wealth Surveys." *Review of Income and Wealth*, 37(1): 33–62.

Kiel, Katherine A., and Jeffrey E. Zabel. 1999. "The Accuracy of Owner-Provided House Values: The 1978–1991 American Housing Survey." *Real Estate Economics*, 27(2): 263–298.

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.

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

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