Beyond Variance:
Asymmetric Labor Income Risk and Portfolio Entry

August 2026

Abstract

Variance treats favorable and adverse earnings realizations as the same risk. I use public CPS histories and strictly later SIPP portfolios to separate lower-tail distance, upper-tail distance, total tail span, and orientation. Among 27,752 annual transitions beginning with no reported stock-and-mutual-fund holdings, a one-cross-cell-standard-deviation wider lower tail predicts a 2.71-percentage-point lower probability of reporting positive holdings one year later; a one-standard-deviation wider upper tail predicts a 1.92-point higher probability. The two distances remain jointly significant after 200 CPS resamples rebuild the first stage ($p = 0.0022$), while the total span is imprecisely estimated. A scale-free incremental Kelley contrast predicts a 1.05-point lower transition relative to an 8.79% rate. Measurement is consequential: zero-inclusive Kelley reliability is 0.74 in CPS but falls to 0.37 beyond zero incidence, mean, and dispersion. The estimates are descriptive: historical group distributions are neither exogenous nor elicited beliefs.

Keywords: labor income risk; skewness; measurement error; household portfolios; public surveys; generated regressors
JEL codes: C18, D14, G11, J31

1 Introduction

Labor income is the largest nontradable asset of many households, but "more labor-income risk" is incomplete when favorable and adverse tails move differently. A wider lower tail expands the set of earnings disasters; a wider upper tail expands the set of earnings opportunities. Variance and the 90–10 span treat both changes as more risk, although they can imply opposite portfolio choices. The paper's question is therefore not whether downside risk discourages stockholding—that result is well established—but whether the two tails separately organize a concrete household decision margin (Viceira, 2001; Cocco, Gomes, and Maenhout, 2005; Chang, Hong, and Karabarbounis, 2018; Catherine, 2022; Catherine, Sodini, and Zhang, 2024; Shen, 2024; Gálvez and Paz-Pardo, 2026).

I find opposite conditional associations for the transition to positive reported holdings. In the exact lagged-nonholder risk set, lower 50–10 and upper 90–50 distances enter with opposite signs conditional on each other, the zero-earnings probability, mean growth, dispersion, household controls, and fixed effects. A one-standard-deviation wider lower tail predicts a 2.71-percentage-point lower transition rate; a one-standard-deviation wider upper tail predicts a 1.92-point higher rate. Under the 200-draw CPS first-stage sensitivity, their joint p-value is 0.0022; the lower coefficient has $p = 0.0015$ and the upper has $p = 0.068$. Rotating the same model into total 90–10 span and lower-minus-upper composition leaves the span imprecisely estimated and the composition contrast significant ($p = 0.014$); a raw-scale test also rejects equality of the two marginal effects ($p = 0.013$). Thus the finance result is not "downside matters." It is an opposite-sign both-tails pattern on a concrete portfolio transition.

Empirical work measures background risk in two ways. One elicits households' own expectations, with applications to portfolio choice and the calibration of earnings risk (Guiso, Jappelli, and Terlizzese, 1996; Wang, 2023; Caplin et al., 2026; Arellano et al., 2026). The other estimates risk from realized histories and assigns the resulting statistic across workers or samples (Angerer and Lam, 2009; Betermier et al., 2012; Chang et al., 2022; Fagereng, Guiso, and Pistaferri, 2018; Ranish, 2013; d'Astous and Shore, 2024). The first targets perceived risk but requires specialized elicitation; the second scales to public data, while making construction and generated-regressor uncertainty part of the empirical object. This paper uses the second design for its primary out-of-sample portfolio test and applies the same decomposition to individual SCE densities as a separate construct check. It does not equate assigned histories with perceptions or claim that perceptions dominate assigned measures.

Three empirical achievements are distinct in such a design. *Reliability* asks whether independent workers reproduce the group ranking. *Incremental content* asks whether the ranking supplies variation in tail orientation conditional on destination-zero incidence, realized cell mean growth, and dispersion. *Downstream uncertainty accounting* asks whether uncertainty from estimating a nonlinear group statistic is carried into the household regression. A split-half correlation answers the first question. It cannot, by itself, answer the other two. Indeed, reliability can converge to one when the reproducible component is entirely a stable function of lower moments.

I separate these objects using matched Current Population Survey Annual Social and Economic Supplement (CPS ASEC) records and the Survey of Income and Program Participation (SIPP). Baseline-positive earners are followed for one year, transitions to zero earnings are retained, and growth is measured by the bounded arc rate. The primary CPS construction uses annual wage-and-salary income and the exact public industry, education, and age crosswalk needed for the portfolio application; it contains 183,809 transitions. SIPP instead measures monthly main-job earnings. Every split is global at the person level, so the same worker never appears in both reliability halves. Parallel SIPP estimates provide a second-survey check, not a claim that the two zero endpoints are identical.

The central measurement fact is a construction reversal. Conditioning on positive earnings at both endpoints produces a weakly reproducible group ranking. Retaining transitions to zero makes downside-oriented Kelley shape reproducible across independent workers. In the pooled exact CPS crosswalk, Spearman–Brown reliability is 0.740 for raw Kelley, 0.646 after the zero-transition share and demographic components are separated, and 0.370 after mean growth and dispersion are also separated; the zero-transition share itself has reliability 0.940. Within SIPP, the one-transition-per-person check gives 0.602, 0.125, 0.151, and 0.809. The zero anatomy makes the interpretation sharper. Among CPS annual wage zeros, 64.5% report no work during the year, but the worked-without-wage component is even more reproducible than the no-work component. Among later SIPP monthly main-job zeros, 64.5% record a contemporaneous self-employment arrangement and only 2.1% record a no-job spell. The conclusion is not that Kelley shape is an invalid descriptive statistic: zeros alter the measured unconditional distribution. It is that raw reliability can rank a survey-specific support-boundary composite without identifying a common shape-specific contrast. Separating location from orientation is not an arbitrary over-control: Harmenberg (2021) reports mean-growth correlations of 0.87–0.88 with skewness measures over time and 0.96–0.97 with Kelley skewness across education groups in Danish registers.

The finance application estimates group distributions using CPS origin years through 2017 and assigns them to SIPP portfolios observed from 2019 through 2024, all strictly after the last CPS destination year. SIPP observes stock-and-mutual-fund balances, not total household equity. Among lagged nonholders, the lower and upper tail distances have opposite conditional associations with the transition to positive holdings, while the total span is imprecisely estimated conditional on dispersion and orientation. The scale-free summary gives the same ordering: a one-standard-deviation increase in incremental downside orientation predicts a 1.05-percentage-point lower one-year transition to positive observed holdings (standard error 0.27 point), compared with a weighted transition rate of 8.79%. The conditional $p = 0.00030$ and three-outcome $q = 0.00089$ become $p = 0.0052$ and $q = 0.0157$ when 200 CPS resamples rebuild support, quantiles, residualization, standardization, assignment, and the transition sample. The estimate is -0.96 point when both SIPP observations remain in the same signal cell, ranges from -1.20 to -0.88 point when each outcome year is omitted in turn, and remains negative for sustained entry and positive-balance thresholds. In contrast, the standing all-worker share coefficient of -0.234 point has conditional $p = 0.019$ but plug-in $p = 0.109$. Independent CPS and SIPP samples remove own-observation mechanical correlation, not sorting into signal cells. These are descriptive early-sample assignment tests, not causal portfolio responses.

This transition margin connects the measurement audit to the household-finance distinction between participation and allocation. Fixed and per-period participation costs can make the acquisition and retention of risky assets state dependent even when a conditional risky share changes little (Vissing-Jorgensen, 2002; Gomes and Michaelides, 2005; Fagereng, Gottlieb, and Guiso, 2017). Beliefs need not determine whether an investor adjusts: Giglio et al. (2021) finds that belief changes do not predict when investors trade, although they predict the direction and magnitude conditional on trading. The empirical outcome here is narrower: a transition from zero to positive observed stock-and-mutual-fund holdings, not first-ever account opening, and the fixed-cost benchmark motivates examining the discrete margin without identifying a cost parameter. The result says that a corrected public-cell tail measure can be informative about a decision margin even when inference for standing portfolio levels remains fragile; it does not show that the assigned history changed beliefs or triggered acquisition.

That fragility does not imply that genuine distributional orientation is economically irrelevant. In a sealed two-period benchmark, I compare an empirical CPS earnings-growth law with its mirror image while holding the complete absolute-shock distribution fixed. The laws have the same mean and every even moment, yet optimal risky investment can differ, especially at fixed-cost participation thresholds. The benchmark illustrates the stakes: the difficulty is not showing that shape can matter in the benchmark, but recovering stable shape-specific variation in public group cells.

Recent work already establishes that non-Gaussian earnings risk can matter economically. Catherine, Sodini, and Zhang (2024) and Shen (2024) connect state-dependent skewness to portfolios, while Gálvez and Paz-Pardo (2026) shows that rich non-normal earnings dynamics help a structural life-cycle model rationalize limited participation and low risky holdings. Very close concurrent work by van den Akker et al. (2026) links subjective expectations to administrative portfolios and reports separately that higher expected income and a longer perceived right tail predict stock and bond investment. Its publicly available abstract does not disclose the exact tail functional or whether total dispersion enters jointly, but it rules out a priority claim that perceived earnings upside affects actual portfolios. Those objects differ from permanent and transitory variance, labor-income covariance with market returns, unconditional group tail orientation, and subjective perceived distributions. In particular, the education-by-industry cyclical-skewness comoment in Catherine, Sodini, and Zhang (2024) is not the unconditional public-cell Kelley ranking studied here: rich administrative evidence can identify its object even when a public-survey proxy fails the present incremental-content gate. The question here is logically prior and narrower: whether low-information public-survey cells provide a reproducible shape-specific contrast that can support a reduced-form household assignment.

The paper makes two contributions. The household-finance contribution is a both-tails result on a decision margin. Conditional lower and upper tail distances predict the transition to positive reported holdings with opposite signs; an exact rotation shows a precise composition contrast and an imprecise total-span coefficient. The scale-free orientation result survives the paper's independently rebuilt CPS first stage, while the analogous standing-share result does not. This separates opportunities from disasters and transitions from accumulated holdings, rather than relabeling a conventional downside-risk result.

The measurement contribution is to show what an assigned tail contrast must clear before that finance interpretation is available. In the literature on non-Gaussian earnings dynamics (Guvenen, Ozkan, and Song, 2014; Arellano, Blundell, and Bonhomme, 2017; Guvenen et al., 2021; Busch et al., 2022), robust quantile statistics can limit outlier sensitivity (Kim and White, 2004), but they do not make a raw group ranking shape-specific. I combine person-level reproducibility, incremental content, cross-survey transport, and nonlinear first-stage rebuilding for the same statistic. A direct household-finance precedent is Ranish (2013), whose PSID-to-SCF design bootstraps estimated labor-income risks and generated regressors; recent factor evidence likewise shows that construction choices can change downstream finance conclusions (Akey, Robertson, and Simutin, 2026). Here, clustering by assigned group handles shared outcome residuals but does not propagate uncertainty from estimated group quantiles (Pagan, 1984; Murphy and Topel, 1985; Moulton, 1990; Cameron, Gelbach, and Miller, 2011). Historical group measures are not elicited beliefs, so the results establish neither how households learn about earnings risk nor that perceived risk dominates assigned histories.

The scope is deliberately narrow. The measured distribution is historical rather than an elicited expectation; annual growth is not decomposed into permanent and transitory shocks; group membership and the lagged-nonholder risk set are selected; survey-specific wage and main-job endpoints are not a common job-loss measure; and modest incremental-signal reliability prevents interpreting an imprecise coefficient as a zero effect of latent shape. The transition family and tail decomposition were selected in an exploratory audit rather than a preregistered design. The reported false-discovery-rate families do not cover every outcome, threshold, or specification examined during the project's history. The supported conclusion is therefore descriptive but sharper than a downside-only claim: assigned upper and lower earnings tails have opposite associations with a subsequent portfolio transition, while a reproducible raw downside ranking need not be a stable shape-specific portfolio regressor.

Section 2 defines the economic objects and diagnostics. Section 3 describes the public measurement and portfolio designs. Section 4 shows what the group rankings measure. Section 5 reports the early-sample finance application. Section 6 reports the CPS-first-stage sensitivity and summarizes additional construct checks. Section 7 gives the implications for similar generated-risk designs. Section 8 concludes. Detailed estimator tournaments, robustness results, and the SCE timing audit are in the Appendix.

2 What a Group Tail Statistic Can Identify

2.1 Three economic objects

Let G_c denote the distribution of one-year earnings growth for workers in cell c, and let $q_{p,c}$ be its pth quantile. I orient every asymmetry measure so that a positive value denotes a longer or heavier downside. The upper and lower 90–50 gaps are

$$U_c = q_{0.90,c} - q_{0.50,c}, \qquad L_c = q_{0.50,c} - q_{0.10,c}. \tag{1}$$

Their sum $S_c = L_c + U_c$ is the 90–10 spread and their difference $D_c = L_c - U_c$ is an unscaled tail imbalance. The primary scale-free statistic is downside-oriented Kelley shape,

$$K_c = \frac{L_c - U_c}{L_c + U_c} = \frac{2q_{0.50,c} - q_{0.10,c} - q_{0.90,c}}{q_{0.90,c} - q_{0.10,c}}. \tag{2}$$

Bowley shape replaces the 10th and 90th percentiles with the 25th and 75th percentiles. I also consider downside-oriented L-skewness, conventional standardized third-moment skewness with its sign reversed, and probability imbalance beyond one within-cell standard deviation.

These estimators answer different questions. Tail gaps measure magnitudes; tail-probability balance measures mass; Kelley and Bowley normalize by an interquantile range; moment skewness is sensitive to the full support. The paper distinguishes three economic objects. The *joint downside environment* is the unconditional distribution G_c, including transitions to zero measured earnings. *Conditional wage-growth shape* is the distribution among workers with positive earnings at both endpoints. *Incremental tail orientation* is the component of a stated tail statistic that varies within the observed support after the zero-transition probability, mean growth, and dispersion are partialled out. The first two objects are distributions. The third is a conditional empirical contrast, not a primitive structural parameter.

A destination-zero mass links the first object to lower moments exactly.

Lemma 1 (A point mass below the positive-earnings support)**.** *Let F have support strictly above a, finite mean μ_F, and variance σ_F^2. For $p \in [0, 1)$, let $G_p = p\delta_a + (1-p)F$. Using* *generalized-inverse quantiles, for $\tau \in (0, 1)$,*

$$q_\tau(G_p) = \begin{cases} a, & \tau \le p, \\ F^{-1}((\tau - p)/(1-p)), & \tau > p, \end{cases}$$

while

$$\mu(p) = pa + (1-p)\mu_F, \qquad \mathrm{Var}(G_p) = (1-p)\sigma_F^2 + p(1-p)(\mu_F - a)^2.$$

Consequently, cross-cell variation in a stable destination-zero share changes the quantile map, mean, and dispersion even when the positive-earnings distribution F is identical. The resulting Kelley statistic remains a legitimate description of G_p. What it cannot do alone is identify a counterfactual comparison that changes tail orientation while holding destination-zero incidence and lower moments fixed. Lemma 1 does not imply that Kelley shape is monotone in p without additional restrictions on F. Appendix A.1 gives the proof.

2.2 Support: conditional wage growth or earnings risk

Log growth requires positive earnings in both periods. Applying it to continuously positive earners estimates the distribution of wage growth conditional on positive measured earnings at both endpoints. That object can be useful, but it excludes transitions to zero measured earnings, which can reflect job loss, nonemployment, changes in work intensity or work arrangement, and survey-specific earnings definitions. Labor-market transitions are economically relevant to perceived earnings risk (Caplin et al., 2026), but the anatomy of a recorded zero must be shown rather than inferred from the endpoint alone. Adding one dollar before taking logs retains zeros numerically but imposes a unit-dependent distance between zero and small positive earnings.

The baseline therefore uses bounded arc growth (Törnqvist, Vartia, and Vartia, 1985),

$$g_{i,t+1} = 2\frac{y_{i,t+1} - y_{i,t}}{y_{i,t+1} + y_{i,t}}, \tag{3}$$

where baseline earnings are positive and destination earnings may be zero. The rate is symmetric in the two earnings levels, lies in $[-2, 2]$, and maps a destination zero to -2 without an arbitrary offset. Support choice is part of the estimand: positive-to-positive log and arc growth are reported as alternative objects, not as cleaning variations around the same object.

2.3 Reliability

For each construction, I assign every unique person globally to one of two halves, estimate cell statistics separately, and calculate the correlation $r_{1/2}$ across cells. Global assignment prevents the same worker from entering both halves after an industry or education change. The full-sample reliability projection is the Spearman–Brown transform (Spearman, 1910; Brown, 1910)

$$\rho_{\text{SB}} = \frac{2r_{1/2}}{1 + r_{1/2}}. \tag{4}$$

The reported value applies equation (4) to the mean half-sample correlation across random person splits. I also report rank correlation and early-versus-late stability. Support is measured using the Kish effective sample size (Kish, 1965), $N_{\text{eff}} = (\sum_i w_i)^2 / \sum_i w_i^2$, not the raw number of records.

Reliability has a narrow interpretation. The following decomposition makes the limit explicit.

Lemma 2 (Reliability of a composite statistic)**.** *Suppose the two independent half-sample estimates of a group statistic satisfy*

$$\widehat{K}_c^{(h)} = a'Z_c + s_c + e_c^{(h)}, \qquad h \in \{1, 2\},$$

where Z_c is a stable vector of lower distributional features, s_c is stable incremental tail variation, and the half-sample errors are mutually independent and independent of (Z_c, s_c). If their error variances are equal to σ_e^2, then

$$\text{Corr}\Big(\widehat{K}_c^{(1)}, \widehat{K}_c^{(2)}\Big) = \frac{\text{Var}(a'Z_c + s_c)}{\text{Var}(a'Z_c + s_c) + \sigma_e^2}.$$

Reliability can therefore converge to one as sampling error vanishes even when $\text{Var}(s_c) = 0$.

The result does not declare $a'Z_c$ spurious. It shows that reliability cannot determine which stable economic feature produces the ranking. That determination requires a stated conditional contrast.

2.4 Incremental content

I estimate a sequential projection across supported cells. The most demanding specification is

$$K_c = a + \gamma p_{0c} + f_2(\mu_c) + h_2(\sigma_c) + \lambda_{e(c)} + \lambda_{a(c)} + u_c, \tag{5}$$

where p_{0c} is the transition-to-zero probability; f_2 and h_2 are quadratic functions; and the fixed components capture education and age bands. The fit is weighted by effective cell support capped at 5,000. I report the raw statistic, the residual after the linear p_{0c} term and demographic components, and the residual after the complete projection. Equation (5) is re-estimated within each person half. Residualizing the full-sample statistic once and then splitting would leak information across the reliability test.

Including mean growth is particularly important given the strong empirical association between location and Kelley shape documented by Harmenberg (2021). The projection defines a conditional contrast; it does not declare location or scale economically spurious.

The residual u_c is an operational incremental signal, not "true skewness." Equation (5) does not recover a latent shock process, prove that correlated distributional features are confounds, or define a universally preferred shape estimator. It asks the narrower question required by a shape-specific portfolio interpretation: after allowing a transparent relation with the positive-measured-earnings boundary, location, and scale, is enough stable cross-cell variation left to support a downstream regression?

2.5 Downstream estimand

Let $\widehat{u}_c^{\mathrm{CPS}}$ be the standardized incremental signal estimated outside the SIPP. For worker i in SIPP year t, the portfolio application estimates

$$P_{it} = \beta\widehat{u}_{c(i)}^{\mathrm{CPS}} + \rho\widehat{p}_{0,c(i)} + f_2(\widehat{\mu}_{c(i)}) + h_2(\widehat{\sigma}_{c(i)}) + X_{it}'\gamma + \delta_t + \delta_j + \delta_e + \delta_a + \varepsilon_{it}. \tag{6}$$

Controls include log income, log wealth, age and age squared, sex, and an unemployment-spell indicator; fixed effects cover year, broad industry, education, and age band. SIPP person-year weights are used. Conditional standard errors are two-way clustered by SIPP person and CPS signal cell; the CPS-first-stage sensitivity also resamples the independent CPS measurement sample and rebuilds equation (5). Coefficients are descriptive associations per one cross-cell standard deviation of $\widehat{u}_c^{\mathrm{CPS}}$. Group assignment is endogenous, and equation (6) is not interpreted causally.

The dynamic application defines $H_{it} = 1\{A_{it} > 0\}$, where A_{it} is the observed stock-and-mutual-fund balance. For consecutive annual observations, signed participation change is $\Delta H_{it} = H_{it} - H_{i,t-1}$; entry is H_{it} in the risk set $H_{i,t-1} = 0$; and exit is $1 - H_{it}$ in the risk set $H_{i,t-1} = 1$. The signal and all cell-distribution controls are assigned from the lagged observation. Household controls and fixed effects are likewise lagged except for the current outcome-year effect, and the current SIPP weight is used. Both observations must lie in 2019–2024, so transition outcomes span 2020–2024. The entry outcome is a nonholder-to-holder transition in observed balances, not necessarily a household's first-ever stock-market entry.

2.6 Why the conditional contrast can matter

A small portfolio benchmark isolates orientation from magnitude. Let F be a centered, unit-variance law formed from corrected CPS arc-growth residuals. The mirror law F^- reverses every realization, while F^0 assigns half of each absolute-shock probability to each sign. The three laws have mean zero and the same complete distribution of absolute shocks; consequently, all even moments are identical. Mirroring swaps upper and lower semivariance.

Households with terminal CRRA utility allocate financial wealth between a safe and a risky asset, subject to a fixed participation cost. Income innovations follow one of the three laws and are independent of the risky return. Financial-wealth-to-income states are the ten weighted deciles of the public 2019 Survey of Consumer Finances. Across risk-aversion values three through ten and ten wealth states, replacing the downside-oriented empirical law with its upside-oriented mirror never lowers the optimal risky share: 24 comparisons are positive and 56 are unchanged. At risk aversion six, the median difference is 0.25 percentage point. One 52-point difference occurs at a discrete participation boundary and is not representative of the intensive margin. This sealed exercise is a mechanism benchmark, not a fitted lifecycle model or evidence about the empirical coefficient. Its purpose is to illustrate that a successful conditional-orientation measure would have an economically coherent portfolio object.

3 Public Data and Frozen Designs

3.1 Measurement samples

The CPS ASEC supplies annual wage and salary income, stable person identifiers across adjacent surveys, industry, education, age, and survey weights. The raw extract covers interview years 2014–2023. I begin with workers aged 25–65 who have positive baseline wage earnings and match them to the next ASEC observation, allowing destination earnings to equal zero. Recorded age must rise by zero, one, or two years. The exact CPS-to-SIPP crosswalk yields 183,809 unique-person transitions in 72 candidate broad-industry-by-education-by-age cells. Education code 73, a high-school diploma in the IPUMS CPS coding, is included in the high-school group. Age is split at 45. Appendix C records the industry and education mapping, support rules, and sample waterfall.

The post-redesign SIPP is a public longitudinal survey with retrospective monthly earnings and annual household balance-sheet measures. For the independent measurement comparison, I match a positive monthly main-job earnings observation to the same calendar month 12 months later and allow the destination value to be zero. Destination-month job-or-business class and no-job-spell fields reveal whether a main-job zero occurs inside an employer, self-employment, or other arrangement and whether any no-job spell is recorded during that month. Repeated observations from one worker always remain in the same split half. Because overlapping monthly transitions may mechanically stabilize cell estimates, the primary SIPP reliability check also keeps one fixed nonoverlapping 12-month-apart transition per person. The repeated-month design is reported as a precision benchmark rather than treated as independent information.

Both surveys use their official person weights to estimate distributions. Effective support is measured by the Kish count. Reliability splits assign each person globally to one half; the CPS first-stage sensitivity instead resamples person clusters within signal cell and is described separately in Section 6. The baseline construction retains destination zeros and uses equation (3); positive-to-positive arc and log growth define separate conditional-wage objects. The broader estimator and support tournaments are in Appendix B.

3.2 Strict-post-measurement portfolio levels and transitions

The primary finance design estimates the group distribution using CPS transitions whose origin year is no later than 2017 and assigns that frozen ranking to SIPP portfolios from 2019 through 2024. The crosswalk uses 12 harmonized nonmilitary broad industries, three education groups, and two age bands. Because the surveys contain different people, no SIPP portfolio record enters the measured CPS distribution. CPS destination earnings extend through 2018, so every primary portfolio outcome strictly postdates the measurement window. A broader-window robustness check restores SIPP portfolios from 2018. Neither design makes group membership exogenous or proves that households perceived the assigned history.

Monthly SIPP balance-sheet observations are collapsed to person-years by averaging the available eligible records within each calendar year. Participation equals one when the observed stock-and-mutual-fund balance is positive. The primary share sets the observed balance to zero for nonparticipants and divides it by reported financial assets. Conditional-positive-asset and holder shares are reported secondary outcomes. SIPP does not observe the allocation inside retirement accounts, mutual funds can contain non-equity securities, and employer-stock compensation cannot be cleanly removed. The outcome is therefore labeled the *observed stock-and-mutual-fund asset share*, not total equity exposure. Appendix C.3 reports denominator support, ratio bounds, eligible-month coverage, and repeated-household-record diagnostics for the exact estimation frame.

The transition frame links consecutive annual observations and assigns the lagged early-CPS signal to the next year's balance state. It contains 35,112 person-year pairs on 23,170 people: 27,752 observations begin with zero holdings and 7,360 begin with positive holdings. The weighted nonholder-to-holder transition rate is 8.79%. Requiring the same signal cell in both years is a migration sensitivity, not part of the primary risk set. Sustained entry additionally requires positive holdings one year after the transition; nominal $1,000 and $5,000 crossings test whether a small positive balance alone drives the result.

Every regression uses SIPP person-year weights and household controls measured at the relevant date. Conditional inference clusters by person and assigned CPS signal cell; a cell-year aggregation and a CPS first-stage resampling sensitivity address the limited number of distinct signals. Appendix E separately applies the measurement logic to public subjective earnings distributions from the New York Fed Survey of Consumer Expectations.

4 What Public Earnings Cells Measure

4.1 The construction reversal

Figure 1 begins with the population choice. In a companion CPS construction audit with the same age-refined cell schema, raw Kelley reliability is 0.018 when arc growth is restricted to positive earnings at both endpoints and 0.735 when destination zeros are retained. The exact crosswalk used below gives 0.740. In SIPP, keeping one 12-month-apart transition per person gives reliability 0.147 for positive-to-positive arc growth and 0.602 when destination zeros are retained. Positive-to-positive log growth is similarly weak at 0.120. Repeated overlapping SIPP months are not responsible for the ordering: the zero-inclusive estimate is 0.622 with repeated months and 0.602 with one transition per person.

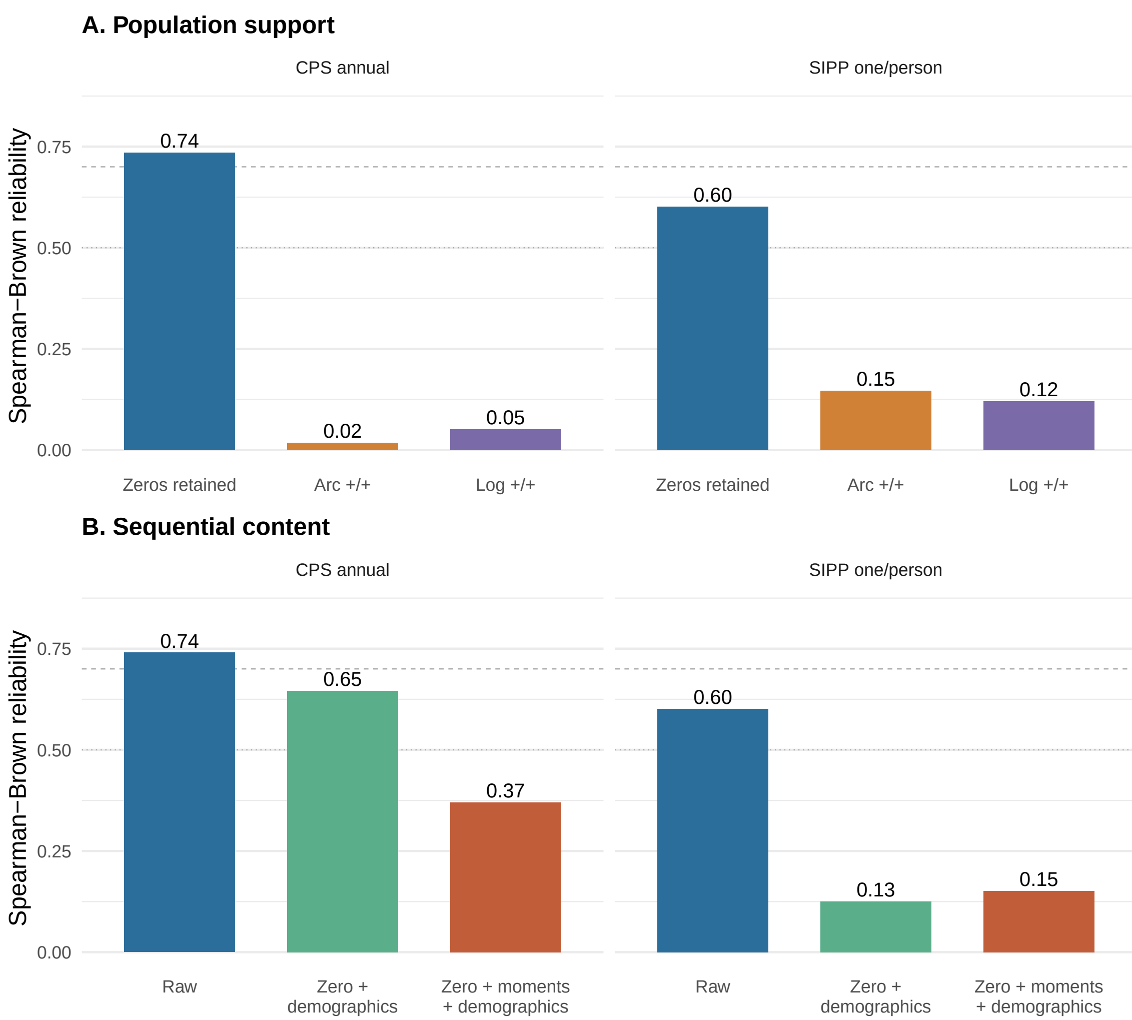

Figure 1: Where the reproducible downside ranking comes from

Notes. Panel A compares raw downside-oriented Kelley reliability when destination zeros are retained with positive-to-positive constructions. Panel B uses the pooled exact CPS crosswalk and matched SIPP: the second stage removes the linear zero share and demographic components, and the third also removes quadratic mean and dispersion. Reliability is the Spearman–Brown transform of the mean independent-person half-sample correlation across 30 deterministic splits. Each projection is re-estimated inside every half.

The construction reversal is economically interpretable. A transition to zero maps to the arc-growth lower bound and changes the unconditional earnings distribution. Conditioning it away estimates wage growth among workers with positive measured earnings at both endpoints, not the same risk object with fewer observations. The positive result is therefore that, within each survey, independent worker splits reproduce a downside-support ranking that includes transitions to zero measured earnings.

4.2 The economic anatomy of the zero endpoint

Table 1 opens the endpoint rather than assigning it an economic label. CPS wage-and-salary income and weeks worked refer to the same destination calendar year. Of 11,828 annual wage-zero destinations, 64.50% report no work, but 35.50% report some work and 26.02% report all 52 weeks. The two components are separately stable across independent people: Spearman–Brown reliability is 0.840 for no-work incidence and 0.940 for worked-without-wage incidence, compared with 0.940 for total zero incidence. The full-year-work-without-wage component has reliability 0.914. Thus the raw tail ranking is not just a measure of full-year nonwork; worked-without-wage incidence is at least as reproducible as total wage-zero incidence.

The SIPP endpoint is different again. In a strict out-of-time frame, one eligible 2018–2023 origin transition is selected uniformly per person without conditioning on destination earnings or job state; 914 of 25,310 selected workers enter zero monthly main-job earnings. Among those zeros, 64.47% have a contemporaneous self-employment arrangement and 34.62% an employer arrangement. For the former, the job-1 field can record zero business earnings rather than wage loss. Only 19 zeros, or a weighted 2.06%, contain any recorded no-job spell; the corresponding weighted share is 0.99% when all eligible overlapping transitions are retained. Arrangement class is mutually exclusive, while the within-month no-job flag can coexist with an arrangement. A numerical zero therefore carries economic content, but not one invariant event across the two public surveys.

Table 1: What the reproducible destination-zero endpoint contains

Endpoint or component	Weighted share (%)	Zero rows	CPS incidence reliability
Panel A. CPS annual wage-and-salary income			
Destination annual wage zero, all linked transitions	6.63	11,828	0.940
No work in destination calendar year	64.50		0.840
Worked in destination calendar year	35.50		0.940
Worked all 52 weeks	26.02		0.914
Panel B. SIPP monthly main-job earnings, one transition/person			
Destination monthly main-job zero, all selected transitions	3.51	914	
Self-employed work arrangement	64.47		
Employer work arrangement	34.62		
Any no-job spell recorded in destination month	2.06		

Overall zero rates use all positive-origin transitions; indented shares condition on a destination zero. Official person weights are used. CPS anatomy rates use 183,809 linked annual transitions across all 72 candidate cells; the reliability entries use unconditional component incidence among 181,946 transitions in the 65 supported cells. Each reliability entry applies the Spearman–Brown transform to the mean independent-person half-sample correlation over 30 splits. The strict SIPP anatomy sample contains 25,310 people and freezes CPS signal information through 2017 before SIPP destination outcomes in 2019–2024. SIPP arrangement and no-job indicators are contemporaneous with the destination month and refer to job 1 rather than total labor income. Arrangement classes are mutually exclusive; the no-job-spell flag can coexist with an arrangement within a month.

An out-of-time prediction audit is directionally consistent but not promotion-grade. Across 30 deterministic choices of one 2018–2023 origin transition per person, the early-CPS raw Kelley coefficient for later SIPP main-job-zero incidence is positive in 29, with median 0.297 percentage point and P10–P90 range 0.068–0.450, but no selection has a four-outcome-adjusted $q < 0.05$. The incremental coefficient is positive in 26 selections, with median 0.154 point and P10–P90 range -0.013–0.338; again none passes the adjusted threshold. Appendix Table 11 reports the complete selection audit. I therefore use the exercise as directional evidence about what the raw composite travels with, not as a new predictive rejection.

4.3 Sequential content and time stability

Table 2 asks what variation supports that ranking. The zero-transition share is highly reproducible: 0.940 in pooled CPS and 0.809 in the matched SIPP design. In pooled CPS, adding the zero share and demographic components changes reliability from 0.740 to 0.646; adding quadratic mean and dispersion terms lowers it to 0.370. The exact early signal later assigned to portfolios has reliability 0.687 raw and 0.397 after the complete projection. In matched SIPP, the corresponding sequence is 0.602, 0.125, and 0.151. Thus the survey-specific zero endpoint and lower moments account for reproducible CPS variation, while the fully incremental SIPP ranking remains weak under either adjustment.

Table 2: Sequential content and time stability of downside Kelley shape

Design	Supported records	Cells	Raw	+Zero/demog.	+Moments	Zero share
CPS pooled exact cross-walk	181,946	65	0.740	0.646	0.370	0.940
CPS early downstream signal	106,404	60	0.687	0.564	0.397	0.909
SIPP one/person	50,137	58	0.602	0.125	0.151	0.809
SIPP repeated months	953,809	60	0.622	0.178	0.070	0.831
Early-versus-late Pearson correlation						
CPS pooled exact cross-walk			0.631	0.591	0.230	0.921
SIPP one/person			0.414	0.153	−0.100	0.691
SIPP repeated months			0.279	0.066	0.054	0.614

Supported records sum observations in the displayed post-support cells; the underlying candidate samples before cell-support screening contain 183,809 pooled CPS transitions, 108,707 early CPS transitions, 51,152 matched SIPP transitions, and 967,563 overlapping SIPP transitions. Entries above the midrule apply the Spearman–Brown transform to the mean global-person split correlation. "+Zero/demog." removes the linear cell zero-transition share plus education and age components. "+Moments" also removes quadratic cell mean and standard deviation and is the exact projection used for the portfolio regressor. Pooled CPS uses Kish support of 200 plus 100 in each period; the early row uses Kish support of 200 in the early period and the same 60 cells assigned downstream. SIPP uses effective-person support of 100 pooled and 50 in each period. The matched SIPP transition is drawn once per person before any destination-positive restriction.

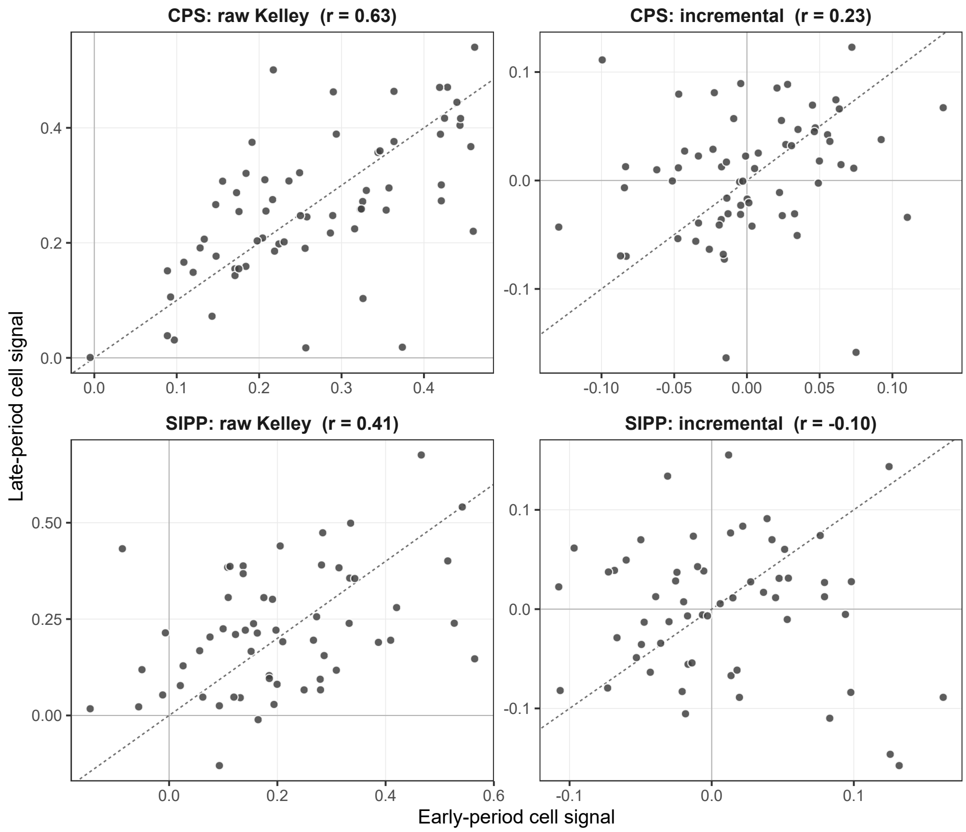

Figure 2: Within-survey time stability of the cell signal, cell by cell

Notes. Each point is one supported cell; the horizontal axis is the early-period signal and the vertical axis the late-period signal, with the dashed line marking equality and Pearson correlations in the panel titles. Raw Kelley uses the zero-inclusive construction; the incremental signal is the residual from equation (5), re-estimated within each period. CPS panels use the exact crosswalk (65 cells); SIPP panels use one transition per person (58 cells).

Within-survey time stability is a separate gate. The CPS early-versus-late correlation is 0.631 for raw Kelley and 0.230 for the full incremental signal. In the matched SIPP design it falls from 0.414 to -0.100. Rank correlations are 0.634 and 0.326 in CPS, and 0.396 and -0.025 in matched SIPP. A larger nominal transition count from overlapping months does not improve the incremental ranking. Figure 2 plots the cell-level signals behind these within-survey correlations: the raw ranking retains a visible diagonal in each survey, while the incremental panels are close to exchangeable across periods. Appendix Table 10 shows the same hierarchy under alternative splits and transition selections, complete-year SIPP main-job earnings, and direct common-year comparison: pooled raw and incremental

CPS–SIPP correlations are 0.482 and 0.053.

4.4 What is and is not measured

The data support two conclusions and rule out a third. First, a survey-specific downside-support environment is measurable: independent people reproduce the zero-inclusive group ranking separately within CPS and SIPP. Second, the stable ranking jointly reflects the zero endpoint, mean growth, dispersion, and whatever shape variation remains. Third, reliability alone cannot apportion that joint object or establish the counterfactual comparison needed for "shape holding the other distributional features fixed."

Equation (5) is deliberately a conditional diagnostic, not a declaration that correlated variation is spurious. Its lower reliability is also a limit on the portfolio application: attenuation can make the measured incremental coefficient small even when a latent shape contrast matters. The estimator tournament, support frontier, and construction details are reported in Appendix B.

5 Consequences for Household-Finance Inference

5.1 Early CPS signal and later SIPP portfolio decisions

The primary application uses 60 early-CPS signal cells assigned to 90,110 SIPP worker-years on 52,115 people observed from 2019 through 2024. Table 3 separates standing portfolio levels from one-year transitions. Raw Kelley is assigned without distributional controls. Incremental Kelley is the standardized residual from equation (5) and enters with the corresponding zero-share, mean, and dispersion controls. The raw-plus-controls coefficient has the same t statistic as the incremental coefficient by Frisch–Waugh–Lovell; it is a rescaling of the same conditional contrast and remains in the replication output rather than being marketed as a separate test.

Table 4 opens the entry coefficient into its two economic sides. The lower distance is $L_c = q_{0.50,c} - q_{0.10,c}$ and the upper distance is $U_c = q_{0.90,c} - q_{0.50,c}$. Panel A enters their standardized early-CPS values jointly. Panel B is an exact rotation into the total span $S_c = L_c + U_c$ and signed composition $D_c = L_c - U_c$; the two raw column spaces produce fitted values that agree to 1.1×10^{-8}, although each displayed component is standardized separately.

Table 3: Early CPS signals and strictly later portfolio levels and transitions

	Raw Kelley			Incremental Kelley			
Outcome	Coef.	S.E.	BH q	Coef.	S.E.	BH q	Obs.
Panel A. Standing outcomes in 2019–2024							
Participation	0.38	0.73	0.880	-0.92^{*}	0.48	0.081	90,110
Share, positive financial assets	0.061	0.185	0.880	-0.261^{***}	0.096	0.035	84,161
Share, all workers	0.080	0.185	0.880	-0.234^{**}	0.097	0.039	90,110
Share, current holders	-0.097	0.638	0.880	-0.153	0.424	0.721	17,578
Panel B. Consecutive-year transitions; both observations in 2019–2024							
Signed participation change	-0.335	0.428	0.437	-0.731^{***}	0.260	0.010	35,112
Nonholder to positive holdings	-1.093^{**}	0.526	0.126	-1.047^{***}	0.272	0.001	27,752
Holder to zero holdings	-2.236	1.538	0.227	-0.813	1.152	0.483	7,360

Coefficients and standard errors are percentage points per one cross-cell standard deviation. SIPP observes combined stock-and-mutual-fund balances, not total equity. Every incremental specification includes the linear CPS zero share and quadratic mean growth and dispersion. Panel A also includes current household controls and fixed effects. Panel B assigns the lagged signal and distributional controls, uses lagged household controls and lagged industry, education, and age-band fixed effects plus current-year effects, and weights by current-year WPFINWGT. Standard errors are two-way clustered by person and the assigned CPS signal cell. Superscripts denote unadjusted two-sided significance based on the standard error in the same row: $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$. Stars omit multiple-testing adjustment and CPS first-stage uncertainty. Benjamini–Hochberg q-values adjust the four standing outcomes within each Panel-A specification and the three transition outcomes within each Panel-B specification (Benjamini and Hochberg, 1995). Panel-B entry and exit use their lagged-state risk sets; the signed-change row includes lagged participation. Every outcome strictly postdates CPS destination earnings through 2018. Assignment is endogenous.

Panel C adds the span directly to the paper's scale-free Kelley contrast. Every panel uses the same 27,752 lagged-nonholder transitions, controls, fixed effects, weights, and clustering as the primary entry row. A separate 200-draw rebuild carries CPS quantile estimation, residualization, and standardization uncertainty through every displayed coefficient while holding the exact SIPP risk set fixed.

Table 4: Both earnings tails and transitions into positive reported holdings

Early-CPS distributional component	Coef.	Cond. S.E.	CPS S.D.	Plug-in S.E.	p
Panel A. Lower and upper distances entered jointly					
Lower distance, $q_{0.50} - q_{0.10}$	-2.712^{***}	0.646	0.562	0.857	0.002
Upper distance, $q_{0.90} - q_{0.50}$	1.917^{*}	0.846	0.619	1.049	0.068
Panel B. Exact rotation: span and signed composition					
Total tail span, $L + U$	2.721	1.850	1.331	2.279	0.232
Downside composition, $L - U$	-4.456^{**}	1.530	0.986	1.820	0.014
Panel C. Scale-free orientation conditional on span					
Total tail span, $L + U$	−0.755	0.989	0.770	1.254	0.547
Incremental downside Kelley orientation	-0.870^{*}	0.376	0.293	0.477	0.068

The outcome is a one-year transition from zero to positive reported stock-and-mutual-fund holdings. Coefficients are percentage points per one early-CPS cross-cell standard deviation of the named regressor. Every specification includes the lagged CPS zero share, quadratic mean and standard deviation, lagged household controls, current-year and lagged industry, education, and age-band fixed effects, current SIPP weights, and two-way clustering by SIPP person and lagged CPS cell. Each panel standardizes its named components separately. The lower and upper coefficients are conditional on each other; the span and composition coefficients are likewise joint.

"CPS S.D." is the across-draw standard deviation from 200 deterministic CPS person-within-cell resamples. Each draw rebuilds the early-cell quantiles, Kish support, zero share, mean, standard deviation, residualization, and cross-cell standardization. All 200 draws retain the exact 27,752-observation, 19,300-person, 60-cell SIPP risk set; 59–60 of the locked cells clear rebuilt support. Plug-in S.E.^{2} = conditional S.E.^{2}+ CPS S.D.^{2}, and the final column uses a two-sided normal reference. Superscripts use that plug-in p-value: $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$. The plug-in joint p-values are 0.0022 in Panel A, 0.0023 in Panel B, and 0.0106 in Panel C. A draw-specific raw-scale test rejects equality of the lower and upper marginal effects ($p = 0.0135$). These are generated-regressor sensitivities, not joint CPS–SIPP bootstraps. Assignment is endogenous, and historical cell distributions are not elicited beliefs.

The level results reproduce the earlier measurement warning. Incremental Kelley is negative for all four outcomes, but the all-worker share's conditional $p = 0.019$ and four-outcome $q = 0.039$ become fragile after rebuilding the CPS first stage. The broader 2018–2024 level window remains in Appendix Table 17. Household-level clustering and alternative decision-unit restrictions preserve the negative conditional direction (Appendix Table 14).

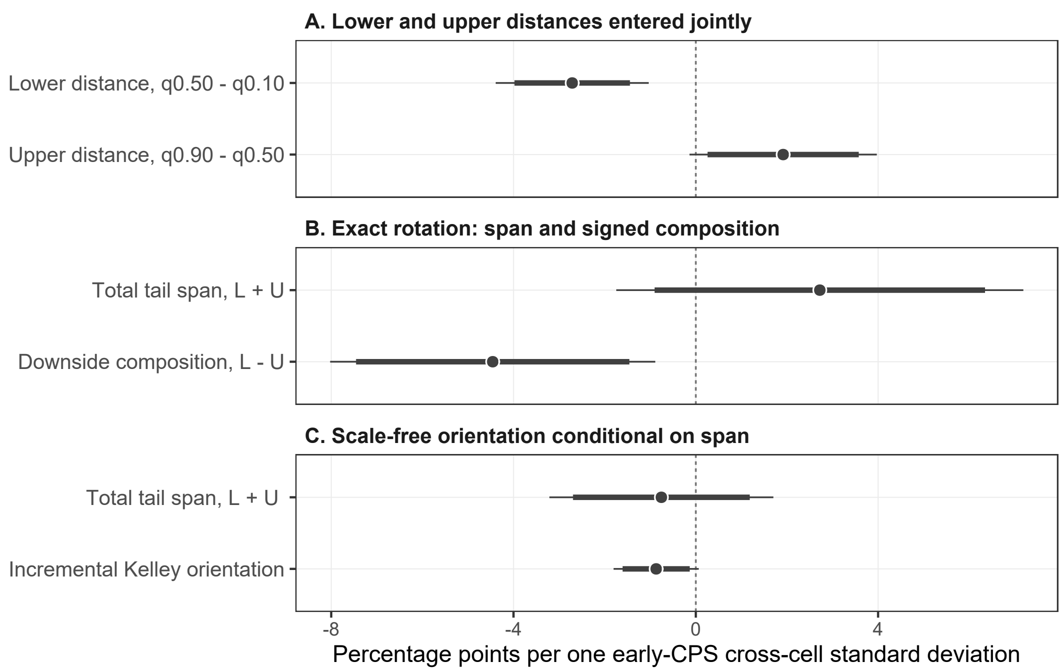

Figure 3: The entry-margin tail decomposition at a glance

Notes. Points are the coefficients of Table 4, in percentage points per one early-CPS cross-cell standard deviation, for the same 27,752-transition lagged-nonholder risk set, controls, fixed effects, weights, and clustering. Thick intervals are ±1.96 conditional two-way-clustered standard errors; thin intervals are ±1.96 plug-in standard errors adding the 200-draw CPS first-stage component. The dashed line marks zero. Each panel standardizes its named components separately, and the coefficients within a panel are conditional on each other.

The transition result is different. Among lagged nonholders, incremental Kelley predicts a 1.047-percentage-point lower probability of positive observed holdings one year later (SE 0.272 point; $p = 0.00030$; three-outcome $q = 0.00089$). Relative to the weighted 8.79% entry rate, the point estimate is 11.9%. Raw Kelley gives a similar point estimate but less precision: its unadjusted $p = 0.042$ does not survive the raw three-outcome adjustment ($q = 0.126$). Table 4 supplies the economic content that a one-index regression cannot: conditional on each other, lower- and upper-tail distances have opposite signs, while total tail span is imprecisely estimated. Rebuilding the CPS first stage leaves the joint lower–upper test at $p = 0.0022$ and the signed composition contrast at $p = 0.014$; the upper coefficient alone is suggestive at $p = 0.068$. The conventional scale check is assigned cell standard deviation. In this locked entry frame, its standardized linear term, evaluated at mean standardized SD, is -0.262 point (SE 1.120; $p = 0.816$) when entered with its square and the same controls; adding Kelley orientation or separate lower and upper distances does not reveal a positive scale coefficient. The result therefore does not reproduce the earlier claim that a positive variance association is an upper-tail artifact. It supports the narrower statement that the two tails have opposite conditional associations, whereas total span and assigned SD are not precisely estimated. Signed participation change is also negative for incremental Kelley; exit is too imprecisely estimated to characterize. The entry outcome is not first-ever account entry, and the table does not establish that effects operate exclusively through the extensive margin. Figure 3 plots the six decomposition coefficients with both interval widths; the opposite-signed tail pair and the imprecise span are visible at a glance.

The result is not driven by observed cell migration or token balances. Requiring the same signal cell in both years gives -0.958 point ($p = 0.0039$; three-outcome $q = 0.0118$). Omitting each outcome year in turn yields five negative coefficients from -1.203 to -0.881 point. Sustained entry—zero holdings at $t-1$ and positive holdings at both t and $t+1$—is -0.697 point (SE 0.293; six-test $q = 0.041$). Nominal crossings above \$1,000 and \$5,000 are -1.104 and -0.453 point, with six-test $q < 0.001$ and $q = 0.045$. These threshold checks are not inflation-adjusted. The stable-household frame gives -1.057 point, while selecting one highest-earning eligible worker per household transition gives -0.869 point; both are conditional decision-unit checks, not household-representative estimates. Appendix Table 15 reports the displayed robustness family. Post hoc influence and specification checks are also stable: every leave-one-cell and leave-one-industry coefficient is negative, cell-year aggregation gives −0.985 point, and adding quadratic or cubic zero-incidence controls gives −0.982 and −0.901 point (Appendix Table 16).

This family was isolated through exploratory post-rejection screening. The reported three-transition and six-robustness-test adjustments describe the displayed families; they do not retroactively adjust for every outcome, threshold, or specification considered during the project's history. Section 6 therefore emphasizes the independently rebuilt first stage and the full pattern rather than treating a small conditional p-value as stand-alone confirmation.

The negative direction is consistent with background-risk intuition, but it is not behavioral identification. Risk-tolerant workers can sort into both volatile jobs and risky portfolios (Ranish, 2013; d'Astous and Shore, 2024); historical group distributions need not equal household beliefs (Guiso, Jappelli, and Terlizzese, 1996; Wang, 2023; Caplin et al., 2026; Arellano et al., 2026); and the observed balance omits retirement-account allocation. The table establishes that the interpretation of the same public-data assignment depends on the distributional contrast, the portfolio decision margin, and the uncertainty propagated.

5.2 Precision and limits

Incremental Kelley reliability is 0.397 for the exact 60-cell early signal assigned to portfolios, 0.370 in pooled CPS, and 0.151 in matched SIPP. These values make attenuation a first-order concern. A coefficient on the measured signal is not an estimate of the effect of latent tail orientation corrected for measurement error, and its confidence interval is not an equivalence bound for such an effect. The effective identifying variation is 60 early-CPS cells, not tens of thousands of independent risk environments. A normal-approximation audit confirms that the standing level family is underpowered once first-stage sensitivity is included: at 5% size and 80% power, the measured-signal minimum detectable effects are 0.409 point for the all-worker share and 1.79 points for participation under the CPS plug-in. Appendix Table 21 reports the full level calculation. The transition result survives the corresponding first-stage rebuilding, but that does not repair endogenous assignment or convert the measured signal into latent risk.

Pooled estimates, the broader construction tournament, alternative centered estimators, and the SCE construct/timing audit are relegated to the Appendix. They do not supply causal variation and do not override the hierarchy in Table 3.

6 First-Stage Uncertainty and Additional Construct Checks

The CPS signal is estimated independently of SIPP, but independence does not make it known. Ranish (2013) provides a direct household-finance precedent for bootstrapping an estimated labor-risk object through a portfolio regression. The primary resampling exercise here draws CPS workers within signal cells, rebuilds the weighted distribution, recomputes support, re-estimates the sequential projection in equation (5), restandardizes the signal, remaps it to both SIPP observations, reconstructs the transition risk sets, and refits the three outcome regressions. Variation across the rebuilt-CPS estimates isolates the component induced by rebuilding the independent measurement sample. I report a normal plug-in sensitivity that adds this CPS component to the conditional two-way-clustered variance. Because SIPP is held fixed, this is not joint inference. It does not correct attenuation or deliver inference for an error-free latent signal. Table 5 reports the complete transition family.

Table 5: Transition inference after rebuilding the CPS first stage

Outcome	Est.	Cond. S.E.	Cond. q	CPS S.D.	Plug-in S.E.	p	q
Signed change	-0.731^{***}	0.260	0.010	0.250	0.360	0.043	0.064
Zero to positive holdings	-1.047^{***}	0.272	0.001	0.258	0.375	0.005	0.016
Positive to zero holdings	-0.813	1.152	0.483	0.845	1.429	0.569	0.569

Entries are percentage points per one early-CPS cross-cell standard deviation. The calculation uses 200 deterministic CPS person-within-cell resamples and rebuilds weighted quantiles, Kish support selection, zero-share and quadratic mean/dispersion residualization, standardization, assignment to both SIPP observations, and the transition sample. Every draw produces all three estimates; supported cells range from 59 to 61. SIPP is held fixed, while conditional two-way clustering supplies the downstream sampling component. The plug-in calculation is $\text{S.E.}^2 = \text{conditional S.E.}^2 + \text{CPS coefficient S.D.}^2$ and uses a two-sided standard-normal reference. It is a sensitivity, not a joint nonparametric bootstrap or inference for an error-free latent signal. Conditional and plug-in q-values each adjust exactly the three displayed transitions. Superscripts refer only to unadjusted conditional tests: $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

The CPS-first-stage sensitivity separates the transition result from the standing levels. For entry, the across-draw CPS standard deviation is 0.258 point and the plug-in standard error is 0.375 point. The resulting $p = 0.0052$ and three-outcome $q = 0.0157$ remain below 5%; replacing the normal reference with t_{59}, the minimum signal-cell cluster degrees of freedom, gives $p = 0.0070$ and $q = 0.0211$. Exactly 198 of the 200 fixed-SIPP CPS remeasurement coefficients are negative. Their mean is -0.650 point, compared with the observed-signal estimate of -1.047 points; the mean-minus-point displacement is 0.397 point. Their 2.5th and 97.5th percentiles are -1.125 and -0.138 points. Those percentiles summarize remeasurement draws and are not confidence intervals or a probability statement. The displacement is why the reported calculation remains a variance sensitivity rather than being presented as a conventional bootstrap distribution. Signed change passes only the 10% plug-in family threshold, and exit remains imprecise. For comparison, the standing all-worker share has plug-in $p = 0.109$ and standing participation has $p = 0.153$.

The tail decomposition uses a second exact-risk-set rebuild. Every draw reconstructs $q_{0.10}$, $q_{0.50}$, $q_{0.90}$, total span, signed composition, and incremental Kelley, while all 27,752 SIPP transitions and 60 assigned cells remain fixed; 59–60 locked cells clear the draw-specific support threshold. The CPS components of the lower- and upper-distance standard errors are 0.562 and 0.619 percentage point, yielding plug-in p-values of 0.0015 and 0.068 and a joint $p = 0.0022$. For the exact span–composition rotation, the CPS components are 1.331 and 0.986 points; signed composition remains significant at $p = 0.014$, whereas span has $p = 0.232$. Transforming every draw back to raw arc-growth units rejects equality of the lower and upper marginal effects at $p = 0.0135$. These calculations propagate estimation of the assigned historical distributions; they do not identify household perceptions or causal responses.

Alternative estimators show the same hierarchy. Raw Spearman–Brown reliability is 0.768 for Bowley shape, 0.794 for moment skewness, and 0.685 for L-skewness; after the full sequential projection the corresponding values are -0.051, 0.233, and 0.178. In the retained broader 2018–2024 tournament, none of their four-outcome portfolio families survives false-discovery-rate adjustment (minimum $q = 0.244$). The Appendix reports the estimator-family summary, and the replication output retains the complete coefficient grid rather than selecting the closest result. It also carries the Survey of Consumer Expectations as a separate construct-and-timing check. The SCE fixed-threshold imbalance is mechanically sensitive to expected growth, and the portfolio module asks retrospectively about the preceding twelve months. In the exploratory timing grid of Appendix E.3, the pre-module coefficient is positive and larger in the small realized-income-fall subgroup, while the post-module estimate has the opposite sign; a window that mixes reports from both sides of the module can therefore return either sign. These exercises test the measurement logic in distinct data; they do not supply behavioral identification. They nonetheless provide a useful scale benchmark: in the pre-module, officially weighted Household Finance sample, a one-standard-deviation increase in perceived earnings dispersion is associated with a 3.63–4.73 percentage-point lower standing defined-contribution equity share across the three shape controls (all $p < 0.001$; Appendix Table 19). Because this outcome is an accumulated portfolio stock and beliefs are not assigned, the association is descriptive rather than causal.

Appendix Table 20 also reconstructs respondent-level centered semimoments. The upside-minus-downside semideviation gap is positively associated with reported riskier reallocation ($p = 0.036$), but its primary-family $q = 0.054$, expanded-support $q = 0.169$, null normalized orientation, and negligible out-of-sample increment keep it exploratory. The separate lower and upper components have the same opposite-sign pattern as the assigned-risk result, but correlate 0.95; the portfolio action is retrospective, and the standing-share shape coefficient is null. This is directional evidence about individual perceived tail magnitude, not evidence that perceptions dominate assigned histories.

7 Implications for Generated Earnings-Risk Measures

The evidence motivates four disclosures for studies that assign a group-estimated risk statistic to households. First, name the population and support: destination zeros, positive-to-positive conditioning, and the growth transformation define different economic objects. Second, split at the person level and show a support frontier, rank reliability, and time stability; a household sample size cannot substitute for the number and precision of the measured groups. Third, state the conditional interpretation and show the statistic's sequential relation with destination-zero incidence, mean, and dispersion. Correlation does not make joint downside risk unimportant, but it limits a shape-specific claim. Fourth, assess uncertainty induced by the measurement step, including support selection, nonlinear quantiles, residualization, and standardization, and state whether the calculation is a sensitivity or fully joint inference.

These disclosures do not prescribe Kelley shape or arc growth as universally preferred estimators. Nor do they turn an independent two-sample assignment into causal identification. They make the empirical object and the supported interpretation auditable. Appendix Table 22 provides a compact implementation checklist.

8 Conclusion

The main result goes beyond downside risk. Among lagged nonholders, assigned lower and upper earnings-tail distances predict a subsequent transition to positive observed holdings with opposite signs. Rotating those two distances into total span and signed composition leaves the span imprecisely estimated and orientation significant. A variance or interquantile range can therefore conceal economically distinct earnings disasters and opportunities. This result concerns reported stock-and-mutual-fund balances, not verified purchases, first-ever account opening, total equity, or a causal response.

Measurement determines whether that interpretation is warranted. Public surveys reliably distinguish some workers' downside-support environments, and retaining transitions to zero is central to that fact. Yet CPS annual wage zeros mix full-year nonwork with highly reproducible worked-without-wage states, while SIPP monthly main-job zeros mostly record a contemporaneous work arrangement. Reliability alone therefore does not identify which feature of the unconditional distribution drives a portfolio association or establish a common job-loss object across surveys.

The corrected CPS and SIPP evidence separates that positive measurement fact from a narrower limit. Once the zero-transition probability, realized cell mean growth, and dispersion are made explicit, the incremental tail ranking is substantially less stable across people and time within each survey. That limitation is consequential for standing holdings: the all-worker share association loses conventional significance when the CPS first stage is rebuilt. The lagged nonholder-to-positive-holdings transition, however, remains negative under same-cell, leave-one-year-out, sustained-entry, balance-threshold, and decision-unit checks, and survives the normal first-stage plug-in sensitivity. Exit is too imprecisely estimated to characterize. This contrast concerns standing portfolio levels versus a transition to positive reported holdings; it does not identify a fixed cost, a first-ever account opening, or a causal response.

Nothing here shows that asymmetric labor-income risk is behaviorally irrelevant. The matched-even-moments benchmark illustrates why swapping upside and downside can change portfolios even when every even moment is held fixed, and richer administrative designs measure persistent and state-contingent risks that public survey cells cannot recover. Because the primary design neither elicits beliefs nor observes household information sets, it does not distinguish statistical learning from experience-based learning or establish that perceived risk dominates assigned histories. The exploratory individual SCE exercise asks that separate question but does not clear its full robustness and multiplicity gates. The bounded conclusion is that both tails have opposite conditional associations with a later portfolio transition, while reproducibility, shape-specific variation, perception, and downstream uncertainty remain separate empirical achievements.

Data Availability

The CPS ASEC extract uses IPUMS CPS microdata (Flood et al., 2025); its license prevents redistribution. Public SIPP files are at https://www.census.gov/programs-surveys/sipp/data/datasets.html (U.S. Census Bureau, 2025), SCE files at https://www.newyorkfed.org/microeconomics/databank (Federal Reserve Bank of New York, 2026), and the 2019 SCF benchmark files at https://www.federalreserve.gov/econres/scfindex.htm (Board of Governors of the Federal Reserve System, 2019).

A Estimator Definitions and Orientation

Let q_p denote a survey-weighted quantile of growth g, and let μ and σ denote its survey-weighted mean and standard deviation. Every estimator is oriented so that a positive value denotes a more adverse downside distribution.

Table 6: Distributional estimators

Estimator	Definition	Interpretation
Upper tail	$U = q_{.90} - q_{.50}$	Magnitude above the median
Lower tail	$L = q_{.50} - q_{.10}$	Magnitude below the median
90–10 spread	$S = L + U = q_{.90} - q_{.10}$	Interquantile scale
Tail difference	$D = L - U$	Unscaled downside-minus-upside length
Kelley shape	$(L - U)/(L + U)$	Tail difference normalized by the 90–10 spread
Bowley shape	$[(q_{.50} - q_{.25}) - (q_{.75} - q_{.50})]/(q_{.75} - q_{.25})$	Central-quantile analogue of Kelley shape
Moment skewness	$-\mathbb{E}[((g - \mu)/\sigma)^3]$	Conventional skewness with sign reversed
Tail probability difference	$\Pr(g \leq \mu - \sigma) - \Pr(g \geq \mu + \sigma)$	Downside-minus-upside mass beyond one SD
Conditional tail balance	Tail probability difference divided by its probability sum	Direction conditional on being in either tail
L-skewness	$-\lambda_3/\lambda_2$ from probability-weighted L-moments (Hosking, 1990)	Robust full-distribution orientation
Incremental estimator	Residual from equation (5), applied to the named estimator	Cell ranking incremental to zero share, location, and scale

Weighted quantiles sort observations by g, cumulate survey weights, and select the first value whose cumulative weight share meets the requested probability. L-moments are approximated on a probability grid from 0.005 through 0.995. Weighted variance divides by the sum of weights because the target is the weighted empirical distribution in the cell.

For Kelley shape, $K = 1$ is the limiting case in which the 90–10 span lies entirely below the median, while $K = -1$ is the limiting upside case. A point mass at $g = -2$ can move the lower quantile to the support bound. This changes the unconditional distribution's shape; it is not treated as an outlier correction.

A.1 Proofs of the measurement lemmas

Proof of Lemma 1. Because a lies below the support of F, the mixture distribution is zero below a, jumps to p at a, and equals $p + (1-p)F(x)$ above a. The generalized inverse is therefore a for $\tau \leq p$. For $\tau > p$, solving $p + (1-p)F(x) \geq \tau$ gives the stated transformed quantile of F.

The mean follows by iterated expectations. Conditioning on the mixture component and applying the law of total variance gives

$$\mathrm{Var}(G_p) = (1-p)\sigma_F^2 + p(1-p)(\mu_F - a)^2.$$

Changing p therefore changes the quantile map and lower moments while leaving the conditional distribution F untouched. The formula imposes no global monotonicity on Kelley shape without additional restrictions on F. □

Proof of Lemma 2. Independence implies that the covariance between the two half-sample estimates is $\mathrm{Var}(a'Z_c + s_c)$. Each half-sample variance is $\mathrm{Var}(a'Z_c + s_c) + \sigma_e^2$. Dividing covariance by the common variance gives the displayed correlation. Setting $\mathrm{Var}(s_c) = 0$ and taking σ_e^2 to zero establishes the final statement. □

A.2 Why arc growth is the baseline

For nonnegative earnings y_0 and y_1, with $y_0 > 0$, arc growth is defined at $y_1 = 0$ and equals -2. Writing $y_1 = y_0(1+x)$,

$$g^{\mathrm{arc}} = \frac{2x}{2+x} = x + O(x^2), \qquad \log(y_1/y_0) = x + O(x^2).$$

Unlike $\log(y+1)$ growth, it is invariant to the reporting unit. Unlike positive-to-positive log growth, it retains a transition out of positive earnings. Its bounded support and lower-bound mass are disclosed features of the estimand.

B Reliability Design and Additional Results

B.1 Person-level splitting

Each repetition draws one half indicator for every unique person, attaches it to all records of that person, rebuilds weighted cell distributions independently in each half, re-estimates any projection separately within each half, and correlates the common eligible cells. Eligibility is frozen from the full sample; each half must also clear its fixed support floor. The same split assignment is used across paired construction comparisons.

The main sequential audit uses 30 deterministic person splits. The alternative-estimator audit uses 50. A row split is not valid in SIPP because one worker can contribute repeated monthly year-apart transitions; the matched audit instead draws one eligible transition per person before applying destination-positive restrictions.

B.2 Population-support construction

Table 7: Positive-to-positive and zero-inclusive Kelley reliability

Data design	Growth construction	Cells	Raw reliability
CPS age-refined cells	Arc, zeros retained	60	0.735
CPS age-refined cells	Arc, positive-to-positive	60	0.018
CPS age-refined cells	Log, positive-to-positive	60	0.051
SIPP one transition/person	Arc, zeros retained	58	0.602
SIPP one transition/person	Log$(y+1)$, zeros retained	58	0.630
SIPP one transition/person	Arc, positive-to-positive	58	0.147
SIPP one transition/person	Log, positive-to-positive	58	0.120

Entries are Spearman–Brown transforms of mean independent-person half-sample correlations. The CPS construction audit uses the corrected IPUMS education mapping. The SIPP comparison selects one year-apart transition per person before construction-specific destination restrictions and uses identical split seeds in every row.

B.3 Support frontier

At support 400, early-versus-late Pearson correlations in the broader CPS tournament are 0.731 for the upper tail, 0.794 for the lower tail, 0.840 for the spread, 0.603 for raw Kelley,

Table 8: Selected CPS reliability-frontier values

Effective support floor	Common cells	90–10 spread	Raw Kelley	Mean/SD residual Kelley
100	63	0.922	0.728	0.379
200	60	0.929	0.746	0.334
400	58	0.921	0.741	0.368
800	51	0.937	0.791	0.297
1,200	34	0.951	0.839	0.344

General CPS construction tournament, 50 global person splits. This robustness frontier predates the sequential zero-share projection and therefore labels its final column explicitly as the mean/SD residual. Main Table 2 uses the exact portfolio crosswalk and the full sequential projection.

and 0.109 for mean/SD-residual Kelley. Rank correlations are 0.648, 0.785, 0.790, 0.622, and 0.225. Time stability and split reliability are distinct: the former changes both people and calendar periods, while the latter changes people holding the period fixed.

B.4 Alternative centered estimators

Table 9: Estimator robustness in the corrected exact crosswalk

Estimator	Raw rel.	Incr. rel.	Raw early/late	Incr. early/late	Min. 2018–24 BH q
Kelley	0.740	0.370	0.631	0.230	0.058
Bowley	0.768	−0.051	0.625	0.249	0.824
Moment skewness	0.794	0.233	0.733	0.290	0.244
L-skewness	0.685	0.178	0.597	0.050	0.381

Reliability is the Spearman–Brown transform of the mean global-person split correlation. Incremental signals remove the linear zero share, quadratic mean and standard deviation, and education and age components, re-estimated inside every split. Early/late entries are Pearson correlations. The last column retains the broader 2018–2024 estimator tournament and is the smallest conditional false-discovery-rate q across participation, positive-financial-assets share, all-worker share, and holder share. The Kelley minimum in that broader window does not pass a 5% family-adjusted threshold and its corresponding positive-assets plug-in first-stage-uncertainty p-value is 0.115. The strict-post Kelley family appears in the main text; alternative estimators were not selected for a new strict-window tournament. The replication output contains the complete estimator-by-outcome coefficient grid.

For Bowley, moment, and L-skewness, the raw-controlled and residualized parameterizations' p-values agree to numerical tolerance, as required by Frisch–Waugh–Lovell. None of the twelve alternative-estimator portfolio coefficients survives within-estimator family adjustment.

The negative Bowley reliability is not interpreted as a negative population reliability; it indicates that no stable incremental component is detected.

B.5 Assignment, transition-selection, and cross-survey sensitivities

Table 10 separates sensitivity to analyst-generated person splits and transition selections from formal sampling uncertainty. It also replaces the overlapping-month SIPP object with complete calendar-year main-job earnings and directly compares CPS and SIPP cells over common origin years.

Table 10: Reliability and cross-survey sensitivity ranges

Design	Raw Kelley	Incremental Kelley	Zero share
Panel A. Spearman–Brown reliability: center [P10, P90]			
CPS exact crosswalk, 30 person splits	0.740 [0.679, 0.802]	0.370 [0.109, 0.549]	0.940 [0.916, 0.956]
SIPP, 30 one-transition selections	0.566 [0.501, 0.630]	0.092 [−0.072, 0.239]	0.829 [0.814, 0.851]
SIPP, complete-year main-job earnings	0.442 [0.213, 0.607]	0.128 [−0.203, 0.369]	0.774 [0.703, 0.859]
Panel B. Direct common-year CPS–SIPP Pearson correlation: mean [P10, P90]			
Pooled common-year cells	0.482 [0.440, 0.522]	0.053 [−0.025, 0.137]	0.619 [0.598, 0.633]

Brackets are analyst-generated sensitivity ranges, not confidence intervals. The CPS row varies the 30 frozen person-half assignments; its center applies the Spearman–Brown transform to the mean half-sample correlation. The second row reports the median and P10–P90 range across 30 deterministic choices of one monthly 12-month-apart transition per person, with 30 person splits per choice. The complete-year row sums twelve finite nonnegative decoded monthly main-job earnings observations, uses December origin characteristics, selects one adjacent-year transition per person, and reports the range across 30 person splits; it contains 40,635 people in 57 supported cells. Panel B compares 58 common cells over the eight origin years shared by both surveys and varies the SIPP transition selection 30 times. CPS measures total annual wage-and-salary earnings, whereas SIPP measures main-job earnings; the direct correlation is therefore a comparability sensitivity, not validation against an identical construct. Incremental Kelley is residualized on the zero share, quadratic mean and standard deviation, education, and age within each design.

Table 11: Strict out-of-time prediction across SIPP transition selections

Later SIPP destination outcome	Early CPS signal	Median [P10, P90]	Positive / $q < 0.05$
Any monthly main-job zero	Raw Kelley	0.297 [0.068, 0.450]	29 / 0
Any monthly main-job zero	Incremental Kelley	0.154 [−0.013, 0.338]	26 / 0
Any monthly main-job zero	Destination-zero share	0.965 [0.104, 1.662]	28 / 0
Employer-arrangement zero	Raw Kelley	0.200 [0.081, 0.328]	30 / 1
Self-employment-arrangement zero	Raw Kelley	0.059 [−0.045, 0.165]	22 / 0
Zero with any no-job spell	Raw Kelley	−0.006 [−0.043, 0.057]	14 / 0

Entries are percentage-point coefficients per one early-CPS signal standard deviation across 30 deterministic choices of one eligible 2018–2023 SIPP origin transition per person, selected without conditioning on destination state. Destination outcomes occur in 2019–2024, after the final destination year entering the CPS signal. Each cell-year regression uses 360 cell-years in 60 signal cells, official person weights, baseline earnings and demographics, destination-year and additive industry, education, and age-band fixed effects, and standard errors clustered by signal cell. The incremental specification also includes the CPS zero share and quadratic mean and dispersion. The final column counts positive coefficients and selections whose four-outcome Benjamini–Hochberg q-value is below 0.05. Brackets are transition-selection sensitivity ranges, not sampling confidence intervals. SIPP arrangement classes are mutually exclusive; the no-job-spell flag can coexist with an arrangement within a month.

C CPS-to-SIPP Crosswalk and Portfolio Construction

C.1 Exact public industry mapping

Table 12: CPS IND1990 to SIPP broad-industry crosswalk

SIPP code	Sector	CPS IND1990 support
1	Agriculture	10–32
2	Mining	40–50
3	Construction	60
4	Manufacturing	100–392
5	Trade	500–691
6	Transportation and utilities	400–432
7	Information	440–472
8	Finance, insurance, and real estate	700–712
9	Professional and business services	721–760 and 861–893
10	Education and health	800–860
11	Leisure and hospitality	761–791
12	Other services	No clean IND1990 analogue; unmatched
13	Public administration	900–932
14	Military	940–960; excluded ex ante

Education is harmonized to high school or less, some college, and college; IPUMS CPS code 73 is a high-school diploma and belongs to the first group. Age is divided into 25–44 and 45–65. The 12 matched industries, three education groups, and two age bands give 72 candidate cells. Base support requires pooled CPS Kish support of at least 200 and at least 100 in each early and late period; 65 cells pass. The early-signal rule requires early CPS Kish support of 200; 60 cells pass.

The SIPP source contains 7,026,846 person-months on 263,280 identifiers. Age, weight, and industry restrictions leave 2,495,083 worker-months on 109,401 people. Annualization produces 224,910 worker-years on 108,606 people before the corrected CPS support merge. The pooled supported merge contains 212,635 worker-years on 104,093 people. The strict-post primary sample contains 90,110 worker-years on 52,115 people during 2019–2024 in 60 signal cells; every outcome therefore postdates the CPS destination years, which end in 2018. The broader 2018–2024 timing-sensitivity window contains 107,555 worker-years on 61,162 people in the same 60 cells.

C.2 Portfolio outcomes

Let M_{it} denote the observed stock-and-mutual-fund balance and F_{it} reported financial assets. The level outcomes are

$$\text{PART}_{it} = \mathbf{1}\{M_{it} > 0\}, \tag{7}$$

$$\text{SMF}_{it}^{+} = 100M_{it}/F_{it} \quad \text{for } F_{it} > 0, \tag{8}$$

$$\text{SMF}_{it}^{\text{all}} = 100M_{it}/F_{it} \quad \text{when } F_{it} > 0, \text{ and 0 otherwise}, \tag{9}$$

$$\text{SMF}_{it}^{\text{holder}} = 100M_{it}/F_{it} \quad \text{for } M_{it} > 0, F_{it} > 0. \tag{10}$$

The ratio is not total household equity exposure. Mutual funds need not be pure equity, retirement balances enter reported financial assets without their internal allocation being observed, and employer-stock compensation cannot be cleanly removed.

C.3 Portfolio sample and outcome audit

The person-year construction averages the available monthly balance-sheet records within a calendar year. Table 13 documents the strict-post early-CPS-to-later-SIPP frame before any outcome-specific deletion and retains the broader-window count for comparison. The cached cleaned file contains no asset-imputation flags, so source imputations cannot be classified; missingness, denominator support, ratio bounds, and repeated records can nevertheless be audited directly.

Table 13: SIPP portfolio-sample audit

Audit item	Value	Interpretation
Strict-post primary frame	90,110	Worker-years on 52,115 people during 2019–2024 in 60 supported cells
Broader timing-sensitivity frame	107,555	Worker-years on 61,162 people during 2018–2024 in the same cells
Strict-post positive financial assets	93.40%	Observations entering the positive-assets share
Strict-post weighted participation	19.49%	Positive observed stock-and-mutual-fund balance
Strict-post all-worker share equal to zero	80.49%	Nonholders, including nonpositive financial-asset years
Strict-post fewer than 12 eligible monthly records	15.77%	Annual mean uses the available within-year records
Strict-post positive-assets ratio outside $[0, 100]$	0.00%	No mechanical ratio-tail deletion is required
Broader-frame person-years in multiworker household-years	61.09%	More than one eligible worker shares the household-year
Broader-frame multiworker household-years with same asset tuple	0.87%	Exact cross-worker duplication is rare

Values are computed before outcome-specific missing-value deletion. Strict-post outcome diagnostics apply the 2019–2024 restriction to the frozen worker-year file. The two explicitly labeled household-record rows come from the broader 2018–2024 raw-record audit. The asset tuple comprises the annual stock-and-mutual-fund balance, financial assets, total assets, and net worth. Monthly asset values are constant within each retained person-year in the cached source. Official person-year weights are used only for the participation rate; the remaining shares are unweighted construction diagnostics. Machine-readable broader-window definitions and weighted ratio quantiles are produced by `15_portfolio_sample_audit.R`; the timing restriction and exact strict-post counts are reproduced by `17_post2019_portfolio_timing.R`.

The strict-post frame maps 90,110 worker-years into 61,658 household-years and 36,333 distinct households, using the public SIPP composite household key SSUID plus SHHADID. SSUID alone would merge multiple current households in 811 sample-unit-years, representing 1,651 distinct household-years. Of the correctly keyed household-years, 36,049 contain exactly one worker who satisfies the age, industry, education, supported-cell, and regression-sample rules. Table 14 therefore changes the inference cluster in the full frame and then removes household-years with multiple supported analytic workers. Neither exercise identifies the household's financial decision maker, but together they test whether repeated workers attached to the same household drive the conditional level results.

Table 14: Household decision-unit sensitivity, strict-post sample

Outcome	Raw Kelley	Raw q	Incremental Kelley	Incremental q	N
Panel A. Full worker frame; household and signal-cell clustering					
Participation	0.0038 (0.0073)	0.879	-0.0092^{*} (0.0048)	0.080	90,110
Share, positive financial assets	0.0608 (0.1836)	0.879	-0.2610^{***} (0.0976)	0.039	84,161
Share, all workers	0.0800 (0.1833)	0.879	-0.2341^{**} (0.0985)	0.042	90,110
Share, current holders	−0.0967 (0.6303)	0.879	−0.1525 (0.4292)	0.724	17,578
Panel B. One supported analytic worker per household-year					
Participation	0.0053 (0.0086)	0.538	-0.0108^{*} (0.0058)	0.091	36,049
Share, positive financial assets	0.2704 (0.2272)	0.344	-0.3123^{**} (0.1501)	0.091	33,228
Share, all workers	0.2648 (0.2189)	0.344	-0.2742^{*} (0.1382)	0.091	36,049
Share, current holders	0.9947 (0.8708)	0.344	−0.2065 (0.8321)	0.805	6,497

Coefficients are per one early-CPS signal standard deviation; share outcomes are percentage points. Standard errors in parentheses cluster by the composite SIPP household key (SSUID plus SHHADID) and early-signal cell. Superscripts use unadjusted two-sided conditional tests: $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$. Benjamini–Hochberg q-values are computed separately across the four outcomes within each panel and signal specification. Panel B retains a household-year only when exactly one worker appears in the supported strict-post analytic frame; this does not assert that the retained person is the household's sole earner or financial decision maker. The selected person's official weight is retained because the public cache contains no dedicated household weight. These conditional estimates do not propagate CPS first-stage uncertainty.

The full-frame all-worker coefficient is unchanged by construction and remains −0.234; household rather than person clustering gives an SE of 0.099 and $q = 0.042$. In the single-supported-worker restriction, the corresponding estimate is −0.274 (SE 0.138; $q = 0.091$). Thus the sign and magnitude do not depend on counting multiple supported workers in a household-year, although the narrower sample loses precision and remains a conditional generated-regressor sensitivity rather than causal or household-representative evidence.

C.4 Transition durability and positive-balance thresholds

Table 15 asks whether the nonholder-to-positive-holdings result is a one-year classification or a token-balance artifact. Sustained entry requires zero holdings at $t-1$ and positive holdings at both t and $t+1$. The two threshold outcomes instead require a nominal stock-and-mutual-fund balance to move from no more than \$1,000 or \$5,000 to above that threshold. Each outcome is estimated in the primary frame and in the same-signal-cell restriction.

Table 15: Durability and positive-balance threshold checks

Outcome	Design	Coef.	S.E.	p	BH q	N
Sustained entry through $t+1$	Primary	-0.697^{**}	0.293	0.0205	0.041	9,331
Sustained entry through $t+1$	Same cell	-0.578^{**}	0.268	0.0349	0.045	7,328
Cross from $\leq \$1,000$ to $> \$1,000$	Primary	-1.104^{***}	0.222	< 0.0001	< 0.001	28,686
Cross from $\leq \$1,000$ to $> \$1,000$	Same cell	-1.131^{***}	0.273	0.0001	< 0.001	22,242
Cross from $\leq \$5,000$ to $> \$5,000$	Primary	-0.453^{**}	0.221	0.0445	0.045	29,823
Cross from $\leq \$5,000$ to $> \$5,000$	Same cell	-0.537^{**}	0.260	0.0436	0.045	23,132

Coefficients and standard errors are percentage points per one early-CPS incremental-signal standard deviation. Specifications retain the lagged signal, lagged controls, current-year weight, fixed effects, and two-way person and lag-cell clustering of the primary transition design. Superscripts use unadjusted conditional two-sided tests: $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$. The BH column adjusts once across all six displayed rows. These are conditional generated-regressor checks and do not add CPS first-stage uncertainty. Thresholds are nominal and are not inflation-adjusted. Sustained-entry samples end in 2023 so that $t+1$ is observed.

The primary entry estimate is also negative in every leave-one-outcome-year regression, from -1.203 to -0.881 point; the largest conditional p-value is 0.0031. In a stable-household transition frame, the estimate is -1.057 point (SE 0.275; three-outcome $q = 0.001$). Selecting the highest-earning eligible worker in each household transition gives -0.869 point (SE 0.360; $q = 0.056$). The latter narrowly misses 5% after its three-outcome adjustment, so these rows support decision-unit stability rather than a household-pooling mechanism.

Table 16: Post hoc entry influence and specification stress tests

Diagnostic	Coef.	S.E.	Conditional p	N
Lag signal-cell by outcome-year aggregation	-0.985^{***}	0.297	0.0016	300
Lag-year rather than current-year weights	-0.975^{***}	0.272	0.0007	27,752
Unweighted	-0.806^{***}	0.242	0.0015	27,752
Twelve eligible months in both years	-0.862^{**}	0.345	0.0152	20,522
Quadratic zero-incidence control	-0.982^{***}	0.273	0.0007	27,752
Quadratic and cubic zero-incidence controls	-0.901^{***}	0.298	0.0037	27,752

Coefficients and standard errors are percentage points per one early-CPS incremental-signal standard deviation. The aggregate regression uses 300 lag-cell-by-outcome-year observations and signal-cell-clustered inference with 59 reference degrees of freedom. Other rows retain the primary person-transition frame except where the row changes weights or monthly coverage. Superscripts use unadjusted two-sided conditional tests: $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$. These diagnostics were conducted post hoc; their p-values are not multiplicity adjusted and omit CPS first-stage uncertainty. The nonlinear-zero rows address misspecification of the outcome equation's zero-incidence control without redefining the first-stage signal.

Leaving out each of the 60 signal cells produces entry coefficients from -1.223 to -0.890 point; all are negative and the largest unadjusted conditional p-value is 0.0020. Leaving out each of the 11 represented broad industries gives a range from -1.316 to -0.700 point, again all negative, with maximum $p = 0.0155$. These omission grids are influence summaries, not 60- or 11-test confirmatory families.

D Portfolio Parameterizations

Table 17: Broader 2018–2024 early-CPS to later-SIPP specification path

Outcome	Raw	Raw + controls	Incremental	BH q
Participation	0.004 (0.007)	-0.022^{*} (0.012)	-0.008^{*} (0.005)	0.106
Share, positive financial assets	0.097 (0.162)	-0.552^{**} (0.219)	-0.209^{**} (0.083)	0.058
Share, all workers	0.115 (0.163)	-0.489^{**} (0.227)	-0.185^{**} (0.086)	0.071
Share, current holders	0.200 (0.590)	0.070 (0.993)	0.027 (0.376)	0.944

This table is the broader 2018–2024 timing-sensitivity window, not the strict-post primary sample. Conditional two-way-clustered standard errors are in parentheses. Superscripts denote unadjusted two-sided conditional significance: $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$; they omit BH adjustment and CPS first-stage uncertainty. Raw Kelley is standardized using early supported cells. "Raw plus controls" adds the linear CPS zero share and quadratic mean and dispersion; Incremental Kelley is the standardized residual from that projection. Those two columns are FWL-equivalent—their t statistics, p-values, and four-outcome q-values agree—and are not independent tests. Plug-in first-stage-uncertainty p-values are 0.182 for participation, 0.115 for the positive-assets share, and 0.168 for the all-worker share.

In this broader person-year window, the raw estimates are positive and imprecise. The conditional comparison changes sign for the broad margins. This sign change is the downstream consequence of changing the estimand; the person-year q-values do not pass a 5% family-adjusted threshold, and the CPS first-stage resampling sensitivity widens inference further. The retained broader-window signal-cell-year aggregation gives an incremental positive-assets-share coefficient of -0.224 (SE 0.085; 420 cell-years in 60 signal cells; four-outcome BH $q = 0.044$). The corresponding participation, all-worker-share, and holder-share coefficients are -0.007, -0.170, and 0.092, with q-values 0.198, 0.125, and 0.816. This isolated conditional result predates the strict timing restriction and does not incorporate CPS first-stage uncertainty, so it is not promoted to primary evidence. Sorting, historical rather than perceived risk, and incomplete portfolio measurement preclude a behavioral interpretation.

E Subjective-Belief Construct and Timing Audits

Several distinct literatures discipline this audit. Guiso, Jappelli, and Terlizzese (1996) relates directly elicited income uncertainty to household portfolios. Wang (2023), Caplin et al. (2026), and Arellano et al. (2026) show why perceived earnings risk cannot be substituted mechanically with risk estimated from realized panels. Drerup, Wibral, and Zimpelmann (2023) finds that allocations respond to elicited asset-return skewness conditional on other expected moments, while Giglio et al. (2021) shows that beliefs can predict adjustment direction conditional on trading without predicting whether an investor trades. Dew-Becker (2024) illustrates a useful construct discipline in a different setting: forward-looking and realized skewness are distinct objects, and relative upper and lower semivariance can separate orientation from total scale. Very close concurrent work by van den Akker et al. (2026) reports that a longer perceived earnings right tail predicts administrative stock and bond investment. None of those results validates the fixed-threshold earnings imbalance used here: the risk object, centering, portfolio outcome, and timing differ.

E.1 Fixed thresholds are not centered shape

Lemma 3 (Translation of a fixed-threshold imbalance)**.** *Let X have distribution function F, with density f positive at the relevant thresholds, and let $c > 0$. For the translated variable $X + m$, define*

$$A_c(m) = \Pr(X + m < -c) - \Pr(X + m > c) = F(-c - m) + F(c - m) - 1.$$

Then $A_c'(m) = -f(-c - m) - f(c - m) < 0$, although every translation-invariant centered shape measure is unchanged.

Proof. Continuity gives $\Pr(X + m < -c) = F(-c - m)$ and $\Pr(X + m > c) = 1 - F(c - m)$. Subtraction yields the expression; differentiation yields the derivative. A translation-invariant centered-shape functional takes the same value for X and $X + m$. □

The public SCE density bins use fixed percentage-growth thresholds. The imbalance subtracts probability above +4 percent from probability below −4 percent. In the state-education-year reconstruction, this threshold imbalance is correlated −0.941 with the forecast mean. Additive cubics in the mean and standard deviation explain 89.4 percent of its variation, leaving 32.6 percent of its raw standard deviation. The residual correlates 0.018 with centered Bowley shape and 0.029 with centered moment skewness. Residualizing a fixed-threshold statistic therefore does not recover the same ranking as a centered shape estimator.

E.2 Prospective public SCE-to-SIPP exercise

SCE beliefs are aggregated by state, education, and year and lagged before a subsequent SIPP portfolio change. The prospective merge contains 42,375 consecutive SIPP transitions on 26,005 people, spanning 31 states and 366 state-education-outcome-year cells. The specification absorbs state-by-year, state-by-education, and education-by-year effects, adds perceived mean and dispersion and baseline household controls, and clusters by state. It is a construct-and-timing comparison, not causal evidence.

Table 18: Prospective SCE-to-SIPP construct comparison

Belief signal for subsequent active share change	Coefficient	S.E.	p-value
Fixed-threshold residual	0.1543	0.2136	0.476
Centered Bowley downside shape	-0.3471^{*}	0.1997	0.093
Centered moment downside skewness	−0.0121	0.1745	0.945

Standard errors are clustered by state; superscripts denote unadjusted two-sided significance: $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$. Split reliability is 0.198 for the raw threshold, 0.147 for its leave-year-out residual, 0.068 for Bowley shape, and 0.121 for moment skewness. The centered estimators do not agree, and the threshold result is null.

E.3 Household Finance timing reconstruction

The Household Finance action is the reported direction of retirement reallocation in the defined-contribution account over the preceding twelve months, coded +1 for a move into riskier assets, −1 for a move into safer assets, and 0 for no change; every belief regressor keeps the paper's downside-positive orientation, so a positive coefficient means that perceived downside predicts moving toward risk. The action window is retrospective. A same-calendar-year average of recurring belief reports is not prospectively ordered: among respondents with a coded reallocation answer and at least one same-calendar-year belief report, 87.1 percent have at least one report recorded after the module month, and no same-calendar-year report can predate the start of the twelve-month action window. The reproducible reconstruction therefore uses only reports one through eleven months before the module, the official module weight, and forecast mean and standard deviation controls.

Table 19: Timing-ordered Household Finance reconstruction

Belief window for the reallocation outcome	Fixed-threshold imbalance	Centered Bowley downside shape	N
Panel A. Timing grid (outcome: +1 riskier, −1 safer, 0 no change)			
Pre-module: months −11 to −1	0.067** (0.031)	−0.024* (0.013)	2,136
no intervening realized shock	0.013 (0.051)	−0.037** (0.017)	1,270
realized large income fall only	0.240*** (0.083)	0.015 (0.035)	190
Post-module: months +1 to +11	−0.052 (0.041)	−0.010 (0.015)	1,805
Joint design: pre-module term	0.109** (0.049)	−0.022 (0.015)	1,749
Joint design: post-module term	−0.091* (0.047)	0.001 (0.017)	1,749
Panel B. Standing DC equity share: perceived-dispersion coefficient			
Perceived s.d. (fixed-threshold control)	−3.633*** (1.002)		3,332
Perceived s.d. (Bowley-shape control)	−3.766*** (1.006)		3,332
Perceived s.d. (moment-skewness control)	−4.734*** (0.961)		2,952

Standard errors, in parentheses, are clustered by state; superscripts denote unadjusted two-sided significance: $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$. This timing and subgroup grid is exploratory, and the reported stars and p-values are not adjusted for the family of comparisons in the table. Every specification uses the official Household Finance quarterly weight and Household Finance year and state fixed effects. Panel A additionally controls for forecast mean, perceived dispersion, age, education, and household income. Post-module beliefs cannot have caused the retrospective action, so the post-module row is a negative control; the joint design enters the pre-module and post-module signals together ($N = 1,749$). N refers to the fixed-threshold column; Bowley samples differ by at most one respondent. Centered moment skewness, run on the same grid, is small and unstable: coefficients range from −0.083 to +0.026, and the only entries with $p < 0.05$ are the realized-fall subsample and the joint post-module term. In Panel B, perceived dispersion is the focal regressor rather than a control; each row includes the named shape control plus forecast mean, age, education, and household income. Panel B uses the pre-module construction and reports percentage-point changes in the standing DC equity share per one-standard-deviation increase in perceived earnings dispersion. The equity share is an accumulated stock and beliefs are not assigned, so Panel B is a descriptive scale benchmark rather than a causal estimate.

The grid locates the fixed-threshold measure's apparent pre-module association. Its point estimate is 0.013 when respondents with an intervening realized shock are excluded and 0.240 in the 190-person realized-large-fall subsample; this contrast is exploratory and is not a decomposition of the full-sample estimate. The post-module estimate is -0.052, and post-module beliefs cannot have caused the retrospective action. Entered jointly, the pre-module and post-module terms carry opposite signs ($+0.109$ and -0.091). A belief window that mixes reports from both sides of the module—as 87.1 percent of same-calendar-year report sets do—therefore blends differently signed components and can return either sign. A realized income fall occurring after a pre-module belief report also cannot have imprinted itself on that earlier report; the subgroup contrast may instead reflect selection, common shocks, or the action window's overlap with the intervening period. Combined with Lemma 3's result that the fixed-threshold statistic is dominated by the forecast mean, the grid shows sensitivity to construct and timing choices, not a prospective orientation signal or an identified belief-formation mechanism.

Centered Bowley shape has the background-risk sign in the pre-module rows, a larger magnitude when realized shocks are excluded, and an estimate near zero in the post-module row. It is not corroborated by moment skewness, the comparisons are exploratory, and its magnitude is modest, so it is reported as a directional diagnostic rather than as validation or evidence about how beliefs are formed. The dispersion benchmark, by contrast, is stable across all three shape controls. Together the exercises show that the linked survey data contain economically relevant scale information while reinforcing the paper's narrower lesson: temporal ordering, centered construction, independent reliability, and outcome inference are separate gates.

E.4 Centered upper and lower semimoments

The fixed-threshold statistic does not isolate tail orientation, so I reconstruct each respondent-month's ten-bin public density. Open tails are assigned ± 16 percentage points, the three reports whose integer bins sum to 99 or 101 are renormalized, and lower and upper semideviations are centered on the respondent's own elicited mean. Let d_i^- and d_i^+ denote those semideviations. The focal regression enters their sum and the tilt $d_i^+ - d_i^-$; a positive tilt coefficient means that shifting perceived tail magnitude from downside toward upside is associated with a riskier reported reallocation. This is an absolute semideviation-gap contrast, not a scale-free orientation measure.

Table 20: Perceived earnings-tail tilt and reported retirement reallocation

Specification or outcome	Coef.	S.E.	p-value	BH q	N
Pre-module reallocation: primary tilt	0.0276**	0.0128	0.036	0.054	2,136
equal weight for each belief month	0.0284**	0.0127	0.030	—	2,136
no reported large change or contribution cut	0.0435***	0.0137	0.003	0.008	1,270
Post-module reallocation: negative control	0.0047	0.0178	0.793	—	1,805
Pre-module reallocation: normalized semivariance orientation	0.0189	0.0137	0.173	0.347	1,893
Standing DC equity share: primary tilt	−0.576	0.664	0.390	0.390	3,332

The reallocation outcome equals +1 for a reported move into riskier assets, −1 for a move into safer assets, and zero for no change over the preceding twelve months. Except for the explicitly equal-month row, repeated pre-module belief reports are averaged within respondent using the public monthly SCE sampling weight. The primary tilt is the standardized upside-minus-downside semideviation gap entered jointly with total semideviation scale, expected earnings growth, age, education, household income, Household Finance year fixed effects, and state fixed effects. The normalized row divides the upper-minus-lower semivariance difference by total variance and controls for total standard deviation. Official Household Finance quarterly weights are used; standard errors are clustered by state. Superscripts denote unadjusted two-sided significance: $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

The primary $q = 0.054$ adjusts the three displayed focal shape tests in its outcome family; expanding to fourteen support and outcome variants gives $q = 0.169$. The no-large-change restriction excludes a reported large income, health, or wealth change and a respondent or spouse retirement-contribution cut; it is a post hoc subgroup and its stronger coefficient is scale dependent. The post-module report cannot have caused the retrospective action, but the formal pre-minus-post contrast is also imprecise ($p = 0.191$). Beliefs recorded before the module usually fall inside the retrospective action window. The separate lower and upper semideviations are correlated 0.95 and are therefore not presented as independent discoveries. The result is exploratory and descriptive: it does not survive the full support family, normalized orientation is null, standing share is null, and adding the tilt changes five-fold cross-validated RMSE by only −0.065%.

The main gap coefficient is positive because the separate components carry the economically coherent signs: lower semideviation is −0.107 (SE 0.047; $p = 0.026$) and upper semideviation is +0.087 (SE 0.045; $p = 0.060$) when entered jointly. Their 0.95 correlation makes the gap-plus-sum rotation the more transparent parameterization. In the full sample, the tilt mainly predicts a lower probability of reporting a move into safer assets rather than a statistically precise increase in moves into riskier assets; among active movers alone it is null. Thus the individual SCE densities supply a directional perceived-upside pattern worth reporting, but not evidence that perceived risk dominates assigned histories and not a second headline result.

F Plug-In First-Stage Uncertainty

Suppose the second-stage regression assigns $\widehat{K}_c = T(\widehat{G}_c)$. A conventional variance conditions on the observed $\widehat{K}_c$. Clustering by c permits shared second-stage residuals but does not integrate over repeated samples of the nonlinear first-stage functional.

The CPS bootstrap resamples unique people within each frozen signal cell. If a draw repeats a person m times, weighted first moments use mw and Kish denominators use mw^2, not $(mw)^2$. Every draw rebuilds weighted mean, standard deviation, zero share, quantiles, support selection, equation (5), standardization, SIPP assignment, and the outcome regression. Across 200 draws, all requested estimates succeed; early support ranges from 59 to 61 cells and pooled support from 63 to 66. The strict-post level coefficient draws are not perfectly centered on the corrected point estimates. For participation, their mean is -0.0062 versus a point estimate of -0.0092 (mean-minus-point 0.0029; 2.5th–97.5th percentiles $[-0.0145, 0.0013]$). For the all-worker share, the corresponding values are -0.150 versus -0.234 (mean-minus-point 0.084; percentiles $[-0.363, 0.045]$).

The transition bootstrap additionally reassigns both annual observations and rebuilds the risk sets in each draw. For nonholder-to-positive-holdings entry, the draw mean is -0.650 point versus a point estimate of -1.047 points, a mean-minus-point displacement of 0.397 point. The across-draw standard deviation is 0.258 point, 198 of 200 coefficients are negative, and the 2.5th–97.5th percentile range is $[-1.125, -0.138]$ points. For signed change the corresponding standard deviation is 0.250 point; for exit it is 0.845 point. These ranges are descriptive summaries of rebuilt-CPS coefficients with SIPP held fixed, not confidence intervals, probability statements, or bias corrections; their displacement further motivates treating the variance combination as a sensitivity.

Because CPS and SIPP are independent, the paper reports the following plug-in first-stage uncertainty calculation:

$$\widehat{\mathrm{Var}}_{\text{plug-in}}(\widehat{\beta}) = \widehat{\mathrm{Var}}_{\mathrm{SIPP}}(\widehat{\beta} \mid \widehat{G}) + \widehat{\mathrm{Var}}_{\mathrm{CPS}}(\widehat{\beta}).$$

The first term is the conditional two-way-clustered SIPP variance evaluated at the observed signal, and the second is the across-draw variance induced by resampling CPS people and rerunning the stated first-stage pipeline. This is not a joint nonparametric bootstrap of CPS and SIPP. It propagates CPS signal-estimation uncertainty within the frozen cell schema, crosswalk, and observed SIPP sample; it does not cover crosswalk or functional-form choice, SIPP sample-reconstruction uncertainty, or a latent-shape correction. The resampler also does not model cross-cell covariance from a person appearing in multiple cells, but no CPS person does so in the frozen annual-link sample. Reported plug-in p-values use a two-sided standard-normal reference. The primary strict-post first-stage results use this calculation; Table 21 reports the corresponding design sensitivity.

Table 21: Strict-post power for the early incremental signal

	80% minimum detectable effect		Power at observed magnitude	
Outcome	Conditional	CPS plug-in	Conditional	CPS plug-in
Participation	1.34	1.79	0.481	0.298
All-worker asset share	0.273	0.409	0.672	0.361

Entries use the strict-post 2019–2024 sample and are per one standard deviation of the early CPS incremental signal. Both MDE columns are in percentage points. The first pair gives two-sided 5%-size, 80%-power minimum detectable effects under a normal approximation; the second gives power at the absolute magnitude of the observed coefficient. "Conditional" treats the estimated CPS signal as fixed. "CPS plug-in" uses the conditional variance plus the CPS first-stage bootstrap variance and remains a sensitivity rather than joint CPS–SIPP inference. The table deliberately reports no latent-signal correction: weak cross-survey transport and possible nonclassical error preclude an equivalence mapping from the assigned signal to latent orientation.

G Scope of the Group-Generated-Risk Workflow

The protocol is motivated by a class of designs rather than by a claim that every paper below makes the same choices or faces the same data constraint. Table 23 records representative finance applications in which a risk object is estimated from one history or grouping and then used in a portfolio, consumption, or performance exercise. Rich administrative designs can make some gates much less binding; the common feature is that construction of the estimated risk object remains part of the downstream estimand.

H Reproducibility Map

All evidence uses public CPS ASEC, SIPP, SCE, and SCF already in the workspace; no experiment or new respondent collection enters. Scripts 01–03, 05–09, and 14–26 accept

Table 22: Implementation checklist for group-estimated earnings-risk regressors

Layer	Required disclosure	Diagnostic reported here
Population and support	Define origin population, destination zeros, growth transformation, cell schema, weights, and effective support.	Positive-origin workers; zero-inclusive arc growth; exact crosswalk; Kish support.
Reproducibility	Split globally by person and report rank and time stability.	Person-split Spearman–Brown reliability and early-versus-late correlations.
Incremental content	State the conditional interpretation and report the sequential relation with zero share, mean, and dispersion.	Raw; after zero share; after zero share, mean, and dispersion.
Downstream inference	Preserve support, nonlinear estimation, projection, standardization, and assignment inside first-stage resampling.	Independent CPS person bootstrap combined with conditional SIPP variance.

the project root as their first argument (the public runner supplies it); utilities 10, 11, 27, and 99 resolve paths locally; structural scripts remain under `replication/structural_v2`. Frozen invariants are 183,809 CPS transitions, 72 candidate cells, 65 pooled reliability cells, 60 early signal cells, 51,152 matched SIPP transitions on 51,152 people, 90,110 strict-post SIPP worker-years on 52,115 people, 35,112 consecutive strict-post pairs on 23,170 people, 27,752 entry-risk and 7,360 exit-risk observations, 107,555 broader-window SIPP worker-years on 61,162 people, global person splits, destination zeros at arc growth -2, and 200 valid strict-post first-stage draws.

Table 23: Illustrative workflows using estimated risk objects

Study	Estimated object	Downstream use	Relation to this protocol
Ranish (2013)	Permanent variance, disaster exposure, and cyclical labor-income risk estimated in PSID histories and mapped to SCF households	Household stockholding and leverage	Direct two-sample group assignment; bootstraps estimated risk, multiple imputation, and generated-regressor uncertainty
Angerer and Lam (2009)	Occupation-by-education permanent and transitory income risk	Household risky-asset participation and shares	Early public group-risk assignment; does not isolate a nonlinear tail-orientation contrast
Fagereng, Guiso, and Pistaferri (2018)	Firm-level shocks to individual wage income	Household risky shares	Administrative and instrumental-variables benchmark for the effect of wage risk rather than a public-cell measurement design
Chang et al. (2022)	Individual income-volatility breaks and firm-volatility instruments	Household portfolio choice	Causal second-moment benchmark; distinct from higher-moment reliability and transport
d'Astous and Shore (2024)	Permanent income volatility shifted by university-admission discontinuities	Household participation and risky shares	Quasi-experimental sorting benchmark; stronger identification but a different risk object
Shen (2024)	Regional earnings-risk skewness	PSID consumption and risky portfolio shares	Nonlinear group risk assigned to households; cell support and incremental shape are relevant interpretation gates
Catherine, Sodini, and Zhang (2024)	Education-by-industry cyclical income skewness in Swedish registers	Household equity participation and shares	High-information administrative benchmark for the same risk-to-portfolio logic
Catherine (2022)	Estimated non-Gaussian, state-contingent income process	Life-cycle portfolio model	Shows the economic object that a reduced-form group statistic seeks to proxy
Gálvez and Paz-Pardo (2026)	Nonlinear earnings process estimated from PSID histories	Structurally estimated consumption, housing, participation, and risky shares	Substantive benchmark showing that rich nonlinear risk can matter; it does not test a public group-assigned shape statistic
Akey, Robertson, and Simutin (2026)	Methodologically revised Fama–French factor vintages	Mutual-fund performance, stock alphas and betas, and corporate-finance conclusions	Finance analogue showing that construction revisions can change downstream inference

The table establishes empirical scope; it is not an allegation that the cited papers fail a universal checklist. Their data, estimands, and identification strategies differ. The present contribution isolates three disclosures that are especially important when nonlinear public-cell statistics are assigned across samples: reproducibility, incremental content, and first-stage uncertainty accounting.

Table 24: Public-data reproduction sequence

Script	Outputs used in the paper
`01_cps_reliability_frontier.R`	General estimator and support tournament; Appendix Table 8
`02_cps_construction_audit.R`	Corrected CPS positive-to-positive and zero-inclusive comparison
`05_sce_sipp_strict_predictive.R`	Prospective SCE-to-SIPP construct comparison
`06_sce_hf_timing_falsification.R`	Timing-ordered Household Finance reconstruction
`07_cps_sipp_full_pipeline_bootstrap.R`	Corrected early and pooled signals; broader-window four-outcome table; cell-year estimates; 200-draw CPS-first-stage resampling and plug-in sensitivity
`08_sequential_content_and_matched_sipp.R`	Main sequential content table and one-transition-per-person SIPP audit
`09_alternative_shape_robustness.R`	Bowley, moment, and L-skewness reliability and portfolio families
`09_build_rof_exhibits.R`	Main construction figure from frozen machine-readable outputs
`10_transport_figure.R`	Cell-level time-transport figure from the frozen sequential-content manifest; self-verifies against the published correlations
`11_window_mixing_share.py`	Window-mixing share of same-calendar-year belief reports around the Household Finance module
`14_reliability_transport_uncertainty.R`	Split-assignment, transition-selection, complete-year SIPP, and direct common-year CPS–SIPP sensitivities
`15_portfolio_sample_audit.R`	Denominator support, ratio bounds, annual coverage, and household-record audit for the broader 2018–2024 portfolio frame
`16_zero_state_predictive_validity.R`	CPS and SIPP destination-zero anatomy, component reliability, and the selection-complete strictly ordered prediction audit
`16_measurement_sampling_bootstrap_power.R`	Exact point validation and the MDE/power table; its separately retained multiplier pilot failed the centering gate and supplies no sampling interval used in the paper
`17_post2019_portfolio_timing.R`	Strict-post-measurement 2019–2024 portfolio timing audit and exact reproduction of the broader 2018–2024 estimates
`18_post2019_first_stage_bootstrap.R`	Strict-post 200-draw CPS first-stage reconstruction, plug-in inference, and exact draw-by-outcome validation against the broader frozen pipeline
`19_household_unit_robustness.R`	Household-plus-signal clustering and the one-supported-worker-per-household-year strict-post sensitivity
`20_strict_entry_margin.R`	Strict-post signed-change, entry, and exit family; same-cell and leave-one-year-out checks; retained exploratory-pipeline bridge
`21_strict_entry_first_stage_bootstrap.R`	Dynamic 200-draw CPS first-stage reconstruction, risk-set rebuilding, plug-in inference, and exact validation
`22_public_sipp_pooling_audit.R`	Composite-household identifier audit and conditional transition decision-unit sensitivities
`23_entry_durability_thresholds.R`	Sustained-entry and nominal $1,000/$5,000 positive-balance threshold checks
`24_entry_specification_stress_tests.R`	Post hoc influence, aggregation, weighting, annual-coverage, and nonlinear-zero-control diagnostics
`25_sce_upside_assigned_comparison.R`	Conditional entry tail decomposition and rotation (Table 4 point estimates), assigned-SD scale check, and the retained SCE head-to-head comparison
`26_entry_tail_decomposition_first_stage_bootstrap.R`	200-draw CPS first-stage rebuild of the tail decomposition: plug-in inference, joint tests, and the raw-scale equality test
`26_sce_centered_semimoments.R`	Respondent-level centered semideviations and the perceived tail-tilt audit of Appendix Table 20
`27_entry_tails_figure.R`	Entry-margin tail-decomposition figure from the frozen first-stage summary; self-verifies against the published coefficients
`99_verify_paper_numbers.py`	Checks every number printed in the paper against the frozen outputs; exits nonzero on any mismatch
`replication/structural_v2/prepare_public_inputs.R`	Public CPS shock laws and 2019 SCF wealth-to-income states for the matched-even-moments benchmark
`replication/structural_v2/run_benchmark.py`	Orientation benchmark, policy comparisons, moment identities, and grid checks

References

Akey, Pat, Adriana Z. Robertson, and Mikhail Simutin, 2026, Noisy factors? The retroactive impact of methodological changes on the Fama–French factors, *Review of Finance*, rfag002, doi:10.1093/rof/rfag002.

Angerer, Xiaohong, and Pok-Sang Lam, 2009, Income risk and portfolio choice: An empirical study, *Journal of Finance* 64, 1037–1055.

Arellano, Manuel, Richard Blundell, and Stéphane Bonhomme, 2017, Earnings and consumption dynamics: A nonlinear panel data framework, *Econometrica* 85, 693–734.

Arellano, Manuel, Orazio Attanasio, Margherita Borella, Mariacristina De Nardi, and Gonzalo Paz-Pardo, 2026, Subjective earnings and employment dynamics, NBER Working Paper No. 35027, doi:10.3386/w35027.

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, 289–300.

Betermier, Sébastien, Thomas Jansson, Christine Parlour, and Johan Walden, 2012, Hedging labor income risk, *Journal of Financial Economics* 105, 622–639.

[dataset] Board of Governors of the Federal Reserve System, 2019, 2019 Survey of Consumer Finances public data, Washington, DC, https://www.federalreserve.gov/econres/scfindex.htm.

Brown, William, 1910, Some experimental results in the correlation of mental abilities, *British Journal of Psychology* 3, 296–322.

Busch, Christopher, David Domeij, Fatih Guvenen, and Rocio Madera, 2022, Skewed idiosyncratic income risk over the business cycle: Sources and insurance, *American Economic Journal: Macroeconomics* 14, 207–242.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, 2011, Robust inference with multiway clustering, *Journal of Business & Economic Statistics* 29, 238–249.

Caplin, Andrew, Victoria Gregory, Eungik Lee, Søren Leth-Petersen, and Johan Sæverud, 2026, Subjective earnings risk, *Review of Economic Studies*, rdag070, doi:10.1093/restud/rdag070.

Catherine, Sylvain, 2022, Countercyclical labor income risk and portfolio choices over the life cycle, *Review of Financial Studies* 35, 4016–4054.

Catherine, Sylvain, Paolo Sodini, and Yapei Zhang, 2024, Countercyclical income risk and portfolio choices: Evidence from Sweden, *Journal of Finance* 79, 1755–1788.

Chang, Yongsung, Jay H. Hong, and Marios Karabarbounis, 2018, Labor market uncertainty and portfolio choice puzzles, *American Economic Journal: Macroeconomics* 10, 222–262.

Chang, Yongsung, Jay H. Hong, Marios Karabarbounis, Yicheng Wang, and Tao Zhang, 2022, Income volatility and portfolio choices, *Review of Economic Dynamics* 44, 65–90.

Cocco, João F., Francisco J. Gomes, and Pascal J. Maenhout, 2005, Consumption and portfolio choice over the life cycle, *Review of Financial Studies* 18, 491–533.

d'Astous, Philippe, and Stephen H. Shore, 2024, Human capital risk and portfolio choices: Evidence from university admission discontinuities, *Journal of Financial Economics* 154, 103793.

Dew-Becker, Ian, 2024, Real-time forward-looking skewness over the business cycle, *Review of Economic Dynamics* 54, 101233, doi:10.1016/j.red.2024.101233.

Drerup, Tilman H., Matthias Wibral, and Christian Zimpelmann, 2023, Skewness expectations and portfolio choice, *Experimental Economics* 26, 107–144, doi:10.1007/s10683-022-09780-9.

Fagereng, Andreas, Charles Gottlieb, and Luigi Guiso, 2017, Asset market participation and portfolio choice over the life-cycle, *Journal of Finance* 72, 705–750, doi:10.1111/jofi.12484.

Fagereng, Andreas, Luigi Guiso, and Luigi Pistaferri, 2018, Portfolio choices, firm shocks, and uninsurable wage risk, *Review of Economic Studies* 85, 437–474.

[dataset] Federal Reserve Bank of New York, 2026, Survey of Consumer Expectations public data bank, New York, NY, https://www.newyorkfed.org/microeconomics/databank.

[dataset] Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, Daniel Backman, Etienne Breton, Grace Cooper, Julia A. Rivera Drew, Stephanie Richards, David Van Riper, and Kari C. W. Williams, 2025, IPUMS CPS: Version 13.0, IPUMS, Minneapolis, MN, doi:10.18128/D030.V13.0.

Gálvez, Julio, and Gonzalo Paz-Pardo, 2026, Richer earnings dynamics, consumption and portfolio choice over the life cycle, *Journal of Financial Economics* 176, 104206, doi:10.1016/j.jfineco.2025.104206.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus, 2021, Five facts about beliefs and portfolios, *American Economic Review* 111, 1481–1522, doi:10.1257/aer.20200243.

Gomes, Francisco, and Alexander Michaelides, 2005, Optimal life-cycle asset allocation: Understanding the empirical evidence, *Journal of Finance* 60, 869–904, doi:10.1111/j.1540-6261.2005.00749.x.

Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese, 1996, Income risk, borrowing constraints, and portfolio choice, *American Economic Review* 86, 158–172.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song, 2021, What do data on millions of U.S. workers reveal about lifecycle earnings dynamics?, *Econometrica* 89, 2303–2339.

Guvenen, Fatih, Serdar Ozkan, and Jae Song, 2014, The nature of countercyclical income risk, *Journal of Political Economy* 122, 621–660.

Harmenberg, Karl, 2021, The labor-market origins of cyclical skewness, Working paper, Copenhagen Business School.

Hosking, J. R. M., 1990, L-moments: Analysis and estimation of distributions using linear combinations of order statistics, *Journal of the Royal Statistical Society: Series B (Methodological)* 52, 105–124.

Kim, Tae-Hwan, and Halbert White, 2004, On more robust estimation of skewness and kurtosis, *Finance Research Letters* 1, 56–73.

Kish, Leslie, 1965, *Survey Sampling* (Wiley, New York, NY).

Moulton, Brent R., 1990, An illustration of a pitfall in estimating the effects of aggregate variables on micro units, *Review of Economics and Statistics* 72, 334–338.

Murphy, Kevin M., and Robert H. Topel, 1985, Estimation and inference in two-step econometric models, *Journal of Business & Economic Statistics* 3, 370–379.

Pagan, Adrian, 1984, Econometric issues in the analysis of regressions with generated regressors, *International Economic Review* 25, 221–247.

Ranish, Benjamin, 2013, Why do households with risky labor income take greater financial risks?, Available at SSRN 2681360, doi:10.2139/ssrn.2681360.

Shen, Jialu, 2024, Countercyclical risks, consumption, and portfolio choice: Theory and evidence, *Management Science* 70, 2862–2881.

Spearman, Charles, 1910, Correlation calculated from faulty data, *British Journal of Psychology* 3, 271–295.

Törnqvist, Leo, Pentti Vartia, and Yrjö O. Vartia, 1985, How should relative changes be measured?, *The American Statistician* 39, 43–46.

[dataset] U.S. Census Bureau, 2025, Survey of Income and Program Participation public-use data, 2014–2023 panels, Washington, DC, https://www.census.gov/programs-surveys/sipp/data/datasets.html.

van den Akker, Ruben, Jochem de Bresser, Federica Teppa, and Serena Trucchi, 2026, Subjective expectations and household portfolio choices, Work in progress presented at the 2026 IAAE Annual Conference, https://editorialexpress.com/conference/IAAE2026/program/IAAE2026.html.

Viceira, Luis M., 2001, Optimal portfolio choice for long-horizon investors with nontradable labor income, *Journal of Finance* 56, 433–470.

Vissing-Jorgensen, Annette, 2002, Towards an explanation of household portfolio choice heterogeneity: Nonfinancial income and participation cost structures, NBER Working Paper 8884, doi:10.3386/w8884.

Wang, Tao, 2023, Perceived versus calibrated income risks in heterogeneous-agent consumption models, Bank of Canada Staff Working Paper 2023-59, doi:10.34989/swp-2023-59.