

Asymmetric Labor Income Risk and Stock Market Participation: Evidence from the US*

Tai Lo Yeung[†]

October 2026

Abstract

I document that upside earnings opportunities predict stock market participation among initial nonholders, beyond mean earnings growth, total variance and measured household resources. I measure earnings distributions in historical Current Population Survey records and link them to subsequent portfolio transitions in the Survey of Income and Program Participation. A one-standard-deviation increase in upside composition predicts 1.6 percentage points more entry into stocks and mutual funds, nearly one fifth of the sample entry rate. The association extends to entry followed by ownership at two subsequent annual observations. The result concerns asymmetry in earnings itself, rather than its covariance with stock returns. A standard participation-cost model shows how favorable realizations can finance entry even at unchanged mean and variance. In an estimated life-cycle model, measured earnings dynamics strengthen the untargeted composition-entry relationship relative to pooled dynamics, mainly through differences in resources available to invest. Controlled experiments show why immediate entry and later access need not move together: the anticipated earnings distribution changes saving and the resources available for future participation.

Keywords: labor income risk; household portfolios; tail composition; stock market participation; measurement error

JEL codes: C18, D14, G11, J31

Online Appendix. Supplementary results and technical details are provided separately through the author's research webpage, <https://tlyeungae.github.io/#research>.

*I thank Lorenz Kueng, Alberto Plazzi, Luise Eisfeld, Fabrizio Colella, Darrell Duffie, Moritz Kuhn, Petra Todd, Harald Uhlig, Sylvain Catherine, Jessica Wachter, and seminar participants for helpful comments. All errors are my own.

[†]Università della Svizzera italiana, Via Giuseppe Buffi 13, Lugano, Switzerland. Email: tai.lo.yeung@usi.ch.

1 Introduction

I document that upside earnings opportunities predict the start of stock market participation. Among workers who initially hold no stocks or mutual funds, a more upside-oriented earnings distribution predicts more entry at a given mean earnings growth, total variance and measured household resources. A one-standard-deviation increase in upside composition predicts 1.6 percentage points more entry, nearly one fifth of the sample's 8.82% annual entry rate. The relationship also extends to entry followed by ownership at later annual observations.

The economic question is how earnings uncertainty affects becoming an investor when it includes opportunities for gains as well as risks of losses. Earnings finance household investment, while uncertainty about those earnings can discourage financial risk taking. Workers facing the same earnings variance may therefore have different prospects for accumulating the resources needed to enter, depending on which side of the earnings distribution accounts for that variation.

The central empirical distinction is between the two sides of earnings variation. I divide the variance of age-adjusted earnings growth into upper and lower partial variances, measuring variation above and below mean growth. Their sum is total variance; their difference is upside composition. Upside composition positively predicts entry when variance and its square root are included in the regression. Both variance terms are individually significant, whereas a linear variance-only specification yields an imprecise coefficient. In a linear specification entering the upper and lower partial variances jointly, their coefficients have opposite signs. The equivalent composition coefficient is 1.11 percentage points per standard deviation, compared with 1.60 when the variance relationship is allowed to curve. The direction survives this simpler specification, although its magnitude and precision depend on the functional form.

Standard saving and costly-participation models give this distinction an economic interpretation. A two-period consumption and saving problem yields a resource threshold for investing: a favorable earnings realization can make paying the fee worthwhile. Mean and variance do not determine how often earnings cover a household's shortfall to that threshold. An analytical example therefore generates different entry rates at unchanged mean and variance without group-specific preferences or costs. In a repeated version, more entry can also produce more subsequent ownership, even if entrants are no more likely to survive conditional on entry. The aggregate sign depends on households' positions relative to the threshold; the example does not imply that every increase in upside composition raises participation.

The empirical comparison uses earnings histories measured before, and separately from, portfolio decisions. I construct earnings distributions from the Current Population Survey (CPS), using histories with origins in 2013–2017 and destinations through 2018. I assign moments from 58 industry–education–age cells to workers in the Survey of Income and Program Participation (SIPP), then follow portfolios from 2019–2023 baselines to the next year. The entry sample contains 27,419 initial-nonholder transitions. Separating the surveys prevents a worker's own realized earnings change from mechanically determining the assigned distribution. The estimates compare historical

earnings environments; they do not require treating those environments as measured individual beliefs.

Measurement matters because adverse earnings outcomes include leaving wage work. I use arc growth, which remains defined when positive wages fall to zero, and remove industry–education-specific age profiles following Catherine, Sodini, and Zhang (2024). Including zero-wage destinations substantially improves the reliability of measured composition and its agreement with independently constructed PSID earnings distributions. Cells with more zero-wage destinations have less upside composition and greater total variance, helping explain why composition and variance are negatively correlated across workers’ earnings environments. The participation association remains after controlling for that frequency and measured cyclical exposure. Partial variances also give the comparison a transparent interpretation: composition changes which side of mean growth accounts for variation while total variance is held fixed. Conventional skewness captures closely related asymmetry. The advantage here is a more reliably measured representation in these survey cells and an exact variance decomposition, rather than an empirically distinct source of asymmetry.

The main competing explanation is sorting into earnings environments. The baseline conditions on resources, demographics and industry, education, age-band and year effects. The entry association remains with joint industry–education effects, controls for cyclical exposure, and stricter portfolio-reporting screens. Adding baseline earnings and realized earnings changes, including separate gains and losses, does not attenuate it on a common sample. In the PSID, reported risk tolerance rises with variance but not detectably with composition. These comparisons address observable differences in resources and risk exposure, while leaving room for unobserved sorting. The paper establishes a conditional association across earnings environments, not a causal response to an isolated change in earnings risk.

The subsequent ownership evidence makes the entry result more economically informative. Greater wage upside composition predicts entry followed by ownership at both subsequent annual observations in the smaller cohort with longer follow-up. Entry followed by exit also rises at the shorter horizon. Both outcomes are measured among initial nonholders: more entry can increase the number who later stay and the number who later leave, without changing entrants’ likelihood of staying. The estimates of that conditional likelihood are imprecise. These exploratory comparisons use selected complete-followup cohorts and measure ownership at annual interviews. Among initial holders, the main wage measure predicts lower retention but no clear average change in risky shares; the split by holding history is imprecise. The strongest evidence concerns becoming an investor, rather than a general increase in portfolio risk taking.

The entry pattern also appears outside the SIPP and before the pandemic. Among younger households, pre-pandemic PSID transitions show positive entry associations with the same opposing tail signs as the SIPP. Ownership associations extend to the 2019 Survey of Consumer Finances and the earlier SIPP panel. The main sample uses recent SIPP releases following the Census Bureau’s revision of asset processing, including stock and mutual-fund ownership imputation. The PSID

supplies an independent comparison under a different survey design and in earlier years. Section 4.7 brings together this evidence and the comparison across SIPP survey vintages.

The PSID also connects earnings opportunities to other household decisions. Among initial stockholders, upside composition predicts greater participation in IRA or private-annuity contributions, with a significantly stronger association than among nonholders. Furnishings and household equipment also have a positive association among initial holders in the primary transformed specification, whereas broad expenditure, selected nondurables and services, and food show no clear association. These comparisons motivate distinguishing financial participation and occasional resource commitments from average spending. They do not establish that the responses differ statistically across expenditure categories. The two-period analytical model and access experiments in the life-cycle model show how a discrete participation margin can move substantially with modest changes in average consumption. The estimated model has one consumption good; the furnishings evidence is complementary to its mechanism.

I then use a life-cycle model to test the empirical relevance of this mechanism when households anticipate future earnings and repeatedly choose consumption, saving and portfolios. The 58 observed earnings cells inform 30 industry–education career processes, with adjustments before and after age 45. These measured laws affect both realized resources and the future earnings prospects used to choose optimal saving and investment rules. Preferences, participation costs and entry opportunities are common across groups. A nonholder can receive a fresh opportunity each year; holders retain access and decide whether the recurring payment is worthwhile. Five financial parameters are estimated from eight entry, holder and wealth outcomes, excluding all earnings-risk coefficients. Observed wages, wealth and portfolios initialize the model, so the exercise evaluates conditional decisions rather than fitting an unconditional wealth distribution.

At common financial parameters and observed starting positions, the composition coefficient is 0.241 percentage point with cell dynamics, -0.005 with pooled dynamics, and 1.816 in the matched data. Pooling removes the cell adjustments while retaining non-Gaussian education-specific processes and resolves household choices without refitting financial parameters. Cell dynamics recover the positive composition association and opposing tail signs in the selected specification, although the magnitude remains much smaller than in the data. The comparison tests what measured earnings heterogeneity adds; it does not establish that standard portfolio models require cells to generate the mechanism. Indeed, the permanent-access alternative produces a positive composition association even under pooled laws.

Crossing the earnings laws with rules previously optimized under either law attributes most of the additional composition gradient to differences in resources and earnings states reached before the entry decision. Rule adjustment partly offsets that increment. This compares measured environments that differ in growth and employment risk as well as asymmetry. A separate controlled experiment reallocates dollar earnings variance between favorable and adverse outcomes while preserving expected disposable earnings and total dollar variance. In this illustration, more upside reduces desired saving at given resources while initially making stockholding slightly more attractive.

A five-year change eventually lowers cumulative entry; a change lasting throughout working life lowers wealth more, but its cumulative-entry response fades toward zero as the investment incentive continues. The two exercises therefore do not imply that expectations are always secondary. Together they show how familiar resource and saving mechanisms can produce different short- and longer-run participation responses, with independent future stock returns. The controlled experiment illustrates this interaction rather than matching the empirical coefficient.

Related literature. The main contribution is to establish the economic importance of upside earnings opportunities for stock market participation. The allocation of earnings variation between favorable and adverse outcomes predicts economically important differences in entry beyond mean growth, variance and measured resources. This finding connects three questions: which features of earnings risk matter for portfolios, how households become investors, and what portfolio models should explain beyond average holdings. The starting point is the role of nontradable income, borrowing constraints and background risk in portfolio demand (Viceira, 2001; Cocco, Gomes, and Maenhout, 2005; Kimball, 1993; Gollier and Pratt, 1996). Empirical work relates subjective risk, earnings histories, employer shocks and human-capital exposure to portfolios (Guiso, Jappelli, and Terlizzese, 1996; Angerer and Lam, 2009; Betermier et al., 2012; Fagereng, Guiso, and Pistaferri, 2018; d'Astous and Shore, 2024). More recent evidence links wrongful-discharge protections to participation (Ahn, Jo, and Lam, 2026) and downside income risk to portfolio skewness (Zhang, 2018). I bring favorable and adverse variation into the same prospective entry comparison, holding mean growth and variance fixed when assessing composition. This reveals an entry pattern that treating earnings risk only as dispersion does not recover.

The distinction from cyclical tail risk concerns the earnings distribution being measured. Catherine (2022) and Shen (2024) model countercyclical earnings risk and portfolio choice. Catherine, Sodini, and Zhang (2024) show that Swedish workers with greater left-tail income exposure when equity markets perform poorly hold less equity, examining both participation and risky shares. I study asymmetry in the earnings distribution itself: whether favorable or adverse variation predicts becoming an investor, conditional on mean growth, total variance and measured resources. My composition measure is constructed from earnings histories without using stock returns; it is not a co-moment between earnings risk and market returns. Controls for cyclical exposure leave its entry association intact. The model then studies a resource-and-participation mechanism that operates with independent future stock returns, complementing the joint-downturn mechanism.

Participation costs make entry distinct from the allocation of an existing portfolio (Vissing-Jorgensen, 2002; Gomes and Michaelides, 2005; Fagereng, Gottlieb, and Guiso, 2017). Gálvez (2023) finds larger participation responses to favorable persistent earnings shocks among low-income households; wealth changes and lottery windfalls also reveal differences between entry and participants' allocations (Brunnermeier and Nagel, 2008; Briggs et al., 2021, 2023). Work on participation dynamics relates income risk and income changes to entry and exit (Bonaparte, Korniotis, and Kumar, 2020; Bonaparte et al., 2025), while Galaasen and Raja (2024) document

short ownership spells and the role of poor initial investment experiences. I compare entire earnings environments measured outside the portfolio sample and distinguish entry followed by ownership from survival conditional on entry. This connects the distribution of earnings opportunities to subsequent market access without interpreting every additional entrant as a lasting investor.

Gálvez and Paz-Pardo (2026) provide the closest quantitative benchmark. Their model combines nonlinear, non-Gaussian earnings dynamics, housing and participation costs to explain stockholding and risky shares with lower estimated risk aversion. I document a composition association with subsequent entry and entry followed by ownership, then apply the empirical specification to conditional life-cycle predictions without targeting these coefficients. The model's role is to provide a familiar theoretical foundation and assess it against this distributional entry evidence. Measured earnings environments add a positive entry gradient in the selected specification, chiefly through resources, while the magnitude gap and controlled saving experiment delimit the quantitative interpretation. The model omits housing choice and distinguishes intermittent entry opportunities from permanent access. The contribution is empirical discipline on established mechanisms beyond their fit to average financial outcomes.

The measurement draws on evidence that earnings risk is non-Gaussian and changes over the life cycle (Guvenen, Ozkan, and Song, 2014; Guvenen et al., 2021; Arellano, Blundell, and Bonhomme, 2017), and that household and government insurance alter the risk borne by individuals (Busch et al., 2022). The decomposition parallels the distinction between upside and downside semivariances in financial markets (Barndorff-Nielsen, Kinnebrock, and Shephard, 2010; Segal, Shaliastovich, and Yaron, 2015). Conditional return asymmetry also matters for international portfolio allocation (Ghysels, Plazzi, and Valkanov, 2016); related work studies firm-level conditional skewness (Dew-Becker, 2024) and labor-market event risk (Schmidt, 2025). I use partial variances to make a constant-variance comparison of entry and assess its measurement across public surveys. Subjective-expectations research offers a complementary route (Wang, 2023; Caplin et al., 2026; Arellano et al., 2026; Giglio et al., 2021; Drerup, Wibral, and Zimpelmann, 2023). I examine perceived risk separately in the supplementary Survey of Consumer Expectations exercise, keeping the distinction between historical earnings distributions and personal beliefs explicit.

Section 2 develops the participation mechanism. Section 3 describes the data and empirical design. Section 4 presents entry, sustained ownership and external comparisons. Section 5 examines retirement-saving participation, expenditure, portfolio-change accounting and subjective risk. Section 6 estimates the conditional life-cycle model and evaluates its entry predictions. Section 7 concludes.

2 Earnings Composition and the Participation Margin

A familiar participation cost gives earnings composition an economic role. A favorable realization can make investing worthwhile by reducing the fraction of resources absorbed by the cost. The probability of reaching that point depends on the earnings distribution, not just its mean and

variance. I first show this mechanism in a two-period consumption and saving problem, then isolate its implications for entry and subsequent ownership. The mechanism uses standard household-portfolio ingredients (Vissing-Jorgensen, 2002; Fagereng, Gottlieb, and Guiso, 2017); it requires neither different preferences across earnings groups nor a particular number of groups. The earnings cells introduced below supply measured environments in which to evaluate it.

2.1 A two-period consumption and investment decision

An earnings innovation Y is realized before the household chooses consumption and its portfolio. Current resources are $m = w + Y > 0$, where w is predetermined wealth. The household can save safely or pay a participation fee $\kappa > 0$ to hold stocks. The safe net return is zero; stock returns are independent of earnings and offer a positive optimized expected log-return gain b . There is no borrowing, short selling or terminal nontradable income. With logarithmic utility and discount factor $\beta > 0$, the value conditional on participation $d \in \{0, 1\}$ is

$$V_d(m) = \max_{0 < c < m - \kappa d} \{\log c + \beta[\log(m - \kappa d - c) + bd]\}. \quad (1)$$

The household optimally divides resources remaining after the fee between consumption and saving. Its consumption and participation choices are

$$c_d(m) = \frac{m - \kappa d}{1 + \beta}, \quad d^*(m) = \mathbf{1}\{m \geq m^*\}, \quad m^* = \frac{\kappa}{1 - \exp[-\beta b/(1 + \beta)]}. \quad (2)$$

Participation is infeasible when $m \leq \kappa$; indifference at m^* is resolved in favor of participating. Appendix A derives the threshold and consumption response.

The household invests when earnings cover its shortfall to the threshold: $\Pr(d^* = 1) = \Pr(Y \geq m^* - w)$. Mean and variance do not determine that probability. Two distributions with the same mean and variance can therefore produce different entry rates in the same decision problem. No group-specific fee or preference is needed. The sign depends on which earnings realizations cross the threshold, so higher upside composition does not universally increase entry. The life-cycle model in Section 6 adds future earnings uncertainty and repeated saving decisions to determine how this resource channel interacts with expectations and accumulation.

2.2 Repeated participation at a resource threshold

To isolate entry and subsequent ownership, consider a repeated version with saving and initial wealth fixed. A household begins with financial wealth w and observes an earnings innovation Y , giving resources $w_1 = w + Y > 0$. It receives safe nontradable income $h > 0$ at the end of the period. Earnings are realized before portfolio choice. Utility is logarithmic; the safe net return is zero; and the independent risky excess return is $\bar{r} \in (0, 1)$ with probability $p \in (1/2, 1)$ and $-\bar{r}$ otherwise. There is no borrowing or short selling. Holding stocks costs $\kappa > 0$ each period. Wealth w does not

accumulate, and innovations are independent across periods. This benchmark holds accumulation fixed to isolate the participation decision.

Let $b > 0$ denote the optimized log-utility gain from risky investment before the cost is deducted. When the allocation at the participation boundary is interior, the household participates if

$$w_1 \geq w^* = \frac{\kappa}{1 - e^{-b}} - h, \quad (3)$$

with indifference resolved in favor of participating. The fee absorbs a smaller fraction of resources after a favorable realization. Unlike a one-time access cost (Gomes and Michaelides, 2005), a recurring cost makes current holders revisit the participation decision (Fagereng, Gottlieb, and Guiso, 2017).

The participation probability of a household with wealth w is therefore

$$\pi(w; F_Y) = \Pr(Y \geq w^* - w) = 1 - F_Y((w^* - w)^-). \quad (4)$$

The relevant feature of the distribution is its probability mass at or above the household's resource shortfall. Independent innovations make π both the household's entry and retention probability. Aggregate rates differ because holders disproportionately have high π : in a stationary population, entry is $\mathbb{E}[\pi(1 - \pi)]/\mathbb{E}[1 - \pi]$, whereas retention is $\mathbb{E}[\pi^2]/\mathbb{E}[\pi]$.

2.3 An upside rotation at fixed mean and variance

Consider two earnings laws, with $d > 0$:

$$Y_0 = \begin{cases} -d, & \text{with probability } 1/2, \\ d, & \text{with probability } 1/2, \end{cases} \quad Y_1 = \begin{cases} -d/2, & \text{with probability } 4/5, \\ 2d, & \text{with probability } 1/5. \end{cases}$$

Both have mean zero and variance d^2 . The second replaces a large possible loss with a smaller, more frequent loss and a larger, less frequent gain. Upper-minus-lower partial variance rises from zero to $3d^2/5$.

For example, set $d = \$1,000$. A household \$1,500 short of either participation threshold enters with probability zero under Y_0 and 20% under Y_1 : only the \$2,000 gain covers its shortfall. A household just \$500 short instead goes from 50% to 20%, because the smaller frequent gain under Y_0 was already sufficient. The aggregate entry response therefore depends on the distribution of resources relative to participation thresholds. The next proposition specifies a population for which the upside rotation raises entry and sustained ownership.

There are $N_m > 0$ marginal households with wealth $w_m > d$ and $w^* - w_m \in (d, 2d]$, and $N_e > 0$ established households with $w_e \geq w^* + d$. Only the largest favorable realization carries a marginal household across the threshold; established households participate in every state. These are wealth types, not definitions based on observed tenure.

Proposition 1 (Participation responses to an upside rotation). *Suppose optimal positions are interior whenever participation is chosen. Comparing stationary populations under Y_0 and Y_1 :*

- (i) *marginal households' participation probability rises from 0 to $1/5$, while established households' probability stays at 1. Entry among baseline nonholders rises from 0 to $1/5$;*
- (ii) *retention among baseline holders falls from 1 to $(N_e + N_m/25)/(N_e + N_m/5)$. All exits are by marginal households, and $4/5$ of exiters entered in the preceding period.*

Under Y_0 , only established households hold stocks. Under Y_1 , the holder pool also includes marginal households that drew $2d$. Each remains only if the favorable state recurs, with probability $1/5$. Entry rises while aggregate retention falls because the holder population changes, not because every incumbent becomes more likely to exit. This comparison is between stationary populations in different earnings environments. It does not rank survival among entrants across the two laws: the symmetric law produces no entrants in this example.

The same rotation increases entry followed by sustained ownership. Let D_t indicate participation, with $t = 0$ the initial date. For $k \geq 1$, independence conditional on household type gives the population probability

$$\Pr(D_1 = \dots = D_{k+1} = 1 \mid D_0 = 0) = \frac{\mathbb{E}[(1 - \pi)\pi^{k+1}]}{\mathbb{E}[1 - \pi]}. \quad (5)$$


For the marginal households, the upside rotation raises this probability from zero to $(1/5)^{k+1}$: 4% for one follow-up period and 0.8% for two. Both sustained entry and entry followed by an interruption become possible. Dividing this joint probability by the entry probability gives survival among entrants, a different object. With persistent earnings or endogenous saving, resources in subsequent periods need not follow the independent, fixed-wealth benchmark, and duration can differ.

Figure 1 illustrates the distributions and the resource threshold. With equal numbers of marginal and established households, the example gives entry rates of 0% and 20%, and retention rates of 100% and 86.7%. Its threshold placement holds for an open interval of costs.

Empirical implications. The model links upside-oriented earnings to entry when favorable realizations bring the participation threshold within reach. This is not a monotone wealth effect: households far below the threshold never enter, those far above always participate, and the displayed rotation lowers participation for some intermediate wealth levels. Entry and sustained entry are therefore the central comparisons across earnings environments. Transitory gains can also generate short spells. Longer holding histories increasingly select established households, but observed tenure is only a proxy for wealth type; conditional on a known entry, the example's exit hazard is constant. Survival, earnings reversals and baseline resources help assess this particular transitory-income mechanism.

The proposition establishes these signs for a specified rotation, not for every increase in upper partial variance. The rotation holds dollar-innovation moments fixed, whereas equal variance of adjusted log or arc growth need not preserve dollar-income variance. Earnings are realized before

Figure 1. Earnings composition and the participation margin at fixed mean and variance



Notes: Panels A and B show earnings-innovation distributions with mean zero and variance one. The dashed line at $Y = 1.5$ is the innovation that carries a household with wealth $w = 3$ to the resource threshold $w^* = 4.5$. Neither realization of the symmetric law crosses the threshold; the upside law crosses it with probability 20%. The same household remains invested only if the favorable state recurs. Earnings are realized before allocation and saving is fixed. These are exact distributions from the illustration in Section 2, not an estimated income process or a calibration to SIPP.

portfolio choice, so participation responds to realized resources. The proposition makes no prediction about unresolved background risk or an age cutoff. Whether participation costs quantitatively account for the evidence is a separate question; windfall evidence and participation dynamics identify limitations of standard cost models (Briggs et al., 2021; Galaasen and Raja, 2024).

2.4 Financial participation and current consumption

A sizable participation response can coexist with a small average consumption response. Access to an investment opportunity is a discrete choice, whereas consumption is chosen continuously. Returning to the two-period problem in equations (1)–(2), consider two earnings environments with the same mean resources. Optimal consumption implies

$$\Delta \mathbb{E}[c] = -\frac{\kappa}{1 + \beta} \Delta \Pr(d^* = 1). \quad (6)$$

An upside rotation that brings households across the access threshold can therefore change participation appreciably while changing mean current consumption by only the consumption-financed portion of the additional fees. Equal variance is not needed for this identity; the specified equal-mean, equal-variance rotation above supplies an example when its support straddles the new threshold.

The identity concerns mean current consumption, not its variance or downside risk. Although the two earnings environments have equal mean resources, they can expose households to different consumption outcomes. Saving and future earnings also affect that exposure in a life-cycle problem. Section 5.1 therefore examines financial participation and observed expenditure together, while Section 6 allows consumption, safe saving and stockholding to adjust jointly over time. An IRA or

private-annuity contribution is another discrete financial decision, although it need not increase equity exposure. The example allows a small consumption response but predicts neither exactly unchanged spending nor increased total saving, and leaves the allocation within retirement accounts unspecified.

The threshold logic also motivates examining occasional durable purchases. Transaction costs or indivisibilities can create an inaction region in which a household postpones a purchase (Eberly, 1994). When a favorable earnings realization carries resources into the purchase region, the timing of expenditure can change without a proportional increase in every spending category. This is an interpretation of the furnishings evidence in Section 5.1, rather than a durable-choice result derived from the present model. Its sign depends on the earnings distribution and the location of the purchase threshold, and it does not imply unchanged nondurable spending.

3 Data and Empirical Design

3.1 Earnings histories and stock market participation

I combine earnings histories from the Current Population Survey (CPS) with subsequent portfolio observations from the Survey of Income and Program Participation (SIPP). The CPS measures historical earnings distributions for different groups of workers; the SIPP records portfolio transitions among workers assigned to those groups. Using separate surveys prevents a worker’s own earnings realization from mechanically determining the assigned risk measure.

The main CPS sample contains 88,862 validated annual earnings histories, including 5,490 with zero destination wages. Workers are aged 25–65 and have positive wages at the origin. Origins run from 2013 through 2017 and destinations through 2018, before the SIPP portfolio window. I use the public Census ASEC amounts at both endpoints and retain the redesigned 2014 questionnaire subsample. The wage measure includes transitions out of wage earnings, whether through nonemployment or a move to another income source; a zero wage is not a separately identified dismissal. Total labor earnings sum wages, net business earnings, and net farm earnings. The original positive-support comparison uses 83,372 histories with positive wages and total earnings at both endpoints.

I group workers by harmonized industry and education. Within each of the 30 supported industry–education groups, the age bands 25–44 and 45–65 yield 58 supported earnings cells. The group determines the age profile of earnings growth; the finer cell determines the distribution around that profile.

The SIPP reports directly held stocks and mutual funds, other financial assets, and retirement balances. Let R_{it} denote stock-and-mutual-fund balances, F_{it} financial assets excluding retirement balances, and $s_{it} = R_{it}/F_{it}$ the nonretirement risky share. The public records do not reveal the asset composition of retirement accounts, and mutual funds need not consist entirely of equities. I aggregate eligible monthly records to annual observations and link consecutive years. The observations come from the 2020–2025 surveys: portfolio baselines refer to 2019–2023 and outcomes to 2020–2024.

These releases follow the revised asset processing introduced with the 2018 SIPP, including model-based imputation of stock and mutual-fund ownership (U.S. Census Bureau, 2020). Section 4.7 compares the results with the earlier SIPP panel and with pre-pandemic PSID portfolios.

The entry sample contains 27,419 transitions beginning without stock or mutual-fund holdings. The prospective allocation sample requires positive broad and nonretirement financial denominators and usable ratios in $[0, 1]$ at both dates. It contains 31,569 transitions: 24,235 initial nonholders and 7,334 initial holders, including 2,165 who subsequently exit. Missing financial assets are not coded as zero shares. Table 1 describes this sample by initial ownership.

Baseline ownership is 23 percent in the analysis sample, compared with 19 percent of 2019 Survey of Consumer Finances (SCF) households with a head aged 25–64. The corresponding rates for college graduates are 34 percent in both surveys; for those without a degree they are 13 and 10 percent. The surveys differ in their population and sample selection. The SIPP unit is a worker transition involving household-reported assets. Changes in balances can reflect transactions, valuation, or household balance-sheet changes; they are not direct observations of stock purchases and sales.

The PSID extends the analysis to spending and contributions to retirement-saving vehicles. I link the same reference person across adjacent biennial waves in 2005–2019, retaining 33,353 spending transitions among 9,183 people aged 25–65 at baseline. Official annualized expenditure components distinguish food, selected nondurable and service spending, and furnishings. Contributions are any reported money put into a private IRA or annuity, not the equity share in those accounts or a total saving rate. The regressions use baseline family weights and control for the earlier outcome, family income, liquid assets, homeownership, family size and demographics, with separate coefficients by initial non-IRA stock ownership. Section 5.1 explains the timing.

3.2 Age-adjusted earnings risk

The main measure uses nominal arc wage growth,

$$g_{i,t+1} = \frac{2(y_{i,t+1} - y_{it})}{y_{i,t+1} + y_{it}},$$

which equals -2 when destination wages are zero. Following the earnings-profile adjustment of Catherine, Sodini, and Zhang (2024), I remove a cubic in origin age separately within each industry–education group:

$$g_{i,t+1} = \theta_{0g} + \theta_{1g}a_{it} + \theta_{2g}a_{it}^2 + \theta_{3g}a_{it}^3 + \epsilon_{i,t+1}. \quad (7)$$

The regression uses CPS person weights. The residual is earnings growth net of the group’s age profile. I apply the same procedure to total earnings and, for the original positive-support comparisons, to nominal arc growth and real log growth. Real log growth removes annual CPI-U inflation. The adjustment follows CSZ’s method, but the object differs: historical pretax earnings composition, rather than the comovement of after-tax disposable income with stock returns.

Table 1. Prospective portfolio sample by initial ownership

	All		Nonholders		Holders	
	Mean	SD	Mean	SD	Mean	SD
Age (years)	43.74	11.27	43.56	11.31	44.34	11.09
Female (%)	47.47	49.94	49.82	50.00	39.71	48.93
College graduate (%)	49.99	50.00	43.12	49.52	72.65	44.57
Nonretirement financial assets (\$1,000)	78.70	317.76	27.67	107.26	247.14	599.47
Retirement balances (\$1,000)	129.46	411.21	83.29	310.00	281.89	616.26
Stocks and mutual funds (\$1,000)	33.96	222.05	0.00	0.00	146.06	442.40
Baseline risky share (%)	10.76	25.12	0.00	0.00	46.27	32.73
Subsequent risky share (%)	11.18	25.77	4.27	16.86	33.99	35.22
Subsequent participation (%)	23.52	42.41	9.60	29.46	69.49	46.05
Person-year pairs	31,569		24,235		7,334	

Notes: Weighted means and standard deviations in the retained prospective sample, using outcome-year SIPP person weights. Ownership is measured at the baseline in 2019–2023; subsequent outcomes are in 2020–2024. Other characteristics and balances refer to the baseline. Dollar balances are nominal. The risky share is stocks and mutual funds divided by financial assets excluding retirement balances. Both dates have positive usable financial denominators and shares in $[0, 1]$; future zero holdings remain. The table describes selected consecutive person-year pairs, not a cross-sectional population of all households.

For each industry–education–age cell c , let μ_c be mean residual growth. The lower and upper centered second partial moments are

$$V_c^- = \mathbb{E}_c[(\epsilon - \mu_c)^2 \mathbf{1}\{\epsilon < \mu_c\}], \quad V_c^+ = \mathbb{E}_c[(\epsilon - \mu_c)^2 \mathbf{1}\{\epsilon > \mu_c\}]. \quad (8)$$

Total variance is $V_c = V_c^- + V_c^+$ and upside composition is $C_c^{\text{up}} = V_c^+ - V_c^-$. At fixed variance, higher composition reallocates variation from below to above mean growth. A specification using only total variance cannot distinguish these distributions. At fixed variance, an increase in composition combines a larger upper partial variance with a smaller lower partial variance. Section 4.2 compares the entry association of composition with specifications based on total variance.

Zero-inclusive wages lead the assigned-risk analysis because they measure cell composition more reliably and reproduce the cross-cell ranking in the PSID. I use the same wage measure for every portfolio outcome and report the positive-support measures for comparison. Section 4.10 reports the evidence behind this choice and when it was made.

3.3 Participation and allocation regressions

Stock market participation means reported ownership of stocks or mutual funds outside retirement accounts. Let H_{it} indicate that ownership. Entry is starting participation: $H_{it} = 1$ conditional on $H_{i,t-1} = 0$. I estimate

$$H_{it} = \beta_E \tilde{C}_{c(i,t-1)}^{\text{up}} + Z'_{i,t-1} \gamma_E + u_{it}, \quad H_{i,t-1} = 0. \quad (9)$$

A tilde denotes standardization using the construction's fixed source scale across CPS cells. For zero-inclusive arc growth and positive-support real log growth, I use the mean and standard deviation of the pre-adjustment cell moments, weighted by Kish effective cell size capped at 1,000. Positive-support nominal arc growth retains the fixed scales of the original positive-support design. The scales are unchanged across outcomes and initial ownership states. Consequently, the reported coefficient compares a one-standard-deviation increase in composition using this fixed scale, not a standard deviation of the variation remaining after controls.

Z includes variance, standard deviation, mean growth and its square, household controls, and outcome-year and baseline industry, education, and age-band effects. Including both variance and standard deviation permits curvature in dispersion. Household controls are age and its square, sex, log gross assets, transformed first-listed-job earnings, and an indicator for any eligible month with nonpositive first-listed-job earnings.¹ All main SIPP regressions use outcome-year person weights.

For allocation, every risk and control coefficient may differ by initial ownership:

$$\begin{aligned} s_{it} = & (1 - H_{i,t-1})(\beta_0 \tilde{C}_{c(i,t-1)}^{\text{up}} + Z'_{i,t-1} \gamma_0) \\ & + H_{i,t-1}(\beta_1 \tilde{C}_{c(i,t-1)}^{\text{up}} + Z'_{i,t-1} \gamma_1 + \phi s_{i,t-1}) + u_{it}. \end{aligned} \quad (10)$$

The common scaling permits a direct test of $\beta_1 - \beta_0$. Holders remain in their initial group if they exit, so the regression does not condition on continued participation. With the initial share controlled, the holder coefficient is also the composition coefficient in a regression of share changes. Replacing s_{it} with H_{it} yields entry and retention coefficients on the same sample.


Main standard errors cluster by the 30 industry–education groups, with CR1 correction and a t_{29} reference distribution; restricted samples use their actual cluster count. Stars report unadjusted two-sided tests. The analysis is exploratory. Predating the main portfolio window separates assigned earnings risk from contemporaneous own realizations, but does not make industry or education assignment exogenous. The coefficients are conditional associations; Section 4.8 examines the sorting explanations most relevant to them.

4 Earnings Composition and Stock Market Participation

Workers assigned greater upside composition are more likely to enter stocks and mutual funds, including entry followed by sustained ownership. I first report the magnitude and compare composition with total variance, then examine differences across households, portfolio outcomes, and surveys. The main estimates use zero-inclusive wage growth; positive-support measures provide comparisons.

¹For mean monthly job-1 earnings $\bar{y}^{(1)}$, the income control is $\log(\max\{\bar{y}^{(1)} + 1, 1\})$. Job 1 is the first-listed job, not necessarily the largest. The nonpositive-earnings indicator is not a separately observed unemployment spell.

Figure 2. Earnings opportunities and portfolio choice by initial ownership



Notes: Panel A reports starting stock market participation among initial nonholders. Panel B reports subsequent risky share by initial ownership and the directly estimated holder–nonholder difference. Points report composition coefficients per fixed pre-adjustment source standard deviation, holding total variance and the baseline controls fixed. The measure is the zero-inclusive wage construction of Section 3.2. Outcomes are dated 2020–2024. Entry has 27,419 initial-nonholder transitions; subsequent-share estimates use 24,235 initial nonholders and 7,334 initial holders, including later exits. Bars are 95% CR1/ t_{29} intervals using 30 industry–education clusters, conditional on estimated CPS moments. Tables 2 and 5 report the specifications.

4.1 Starting stock market participation

At a given total earnings variance, a greater share of variation above mean growth predicts more entry. In Table 2, a one-standard-deviation increase in upside composition predicts 1.601 percentage points more entry across 27,419 baseline nonholder transitions. This is nearly a fifth of the 8.82 percent annual entry rate.

The positive association appears under all four earnings constructions. Zero-inclusive total earnings give 1.042 points; the original positive-support constructions give 0.875 point with real log growth and 0.695 with nominal arc growth. Observations, controls, and inference are unchanged across columns. Wild-cluster bootstrap inference supports the wage and positive-support entry results, with weaker evidence for zero-inclusive total earnings. This comparison addresses the modest number of industry–education groups.

The magnitudes depend on the increment in composition being considered. The reported coefficients use its unconditional standard deviation across CPS cells. Scaling instead by the standard deviation remaining after the other regressors gives 0.590 point for wages, 0.506 for total earnings, and 0.573 and 0.534 for the positive-support real-log and nominal-arc constructions. For comparison, in the prospective regression using the positive-support real-log construction, a standard deviation of log financial assets predicts 3.2 points more entry and a standard deviation of log income 1.0 point, against 0.97 point for composition.

Table 2. Earnings composition and starting stock market participation

	(1)	(2)	(3)	(4)
	Zero-inclusive wages Wages Arc, zeros kept	Zero-inclusive total earnings Total Arc, zeros kept	Original real log Total Real log, positive	Original nominal arc Total Arc, positive
Earnings Growth				
Earnings origin years	2013–2017	2013–2017	2013–2017	2013–2017
Upside composition	1.601*** (0.506)	1.042** (0.405)	0.875*** (0.279)	0.695*** (0.248)
Conditional-SD magnitude (pp)	0.590	0.506	0.573	0.534
Split-half reliability of composition	0.73	0.62	0.28	0.24
Group-specific cubic age adjustment	Yes	Yes	Yes	Yes
Risk and baseline controls	Yes	Yes	Yes	Yes
Industry, education, age-band, year FE	Yes	Yes	Yes	Yes
Nonholder transitions	27,419	27,419	27,419	27,419
Observed entry rate (%)	8.82	8.82	8.82	8.82
Assignment cells / inference groups	58 / 30	58 / 30	58 / 30	58 / 30

Notes: Weighted linear probability models of next-year stock-or-mutual-fund entry among baseline nonholders, using outcome-year SIPP person weights; baselines are 2019–2023 and outcomes 2020–2024. All four constructions use CPS earnings origins in 2013–2017 and destinations in 2014–2018, before every portfolio observation. Columns (1)–(2) use nominal arc growth of wages and of total earnings that keeps transitions to zero earnings (a zero destination contributes –2); columns (3)–(4) use the original construction on positive endpoints (Section 3.2). Survey-weighted group-specific origin-age cubics are removed before constructing centered moments in 58 industry–education–age cells; composition is upper minus lower partial variance. Coefficients are percentage points per fixed pre-adjustment source SD of each construction; the conditional-SD row uses residual composition after controls. Variance, SD, mean growth and its square are controlled; Section 3.3 specifies household controls and FE. Split-half reliability is the Spearman–Brown reliability of the cell composition across the 58 cells. Stars use unadjusted two-sided tests with CR1 standard errors clustered by 30 industry–education groups and t_{29} : *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Inference conditions on the CPS moments. Wild-cluster bootstrap p -values for the composition coefficients (null imposed, 1,999 Rademacher draws over the 30 groups) are 0.039, 0.080, 0.022 and 0.019 for columns (1)–(4).

4.2 Composition beyond total variance

A linear specification containing total variance but omitting composition and the standard deviation gives an entry coefficient of -0.457 percentage point per source standard deviation (SE 0.584). Entering composition instead gives 1.091 points (SE 0.532). In the full specification, which includes composition, variance, and the standard deviation, the composition coefficient is 1.601 points (SE 0.506). Upside composition thus predicts entry conditional on earnings dispersion.

Dispersion also has a nonlinear relationship with entry. In the full wage regression, the coefficients on V and $\sqrt{V}$ are -15.380 (SE 3.505) and 15.762 (SE 3.456) percentage points per respective source standard deviation, both with $p < 0.001$. Together they describe a concave variance–entry relationship. The imprecise coefficient in the linear variance-only specification therefore does not establish that dispersion is unrelated to entry. The composition coefficient measures the association with a more upside-oriented earnings distribution at a given level of dispersion.

Table 3 asks whether the positive composition association depends on allowing curvature in variance. Column (1) retains the baseline specification. Column (2) drops only $\sqrt{V}$, leaving composition and total variance with the same sample, weights, controls and fixed effects. Composition predicts 1.111 percentage points more participation among initial nonholders (SE 0.526, $p = 0.044$), while the variance coefficient is 0.053 point (SE 0.548). The positive association therefore remains in the simpler linear specification, with a smaller magnitude and less precision.

Column (3) writes that linear specification directly in the upper and lower partial variances, without composition, total variance or its square root as additional regressors. An increase of 0.01 in upper partial variance, holding lower partial variance fixed, predicts 0.331 percentage point more participation (SE 0.167, $p = 0.056$); the corresponding lower-partial-variance coefficient is -0.324 point (SE 0.154, $p = 0.044$). The common increment is measured in squared units of arc growth. Columns (2) and (3) are equivalent linear specifications with identical fitted values, not independent evidence. These coefficients come from the linear comparison, rather than local derivatives of the nonlinear baseline. Their signs support the distinction between favorable and adverse earnings variation, while the change from 1.60 to 1.11 points shows why the magnitude should not be treated as invariant to the variance specification.

Composition and conventional skewness capture closely related asymmetry. Their partial correlation is 0.91 after controlling for wage dispersion, against an unconditional correlation of 0.05; split-half reliability is 0.73 for composition and 0.39 for skewness. Keeping the headline variance and SD controls, replacing composition with skewness gives 0.91 point of entry. Entering both gives a composition coefficient of 2.58 points, with both slopes imprecise. The 1.32-point composition estimate with skewness omits the variance and SD controls. Partial variances are more reliably measured and sum exactly to total variance. The joint regression does not establish that composition contains information beyond skewness.

4.3 Economic drivers of earnings asymmetry

Leaving wage work helps account for the cross-cell pattern in earnings risk. The share of CPS histories ending at zero wages correlates at -0.855 with upside composition and 0.924 with total variance. Cells with more zero destinations therefore tend to combine greater variance with more variation below the mean. This is a descriptive feature of the measured distribution: zero annual wages can reflect nonemployment or a move to another income source and do not separately identify dismissal.

The participation association is not absorbed by the frequency of zero wages. Adding the cell's zero-destination share to the full design leaves a composition coefficient of 2.420 percentage points (SE 0.911); the zero-share coefficient is imprecise at 2.630 points (SE 2.030). The zero-share-only coefficient is also imprecise, at -1.163 points. Zero-wage transitions help describe the variation in earnings environments, but their frequency alone does not account for the positive composition association. This comparison differs from the dismissal-risk quasi-experiment of [Ahn, Jo, and Lam \(2026\)](#).

Table 3. Earnings composition, variance, and direct tail coefficients

	(1) Curved variance	(2) Linear variance	(3) Direct tails same fit as (2)
Upside composition C (per source SD)	1.601*** (0.506)	1.111** (0.526)	
Total variance V (per source SD)	−15.380*** (3.505)	0.053 (0.548)	
Standard deviation $\sqrt{V}$ (per source SD)	15.762*** (3.456)		
Upper partial variance V^+ (per 0.01 growth-squared unit)			0.331* (0.167)
Lower partial variance V^- (per 0.01 growth-squared unit)			−0.324** (0.154)
Mean growth and its square	Yes	Yes	Yes
Baseline controls and fixed effects	Yes	Yes	Yes
Initial-nonholder transitions	27,419	27,419	27,419
Industry–education clusters	30	30	30
Annual entry rate (%)	8.82	8.82	8.82

Notes: Weighted linear probability models of next-year stock-or-mutual-fund participation among initial nonparticipants; coefficients are percentage points. The SIPP observations, outcome-year person weights, household controls and industry, education, age-band and year effects are identical across columns, with baselines in 2019–2023 and outcomes in 2020–2024. Risk measures use zero-inclusive, age-adjusted CPS wage growth from 2013–2017 origins. Column (1) is the main specification. Column (2) omits only $\sqrt{V}$. Column (3) reports the coefficients on V^+ and V^- of that same linear specification: because $C = V^+ - V^-$ and $V = V^+ + V^-$, columns (2) and (3) have identical fitted values. Thus column (3) is neither an additional regression nor a local derivative of column (1). Composition and variance in columns (1)–(2) use their fixed source SDs, 0.033913 and 0.135425; the tail rows use the common raw increment 0.01, rather than tail-specific SDs. Standard errors in parentheses use industry–education CR1 clustering and t_{29} inference, conditional on the CPS moments. Stars denote unadjusted two-sided tests against zero: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The linear comparison was added in response to the functional-form question; it does not replace the main specification.

Measured cyclical exposure does not absorb the relationship either. Adding the cell’s mean wage growth in 2020 gives a composition coefficient of 1.783 points (SE 0.528). Controlling separately for historical PSID exposures of mean growth, lower partial variance, or composition to changes in national unemployment gives 1.612, 1.701 and 1.662 points; including all three gives 2.104 points (SE 0.452). The exposures use biennial PSID cell moments over 1999–2019. These comparisons retain the headline sample, controls and weights. They distinguish the association from these measured cyclical exposures without identifying a causal contribution of particular labor-market events.

The evidence is less decisive about persistent versus transitory earnings changes. An exploratory PSID variance-decomposition proxy initially gives a larger composition association in cells with a greater transitory share, but the interaction becomes smaller and imprecise after excluding very large log-wage changes. These results do not establish that promotions, persistent job losses or transitory shocks drive a particular share of the participation association. The life-cycle model separately examines how realized earnings and the anticipated earnings distribution can affect resources, saving and participation.

4.4 Entry and sustained market access

The entry association extends to ownership at later annual observations. Table 4 divides entry into two joint outcomes: entry followed by sustained ownership and entry followed by an interruption. Within each panel, all outcomes use the same initial nonholders, weights and controls. Both outcomes are measured among all eligible initial nonholders, so differences do not arise from selecting different groups of entrants.

Table 4. Earnings opportunities and sustained market access

	(1) Entry	(2) Sustained entry	(3) Interrupted entry
<i>A. One annual follow-up</i>			
Upside composition	2.086*** (0.697)	0.978* (0.501)	1.108** (0.512)
Outcome mean (percent)	9.26	4.74	4.53
Observations	9,236	9,236	9,236
Sustained minus interrupted		-0.131 (0.735)	
Conditional-survival gradient		-0.959 (4.036)	
<i>B. Two annual follow-ups</i>			
Upside composition	3.315** (1.430)	1.894** (0.786)	1.421 (1.053)
Outcome mean (percent)	10.17	3.53	6.63
Observations	3,254	3,254	3,254
Sustained minus interrupted		0.473 (1.187)	
Conditional-survival gradient		7.296 (5.674)	
Controls and fixed effects	Yes	Yes	Yes
Industry–education clusters	30	30	30

Notes: Initial nonholders with ownership observed at entry and each required annual follow-up. Panel A covers entry years 2020–2023; Panel B covers 2020–2022 and a smaller complete-follow-up cohort. Sustained ownership requires holdings at every follow-up; interrupted ownership includes any observed interruption after entry. Nonentrants have zero for all three outcomes. Coefficients are percentage points per one-standard-deviation increase in predetermined zero-inclusive wage composition, holding variance fixed. All models retain the controls, fixed effects and outcome-year survey weights of equation (9). Columns (2) and (3) add to column (1); the first comparison subtracts their coefficients. The survival gradient is the local derivative of fitted sustained entry divided by fitted entry at the weighted mean design, not a finite one-standard-deviation change. Both comparisons use full joint covariance, including fitted-level uncertainty for the derivative. Standard errors in parentheses use CR1 clustering by 30 industry–education groups, conditional on the assigned CPS moments. Stars use unadjusted two-sided t_{29} tests: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Follow-up availability is related to composition; missing follow-up is not coded as exit. The panels do not identify a duration effect.

In Panel A’s 9,236 initial-nonholder transitions, a one-standard-deviation increase in upside composition predicts 2.086 percentage points more entry. Of this, 0.978 point is entry followed by ownership at the next annual observation, and 1.108 points is entry followed by exit. The sustained

component is suggestive, and the difference of -0.131 point between the two components is imprecise. The data do not establish that the association is concentrated in short spells.

Panel B follows 3,254 initial-nonholder transitions through two post-entry observations. Composition predicts 1.894 percentage points more entry with ownership at both dates. The remaining entry coefficient is 1.421 points, including any interruption or exit followed by reentry. This smaller, earlier complete-followup cohort extends the finding to successive annual dates; its larger coefficient does not establish an increasing association with duration. Annual observations also cannot establish uninterrupted ownership between interviews.

More sustained entry need not mean higher survival among entrants. The conditional-survival estimates in Table 4 are imprecise at both horizons and establish neither a survival advantage nor a penalty. As equation (5) illustrates, more households can enter and remain invested without a detectable change in entrants' survival rate.

Follow-up selection limits these exploratory comparisons. Missing ownership is not coded as exit, and follow-up must meet the worker-sample eligibility rules. Composition predicts eligible-followup availability by 1.820 percentage points, so complete cases need not represent all initial nonholders. Requiring reported ownership and balances at all three dates lowers the entry estimate to 0.921 point; both entry and sustained entry are then imprecisely estimated. This is a different restriction from the headline's two-date reporting checks. Excluding sequences with a first follow-up in 2022 retains positive entry and no established survival penalty. Composition predicts entry followed by sustained ownership; its association with survival conditional on entry remains imprecise.

4.5 Participation costs and household resources

Participation costs make the resource distribution relevant for entry. In the model, favorable earnings realizations allow households near a resource threshold to enter. Financially wealthy nonholders provide a comparison: a fixed participation cost is small relative to their resources.

The positive-support real-log and nominal-arc constructions give larger entry coefficients among less wealthy nonholders. Between the 25th and 75th percentiles of financial assets, the real-log coefficient declines from 1.97 to 1.27 percentage points and the arc-growth coefficient from 0.68 to 0.18 point. The wealth interactions are imprecise, so the estimates are consistent with a participation threshold but also allow similar associations across wealth levels.

Education gives a similar ordering. Among workers without a college degree, whose entry rate is 5.4 percent, the wage-composition coefficient is 1.98 points. Among college graduates, whose entry rate is 13.9 percent, the estimate is imprecise at 0.13 point. The wage difference of -1.85 points is also imprecise; the positive-support real-log difference of -1.34 points is only marginally significant. The larger estimate among workers without a degree is consistent with participation costs mattering more at low resources, although education also differs with other household characteristics.

Age gives no uniform ordering within the SIPP. In the positive-support interaction design, real log growth gives a larger entry coefficient among younger households, but nominal arc growth gives a larger coefficient among older households. Separate wage regressions give positive coefficients in

both bands, 2.65 points at ages 25–44 and 2.18 at ages 45–65. The estimated age pattern therefore depends on the earnings measure. Section 4.7 compares these age-band results with the PSID.

Observed earnings progression does not absorb the entry association. On 23,261 initial nonholder transitions with annual personal earnings at both dates, adding baseline earnings and the realized change moves the wage-composition coefficient from 1.518 to 1.574 percentage points. Separate gain and loss slopes give 1.569 points. The samples are identical and retain zero and negative destination earnings. The corresponding coefficient changes are 0.056 and 0.051 point. Since realized earnings may themselves reflect the earnings environment, these conditional comparisons do not identify a direct risk effect or households' perceptions.

Entry accompanied by a large earnings gain is one part of this relationship. Defining a gain as at least 25 percent growth in real annual personal earnings, it accounts for 0.546 point of the matched sample's 1.518-point composition coefficient; other entry accounts for 0.972 point, and their difference is imprecise. These joint outcomes do not identify mediation: gain-entry can rise because entry rises. In the model, gains and composition are complementary explanations, since composition changes the probability of crossing a resource threshold at fixed mean and variance.

The income-reversal explanation for short spells remains unconfirmed. Among 802 first observed entrants with complete follow-up, mean real earnings rise after entry for both exiters and retainers. Exiters' change exceeds retainers' by \$3,001, but the difference is imprecise. Among entrants with a real earnings gain of at least 25 percent, retaining that income elevation predicts an imprecise -0.006 percentage point difference in retention. The composition interaction with pre-entry nonretirement financial assets is also imprecise. The earnings profiles do not establish income reversal among entrants who exit. They concern a selected group of entrants, whereas the sustained-entry estimates compare all eligible initial nonholders.

4.6 Initial holders: allocation and exit

Does greater entry accompany greater investment by existing holders? Table 5 allows all risk coefficients, controls, and fixed effects to differ between 24,235 initial nonholders and 7,334 initial holders. Baseline share also enters the holder equation.

The participation associations differ by initial ownership. Wage composition predicts 1.646 percentage points more entry and 3.720 points less holder retention, with a directly estimated retention-minus-entry contrast of -5.365 points. The positive-support real-log construction gives entry of 0.969 point, retention of -2.650 points, and a contrast of -3.619 points. The wage retention association is not distinguishable from zero under the wild-cluster bootstrap, whereas the positive-support real-log association remains significant. Greater entry therefore coexists with lower estimated retention among initial holders.

The wage measure does not establish a decline in holders' risky shares. It predicts a 0.622-point higher subsequent share among nonholders and a 0.514-point lower share among holders; the latter and their difference are imprecise. The estimates are too imprecise to rule out a moderate decline. Positive-support real log growth gives 0.360 and -1.621 points, with a difference of -1.980 points.

Table 5. Earnings composition and subsequent portfolio allocation

Composition	(1) Zero-inclusive wages	(2) Original real log
<i>Outcome: next-year nonretirement risky share</i>		
Upside composition: baseline nonholders	0.622* (0.345)	0.360** (0.142)
Upside composition: baseline holders	−0.514 (0.941)	−1.621** (0.715)
Difference: holder minus nonholder	−1.136 (1.154)	−1.980** (0.737)
<i>Outcome: next-year ownership</i>		
Nonholder entry	1.646** (0.599)	0.969*** (0.302)
Holder retention	−3.720** (1.708)	−2.650*** (0.952)
Difference: holder minus nonholder	−5.365** (2.098)	−3.619*** (1.070)
Baseline-nonholder observations	24,235	24,235
Baseline-holder observations	7,334	7,334
All risk and other controls × baseline status	Yes	Yes
Baseline share in holder block	Yes	Yes

Notes: Weighted OLS on the prospective sample of 31,569 consecutive person-year pairs (baselines 2019–2023, outcomes 2020–2024) with positive financial denominators and valid ratios in $[0, 1]$ at both dates; genuine future zeros, including exits, remain. The share is stock-and-mutual-fund balances divided by financial assets net of retirement balances. Every risk coefficient, household control and fixed effect is interacted with baseline ownership, and the baseline share enters the holder block. Coefficients are percentage points per fixed source SD of upside composition under each construction (Table 2); the difference row tests the holder-minus-nonholder slope directly using the saved covariance. Weights are destination-year person weights. Parentheses give CR1 standard errors clustered by 30 industry–education groups (t_{29} , conditional on the CPS moments): *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Positive-support nominal arc growth also predicts a higher nonholder share, but not a precise holder response. These regressions retain future zeros, so even the positive-support real-log construction’s negative holder coefficient combines exit, reallocation, and valuation changes rather than identifying sales by continuing investors.

Exiters differ substantially from continuing holders. Their median risky balance is \$12,000 rather than \$25,000, and their median risky share is 0.375 rather than 0.456. Among those with prior history, 49.6 percent of exiters entered in the previous year, compared with 22.8 percent of stayers.

Comparisons between recent entrants and established holders depend on the earnings construction and are inconclusive for zero-inclusive wages. Established holders’ positive-support real-log share estimate is too imprecise to rule out a modest decline.

Short spells are common. The exit hazard is 46.6 percent for last year’s entrants, 33.0 percent in the second year, and 15.9 percent after three or more years, an ordering also documented by [Galaasen and Raja \(2024\)](#). Even the strictest reporting sample has 41.4 percent of entrants leave within a year. The composition coefficients for entry followed by exit and for persistent entry are

not precisely distinguishable. Neither this comparison nor initial-holder retention establishes lower survival among upside-related entrants; Table 4 shows that wage composition also predicts sustained entry.

Most exits do not coincide with the acquisition of another large risky asset: 2,042 of the 2,165 exits involve neither business entry nor home purchase. Their small starting balances and infrequent business entry or home purchase suggest that many exits close small stock positions.

4.7 External evidence: the PSID and pre-pandemic surveys

Pre-pandemic PSID portfolios provide an independent test of whether the entry pattern is specific to the SIPP or to the years after COVID-19. The PSID records previous-calendar-year labor income and biennial direct, non-IRA stock ownership ([Panel Study of Income Dynamics, 2023](#)). Its 1999–2019 reference-person sample provides 45,840 two-year transitions at ages 25–65. The CPS earnings histories used for the ranking overlap or postdate many of these earlier portfolio transitions. The clearest agreement with the SIPP is among younger households. The initial design preceded estimation; the age-band and window comparisons followed the pooled results and are exploratory.

Among households aged 25–44, upside composition predicts entry in both surveys. Wage coefficients are 2.65 percentage points in annual SIPP transitions and 3.56 points in PSID transitions with 2009–2017 baselines. Original real log growth gives 2.33 and 1.31 points; zero-inclusive total earnings give 2.95 and 2.13. Restricting the PSID to 2013–2017 baselines retains positive entry associations of 2.90 points for wages and 1.50 for original real log growth. Young-household entry thus has a positive composition association in both surveys, including before the pandemic.

Ownership levels extend the comparison to other pre-pandemic sources. Among households under 45, wage composition predicts 3.11 points more ownership in the SIPP and 6.32 in the PSID, against rates of 23 and 18 percent. In the 2019 SCF, it predicts 5.0 points more ownership of any equity, including retirement accounts, and 3.1 points more direct stock-or-stock-fund ownership among working heads aged 25–64; the direct-ownership association is marginally significant. In SIPP 2014 reference years 2013–2016, ownership among workers under 45 is positive in all four constructions, including 1.80 points with wages and 1.43 with original real log growth.

The PSID also reproduces the cross-cell ranking of zero-inclusive wage composition. Of 50,013 candidate PSID growth pairs, 43,805 contribute to pooled wage moments in the 58 matched cells. Their composition correlates at 0.33 with CPS composition, or 0.47 after correction for the two split-half reliabilities. The recent window gives raw and corrected correlations of 0.33 and 0.53. Original real-log composition is not reproduced: correlations are -0.19 pooled and -0.32 in the recent window. Including zero wages thus changes which earnings environments are ranked as having greater upside composition.

The comparison is strongest for younger households' entry and ownership. At ages 45–65, ownership estimates are imprecise and PSID entry estimates are small, whereas SIPP wage and total-earnings entry coefficients are positive, at 2.18 and 2.96 points. The all-age PSID entry coefficient is an imprecise -0.18 point, against 1.60 in the main SIPP sample. Holder transitions also differ:

Table 6. Entry and stock ownership by age band in two surveys

	(1) Zero-inclusive wages 2013–2017	(2) Zero-inclusive total earnings 2013–2017	(3) Original construction positive-support real log
<i>A. Entry, percentage points per source SD of composition</i>			
SIPP, ages 25–44 (annual, 2019–2024)	2.648*** (0.526)	2.946*** (0.350)	2.331*** (0.483)
PSID, ages 25–44 (two-year, 2009–2017 baselines)	3.564*** (0.728)	2.134*** (0.448)	1.308*** (0.377)
SIPP, ages 45–65	2.177*** (0.772)	2.963*** (0.880)	0.424 (0.530)
PSID, ages 45–65 (two-year, 2009–2017 baselines)	0.442 (1.302)	–0.024 (1.077)	–0.036 (0.634)
<i>B. Stock market ownership at the baseline, percentage points per source SD</i>			
SIPP, ages 25–44	3.109* (1.710)	2.791** (1.024)	1.803 (1.110)
PSID, ages 25–44	6.317*** (0.809)	4.893*** (0.615)	2.026*** (0.515)
SIPP, ages 45–65	3.083 (2.425)	–1.388 (2.258)	0.112 (1.293)
PSID, ages 45–65	1.801 (2.163)	–0.047 (1.624)	–0.238 (0.829)

Notes: SIPP rows refit the entry design of Table 2 (Panel A) and a baseline-ownership regression with the same controls (Panel B) within age bands; ownership refers to stocks and mutual funds. PSID rows assign the CPS moments to reference persons aged 25–65 in the public PSID family files and fit the prospective design to two-year transitions in imputed direct (non-IRA) stock ownership (Panel A) and baseline ownership on the same moments with household controls, wave and industry effects (Panel B, baselines 1999–2017). Percentage points per fixed source SD of composition; CR1 standard errors clustered by the 29 industry–education groups represented in each age band appear in parentheses. Stars use unadjusted two-sided t_{28} tests conditional on the CPS moments: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The age-band comparison was specified after a pooled null result and is exploratory.

young-holder retention falls with composition in the SIPP but rises in the PSID, which does not reproduce recent-entrant exit. Within-person PSID ownership is positive but marginally significant, at 0.98 point per standard deviation; the sibling estimate is an imprecise –0.57 point. The age comparisons are exploratory and do not establish a sharp age gradient in SIPP wage entry alone.

Comparing annual entry across SIPP vintages requires attention to how asset ownership is recorded. The recent sample follows the Census Bureau’s introduction of model-based imputation for stock and mutual-fund ownership in the 2018 SIPP (U.S. Census Bureau, 2020). These processing changes are a reason to examine survey comparability, rather than attribute every difference between panels to household behavior. They do not by themselves establish that measurement explains the earlier result. In the SIPP 2014 panel, the primary holdout test, specified before estimation for wages at ages 25–44, gives an imprecise –0.50 point, against 2.65 points in the later sample; their 3.14-point difference is significant.

Common annual ownership, income, wealth and record-aggregation definitions give an imprecise -0.15 point in the earlier SIPP and 2.73 points in the later one, with a significant difference of 2.88 points. Three symmetric comparisons—a liquid-asset control, ownership answers or positive balances, and December worker snapshots—also retain significant differences. Reweighting toward a common overlapping population balances earnings, gross and liquid assets, demographics and earnings-group representation. Coefficients become -0.03 point in the earlier panel and 2.60 points in the later panel. The 2.62-point gap is marginally significant, while the change in the gap is imprecise. The earlier SIPP entry result therefore remains different after harmonizing observed definitions and characteristics.

The PSID is especially useful for interpreting this comparison: it supplies positive young-household entry associations in overlapping pre-pandemic years. The SCF and earlier SIPP extend the ownership evidence. Together, these sources show that the relationship between earnings composition and market access reaches beyond the main SIPP sample and its pandemic-era portfolio window.

4.8 Stability and occupational sorting

Common industry-year developments and persistent occupational differences are two alternatives to the earnings-opportunity interpretation. Adding industry-by-year and education-by-year effects leaves the wage entry coefficient at 1.663 percentage points; joint industry–education effects give 1.865 points (Table 7). The latter uses within-group age-band differences, controlling for the common age-band effect. Age-specific sorting into industries remains possible.

Portfolio reporting is a separate concern. Restricting both dates to reported ownership and balance components gives a wage entry coefficient of 1.699 points; additionally requiring self-reported interviews and stable household and residence records gives 2.557 points. Positive-support real-log estimates are 0.699 and 1.008 points. Holder retention becomes smaller and imprecise in both screens; wage holder-share coefficients are -0.480 and -2.275 points, the latter marginally significant. Since these restrictions also select households, their sensitivity does not establish that full-sample exits are erroneous.

Previous switching matters more for exit than entry. In the positive-support real-log construction, controlling for a prior switch on identical observations reduces the exit association by 0.814 point, with little change in entry. Every outcome-year deletion retains positive entry and negative retention, and early- versus late-period differences are imprecise.

Observed risk tolerance does not absorb the PSID composition association. Among reference persons who answered the 1996 hypothetical income gambles, tolerance rises with assigned earnings variance, but its composition coefficient is small and imprecise at 0.05. Adding the tolerance categories moves the composition coefficients by less than 0.1 point. The estimate still allows economically relevant sorting. The within-person ownership estimate is positive but only marginally significant, at 0.98 point per standard deviation, while sibling estimates are imprecise.

A common preference for greater earnings dispersion does not by itself explain the positive composition association conditional on dispersion or its persistence within industry–education groups.

Table 7. Stability of the entry and initial-holder associations

	(1) Nonholder entry	(2) Holder retention	(3) Initial-holder share	(4) <i>N</i>
<i>A. Zero-inclusive wage composition (headline construction)</i>				
Baseline prospective sample	1.646** (0.599)	−3.720** (1.708)	−0.514 (0.941)	31,569
Group-by-year effects	1.663** (0.619)	−4.098** (1.765)	−0.854 (0.946)	31,569
Joint industry–education effects	1.865** (0.875)	−2.362 (1.658)	−0.095 (0.960)	31,569
Reported components	1.699*** (0.449)	−2.054 (1.752)	−0.480 (1.232)	24,114
Reported and stable households	2.557*** (0.555)	−1.457 (2.021)	−2.275* (1.298)	13,939
<i>B. Original real log composition</i>				
Baseline prospective sample	0.969*** (0.302)	−2.650*** (0.952)	−1.621** (0.715)	31,569
Group-by-year effects	0.988*** (0.303)	−2.707** (1.016)	−1.599** (0.740)	31,569
Joint industry–education effects	0.955* (0.469)	−2.713*** (0.758)	−1.345* (0.779)	31,569
Reported components	0.699** (0.303)	−0.776 (1.129)	−0.788 (0.932)	24,114
Reported and stable households	1.008** (0.375)	−0.618 (1.115)	−0.593 (0.887)	13,939
Industry–education clusters	30	30	30	

Notes: Columns (1)–(2) reestimate the next-year ownership regression and column (3) the initial-holder block of the next-year share regression of Table 5; the first row of each panel reproduces the published estimates. Row 2 adds industry-by-year and education-by-year effects by initial ownership on the unchanged sample; row 3 adds joint industry–education effects, so that composition is identified only from the difference between the two age bands within each industry–education group. Rows 4–5 keep transitions whose stock and mutual-fund ownership and balance components are reported rather than allocated at both dates, and additionally interviews coded as self-reported with stable household and residence records; initial holders number 7,334, 4,081 and 2,710. Panel A replays the saved designs with the zero-inclusive wage moments; Panel B is the original construction. Holder retention is next-year ownership among initial holders; exit is its negative with the same standard error. The initial-holder share includes exits. Coefficients are percentage points per fixed source SD of upside composition, with the original controls and destination weights; parentheses are industry–education CR1 standard errors ($G = 30$, t_{29} , conditional on the CPS moments). Stars use unadjusted two-sided tests: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Preferences for upside-oriented earnings, financial sophistication, and age-specific sorting nevertheless remain possible. Earnings cells are not randomly assigned.

4.9 Construction comparisons

Entry is more consistent across constructions than the holder-share association. In the positive-support family, entry remains positive after the group-specific age adjustment for both earnings sources and growth transformations. The real-log holder-share contrast remains negative with a broader financial

denominator, alternative weights, a \$1,000 initial-asset floor, and every group and year deletion. Arc holder-share estimates remain imprecise.

A richer prediction using earnings ranks attenuates the original entry estimates and leaves them imprecise, as are the paired changes from the raw estimates. This adjustment removes more predictable progression than the CSZ age profile. It limits a pure-innovation interpretation without establishing that promotions explain the result.

The different coefficients by initial ownership also matter for the regression specification. A pooled Tobit imposes a common latent equation on zero holdings and positive shares, while the prospective specifications allow these associations to differ.

4.10 Measurement reliability

Including zero earnings substantially improves the measurement of composition. Split-half reliability across CPS cells is 0.73 for wages and 0.62 for total earnings, against 0.28 and 0.24 for the original real-log and arc measures. Extending wage histories through origin year 2021 raises reliability to 0.85 but overlaps the portfolio window, so it is not the headline measure. The choice to lead with wages followed the reliability and external-survey comparisons; it did not precede estimation. Total earnings were initially the primary construction.

Including zero earnings also changes the ranking of cells. Wage composition correlates only 0.22 with original real-log composition and 0.35 with original arc composition, since it also describes transitions out of wage earnings. Composition and variance are more collinear: their correlation is -0.81 for pre-portfolio wages and -0.86 for pooled wages. Higher raw reliability therefore need not deliver smaller regression standard errors. The pooled-wage entry estimate is 1.23 points and imprecise.

Sampling uncertainty in the CPS affects the estimated entry association. Resampling CPS households while holding the SIPP fixed gives positive coefficients in 99.4 percent of 499 wage draws and 97.6 percent of total-earnings draws. Their 2.5th to 97.5th percentiles are 0.13 to 1.51 and 0.03 to 1.11 points. These are descriptive remeasurement ranges, not confidence intervals: the reconstructed coefficients are centered below the original estimates. Joint CPS–SIPP remeasurement of the positive-support family spans zero, and prospective allocation has not been jointly remeasured. Conventional tests remain conditional on the estimated cell moments.

Simulation and extrapolation (SIMEX) illustrates the sensitivity of the magnitude. Adding noise reduces the mean coefficient; linear and quadratic extrapolations give wage estimates of 1.96 and 2.82 points. This suggests attenuation, but depends on the extrapolant and a measurement-error model for jointly estimated moments. Neither is a corrected estimate or a lower bound. I therefore report the uncorrected estimate based on the pre-portfolio wage construction.

5 Saving Participation, Expenditure, and Portfolio Choice

The entry finding raises a broader question about household responses to earnings opportunities: do they accompany higher current spending, greater use of saving vehicles, or changes in existing portfolios? I first compare retirement-contribution participation and expenditure in the PSID, distinguishing furnishings and equipment from more regular spending. The SIPP then separates changes in risky balances from changes in other financial assets, while the SCE relates earnings beliefs to retirement portfolios. These comparisons connect the entry evidence to the joint consumption, saving and participation choices in Section 6.

5.1 Consumption and retirement-saving participation

Earnings risk matters for portfolios through the resources available to finance consumption and investment. Catherine, Sodini, and Zhang (2024) make this connection explicit by constructing lifetime consumption exposure from persistent disposable-income risk and human capital's share of total resources. Their consumption-risk measure is implied by earnings and wealth rather than observed expenditure. The distinction matters here as well: the persistence of earnings shocks, other household income and accumulated assets determine how individual wage risk translates into household consumption risk.

The PSID permits a complementary comparison using observed expenditure and retirement contributions. I ask whether the earnings composition associated with financial participation is also accompanied by a broad expansion in current spending. Table 8 uses the headline wage measure, conditions on the earlier outcome and baseline resources, and allows the full specification to differ between initial stockholders and nonholders. These regressions compare expenditure levels across earnings environments; they do not estimate the risk of consumption growth.

Greater upside composition predicts IRA or private-annuity contribution participation among initial stockholders. A one-standard-deviation increase corresponds to about 6.4 percentage points more participation, or 17% of their mean contribution rate. The nonholder estimate is close to zero and imprecise, and the directly tested difference supports a stronger association among holders. This outcome measures use of retirement-saving vehicles, without observing how their contributions are invested.

The tail signs agree with those for entry. Among initial holders, higher upper partial variance predicts more contribution participation and higher lower partial variance predicts less, each holding the other tail fixed. The dispersion component is small and imprecise. Opposing tail associations therefore also appear in a financial participation decision among households already holding stocks.

Broad expenditure has a small and imprecisely estimated association. The baseline-weight-share average in column (2) is roughly \$450 annually against mean expenditure of about \$57,000. The selected nondurable and service basket and food spending likewise show no clear association, as do the inverse-hyperbolic-sine specifications. Broad expenditure includes financing payments and is not a pure measure of consumption services; the narrower baskets help assess whether the result

Table 8. Earnings opportunities, retirement contributions, and expenditure

	(1) IRA/annuity contribution	(2) Broad expenditure	(3) Selected nondurables and services	(4) Food
Outcome units	Percentage points	2019 dollars	2019 dollars	2019 dollars
Initial nonholders	0.215 (0.331)	504 (449)	124 (273)	53 (91)
Initial holders	6.356*** (2.191)	225 (1,729)	508 (937)	-23 (303)
Holder minus nonholder	6.141** (2.325)	-279 (1,730)	384 (950)	-77 (335)
Baseline-weight-share average	1.365*** (0.407)	452 (519)	196 (296)	39 (83)
Nonholder outcome mean	10.10	51,004	29,885	8,841
Holder outcome mean	37.65	82,694	45,954	11,613
Overall outcome mean	15.26	56,966	32,909	9,362
Nonholder observations	28,337	28,569	28,569	28,569
Holder observations	4,719	4,784	4,784	4,784
Total observations	33,056	33,353	33,353	33,353
Distinct reference persons	9,148	9,183	9,183	9,183
Industry–education clusters	30	30	30	30

Notes: Coefficients per one-standard-deviation increase in wage upside composition, holding variance fixed. PSID baseline waves 2005–2017 are linked to 2007–2019. Initial ownership refers to non-IRA stocks. The fully interacted models include the earlier outcome, baseline resources and demographics, other earnings moments, and industry and wave fixed effects. Baseline family weights also determine the average of the two ownership-state slopes. Column (1) records any private IRA/annuity contribution, not its amount or asset allocation; means are percentages. Its reporting window starts in January of the destination year minus two and overlaps the baseline interview year. Columns (2)–(4) are declared dollar-level checks retaining the transformed earlier outcome. Spending is annualized but mixes current and previous-calendar-year recall periods; the overlapping baskets do not add. Standard errors in parentheses use CR1 clustering by 30 industry–education groups, conditional on the saved CPS moments. Stars use unadjusted two-sided t_{29} tests: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The CPS window overlaps or postdates some observations; these are descriptive external associations.

depends on those components. Taken together, the estimates allow modest spending changes but provide no clear evidence of a general spending expansion.

Furnishings and household equipment. Furnishings expenditure provides evidence on a different spending margin. Among initial stockholders, greater upside composition is positively associated with spending on furnishings and household equipment in the primary transformed specification. Table 9 presents this finding alongside selected nondurables and services and food. The directly tested holder–nonholder difference is positive. The holder association remains similar when realized family income is added, and its upper- and lower-tail derivatives have opposing signs. Furnishings also have a positive dispersion association, so composition is an additional dimension of their relationship with earnings risk.

The spending margin is economically relevant because durable purchases involve an outlay that can be postponed. Models of infrequent durable adjustment link purchase decisions to income

Table 9. Earnings composition, furnishings and other expenditure

	(1) Furnishings and equipment	(2) Selected nondurables and services	(3) Food
Asinh dollar scale	\$1,000	\$10,000	\$10,000
<i>Panel A. Wage composition, holding variance fixed</i>			
Initial nonholders	0.007 (0.013)	0.001 (0.008)	0.002 (0.005)
Initial holders	0.114*** (0.041)	0.012 (0.020)	0.009 (0.017)
Holder minus nonholder	0.107** (0.046)	0.011 (0.024)	0.006 (0.019)
Baseline-weight-share average	0.027** (0.012)	0.003 (0.006)	0.004 (0.004)
<i>Panel B. Composition checks among initial holders</i>			
Add realized family income (asinh)	0.107** (0.039)	0.008 (0.019)	0.006 (0.016)
Positive-support real-log exposure (asinh)	0.027 (0.017)	0.007 (0.010)	0.003 (0.007)
Wage exposure: outcome in 2019 dollars	188 (266)	508 (937)	-23 (303)
<i>Panel C. Wage dispersion among initial holders</i>			
Dispersion, holding composition fixed	0.021*** (0.006)	0.003 (0.004)	0.001 (0.003)
Holder outcome mean (2019 dollars)	2,433	45,954	11,613
Overall outcome mean (2019 dollars)	1,368	32,909	9,362
33,353 observations (28,569 nonholders; 4,784 holders)			
9,183 reference persons; 30 industry–education clusters			

Notes: Controls, weights, timing and inference follow Table 8. Composition is per construction-specific source SD at fixed variance; dispersion raises both partial variances equally at fixed composition using the wage-composition SD. Outcomes are $\text{asinh}(x/s)$ at the displayed scales, except the dollar row, which retains the transformed earlier outcome. Transformed coefficients are not exact percentage changes or directly comparable across baskets. Furnishings measure narrow durable expenditure; the other baskets overlap. Stars use unadjusted t_{29} tests (* $p < .10$, ** $p < .05$, *** $p < .01$). Furnishings holder/average BH q -values are 0.058/0.092 in the twelve ownership-state/six average composition tests; neither family covers dispersion or the ownership difference. No direct cross-outcome difference test was estimated.

growth, uncertainty and the household's position relative to an adjustment threshold (Eberly, 1994). The resource-threshold argument in Section 2.4 suggests a related interpretation: an earnings environment that changes the probability of reaching a resource threshold can matter for occasional purchases as well as financial participation. The positive furnishings association and imprecise nondurable estimates are consistent with this distinction, without establishing a difference in their economic responses.

The finding concerns a narrow category of durable expenditure, not the stock of durables or its service flow. Transformed spending combines whether a purchase occurs with its amount; these margins are not separately estimated. The dollar-level furnishings estimate and the composition estimate using the original positive-support real-log earnings measure are imprecise. Table 9 reports

both checks together with the primary estimates. The association with net non-IRA stock acquisitions is also imprecisely estimated.

The joint pattern is consistent with financial participation and occasional purchases responding without a broad expansion in average spending. In the analytical example, equal mean resources leave mean consumption affected only by the consumption-financed portion of additional participation fees (equation 6). The life-cycle model in Section 6 lets consumption and saving adjust to the entire future earnings process. Its access experiments generate large ownership differences alongside modest consumption differences (Figure 9). Those experiments change access restrictions rather than earnings composition. They illustrate the distinction between participation and aggregate consumption; the model does not separately choose durable purchases, nondurable spending or IRA contributions.

The comparison uses 33,353 spending transitions in 2005–2019 and 33,056 contribution observations. CPS moments are assigned by earlier PSID industry, education and age, but overlap or postdate some historical observations. Contribution reports run from January of the earlier interview year through the next interview, overlapping the baseline year’s beginning. These estimates therefore describe associations across earnings environments rather than responses to new information. The original positive-support real-log construction does not reproduce contribution participation. The evidence concerns financial participation and expenditure. It does not establish consumption insurance, changes in total saving, taxable-to-retirement substitution, or equity exposure within retirement accounts.

5.2 Risky balances and the financial-asset denominator

A risky share can fall because risky holdings decline or other financial assets grow. The original real-log construction’s negative holder-share coefficient permits an accounting comparison of these explanations. For nonretirement financial assets F and risky balances R , define $\bar{F} = (F_t + F_{t-1})/2$ and $\bar{s} = (s_t + s_{t-1})/2$. Then


$$\Delta s = \frac{\Delta R}{\bar{F}} - \bar{s} \frac{\Delta F}{\bar{F}}. \quad (11)$$

The identity separates risky-balance and financial-asset components using midpoint weights. It accommodates entry and exit when both financial denominators are positive. On the same sample, weights, and regression design, the component coefficients sum to the share-change coefficient.

Lower risky balances account for the negative holder-share association (Figure 3). A one-standard-deviation increase in upside composition gives an estimated risky-balance contribution of -3.68 percentage points, partly offset by an imprecise $+2.06$ -point denominator contribution. Their sum is the -1.62 -point share coefficient. The risky-balance contribution is -4.57 points for higher upper partial variance and an imprecise $+2.79$ points for higher lower partial variance, each holding the other fixed.

The regressions retain the prospective-allocation controls. They show which balance-sheet component accounts for the share association, but cannot separate trades from returns or exits from

Figure 3. Accounting for the lower risky share of initial holders



Notes: The midpoint identity in equation (11) separates share changes into a risky-balance and a financial-asset-denominator contribution. Points report coefficients per fixed source standard deviation of age-adjusted real-log earnings composition (the original construction). The three regressions use the same 31,569 transitions and fully interacted design; plotted coefficients refer to 7,334 initial holders, retaining later exits. Point coefficients add exactly; the total uses its own estimated standard error. Bars are 95% CR1/ t_{29} intervals using 30 industry–education clusters, conditional on CPS moments. Balance changes include valuation and are not measured net trades.

changes among continuing holders. This result uses the positive-support real-log construction; the wage share coefficient is imprecise. Section 4.6 separates the participation margins.

5.3 Elicited earnings beliefs and retirement portfolios

Respondents’ own earnings beliefs provide a different comparison. The SCE relates individual perceived distributions to defined-contribution retirement portfolios, distinguishing standing equity shares from reported reallocation (Table 10). A one-standard-deviation increase in perceived dispersion accompanies a 4.75-percentage-point lower equity share conditional on normalized orientation (column (4)); using the sum of semideviations as the scale gives 3.80 points (column (2)). Orientation coefficients for the standing share are imprecise.

For reallocation, a one-standard-deviation increase in centered upside-minus-downside tilt accompanies a 0.028-unit higher riskier-minus-safer score at fixed total semideviation scale (column (1)). The score is one for riskier assets, minus one for safer assets, and zero for no change, not a change in equity share. Normalized orientation is imprecise (column (3)), and the tilt result does not survive the broader family adjustment.

Beliefs precede the Household Finance interview but generally fall within its twelve-month retrospective action window. They need not precede the reported action, so the comparison does not identify an effect of beliefs on adjustment. Repeated-belief tests and annual SIPP tail-change estimates also do not establish that tail revisions move shares. The clearest SCE pattern is the negative association between perceived dispersion and standing retirement equity shares.

Table 10. Perceived earnings risk, portfolio shares, and reported adjustment

	Semideviation tilt and scale		Normalized orientation and SD	
	(1) Reallocation	(2) DC share	(3) Reallocation	(4) DC share
Upside-minus-downside tilt	0.028** (0.013)	−0.576 (0.664)		
Sum of semideviations	−0.017 (0.018)	−3.805*** (1.011)		
Normalized upside orientation			0.019 (0.014)	0.059 (0.937)
Total standard deviation			−0.016 (0.019)	−4.754*** (0.970)
Shape coefficient: marginal p	0.036	0.390	0.173	0.950
Mean and demographics	Yes	Yes	Yes	Yes
State and module-year FE	Yes	Yes	Yes	Yes
Respondents	2,136	3,332	1,893	2,952
State clusters	50	50	50	50

Reallocation is the reported direction of a DC-account adjustment over the preceding twelve months: +1 for riskier assets, −1 for safer assets, and zero for no change. DC share is the standing equity allocation in percentage points. Beliefs are averaged from reports one to eleven months before the Household Finance module, using monthly sampling weights; these reports generally overlap the retrospective action window. All regressions use official module weights and control for expected earnings growth, age, education, and household income.

Each risk regressor uses its unweighted strict-pre respondent standard deviation. Columns (1)–(2) center lower and upper semideviations on the elicited mean and retain degenerate densities with zero tilt and scale. Columns (3)–(4) normalize the upper-minus-lower semivariance difference by total variance, requiring positive variance. Their samples and risk scales therefore differ. Standard errors are clustered by state. Stars retain unadjusted two-sided t_{49} tests of a zero coefficient:

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

5.4 Implications for household-finance models

At a given earnings variance, portfolio models should be able to explain why greater upside composition predicts entry and ownership at successive annual dates. Raising both partial variances produces a much smaller estimated entry response than changing their composition. This adds to evidence on non-Gaussian earnings dynamics and portfolio choice (Catherine, 2022; Shen, 2024; Gálvez and Paz-Pardo, 2026). Matching total variance and average risky shares alone need not reproduce the participation pattern.

The PSID evidence adds a complementary comparison: retirement-contribution participation and transformed furnishings expenditure are positively associated with earnings composition among initial holders, while average broad expenditure has no clear association. This pattern motivates attention to discrete financial decisions and infrequent purchases alongside ordinary spending. Section 2.4 shows why a discrete access decision can respond when average consumption changes little. In the life-cycle model, that distinction emerges when access restrictions change at fixed preferences and earnings laws. Its single consumption good does not deliver a separate prediction for the durable margin.

The participation-cost interpretation depends on the specified earnings distribution and resource threshold. Standard or risk-vulnerable preferences can make uninsurable risk reduce financial risk

taking (Kimball, 1993; Gollier and Pratt, 1996), without signing upside composition. Section 2 specifies a fixed-mean, fixed-variance change that creates both sustained and interrupted entries. A recurring cost (Vissing-Jorgensen, 2002; Fagereng, Gottlieb, and Guiso, 2017) permits exit after a transitory gain reverses; greater entry itself need not lower entrant survival. The resource responses in Briggs et al. (2021, 2023) connect to this mechanism, although standard models overpredict the windfall entry response in Briggs et al. (2021).

Several comparisons leave the participation-cost interpretation unresolved. Original real-log results link the upper tail more strongly to exit among recent entrants, but wage estimates are imprecise and the PSID does not reproduce the exit pattern. Earnings profiles, gain persistence and resource interactions do not establish the income reversal, while joint entry and gain does not identify mediation. Nor do the data establish a common intensive-margin response among continuing investors. Direct work-horizon and prior-participation comparisons are inconclusive; the analysis neither estimates household-specific human capital nor reproduces the twin design of Catherine, Sodini, and Zhang (2024). These results therefore need to be evaluated separately for entry, sustained entry, and incumbent retention.

Endogenous saving makes income dynamics and household resources central to this interpretation. The illustrative endogenous-saving benchmark, which holds mean and variance of earnings levels fixed, reduces precautionary saving under more upside-oriented earnings but does not reproduce the positive entry association and implies nearly universal young-household participation. Its refined solution gives one-year survival of 97.48% for upside-law entrants versus 96.08% under the downside law: carried wealth cushions later fluctuations. Section 6 therefore evaluates optimal saving and funded participation from observed SIPP financial positions, using earnings laws fitted to CPS and PSID moments, own employee contributions and retirement benefits, and calendar returns on inherited portfolios. I estimate preferences, participation costs and opportunities from entry, retention, portfolio shares and wealth changes, requiring positive wealth growth. The model then predicts entry-risk gradients on the same regression design, decomposes the contributions of earnings laws and decision rules, and evaluates sustained ownership on common model and data cohorts.

6 Earnings Risk and Portfolio Choice over the Life Cycle

Standard saving and costly-participation mechanisms provide a theoretical foundation for the empirical composition association. Section 2.1 shows why earnings distributions with the same mean and variance can put different numbers of households across an investment threshold. A life-cycle model makes that argument more demanding: households can consume an earnings gain, save it safely or invest it, and anticipated earnings affect these choices before income arrives. I estimate a model with recurring participation costs and intermittent opportunities to enter, then ask whether measured earnings environments generate the entry association without targeting it. Cell-specific dynamics yield a composition gradient of 0.24 percentage point, compared with approximately zero under pooled dynamics and 1.82 points in the matched data. The comparison evaluates what mea-

sured earnings heterogeneity adds to this familiar household problem, with a substantial remaining magnitude gap.

6.1 The workhorse mechanism and the earnings inputs

A Gaussian earnings innovation is determined by its conditional mean and variance. At fixed values of those moments, it leaves no independent composition change between favorable and adverse outcomes. Non-Gaussian processes, including those in [Catherine \(2022\)](#) and [Gálvez and Paz-Pardo \(2026\)](#), already allow distributional shape to affect portfolios. The theoretical ingredient is therefore an earnings distribution whose shape can vary beyond its first two moments; 58 distinct cells are not necessary for the mechanism. The cells supply measured heterogeneity with which to evaluate its relevance for entry.

Costless, continuously available investment also removes much of the distinction between choosing a risky share and becoming a stockholder. With a positive equity premium and independent earnings and stock returns, households with positive investable savings generally benefit from a small risky position. More cautious households can reduce that position without leaving the market. A recurring cost makes participation depend on whether its gain justifies the resources used to invest. Intermittent entry opportunities further distinguish an attractive investment from an investment the household can implement today. The experiments below remove these restrictions separately at fixed preferences and earnings laws.

Pooling earnings laws asks what measured cell differences add to common career dynamics and observed starting positions. The pooled benchmark retains non-Gaussian education-specific career dynamics, group wage levels and observed initial portfolios, but sets the cell adjustments to zero. I resolve the saving and portfolio rules under those pooled future laws at the same financial parameters. Workers retain their different wealth, wages and filtered earnings states, so predicted entry can still be associated with the original CPS composition measure. This fixed-parameter comparison is not a separately estimated pooled model, and it does not establish that workhorse models cannot generate the empirical signs. The permanent-access alternative below, for example, already has a positive composition gradient under pooled laws.

6.2 Households, participation costs and opportunities

Workers choose consumption, financial saving and a risky share through age 83, retire at 65, and consume terminal resources at 84. Preferences are Epstein–Zin, with discount factor β , risk aversion γ and elasticity of intertemporal substitution $\psi = 0.5$. Let M_t denote cash on hand, A_t end-of-period financial assets, α_t the risky share and D_t ownership. If the worker participates, a recurring cost $\kappa_R \xi_t$

is paid from current resources:

$$\begin{aligned}
c_t + A_t + D_t \kappa_R \xi_t &= M_t, \\
M_{t+1} &= A_t \{R_{f,t+1} + \alpha_t (R_{t+1}^s - R_{f,t+1})\} + Y_{t+1}, \\
c_t > 0, \quad A_t &\geq 0, \quad D_t \in \{0, 1\}, \quad 0 \leq \alpha_t \leq D_t, \quad \alpha_t A_t \geq \varepsilon D_t.
\end{aligned} \tag{12}$$

The minimum funded position is $\varepsilon = \$1$. Financial assets are gross of liabilities and exclude retirement accounts. The empirical participation margin is separately reported stocks and mutual funds outside retirement accounts; it is not total equity exposure through pensions and retirement accounts. Working resources incorporate taxes, own employee contributions and income outside wage work. Retirement resources follow PSID reference persons' own Social Security and pension income by education and age.

After income and portfolio returns are realized, an opportunity indicator O_t is observed before the investment decision. Its probability is

$$\Pr(O_t = 1 \mid D_{t-1}) = q(D_{t-1}) = D_{t-1} + (1 - D_{t-1})q, \quad D_t \leq O_t. \tag{13}$$

A nonholder receives a new opportunity with probability q each year. A current holder has access but may choose to exit; an exit exposes the worker to the opportunity restriction in the next period. A closed opportunity requires safe saving and consumes no money. When access is open, the worker observes the recurring-cost offer and decides whether to participate. The multiplier is iid lognormal, with $E\xi_t = 1$ and log standard deviation σ_κ .

The opportunity restriction represents occasions on which a nonholder can act on an investment decision. It can accommodate intermittent consideration or an opportunity to arrange participation, but the financial moments do not identify a particular attention, information or trust mechanism. There is no additional monetary entry charge. The recurring fee instead represents resources used while investing and is paid only upon funded holding. Both q and the cost distribution are common across earnings cells and independent of current earnings shocks. Differences in the estimated tail gradients therefore do not arise from assigning more opportunities to high-upside cells.

Let $V_t^0(m, z)$ be the optimized value of safe saving at cash m and earnings state z , and $V_t^1(m - k, z)$ the optimized value of funded holding after payment k . The worker accepts an open opportunity when

$$\Delta_t(m, z, k) = V_t^1(m - k, z) - V_t^0(m, z) > 0. \tag{14}$$

Both values allow future saving, consumption and portfolio decisions to adjust. Opportunities enter the entire lifetime problem, including the value of maintaining access. They are not applied as an entry-rate adjustment after solving an unrestricted model. The recursion in Appendix B integrates open and closed opportunities in the Epstein–Zin risk certainty equivalent.

6.3 Earnings dynamics, asymmetry and the cell inputs

An earnings law specifies the amounts and probabilities of future wages and wage-work interruptions. While employed, a worker in industry–education group g and education type h earns

$$\begin{aligned} \log W_{gt} &= f_h(t) + \ell_g + B_g(t) + \eta_t + \epsilon_{gt}, \\ \eta_{t+1} &= \rho\eta_t + \zeta_{t+1} + J^+ I_{t+1}^+ - J_b^- I_{t+1}^- - m_{gb}, \quad \zeta_{t+1} \sim N(0, s_0^2). \end{aligned} \tag{15}$$

Here $f_h(t)$ is the education-specific age profile, ℓ_g a group wage level and $B_g(t)$ a group growth adjustment. The persistence parameter ρ carries earnings states forward. Ordinary innovations and occasional promotions or losses change η_t , with the latter indicated by the mutually exclusive I^+ and I^- . Their probabilities depend on age and the current earnings state; b indexes ten-year age bands. The compensating drift m_{gb} centers each band’s stationary log-state distribution at zero. Temporary earnings ϵ_{gt} start from a two-component mixture with cell-adjusted node probabilities. A separate transition governs departure from and return to wage work. Unequal promotion and loss probabilities or sizes, the temporary-shock mixture, and the possibility of losing wage income all make the earnings distribution non-Gaussian. Survey reporting error is added to observed wages but does not enter household resources.

Two base processes, for college and noncollege workers, are calibrated to CPS one-year and PSID two-, four- and six-year earnings moments. The 58 observed industry–education–age cells then inform 30 industry–education career processes. Each has younger and older adjustments, corresponding to empirical ages 25–44 and 45–65; model workers switch adjustments at 45 and retire at 65. Two unsupported age bands borrow their group’s supported-band adjustment. High-school and some-college groups share the noncollege base but have distinct adjustments. Thus the cells add measured heterogeneity to familiar earnings dynamics; their number is not a numerical setting of the household solver.

Each cell has four adjustments: the wage-work departure hazard, the growth profile $B_g(t)$, the scale of temporary earnings and reporting variation, and an asymmetry tilt a_c . The tilt shifts probability toward favorable temporary shocks and promotions, and away from adverse shocks and losses, when $a_c > 0$. In the persistent component, promotion and loss probabilities are multiplied by e^{a_c} and e^{-a_c} before a common normalization, state dependence and probability caps. The same tilt acts on temporary-shock probabilities. The four adjustments are chosen jointly to fit each cell’s zero-wage share, mean growth, upper partial variance and lower partial variance after empirical-Bayes shrinkage. There is no one-to-one mapping from the empirical composition measure to a_c : departures, growth and scale also generate composition differences. The data discipline the cell earnings moments more tightly than the allocation of their asymmetry between persistent and temporary shocks.

The law enters choices through two channels. Realized wages, net of taxes and employee contributions, supply current cash; workers outside wage employment receive the model’s nonwage resources. These cash flows and the law’s transition probabilities enter the expectation over future resources in the household’s continuation values. Solving that problem produces a rule mapping

current cash, earnings state, age, ownership and access into consumption, saving and entry. Thus realized income changes what a household can invest today, while anticipated earnings change what it chooses to consume or save at the same current resources. Preferences, costs and entry opportunities remain common across groups. Future stock returns are independent of earnings shocks, distinguishing this mechanism from the joint stock–earnings downside risk emphasized by Catherine (2022).

6.4 Financial estimation with fixed earnings laws

Earnings moments, including upper and lower partial variances, are targets of the earnings calibration. The coefficients relating earnings risk to stock market participation are not targets. I fix the calibrated earnings laws before estimating the financial parameters and evaluating those portfolio coefficients.

Following the conditional approach of Briggs et al. (2021), I initialize the model from observed wages, wealth and portfolios. The financial sample contains 21,188 SIPP worker-year pairs aged 25–63. Wages filter latent earnings states. Calendar-year market and safe-return proxies apply to inherited portfolios in observed transitions, while decision rules use stationary future return expectations. Predictions condition within each observed origin on subsequent wage work and positive financial assets before aggregation with survey weights. Initial wealth, ownership and dated returns are inputs, not distributions fitted by the model.

I estimate $(\gamma, \beta, \kappa_R, \sigma_\kappa, q)$ from eight financial outcomes. Three entry rates by initial wealth discipline how participation changes with resources. Two retention rates and the next-year share of all initial holders, including exits at zero, discipline the holder margin. Mean financial-wealth changes for initial holders and nonholders discipline accumulation. These moments jointly inform preferences, recurring costs and opportunities; none is assumed to identify one parameter in isolation.

Table 11. Opportunity access: parameters

Parameter	Value
Risk aversion γ	8.801
Annual discount factor β	0.811
Mean recurring offer (2018 dollars)	524
Log-cost standard deviation	0.944
Opportunity probability q	0.176
Elasticity of intertemporal substitution (fixed)	0.50

Notes: Five parameters are estimated from the eight financial moments in Table 12. The elasticity of intertemporal substitution is fixed. The recurring offer has a mean-one lognormal multiplier; the reported mean is not the selected payment. There is no additional entry payment. The access probability is a model restriction, not a measured attention or trust variable.

The criterion is $Q_E + 0.1Q_H + Q_W$, with covariance-weighted entry, holder and wealth blocks. Both wealth changes must remain positive. The search places an exploratory upper bound of \$1,000 on the *mean* recurring offer, not on every lognormal draw. The estimated mean is about \$524 and the annual opportunity probability is 17.6%; no parameter bound binds. The bounded search included

paths that exhausted their evaluation budgets. Parameters were fixed before the earnings-risk, duration and counterfactual comparisons. These are outcomes outside the fitting criterion within exploratory model development, rather than a claim of untouched external validation.

Table 12. Opportunity access: targeted financial outcomes

Financial outcome	Data	Model
	2.93	
Entry: lower wealth tercile (%)	(0.26)	3.18
	5.84	
Entry: middle wealth tercile (%)	(0.36)	5.51
	10.36	
Entry: upper wealth tercile (%)	(0.48)	10.44
	66.19	
Retention: lower initial risky-share bin (%)	(1.43)	67.75
	76.31	
Retention: upper initial risky-share bin (%)	(1.23)	74.34
	33.38	
Next risky share: all initial holders (%)	(0.84)	34.40
	3,939	
Annual financial wealth growth: initial nonholders (dollars)	(778)	9,941
	1,946	
Annual financial wealth growth: initial holders (dollars)	(5,177)	511

Notes: All eight outcomes are targets. Rates and portfolio shares are percentages; financial-wealth changes are annual real 2018 dollars and include portfolio gains. Data standard errors, in parentheses, use 299 household-lineage multiplier draws. Initial holders who exit have zero next-year shares. The conditional sample contains 17,606 initial nonholders and 3,582 holders. Observed initial states are inputs.

Table 12 shows that the model reproduces the steep increase in entry across wealth groups and broadly matches retention and the next-year holder share. Both wealth changes are positive, including under the finer numerical grid. The main remaining financial discrepancy is accumulation: nonholder wealth rises by roughly \$10,000 against \$4,000 in the data, while holder growth is about \$500. These changes include capital gains. Positive holder wealth growth therefore does not imply positive saving net of returns.

The mean offer is also different from the payment made by a selected participant. Entrants pay about \$235 and retained holders about \$359 on average, compared with the potential mean of \$524. The entrant calculation covers 99.994% of initial-nonholder survey weight; one origin lacks simulation support for a selected-payment mean. The monetary cost has a direct budget interpretation, while the opportunity restriction governs whether the participation choice is available.

6.5 Entry gradients under cell and pooled earnings laws

I apply the empirical specification to predicted entry for the same 17,606 initial nonholders. In both calculations, predicted entry is related to the same original CPS risk measures. The cell and pooled calculations retain identical financial parameters, observed initial positions, cell-filtered earnings posteriors, calendar returns, projection controls and survey weights. The pooled calculation removes

future cell adjustments and resolves the optimal rules; it does not refit financial parameters. Table 13 reports the data and both model projections, together with their difference.

Table 13. Opportunity access: earnings-risk gradients

	Data	Data SE	Cell laws	Pooled laws	Difference
Upside composition	1.816***	0.480	0.241	-0.005	0.247
Upper-tail derivative	1.967***	0.513	0.115	-0.116	0.230
Lower-tail derivative	-1.665***	0.470	-0.368	-0.105	-0.263
Dispersion at mean volatility	0.151	0.104	-0.127	-0.110	-0.016

Notes: All columns use the same 17,606 initial nonholders, weights, controls and source standardization. Units are percentage points per source standard deviation. Data use industry–education CR1 standard errors and two-sided t_{29} tests: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Model columns are deterministic conditional projections without sampling intervals or stars. No earnings-risk coefficient enters estimation. Tail rows are local partial-variance derivatives; dispersion is evaluated at common mean volatility. The difference is cell minus pooled laws at common parameters, not a causal share explained.




Figure 4. Earnings laws and the opposing entry gradients

Notes: Panel A applies the same projection to observed entry and model probabilities for 17,606 initial nonholders. Cell and pooled calculations share financial parameters, initial states, dated returns, controls and survey weights; only future earnings laws differ. Upper- and lower-tail derivatives use the same composition-standard-deviation normalization, with dispersion evaluated at common weighted mean volatility, as in Table 13. Data intervals are CR1/ t_{29} 95% intervals; deterministic model projections have no sampling intervals. Panel B reports the signed earnings-law and decision-rule contributions to the cell-minus-pooled gradients, averaging the two change orders in Table 14. Its horizontal scale is expanded relative to Panel A. Each counterfactual conditions on its own future wage work and positive assets. The accounting is not causal mediation or an isolated change in perceived risk.

The composition coefficient is -0.005 percentage point under pooled laws and 0.241 under cell laws, compared with 1.816 points in the common data sample. The cell-minus-pooled difference, calculated before rounding, is 0.247 point and remains approximately 0.25 on the finer grid. Measured earnings heterogeneity recovers the positive association in this fitted specification, but most of the empirical magnitude remains unexplained. The 1.816 -point common-sample benchmark is distinct

from the 1.601-point full empirical-sample finding. Neither coefficient is used to choose financial parameters, and the ratio of model to data is not a causal share explained.

Figure 4, Panel A, provides the corresponding tail comparison. Pooled laws predict negative entry gradients for both tails. Cell-specific laws reverse the upper-tail sign and strengthen the negative lower-tail gradient. The cell-minus-pooled upper- and lower-tail gradients are approximately 0.23 and -0.26 percentage point, with little net dispersion gradient. This shows what the full measured earnings environment adds in the selected opportunity model. Because cell adjustments change growth, dispersion and employment risk as well as asymmetry, it does not isolate the contribution of an asymmetry parameter.

6.6 Earnings laws and household decision rules

Table 14. Opportunity access: earnings laws and decision rules

Projection	Earnings laws	Decision rules	Total
Upside composition	0.274	-0.027	0.247
Upper-tail derivative	0.302	-0.072	0.230
Lower-tail derivative	-0.245	-0.018	-0.263
Dispersion at mean volatility	0.029	-0.045	-0.016

Notes: The four calculations cross cell and pooled earnings laws with their complete value functions and saving/portfolio rules. Contributions average the two orders of changing laws and rules. Each calculation conditions within origin on its own future wage-work and positive-asset events. Units are percentage points per source standard deviation. The total equals the cell-minus-pooled projection. This accounting is not causal mediation or a pure perception effect.

An earnings law and a decision rule are different objects. The law gives the probabilities of future earnings and employment states; the rule maps the household's current cash, earnings state, age, ownership and access offer into consumption, saving and stockholding. I solve one rule using cell-specific future laws and another using pooled future laws. Table 14 then crosses the law generating next-year earnings with the previously solved rule applied after those earnings arrive. For example, cell-law earnings and employment outcomes can be fed into the pooled-law saving and portfolio plan. That plan continues to evaluate future choices using pooled earnings dynamics; it is held fixed rather than reoptimized under the cell law. The crossed cases are imposed accounting comparisons, while the cell/cell and pooled/pooled cases each use one earnings law throughout.

Averaging the two change orders assigns 0.274 percentage point of the composition gradient to earnings laws and -0.027 to decision rules, totaling 0.247 point. The first component changes the resources and earnings states households reach; the second changes their response to those states because future earnings prospects differ. Panel B of Figure 4 shows both contributions. Rule adjustment dampens the positive upper-tail increment and reinforces the negative lower-tail increment. Earnings outcomes and resources dominate the net composition difference, although both plans already anticipate future income. The accounting also changes which future states satisfy the work and positive-asset conditions; it is not causal mediation or an isolated change in beliefs.

This accounting compares two measured earnings environments, not realized income against an absence of expectations. Both decision rules anticipate future earnings, and the changed-rule component measures how choices differ at a given state when those prospects change. Its small contribution here does not imply that expectations are generally unimportant. The next experiment transfers variation between upper and lower dollar tails at fixed dollar means and variances; it shows how a saving response can reverse the entry effect over time.

6.7 Composition, saving and the duration of earnings opportunities

I shift an equal amount of dollar earnings variance toward favorable or adverse outcomes. Within each age, earnings group and persistent destination state, the upside law adds to the upper partial variance exactly what it removes from the lower partial variance; the downside law reverses that transfer. Outcome probabilities and each tail's conditional mean remain fixed. Expected disposable earnings, total dollar variance and expected discounted disposable resources are therefore unchanged. This is a small mechanism illustration, not a shock calibrated to an empirical standard deviation or to the entry coefficient. Arc-growth means and variances are not held fixed. Financial parameters remain fixed and earnings-risk coefficients remain untargeted.

The duration comparison separates a lasting change in desired saving from an entry response that depends on when the earnings change ends. I apply the same annual perturbation for one year (age 26), five years (ages 26–30), or every remaining working year (ages 26–64). Retirement resources are unchanged from age 65, and households know each schedule. Five years is an illustrative window, not an estimate of persistence. Each experiment starts from the same observed age-25 portfolios, and the first new decision follows age-26 earnings. Consequently, the one-year change leaves no altered future earnings at the first new decision; its effects arise through realized resources. This timing does not measure a saving response at an earlier announcement date. Appendix C gives the exact rotation.

Figure 5 shows that the later entry reversal is duration-dependent. Under the five-year upside change, a small initial increase gives way to a cumulative-entry decline of 0.0138 percentage point after ten years. With upside lasting throughout working life, the initial increase is 0.0175 point but the year-ten response is near zero: -0.0042 point, with a paired simulation standard error of 0.0035 point. This is not a robust negative response. Persistent downside remains below baseline over the displayed decade, whereas five-year downside eventually raises cumulative entry. The one-year entry responses are near zero and their signs are sensitive to the numerical grid.

Wealth responds more strongly when the earnings change lasts longer (Figure 6). At year ten, persistent upside reduces mean financial wealth by about \$192, compared with \$13 under five years; persistent downside raises it by \$262, compared with \$22. The matched-state comparison holds current cash and earnings states fixed. Under persistent upside, saving in the no-stock alternative falls by about \$69 in year one and \$41 in year ten. Under the five-year change, the same-state saving response is zero from year five, once the final affected earnings have arrived. Higher cumulative consumption accounts for \$163 of the persistent upside wealth decline by year ten, lower investment




Figure 5. Duration of earnings composition and cumulative first entry

Notes: Differences from baseline for 395 initial nonholders drawn from 452 observed age-25 origins in 29 earnings groups. Both directions use mirrored transfers between upper and lower dollar partial variances, preserving conditional expected disposable earnings and dollar variance. The three durations change earnings at ages 26, 26–30 or 26–64; retirement resources are unchanged. Curves use perturbed earnings and their optimized rules and display ages 26–35. Cumulative first entry is the fraction of initial nonholders who have ever entered, regardless of subsequent exit. It is neither current ownership nor an annual hazard. Differences are percentage points, using the 201-node calculation and common draws. Paths condition on survival without future employment or positive-asset selection.




Figure 6. Duration of earnings composition and financial wealth

Notes: Mean financial-wealth differences in real 2018 dollars for the same 395 initial nonholders, earnings changes and full optimized paths as Figure 5. All initial nonholders remain in the wealth average, whether or not they subsequently enter. The persistent change covers ages 26–64, although the figure displays only ages 26–35. These magnitudes illustrate the model's mechanism and are not calibrated to an empirical composition change.

gains for \$28, and participation fees for less than one dollar. This supports lower buffer saving, rather than participation payments, as the principal source of the wealth loss.

A larger wealth decline need not produce a larger decline in cumulative entry. Wealth measures dollars accumulated by the whole initial cohort; first entry records whether a household has crossed the investment threshold at least once. A household can consume more and carry fewer assets forward while still finding a small stock position worthwhile. The persistent upside law continues to raise the same-state probability of taking a stock position among baseline current nonholders, even while it lowers desired saving. In the five-year case, that direct encouragement disappears after

the final affected earnings arrive, but the lower assets accumulated earlier remain. Under persistent upside, ongoing encouragement and resource erosion work in opposite directions. The same-state comparison includes possible re-entry and does not by itself isolate a first-entry mechanism.

The distinction between entry flows and cumulative entry also matters. Under persistent upside, annual first-entry differences turn negative from year three onward. These later shortfalls largely erase the extra entries in the first two years. Earlier entrants remain counted in cumulative entry even after a later exit, so the cumulative difference returns toward zero without becoming clearly negative. The resource-rule accounting lets these entry and accumulation histories evolve rather than holding cash fixed.

These comparisons do not imply a universal entry sign. The no-stock saving comparison removes the current participation payment but retains future investment opportunities; it therefore does not uniquely identify pure prudence. The path decomposition changes saving and investment together and is not saving-only mediation. The lesson is that anticipated earnings composition changes both the desire to invest and the resources carried forward: the lower-saving and wealth response persists with a lasting change, while a clear reversal in cumulative entry depends on duration.

6.8 Earnings laws over working life

Figure 7 follows the same workers under cell-specific and pooled future earnings laws. Cell laws generate lower ownership early in working life and higher ownership later: the differences are about -4.1 percentage points at age 30 and $+6.2$ at age 60. Financial wealth follows a similar crossover. The earnings environment therefore changes the timing of accumulation and participation, rather than shifting ownership uniformly at every age. Together with the four-corner accounting, these paths illustrate how different earnings histories translate into different participation histories. The paths alone do not separate the contribution of accumulated resources from changes in decision rules.

6.9 Entry decisions at common resources

The earnings environment can also change the incentive to enter before cash balances evolve. Figure 8 holds current cash, age and earnings state fixed and compares the entry rules. The upper panels show participation probabilities, and the lower panels make their differences visible. The differences change sign across cash positions and earnings groups: cell laws do not imply a uniform increase in entry at a given resource level. This distinguishes a change in the willingness to invest from a change in the resources available to invest. The distinction is consistent with the accounting result that earnings and resource paths contribute more than decision rules to the composition gradient.

6.10 Why the participation restrictions matter

Figure 9 removes the opportunity restriction, the recurring fee, and then both. Each economy reoptimizes saving and investment at the same preferences and earnings laws. Opening access raises




Figure 7. Earnings laws, accumulation and participation over working life

Notes: The same 135 observed age-25 origins in the two cells with the largest initial survey-weight totals: college graduates in education and health, and in professional and business services. The cohort includes 112 initial nonholders and 23 holders. Financial parameters, observed initial portfolios and earnings-state posteriors are common inputs. The first new decision is at age 26. Paths condition on survival, without future employment or asset screening. Panels A and B report ownership and mean financial wealth; Panels C and D show cell minus pooled differences. Ownership is distinct from first entry. Pooling changes the full future earnings laws, including means and persistence. Ages 26–64 are shown; the saved paths extend through age 84. Wealth is in thousands of real 2018 dollars. These conditional cohort projections are not fitted population age profiles.

participation, but recurring costs still make nonholding attractive for some households. Removing the recurring fee while retaining intermittent opportunities instead makes holding attractive to almost everyone who has obtained access. With open, costless investment, ownership is virtually universal in the displayed working-age cohort.

At age 45, removing both restrictions raises ownership from about 42% to nearly 100%, while mean annual consumption rises by only about 1.3%. Opening opportunities alone raises ownership to about 66% with a consumption increase of about 0.4%. Households can save safely when access is closed, so nonparticipation need not imply that resources are consumed instead. Financial wealth also changes much less than participation. Fitting accumulation or conditional risky shares alone therefore leaves an important part of the investment decision unresolved.

This comparison connects the model to Section 5.1: a large difference in financial participation can coexist with a modest difference in current consumption when both choices are optimized. It provides a quantitative example of the distinction in equation 6, with endogenous saving and future earnings.




Figure 8. Entry under different earnings laws at the same current resources

Notes: Initially nonholding college workers aged 45 in education and health, and in professional and business services, at a common employed persistent earnings state (index 8). Each comparison fixes current cash, financial parameters, the opportunity probability and the recurring-cost distribution. The upper panels report entry probabilities under cell and pooled future earnings laws; the lower panels report their difference in percentage points. Cash is in thousands of real 2018 dollars. Probabilities integrate opportunities and cost offers without the future-work and positive-asset conditioning used in the observed-transition comparison. All future cell adjustments change jointly; this is not isolated upside news.

The access experiments hold earnings laws fixed, whereas the PSID regressions compare earnings environments, so their magnitudes are not a matched consumption test. Consumption outcomes were not targets in financial estimation, and retirement contributions enter working resources rather than being modeled as a separate participation choice. The experiment does not estimate a new earnings-tail regression for each restriction.

6.11 Accumulation, subsequent ownership and alternative access

Figure 10 follows observed age-25 origins forward with the estimated parameters. Ownership rises from about 15% at age 30 to 32% at 45 and 37% at 60 as savings and repeated opportunities support entry. Risky shares among holders need not move with ownership. The figure separates the extensive margin of access from the allocation decision conditional on holding.

The long-run exercise also shows the limits of the financial calibration. Wealth and participation fall sharply after retirement, with no bequest motive and low effective patience under the specified resource and terminal conventions. The projected life cycle is not a claim to match population




Figure 9. Entry opportunities and recurring participation costs

Notes: Fixed-parameter counterfactuals for the two age-25 cells with the largest initial survey-weight totals: college graduates in education and health, and in professional and business services. Observed initial portfolios and earnings-state distributions are common inputs. Paths condition on survival but impose no future employment or asset screen. Open access sets $q = 1$; no recurring fee sets $\kappa_R = 0$; both restrictions are removed in the final comparison. Other preferences and earnings laws remain fixed. These are conditional cohort projections, not reestimated specifications or population policy effects. Displayed paths stop at age 80; terminal liquidation occurs at 84.




Figure 10. Participation, accumulation and portfolio allocation over the life cycle

Notes: Projections start from 452 observed age-25 origins in 29 earnings groups, with survey-weighted initial wealth, ownership, shares and earnings-state posteriors as inputs. The first new decision is at age 26. Paths condition on survival without future employment or positive-asset selection. Curves use the accepted 141-node calculation; the vertical line marks retirement at 65. These are conditional cohort projections, not a fitted unconditional distribution or an age-period-cohort comparison. There is no bequest motive and remaining resources are consumed at 84.

retirement portfolios. In the observed-origin comparisons, sustained-entry levels are about 4.2% and 3.3% at two and three annual follow-ups, above 3.48% and 2.27% in the corresponding data cohorts. Their empirical composition gradients are imprecise.

Participation history provides an additional distinction. Among 408 nonholders with quality-screened reported prior ownership, the opportunity model predicts re-entry of about 10%, compared with 27% in the data. It also generates far fewer small entrant accounts than observed. These untargeted outcomes indicate where a common opportunity probability and the investment problem remain restrictive.

An alternative specification replaces recurring opportunities with permanent access types known to households and estimates the same five dimensions from the same financial targets. Its financial fit is close and its longer sustained-ownership level fits better. Under permanent access, pooled laws already produce a positive composition coefficient of 0.211 percentage point and opposing tail signs; cell-specific laws raise the composition coefficient to 0.382 point. However, exclusion imposes a fixed ceiling on eventual participation, and its treatment of reported prior holders overpredicts their re-entry. The opportunity specification allows changing resources to translate into entry over time without declaring initially inactive workers permanently outside the market. Both specifications leave a substantial gap between the observed and predicted earnings-risk magnitudes. Their comparison separates the economic role of access from the stronger claim that one particular access mechanism has been identified.

7 Conclusion

Upside earnings opportunities predict the start of stock market participation. At a given mean earnings growth, total variance and measured household resources, workers in more upside-oriented earnings environments are more likely to begin holding stocks or mutual funds. A one-standard-deviation increase predicts 1.6 percentage points more entry, nearly one fifth of the sample's annual entry rate. Upside composition predicts entry beyond the nonlinear relationship captured by variance and its square root. Its positive association also remains in the simpler linear partial-variance specification, at 1.11 percentage points per standard deviation.

The main comparison relates asymmetry in the earnings distribution itself to entry, rather than measuring how that asymmetry covaries with stock returns. Earnings distributions are measured before and separately from portfolio decisions. It relates historical earnings environments to participation after accounting for observed differences in resources and risk exposure. Transitions to zero wages help explain why more downside-oriented cells also have greater variance, but controlling for their frequency and measured cyclical exposure does not absorb the composition association. The finding also extends beyond the opening of a stock position: greater upside composition predicts entry followed by ownership at later annual observations among initial nonholders in the complete-followup cohorts. Among younger households, pre-pandemic PSID transitions show the same positive composition association and opposing tail signs. Ownership associations also appear in the SCF

and the earlier SIPP panel. Together, these findings connect earnings composition to stock market participation across survey settings and periods.

The model grounds this evidence in familiar saving and costly-participation mechanisms. In the two-period example, earnings composition matters because mean and variance do not determine the probability of covering a household's shortfall to the investment threshold. The 58 observed earnings cells supply empirical inputs to 30 career processes; they are not required for the theoretical mechanism. In the estimated life-cycle model, measured cell dynamics raise the untargeted composition gradient from approximately zero under pooled dynamics to 0.24 percentage point, against 1.82 in the matched data. This fixed-parameter comparison evaluates measured earnings heterogeneity, not the ability of every standard model to fit the association. Differences in realized resources dominate the additional gradient, while rules optimized for the different future earnings prospects modestly offset it. The mechanism operates with future stock returns independent of earnings shocks.

A controlled composition experiment makes the saving margin explicit. The experiment shifts an equal amount of variation between earnings outcomes above and below their mean, holding expected disposable earnings and total variance fixed. Here, variance is calculated from earnings amounts rather than from earnings growth rates and is measured in squared dollars. In this illustration, more upside lowers desired saving at given resources while initially making stockholding slightly more attractive. When the change lasts five years, the incentive ends while lower accumulated resources remain, and cumulative entry eventually falls. When the change lasts throughout working life, wealth falls more but the ongoing incentive offsets resource erosion sufficiently that the cumulative-entry response fades toward zero. The small response is a mechanism illustration, not a quantitative explanation of the empirical coefficient. The model therefore distinguishes realized earnings that finance entry from future earnings risk that changes saving and the resources available for later entry. This interaction can separate the immediate participation response from subsequent market access.

The empirical and model results place earnings opportunities alongside downside risk in explaining who becomes an investor. Workers with similar measured mean earnings growth, variance and current resources can have different entry patterns, depending on how earnings variation is divided between favorable and adverse outcomes. These entry relationships provide an additional test of the model beyond its fit to the targeted financial outcomes. How earnings vary, as well as how much they vary, helps explain who becomes an investor.

Data Availability. The CPS ASEC extract uses licensed IPUMS CPS microdata ([Flood et al., 2025](#)). SIPP ([U.S. Census Bureau, 2025](#)), PSID ([Panel Study of Income Dynamics, 2023](#)), SCE ([Federal Reserve Bank of New York, 2026](#)), and SCF ([Board of Governors of the Federal Reserve System, 2019](#)) are public. The replication package supplies download links, raw-data build scripts, frozen analysis inputs, and source-to-paper validation; redistribution of the IPUMS extract is subject to its license.

Disclosures. The author reports no research funding or conflicts of interest.

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.
- Ahn, Changhyun, Chanik Jo, and Sing-Sen Lam, 2026, Left-tail income risk and stock market participation: Evidence from U.S. wrongful discharge laws, Working paper, SSRN 4163869 (revised March 2026).
- 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.
- Barndorff-Nielsen, Ole E., Silja Kinnebrock, and Neil Shephard, 2010, Measuring downside risk: Realised semivariance, in Tim Bollerslev, Jeffrey R. Russell, and Mark W. Watson, eds.: *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle* (Oxford University Press, Oxford), 117–136.
- 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>.
- Bonaparte, Yosef, George M. Korniotis, and Alok Kumar, 2020, Income risk and stock market entry/exit decisions, CEPR Discussion Paper 15370.
- Bonaparte, Yosef, George M. Korniotis, Alok Kumar, Alexander Michaelides, and Yuxin Zhang, 2025, Stock market ownership transitions, *Management Science* 71, 4977–5000, doi:10.1287/mnsc.2023.00290.
- Briggs, Joseph, David Cesarini, Erik Lindqvist, and Robert Östling, 2021, Windfall gains and stock market participation, *Journal of Financial Economics* 139, 57–83, doi:10.1016/j.jfineco.2020.07.014.
- Briggs, Joseph S., David Cesarini, Sean Chanwook Lee, Erik Lindqvist, and Robert Östling, 2023, Financial windfalls, portfolio allocations, and risk preferences, NBER Working Paper No. 31864, doi:10.3386/w31864.
- Brown, William, 1910, Some experimental results in the correlation of mental abilities, *British Journal of Psychology* 3, 296–322.
- Brunnermeier, Markus K., and Stefan Nagel, 2008, Do wealth fluctuations generate time-varying risk aversion? Micro-evidence on individuals’ asset allocation, *American Economic Review* 98, 713–736.
- 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.

- 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, Ellen Jiayang Lu, and James D. Paron, 2026, A dynamic model of the racial wealth gap, *Review of Financial Studies*, hhag047, doi:10.1093/rfs/hhag047.
- 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, 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 Wibrall, and Christian Zimpelmann, 2023, Skewness expectations and portfolio choice, *Experimental Economics* 26, 107–144, doi:10.1007/s10683-022-09780-9.
- Eberly, Janice C., 1994, Adjustment of consumers' durables stocks: Evidence from automobile purchases, *Journal of Political Economy* 102, 403–436, doi:10.1086/261940.
- 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.
- Galaasen, Sigurd Mølster, and Akash Raja, 2024, The dynamics of stock market participation, Norges Bank Working Paper 5/2024.
- Gálvez, Julio, 2023, Household portfolio choices under (non-)linear income risk: An empirical framework, Banco de España Documentos de Trabajo 2327.
- 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.

- Ghysels, Eric, Alberto Plazzi, and Rossen Valkanov, 2016, Why invest in emerging markets? The role of conditional return asymmetry, *Journal of Finance* 71, 2145–2192, doi:10.1111/jofi.12420.
- 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.
- Gollier, Christian, and John W. Pratt, 1996, Risk vulnerability and the tempering effect of background risk, *Econometrica* 64, 1109–1123.
- 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.
- Inkmann, Joachim, Alexander Michaelides, and Yuxin Zhang, 2026, Intra-household risk sharing in collective portfolio choice models, *Journal of Financial and Quantitative Analysis*, FirstView, published online April 27, 2026, doi:10.1017/S0022109026102920.
- Kimball, Miles S., 1993, Standard risk aversion, *Econometrica* 61, 589–611.
- Kish, Leslie, 1965, *Survey Sampling* (Wiley, New York, NY).
- Li, Fan, Kari Lock Morgan, and Alan M. Zaslavsky, 2018, Balancing covariates via propensity score weighting, *Journal of the American Statistical Association* 113, 390–400, doi:10.1080/01621459.2016.1260466.
- MacKinnon, James G., Morten Ørregaard Nielsen, and Matthew D. Webb, 2023, Fast and reliable jackknife and bootstrap methods for cluster-robust inference, *Journal of Applied Econometrics* 38, 671–694, doi:10.1002/jae.2969.
- Owen, Art B., and Dean Eckles, 2012, Bootstrapping data arrays of arbitrary order, *Annals of Applied Statistics* 6, 895–927, doi:10.1214/12-AOAS547.
- [dataset] Panel Study of Income Dynamics, 2023, public use dataset: family files 1999–2019, cross-year individual file and wealth supplements, produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI, <https://psidonline.isr.umich.edu>.
- Ranish, Benjamin, 2013, Why do households with risky labor income take greater financial risks?, Available at SSRN 2681360, doi:10.2139/ssrn.2681360.

- Roodman, David, James G. MacKinnon, Morten Ørregaard Nielsen, and Matthew D. Webb, 2019, Fast and wild: Bootstrap inference in Stata using boottest, *Stata Journal* 19, 4–60, doi:10.1177/1536867X19830877.
- Schmidt, Lawrence D. W., 2025, Climbing and falling off the ladder: Asset pricing implications of labor market event risk, *Journal of Financial Economics* 172, 104131.
- Segal, Gill, Ivan Shaliastovich, and Amir Yaron, 2015, Good and bad uncertainty: Macroeconomic and financial market implications, *Journal of Financial Economics* 117, 369–397, doi:10.1016/j.jfineco.2015.05.004.
- 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.
- U.S. Census Bureau, 2020, 2018 SIPP: Model-based imputation of asset variables, SIPP User Note, August 14, <https://www.census.gov/programs-surveys/sipp/tech-documentation/user-notes/2018-usernotes/2018-mod-bas-impute-ass-var.html>.
- [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>.
- 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.
- Zhang, Yapei, 2018, Seeking skewness, Working paper, HEC Paris.

A Participation thresholds

A.1 Endogenous consumption and the participation decision

In the two-period problem of Section 2.1, current resources are $m > 0$, the discount factor is $\beta > 0$, and participation costs $\kappa > 0$. The safe gross return is one. Independent risky excess returns are $\bar{r}$ with probability p and $-\bar{r}$ otherwise, where $p \in (1/2, 1)$ and $\bar{r} \in (0, 1)$. With no borrowing or short sales, define the optimized expected log-return gain

$$b = \max_{0 \leq \alpha \leq 1} \mathbb{E}[\log(1 + \alpha r)] > 0.$$

Logarithmic utility makes this gain independent of saved principal. Conditional on participation $d \in \{0, 1\}$ and $m > \kappa d$, optimal consumption and saving are

$$c_d = \frac{m - \kappa d}{1 + \beta}, \quad s_d = \frac{\beta(m - \kappa d)}{1 + \beta}.$$

Substitution into the objective gives

$$V_d(m) = (1 + \beta) \log(m - \kappa d) + \beta b d + K_\beta, \quad K_\beta = \beta \log \beta - (1 + \beta) \log(1 + \beta). \quad (\text{A.1})$$

The value difference increases with resources on the feasible domain. Setting it to zero gives the threshold m^* in equation (2):

$$m^* = \frac{\kappa}{1 - \exp[-\beta b / (1 + \beta)]} > \kappa. \quad (\text{A.2})$$

With indifference resolved toward participation, $d^* = \mathbf{1}\{m \geq m^*\}$. Therefore

$$\mathbb{E}[c] = \frac{\mathbb{E}[m] - \kappa \Pr(d^* = 1)}{1 + \beta}.$$

For two earnings laws with the same mean resources, this proves

$$\Delta \mathbb{E}[c] = -\frac{\kappa}{1 + \beta} \Delta \Pr(d^* = 1). \quad (\text{A.3})$$

The fixed-mean, fixed-variance example in Section 2 changes the probability of crossing the threshold. Its direction depends on the threshold's location. The identity concerns current consumption after earnings arrive; terminal consumption also depends on returns. It does not imply greater total saving when participation rises.

A.2 Proof of the stationary participation comparison

The repeated fixed-saving problem instead has safe terminal nontradable income $h > 0$. For an interior stock position, its participating value is $\log(w_1 + h - \kappa) + b$, with $w_1 = w + Y$ and

$$b = p \log(1 + \eta \bar{r}) + (1 - p) \log(1 - \eta \bar{r}) > 0, \quad \eta = \frac{2p - 1}{\bar{r}}.$$

The nonparticipating value is $\log(w_1 + h)$. Their comparison yields $w^* = \kappa/(1 - e^{-b}) - h$, as in equation (3). A sufficient interiority condition at and above the threshold is $0 < \eta < 1$ and $\kappa > h(e^b - 1)/(1 - \eta)$.

Proof of Proposition 1. The symmetric earnings law Y_0 has upper and lower partial variances $d^2/2$ each; the upside-oriented law Y_1 has $4d^2/5$ above zero and $d^2/5$ below. Both have mean zero and variance d^2 . For each of the N_m marginal households, $w^* - w_m \in (d, 2d]$ means that neither realization of Y_0 reaches the threshold and only $Y_1 = 2d$ does. For each of the N_e established households, $w_e - d \geq w^*$ means every realization reaches it. Participation probabilities are thus $\pi_m = 0$ under Y_0 , $\pi_m = 1/5$ under Y_1 , and $\pi_e = 1$ under both laws. Baseline nonholders are marginal households, so entry rises from zero to $1/5$.

Under Y_0 , the baseline holders are the N_e established households and all remain. Under Y_1 , they also include $N_m/5$ marginal households that drew $2d$. Established households remain with probability one and marginal holders with probability $1/5$, yielding retention $(N_e + N_m/25)/(N_e + N_m/5) < 1$. Every exit is by a marginal holder. A marginal holder at the baseline was a nonholder one period earlier with probability $4/5$. Independence makes this fraction the same among those who subsequently exit. This proves both parts of the proposition. $\square$

The comparison is between stationary populations. The symmetric law produces no entrants, so the proposition does not compare entrant survival across the two laws. The joint sustained-entry probability in equation (5) follows from the same independent draws conditional on household type.

B Conditional model and estimation

B.1 Timing and recursive preferences

In the life-cycle model, earnings and returns on inherited portfolios arrive before consumption, saving and participation choices. Let z denote the persistent earnings and wage-work state, d prior ownership, and m cash before any participation payment. A previous nonholder obtains an opportunity with probability q ; a previous holder has access with probability one, so $q(d) = q + (1 - q)d$. When access is open, the worker observes the recurring-cost offer

$$K_t = \kappa_R \xi_t, \quad \xi_t = \exp(\sigma_\kappa \nu_t - \sigma_\kappa^2/2),$$

where ν_t is iid standard normal. The mean offered cost is κ_R ; it is paid only for a funded stock position. A closed opportunity permits safe saving and consumes no resources. Primitive opportunity and cost draws are independent across dates and of earnings and returns.

Let $\rho = 1 - 1/\psi$ and $\omega = 1 - \gamma$, where ψ is the elasticity of intertemporal substitution and γ risk aversion. Write $V_{jt}(z, x)$ for optimized value conditional on current holding choice $j \in \{0, 1\}$ and cash x after any payment. With discount factor β and survival probability s_t , the consumption-equivalent recursion is

$$\begin{aligned} V_{jt}(z, x) &= \max_{(A, \alpha) \in \mathcal{A}_j(x)} \left[(x - A)^\rho + \beta s_t \left\{ \mathbb{E}[\bar{V}_{t+1}(z', j, AR(\alpha) + Y')^\omega \mid z] \right\}^{\rho/\omega} \right]^{1/\rho}, \\ \bar{V}_t(z, d, m) &= \left\{ [1 - q(d)] V_{0t}(z, m)^\omega \right. \\ &\quad \left. + q(d) \mathbb{E}_\xi \left[\max\{V_{0t}(z, m), V_{1t}(z, m - \kappa_R \xi)\}^\omega \right] \right\}^{1/\omega}. \end{aligned} \tag{B.1}$$

Here A is saving, α its risky share, $R(\alpha)$ its gross portfolio return, and Y' next-period disposable resources. The feasible set $\mathcal{A}_j(x)$ imposes equation (12), including positive consumption $x - A$, no borrowing or short sales, and the minimum funded equity position $\varepsilon = \$1$. An infeasible participating branch is excluded. The first expectation integrates next earnings and returns. The second integrates opportunities and observed cost offers in the same risk certainty equivalent; it is not an entry adjustment applied after solving an unrestricted problem. Exiting exposes a current holder to the opportunity restriction next period.

Working life ends at 65. At age 84, $\bar{V}_{84}(z, d, m) = m$: all terminal resources are consumed, without a portfolio choice or bequest motive. The recursion omits the usual flow multiplier $1 - \beta$. Multiplying all values, including the terminal value, by $(1 - \beta)^{1/\rho}$ restores that normalization with the same β .

B.2 Conditional predictions and financial targets

Observed age, wages, assets, ownership and risky shares initialize each worker. The earnings posterior describes the researcher's uncertainty; workers observe their earnings and work states when choosing. Dated return proxies apply to inherited portfolios in the observed transitions, while decision rules use stationary future return expectations. Reporting error affects wage measurement and initialization, but does not enter the resource budget.

The eight financial targets are three entry rates by initial-nonholder wealth tercile, two retention rates by initial-holder share bin, the next-year mean share of all initial holders including exits at zero, and mean annual financial-wealth growth for initial holders and nonholders. Empirical bin memberships and survey weights remain fixed during estimation. Financial wealth is gross of liabilities and excludes retirement accounts; wealth changes include portfolio returns.

Predictions condition within each observed origin on positive next-year wage work and financial assets. Let $\mathcal{E}_{i1}$ denote this eligibility event, $\mathcal{I}_i$ the observed origin and earnings posterior, h_j the

financial outcome, and $\mathcal{B}_j$ its empirical bin. With destination survey weight w_i , the model moment is

$$m_j(\theta) = \frac{\sum_{i \in \mathcal{B}_j} w_i \frac{\mathbb{E}_\theta[h_j \mathbf{1}\{\mathcal{E}_{i1}\} \mid \mathcal{I}_i]}{\Pr_\theta(\mathcal{E}_{i1} \mid \mathcal{I}_i)}}{\sum_{i \in \mathcal{B}_j} w_i}. \quad (\text{B.2})$$

The dated return input is fixed in both expectations. Open and closed opportunity branches enter the outcome and eligibility masses before taking their ratio. Conditioning each branch separately and then weighting its conditional mean would generally give a different moment. The origin distribution is not selected again; a zero eligibility probability makes its prediction undefined rather than causing deletion. For wealth growth, $h_j = A_{i,t+1} - A_{it}$. Holder predictions condition on initial holding, not continued ownership.

Let $r_E(\theta)$ collect the three entry-rate residuals, $r_H(\theta)$ the two retention residuals and holder-share residual, and $r_W(\theta)$ the two wealth-growth residuals. The parameter vector is $\theta = (\gamma, \beta, \kappa_R, \sigma_\kappa, q)$. With $\widehat{\Sigma}_j$ the marginal covariance block and I_k the k -dimensional identity matrix, the criterion is

$$\begin{aligned} Q(\theta) = & r_E(\theta)' [\widehat{\Sigma}_E + s_E^2 I_3]^{-1} r_E(\theta) \\ & + 0.1 r_H(\theta)' [\widehat{\Sigma}_H + s_E^2 I_3]^{-1} r_H(\theta) \\ & + r_W(\theta)' [\widehat{\Sigma}_W + s_W^2 I_2]^{-1} r_W(\theta), \quad s_E = 0.0025, \quad s_W = 2. \end{aligned} \quad (\text{B.3})$$

Rates and shares enter as proportions; wealth changes use thousands of 2018 dollars. Covariance blocks use 299 common household-lineage multiplier draws. Within-block covariance is retained and cross-block covariance omitted. The weights reflect economic priorities chosen during model development; the criterion is neither efficient SMM nor a specification-test statistic. The search imposes a \$500 positive-growth buffer for both initial ownership groups. Acceptance requires positive growth at both numerical resolutions, with wealth differences below \$250; the buffer is not a target of exactly \$500 or a restriction on every subgroup.

Backward induction uses 93, 141 and 201 cash nodes for coarse, fine and accuracy calculations, with the funded-position boundary included. The selected parameters are a feasible, converged terminal solution. Conditional cost integration and finite earnings-law integration determine financial moments without a simulation seed. Fixed-parameter numerical checks cover budgets, funded ownership, consistent shares, eligibility and the eight financial moments. No earnings-risk, age, duration, windfall or account-size outcome enters financial estimation or numerical acceptance.

C Controlled earnings-composition experiment

The experiments in Section 6.7 change transitory employed earnings for one year (age 26), five years (ages 26–30), or every remaining working year (ages 26–64). The annual perturbation at each shared age and state is identical across durations. Retirement resources are unchanged from age 65.

For each earnings group, age and persistent destination state, let μ be baseline mean dollar wages and U, L the upper and lower partial variances around μ . With direction s and nonnegative transfer d , the two laws satisfy

$$(U^s, L^s) = (U + sd, L - sd), \quad s \in \{-1, +1\}, \quad d \geq 0. \quad (\text{C.1})$$

Each tail keeps its probability mass and conditional mean. Within-tail deviations are expanded or contracted around that tail's conditional mean, using a square-root scale to deliver equation (C.1). The transfer is 10% of the smaller adjustable within-tail variance, reduced by the same amount in both directions when positivity or preservation of tail membership requires it. Unavailable rotations leave the state unchanged. The symmetry concerns dollar partial-variance changes, not equal and opposite movements of every wage outcome.

The wage-to-disposable-income map and persistent and employment transitions are unchanged. Fixed destination means and variances therefore preserve each current state's conditional disposable-income mean and variance and its expected discounted disposable resources. Arc-growth moments and risk-adjusted human-capital values are not constrained to remain fixed. The change is a mechanism illustration, not a perturbation calibrated to the empirical composition coefficient or one empirical standard deviation.

All scenarios share the 452 observed age-25 origins in 29 earnings groups, including 395 initial nonholders and 57 holders, and their financial states, earnings posteriors and weights. Paths condition on survival without selecting on future work or positive assets. The first new choice follows age-26 earnings. For the one-year change, no future earnings change remains at any reported choice, so matched-state rules equal baseline. For the five-year change, rules coincide from the age-30 choice after the last affected earnings. Under the persistent change, policies anticipate altered earnings through age 64. Paths can continue to differ after the laws coincide because they inherit different assets and ownership histories. These timings distinguish realized-resource effects from responses to altered future earnings; they do not measure a response at an earlier announcement date.