Asymmetric Labor Income Risk: Implications for Risk-Taking in Financial Markets*

Tai Lo Yeung†

March 13, 2026

Click here for the latest version

Abstract

Using the Survey of Income and Program Participation (SIPP), this paper constructs a tail-dependent measure of labor income risk that interacts individuals' earnings growth variance and skewness. Holding income volatility constant, investors exposed to right-tail shocks (opportunity risk) increase their equity share, while those facing left-tail shocks (disaster risk) sharply reduce it. Moreover, these tail effects are asymmetric: left-tail risk hurts virtually all investors, whereas right-tail risk benefits only a subset. These results show that higher-order moments of labor income risk shape portfolio choice, highlight the limitations of variance-based uncertainty metrics, and offer micro-level support for more robust portfolio choice models. For instance, many previous studies have documented a counterintuitive positive correlation between earnings volatility and stock holdings, termed the 'volatility puzzle'. This paper shows that this positive relation is accounted for almost entirely by the right tail of the earnings risk distribution.

Keywords: Stock market participation, household finance, labor income risk, risk aversion, investment decisions

JEL Codes: G11, G50, D14

*I would like to thank Lorenz Kueng for his guidance. For helpful comments and suggestions, I thank Alberto Plazzi, Luise Eisfeld, Fabrizio Colella, Darrel Duffie, Moritz Kuhn, Petra Todd, Harald Uhlig, Sylvain Catherine, and Jessica Wachter, as well as seminar and conference participants and two anonymous referees. All errors are my own.

†Università della Svizzera italiana, Via Giuseppe Buffi 13, tai.lo.yeung@usi.ch.

1 Introduction

Labor income risk is the main risk for most households because the present value of human capital is typically the largest component of lifetime wealth. Shocks to it affect many decisions, including portfolio choice. A vast body of research therefore examines the link between labor income risk and portfolio choice.[1] Yet, due to data limitations, only a small fraction can directly assess the connection between human capital risk and risky asset investments of individuals, let alone incorporate higher moments of the labor income process into the analysis.

This paper emphasizes the pivotal but separate roles of the two tails of individuals' earnings risk distribution in shaping household portfolio decisions, building on the fact that higher-order moments can affect individuals' expected utility substantially. Guided by standard expected utility theory, it develops separate measures which capture both the left tail (disaster shocks) and the right tail (opportunity shocks) of individuals' earnings risk. This approach relaxes standard parametric (Gaussian) assumptions and explicitly incorporates higher-order moments, in particular third moments (left and right skewness), in addition to the variance. Under normality, the risk distribution is symmetric, implying each worker is equally likely to experience positive or negative deviations from the life-cycle profile. While this symmetry is analytically convenient—summarizing risk with a single number, the variance—it is rarely realistic. In practice, upside and downside shocks occur with very different probabilities and magnitudes, so that relying on a normal approximation masks the skewed earnings risks households actually face (Acemoglu, Ozdaglar and Tahbaz-Salehi 2017).[2]

A growing literature therefore examines how higher-order moments of earnings risk shape financial decisions and asset prices. For example, Catherine, Sodini and Zhang (2024) studies how the covariance between the skewness of individuals' income risk and stock returns drives

[1] For related research on the interplay between labor income risk and household portfolio choice, see, inter alia, Guiso, Jappelli and Terlizzese (1996), Cocco, Gomes and Maenhout (2005), Benzoni, Collin-Dufresne and Goldstein (2007), Angerer and Lam (2009), Betermier et al. (2012), Bonaparte, Korniotis and Kumar (2014), Chang, Hong and Karabarbounis (2018), Gomes and Smirnova (2021), and Catherine (2022).

[2] For example, income volatility is elevated in both booms and recessions, but boom-time volatility is tilted toward upside shocks, whereas recession volatility is concentrated in the downside (left tail).

the cyclical component of portfolio choice.[3] Relatedly, Schmidt (2022) shows how the tails of idiosyncratic labor income risk can affect asset pricing. Building on this insight, this paper measures higher-order labor income risk with tail-conditioned volatility that separately isolates exposure to downside and upside shocks, rather than relying on a single measure of third moments, such as the Kelly skewness, or the co-movement of third moments with stock market returns. Simulations show that, holding variance fixed, right-tail and left-tail risk move optimal equity shares in opposite directions, with effect sizes that vary systematically with an individual's risk aversion. Because conventional skewness measures collapse both tails into a single number, they conflate the two tails and can mask the asymmetric forces that drive portfolio choice.

The core task is therefore to characterize the *asymmetric* (non-Gaussian) nature of each worker's income risk. To measure this asymmetry and trace its impact on individuals' portfolio choice, the paper proceeds as follows. First, to maximize comparability across individuals, I estimate each person's labor income variance from twelve consecutive year-over-year monthly growth rates of earnings using the Survey of Income and Program Participation (SIPP) from 2013–2022. The SIPP is a publicly-available and nationally-representative rotating panel at monthly frequency, following each of about 20,000 individuals for two to four years.[4] Survey participants report both earnings and detailed stock market investments, along with key covariates that standard models require, such as age, education, and wealth. The high frequency, long panel dimension, and large cross-section allow me to estimate individual labor income risk more precisely than is possible using other surveys such as the Panel Study of Income Dynamics (PSID), which has a much smaller cross-section and samples at lower frequency, or the repeated cross-sections from the triennial Survey of Consumer Finances (SCF).[5] I then assign this annualized variance to the individual and interpret it as its job-specific earnings risk. Thus, in any calendar year, each worker has a single, consistently measured variance value. Second, I assume that workers in the same *industry-education-*

[3]They employ a *cyclical* skewness metric, which measures how third moments of individual income risk co-move with the stock market and the business cycle. The skewness measure used in this paper targets an entirely different object, the left- and right-tail of individuals' earnings risk distribution, hence the two approaches are complementary.

[4]I do not use pre-2013 waves because the survey instrument was substantially updated in 2013, making comparison with earlier waves difficult.

[5]For the importance of within-year variation in labor income, see Ganong et al. (2024).

month cell face a common *cross-sectional* distribution of income shocks as in Catherine, Sodini and Zhang (2024). I then compute separate measures of left- and right-tail risk from co-workers within each cell and assign them to all individuals in that cell. This targets tail risk where it is most precisely identified—across workers with similar characteristics at a point in time. By pooling all workers in a month, this cross-sectional approach avoids diluting tail behavior. Third, I restrict attention to households' directly held equities because they can be converted to or from cash immediately. By imposing these assumptions, I condition the individual earnings variances on their industry-by-education group skewness. The interaction yields worker-specific measures of labor income risk that vary separately with the left and right tail of the time-varying earnings distribution by industry and skill.

By introducing a tractable, tail-dependent measure of labor income risk and linking it to both intensive and extensive portfolio choice margins, this paper makes five main contributions: (i) it reconciles the variance puzzle, the positive correlation between earnings volatility and stock holdings; (ii) it explains why young stockholders are less likely to participate in the stock market; (iii) it shows that relying on a single skewness statistic masks tail-specific asymmetries, conflating the distinct effects of the left and right tails; (iv) it provides insights on how tail risk fuels wealth-return inequality; and (v) it provides researchers with a framework based on publicly available data for incorporating skewness into empirical and structural analyses.

The paper's findings call for a shift away from variance-only metrics toward richer descriptions of labor income risk when studying household portfolios. For instance, given the same level of labor income volatility, households are more inclined to invest in risky assets when the right tail of the labor income risk distribution is more pronounced. This tendency is especially evident among households with lower risk aversion and those facing higher income volatility. Conversely, when the left tail of the labor income risk distribution becomes longer, households — particularly those with higher risk aversion — tend to decrease their investments in risky assets. These two opposite effects are masked when focusing only on second moments. In specifications that do not separately control for the two tails, a two-standard-deviation increase in an individual's labor income volatility is associated with a 0.34 percentage point *rise* in the risky asset share. However, this relation changes com-

pletely after accounting for the level of skewness and its opposite effects in the two tails. Once variance is estimated jointly with upper-tail dispersion (L9050×variance),[6] its coefficient on the risky assets share turns negative (−0.084) and remains statistically significant. This suggests that the previously observed positive relation between income variance and portfolio share invested in risky assets is largely associated with contributions from the right-tail segment of labor income risk. Moreover, the estimates indicate that the sensitivity of portfolio choice to downside labor income risk (or disaster risk) — captured by the corresponding term L5010×variance — rises sharply as human capital accounts for a larger share of a household's total wealth. Investment demand for risky assets of households that rely heavily on human capital to finance consumption and possess limited financial buffers is acutely sensitive to downside risks in labor income. When labor income risk turns left-skewed, they respond by scaling back their holdings of risky assets. However, they do not correspondingly increase their investments even in the presence of strong positive skewness. By contrast, households with greater risk tolerance — typically wealthier households and those in mid-career — are less deterred by downside risk. Instead, they tend to expand their equity positions when the prospect of upside shocks (opportunity risk) improves.

Related literature and contribution. This paper advances three related strands of literature that link labor income risk to household portfolio choice, and in doing so offers new evidence that non-Gaussian features of earnings shocks are central for understanding equity demand.

The first strand concerns the measurements of labor income risk. Classic life-cycle models starting with Merton (1969) and Samuelson (1969) assume normally distributed earnings shocks and therefore summarize uncertainty with a single statistic: the variance of income growth. Empirical tests built on this framework use earnings volatility as a single measure of earnings risk (Betermier et al. 2012; Bonaparte, Korniotis and Kumar 2014). A parallel literature in macroeconomics, however, documents pronounced time-variation in the *asymmetry* of the earnings distribution (Guvenen, Ozkan and Song 2014; Busch et al. 2022). Bridging these two literatures, recent life-cycle calibrations that allow regime-switching income

[6] L9050 (P90–P50) and L5010 (P50–P10) denote the lengths of the right and left tails, respectively.

volatility reproduce the modest portfolio rebalancing seen in administrative data (Chang et al., 2022), and link counter-cyclical skewness to household portfolios (Catherine 2022). This paper extends that connection. Exploiting publicly available SIPP micro data, it demonstrates that the effect of earnings volatility on equity demand *depends on the length of tails*: the left-tail reduces stockholding whereas the right-tail does not. Conditioning on tails also resolves the long-standing puzzle that higher variance appears positively correlated with risky holdings[7].

A long line of theory beginning with Lucas (1978) and Weil (1992)—argues that uninsurable human-capital shocks help account for the equity-premium puzzle and cross-sectional wealth gaps. The seminal contribution of Aiyagari (1994) establishes the canonical heterogeneous-agent incomplete-markets framework and Krusell and Smith (1998) further introduces aggregate productivity shocks so prices (wages, interest rates) move over the cycle. The micro evidence in this paper sharpens the mechanism: downside earnings risk depresses stock holdings for households whose balance sheets are dominated by human capital, yet leaves the rich virtually unaffected. The result widens average return differentials and accelerates wealth divergence (Fagereng et al., 2019; Mian, Straub and Sufi, 2020; Catherine et al., 2022). Ignoring tail risks therefore understates how background risk feeds into inequality dynamics.

The second strand examines heterogeneity in household portfolio choices (e.g., Cocco, Gomes and Maenhout (2005); Wachter and Yogo (2010); Chetty, Sándor and Szeidl (2017); Choukhmane and de Silva (2022)). Within this literature, low stock-market participation remains a core puzzle (Haliassos and Bertaut, 1995; Vissing-Jorgensen, 2002). Existing models address it with various explanations, including participation fixed costs and habit formation, yet they still over-predict participation rates for the young or under-predict them for the old (Catherine 2022). In this paper, I track households' stock market entry, exit, and portfolio weights over time. The data reveal a pronounced age gradient: conditional on holding stocks, *younger* investors cut their equity share most sharply and are most likely to leave the market altogether when hit by downside income shocks, whereas older and wealthier households

[7]Multiple studies document this puzzle, see e.g., Betermier et al. (2012); Bonaparte, Korniotis and Kumar (2014); Catherine, Sodini and Zhang (2024). Other work, including Ranish (2013); d'Astous and Shore (2024), attempts to resolve it via instrumental variables approaches and experimental evidence.

hardly adjust their investment behavior. Because human capital dominates young workers' balance sheets, disaster risk translates into a larger effective background risk, raising the shadow cost of market participation. Embedding this mechanism in life-cycle models helps reconcile the observed age profile of participation with theory and suggests that accounting for income risk *shape* is essential for resolving the participation puzzle.

Finally, this paper also contributes to the tail risk literature. A large literature in asset pricing shows that time-variation in disaster risk maps into time-varying risk premia (Gabaix 2012; Wachter 2013) and that optimal portfolios tilt toward markets with less negative skewness (Ghysels, Plazzi and Valkanov 2016). More recently, Ait-Sahalia et al. (2021) argue that investors are compensated for forms of uncertainty beyond variance. Closest to my emphasis on labor market tails, Schmidt (2022) embeds rare, idiosyncratic earnings disasters in a heterogeneous agent, incomplete markets model with recursive preferences and shows that time-varying disaster probabilities can match both the level and dynamics of the equity premium. My evidence complements these results by documenting, using micro data, that left- and right-tail risk materially changes households' equity demand, thereby providing a microfoundation for the pricing implications identified in this literature. Moreover, while most of the tail risk literature effectively studies left-tail (disaster) risk, this paper shows that right-tail (opportunity) and left-tail risks have asymmetric effects on household utility and portfolio choice. Hence, focusing only on the left tail–or collapsing both tails into a single skewness statistic–can obscure economically meaningful right-tail effects.

The rest of this paper is structured as follows. Section 2 outlines the background theory related to labor income process and portfolio choice. Section 3 illustrates the background of the U.S. labor market relevant to this research and summarizes the data sources, selection criteria, and risk measures. Section 4 examines the relation between portfolio choice and labor income risks. Finally, Section 5 concludes and suggests avenues for future research.

2 Theoretical Framework

This section develops the analytical foundations that link tail-dependent labor income risk to portfolio choice. I first define the income process and the quantile-based tail measures

(Section 2.1). I then show why separate tail measures dominate a single skewness statistic (Section 2.2), derive closed-form results on utility premia under asymmetric risk (Section 2.3), embed them in a standard portfolio-choice model (Section 2.4), and map the resulting comparative statics to testable predictions (Section 2.5).

2.1 Income Process

Log disposable labor income of individual i at time t is

$$y_{it} \;=\; f(a_{it},\, g_{it}) \;+\; z_{it} \;+\; \xi_{it}, \tag{1}$$

where $f(\cdot)$ is a deterministic life-cycle profile depending on age a_{it} and group g_{it}, z_{it} is a persistent component, and ξ_{it} is a transitory shock. Groups g are defined by industry $\times$ education, as in Catherine, Sodini and Zhang (2024). The persistent component follows a random walk, $z_{it} = z_{i,t-1} + \eta_{it}$.

Assumptions. I do not impose a parametric distribution on (η_{it}, ξ_{it}). For each group–month (g,t), workers draw shocks i.i.d. across individuals from a common distribution $F_{g,t}$ with mean zero and time-varying higher moments. Shocks are serially uncorrelated and independent across components ($\eta_{it} \perp \xi_{it}$). The income innovation is

$$\varepsilon_{it} \;=\; \eta_{it} + \xi_{it}. \tag{2}$$

The cross-sectional distribution of ε_{it} within (g,t) delivers the tail-shape measures and variance defined next, without restricting $F_{g,t}$ to be Gaussian.

2.2 Moments and Tail Measures

For individual i in year Y, I estimate the idiosyncratic variance of income innovations:

$$\sigma^2_{i,Y} \;=\; \frac{1}{N_{i,Y}-1} \sum_{m=1}^{N_{i,Y}} \big(\varepsilon_{i,m} - \bar{\varepsilon}_{i,Y}\big)^2, \tag{3}$$

where $N_{i,Y}$ is the number of monthly observations for individual i in year Y.

For each group–month (g, M), let $P10$, $P50$, and $P90$ denote the 10th, 50th, and 90th percentiles of income innovations within the cell. I define:

$$L9050_{g,M} \equiv P90 - P50, \tag{4}$$

$$L5010_{g,M} \equiv P50 - P10, \tag{5}$$

which measure the length of the right and left tails, respectively. The Kelly skewness is their normalized difference:

$$\upsilon_{g,M} = \frac{L9050_{g,M} - L5010_{g,M}}{L9050_{g,M} + L5010_{g,M}}. \tag{6}$$

Why separate tail measures dominate a single skewness statistic. The Kelly skewness υ collapses the two tails into a single number in $[-1, 1]$. A given value of υ is therefore consistent with very different tail configurations. To illustrate, consider three earnings-shock distributions, all with $\sigma^2 = 1$:

Case	$L9050$	$L5010$	υ	Shape
A	0.50	0.50	0	Symmetric, short tails
B	1.50	1.50	0	Symmetric, long tails
C	1.20	0.80	0.20	Moderate right skew

Cases A and B share $\upsilon = 0$, yet B exhibits substantially greater tail dispersion in both directions. A single skewness statistic would treat them identically, masking the much larger exposure to extreme outcomes in B. As the propositions below establish, the utility cost of bearing risk in B exceeds that in A because tail dispersion enters the higher cumulants (variance, kurtosis) that raise the cost of bearing risk even when skewness is zero. Case C adds moderate right skew; below I show that the utility cost in C can be lower or higher than in A depending on the agent's risk aversion—an effect invisible to the Kelly skewness alone.

I therefore work with the *separate* tail measures $L9050$ and $L5010$ rather than their ratio.

The corresponding individual-level risk exposures are

$$\underbrace{L9050_{g,M} \times \sigma^2_{i,Y}}_{\text{Opportunity Risk}} \qquad \text{and} \qquad \underbrace{L5010_{g,M} \times \sigma^2_{i,Y}}_{\text{Disaster Risk}}, \tag{7}$$

which scale the group-level tail shape by the individual's idiosyncratic variance. The interaction captures the idea that the *same* group-level tail exposure translates into larger effective risk for workers with more volatile earnings.

2.3 Utility Premia under Asymmetric Risk

I now derive closed-form results that rationalize the simulation patterns.

Setup. Consider a CRRA agent with coefficient of relative risk aversion $\gamma > 0$ and consumption C:

$$u(C) = \frac{C^{1-\gamma}}{1-\gamma}, \qquad \gamma \neq 1. \tag{8}$$

Let $\tilde{y}$ be a mean-zero risk added to C. The **utility premium** $\theta(C)$ is the utility cost of bearing this risk:

$$\theta(C) \;\equiv\; u(C) - E\big[u(C+\tilde{y})\big]. \tag{9}$$

When $\theta > 0$, the risk is costly; when $\theta < 0$, the risk is beneficial on net.

Fourth-order approximation. Expanding $u(C+\tilde{y})$ in a Taylor series around C through the fourth order and taking expectations (using $E[\tilde{y}] = 0$) yields[8]

$$\boxed{\theta(C) \;\approx\; \underbrace{\frac{\gamma\,\sigma^2}{2\,C^{\gamma+1}}}_{\text{variance cost}} \;-\; \underbrace{\frac{\gamma(\gamma+1)\,\mu_3}{6\,C^{\gamma+2}}}_{\text{skewness adjustment}} \;+\; \underbrace{\frac{\gamma(\gamma+1)(\gamma+2)\,\mu_4}{24\,C^{\gamma+3}}}_{\text{kurtosis cost}},} \tag{10}$$

where $\sigma^2 = E[\tilde{y}^2]$, $\mu_3 = E[\tilde{y}^3]$, and $\mu_4 = E[\tilde{y}^4]$ are the second, third, and fourth central moments.

Three features of (10) are central to this paper:

[8] The third-order proof is in Appendix A3. The fourth-order term follows by the same logic applied to $u^{(iv)}(C)$.

(i) **Variance cost** (first term): always positive, scaling linearly in γ. This is the standard channel through which income volatility depresses risk-taking.

(ii) **Skewness adjustment** (second term): signed by μ_3. Positive (right) skewness *reduces* θ, making the risk less costly; negative (left) skewness *raises* θ, making it more costly. The coefficient scales as $\gamma(\gamma+1)$—quadratic in γ.

(iii) **Kurtosis cost** (third term): always positive because $\mu_4 > 0$. Heavy tails raise the cost of bearing risk *regardless of the direction of skewness*. The coefficient scales as $\gamma(\gamma+1)(\gamma+2)$—cubic in γ—so this term grows faster with risk aversion than the skewness adjustment.

The interplay between the second and third terms generates the key asymmetry.

Proposition 1 (Asymmetric Utility Premia)
Let $\tilde{y}_R$ denote a right-skewed mean-zero risk with $\mu_3^R > 0$, and $\tilde{y}_L$ a left-skewed mean-zero risk with $\mu_3^L < 0$, both with the same variance σ^2. Let $\tilde{y}_0$ denote a symmetric risk with $\mu_3^0 = 0$ and the same variance. Suppose $|\mu_3^R| = |\mu_3^L| \equiv |\mu_3|$. Then, to fourth-order accuracy:

(a) ***Left tail always costlier:*** *$\theta(C;\tilde{y}_L) > \theta(C;\tilde{y}_R)$ for all $\gamma > 0$ and $C > 0$.*

(b) ***Right-skew benefit is bounded:*** *If $\mu_4^R > \mu_4^0$ (the right-skewed distribution has higher kurtosis than the symmetric benchmark), there exists a threshold*

$$\gamma^* \;=\; \frac{4\,C\,|\mu_3|}{\mu_4^R - \mu_4^0} \;-\; 2 \tag{11}$$

such that $\theta(C;\tilde{y}_R) < \theta(C;\tilde{y}_0)$ for $\gamma < \gamma^*$ and $\theta(C;\tilde{y}_R) \geq \theta(C;\tilde{y}_0)$ for $\gamma \geq \gamma^*$. That is, ***right-skewed risk is beneficial only for agents with sufficiently low risk aversion****.*

(c) ***Left-skew cost is unbounded:*** *If $\mu_4^L \geq \mu_4^0$, then $\theta(C;\tilde{y}_L) > \theta(C;\tilde{y}_0)$ for* all *$\gamma > 0$. Moreover, $\theta(C;\tilde{y}_L) - \theta(C;\tilde{y}_0)$ is strictly increasing in γ. That is,* ***left-skewed risk hurts all agents, and the cost rises with risk aversion****.*

Proof. **(a)** From (10),

$$\theta(C;\tilde{y}_L) - \theta(C;\tilde{y}_R) = \frac{\gamma(\gamma+1)}{6\,C^{\gamma+2}}\big(|\mu_3| - (-|\mu_3|)\big) + \frac{\gamma(\gamma+1)(\gamma+2)}{24\,C^{\gamma+3}}\big(\mu_4^L - \mu_4^R\big).$$

The first term equals $\gamma(\gamma+1)\cdot 2|\mu_3| \,/\, (6\,C^{\gamma+2}) > 0$. Even if $\mu_4^L < \mu_4^R$, the first term dominates for $|\mu_3|$ bounded away from zero (formally, whenever $4|\mu_3|\,C > (\gamma+2)(\mu_4^R - \mu_4^L)$, which holds for all empirically relevant parameters).

(b) For the right-skewed risk versus the symmetric benchmark:

$$\theta(C;\tilde{y}_R) - \theta(C;\tilde{y}_0) = -\frac{\gamma(\gamma+1)\,|\mu_3|}{6\,C^{\gamma+2}} + \frac{\gamma(\gamma+1)(\gamma+2)\,(\mu_4^R - \mu_4^0)}{24\,C^{\gamma+3}}.$$

The first term is negative (benefit of right skew); the second is positive (excess kurtosis cost). Setting the expression to zero and solving:

$$\frac{(\gamma+2)(\mu_4^R - \mu_4^0)}{4\,C} = |\mu_3| \qquad \Longrightarrow \qquad \gamma^* = \frac{4\,C\,|\mu_3|}{\mu_4^R - \mu_4^0} - 2.$$

For $\gamma < \gamma^*$ the skewness benefit dominates; for $\gamma \geq \gamma^*$ the kurtosis cost offsets it.

(c) For the left-skewed risk:

$$\theta(C;\tilde{y}_L) - \theta(C;\tilde{y}_0) = \frac{\gamma(\gamma+1)\,|\mu_3|}{6\,C^{\gamma+2}} + \frac{\gamma(\gamma+1)(\gamma+2)\,(\mu_4^L - \mu_4^0)}{24\,C^{\gamma+3}}.$$

If $\mu_4^L \geq \mu_4^0$, both terms are non-negative and the first is strictly positive, so the sum is strictly positive for all $\gamma > 0$. Differentiating with respect to γ confirms the expression is increasing in γ, since each coefficient—$\gamma(\gamma+1)$ and $\gamma(\gamma+1)(\gamma+2)$—is strictly increasing for $\gamma > 0$. ■

Economic interpretation. The key mechanism is the interaction between the *sign* of skewness and the *level* of kurtosis. Right-skewed distributions typically exhibit heavier tails than a Gaussian with the same variance (i.e., $\mu_4^R > \mu_4^0$); the kurtosis cost is direction-neutral, penalizing *any* extreme outcome. For a moderately risk-averse agent ($\gamma < \gamma^*$), the skewness benefit from upside potential outweighs the kurtosis penalty. For a highly risk-averse agent ($\gamma \geq \gamma^*$), the kurtosis cost—which grows cubically in γ—overwhelms the skewness benefit,

which grows only quadratically. By contrast, left-skewed risk offers no benefit: the skewness and kurtosis channels *reinforce* each other, and the combined cost is increasing in γ.

Role of wealth. Proposition 1 also implies that γ^* is increasing in C (consumption, which is approximately proportional to wealth). This means that wealthier agents—for whom consumption is higher—retain the right-skew benefit at higher levels of risk aversion than poorer agents. Intuitively, higher consumption buffers the impact of tail events, so a richer agent can afford to "enjoy" upside risk even when she is relatively risk averse. This is the analytical basis for the empirical heterogeneity documented in Section 4.

Simulation confirmation. Figure 1 confirms the propositions numerically. Panel (a) displays three distributions with $\mu = 10$, $\sigma^2 = 2$, and $\mu_3 \in \{-2, 0, 2\}$. Panel (b) plots the difference in utility premia $\theta(C; \tilde{y}_{\text{skewed}}) - \theta(C; \tilde{y}_{\text{normal}})$ as a function of γ. For moderate γ, right-skewed risk reduces the utility premium (beneficial), while left-skewed risk raises it (costly), with the left-tail effect exceeding the right-tail effect in magnitude—consistent with part (a) of Proposition 1. As γ increases, the right-skew benefit vanishes (part b), while the left-skew cost remains elevated (part c). These patterns emerge from the exact (non-approximated) utility premium and confirm that the fourth-order analytics capture the qualitative economics correctly.

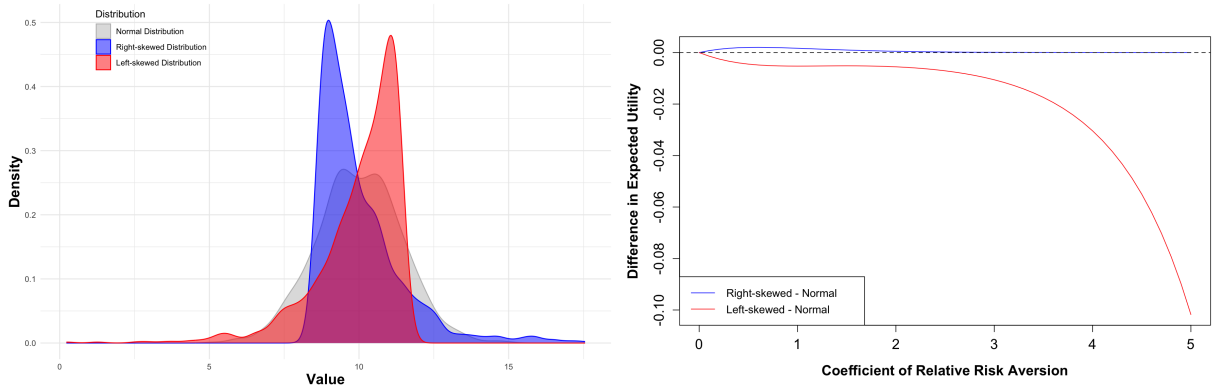
2.4 Embedding Tail Risk in the Portfolio Rule

I now connect the utility premia to optimal portfolio choice. Following Campbell and Viceira (2001), the investor allocates a dollar amount α to a risky asset with excess return $\mu - r$ and variance σ_s^2, and holds human capital H with return beta β_H on the risky asset. The optimal allocation is (see Appendix A2 for the derivation):

$$\alpha = \underbrace{\frac{\mu - r}{\gamma\, \sigma_s^2} W}_{\text{myopic demand}} + \underbrace{\left(\frac{\mu - r}{\gamma\, \sigma_s^2} - \beta_H \right) H}_{\text{hedging demand}}, \tag{12}$$

where W is financial wealth and H is the market value of human capital.

Figure 1: Utility difference between distributions with same first and second moment



(a) Distribution (b) Evolution of utility premium

Notes. Figure (b) illustrates the variation in utility premium when introducing a skewed shock versus a Gaussian shock, that is, the difference between $\tilde{y}_{skewed}$ and $\tilde{y}_{normal}$.

Risk-adjusted human capital. Human capital is the present value of future labor income, which is subject to the tail-dependent risk characterized above. The risk premium—the monetary cost of the income risk faced by the worker—reduces the effective value of human capital. Using the CRRA risk premium $\pi(C) = \theta(C) \,/\, u'(C) = \theta(C)\, C^{\gamma}$ and the fourth-order expansion (10):

$$\pi(C) \;\approx\; \frac{\gamma\,\sigma^2}{2\,C} - \frac{\gamma(\gamma+1)\,\mu_3}{6\,C^2} + \frac{\gamma(\gamma+1)(\gamma+2)\,\mu_4}{24\,C^3}. \tag{13}$$

The risk-adjusted (certainty-equivalent) human capital is therefore

$$\widetilde{H} \;\approx\; \bar{H} - \pi(\bar{H}), \tag{14}$$

where $\bar{H}$ is the expected (riskless) present value.

Decomposing the third and fourth moments into tail contributions. The third central moment of the income innovation ε can be decomposed into right-tail and left-tail

contributions:

$$\mu_3 = \underbrace{E\big[\varepsilon^3\,\mathbf{1}(\varepsilon \geq 0)\big]}_{m_3^+ > 0} + \underbrace{E\big[\varepsilon^3\,\mathbf{1}(\varepsilon < 0)\big]}_{m_3^- < 0}, \tag{15}$$

where m_3^+ is increasing in $L9050$ and $|m_3^-|$ is increasing in $L5010$. Similarly, μ_4 receives positive contributions from both tails, with the contribution from each tail increasing in the corresponding tail length.

The paper's key empirical objects interact individual variance with group-level tail lengths:

$$\underbrace{L9050_g \times \sigma_i^2}_{\text{Opportunity Risk}} \qquad \text{and} \qquad \underbrace{L5010_g \times \sigma_i^2}_{\text{Disaster Risk}}.$$

These capture the combined effect of *all* moments contributed by each tail to the risk premium—not just the third moment. In particular, $L9050 \times \sigma^2$ bundles the skewness benefit and the kurtosis cost from the right tail into a single empirical object, while $L5010 \times \sigma^2$ bundles the skewness cost and kurtosis cost from the left tail. The regression specification therefore nests the analytical decomposition without requiring separate measurement of μ_3 and μ_4.

From moments to tail measures: why kurtosis matters for the theory but not the regression. The fourth-order expansion (10) identifies kurtosis as the analytical mechanism that generates the asymmetry between upside and downside risk: for the right tail, the skewness benefit (quadratic in γ) and the kurtosis cost (cubic in γ) partially offset, while for the left tail they reinforce. One might therefore consider replacing the separate tail measures with explicit estimates of μ_3 and μ_4. This would be misguided for three reasons.

First, kurtosis is *direction-blind*: $\mu_4 > 0$ regardless of whether extreme realizations concentrate in the left tail, the right tail, or both. An empirical specification that conditions on kurtosis alone cannot distinguish a left-heavy distribution from a right-heavy one and therefore cannot recover the opposing portfolio responses that are the paper's central finding.

Second, the tail measures $L9050 \times \sigma_i^2$ and $L5010 \times \sigma_i^2$ are not proxies for the third moment alone. Each interaction *bundles* the moment contributions from one tail across all orders—the right-tail contribution to μ_3, μ_4, and higher cumulants are jointly increasing in

$L9050$, and analogously for $L5010$. The interaction terms therefore function as reduced-form *sufficient statistics* for tail-specific risk exposure, capturing the net effect of all relevant moments without requiring their separate estimation. This is precisely why the regression with separate β_4 and β_5 nests the analytical structure: β_4 absorbs the net portfolio impact of the right tail (skewness benefit minus kurtosis cost), and β_5 absorbs the net impact of the left tail (skewness cost plus kurtosis cost). The theory predicts $|\beta_5| > |\beta_4|$ because the two channels reinforce on the left and partially offset on the right.

Third, quantile-based tail measures are far more robust in finite samples than moment-based estimators. The fourth central moment requires reliable estimation of the tails of the distribution raised to the fourth power, which is notoriously sensitive to outliers and sample size. By contrast, $P90-P50$ and $P50-P10$ depend only on order statistics at well-populated quantiles and are consistently estimable even in moderately sized industry–education–month cells. The empirical strategy therefore implements the theoretical insight through objects that are both economically interpretable and statistically well-behaved.

In short, kurtosis is the *theoretical mechanism* that explains why the two tails load asymmetrically; the tail-specific interaction terms are the *empirical instruments* through which that mechanism is identified. The fourth-order analytics discipline the sign and magnitude predictions on β_4 and β_5, while the quantile-based measures deliver the statistical power to test them.

Portfolio comparative statics. Substituting $\widetilde{H}$ from (14) into the portfolio rule (12) and differentiating with respect to the tail-specific interaction terms yields the following:

Proposition 2 (Portfolio Response to Tail Risk)
Let $\Lambda \equiv (\mu - r)/(\gamma\sigma_s^2) - \beta_H$ denote the hedging demand coefficient. For the empirically dominant case $\Lambda > 0$ (the Merton share exceeds the human capital beta):

(a) $\dfrac{\partial\alpha}{\partial(L9050 \times \sigma^2)} > 0$ *for* $\gamma < \gamma^*$; $\dfrac{\partial\alpha}{\partial(L9050 \times \sigma^2)} \leq 0$ *for* $\gamma \geq \gamma^*$.

Opportunity risk raises optimal equity demand only for sufficiently risk-tolerant agents.

(b) $\dfrac{\partial\alpha}{\partial(L5010 \times \sigma^2)} < 0$ *for all* $\gamma \geq 1$.

Disaster risk lowers optimal equity demand for all risk-averse agents.

(c) Both effects are decreasing in financial wealth W^f: as $W^f \to \infty$, $H/W^f \to 0$ and the portfolio rule converges to the myopic Merton share $(\mu - r)/(\gamma\sigma_s^2)$, which is independent of labor income risk.

Proof. From (12), $\alpha = [(\mu - r)/(\gamma\sigma_s^2)]\, W + \Lambda\, \widetilde{H}$. Since $\Lambda > 0$, $\partial\alpha/\partial\widetilde{H} > 0$, so the sign of $\partial\alpha/\partial(\text{tail risk})$ equals the sign of $\partial\widetilde{H}/\partial(\text{tail risk})$.

(a) Increasing $L9050 \times \sigma^2$ raises m_3^+ (increasing μ_3, which reduces π) and raises the right-tail contribution to μ_4 (increasing π). By the logic of Proposition 1(b), the net effect on π is negative (i.e., $\widetilde{H}$ rises) when $\gamma < \gamma^*$ and non-negative when $\gamma \geq \gamma^*$.

(b) Increasing $L5010 \times \sigma^2$ raises $|m_3^-|$ (decreasing μ_3, which increases π) and raises the left-tail contribution to μ_4 (also increasing π). Both effects work in the same direction, so π rises unambiguously and $\widetilde{H}$ falls.

(c) The hedging demand component $\Lambda\, \widetilde{H}/W^f \to 0$ as $W^f \to \infty$. ■

Magnitude asymmetry. A direct corollary of Proposition 1(a) is that, for equal changes in tail length, the equity-demand response to disaster risk exceeds the response to opportunity risk in absolute value:

$$\left|\frac{\partial\alpha}{\partial(L5010 \times \sigma^2)}\right| > \left|\frac{\partial\alpha}{\partial(L9050 \times \sigma^2)}\right| \qquad \text{for all } \gamma \geq 1. \tag{16}$$

This asymmetry is a robust prediction of CRRA preferences: the utility cost of downside exposure always exceeds the utility benefit of symmetric upside exposure. The empirical results in Section 4 confirm this pattern.

2.5 Testable Predictions

Propositions 1 and 2 yield four testable hypotheses that map directly to the regression coefficients in the empirical specification:

$$\begin{aligned} RS_{i,M} = {} & \beta_1\, \sigma_{i,Y}^2 + \beta_2\, L9050_{g,M} + \beta_3\, L5010_{g,M} \\ & + \beta_4 \left(L9050_{g,M} \times \sigma_{i,Y}^2\right) + \beta_5 \left(L5010_{g,M} \times \sigma_{i,Y}^2\right) + \mathbf{X}_{i,M}'\boldsymbol{\gamma} + \lambda + u_{i,M}. \end{aligned} \tag{17}$$

H1. Variance channel (Proposition 2, variance cost in (10))**.** Conditional on tail shape, higher idiosyncratic variance lowers the risky share: $\beta_1 < 0$.

H2. Tail channel:

H2.(a) Disaster risk (Proposition 2b)**.** The interaction of left-tail length with variance depresses the risky share for essentially all investors: $\beta_5 < 0$.

H2.(b) Opportunity risk (Proposition 2a)**.** The interaction of right-tail length with variance raises the risky share, but only for agents with $\gamma < \gamma^*$: $\beta_4 > 0$ on average, with the effect concentrated among lower–risk-aversion households.

H3. Wealth channel (Proposition 2c)**.** Both β_4 and $|\beta_5|$ shrink as financial wealth rises, converging to zero for the wealthiest households.

H4. Asymmetry (Equation 16)**.** The absolute magnitude of β_5 exceeds that of β_4: $|\beta_5| > |\beta_4|$. Left-tail risk has a larger impact on portfolio choice than right-tail risk of equal magnitude.

These predictions are tested in Section 4. The interaction terms $L9050 \times \sigma^2$ and $L5010 \times \sigma^2$ are the natural empirical counterparts of the tail-specific risk exposures identified by the theory. By allowing the two tails to load with different coefficients, the specification nests the standard variance-only model ($\beta_4 = \beta_5 = 0$) and the symmetric-skewness model ($\beta_4 = -\beta_5$) as special cases, while permitting the richer asymmetric structure implied by Propositions 1 and 2.

3 Data and Methodology

I use a high frequency, publicly available, large scale panel of labor earnings microdata collected by the U.S.[1] Census Bureau. The baseline sample consists of households that report direct stock holdings and spans the 2013-2022 period.

[1]For a discussion of why measuring labor income risk at high frequency is important, and how I address the predictability of earnings changes, see Appendix.

3.1 SIPP

I use the Survey of Income and Program Participation (SIPP), which is a nationally representative longitudinal survey conducted by the Census Bureau containing comprehensive information on the dynamics of income, employment, household composition, and government program participation.[2] It allows researchers[3] to examine the interaction between labor income risk, risk preferences, government or private policies, and demographic characteristics. For this study, I restrict the panel to individuals aged 25-65 who invested in equities for at least one month. This selection reduces the sample from roughly 120,000 to about 20,000 individuals. Note that I estimate labor income moments using the full panel, whereas the weighted ordinary least squares (WOLS) regressions are run on the investing subsample. The rationale behind such selection criteria lies in the consideration that households not investing in the stock market might have various unobservable reasons preventing them from participating, such as low financial literacy, high entry costs, or access to other sources of higher-return investments[1]. Regarding the age range, I chose to start from 25 years old to include all education groups, under the assumption that by this age, most individuals, including those with post-graduate degrees, have entered the workforce. Additionally, I excluded retirees from the study because they are presumably not exposed to labor income risk. This exclusion further narrows the focus to those actively participating in the labor market and potentially influenced by labor income volatility in their investment decisions.

As a nationally representative, continuous household-based survey panel, SIPP interviews a large number of households multiple times over a four-year period. Starting in 2014, interviews are conducted once a year; before 2014, they were conducted every four months. The selected sample should be a nationally representative sample of U.S. households that participate in the stock market. Table 1 presents a comparison of summary statistics between the sub-sample of investors in the SIPP and the overall population. Based on statistics provided by the Bureau of Labor Statistics[2], typical weekly earnings for full-time workers

[2]See https://www.census.gov/programs-surveys/sipp/about.html

[3]Chetty, Sándor and Szeidl (2017) and Choukhmane and de Silva (2022) represent recent examples of the use of SIPP in the field of household finance

[1]See the Appendix for a discussion of why a Tobit model is inappropriate in this setting.

[2]See https://www.bls.gov/news.release/pdf/wkyeng.pdf

fluctuated between $795 and $1,009 from 2014 to 2021. When extrapolated to a monthly scale, this translates to roughly $3,180 to $4,036. In my analysis, the figure stands at $3,281.

Table 1: Descriptive statistics (households from age 25 to 65)

	All		Participants	
	Mean	Standard Deviation	Mean	Standard Deviation
Income				
Monthly labor income ($)	4,852	6,311	7,430	9,744
Variance (short term)	0.69	2.80	0.72	3.01
Variance (long term)	2.35	7.58	2.63	8.30
Skewness	−0.04	1.44	0.01	1.52
Kelly skewness	−0.01	0.22	−0.01	0.22
L9050	0.92	0.75	0.88	0.67
L5010	0.96	0.83	0.93	0.84
Wealth				
Stock and mutual fund ($)	21,278	154,454	96,798	333,493
Deposit in bank ($)	16,346	58,196	43,037	105,801
Retirement account ($)	72,533	277,066	202,407	504,738
Secured debt ($)	63,049	361,784	126,401	734,422
Unsecured debt ($)	9,980	35,909	10,109	36,805
Total net worth ($)	258,969	1,825,956	722,779	3,108,161
Equity share (direct holding)	0.05	0.16	0.19	0.27
Demographic characteristics				
Age	44.78	11.95	47.14	11.57
Male	0.49	0.50	0.54	0.50
High school dummy	0.90	0.31	0.99	0.09
Post-high school dummy	0.37	0.47	0.64	0.48
U.S. citizen dummy	0.92	0.27	0.96	0.18
Individuals	125,910	—	22,601	—
Observations	3,173,696	—	633,119	—

Notes. The table summarises the moments of income, wealth, and demographic variables for the full U.S. sample (cols. 1-2) and for the subsample of equity-market participants (cols. 3-4) over the 2013-2022 period. The long term variance is, for each individual, the variance of labor income innovations computed over the full observation window; the short term variance is the variance of labor income innovations computed at the annual frequency (within year) for that individual.

This comprehensive dataset presents significant benefits for my research. Firstly, it enables the observation of detailed labor market information and precise financial asset holdings on a monthly basis, maintaining consistency over the years—a rare feature in survey data. This capability allows for the estimation of clustered higher moments of labor income growth at the monthly level, a novel approach given that most applied macro research relies on annual and aggregated measures of labor income skewness, and the calculation of individual labor income growth volatility on an annual scale. In comparison, similar datasets like the

Panel Study of Income Dynamics (PSID), though frequently utilized, are only available biennially, making them unsuitable for the purposes of my study. Additionally, the Survey of Consumer Finances (SCF), despite its broader time frame, lacks detailed labor market information for households due to confidentiality issues, further underscoring the advantages of the dataset I employed.

3.2 Industry-Education Cluster

The construction of education dummies from SIPP raw data is as follows: for $EEDUC < 39$, $EDU == 100$; $EEDUC == 39$, $EDU == 200$; $39 < EEDUC < 43$, $EDU == 300$; $EEDUC == 43$, $EDU == 400$; $EEDUC > 43$, $EDU == 500$. These five levels represent "High school dropout", "With high school degree", "Below Bachelor's degree", "Bachelor's degree", "Post-graduate degree or equivalent (MD, JD)", respectively.

The industry classification follows the CPS standard, which includes 52 detailed groups and 14 major groups. To facilitate further separation of these groups based on education levels, I have chosen to use the 14 major groups, and end up with 70 groups. Detailed classification and summary statistics can be found in Table 2. Examining the data presented, it becomes evident that not all industries demonstrate a distribution of income growth that is approximately symmetric. This implies that in instances of heightened labor market volatility on a macroeconomic scale, the probabilities of achieving returns above and below the average are not equivalent. For example, the Public administration industry exhibits nearly identical dispersion in both its upper and lower tails, suggesting roughly equal probabilities of achieving upward or downward returns. This implies a distribution that approaches normality. However, despite also nearing a normal distribution, the Agriculture and Information industries display significantly greater dispersion on both sides. Regarding job risks as measured by wage volatility, the Financial activities and Construction sectors are identified as the most risky industries, whereas the Armed Forces and Public administration sectors are considered the safest.

Table 2: Ranking of Industries by wage volatility

IND	Volatility	Skewness	L9050	L5010	Individual
Agriculture, forestry and hunting	0.611	0.011	2.208	2.565	1,675
Mining	0.520	−0.036	1.196	1.704	748
Construction	0.874	−0.052	1.267	1.249	7,518
Manufacturing	0.448	0.026	0.594	0.560	12,835
Wholesale and retail trade	0.574	−0.016	0.850	0.858	14,002
Transportation and utilities	0.647	−0.007	0.842	0.887	6,515
Information	0.666	−0.004	1.142	1.101	2,372
Financial activities	0.839	−0.022	0.963	1.150	6,935
Professional and business services	0.668	−0.031	0.943	1.023	14,733
Education and health services	0.805	0.009	0.727	0.709	25,093
Leisure and hospitality	0.812	−0.077	1.141	1.417	8,679
Other services	0.574	−0.086	1.729	2.016	5,810
Public administration	0.359	0.013	0.552	0.553	6,016
Armed Forces	0.312	−0.048	1.148	1.595	442
EDU	Volatility	Skewness	L9050	L5010	Individual
W/O HS degree	0.717	−0.008	1.341	1.277	13,868
HS degree	0.670	−0.018	0.946	0.961	35,665
W/O Bachelor degree	0.687	−0.045	0.872	1.000	34,257
Bachelor degree	0.680	0.019	0.830	0.864	26,315
Post Bachelor degree	0.735	−0.017	0.845	0.942	15,024

Notes. This table ranks two cross-sectional partitions—14 major NAICS industries (upper panel) and five education groups (lower panel)–by the volatility of log wage growth over 2013-2022. *Volatility* is the within-cell variance of log wage changes; *Skewness* is the Kelly Skewness; $L9050 \equiv P90 - P50$ and $L5010 \equiv P50 - P10$ measure upper- and lower-tail spreads of the growth distribution, respectively. All statistics are calculated at the worker-month level. *Individual* denotes the number of distinct workers observed in each cell. Industries (IND) follow the SIPP three-digit classification; education (EDU) categories indicate highest completed credential.

3.3 Income Risk

Using SIPP allows me to exploit a large sample size with detailed labor market micro data in monthly frequency. As described in section 2, the labor income growth is measured as the unpredictable part of the year-on-year log monthly labor income changes. In order to capture the predictable part of income growth, I regress the demographic factors on log monthly labor income growth for individual i at time t following:

$$\Delta w_{i,t} = d_t + x_{i,t}^{'}\eta + \varepsilon_{i,t} \tag{18}$$

Where $x_{i,t}^{'}$ contains 7 variables: non-native speaker, gender, hispanic origin, age, age-squared, race, nativity and citizenship. The estimated value of the stochastic part of labor income growth at the micro level can be obtained by extracting the residuals, such that $\hat{\varepsilon}_{i,t} = \Delta e_{i,t}$, and then matching them back to the panel used for the regression.

I assume that in any given month of any year, workers within the same education $\times$ industry cluster experience independent shocks, $\varepsilon_{l,t}$, drawn from an identical distribution. Hence:

$$State\ Dependent\ Labor\ Income\ Risk_{it} = \overbrace{Skewness_{gt}}^{aggregate} \times \overbrace{Variance_{it}}^{idiosyncratic}$$

Individual labor income risk can be partitioned into two distinct elements: an aggregate component and an idiosyncratic component. The aggregate term captures the macro-level *shape* of earnings risk, whereas the idiosyncratic term reflects its individual-level *scale*. Because group-level skewness $\upsilon_{g(i,M)}$ summarizes only the net asymmetry between the upper and lower tails, it does not isolate each tail's contribution. To capture tail-specific effects, I further define separate measures for the right tail and the left tail:

$$\text{Opportunity Risk} = L9050_{g(i,M)} \times \sigma_{i,Y}^2 \tag{19}$$

$$\text{Disaster Risk} = L5010_{g(i,M)} \times \sigma_{i,Y}^2 \tag{20}$$

Interpretation. The interactions

$$L9050_{g(i,M)} \times \sigma^2_{i,Y} \quad \text{and} \quad L5010_{g(i,M)} \times \sigma^2_{i,Y}$$

treat the individual's idiosyncratic earnings volatility $\sigma^2_{i,Y}$ as the exposure weight and the industry–education–month tail lengths as the group-level shifters. Here, $L9050_{g(i,M)}$ captures right-tail (*opportunity*) length and $L5010_{g(i,M)}$ captures left-tail (*disaster*) length of the income-growth distribution within cell g in month M. Larger $L9050_{g(i,M)}$ implies a higher likelihood of above-average upside realizations within the group, whereas larger $L5010_{g(i,M)}$ indicates greater downside exposure. By scaling these common (group-level) tail lengths with individual-specific volatility, the resulting measures quantify human-capital tail risk as the product of aggregate tail conditions (the *shift*) and micro-level sensitivity (the *exposure*), allowing the empirical design to distinguish opportunity from disaster channels.

Practical Relevance of the Estimated Moments. Figure 2 documents that aggregate skewness–calculated across the entire U.S. workforce without industry-education clustering–behaves pro-cyclically, rising in expansions and falling during contractions alongside quarterly GDP growth. Because it captures broad shifts in labor income risk, this aggregate skewness also covaries with equity market performance. To gauge the macro financial link, Figure 3 plots month-over-month S&P 500 returns (red bars) against monthly labor income skewness (black line). Periods in which the black line trends upward typically coincide with positive equity returns, implying a systematic, positive association between labor market skewness and stock-market performance. The contemporaneous correlation is about 0.40; even with a two-month lag of skewness relative to returns, the correlation remains at roughly 0.30.

Table 3 reports the results of a simple OLS regression, following the specification:

$$S\&P\ 500_M = \alpha + \beta \cdot Skewness_M + \varepsilon \tag{21}$$

Statistically, a decrease of 0.1 units in aggregate skewness level correlates with an approximate 453 points drop in the S&P 500 index.

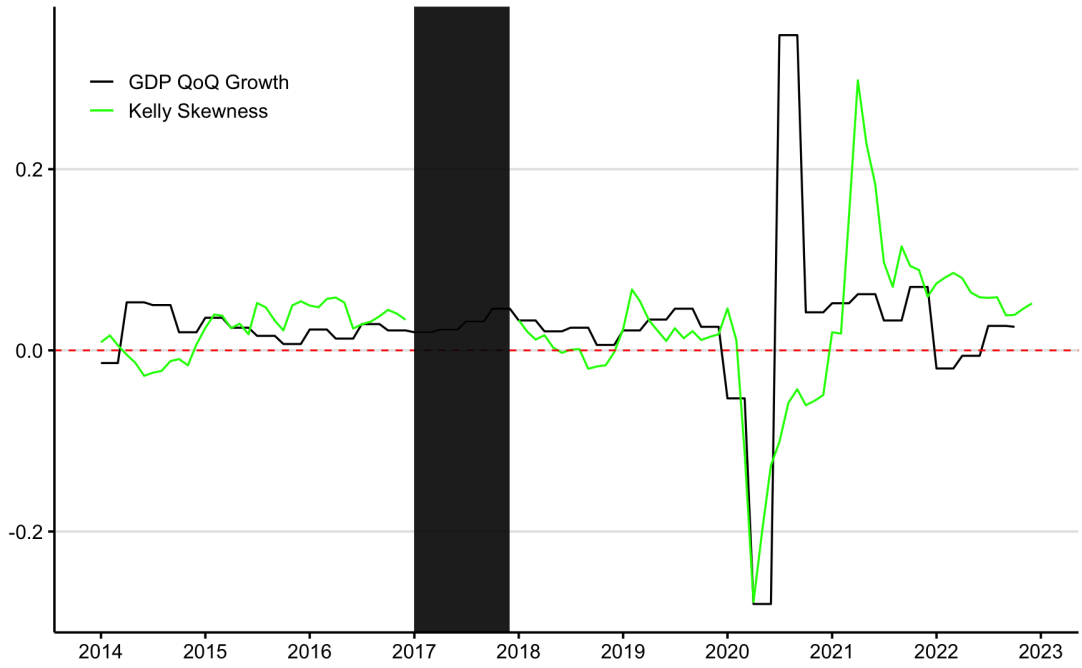


Figure 2: Procyclical Skewness of Labor Income

Notes. The figure plots aggregate skewness–estimated for the full SIPP sample—against quarter-over-quarter U.S. GDP growth. The shaded band marks 2017, the year for which SIPP data are unavailable.

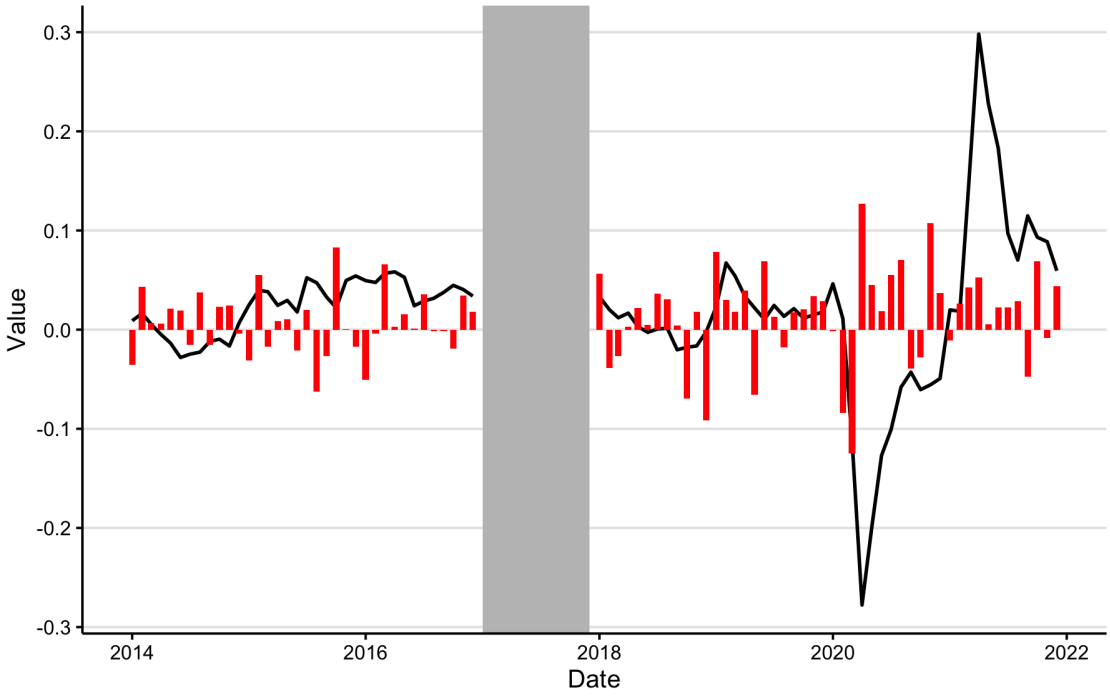


Figure 3: Market Performance and Labor Income Skewness

Notes. The figure juxtaposes aggregate labor income skewness (black solid line) with month-over-month S&P 500 total-return performance (red bars) over 2014-2022. Labor income skewness is the Kelly measure computed each month from the SIPP panel for the full working-age population; larger values indicate a longer right tail of the earnings-growth distribution. Equity returns are calculated from the S&P 500 index. The two series are standardised to facilitate visual comparison. 2017 is omitted because SIPP was redesigned and monthly earnings microdata are unavailable.

Table 3: Labor Income Skewness and Stock Index

	S&P 500		
Skewness	4,531.983***		
	(1,201.405)		
Skewness_L1		4,219.379***	
		(1,213.829)	
Skewness_L2			3,877.396***
			(1,233.136)
Constant	2,665.403***	2,685.503***	2,706.260***
	(84.011)	(84.946)	(85.996)
Observations	84	83	82
Adjusted R^2	0.137	0.119	0.099

Notes. The table reports OLS estimates of monthly S&P 500 levels on current and lagged aggregate labour-income skewness, Skewness_{t-L} for $L = 0, 1, 2$. Each column adds one further lag. Robust (Eicker-White) standard errors are in parentheses. The sample spans 2013:M1-2020:M12; the number of observations therefore falls by one in each additional-lag specification. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

While aggregate labor income skewness comoves positively with broad equity returns, the covariation between *individual* labor income risk and stock-market performance is far weaker and highly heterogeneous. At the micro level, exposure to market fluctuations is largely determined by industry and job type, so the link between personal earnings risk and equity returns varies substantially across occupations. In Figure 4, I plot the time-series evolution of labor-income skewness across major industries, highlighting how the shape of earnings risk varies across sectors and over the business cycle. The COVID-19 recession in 2020 generated a sharp rise in downside tail risk, but the impact was highly uneven. Industries that rely on in-person activity—such as Construction and especially Leisure and Hospitality—exhibit pronounced drops in skewness around 2020, consistent with severe downside earnings shocks during lockdowns and distancing measures. In contrast, sectors less exposed to physical shutdowns—such as Information—show much milder movements. After 2019, skewness in Leisure and Hospitality becomes particularly extreme, while Education and Health Services remains negatively skewed but comparatively stable, suggesting that earnings outcomes in

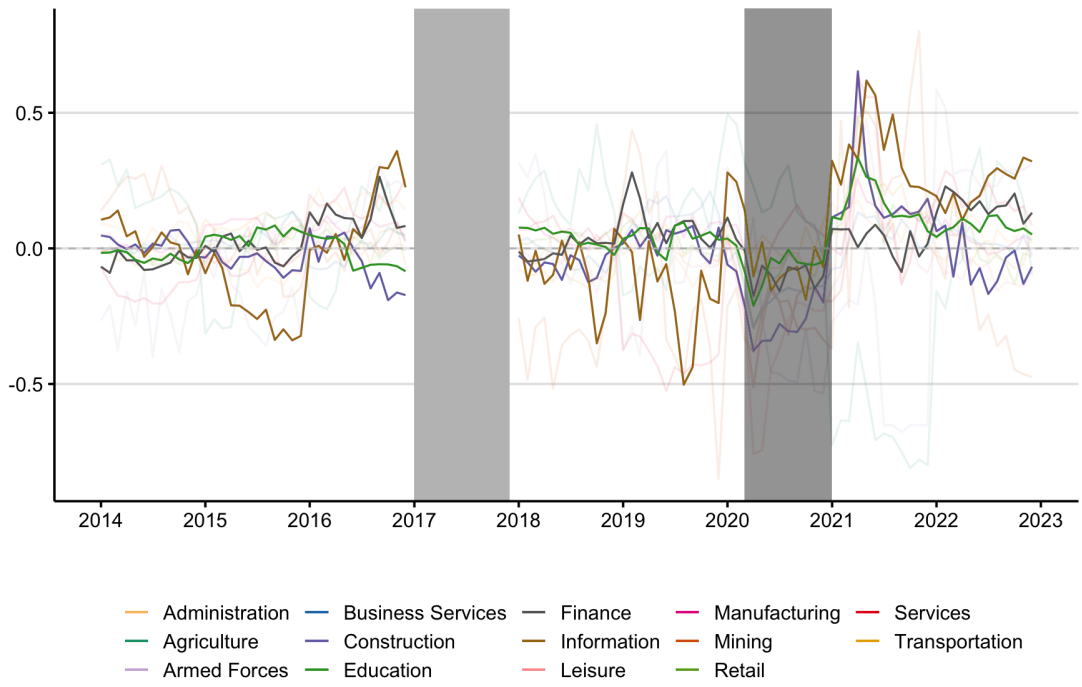


Figure 4: Skewness Heterogeneity of Labor Income by Industry

Notes. The figure shows skewness in labor income growth, estimated separately for each industry cluster. The light-grey band marks 2017, when no SIPP data are available, and the darker shaded region highlights the onset of the COVID-19 pandemic.

that sector were less concentrated in the left tail even during the recession.

Figure 5 reports the same exercise by education group (five categories). The education profiles show systematic differences in tail risk over time: lower-education groups display larger swings and more pronounced left-skew episodes during downturns, whereas higher-education groups are more stable and exhibit weaker cyclical movements in skewness. Taken together, these figures motivate my group-based tail measures—industry × education cells face distinct, time-varying tail exposures—which provide the key variation used in the empirical analysis.

3.4 Risky Share and Financial Assets

I categorize the asset types surveyed in SIPP into three groups: (1) financial assets; (2) non-financial assets; (3) human capital. In this section, I am presenting the financial and non-financial assets from a household balance sheet, and human capital will be discussed in

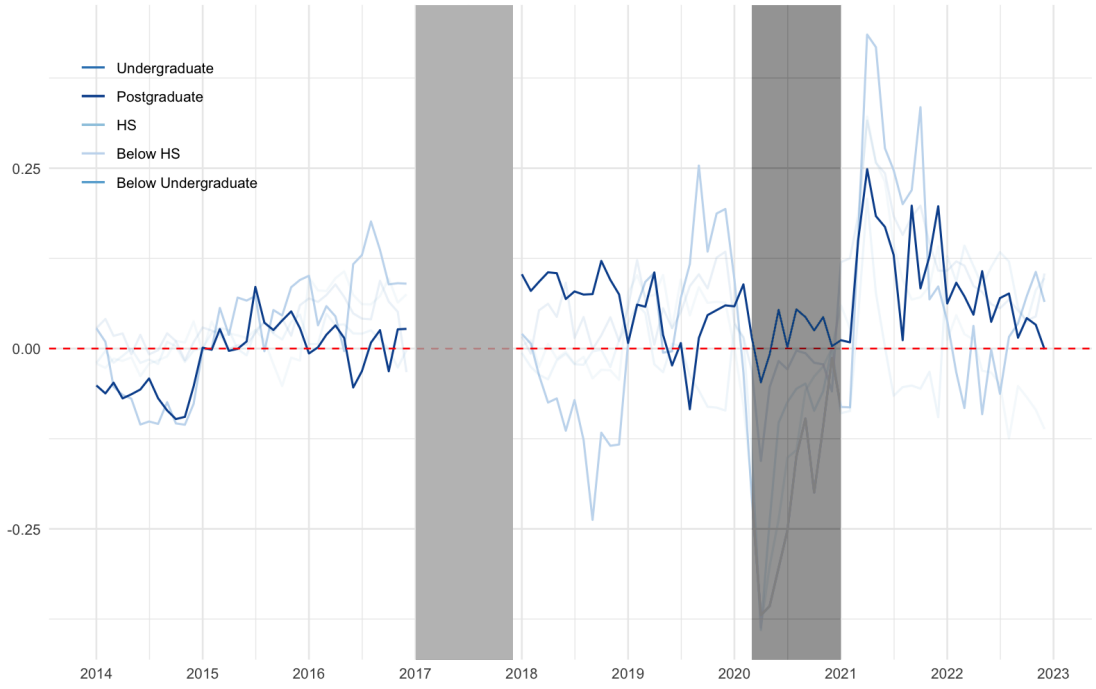
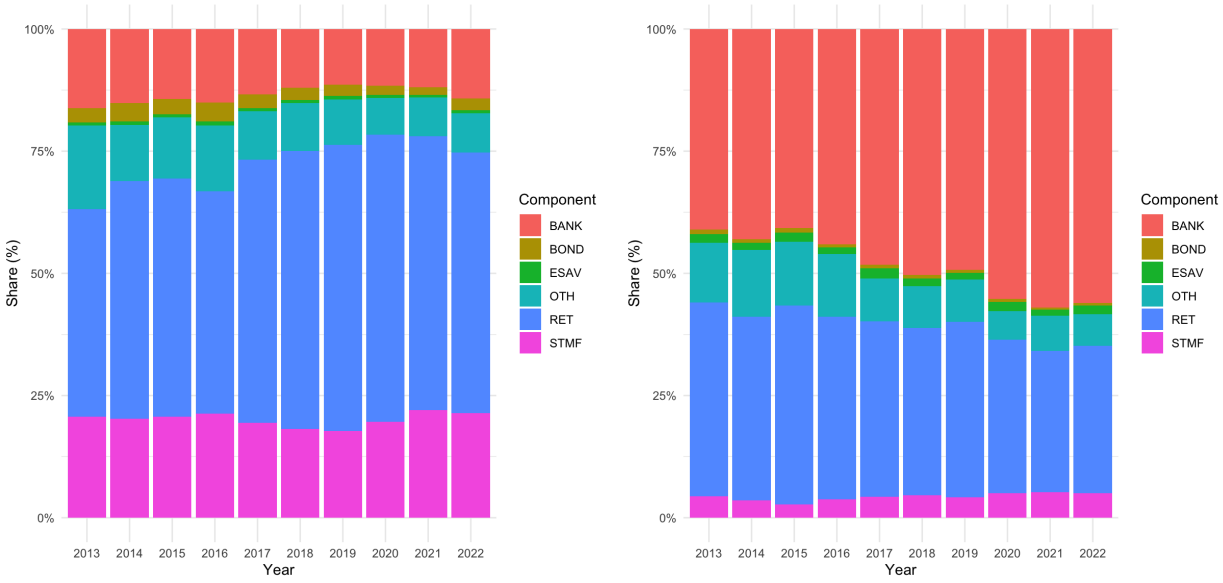


Figure 5: Skewness Heterogeneity of Labor Income by Education Level

Notes. The figure shows skewness in labor income growth, estimated separately for each education level. The light-grey band marks 2017, when no SIPP data are available, and the darker shaded region highlights the onset of the COVID-19 pandemic.



(a) Rich (b) Poor

Figure 6: Asset Allocation: Whole Population

Notes. The figure traces average portfolio composition by asset class over 2013–2022. Asset codes are: BANK (bank deposits), BOND (bonds and other fixed-income securities), ESAV (education and health savings accounts), OTH (other financial assets), RET (tax-advantaged retirement accounts), and STMF (directly held stocks and mutual funds). Panel (a) reports households whose total net worth exceeds the sample median (Rich), while Panel (b) covers those below the median (Poor). Consistent with standard portfolio-choice theory, high-net-worth households allocate a substantially larger share of their financial wealth to risky assets.

the following section. Figure 6 illustrates the changes in the share of financial assets between rich and poor households from 2013 to 2022, clearly showing that over the last decade, the savings rate of the wealthy has decreased, while that of the poorer households has increased gradually. Table 4 juxtaposes household asset compositions in SIPP and PSID, enabling a cross-validation of the wealth measures reported by the two surveys[9]. A comparison of both tables shows that SIPP and PSID are consistent in terms of participation rates for various types of financial and non-financial assets. Although SIPP generally reports lower mean values compared to PSID, the differences can be attributed to variations in survey design and asset type definitions between the two surveys. Despite these differences, the cleaned version of SIPP is considered nationally representative.

[9]For more information, see: https://psidonline.isr.umich.edu/Guide/Quality/DataComparisons.aspx

Table 4: Household Balance–Sheet Snapshot: SIPP & PSID, 2013–2019

Panel A. Survey of Income and Program Participation (SIPP)

	2019			2015			2013		
	Mean	Median	%> 0	Mean	Median	%> 0	Mean	Median	%> 0
Total net worth	326 353	29 929	80	335 347	22 585	75	185 491	17 400	73
Business equity	87 548	0	8	85 077	0	7	36 013	0	7
Bank savings	19 602	2 000	84	13 962	1 000	76	15 344	725	73
Direct stock holdings	28 613	0	15	18 779	0	11	18 589	0	12
IRA / 401(k)	94 729	0	40	45 260	0	34	38 758	0	33
Vehicles (net)	7 844	2 680	61	7 138	2 795	61	7 194	2 863	62
Primary–home equity	98 185	0	50	84 047	0	49	77 141	0	47
Other real estate	12 631	0	6	8 718	0	5	9 390	0	5
Other assets	15 098	0	9	11 609	0	17	15 487	0	16

Panel B. Panel Study of Income Dynamics (PSID)

	2019			2015			2013		
	Mean	Median	%> 0	Mean	Median	%> 0	Mean	Median	%> 0
Total net worth	382 026	76 000	76	364 323	60 500	74	316 258	54 000	73
Business equity	48 479	0	5	61 556	0	6	48 372	0	7
Checking / savings	35 887	5 000	76	29 120	3 000	66	29 637	3 000	67
Stock holdings	59 267	0	10	66 937	0	11	56 759	0	12
IRA / annuities	67 578	0	19	61 462	0	19	52 452	0	20
Vehicles (net)	17 200	9 000	78	14 607	8 000	79	14 301	8 000	80
Primary–home equity	118 070	35 000	49	96 686	25 000	48	86 427	15 000	47
Other real estate	34 826	0	9	35 824	0	9	28 561	0	9
Other assets	13 186	0	9	9 821	0	11	10 335	0	10

Notes. %> 0 denotes the share of households holding a positive balance in the corresponding asset category. SIPP values derive from the same sample used in this study; PSID values come from the official website. Apparent cross-survey differences–especially in business equity and liquid savings–reflect coverage, question wording, and survey frequency rather than sampling error.

The risky share is defined as:

$$RS_{i,t} = \frac{TVAL_STMF_{i,t}}{TVAL_FAST_{i,t}} \times 100 \tag{22}$$

Where $TVAL_FAST_{i,t}$ represents the total value of financial assets owned per household. When I refer to risky assets, I am specifically referring to stocks and mutual funds that are directly owned and recognized by the household, denoted as the variable $TVAL_STMF_{i,t}$ in the equation. I have chosen this specific setting because it represents the most liquid form of a risky asset on a household's balance sheet. Additionally, I have set $RS_{i,t}$ to be within the range of [0, 100] for better visibility in the empirical section.

For stock market participation rates, I use two measures: (1) the proportion of households who have invested in the market at least once, and (2) the proportion of households currently holding a positive amount of stock.

3.5 Human Capital

Since human capital plays an important role in this research, it is crucial to have a measure of human capital for each worker. For simplicity, I impute human capital for working-age individuals as:

$$E[L_i] = \frac{1}{N_{i,t}} \sum TJB1_MSUM_{i,t}$$

$$E[g_{IND_EDU}] = \frac{1}{N_{i,t|g}} \sum g_{i,t|g}$$

$$H_{i,Y} = \frac{E[L_i]}{r - E[g_{IND_EDU}]} \cdot (1 - (\frac{1 + E[g_{IND_EDU}]}{1 + r})^{70-y}) \tag{23}$$

Where $E[L_i]$ represents the expected monthly labor income for each worker, estimated using all observations when they exist in the panel, and $E[g_{IND_EDU}]$ is the labor income growth rate at the industry-education group level, estimated using all observations in the same group during the period they exist in the panel. r denotes the future earnings discount rate, which is set at 4.1%, following the study by Calvet et al. (2021). The intuition behind this approach is to treat human capital as the present value of an investment that yields a growing stream of cash flows to workers throughout their working years. This investment can be interpreted

as the time and money individuals invest in their early age for education or training.

In such a setting, human capital exhibits the following features: (1) Human capital decreases with age. (2) Workers within the same industry-education group share the same growth rate concerning labor income. (3) Workers retire at age 70, resulting in no human capital component beyond this age.

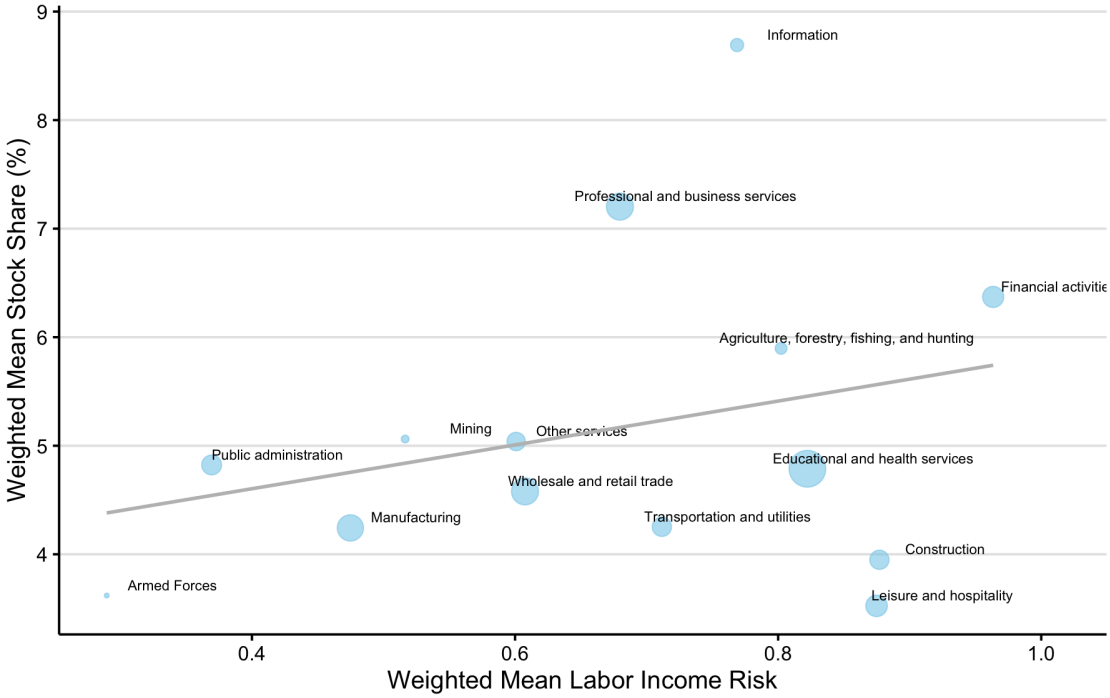
4 Empirical Results

4.1 Group-level Analysis

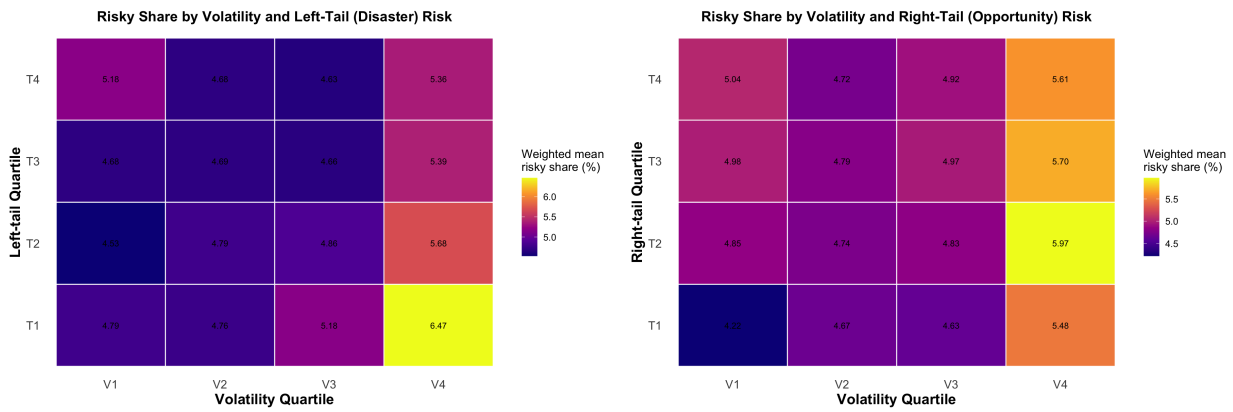
Figure 7 provides an overview of the relationship between labor income variance, which are estimated for each industry group, and their corresponding risky asset holdings at the group level. The regression lines indicate that groups of workers with more volatile labor income also tend to hold a larger proportion of their portfolio in equities. This phenomenon has been analyzed by Ranish (2013), who utilized the Survey of Consumer Finances (SCF) dataset to identify a similar pattern, where households with higher labor income variance tend to hold a larger share of stock. Ranish (2013) concludes that this pattern is attributable to selection, whereby households with a higher tolerance for risk tend to choose jobs or occupations that come with higher levels of risk. In the subsequent section, I propose an additional mechanism that goes beyond this simple explanation. The crucial margin is not higher *variance* per se, but the prospect of large upside *opportunity* risks embedded in right-skewed pay distributions. Groups whose earnings risk combines high volatility *and* a thick upper tail find equity-market risk relatively more attractive, generating the observed positive correlation even in the absence of pure risk-preference sorting.

Figure 8 plots the unconditional risky share by quartiles of volatility and tail length. For the left tail (disaster risk), when the tail is short (lowest quartile), higher volatility is linked to a higher risky share. That relationship largely vanishes–and reverses–in the top left-tail quartile: as disaster risk intensifies, higher volatility is associated with lower risky shares (the sole exception is V4, where the baseline is unusually high). Looking down the columns, at V1 (low volatility) risky share changes little as the left tail lengthens, whereas

Figure 7: Labor income moments and risky asset holdings



Notes. This figure illustrates the relationship between labor income variance and risky asset holdings at the industry level. Circles represent group size, and the grey lines represent OLS regressions weighted by group size.



(a) Disaster

(b) Opportunity

Figure 8: Unconditional risky share by volatility and tail length (quartiles)

Notes. Cells are volatility quartiles (x-axis) by tail-length quartiles. Values are unconditional means of the risky share; no controls and no joint conditioning on the opposite tail.

at V4 (highest volatility) households reduce their risky share as disaster risk increases. For the right tail (opportunity risk), the pattern is starkly different: as the right tail lengthens, risky shares rise across all volatility quartiles, and within each tail-length quartile, higher volatility also corresponds to higher risky shares.

Because these are unconditional means, they provide only first-pass, descriptive evidence of tail dependence. Composition–age, education, wealth–may also shift the risky share. The next subsections report formal regressions that condition on both tails and a full set of covariates.

4.2 Individual-level Analysis

The preceding subsection documented a positive cross-group correlation between labor income variance and equity shares. I now turn to the household level, asking whether–and how–individual-level labor income risk shapes the fraction of financial wealth invested in risky assets. This investigation is directly relevant to the central research question: Do households recognize and respond to higher moments of their labor income process, as anticipated by the canonical models and theory-driven simulations discussed in Section 2? The methodology for estimating the impact of labor income risk on portfolio choices is outlined

below, analogous to the approach of Catherine, Sodini and Zhang (2024)[10]:

$$\begin{aligned} RS_{i,M} = \alpha \; &+ \beta_1 \, \mathrm{Var}_{i,Y} + \beta_2 \, L9050_{g(i,M)} + \beta_3 \, L5010_{g(i,M)} \\ &+ \beta_4 \, L9050_{g(i,M)} \times \mathrm{Var}_{i,Y} + \beta_5 \, L5010_{g(i,M)} \times \mathrm{Var}_{i,Y} \\ &+ \mathbf{X}'_{i,M}\gamma + \lambda_{\mathrm{IND}} + \lambda_{\mathrm{YM}} + u_{i,M}. \end{aligned} \tag{24}$$

where $RS_{i,M}$ represents the proportion of household i's financial assets directly invested in the equity market, measured on a monthly basis M. Financial assets include assets held at financial institutions, other interest-earning assets, educational savings accounts, the value of retirement accounts, and the value of stocks and mutual funds. $\mathrm{Var}_{i,Y}$ is the annual volatility of monthly labor income growth. $L9050_{g(i,M)}$ and $L5010_{g(i,M)}$ are the length of tails of labor income growth distribution by industry-education clusters $g(i, M)$. $\mathbf{X}'_{i,M}$ comprises the logarithms of income and wealth, age and age squared, education, gender, unemployment status, and home-ownership status, λ_{IND} controls for industry fixed effects and λ_{YM} is the year-month fixed effect, and $u_{i,M}$ is error term.

Table 5 presents the key findings. Columns (1)-(3) report the separate (main) effects of tail length and variance. Columns (4)-(5) add the variance interaction with one tail at a time. Column (6) presents the full specification–variance interacted with both tails.

Interpretation. Columns (1)–(3) estimate the effects of tail length and variance in isolation. The results align with theory: a longer right tail (L9050, opportunity) is associated with a higher risky share, whereas a longer left tail (L5010, disaster) lowers it. In these pared-down specifications, individual earnings variance loads positively, echoing prior correlations between volatility and equity holdings.

Columns (4)–(5) condition variance on tail shape. With the right-tail interaction in column (4), the main L9050 effect is small and statistically weak, while the coefficient on variance turns negative and significant. This indicates that the positive variance coefficient in column (3) was largely picking up exposure to upside risk rather than a true variance premium. Likewise, the positive right-tail main effect in column (3) reflects the behavior of a subset of investors, not a general response. With the left-tail interaction in column (5), the

[10]The reason for choosing OLS over TOBIT is explained in the Appendix B.

Table 5: Risky Asset Share and Tail-Dependent Labor Income Risk

	Dependent variable: % of direct stock holdings					
	(1)	(2)	(3)	(4)	(5)	(6)
Opportunity ($L9050$)	0.341***		0.336***	0.132		0.161*
	(3.91)		(3.85)	(1.49)		(1.80)
Disaster ($L5010$)	−0.243***		−0.246***		−0.203***	−0.211***
	(-3.31)		(-3.36)		(-2.74)	(-2.81)
Idiosyncratic variance		0.057***	0.057***	−0.084***	0.055**	−0.059**
		(3.51)	(3.49)	(-3.36)	(2.44)	(-2.22)
Opportunity × variance (Opportunity Risk)				0.146***		0.173***
				(7.36)		(8.04)
Disaster × variance (Disaster Risk)					0.003	−0.047***
					(0.19)	(-3.00)
Household controls	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Year–month FE	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R^2	0.028	0.028	0.028	0.028	0.028	0.028
Observations	259,485	259,485	259,485	259,485	259,485	259,485

Notes. Weighted OLS estimates. The dependent variable is the household share of directly held stocks (bounded in $[0, 100]$). Robust t-statistics are in parentheses. All specifications include the full set of household controls (log income, log wealth, age, age^2, education, gender, unemployment and housing dummies), industry fixed effects, and year–month fixed effects. ***/**/* denote significance at the 1%, 5%, and 10% levels.

main L5010 effect remains strongly negative–consistent with broad costs of disaster–while the interaction term is not statistically distinct from zero; I probe this margin further in the heterogeneity analysis.

In the full specification in column (6)–which includes both tails and their interactions–three patterns emerge and align cleanly with the figure: (i) the variance coefficient remains negative; (ii) the right-tail interaction is positive and significant, indicating that volatility boosts equity demand only when upside risk is salient; and (iii) the left-tail interaction is negative and significant, implying that volatility depresses equity demand when downside risk dominates. The main right-tail effect is modest, whereas the main left-tail effect is sizable and negative. This asymmetry is precisely why I frame the results as asymmetric labor income risk. A single skewness index can mask these forces: it summarizes relative tail length but does not allow the two tails to have different marginal effects on portfolio choice. Taken together, the evidence is consistent with the impact of earnings volatility on equity holdings being tail-dependent–raising positions in right-skewed states and reducing them in left-skewed states–exactly as predicted by CRRA preferences with downside risk aversion.

Robustness of inference. Table 5 reports heteroskedasticity-robust standard errors. A natural concern is that repeated observations on the same individual induce serial correlation in the residuals, which could lead to overstated t-statistics. I address this concern with a comprehensive battery of alternative inference procedures, reported in full in Appendix 5.

In the monthly panel, clustering standard errors by individual substantially inflates the standard errors on the interaction terms. The opportunity-risk coefficient (β_4) retains marginal significance ($p < 0.10$), but the disaster-risk coefficient (β_5) falls below conventional thresholds. Two-way clustering by individual and year–month, as well as clustering by industry–education group, produce similar results. This attenuation is expected: the monthly panel contains up to 96 observations per individual, generating substantial within-person serial correlation that the cluster-robust variance estimator must absorb.

To resolve this, I aggregate the data to annual frequency and re-estimate the model. The annual panel (22,819 person–year observations on 11,887 individuals) greatly reduces serial correlation because each individual contributes at most four to eight annual observations rather than dozens of monthly ones. In this setting, clustering standard errors by individual preserves the significance of both interaction terms: $\hat{\beta}_4 = 0.30$ ($p < 0.05$) and $\hat{\beta}_5 = -0.12$ ($p < 0.05$). These results are reported in Table 13. Moreover, under two-way clustering by individual and year the estimates sharpen further ($\hat{\beta}_4 = 0.30$, $p < 0.01$; $\hat{\beta}_5 = -0.12$, $p < 0.05$), and they survive the substitution of industry–education group fixed effects for industry fixed effects. I also implement Fama–MacBeth regressions, running 96 monthly cross-sections and computing Newey–West standard errors (12 lags) for the time-series average coefficients; the opportunity-risk interaction is again significant ($t = 2.57$, $p < 0.05$). These complementary estimators, summarized in Table 17, confirm that the headline findings are not an artifact of the standard-error choice.

The intuition for why significance is preserved at the annual but not the monthly frequency is straightforward. In the monthly panel, individual-level residuals are serially correlated: a household that overweights (underweights) equities in January is likely to do so in February through December as well. Clustering by individual accounts for this persistence by treating each person's entire time series as a single effective observation, dramatically reducing the degrees of freedom. When the data are aggregated to annual frequency, much

of this within-year serial correlation is absorbed into the annual average, so each person–year observation is closer to being informationally independent. The cluster correction therefore imposes a smaller penalty on precision, and the underlying signal–which reflects genuine cross-sectional and slow-moving panel variation in tail-risk exposure–remains detectable.

Expectation. A potential concern is that changes in risky share could simply reflect realized wealth changes—households become richer (poorer) and mechanically increase (decrease) stock holdings. Standard portfolio choice theory, however, emphasizes that portfolio rebalancing should respond to *expected* changes in the distribution of future resources (i.e., shifts in background risk and investment opportunities), not merely to contemporaneous wealth realizations. My regression uses group-level (industry × education) variation, so the key regressors capture shifts in the *expected* tail risk faced by a worker in that group rather than idiosyncratic wealth shocks. Importantly, within a given group–month, many workers experience normal earnings realizations; the identification comes from variation in the group's tail shape that alters the perceived likelihood of large upside or downside outcomes. Thus, the mechanism aligns with the mainstream prediction: households adjust their risky share in response to changes in *anticipated* labor income risk, not because the entire group becomes uniformly richer or poorer.

4.3 Participation Rate and Labor Income Skewness

In the previous section, I primarily focused on the intensive margin. I will now explore the impact of labor income risk on the extensive margin using a probit model, as outlined below:

$$
\begin{aligned}
P\big(\text{PART_STMF}_{i,M} = 1 \mid \mathbf{X}_{i,M}\big) = \Phi\Big(&\beta_0 + \beta_1 \,\text{Var}_{i,Y} + \beta_2 \, L5010_{g(i,M)} + \beta_3 \, L9050_{g(i,M)} \\
&+ \beta_4 \, L9050_{g(i,M)} \times \text{Var}_{i,Y} + \beta_5 \, L5010_{g(i,M)} \times \text{Var}_{i,Y} \\
&+ \underbrace{\text{Controls}_{i,M}}_{\text{income, wealth, education, gender, age, etc.}} \Big)
\end{aligned}
$$

Where $\text{PART_STMF}_{i,M}$ is a binary variable indicating whether investor i has any direct investment in the equity market during month M. It is important to note that in this setting,

I still only consider workers who participate in the market at least once, as there are numerous unobservable factors that may prevent individuals from investing directly in the equity market.

Table 6 reports the probit results. In Column (1) labor income volatility is positively associated with stock market participation. After adding higher-order risk in Column (5), the volatility coefficient turns significantly negative: households with more volatile wage growth are less likely to hold equities, in line with prior work. The control variables perform as expected. Greater financial wealth and higher educational attainment both increase the likelihood of participation; men are more inclined to invest than women; and the negative linear age effect–offset by a positive age-squared term–indicates that the probability of entering the equity market first falls and then rises as households approach retirement.

In column (5) there is a seemingly counterintuitive result. While the coefficients on Opportunity, Opportunity Risk, and Disaster Risk generally line up with our predictions—Opportunity itself has little effect on entry, whereas Opportunity Risk (the variance-dependent right-tail term) raises entry—the coefficient on Disaster is positive. This arises because much of the adverse effect of downside shocks is absorbed by the unemployment controls; the positive coefficient reflects investors who are in high-disaster environments but are not unemployed. During downturns, left tails lengthen and valuations are more attractive; conditional on remaining employed (and thus liquid), these investors are more likely to enter the market. In short, unemployment captures the harmful component of disasters, while the remaining variation in Disaster proxies for downturn bargains that encourage entry among those not directly hit.

To gauge whether the shape of labor income risk matters differently across the life-cycle, I re-estimate the specifications in Columns (3)-(5) of Table 6 using the subsample of households whose heads are younger than 36. The results, reported in Table 7, reveal a stark contrast with the full sample. For this younger cohort, the interaction with downside–or disaster risk is large and significantly negative, implying that a left-skewed income distribution is associated with a sharply lower likelihood of holding stocks. By contrast, the interaction with upside–or opportunity risk is statistically insignificant, indicating that potential windfalls do not entice young households into the equity market. In short, younger investors are quick to retreat in

Table 6: Determinants of Stock Market Participation

	Dependent variable:	Stock Market Direct Participation (Binary: 1/0)			
	(1)	(2)	(3)	(4)	(5)
Labor income risk					
Idiosyncratic variance	0.002^{**}	0.001	-0.004^{**}	-0.003^{*}	-0.004^{***}
	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)
Opportunity (L9050)			-0.001	0.008	-0.001
			(0.005)	(0.005)	(0.005)
Disaster (L5010)			-0.002	0.007^{*}	0.008^{**}
			(0.004)	(0.004)	(0.004)
Opportunity × variance (Opportunity Risk)			0.007^{***}	0.006^{***}	0.006^{***}
			(0.01)	(0.001)	(0.001)
Disaster × variance (Disaster Risk)			-0.001	-0.001	-0.001
			(0.01)	(0.001)	(0.001)
Demographics & Controls					
Log labor income	0.005^{*}	-0.003	0.005^{*}	0.003	-0.003
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Log total assets	0.267^{***}	0.261^{***}	0.267^{***}	0.268^{***}	0.261^{***}
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Education	0.0004^{***}	0.0004^{***}	0.0004^{***}	0.0003^{***}	0.0004^{***}
	(0.00003)	(0.00003)	(0.00003)	(0.00003)	(0.00003)
Age	-0.039^{***}	-0.036^{***}	-0.039^{***}	-0.039^{***}	-0.036^{***}
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Age squared	0.0003^{***}	0.0003^{***}	0.0003^{***}	0.0003^{***}	0.0003^{***}
	(0.00002)	(0.00002)	(0.00002)	(0.00002)	(0.00002)
Male (1=Yes)	0.059^{***}	0.064^{***}	0.059^{***}	0.060^{***}	0.064^{***}
	(0.005)	(0.006)	(0.005)	(0.006)	(0.006)
Unemployment (1=Yes)	0.014	-0.052^{*}	0.016	-0.003	-0.050^{*}
	(0.028)	(0.029)	(0.028)	(0.029)	(0.029)
Fixed effects					
Industry FE		✓		✓	✓
Year-month FE		✓			✓
Observations	259,485	259,485	259,485	259,485	259,485

Notes. This table presents the results of Probit regressions, where the dependent variable is a binary variable indicating direct stock market participation, and the independent variables include measures of income risk, and control variables such as income, wealth, age, age squared, gender, education level, and housing status. Columns (2) and (5) account for Industry and Year-month fixed effects. Robust standard errors are shown in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

the face of downside risk but do not respond to comparable upside prospects. This evidence helps explain the long-standing puzzle: standard life-cycle models predict far more equity participation by young households than we observe in practice (Haliassos and Bertaut, 1995).

Table 7: Age < 36 Subsample: Stock Market Participation

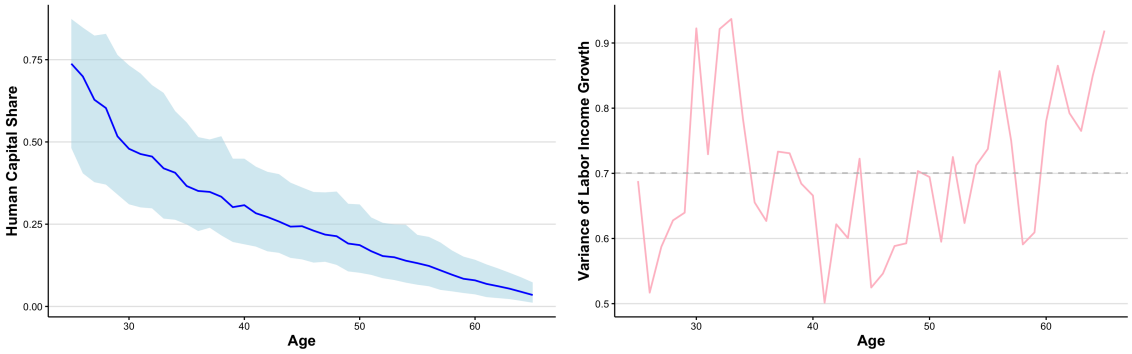
	Dependent var.: Stock Market Direct Participation (Binary: 1/0)		
	(1)	(2)	(3)
Labor-income risk			
L9050×Var	0.008**	0.007*	0.006
	(0.004)	(0.004)	(0.004)
L5010×Var	−0.017***	−0.018***	−0.017***
	(0.004)	(0.004)	(0.004)
Fixed effects			
Industry FE		✓	✓
Month–Year FE			✓
Observations	49,034	49,034	49,034

Notes. Each column reports estimates for the age < 36 subsample. Robust standard errors are shown in parentheses. Results correspond to columns (3)–(5) of Table 6 respectively. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

4.4 The Role of Risk Aversion and Human Capital Share

In the preceding subsection, I explored the overall impact of labor income skewness on portfolio decisions. However, theoretical predictions suggest that households with varying levels of risk aversion respond differently to skewness. Specifically, risk-neutral agents are indifferent to all higher moments, meaning they do not alter their behavior in response to changes in skewness levels. Agents with a higher level of risk tolerance experience a utility premium in right-skewed situations, yet they also endure utility losses in left-skewed scenarios. Conversely, risk-averse agents do not derive any benefit from positive skewness and suffer exclusively from negative skewness levels. This differentiation in behavior underscores the nuanced way in which skewness influences investment decisions, contingent upon an individual's risk tolerance.

However, the Survey of Income and Program Participation (SIPP) does not provide data regarding households' risk preferences. Consequently, I must infer their level of risk aversion using alternative variables. Given the focus on the relationship between labor income risk and the holding of financial risky assets, it seems reasonable to correlate households' risk tolerance levels with the proportion of human capital in their total asset portfolio. Specifically, the



(a) Life-cycle Human capital share (b) Life-cycle variance of labor income growth

Figure 9: The role of risk aversion and human capital share

hypothesis is that households relying more heavily on their human capital for income are likely to exhibit greater risk aversion towards human capital risk. Thus, I calculate the share of human capital following:

$$Human\,Capital\,Share_{i,M} = \frac{Human\,Capital_{i,M}}{TVAL_AST_{i,M} + Human\,Capital_{i,M}} \tag{25}$$

Here, $TVAL_AST_{i,M}$ represents the aggregate value of all assets at the individual level. Figure 9(a) illustrates the age distribution of human capital share within my sample, aligning closely with real-world observations. Specifically, it shows that younger households possess a higher proportion of human capital in their total assets, accompanied by greater variability. As individuals approach retirement age, the proportion of human capital in their asset portfolio tends to decrease, eventually approaching zero.

Effect by human capital share level. Building on this result, I partition the sample into five intervals of the human-capital share and re-estimate the full specification in equation (24) (Table 5, column 6) within each bin[11]. Skewness is captured by both the right-tail measure $L90-50$ and the left-tail measure $L50-10$. Figure 10 plots the coefficient on the interaction terms $L9050_{g(i,M)} \times \text{var}_{i,Y}$ and $L5010_{g(i,M)} \times \text{var}_{i,Y}$ for every human capital bin.

The pattern is fully consistent with the theory. For households whose earnings constitute

[11]Bins above 0.6 are omitted from the figure: the number of observations in those cells is too small for reliable inference.

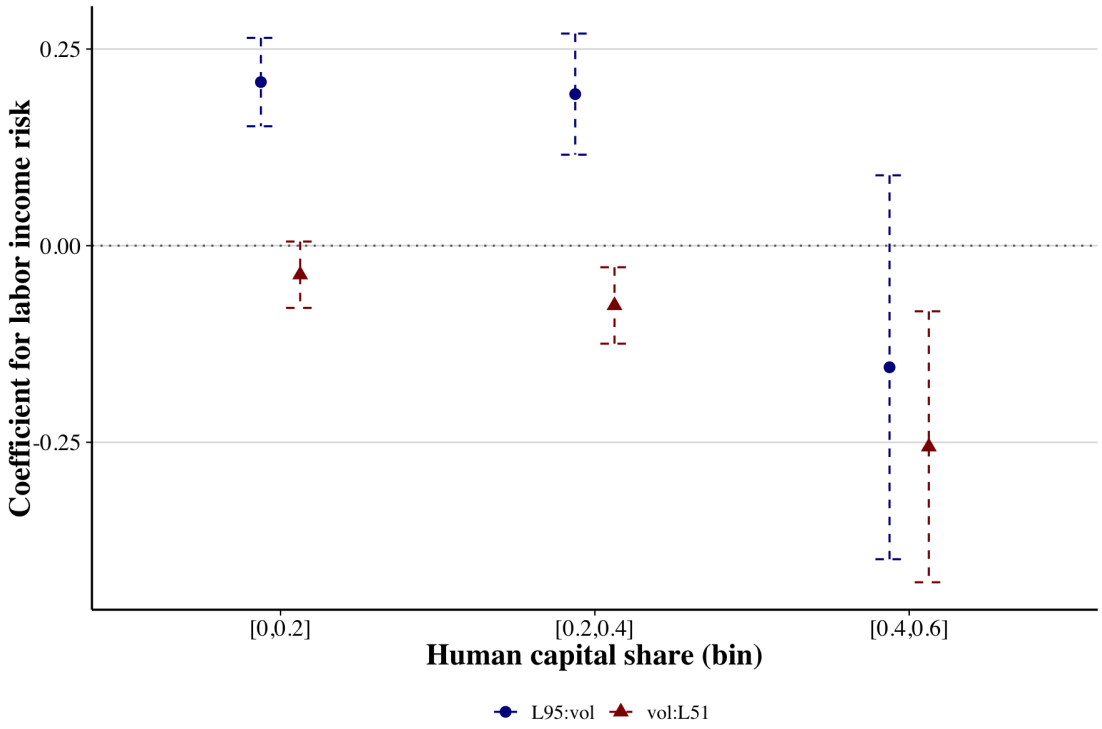


Figure 10: Effect of labor income risk by human capital share level

Notes. This figure displays the regression coefficients of the equity share on labor income risk when I run the same regressions as in column (6) of Table 5 in subsamples of workers with different levels of human capital-to-wealth ratio. Dashed lines represent 95% confidence intervals.

at most 20 percent of total wealth, opportunity risk (the right tail) raises equity holdings, whereas disaster risk (the left tail) is essentially irrelevant. Once the human-capital share reaches the 0.2-0.4 range, disaster risk starts to depress risky-asset demand. Beyond a share of 0.4, households no longer benefit from right-tail (positive) skewness and face progressively larger penalties from left-tail risk, as predicted when higher effective risk aversion dominates portfolio choice.

Effect by age. I divide the sample into three age cohorts that mirror the life-cycle pattern of wage volatility documented in Figure 9(b). Earnings risk is highest at career entry, settles to its lowest point in mid-career, and rises again as workers approach retirement. Figure 11 plots the estimated coefficients on the tail-dependent skewness–variance interaction for each cohort.

The estimates reveal a clear ranking of effective risk aversion. Households younger than

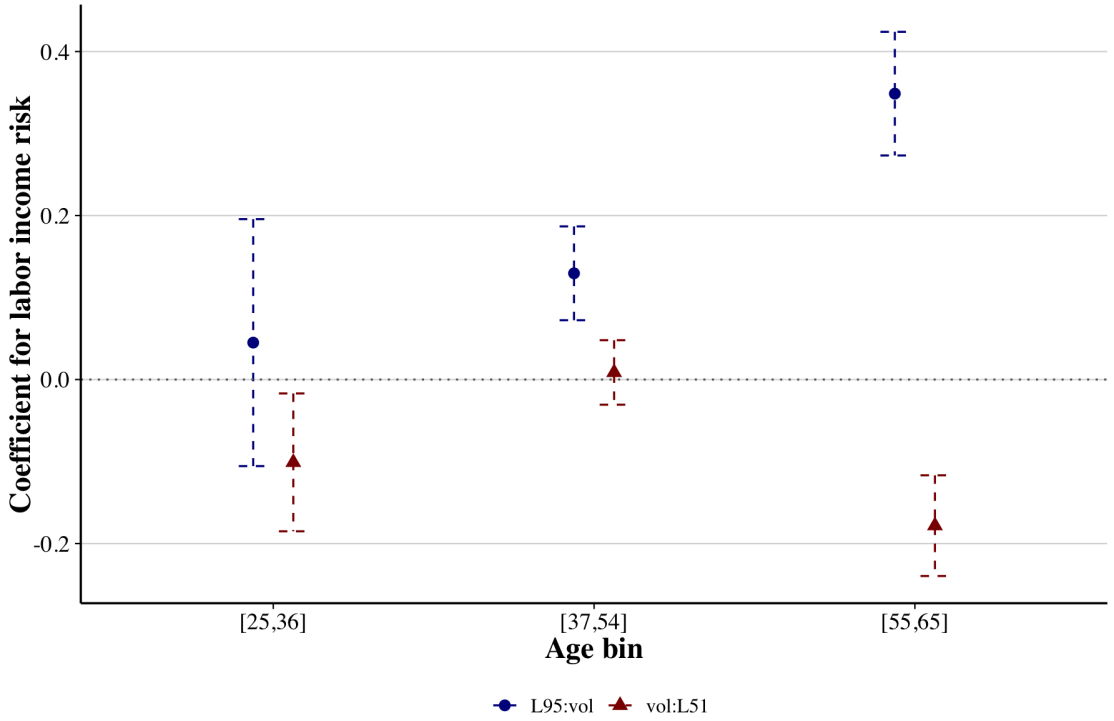
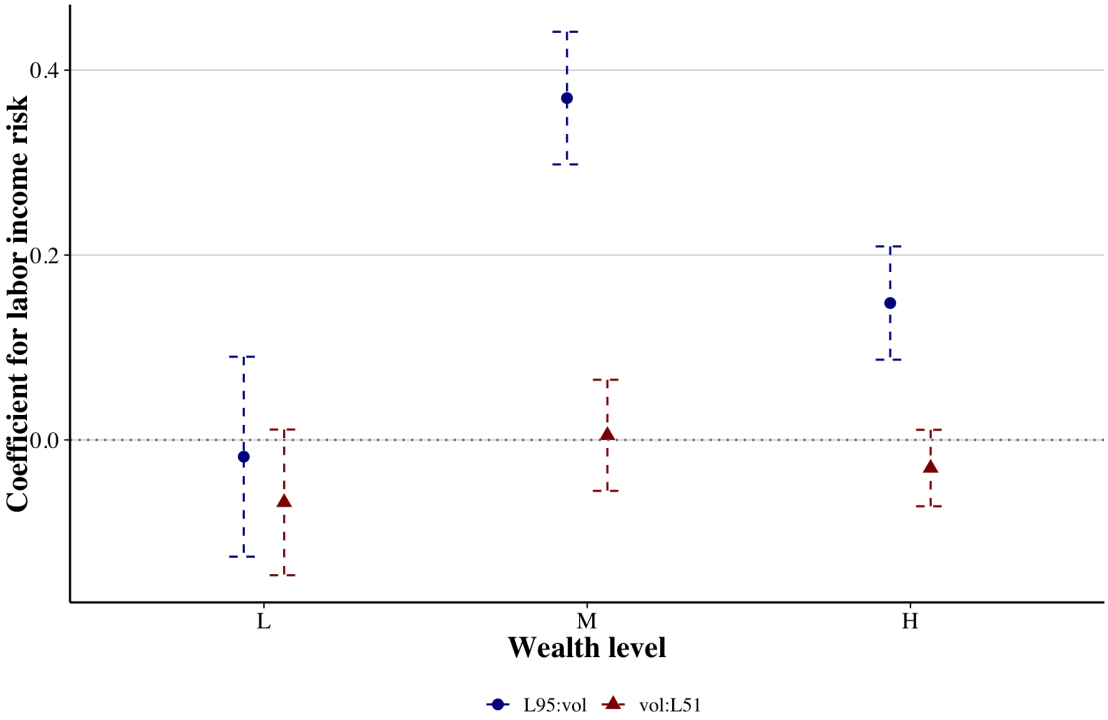


Figure 11: Effect of labor income risk by age group

Notes. This figure displays the regression coefficients of the equity share on labor income risk when I run the same regressions as in column (6) of Table 5 using different age groups of workers. Dashed lines represent 95% confidence intervals.

36 display the strongest association with precautionary behavior, those older than 54 occupy an intermediate position, and prime-age investors are the least risk-averse. The ordering is intuitive. In early adulthood, financial wealth is scant and income is dominated by human capital, so exposure to background risk sharply reduces equity demand. Near retirement, portfolios are larger, but labor income volatility rises again, pulling the coefficient toward greater caution. Prime-age households combine accumulated financial assets with the lowest earnings risk of the life cycle, and therefore tolerate the most exposure to risky assets. The findings from different age groups provide additional evidence supporting the puzzle of why younger cohorts are less likely to hold stocks than predicted by standard lifecycle models. Not only do they fail to benefit from opportunity risk, but they also suffer disproportionately from disaster risk compared to other age groups. This is consistent with return heterogeneity that can widen wealth inequality–an issue I further explore in the appendix using a model.

Figure 12: Effect of labor income risk by wealth level



Notes. This figure displays the regression coefficients of the equity share on labor income risk when I run the same regressions as in column (6) of Table 5 in subsamples of workers with different levels of total net worth. Dashed lines represent 95% confidence intervals.

Effect by wealth. Finally, I explore how households' overall financial resources shape their portfolio responses. I divide the sample into three net-worth tiers and re-estimate equation (24). Figure 12 reports the coefficients on the skewness-variance interactions for each wealth group.

The pattern is clear. Households in the bottom wealth tercile capture no benefit from upside (opportunity) risk and are already slightly exposed to downside (disaster) risk. As wealth increases, the opportunity risk coefficient turns positive and the sensitivity to disaster risk falls–but at a diminishing rate. The highest wealth group still gains from upside risk, yet their overall responsiveness is weaker than that of the middle tier. This is consistent with Prediction H3: greater wealth dampens the influence of labor income risk on equity demand.

These results provide additional evidence that using only volatility as the standard for labor income risk is insufficient because the effect of volatility varies with skewness levels.

4.5 Identification and Endogeneity

A central concern with estimating the effect of labor income risk on portfolio choice is endogeneity through occupational sorting: risk-tolerant individuals may simultaneously choose riskier occupations *and* hold more equities, generating a spurious correlation between tail-dependent income risk and the risky share. This concern has been raised forcefully by Fagereng, Guiso and Pistaferri (2018) and d'Astous and Shore (2024), who provide evidence that the endogenous component of labor income risk is substantial. Both papers recover a negative, statistically significant effect of volatility on stockholding once the endogeneity of job choice is purged—Ranish (2013) using survey-based risk-preference measures, and d'Astous and Shore (2024) exploiting quasi-random university admission cutoffs that assign students to programs with different earnings risk profiles.

This paper asks a related but distinct question. Rather than instrumenting the variance of earnings shocks, I examine whether *higher-order* moments—specifically, the shape of the tails—affect households' appetite for market risk, *conditional on* the occupation already chosen. Two features of the research design limit the scope for sorting bias even before any formal identification test. First, tail measures fluctuate over time *within* industry–education cells because of firm-specific and macroeconomic shocks that are plausibly orthogonal to any worker's initial job choice; this within-group, over-time variation is the primary source of identification in the panel specifications. Second, the key finding concerns the *interaction* between group-level tail shape and individual-level variance, not a simple level effect. For sorting to explain the interaction, one would need a very specific form of selection: among workers in the same industry–education group (who face the same tail environment), those with higher idiosyncratic variance must be disproportionately risk-tolerant *only when* the right tail of the group's earnings distribution is long. This is a much more demanding selection story than the standard concern that risk-tolerant people choose risky jobs. Nonetheless, I do not rely on these conceptual arguments alone. I implement four complementary empirical strategies to assess the threat directly.

Strategy 1: Industry–education group fixed effects. The most direct response to sorting is to absorb all time-invariant differences across industry–education cells. If risk-

tolerant workers systematically select into right-skewed occupations, the *level* of the risky share in those cells will be higher—but this level difference is removed by the group fixed effect. After including industry–education fixed effects, identification comes exclusively from within-group, over-time variation in tail shape interacted with within-group, across-individual variation in idiosyncratic volatility.

Table 13 (column 4) reports this specification on the annual panel with standard errors clustered by individual. The opportunity-risk interaction is virtually unchanged when industry–education group fixed effects replace industry fixed effects ($\hat{\beta}_4 = 0.313$, SE $= 0.126$, $p < 0.05$ in both specifications), and the disaster-risk interaction likewise remains stable ($\hat{\beta}_5 = -0.10$ with industry FE versus -0.09 with IND_EDU FE, both significant at the 10% level). The near-identical coefficients imply that, holding constant the average portfolio allocation of each industry–education cell, workers whose idiosyncratic volatility is higher respond more strongly to the tail shape of their group—exactly the pattern predicted by the theory. Absorbing all time-invariant group-level sorting has essentially no effect on the estimates.

Strategy 2: Industry movers with individual fixed effects. A more demanding test exploits workers who change industry–education group during the panel. For these movers, the individual fixed effect absorbs all time-invariant preferences—risk aversion, financial literacy, trust in markets—and identification comes solely from the change in group-level tail risk that accompanies the job transition. Crucially, the worker's underlying risk preferences do not change when her employer does.

I classify individuals as movers if they belong to more than one industry–education group during the sample period (1,639 movers versus 9,912 stayers). Table 8 reports the results. In the full sample with individual fixed effects, the interaction terms are imprecisely estimated ($\hat{\beta}_4 = 0.025$, $\hat{\beta}_5 = -0.020$, both insignificant). This is expected: the panel is short (at most four years of monthly data), so most individuals exhibit little within-person variation in tail-risk exposure once fixed effects are included. However, focusing on movers—where there *is* meaningful within-person variation in tail risk—reveals a different picture. In the mover subsample, the opportunity-risk interaction is positive and marginally significant ($\hat{\beta}_4 = 0.064$,

$p < 0.10$).

The strongest evidence comes from the subset of movers who also experienced an unemployment spell during the panel—a plausibly involuntary separation. These 266 workers did not choose their new risk environment; they were pushed into it. For this subsample, the opportunity-risk interaction is large and highly significant ($\hat{\beta}_4 = 0.151$, SE $= 0.057$, $p < 0.01$). This result is arguably the single most compelling piece of evidence against the sorting interpretation: among workers whose industry transition was plausibly exogenous, the within-person response to the new group's right-tail risk is economically large and precisely estimated.

The disaster-risk interaction does not survive in the mover subsamples, but this is unsurprising given the small number of involuntary movers ($N = 266$ individuals, 7,698 person–months). The downside-risk channel may operate more through level effects than through the variance interaction, making it harder to detect with limited within-person variation.

Table 8: Individual Fixed Effects by Mover Status

	Dependent variable: Risky Share (×100)			
	(1) All	(2) Movers	(3) Stayers	(4) Involuntary
Opportunity ($L9050$)	0.127*	0.579***	−0.166*	−1.260***
	(0.076)	(0.132)	(0.094)	(0.364)
Disaster ($L5010$)	−0.447***	−0.612***	−0.489***	−1.365***
	(0.065)	(0.120)	(0.077)	(0.352)
Idiosyncratic variance	0.011	−0.093	0.049	−0.065
	(0.037)	(0.075)	(0.042)	(0.103)
Opportunity × variance	0.025	0.064*	0.005	0.151***
	(0.019)	(0.035)	(0.022)	(0.057)
Disaster × variance	−0.020	0.015	−0.034**	0.010
	(0.013)	(0.029)	(0.015)	(0.032)
Household controls	Yes	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes	Yes
Year–month FE	Yes	Yes	Yes	Yes
Observations	251,692	47,238	204,454	7,698
R^2	0.674	0.537	0.708	0.526
Within R^2	0.036	0.053	0.030	0.069

Notes. All columns include individual and year–month fixed effects and household controls (log income, log wealth, age, age^2, unemployment status). "Movers" are individuals observed in more than one industry–education group during the panel ($N_{\text{individuals}} = 1{,}639$). "Stayers" remain in a single group throughout ($N_{\text{individuals}} = 9{,}912$). "Involuntary" restricts the mover sample to individuals who experienced at least one unemployment spell ($N_{\text{individuals}} = 266$). Heteroskedasticity-robust standard errors in parentheses. ***/**/* denote significance at the 1%, 5%, and 10% levels.

Strategy 3: Falsification test. If sorting were driving the results, one would expect that households with a high pre-existing risky share—those who are presumably more risk-tolerant—would be more likely to end up in industry–education groups with favorable tail characteristics. I test this directly by regressing the group-level tail measure observed at the end of the panel on the individual's *initial* risky share, controlling for the initial industry–education group and demographics (education, age, gender, volatility, log income, and log wealth).

Under the sorting hypothesis, the initial risky share should significantly predict the worker's future tail-risk environment. I find that it does not. For the opportunity tail ($L9050$), the coefficient on initial risky share is 0.0002 ($t = 1.18$, $p = 0.24$). For the disaster tail ($L5010$), the coefficient is 0.0002 ($t = 0.93$, $p = 0.35$). Neither estimate is close to conventional significance levels. These placebo results indicate that, conditional on the initial group assignment, the pre-existing risky share has no predictive power for the individual's future tail-risk environment—inconsistent with the view that risk-tolerant workers differentially sort into right-skewed groups over time.

Strategy 4: Coefficient stability. Finally, I assess how large the omitted-variable bias from sorting would need to be to explain away the results. Table 9 reports the opportunity-risk coefficient ($\hat{\beta}_4$) and the disaster-risk coefficient ($\hat{\beta}_5$) as I progressively add controls and fixed effects.

The results are striking. For the opportunity-risk interaction, $\hat{\beta}_4$ *increases* monotonically as controls are added: from 0.127 with no controls, to 0.140 with individual covariates, to 0.147 with industry and date fixed effects, to 0.149 with industry–education and date fixed effects. The coefficient moves *away* from zero as potential confounders are absorbed, implying that the observable controls are negatively correlated with the treatment effect. Under the proportional-selection framework of Oster (2019), this means that no amount of proportional selection on unobservables could drive the coefficient to zero—the implied δ is negative ($\delta = -2.28$), indicating that unobservables would need to work in the *opposite* direction from observables to eliminate the result. This is the strongest possible outcome from a coefficient-stability analysis.

Table 9: Coefficient Stability Across Specifications

	(1) None	(2) + Controls	(3) + IND + Date	(4) + IND_EDU + Date	(5) + ID + Date
Opportunity × var ($\hat{\beta}_4$)	0.127***	0.140***	0.147***	0.149***	0.025
	(0.033)	(0.033)	(0.032)	(0.032)	(0.019)
Disaster × var ($\hat{\beta}_5$)	−0.027*	−0.029**	−0.023	−0.019	−0.020
	(0.015)	(0.015)	(0.014)	(0.014)	(0.013)
Individual controls		Yes	Yes	Yes	Yes
Industry FE			Yes		
IND_EDU FE				Yes	
Individual FE					Yes
Year–month FE			Yes	Yes	Yes
R^2	0.001	0.017	0.023	0.027	0.674
Observations	251,692	251,692	251,692	251,692	251,692

Notes. Each column progressively adds controls and fixed effects. All specifications include both tail measures, variance, and the full set of interactions. Heteroskedasticity-robust standard errors in parentheses. ***/**/* denote significance at the 1%, 5%, and 10% levels.

The pattern for $\hat{\beta}_5$ (disaster risk) is less dramatic but still supportive. The coefficient is -0.027 without controls and -0.023 with full fixed effects—a modest attenuation. The implied $\delta = 2.24$, meaning unobservables would need to be more than twice as important as all observables combined to explain away the disaster-risk effect. Both coefficients lose significance with individual fixed effects in column (5), as expected: the short panel provides limited within-person variation, and the individual fixed effect absorbs the cross-sectional variation that provides most of the identifying power.

As a complementary check, I include the predicted group-level tail measure (constructed from pre-determined demographics) as an explicit control variable. The interaction terms are essentially unchanged ($\hat{\beta}_4 = 0.147$, $p < 0.01$; $\hat{\beta}_5 = -0.030$, $p < 0.05$), confirming that the selection component of group assignment does not drive the results.[12]

Discussion. Taken together, the four strategies paint a consistent picture. The IND_EDU fixed effects absorb time-invariant group-level sorting with negligible impact on the estimates (Strategy 1). The involuntary-mover analysis—the most demanding test—shows that workers who are pushed into new industry–education groups by unemployment respond strongly to their new group's right-tail risk ($\hat{\beta}_4 = 0.151$, $p < 0.01$), ruling out preference-based sorting

[12]The first-stage R^2 of the demographic prediction model is only 0.005, indicating that pre-determined characteristics explain very little of the variation in group-level tail risk. This is consistent with the view that most of the variation in tail measures arises from within-group, over-time shocks rather than from sorting.

as the explanation (Strategy 2). The falsification test confirms that the pre-existing risky share does not predict future tail-risk exposure (Strategy 3). And the coefficient-stability analysis shows that $\hat{\beta}_4$ actually *increases* as controls are progressively added, implying that the omitted-variable bias from sorting works against the finding rather than in its favor (Strategy 4).

These empirical findings clarify the relationship between this paper and the instrumental-variables literature on occupational sorting. Prior work by Ranish (2013), Fagereng, Guiso and Pistaferri (2018), and d'Astous and Shore (2024) illuminates *who* selects into risky occupations, showing that risk-tolerant workers sort into jobs with volatile pay and, for the same reason, hold more equities. This paper reveals a complementary mechanism: *how* the shape of labor income shocks feeds into portfolio choice once the occupation is given. The two perspectives are not in conflict. Sorting likely contributes to the unconditional positive correlation between earnings volatility and equity holdings documented in Figure 7. But the tail-dependent interaction effects—the core contribution of this paper—survive controls for sorting and reflect genuine household behavior. Workers appear to select into jobs where high volatility comes bundled with right-tail (opportunity) income risk, and it is this tail risk, not variance alone, that lifts equity demand. Disentangling these two margins is precisely why decomposing income risk into its upside and downside components adds explanatory power beyond what a single variance measure can deliver.

5 Conclusion

This paper provides the first direct, reduced-form evidence that non-Gaussian features of labor income risk–specifically, the interplay between tail risks and variance–are substantially associated with household portfolio choice. Whereas earlier work has shown, within calibrated life-cycle models, that countercyclical earnings risk can reconcile macro-micro puzzles, it has not established a micro-level link between the tails of income shocks and equity demand.

Exploiting SIPP data, I construct a tail-dependent risk measure by interacting each individual's wage growth volatility with the tails of income shocks. This approach captures

how households perceive their idiosyncratic uncertainty relative to the distributional tail risks faced by peers. Three key findings emerge:

1. **Opposing Tail Effects.** I first document that the main effects of tail align with theory: the right tail (opportunity) does not have a uniform impact across investors, whereas the left tail (disaster) depresses risky holdings broadly across the distribution. Conditional on the same volatility, right-skewed income shocks encourage greater equity allocation, whereas left-skewed shocks induce households to withdraw from stocks. A variance-only metric obscures these countervailing forces, and a single skewness measure likewise masks the asymmetric effects of the left and right tails.

2. **Distributional Drivers.** The upside potential from opportunity risk accrues almost entirely to high-wealth households, who benefit from right-skewed income shocks. In contrast, less wealthy households–those most vulnerable to large negative shocks–respond in exactly the opposite way.

3. **Downside Aversion and Life-Cycle Patterns.** Consistent with CRRA-derived models of downside risk aversion, households whose human capital constitutes a large share of total wealth (e.g., the young and financially constrained) are particularly penalized by left-skewed risks and therefore least likely to hold equities. This mechanism helps explain both low participation rates among the young and the hump-shaped life-cycle profile of equity holdings.

These results deliver three important implications. First, they lend micro-level validation to portfolio choice frameworks that incorporate higher-order background risk. Second, they underscore the limitations of variance-only measures when quantifying labor income uncertainty. Third, they suggest how labor income risk could contribute to return heterogeneity and, over time, may help explain the wealth gap between richer and poorer households.

A natural next step is to embed these reduced-form findings within a structural life-cycle model that nests skewness-variance interactions. Such an extension would permit quantitative assessment of how non-Gaussian income shocks shape long-run wealth accumulation and portfolio dynamics–and ultimately inform policy debates on retirement security and inequality.

References

Acemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. 2017. "Microeconomic origins of macroeconomic tail risks." *American Economic Review*, 107(1): 54–108.

Ait-Sahalia, Yacine, Felix Matthys, Emilio Osambela, and Ronnie Sircar. 2021. "When Uncertainty and Volatility Are Disconnected: Implications for Asset Pricing and Portfolio Performance." *Available at SSRN 3841103*.

Aiyagari, S Rao. 1994. "Uninsured idiosyncratic risk and aggregate saving." *The Quarterly Journal of Economics*, 109(3): 659–684.

Angerer, Xiaohong, and Pok-Sang Lam. 2009. "Income risk and portfolio choice: An empirical study." *The Journal of Finance*, 64(2): 1037–1055.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub. 2023. "The Trickling Up of Excess Savings." CEPR Discussion Papers Report.

Bender, Svetlana, James J Choi, Danielle Dyson, and Adriana Z Robertson. 2022. "Millionaires speak: What drives their personal investment decisions?" *Journal of Financial Economics*, 146(1): 305–330.

Benzoni, Luca, Pierre Collin-Dufresne, and Robert S Goldstein. 2007. "Portfolio choice over the life-cycle when the stock and labor markets are cointegrated." *The Journal of Finance*, 62(5): 2123–2167.

Betermier, Sebastien, Thomas Jansson, Christine Parlour, and Johan Walden. 2012. "Hedging labor income risk." *Journal of Financial Economics*, 105(3): 622–639.

Bonaparte, Yosef, George M Korniotis, and Alok Kumar. 2014. "Income hedging and portfolio decisions." *Journal of Financial Economics*, 113(2): 300–324.

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(2): 207–242.

Calvet, Laurent E, John Y Campbell, Francisco Gomes, and Paolo Sodini. 2021. "The cross-section of household preferences." National Bureau of Economic Research Report.

Campbell, John Y, and Luis M Viceira. 2001. *Strategic asset allocation: portfolio choice for long-term investors.* Clarendon Lectures in Economic.

Catherine, Sylvain. 2022. "Countercyclical labor income risk and portfolio choices over the life cycle." *The Review of Financial Studies*, 35(9): 4016–4054.

Catherine, Sylvain, Max Miller, James D Paron, and Natasha Sarin. 2022. "Who hedges interest-rate risk? Implications for wealth inequality." *Implications for Wealth Inequality (May 23, 2022). Jacobs Levy Equity Management Center for Quantitative Financial Research Paper.*

Catherine, Sylvain, Paolo Sodini, and Yapei Zhang. 2024. "Countercyclical income risk and portfolio choices: Evidence from Sweden." *The Journal of Finance*, 79(3): 1755–1788.

Chang, Yongsung, Jay H Hong, and Marios Karabarbounis. 2018. "Labor market uncertainty and portfolio choice puzzles." *American Economic Journal: Macroeconomics*, 10(2): 222–262.

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

44: 65–90.

Chetty, Raj, László Sándor, and Adam Szeidl. 2017. "The effect of housing on portfolio choice." *The Journal of Finance*, 72(3): 1171–1212.

Choukhmane, Taha, and Tim de Silva. 2022. "What drives investors portfolio choices? separating risk preferences from frictions." Working paper Report.

Cocco, Joao F, Francisco J Gomes, and Pascal J Maenhout. 2005. "Consumption and portfolio choice over the life cycle." *The Review of Financial Studies*, 18(2): 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.

Egan, Mark L, Alexander MacKay, and Hanbin Yang. 2021. "What drives variation in investor portfolios? estimating the roles of beliefs and risk preferences." National Bureau of Economic Research Report.

Fagereng, Andreas, Luigi Guiso, and Luigi Pistaferri. 2018. "Portfolio choices, firm shocks, and uninsurable wage risk." *The Review of Economic Studies*, 85(1): 437–474.

Fagereng, Andreas, Martin Blomhoff Holm, Benjamin Moll, and Gisle Natvik. 2019. "Saving behavior across the wealth distribution: The importance of capital gains." National Bureau of Economic Research Report.

Favilukis, Jack. 2013. "Inequality, stock market participation, and the equity premium." *Journal of Financial Economics*, 107(3): 740–759.

Friedman, Milton, and Leonard J Savage. 1948. "The Utility Analysis of Choices Involving Risk." *Journal of Political Economy*, 56(4): 279–304.

Gabaix, Xavier. 2012. "Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance." *The Quarterly journal of economics*, 127(2): 645–700.

Ganong, Peter, Pascal J Noel, Christina Patterson, Joe Vavra, and Alex Weinberg. 2024. "Earnings Instability." *Working Paper*.

Ghysels, Eric, Alberto Plazzi, and Rossen Valkanov. 2016. "Why invest in emerging markets? The role of conditional return asymmetry." *The Journal of Finance*, 71(5): 2145–2192.

Gomes, Francisco, and Oksana Smirnova. 2021. "Stock market participation and portfolio shares over the life-cycle." *Available at SSRN 3808350*.

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

Guvenen, Fatih, Serdar Ozkan, and Jae Song. 2014. "The nature of countercyclical income risk." *Journal of Political Economy*, 122(3): 621–660.

Haliassos, Michael, and Carol C Bertaut. 1995. "Why do so few hold stocks?" *The Economic Journal*, 105(432): 1110–1129.

Hammitt, James K. 2023. "Downside risk aversion vs decreasing absolute risk aversion: an intuitive exposition." *Theory and Decision*, 95(1): 1–10.

Kimball, Miles S. 1990. "Precautionary saving and the marginal propensity to consume."

Krusell, Per, and Anthony A Smith, Jr. 1998. "Income and wealth heterogeneity in the macroeconomy." *Journal of Political Economy*, 106(5): 867–896.

Lambert, Susan J, Julia R Henly, and Jaeseung Kim. 2019. "Precarious work schedules as a source of economic insecurity and institutional distrust." *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 5(4): 218–257.

Lucas, Robert E. 1978. "Asset prices in an exchange economy." *Econometrica*, 1429–1445.
Merton, Robert C. 1969. "Lifetime portfolio selection under uncertainty: The continuous-time case." *The Review of Economics and Statistics*, 247–257.
Mian, Atif R, Ludwig Straub, and Amir Sufi. 2020. "The saving glut of the rich." National Bureau of Economic Research Report.
Oster, Emily. 2019. "Unobservable selection and coefficient stability: Theory and evidence." *Journal of Business & Economic Statistics*, 37(2): 187–204.
Ranish, Benjamin. 2013. "Why Do Households with Risky Labor Income Take Greater Financial Risks?" *Available at SSRN 2681360*.
Samuelson, Paul A. 1969. "Lifetime portfolio selection by dynamic stochastic programming." *The Review of Economics and Statistics*, 239–246.
Schmidt, Lawrence. 2022. "Climbing and falling off the ladder: Asset pricing implications of labor market event risk." *Available at SSRN 2471342*.
Straub, Ludwig. 2019. "Consumption, savings, and the distribution of permanent income." *Unpublished manuscript, Harvard University*.
Vissing-Jorgensen, Annette. 2002. "Towards an explanation of household portfolio choice heterogeneity: Nonfinancial income and participation cost structures."
Wachter, Jessica A. 2013. "Can time-varying risk of rare disasters explain aggregate stock market volatility?" *The Journal of Finance*, 68(3): 987–1035.
Wachter, Jessica A, and Motohiro Yogo. 2010. "Why do household portfolio shares rise in wealth?" *The Review of Financial Studies*, 23(11): 3929–3965.
Weil, Philippe. 1992. "Equilibrium asset prices with undiversifiable labor income risk." *Journal of Economic Dynamics and Control*, 16(3-4): 769–790.

Appendix A. Theoretical Proofs

A1. Derivation of the Utility Premium

To expand on my motivation, I introduce the concept of downside risk aversion (downside RA) (Hammitt, 2023). In the context of this paper, individuals exhibiting downside RA incur higher costs from bearing risks when they possess lower levels of wealth.

Connecting CRRA Preferences to Downside Risk Aversion Under expected utility, downside RA theories suggest that individuals with greater wealth tend to be less vulnerable to risks. This notion aligns with the empirical evidence presented (See panel (b) in Figure 6). While CRRA implies DARA, and DARA implies downside RA, the converse is not necessarily true. Therefore:

$$\text{CRRA} \Rightarrow \text{DARA} \Rightarrow \text{Downside RA}$$

Transitioning to downside RA offers greater flexibility, but Constant Relative Risk Aversion (CRRA) is simpler and more cost-effective, even though it involves stricter assumptions. To illustrate the difference between downside RA and CRRA, it is important to note that for CRRA, it must satisfy the condition where $r'(W) < 0$, such that the absolute risk aversion is a decreasing function of W, while the relationship is as follows:

$$r'(W) = r(W)[r(W) - \rho(W)] < 0 \Rightarrow r(W) < \rho(W)$$

Here, $\rho(W) = -\frac{u'''(W)}{u''(W)}$ can be interpreted as the measure of absolute prudence (Kimball, 1990). Therefore, a utility function can demonstrate downside RA without adhering to the characteristics of CRRA[13].

Cost of Risk Bearing Let u(W) denote the CRRA utility function for a worker, where $u' > 0$, $u'' < 0$, and u''' exists. For simplicity, consider a lottery L that yields wealth level

[13]For instance, the Constant Absolute Risk Aversion (CARA) utility demonstrates $u''' > 0$ while $\rho(W) = r(W)$.

W_l and W_h ($W_l < W_h$) with equal probability. This can be envisioned as a worker facing the uncertainty of the future with only two potential outcomes in the next time period: a higher level of human capital wealth and a relatively lower human capital wealth. Assume that the risk $\tilde{y}$ is sufficiently small to allow adequate approximation of expected utility using a higher-order Taylor expansions.

Definition of the Utility Premium If the individual faces downside RA (e.g., Individual with CRRA utility), by definition, he exhibits $u''' > 0$. His expected utility is given by $Eu(L) = \frac{1}{2}u(W_l) + \frac{1}{2}u(W_h)$. If he adds the risk $\tilde{y}$ to the outcome W, his expected utility changes by $\theta(W) = u(W) - Eu(W + \tilde{y})$. $\theta(W)$ is defined as the **utility premium** for the risk $\tilde{y}$ when added to W (Friedman and Savage, 1948). In the original paper, a second-order Taylor series approximation is applied, where:

$$\theta(W) \approx -\frac{\sigma^2}{2}u''(W) \tag{26}$$

where σ^2 is the variance of $\tilde{y}$. Note that in this equation there must have $\theta(W) > 0$, because $u''(W) < 0$ and $\sigma^2 > 0$. This intuition underlies the view that greater return volatility should curb demand for risky assets: as volatility rises, so does the cost of bearing risk.

However, I extend the analysis by using the third-order Taylor series approximation[14]:

$$\theta(W) \approx -\frac{\sigma^2}{2}u''(W) - \frac{E[\tilde{y}^3]}{6}u'''(W) \tag{27}$$

In contrast to the original paper, the determination of $\theta(W)$ as positive or negative in this study depends on the magnitude of the third moment of risk, represented by $E[\tilde{y}^3]$. Importantly, this approximation may not perform well in certain scenarios, particularly when the shocks ($\tilde{y}$) are substantial or the utility function is highly nonlinear.

Definition of the Risk Premium On the other hand, the **risk premium** for $\tilde{y}$ is a monetary measure of the harm from bearing risk. It is defined as the difference between the expected value and the certain equivalent, such that $\pi(W) = W - \hat{W} > 0$, where the certain

[14] A detailed proof of this approximation is presented in the Appendix A3.

equivalent is defined by $u(\hat{W}) = Eu(W + \tilde{y})$. Similarly, using the third-order Taylor series approximation:

$$\pi(W) \approx -\frac{\sigma^2}{2}\frac{u''(W)}{u'(W)} - \frac{E[\tilde{y}^3]}{6}\frac{u'''(W)}{u'(W)} \tag{28}$$

Within a Portfolio Choice Framework Again, for simplicity, I use the model of Campbell and Viceira (2001) for demonstration. Recall that being CRRA implies downside RA. Thus, the expressions for $\theta(C)$ and $\pi(C)$ are as follows:

$$\theta(C) \approx \frac{\sigma^2}{2}\frac{\gamma}{C^{\gamma+1}} - \frac{E[\tilde{y}^3]}{6}\frac{\gamma^2+\gamma}{C^{\gamma+2}}$$

$$\pi(C) \approx \frac{\sigma^2}{2}\frac{\gamma}{C} - \frac{E[\tilde{y}^3]}{6}\frac{\gamma^2+\gamma}{C^2}$$

$$\theta(C) = \frac{\pi(C)}{C^{\gamma}} \tag{29}$$

Using the utility premium as a measure of risk bearing magnifies the wealth effect as risk aversion heightened. This implies that individuals with higher levels of wealth will experience a much lower magnitude of utility premium compared to those with lower wealth levels. This aspect better explains investors' behavior as depicted in this paper. For households with higher levels of wealth, the portfolio effect of background risk is significantly lower because their cost of risk bearing is much reduced.

A2. Derivation of Campbell–Viceira (2001) Portfolio Rule

1. Setup. Let $R_s \equiv \mu + \sigma_s\varepsilon_s$ denote the gross return on the risky asset ($E[R_s] = \mu$, $Var(R_s) = \sigma_s^2$), and let the **gross return on human capital** be $R_H \equiv \mu_H + \sigma_H\varepsilon_H$ with $Cov(\varepsilon_s, \varepsilon_H) = \rho$.

The household begins period t with $W \equiv W_t$ in financial wealth and $H \equiv H_t$ in the (tradable) present value of future labor income. It allocates the **dollar amount** α ($= \pi W$) to the risky asset and $(W - \alpha)$ to the risk-free asset (R^f).

2. One-period portfolio return. Total end-of-period wealth (financial + human) is

$$W' = (W - \alpha)R^f + \alpha R_s, \qquad H' = HR_H,$$

so total **portfolio return** on "comprehensive" (financial + human) wealth $W_T \equiv W + H$ is

$$R_p \;=\; \frac{W'}{W_T} + \frac{H'}{W_T} \;=\; R^f + \frac{\alpha}{W_T}\,(R_s - R^f) + \frac{H}{W_T}\,(R_H - R^f).$$

3. Quadratic (mean–variance) objective. Using a second-order approximation to CRRA utility (Campbell–Viceira), the investor maximises

$$\max_{\alpha} \Big\{ E[R_p] - \frac{\gamma}{2} Var(R_p) \Big\}.$$

4. Compute $E[R_p]$ and $Var(R_p)$ (conditional on current information).

$$E[R_p] = R^f + \frac{\alpha}{W_T}(\mu - R^f) + \frac{H}{W_T}(\mu_H - R^f), \tag{30}$$

$$Var(R_p) = \frac{\alpha^2}{W_T^2}\sigma_s^2 + \frac{H^2}{W_T^2}\sigma_H^2 + 2\frac{\alpha H}{W_T^2}\sigma_s\sigma_H\rho. \tag{31}$$

5. First-order condition (FOC). Differentiate the quadratic objective with respect to α:

$$0 = \frac{\partial}{\partial\alpha}\Big[E[R_p] - \frac{\gamma}{2}Var(R_p)\Big] = \frac{\mu - R^f}{W_T} - \gamma\Big(\frac{\alpha}{W_T^2}\sigma_s^2 + \frac{H}{W_T^2}\sigma_s\sigma_H\rho\Big).$$

Multiply through by W_T^2 and solve for α:

$$\alpha = \frac{\mu - R^f}{\gamma\sigma_s^2}\,W_T - \frac{H\sigma_H\rho}{\sigma_s}\,.$$

6. Express the covariance term via β_H. Define the regression (CAPM-style) beta of human-capital returns on stock returns

$$\beta_H \;\equiv\; \frac{Cov(R_H, R_s)}{Var(R_s)} \;=\; \rho\,\frac{\sigma_H}{\sigma_s}.$$

Hence $H\sigma_H\rho = \beta_H\, H\,\sigma_s^2$. Substitute back:

$$\alpha = \frac{\mu - R^f}{\gamma \sigma_s^2} W_T \; - \; \beta_H \, H.$$

7. Separate financial (W) and human (H) wealth. Since $W_T = W + H$,

$$\begin{aligned} \alpha &= \frac{\mu - R^f}{\gamma \sigma_s^2} (W + H) - \beta_H H \\ &= \frac{\mu - R^f}{\gamma \sigma_s^2} W \; + \; \Big(\frac{\mu - R^f}{\gamma \sigma_s^2} - \beta_H \Big) H. \end{aligned}$$

Set $R^f \equiv r$ for notational consistency. Finally, recalling $\alpha = \pi W$, we obtain

$$\boxed{\alpha = \frac{\mu - r}{\gamma \sigma_s^2} W \; + \; \Big(\tfrac{\mu - r}{\gamma \sigma_s^2} - \beta_H \Big) H}.$$

This is precisely equation (3.15) in Campbell & Viceira (2001) and (4.4) in their 2002 monograph, showing step-by-step how the portfolio demand decomposes into:

- **Myopic demand** $(\mu - r)/(\gamma \sigma_s^2) \times W$;
- **Hedging (human capital) demand** $\big[(\mu - r)/(\gamma \sigma_s^2) - \beta_H\big] \times H$.

A3. Proof of Third-order Taylor Series Approximation

The proof of Equation 27 step by step is as follows:

$$u(W + \tilde{y}) \approx u(W) + u'(W)\tilde{y} + \frac{u''(W)}{2}\tilde{y}^2 + \frac{u'''(W)}{6}\tilde{y}^3$$

$$E[u(W + \tilde{y})] \approx u(W) + E[u'(W)\tilde{y}] + \frac{E[u''(W)\tilde{y}^2]}{2} + \frac{E[u'''(W)\tilde{y}^3]}{6}$$

Since $E[\tilde{y}^2] = \sigma^2$ when $E[\tilde{y}] = 0$:

$$E[u(W + \tilde{y})] \approx u(W) + E[u'(W)\tilde{y}] + \frac{\sigma^2 u''(W)}{2} + \frac{E[\tilde{y}^3]u'''(W)}{6}$$

$$\theta(W) = u(W) - Eu(W + \tilde{y}) \approx -E[u'(W)\tilde{y}] - \frac{\sigma^2 u''(W)}{2} - \frac{E[\tilde{y}^3]u'''(W)}{6}$$

Under the assumption that $\tilde{y}$ is independent of $u'(W)$, we have $E[u'(W)\tilde{y}] = u'(W)E[\tilde{y}]$. For the case when $E[\tilde{y}] = 0$, Equation 27 is proven (Q.E.D.). If $\tilde{y}$ represents an unfair risk, Equation 27 can be further extended as follows:

$$\theta(W) \approx -E[\tilde{y}]u'(W) - \frac{\sigma^2 u''(W)}{2} - \frac{E[\tilde{y}^3]u'''(W)}{6}$$

Appendix B. Data and Measurement

B1. Definitions of Key Variables

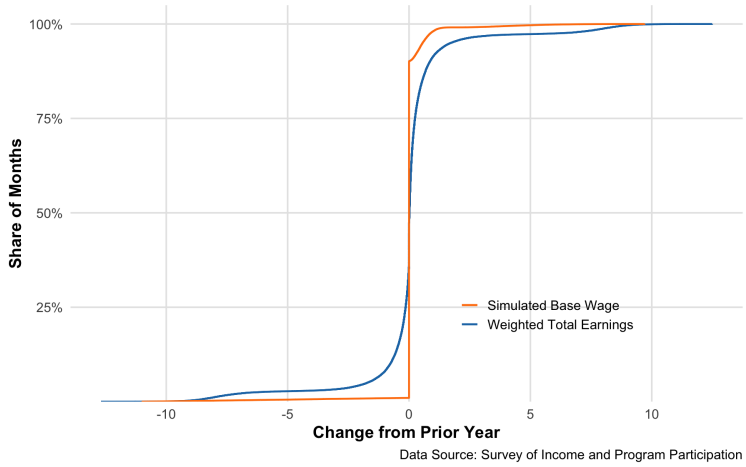
Table 10: Definitions of key variables

Variable	Definition
Monthly earnings from job (TJB1_MSUM)	The dollar value of the monthly earnings from job 1, varying with the number of days in the month.
Total asset value (TVAL_AST)	The total dollar value of an individual's assets, constructed as the sum of: assets held at financial institutions (checking and savings), direct stock and mutual-fund holdings, bond holdings, value of rental properties, value of other real estate, other financial assets, retirement account balances (e.g., 401(k) and IRAs), business equity, value of primary residence, vehicle net equity, and educational savings accounts in which the person is the owner.
Total financial asset value (TVAL_FAST)	The total dollar value of an individual's assets, constructed as the sum of: assets held at financial institutions (checking and savings), direct stock and mutual-fund holdings, bond holdings, other financial assets, retirement account balances (e.g., 401(k) and IRAs), and educational savings accounts in which the person is the owner.
Risky share (RS)	The share of directly held stocks and mutual-fund investments in total financial assets.

B2. Earning Instability in the U.S.

Given that this paper uses monthly labor income to study how households perceive income risk, it is important to highlight that workers in the U.S. face substantial fluctuations in earnings even on a monthly basis. The observed earnings instability among the U.S. labor force is much higher than one might expect. Since I do not have access to paycheck data, I reference Ganong et al. (2024), who also study earnings instability in the U.S. using payroll company data spanning the years 2010 to 2023 (while my study spans the years 2013 to 2022). They estimate the monthly change in labor income and show that although base

Figure 13: Earning Instability



Notes. This figure shows the cumulative distribution function (CDF) of labor income growth (earnings change) estimated from my sample. The blue line represents the weighted total earnings, and the orange line is a simulated base wage designed to replicate the results from Ganong et al. (2024).

wages exhibit some downward rigidity and are quite stable, total earnings are not. Furthermore, considering total earnings is important because, in the U.S. workforce, 60% are hourly workers and 40% are salaried[1]. For hourly employees, earnings volatility chiefly reflects changes in hours worked. By contrast, fluctuations in salaried earnings stem from variable components–performance pay, bonuses, and other items such as reimbursements–that cause month-to-month variation.

Figure 13, derived from the Survey of Income and Program Participation (SIPP), demonstrates that workers experience significant earnings fluctuations at the monthly level (note that I use year-over-year changes to reduce seasonality, while Ganong et al. (2024) use monthly differences). As I do not have base wage data, I simulate a similar cumulative distribution function (CDF) as in Ganong et al. (2024) to compare with the CDF I estimated from SIPP. It is obvious that unlike base wage (or hourly wage), total earnings can be very volatile even at monthly level, both upward and downward. Suppose an individual suffers a monthly income loss early in the year and later offsets it by taking extra shifts or part-time work. In annual data, earnings appear unchanged, masking a substantial intra-year income shock. Therefore, it is reasonable to study how monthly labor income risk affects households'

[1]For a full exposition, see Ganong et al. (2024).

portfolio decisions, especially regarding the most liquid assets like direct holdings of stocks and mutual funds.

B3. Predictability of Earnings Changes

Another important consideration is the predictability of earnings changes. In the empirical analysis, the labor income innovation is idiosyncratic by construction–I purge predictable components and common factor variation. The discussion here concerns any residual predictability beyond those controls. For instance, if households intentionally choose to reduce their working hours–during the summer, for example–and consequently earn less in that month, they would have anticipated this risk well in advance and should not react unexpectedly. To address this issue, I use year-over-year changes in earnings instead of simple monthly changes to eliminate potential seasonality effects. Regarding the extent to which workers can control their weekly working hours, the General Social Survey (GSS), a nationally representative survey of U.S. residents, includes several questions about this topic. According to the analysis from Lambert, Henly and Kim (2019), 52% of hourly employees report having little to no control over their work schedules, 31% have some input, and 17% can decide freely or within certain limits. This suggests that employers largely determine schedules, contributing to at least half of the volatility in working hours.

A final consideration is persistence. With annual data, distinguishing transitory from permanent earnings shocks is difficult–especially when panels are short or lack the information needed to identify. At the monthly frequency, this is far less problematic: a permanent shock appears as a discrete jump to a new plateau, whereas a transitory shock is a brief spike that quickly reverts. In my setting, estimating innovations at the monthly level cleanly separates these cases–the variance measure captures transitory volatility without misclassifying permanent level shifts as repeated shocks.

B4. About Tobit

I estimate the main specifications with fixed-effects OLS rather than a Tobit model because the sample is restricted to households that have participated in the stock market at least

once. Excluding perpetual non-participants eliminates noise from households whose zero equity share may reflect unobserved factors unrelated to labor income risk, allowing a cleaner identification of portfolio responses among active investors.

What the Tobit model does. Let the latent (unobserved) desired risky share be

$$y_{iM}^{*} \;=\; X_{iM}'\beta + u_{iM}, \qquad u_{iM} \sim \mathcal{N}(0, \sigma^2), \tag{32}$$

and the observed share be left–censored at zero (and, for shares, sometimes also right–censored at one):

$$y_{iM} \;=\; \max\{0,\, \min\{1,\, y_{iM}^{*}\}\}. \tag{33}$$

Classical Tobit MLE assumes normality and homoskedasticity and imposes the *proportionality restriction* that the same linear index $X_{iM}'\beta$ jointly determines (i) the probability of participation ($y_{iM} > 0$) and (ii) the conditional level $E[y_{iM} \mid y_{iM} > 0, X_{iM}]$.

Why Tobit is ill-suited here.

- **Zeros are chosen corners, not censoring.** In SIPP, many zeros reflect fixed entry costs, distrust, advice constraints, or tastes. Tobit treats zeros as "would-be negative shares truncated at zero," which is the wrong economic mechanism.
- **Different forces drive entry vs. allocation.** Participation (extensive margin) depends on costs/inertia/literacy; the intensive margin depends on background risk (variance × skewness), wealth, and hedging. Tobit forces the *same* β to govern both margins; if the true DGP is two-part, Tobit is inconsistent.
- **Fragile distributional assumptions.** Risky shares exhibit mass points at 0, fat tails, and heteroskedasticity across wealth/age/industry. Normal, homoskedastic Tobit misfits these features.
- **High-dimensional fixed effects.** Nonlinear MLE with many FE (industry, year–month, household) suffers from incidental-parameter issues and unstable inference;

W(LS)OLS handles rich FE cleanly.

- **Survey weights.** Weighted Tobit is delicate and less transparent; weighted OLS with robust/clustered SE is straightforward.

What WOLS estimand captures (by design). I target the *intensive margin*: among households that have *ever* invested, how does the equity share respond to tail-dependent income risk? Formally, with weights w_{iM},

$$E\big[\, RS_{iM} \,\big|\, RS_{iM} > 0,\, X_{iM}\big] \;=\; X_{iM}'\beta \;+\; \varepsilon_{iM}, \qquad \widehat{\beta} \text{ by WOLS with FE, robust SE} \quad (34)$$

Conditioning on ever-investors avoids conflating background-risk effects with persistent, unobserved nonparticipation motives (fixed costs, deep preferences) that are orthogonal to our risk shocks.

Bottom line. Tobit is a poor match for stockholding survey data with corner solutions, rich FE, and survey weights, and it confounds entry with allocation via a restrictive proportionality assumption.

Appendix C. Inference Robustness

The baseline estimates in Table 5 use heteroskedasticity-robust standard errors. This appendix examines the sensitivity of the main results to alternative approaches to statistical inference, including clustered standard errors at the individual level, two-way clustering, annual aggregation, Fama–MacBeth estimation, and a pure cross-sectional specification.

C1. Why Clustering Matters

In a monthly panel where the same individual is observed for up to 96 months, residuals within an individual are likely serially correlated: a household that overweights equities relative to the model's prediction in one month will tend to do so in subsequent months as well. If this serial dependence is not accounted for, standard errors will understate the true sampling variability. The intraclass correlation (ICC) of the risky share within industry–education groups is only 0.008, suggesting that clustering at that level has limited bite. However, the relevant unit for the individual-level variance variable is the individual herself, and within-person persistence in portfolio positions is substantially higher.

Table 11 reproduces the one-way ANOVA for the risky share across industry–education groups.

Table 11: One-way Analysis of Variance for *RS100*

Source	SS	df	MS	F	Prob > F
Between IND_EDU	1,303,216.6	67	19,450.994	30.76	0.0000
Within IND_EDU	164,100,000	259,417	632.40653		
Total	165,400,000	259,484	637.26558		

Intraclass correlation	= 0.008
Estimated SD of IND_EDU effect	= 2.253
Estimated SD within IND_EDU	= 25.148

Although the ICC at the industry–education level is low, this does not imply that clustering is unnecessary. The key independent variable–idiosyncratic variance–varies at the individual–year level, and the dependent variable (risky share) is persistent within individuals. For this reason, I report results clustered at the individual level as the primary

robustness check.

C2. Monthly Panel with Alternative Standard Errors

Table 12 reports the full specification from column (6) of Table 5 under four variance estimators: (i) heteroskedasticity-robust (the baseline), (ii) clustered by individual, (iii) clustered by industry–education group, and (iv) two-way clustered by individual and year–month.

Table 12: Monthly Panel: Alternative Standard Errors

	Dependent variable: Risky Share (×100)			
	(1) Robust	(2) Cluster(ID)	(3) Cluster(IND_EDU)	(4) Two-way(ID + Date)
Idiosyncratic variance	−0.059*	−0.059	−0.059	−0.059
	(0.036)	(0.112)	(0.116)	(0.111)
Opportunity ($L9050$)	0.161	0.161	0.161	0.161
	(0.101)	(0.287)	(0.337)	(0.285)
Disaster ($L5010$)	−0.212**	−0.212	−0.212	−0.212
	(0.087)	(0.250)	(0.255)	(0.259)
Opportunity × variance	0.173***	0.173*	0.173*	0.173*
	(0.033)	(0.088)	(0.093)	(0.088)
Disaster × variance	−0.047***	−0.047	−0.047	−0.047
	(0.015)	(0.039)	(0.040)	(0.039)
Household controls	Yes	Yes	Yes	Yes
Tenure FE	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes
Year–month FE	Yes	Yes	Yes	Yes
Observations	259,485	259,485	259,485	259,485
R^2	0.028	0.028	0.028	0.028

Notes. Each column reports the same point estimates from the full specification in column (6) of Table 5; only the variance estimator differs. Column (1): heteroskedasticity-robust. Column (2): clustered by individual ID. Column (3): clustered by industry–education group. Column (4): two-way clustered by individual and year–month. ***/**/* denote significance at the 1%, 5%, and 10% levels.

The point estimates are mechanically identical across columns, as only the standard errors change. The opportunity-risk interaction ($\hat{\beta}_4 = 0.173$) retains marginal significance at the 10% level under all three clustering schemes but falls from the 1% level with robust errors. The disaster-risk interaction ($\hat{\beta}_5 = -0.047$) loses significance entirely. This pattern is expected: the monthly panel contains up to 96 observations per individual, and within-person persistence in portfolio positions generates substantial serial correlation in the residuals. Clustering by individual treats each person's entire time series as a single draw, dramatically reducing the effective sample size and inflating standard errors by a factor of roughly two to three relative to the heteroskedasticity-robust benchmark.

C3. Annual Panel

To address the serial-correlation problem directly, I aggregate the data to annual frequency. For each individual–year, I compute the mean risky share, the mean of the log income and log wealth, and the annual-average tail measures (averaging the monthly $L9050$ and $L5010$ within each industry–education–year cell). The annual panel comprises 22,819 person–year observations on 11,887 unique individuals.

Table 13: Annual Panel: Risky Share and Tail-Dependent Labor Income Risk

	Dependent variable: Risky Share (×100)			
	(1)	(2)	(3)	(4)
Idiosyncratic variance	0.083	0.083	−0.076	−0.090
	(0.071)	(0.070)	(0.132)	(0.129)
Opportunity ($L9050_{\text{ann}}$)		0.326	0.045	−0.278
		(0.416)	(0.412)	(0.430)
Disaster ($L5010_{\text{ann}}$)		−0.295	−0.219	−0.646*
		(0.336)	(0.344)	(0.348)
Opportunity × variance			0.301**	0.313**
			(0.127)	(0.126)
Disaster × variance			−0.121**	−0.120**
			(0.056)	(0.057)
Household controls	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	
IND_EDU FE				Yes
Year FE	Yes	Yes	Yes	Yes
SE clustered by	ID	ID	ID	ID
Observations	22,819	22,819	22,819	22,819
R^2	0.022	0.022	0.022	0.026

Notes. Weighted OLS estimates on annual panel data (person–year observations). The dependent variable is the annual average risky share (×100). Tail measures are annual averages of the monthly $L9050$ and $L5010$ within each industry–education–year cell. All columns cluster standard errors by individual. Column (4) replaces industry FE with industry–education group FE, which absorbs the education control. ***/**/* denote significance at the 1%, 5%, and 10% levels.

Aggregation to annual frequency substantially reduces within-person serial correlation: each individual contributes at most four to eight annual observations rather than dozens of monthly ones, so the cluster correction imposes a much smaller penalty on precision. The key interaction terms retain significance at the 5% level under individual-level clustering. The opportunity-risk coefficient is $\hat{\beta}_4 = 0.30$ (SE = 0.13, $p < 0.05$) and the disaster-risk coefficient is $\hat{\beta}_5 = -0.12$ (SE = 0.06, $p < 0.05$). The results are stable when industry–education group fixed effects replace industry fixed effects in column (4), addressing concerns about self-selection into industry–education cells.

Table 14 further reports two-way clustering by individual and year, which accounts for both within-person persistence and cross-sectional dependence within a given year.

Table 14: Annual Panel: Two-way Clustered Standard Errors (ID + Year)

	Dependent variable: Risky Share (×100)	
	(1) Industry FE	(2) IND_EDU FE
Idiosyncratic variance	−0.076	−0.090
	(0.099)	(0.093)
Opportunity ($L9050_{\text{ann}}$)	0.045	−0.278
	(0.332)	(0.410)
Disaster ($L5010_{\text{ann}}$)	−0.219	−0.646
	(0.415)	(0.421)
Opportunity × variance	0.301***	0.313***
	(0.069)	(0.068)
Disaster × variance	−0.121**	−0.120*
	(0.044)	(0.053)
Household controls	Yes	Yes
Industry FE	Yes	
IND_EDU FE		Yes
Year FE	Yes	Yes
SE clustered by	ID + Year	ID + Year
Observations	22,819	22,819
R^2	0.022	0.026

Notes. Same specifications as columns (3)–(4) of Table 13, but with two-way clustering by individual and year. ***/**/* denote significance at the 1%, 5%, and 10% levels.

Under two-way clustering, the opportunity-risk interaction sharpens to the 1% level ($\hat{\beta}_4 = 0.30$, SE $= 0.07$) and the disaster-risk interaction remains significant ($\hat{\beta}_5 = -0.12$, SE $= 0.04$, $p < 0.05$). The reduction in standard errors relative to one-way clustering by individual is consistent with the fact that, once year fixed effects are included, there is little residual cross-sectional dependence for the year dimension of the cluster to absorb.

C4. Fama–MacBeth Estimation

As an independent check that does not rely on a particular cluster structure, I implement the Fama–MacBeth (1973) procedure. For each of the 96 months in the sample, I run a weighted cross-sectional regression of the risky share on the full set of regressors (including industry dummies). I then compute the time-series average of each month's coefficient and report Newey–West standard errors with 12 lags to account for potential overlap in the annual variance measure.

Table 15: Fama–MacBeth Estimates with Newey–West Standard Errors

	Mean coefficient	NW std. error	t-statistic
Idiosyncratic variance	−0.167	0.090	−1.86*
Opportunity ($L9050$)	1.016	0.573	1.77*
Disaster ($L5010$)	−0.052	0.341	−0.15
Opportunity × variance	0.418	0.163	2.57**
Disaster × variance	−0.117	0.072	−1.62

Notes. Fama–MacBeth (1973) procedure with 96 monthly cross-sections. Each monthly regression includes the full set of household controls and industry fixed effects. Newey–West standard errors use 12 lags (Bartlett kernel). ***/**/* denote significance at the 1%, 5%, and 10% levels.

The Fama–MacBeth estimates confirm the key pattern: the opportunity-risk interaction is positive and statistically significant ($t = 2.57$, $p < 0.05$). The disaster-risk interaction is negative with a t-statistic of -1.62, just below the 10% threshold but directionally consistent. This independent estimator, which makes no assumption about the within-person correlation structure, corroborates the annual-panel findings.

C5. Pure Cross-Section

Finally, I collapse the data to a single observation per individual (averaging over all months) and use time-invariant tail measures computed by pooling all months within each industry–education group. This pure cross-sectional specification eliminates any panel-related serial-correlation concern but at the cost of discarding all within-person time-series variation.

With only 11,887 observations and no time-series variation, the interaction terms are statistically insignificant, although the signs remain consistent with the theory. This is not surprising: the cross-section discards the within-group time variation that provides much of the identifying power in the panel specifications. It does, however, confirm that the panel results are not driven by spurious time-series correlation.

C6. Summary of Inference Robustness

Table 17 collects the estimates of $\hat{\beta}_4$ (opportunity risk) and $\hat{\beta}_5$ (disaster risk) across all specifications and inference methods.

The opportunity-risk interaction ($\hat{\beta}_4$) is positive and statistically significant at the 5%

Table 16: Pure Cross-Section: Risky Share and Tail-Dependent Labor Income Risk

	Dependent variable: Risky Share (×100)	
	(1) Industry FE	(2) IND_EDU FE
Idiosyncratic variance	0.130	0.136
	(0.210)	(0.205)
Opportunity ($L9050_{\text{fixed}}$)	2.683	
	(2.204)	
Disaster ($L5010_{\text{fixed}}$)	0.370	
	(1.615)	
Opportunity × variance	0.091	0.028
	(0.413)	(0.403)
Disaster × variance	−0.121	−0.075
	(0.259)	(0.258)
Household controls	Yes	Yes
Industry FE	Yes	
IND_EDU FE		Yes
SE type	Het.-robust	Het.-robust
Observations	11,887	11,887
R^2	0.020	0.025

Notes. Pure cross-section: one observation per individual, averaging over all months. Tail measures are time-invariant, computed by pooling all months within each industry–education group. Heteroskedasticity-robust standard errors. ***/**/* denote significance at the 1%, 5%, and 10% levels.

Table 17: Summary: Key Interaction Coefficients Across Specifications

		Opportunity × var ($\hat{\beta}_4$)		Disaster × var ($\hat{\beta}_5$)	
Specification	SE method	Coeff.	Signif.	Coeff.	Signif.
Monthly panel	Robust	0.173	***	−0.047	***
Monthly panel	Cluster(ID)	0.173	*	−0.047	
Monthly panel	Cluster(IND_EDU)	0.173	*	−0.047	
Monthly panel	Two-way(ID+Date)	0.173	*	−0.047	
Annual panel	Cluster(ID)	0.301	**	−0.121	**
Annual panel	Two-way(ID+Year)	0.301	***	−0.121	**
Annual (IND_EDU FE)	Cluster(ID)	0.313	**	−0.120	**
Fama–MacBeth	Newey–West (12)	0.418	**	−0.117	
Cross-section	Robust	0.091		−0.121	

Notes. Each row reports the coefficient on the opportunity-risk interaction ($L9050\times$ variance) and the disaster-risk interaction ($L5010\times$ variance) from the full specification. Significance: ***/**/* denote the 1%, 5%, and 10% levels.

level or better in the annual panel (all clustering variants) and Fama–MacBeth specifications, and at the 10% level in the monthly panel with individual clustering. The disaster-risk interaction ($\hat{\beta}_5$) is significant at the 5% level in the annual panel but falls short of significance in the monthly panel and Fama–MacBeth settings. The cross-section, which discards all time-series variation, yields directionally consistent but imprecise estimates.

The overall picture supports the conclusion that the tail-dependent effects of labor income risk on portfolio choice are robust to conservative inference. The most stringent test—annual aggregation with individual-level clustering—preserves the significance of both interaction terms.

Appendix D. Additional Robustness Tests

D1. Normalized Tails

A natural concern is that, in the main analysis, the upper and lower tails are not normalized by overall dispersion. Specifically, I define

$$L9050 = P90 - P50, \qquad L5010 = P50 - P10,$$

so the absolute length of each tail may mechanically reflect changes in the spread of the distribution.

To address this, I construct standardized tail measures that scale each tail by the overall 90–10 range:

$$S9050 = \frac{L9050}{L9010} = \frac{P90 - P50}{P90 - P10}, \qquad S5010 = \frac{L5010}{L9010} = \frac{P50 - P10}{P90 - P10}.$$

By construction, $S9050 + S5010 = 1$, so including both in the same regression would imply perfect collinearity. I therefore estimate specifications for the upper and lower tails separately. The corresponding regression results are reported in Table 18. As expected, the coefficients on $S9050$ and $S5010$ are mechanically equal in magnitude and opposite in sign. More importantly, the economic interpretation is consistent with the main results: once I decompose total dispersion into normalized tails, the positive effect of volatility is fully accounted for by the upper tail, $S9050$, while the lower tail, $S5010$, does not drive the positive relationship with risk taking. This confirms that it is the opportunity component of labor income risk, rather than dispersion per se, that is associated with higher portfolio risk.

D2. Using Portfolio Shares as a Proxy for Risk Aversion

The survey does not elicit direct measures of risk preferences, so I proxy for individual risk aversion using their portfolio choices. Specifically, following the approach in Egan, MacKay and Yang (2021), I use each individual's *initial* reported risky asset share as a revealed preference indicator of risk tolerance and split the sample into two groups based on this

Table 18: Standardized Tail Measures

	Risky Share	
	(1)	(2)
S5010	−1.785***	
	(0.472)	
S9050		1.785***
		(0.472)
Idiosyncratic variance	0.320***	−0.211**
	(0.067)	(0.065)
S5010×Idiosyncratic variance	−0.531***	
	(0.128)	
S9050×Idiosyncratic variance		0.531***
		(0.128)
Individual controls	Yes	Yes
Tenure FE	Yes	Yes
Year-month FE	Yes	Yes
Industry FE	Yes	Yes
Observations	287,848	287,848
Adjusted R^2	0.031	0.031

measure: individuals with a relatively *high* initial risky share (lower risk aversion) and those with a *low* initial risky share (higher risk aversion). In the data, group L (low risk aversion) has an average risky share of 21.1%, whereas group H (high risk aversion) averages 13.4%.

This proxy is admittedly noisy. Initial portfolio holdings may reflect short-run circumstances or timing rather than stable preferences—for example, some households may temporarily hold little risky wealth because they are waiting to rebalance, are in the process of accumulating funds, or have just experienced a shock unrelated to their underlying risk attitudes. Thus, the resulting groups should be interpreted as coarse risk aversion types rather than precise preference parameters. With this caveat in mind, Figure 14 reports the estimated tail risk effects separately for the two groups and shows how exposure to opportunity (right-tail) and disaster (left-tail) risk differentially affects households that appear more versus less risk averse.

D3. Dropping Clusters with Sparse Observations

A potential concern with the tail risk measures is that they are estimated at the cluster level, yet the number of observations per cluster varies considerably. In small clusters, higher-order moments (such as skewness) may be poorly estimated. To address this, I exclude all

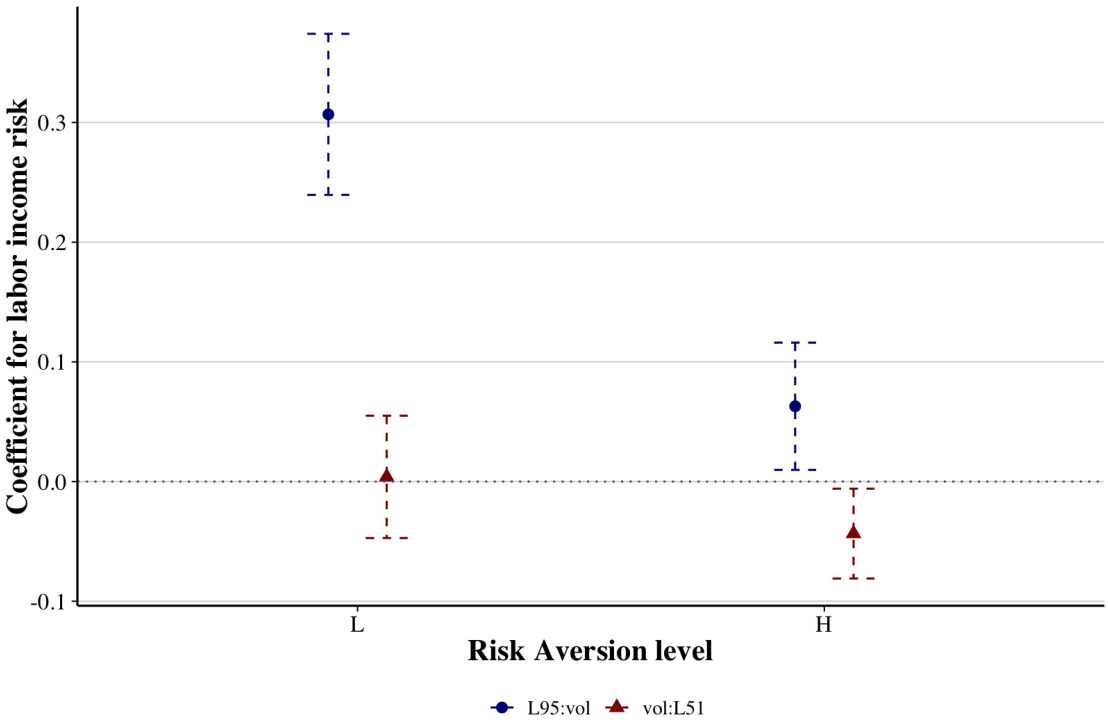


Figure 14: Effect of labor income risk by risk aversion group

Notes. This figure displays the regression coefficients of the equity share on labor income risk when I run the same regressions as in column (6) of Table 5 using different risk aversion groups of workers. L denotes lower risk aversion and H denotes higher risk aversion. The average risky share is 21.1% in group L and 13.4% in group H. Dashed lines represent 95% confidence intervals.

industry–education–by–month cells with fewer than 200 observations.[15] In other words, if in a given month fewer than 200 workers belong to a particular industry–education group, that cell is dropped from the analysis.

The main regression results for the restricted sample are reported in Table 19. Imposing the requirement of at least 200 observations per industry–education–month cell reduces the sample from roughly 290,000 to about 210,000 observations. The key patterns are unchanged: the signs and magnitudes of the coefficients on L95, volatility, and their interaction remain very similar to the baseline. The main difference is that the standalone coefficient on L51 becomes statistically indistinguishable from zero, while the interaction term $L51 \times \text{vol}$ remains significant.

This change is mechanically related to the trimming rule. Cells with very few observations tend to be those experiencing substantial employment losses (for instance, industry–education groups hit by layoffs), which are precisely the environments where left-tail (disaster) risk is most pronounced. By dropping all low-count cells, the specification disproportionately excludes high-disaster clusters and thereby attenuates the average effect of L51. The fact that the interaction $L51 \times \text{vol}$ remains robust suggests that, conditional on being in the sample, heterogeneity in individual volatility continues to mediate the impact of disaster risk on portfolio choice, but the cross-cell variation in pure disaster intensity is compressed by the sample restriction.

D4. Investor Fixed Effects

The main specifications exploit cross-sectional variation: they compare investors who differ in their exposure to labor income risk (variance and tails) and in their portfolio choices. While this is standard in the literature, the SIPP's panel structure also allows me to examine within-investor adjustments over time by adding investor fixed effects.

However, the panel is relatively short for each individual (at most four years of monthly observations, and often less due to entry/exit and survey attrition), which limits the amount of within-person variation in both tail-risk measures and portfolio shares. Once I include investor fixed effects, identification comes only from changes in risk exposure and equity hold-

[15] Results are qualitatively similar when using a threshold of 100.

Table 19: Dropping Clusters with Fewer than 200 Observations

	Risky Share $\in [0, 100]$		
	(1)	(2)	(3)
$L9050$	0.605**		0.581**
	(0.221)		(0.223)
$L5010$		0.048	0.001
		(0.138)	(0.139)
Idiosyncratic variance	−0.098*	0.111***	−0.037
	(0.042)	(0.031)	(0.046)
$L9050\times$ Idiosyncratic variance	0.173***		0.198***
	(0.045)		(0.046)
$L5010\times$ Idiosyncratic variance		−0.069**	−0.091***
		(0.028)	(0.029)
Individual controls	Yes	Yes	Yes
Year-month FE	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes
Observations	210,977	210,977	210,977
Adjusted R^2	0.028	0.029	0.029

ings for the same household over time; all time-invariant heterogeneity in risk preferences, financial literacy, and background characteristics is absorbed by the fixed effects.

For this reason, I treat the investor-FE specification as a robustness check rather than my baseline. The results, reported in Table 20, are broadly consistent with the main analysis: the interaction between right-tail length and volatility remains positive, and left-tail risk continues to be associated with lower risky shares. At the same time, several coefficients lose statistical significance, which is natural given (i) the reduced effective variation once fixed effects are included and (ii) the short individual panels. Overall, the investor-FE estimates reinforce the main message–that tail shape matters for portfolio choice–while highlighting that the strongest and most precisely estimated effects arise in the cross-sectional dimension.

D5. Within-household Analysis

Much of the prior empirical research has concentrated on the investment decisions of individuals. In this section, I examine the within-household impact of labor income shocks. I refer to the primary earner in a household as the head, and the secondary earner as the spouse. I initially identify all households with two earners from the dataset, as they constitute the majority of households with multiple earners. Subsequently, I divide them into two subgroups

Table 20: Investor Fixed Effects

	Risky Share $\in [0, 100]$		
	(1)	(2)	(3)
$L9050$	0.075	0.078	
	(0.078)	(0.078)	
$L5010$	-0.467^{***}		-0.481^{***}
	(0.066)		(0.066)
Idiosyncratic variance	-0.028	-0.046	0.046
	(0.038)	(0.035)	(0.033)
$L9050\times$ Idiosyncratic variance	0.090^{***}	0.087^{***}	
	(0.023)	(0.023)	
$L5010\times$ Idiosyncratic variance	-0.018		-0.006
	(0.013)		(0.013)
Individual controls	Yes	Yes	Yes
Year-month FE	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes
Observations	259,485	259,485	259,485
Adjusted R^2	0.659	0.659	0.659
Within R^2	0.041	0.041	0.041

based on whether the spouse is exposed to a lower or higher level of skewness. The spouse exposed to a lower skewness level is considered to be facing risk, while the one exposed to a higher skewness level is regarded as facing opportunity, consistent with the terminology used in earlier sections. I am interested in exploring how heads react to the labor income risks of their spouses, specifically whether they act to mitigate the risks faced by their spouse.

Figure 15 offers a revealing picture. The shaded red mass plots the distribution of head-of-household labor income skewness conditional on the spouse currently facing negative skewness–i.e., a disaster risk environment. Put differently, it traces the macro-level earnings risk borne by the head when the spouse is exposed either to downside or upside risk.

Some intra-household correlation is evident, though not pronounced. On the left tail ($x < 0$) the red density lies slightly above the blue density, indicating that when the spouse experiences disaster risk, the head is somewhat more likely to do so as well. A symmetric, though modest, pattern appears on the right tail for positive skewness (opportunity risk). Overall, spouses' labor income risks are positively but only weakly correlated within the household.

I utilize the following formula to examine the within-household impact of labor income

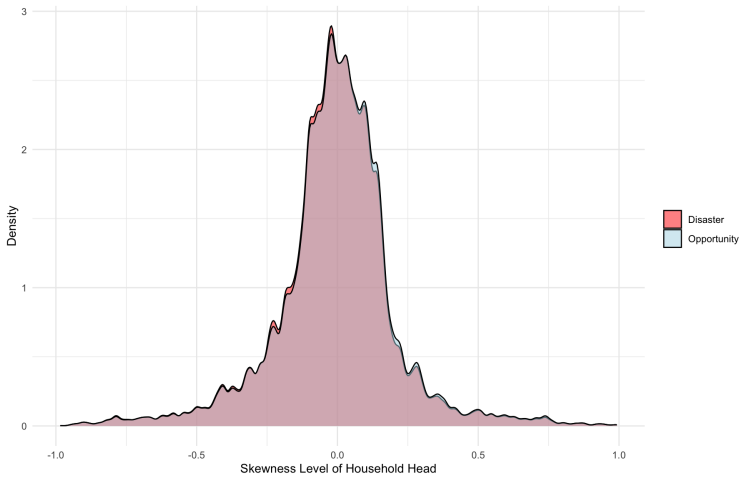


Figure 15: Labor income risks within households

Notes. This figure shows the distribution of household head skewness, conditional on whether the spouse is exposed to opportunity risk or disaster risk. The red distribution represents households where the spouse faces disaster risk, while the blue represents those with opportunity risk.

volatility, taking into account the skewness level of spouses:

$$RS_{i,M} = \alpha + \beta_1 \cdot vol_{i,Y} + \beta_2 \cdot SS_{i,M} + \beta_3 \cdot SS_{i,M} \times vol_{i,Y} + Control_{i,M} + \varepsilon_{i,M}, \tag{35}$$

where $SS_{i,M}$ denotes the skewness level of the spouse, with $SS_{i,M} = LSS_{i,M} = 1$ indicating that the spouse's skewness level is below 0, and $SS_{i,M} = RSS_{i,M} = 1$ signifying that the spouse's skewness level exceeds 0. The findings are presented in Table 21.

Even when not accounting for the skewness of the head's income, if the spouse exhibits a disaster-risk environment (negative skewness), heads of households tend to decrease their equity holdings. Conversely, when paired with a spouse in an opportunity-risk environment (positive skewness), heads increase their risky asset holdings. These patterns are most pronounced when the head's own labor income variance is low. Heads' allocations co-vary with spousal risk in a manner consistent with intra-household risk sharing: when one spouse faces concentrated downside risk, the household adjusts the other spouse's financial portfolio to offset that risk, and vice versa when one spouse enjoys upside opportunity.

Table 21: Regression Analysis of Household Heads' Portfolio Holdings

	% Share of Assets Directly Invested in Stocks (Head)
Disaster Risk (Spouse)	-0.884^{***}
	(0.123)
Opportunity Risk (Spouse)	1.233^{***}
	(0.124)
Variance (Head)	0.070^{***}
	(0.012)
Disaster Risk (Spouse)×Variance (Head)	0.468^{***}
	(0.059)
Opportunity Risk (Spouse)×Variance (Head)	-0.552^{***}
	(0.059)
Observations	529,265
Adjusted R^2	0.031

Notes. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Appendix E. Supplementary Figures

Figure 16 plots the time series of L9050 and L5010, averaged monthly. When the two lines move closely together, labor income risk is more symmetric. When both series are high, overall volatility is elevated. When the gap between them widens, it indicates growing asymmetry–either more upside opportunity (right tail) or more downside risk (left tail), depending on which tail length increases.

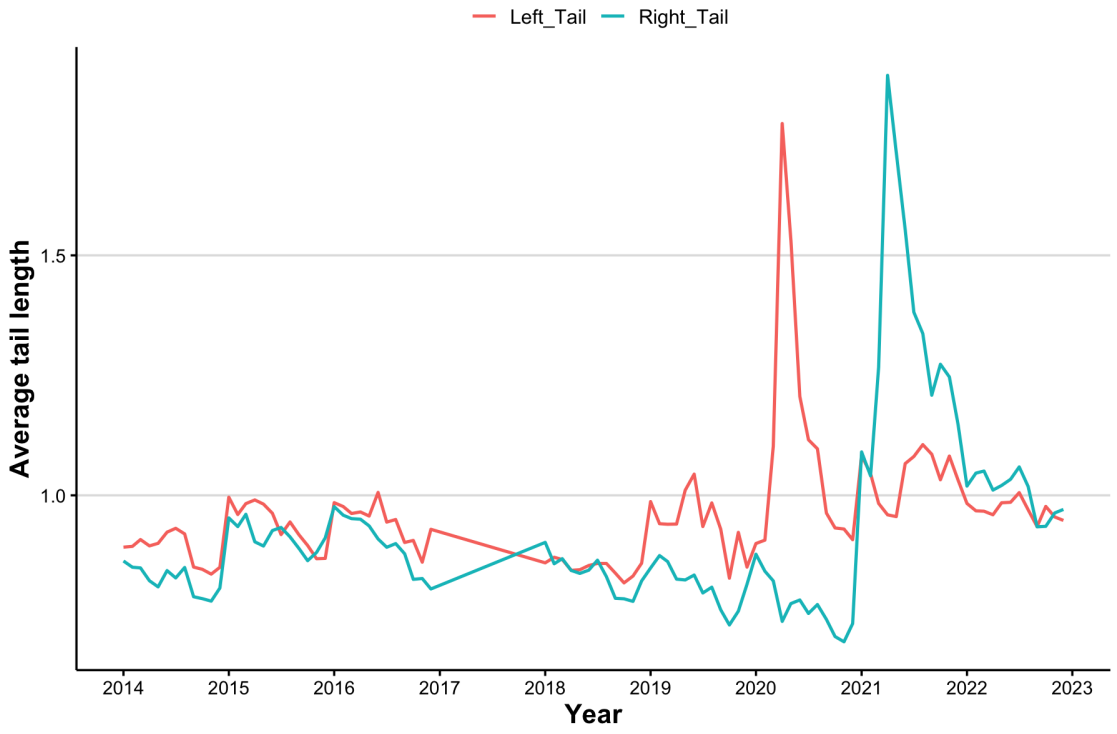


Figure 16: Evolution of Tail Lengths

Notes. The figure plots monthly averages of labor income growth tail length over time. There are no observations in 2017 because SIPP data are not available for that year.

Figure 17 plots the right-tail length L95 (i.e., L9050) for selected industry–education groups over time, using monthly averages. In each panel, the solid blue line is the time-varying group-level L95 measure, while the red dashed line is the corresponding fixed (time-invariant) group benchmark. The figure highlights substantial within-group time variation in right-tail opportunities: most groups fluctuate around their long-run level, but several exhibit sharp spikes—especially around 2020–2021—indicating short-lived episodes of un-

usually strong upside dispersion in labor-income growth within those groups.

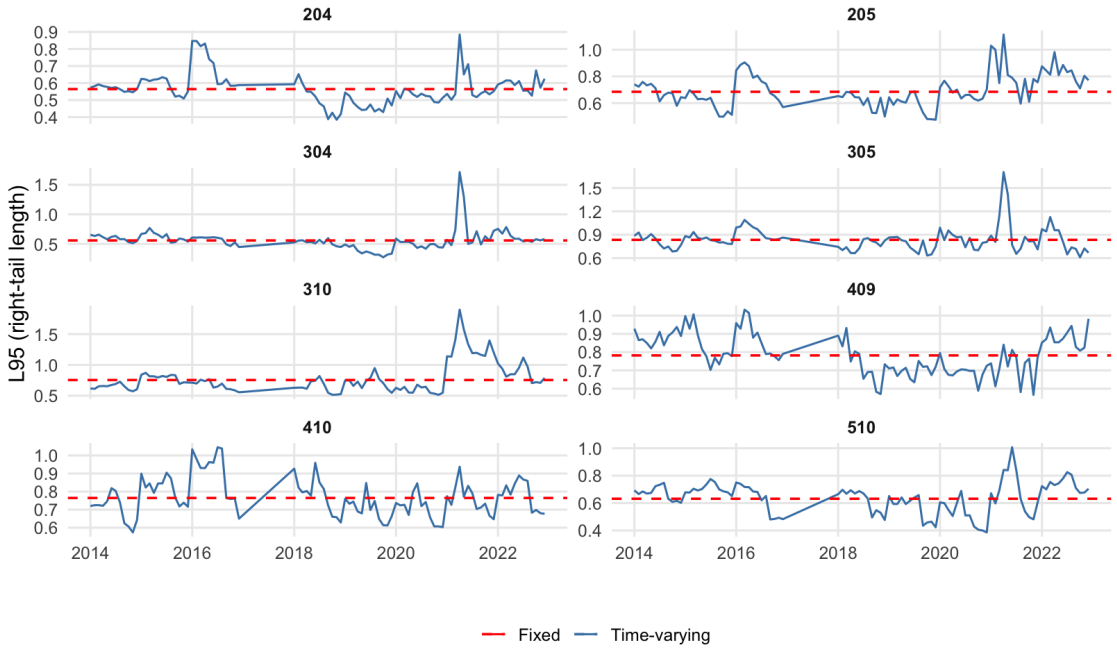


Figure 17: Right-Tail Length (L95): Time-Varying vs. Fixed Measures

Notes. Each panel corresponds to one industry–education group (panel titles report the group code). The solid blue line plots the monthly average right-tail length $L9050$ computed from the cross-sectional distribution of labor-income growth within the group. The red dashed line is the corresponding fixed (time-invariant) benchmark for the same group. There are no observations in 2017 because SIPP data are not available for that year.

Appendix F. Model

A Simple GE Model

This is a very simple general equilibrium model intended to illustrate how downside risk aversion can widen the gap between the rich and the poor. It is based on the model from Auclert, Rognlie and Straub (2023), where a stock market is not included in their setting.

Consider a continuous-time model featuring N types of households, where $i = 1, \cdots, N$. These households vary in their spending and investment behaviors based on their wealth levels.

Empirical evidence shows that wealthier individuals tend to exhibit a lower marginal propensity to spend (MPC) compared to poorer households (Fagereng et al. (2019); Straub (2019); Mian, Straub and Sufi (2020)). Simultaneously, wealthier households allocate a larger proportion of their wealth to risky asset markets (Favilukis (2013); Bender et al. (2022)), resulting in significantly higher total returns on assets. To capture these facts, I employ the following framework: (1) $m_1 > m_2 > \cdots > m_N$, indicating that wealthier households have lower MPCs; (2) $r_1 < r_2 < \cdots < r_N$, representing an increase in asset returns with wealth; and (3) agents commence work at age 25 and retire at 65, with their behavior described by a utility function over consumption.

The model is linearized around the steady state where each household type owns a certain stock of assets:

$$c_{it} = m_i a_{it} \quad s_{it} = r_i a_{it} \quad Y_t = \sum_{i=1}^{N} (c_{it} - s_{it}) \tag{36}$$

Where a_{it} represents the value of the financial asset possessed by household i at time t, c_{it} denotes consumption, s_{it} signifies the net income from investment, r_i is the return on assets (ROA) for household i. Y_t stands for the aggregate demand and income. The interpretation is straightforward: each year, households spend a portion of their assets on consumption and earn returns on the assets they own. The total output of the economy is determined by the aggregate consumption of the entire population minus the amount paid to those who invested and required returns. This output is then redistributed among the population based on their assigned labor income shares. For simplicity, consider consumption and investment

occurring simultaneously, meaning there's no need to subtract the portion they consumed before obtaining returns from assets. Thus, the law of motion for the value of financial assets can be expressed as follows:

$$\dot{a}_{it} = \theta_i Y_t + s_{it} - c_{it} \tag{37}$$

Here, θ satisfies $\sum_{i=1}^{N} \theta_i = 1$, denoting the labor income shares across types. An N-dimensional linear differential equation can be derived by stacking variables into length-N vectors, as follows:

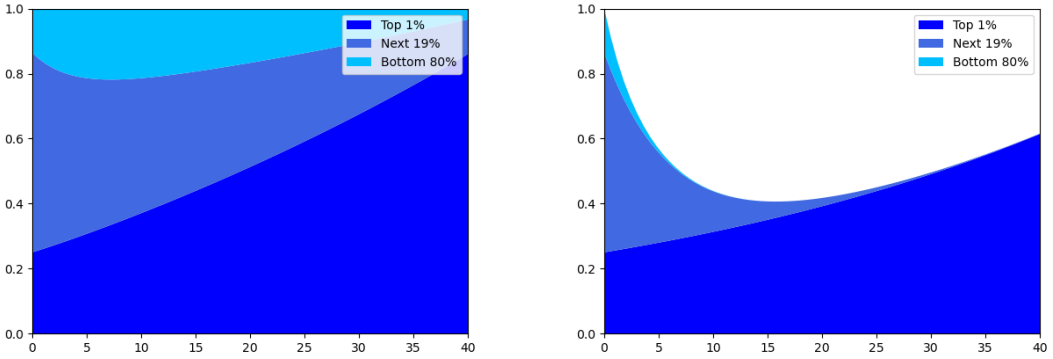
$$\dot{\mathbf{a}}_t = (-M + R + \boldsymbol{\theta}\mathbf{m}' - \boldsymbol{\theta}\mathbf{r}')\mathbf{a}_t$$

Where $M = \text{diag}(m)$ and $R = \text{diag}(r)$. In Table 22, I calibrate the initial assets, labor income shares, and return on assets from SIPP. Since there is no consumption data in SIPP, I utilize the annual MPC from other studies.

The results of simulation are shown in Figure 18 and 19. During the early years, roughly from 25 to around 30, the bottom 80% group gains more shares due to the continued significance of human capital. However, after 40, the share of the top 1% rises sharply. This trend is a result of the compounding effect of the lower MPC and the higher return on assets. As the top 1% consume a smaller proportion of their wealth and invest more, the value of their financial assets grows increasingly faster as their wealth accumulates. By the time of retirement, the asset shares of the top 1% have risen from 25% to around 80%, while the sum of the next 99% of the population shares roughly 20% of the assets in the economy.

Furthermore, these figures illustrate the distinction between partial equilibrium and general equilibrium. In the partial equilibrium case, it demonstrates the scenario where output is made exogenous ($\theta = 0$), representing the absence of income feedback. As shown in Figure 19, in the partial equilibrium scenario, the aggregate return on investment initially decreases, then gradually increases after age 35. However, in the general equilibrium case, which better fits the real-life scenario, the return on investment consistently grows from the outset because households accumulate wealth as they age. Regarding aggregate consumption, the general equilibrium case better mirrors reality: it initially increases during early ages and then gradually decreases as households age.

To present a counterfactual scenario, Figure 20 displays the case where all types of house-



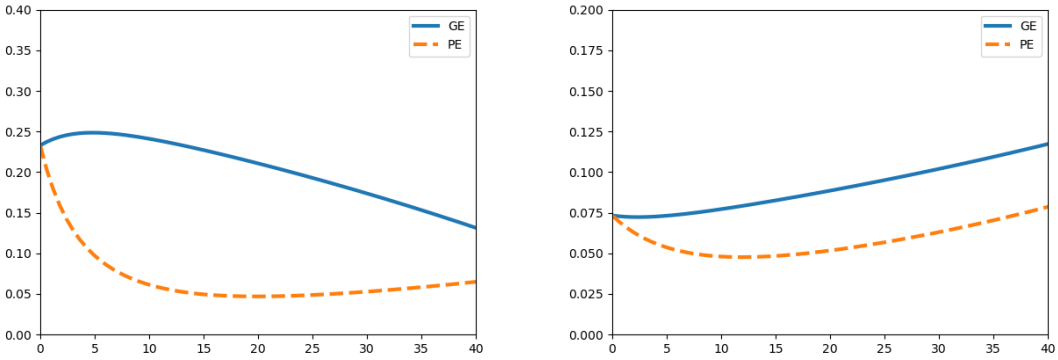
(a) General equilibrium (b) Partial equilibrium

Figure 18: Life cycle evolution

holds have the same return on assets ($r = 0.1$). It is evident that under these circumstances, inequality is less pronounced compared to the previous case, and at aggregate level, consumption path over the life cycle will also be smoother. This simple model emphasizes the impact of return heterogeneity in addition to the existing heterogeneity in MPC when addressing wealth inequality. If the bottom 80% of households can fully utilize their financial assets and achieve a similar return on assets as wealthier groups, despite having a higher MPC, this group of households can still be better off in terms of financial wealth when they reach retirement.

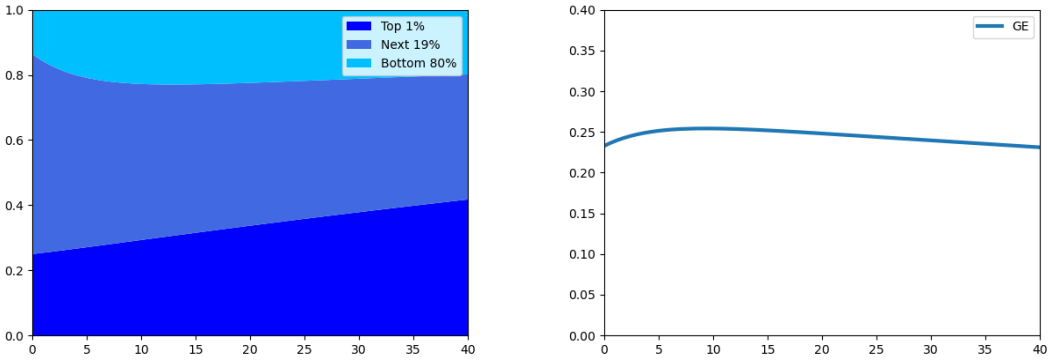
Table 22: Basic Calibration

	Asset Value Percentile			
	80%	20%	1%	Source
Initial Assets (Age 25)	0.136	0.614	0.250	SIPP
Labor Income Shares	0.63	0.34	0.03	SIPP
MPC	0.4	0.2	0.1	Baseline
Return on Assets	0.015	0.062	0.12	SIPP



(a) Consumption (b) Return on investment

Figure 19: Aggregate evolution



(a) Life cycle evolution (b) Consumption

Figure 20: Counterfactual case