Countercyclical Income Risk and Portfolio Choices: Evidence from Sweden

Sylvain Catherine
Wharton

Paolo Sodini
Stockholm School of Economics

Yapei Zhang
ShanghaiTech University

May 26, 2022

Abstract

Using Swedish administrative panel data, we document that workers facing higher left-tail income risk when equity markets perform poorly are less likely to participate in the stock market and, conditional on participation, have lower equity shares. In line with theory, the relationship between cyclical skewness and stock holdings is proportional to the share of human capital in a worker's total wealth and vanishes as workers get closer to retirement. Cyclical skewness also predicts portfolio differences within pairs of identical twins. Our findings show that households hedge against correlated tail risks, an important mechanism in asset pricing and portfolio choice models.

Keywords: Household finance, Asset pricing, Disaster risk, Labor income risk, Portfolio choices

JEL codes: G11, G50, D14

How does labor income affect portfolios? Early life-cycle models suggest that human capital increases the demand for equity because labor income risk is largely uncorrelated with stock returns (Viceira, 2001; Cocco, Gomes and Maenhout, 2005). This prediction is at odds with the reluctance of young workers to buy stocks and makes it more difficult to explain the equity premium. The level and cross-section of stock holdings are better explained in models in which the variance (Storesletten, Telmer and Yaron, 2007; Lynch and Tan, 2011) or skewness (Catherine, 2021) of income risk varies over the business cycle. Furthermore, countercyclical income risk helps asset pricing models match the level, volatility, and cross-section of prices (Constantinides and Duffie, 1996; Constantinides and Ghosh, 2017; Schmidt, 2016; Ebrahimian and Wachter, 2020; Paron, 2022). Overall, quantitative models with countercyclical income risk can solve key empirical puzzles in finance. Yet, even though this idea dates back to Mankiw (1986), there is no direct evidence that households facing higher income risk during recessions invest less in stocks.

We fill this gap in the literature by studying portfolios and earnings trajectories in administrative panel data from Sweden. We document that workers facing higher cyclical skewness – tail risk during recessions – are less likely to participate in the stock market and, when they do, invest a lower share of their financial wealth in stocks. Quantitatively, a two-standard-deviation increase in cyclical skewness is associated with a 2 to 8 percentage point decline in equity shares, an 8% to 29% drop relative to the sample average. This effect is largely driven by the decision not to participate.

Is the negative relationship between cyclical skewness and stock holdings evidence of hedging? We explore this interpretation of our findings in two ways. First, following Calvet and Sodini (2014)'s empirical strategy, we study portfolio differences within pairs of identical twins to address the concern that stock holdings and occupations are jointly determined by latent variables such as risk preference, general intelligence, upbringing and expected inheritance. We find the slope of the relationship between cyclical skewness and equity shares to be quantitatively similar within twin–year pairs and across the entire working population.

Second, we test whether this interpretation is consistent with theory. Theory predicts the hedging demand for equity to be proportional to the human capital share of wealth. As predicted, in the data, the negative relationship between cyclical skewness and equity shares is linear in the human capital-to-wealth ratio with an intercept of zero. The moderating effect of the human capital share supports the interpretation of our findings as evidence of hedging. Indeed, an omitted variable could only explain our results if its effect on equity shares were similarly moderated by the human capital share. In particular, workers whose human capital-to-wealth ratio approaches zero constitute a theoretical placebo group as theory predicts that cyclical skewness should not affect their portfolio. Relatedly, theory predicts that the relationship between cyclical skewness and equity shares converges to zero as workers approach retirement. These two predictions are verified in the data.

Theoretically, the human capital share plays a moderating role because it determines the permanent income elasticity of consumption. It measures how human capital returns translate into changes in lifetime consumption. We build on this insight to construct measures of labor-market-implied consumption risk. To do so, we compute moments of the permanent income shock distribution by filtering out transitory shocks and adjust these moments to account for the diversification effect of non-human wealth. The resulting labor-market-implied consumption risk moments present two advantages. Strongly guided by theory, they offer an additional challenge to the economic interpretation of our findings. Moreover, they provide an additional source of variation, as workers with the same education and industry of employment have different labor-market-implied consumption risk depending on whether their human capital risk is diversified by other sources of wealth. Our findings become more significant when cyclical skewness is measured in terms of implied consumption risk and are robust to controlling for both education and industry fixed effects.

Our methodology is as follows. In a first step, we build measures of countercyclical risk for 321 groups of workers, sorted by industry of employment and level of education. For each of these groups, and for each year from 1983 to 2014, we compute the cross-sectional mean, variance, and skewness of yearly changes in non-financial log disposable income. Then, we regress the time-series of these three moments on stock market returns for each group. The coefficients of these regressions measure countercyclical risk: i.e. the level of covariance, countercyclical variance, and cyclical skewness risk for each group. For example, a positive relationship between returns and skewness indicates that workers face higher left-tail income risk when the stock market performs poorly. These measures indicate the extent to which they can hedge against increases in human capital risk by short-selling the market portfolio.

In a second step, we attribute these measures of risk to workers based on their education level and current industry of employment. Then, we regress the share of their financial wealth invested in equity on these measures of risk. We find that cyclical skewness is the only measure of countercyclical risk that predicts stock market participation and conditional and unconditional equity shares in univariate regressions. Most traditional control variables do not reduce our point estimate. One exception is education: our results remain statistically and economically significant when we introduce education fixed effects, but our point estimates gets smaller. This is expected, as more educated workers face less cyclical skewness and invest more in stocks. In fact, Alan (2012) shows that higher exposure to tail risk in recessions significantly reduces the optimal equity share of less educated workers. Importantly, our measure of cyclical skewness is based on only four recessions. Consequently, our estimates of the relationship between stock holdings and cyclical skewness are probably biased towards zero due to measurement error. This bias is likely stronger when we control for a good proxy of exposure to tail income risk such as education.

What are the implications of our findings for quantitative models? Because its effect declines with the human capital share, cyclical skewness helps explain why the equity share is not strongly decreasing with age as many life-cycle models predict. Regarding asset prices, our main finding is consistent with the intuition of models in which time-varying tail risk reduces investor's demand for equity. However, we find no relationship between cyclical skewness and equity shares in the top three deciles of the financial wealth distribution. Given the concentration of financial wealth, countercyclical labor income risk is unlikely to have a large effect on the aggregate demand for equity and asset prices. This finding does not necessarily reject the core logic of these asset pricing models: rather, it suggests that researchers need to focus on other sources of countercyclical risk. For example, wealthy entrepreneurs are reluctant to invest in the stock market because they own undiversified stakes in a single private business (Heaton and Lucas, 2000). Recent work shows that, just as workers, businesses face substantial left-tail risk during recessions (Salgado et al., 2020).

Contribution to the literature Our paper bridges a gap between two branches of the portfolio choice literature. The first strand of papers tries to rationalize the cross-section of equity holdings using calibrated portfolio choice models. An important takeaway of this literature is that earnings risk primarily reduces equity shares when it is not independently distributed from stock returns, an intuition that extends to asset prices (Krueger and Lustig, 2010). In particular, optimal equity shares are reduced by the covariance of income growth with returns (Viceira, 2001; Cocco et al., 2005; Benzoni et al., 2007) and countercyclical income risk (Storesletten et al., 2007; Lynch and Tan, 2011; Catherine, 2021).

We seek to connect this literature to reduced-form studies documenting relationships between measures of income risk and equity holdings. Several of these papers focus on variance (Betermier et al., 2012; Fagereng et al., 2018; Chang et al., 2021; d'Astous and Shore, 2021). Studies considering the role of statistical dependence focus on the covariance between income shocks and stock market returns and find no evidence of hedging (Vissing-Jorgensen, 2002; Massa and Simonov, 2006; Calvet and Sodini, 2014), Bonaparte et al. (2014) being the notable exception.

This second strand of papers is somewhat disconnected from quantitative models in which the covariance between income shocks and returns does not play a key role. Theory predicts that covariance reduces the optimal equity share but the covariance observed in the data is close to zero (Cocco et al., 2005) and is therefore unlikely to have first-order effects. By contrast, the high cyclicality of skewness in income risk has been documented in administrative data from the United States (Guvenen et al., 2014), Germany, Sweden, France (Busch et al., 2021), Canada (Bowlus et al., 2021), and Spain (Bonhomme et al., 2021). Huggett and Kaplan (2016) estimates that countercyclical risk explains two thirds of the market beta of human capital. Therefore, our contribution to the reduced-form literature is to study a measure of income risk that is a salient feature of the data and quantitatively matters in portfolio choice models.

Our paper also relates to recent studies arguing that cyclical skewness in labor income risk can explain the equity premium (Constantinides and Ghosh, 2017; Schmidt, 2016). Our findings suggest that cyclical skewness is unlikely to have large asset pricing implications. Indeed, the magnitude of our results is much smaller for households close to retirement, and we find no correlation between cyclical skewness and equity shares in the highest deciles of financial wealth. Importantly, our study focuses on labor income, but many rich households hold most of their wealth in a single private business whose fate is also exposed to cyclical skewness (Salgado et al., 2020).

A closely related strand of literature explains asset pricing puzzles by assuming rare macroeconomic disasters causing large consumption drops across the entire population (Rietz, 1988; Barro, 2006; Barro and Ursúa, 2012; Gabaix, 2012; Gourio, 2012; Martin, 2012; Wachter, 2013; Nakamura et al., 2013). Rare disaster models also rely on the coincidence of consumption disasters and stock market crashes to reduce investors' demand for equity. These models provide no testable cross-sectional predictions in our setting. Sweden did not experience any disaster over our sample period, so we cannot measure individuals' exposure to these events. Nevertheless, our findings support a key mechanism of these models: households hedge against correlated tail risks.

1 Theoretical framework

In this section, we introduce three measures of human capital exposure to stock market returns: covariance risk, countercyclical variance risk, and cyclical skewness risk. We then discuss why workers exposed to these risks should invest less in the stock market.

1.1 Income process

We assume that the log disposable income y of worker i is the sum of three components:

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

where z_{it} is a permanent component, ξ are transitory shocks that fully mean-revert within a year, and f is a deterministic function of the agent's age a and group g. We think of these groups as workers with the same level of education and industry of employment. The permanent component follows a random walk with innovation η:

$$z_{it} = z_{it-1} + \eta_{it}. \tag{2}$$

We do not assume any specific distribution for ξ or η. However, we assume that, in any given year, workers of the same group draw ξ and η from the same distributions. We denote ε_{it} the unexpected shock to disposable income, which is the sum of the permanent shock η and transitory shock ξ:

$$\varepsilon_{it} = \eta_{it} + \xi_{it}. \tag{3}$$

1.2 Moments of the income shock distribution

For each year and for each group of workers, we define the unconditional mean, variance, and skewness of income shocks as:

$$\text{Mean}(\varepsilon)_{gt} = \mu_{1,gt}(\varepsilon) = \frac{1}{N_{gt}} \sum_{i \in g} \varepsilon_{it} \tag{4}$$

$$\text{Var}(\varepsilon)_{gt} = \mu_{2,gt}(\varepsilon) = \frac{1}{N_{gt}} \sum_{i \in g} \left(\varepsilon_{it} - \mu_{1,gt}(\varepsilon)\right)^2 \tag{5}$$

$$\text{Skew}(\varepsilon)_{gt} = \mu_{3,gt}(\varepsilon) = \frac{1}{N_{gt}} \sum_{i \in g} \left(\varepsilon_{it} - \mu_{1,gt}(\varepsilon)\right)^3, \tag{6}$$

where N_{gt} is the number of workers in group g in year t. We do not standardize the third moment because standardized skewness is not a meaningful measure of risk in and of itself: a large negative standardized skewness is not worrisome if variance is small.

In the data, these moments may covary with stock market returns. For example, negative stock returns might be associated with lower earnings growth for an entire group, higher volatility of individual earnings in that group, or higher left tail risk. To capture these correlations, we construct three additional measures: covariance, countercyclical variance, and cyclical skewness. We defined these co-moments as follows:

$$\text{Covariance risk}(\varepsilon)_g = \frac{\text{cov}(\mu_{1,g}(\varepsilon), r_s)}{\text{Var}(r_s)} \tag{7}$$

$$\text{Countercyclical Variance risk}(\varepsilon)_g = -\frac{\text{cov}(\mu_{2,g}(\varepsilon), r_s)}{\text{Var}(r_s)} \tag{8}$$

$$\text{Cyclical Skewness risk}(\varepsilon)_g = \frac{\text{cov}(\mu_{3,g}(\varepsilon), r_s)}{\text{Var}(r_s)}, \tag{9}$$

where r_s denotes stock market returns. Covariance risk captures the relationship between stock returns and income shocks. Variance is countercyclical if it increases when the stock market underperforms. Finally, skewness is cyclical if left-tail income risk is higher when market returns are low.

1.3 Relation to portfolio choices

Campbell and Viceira (2002) provide a useful formula to discuss how these moments affect portfolio choices. Specifically, the optimal share of wealth invested in the stock market portfolio by an agent with constant relative risk aversion (CRRA) is:

$$\pi = \frac{\mu - r}{\gamma \sigma_s^2} + \left(\frac{\mu - r}{\gamma \sigma_s^2} - \beta_H \right) \frac{H}{W} \tag{10}$$

$$\beta_H = \frac{\text{Cov}(r_{\text{H}}, r_s)}{\sigma_s^2}, \tag{11}$$

where W is financial wealth, H the certainty equivalent of future earnings, $\text{Cov}(r_{\text{H}}, r_s)$ is the covariance between stock market and human capital returns r_H, β_H the market beta of human capital, γ the coefficient of relative risk aversion, $\mu - r$ the equity premium, and σ_s^2 the variance of stock market returns. In dollar terms, Equation (10) becomes:

$$\pi W = \frac{\mu - r}{\gamma \sigma_s^2} W + \frac{\mu - r}{\gamma \sigma_s^2} H - \beta_H H. \tag{12}$$

Optimal equity share without income risk The first term of Equation (12) is the optimal equity investment in Merton (1969)'s portfolio problem when the agent's only endowment is financial wealth W. The second term represents the effect of risk-less human capital. As Merton (1971) explains: "in computing the optimal decision rules, the individual capitalizes the lifetime flow of wage income at the market (risk-free) rate of interest and then treats the capitalized value as an addition to the current stock of wealth." Essentially, the agent's optimal stock holdings is a share $\frac{\mu-r}{\gamma\sigma_s^2}$ of his total wealth $W + H$, inclusive of human capital.

Income risk Labor income risk changes two factors of the presented portfolio choice problem. First, risk reduces the certainty equivalent value of human capital H. Second, equity holdings must be adjusted to offset the share of stocks in the replicating portfolio of human capital measured by β_H, the slope of the regression of human capital returns onto stock returns. Hence, the third term of equation (12) represents how many dollars of "stocks" are embedded in human capital. If a 10% return on the stock market portfolio translates into a 1% increase in the value of human capital ($\beta_H = .1$), then each dollar of human capital already incorporates ten cents of the stock market portfolio. As they are substitutes, dollars of stocks embedded in human capital reduce investment in stocks by the same amount.

Measuring correlated risk Empirically, β_H is difficult to estimate because returns on human capital are not observable. Indeed, the return on human capital is:

$$r_{\mathrm{H},it} = \frac{H_{it+1} - H_{it} + Y_{it}}{H_{it}}, \tag{13}$$

which cannot easily be computed, as H is not directly observable either.

The existing reduced-form literature implicitly assumes a stationary distribution of income shocks and a time-invariant discount factor. Under these assumptions, returns to human capital can be proxied by permanent income shocks: if earnings permanently drop by x%, so does the present value of human capital. Furthermore, the market beta of human capital β_H is the regression slope of permanent income shocks on stock market returns, in the spirit of Equation (7).

Unlike existing reduced-form studies, we assume that the distribution of income shocks is not stationary and non log-normal. In that case, the market beta of human capital is no longer exclusively determined by the covariance between permanent income shocks and stock returns. In particular, increases in volatility (Storesletten et al., 2007; Lynch and Tan, 2011) or higher left-tail risk (Catherine, 2021) also register as negative human capital returns. Therefore, if they coincide with or follow low stock market returns, they also contribute to a greater β_H. Empirically, we measure these two other contributors – countercyclical variance and cyclical skewness – using Equations (8) and (9). These two contributors are not of second-order importance. As mentioned previously, Huggett and Kaplan (2016) estimate that countercyclical risk explain two thirds of β_H.

1.4 Predictions

Based on this discussion, we derive two theoretical predictions to guide our empirical analysis.

Prediction 1: Higher covariance, countercyclical variance, and cyclical skewness risks reduce equity shares. Indeed a negative shock to earnings, an increase in the variance of income shocks, or a decrease in its skewness all reduce the certainty equivalent of human capital: they imply negative human capital returns ($\frac{\partial H}{\partial y} > 0$, $\frac{\partial H}{\partial \text{Var}} < 0$, and $\frac{\partial H}{\partial \text{Skew}} > 0$). Consequently, higher covariance, countercyclical variance, and cyclical skewness risks increase the linear relationship between human capital and stock market returns β_H and should unambiguously be associated with lower equity shares.

Prediction 2: Portfolio effects to human capital share of total wealth. This directly follows from Equation (10) and the decline of $\frac{H}{W}$ over the life-cycle.

By contrast, there is no clear theoretical prediction regarding the effect of unconditional variance and skewness. Neither variance nor skewness directly enters Equation (10). However, they affect the human capital-to-wealth ratio $\frac{H}{W}$. Indeed, all else equal, future income streams are less valuable if they are more volatile. But the relationship between the optimal equity share and $\frac{H}{W}$ is ambiguous. If $\frac{\mu-r}{\gamma\sigma_s^2} > \beta_H$, higher variance should be associated with a lower equity share because it reduces the weight of "bond-like" human capital in the worker's overall portfolio. On the other hand, if $\frac{\mu-r}{\gamma\sigma_s^2} < \beta_H$, workers facing greater variance would have higher equity shares because variance reduces the weight of "stock-like" human capital in their overall portfolio. Similarly, because unconditional skewness increases the value of human capital, higher skewness increases the optimal equity share if human capital is "bond-like" ($\frac{\mu-r}{\gamma\sigma_s^2} > \beta_H$). Overall, we do not have unambiguous predictions regarding the effects of unconditional variance and skewness on equity shares. Cocco et al. (2005) provide an example in which human capital has a market beta of zero and in which higher idiosyncratic volatility reduces the optimal equity share. By contrast, in Benzoni et al. (2007), because cointegration between the stock and labor markets implies a larger beta for human capital, the optimal equity share increases with unconditional volatility.[1] For similar reasons, optimal

[1]Their Figure 9 shows that a higher variance of permanent idiosyncratic shocks increases the demand for equity for workers below age 50. In their model, cointegration between the labor and stock markets makes human capital stock-like for workers who are more than a decade away from retirement. For them, $\beta_H > \frac{\mu-r}{\gamma\sigma_s^2}$, and therefore higher volatility reduces the

equity shares also increase with unconditional volatility in Catherine (2021).

2 Data

2.1 Swedish Wealth and Income Registry

The Swedish Wealth and Income Registry is an administrative panel of households. Because Swedes paid taxes on wealth until 2007, the national Statistics Central Bureau (SCB) had a parliamentary mandate to collect highly detailed information on every resident's income and wealth. We observe households' balance sheet from 1999 to 2007 and income data from 1982 to 2015.

We observe individual-level data at the end of each year. Demographic information includes age, gender, marital status, nationality, birthplace, and education level. For labor income, the database reports gross labor income, business sector, unemployment benefits, and pensions. The disaggregated wealth data contains the assets owned worldwide by each resident on December 31, including bank accounts, mutual funds, and holdings of stocks, bonds, derivatives, and debts.

This comprehensive dataset offers significant advantages for our study. The main advantage is the ability to observe labor income trajectories for millions of workers over a several decades, which allows us to compute the cross-sectional skewness of income shocks for many sub-groups of the population. The size of our dataset is a critical advantage because higher moments can be highly sensitive to outliers. Similar datasets such as the US Social Security Master

optimal equity share by reducing H. Older workers have shorter horizons on the labor market and are not that exposed to cointegration. For them, $\beta_H < \frac{\mu - r}{\gamma \sigma_s^2}$ and a reduction of H caused by higher idiosyncratic variance reduces their equity share. As the Panel B of their Figure 5 illustrates, the correlation between stock returns and returns to human capital drops rapidly after age 50.

Earnings File (MEF) exist in other countries but, in general, do not include important demographic variables such as education and are restricted to wage income. Our ability to observe government transfers allows us to take into account the social safety net when measuring income tail risk. More importantly, other administrative panel data such as the MEF are not matched with portfolio data. The Swedish portfolio data is also significantly better than surveys used in other studies. For example, the US Survey of Consumer Finances does not provide detailed holdings on each asset, and many non-responses are imputed. Compared to the SCF data, the Swedish data covers accurate individual asset holdings, such as stocks and funds.

2.2 Portfolios and returns

Returns The risk-free rate is represented by the monthly average yield on the one-month Swedish Treasury bill. We use the All Country World Index (henceforth 'world index') compiled by Morgan Stanley Capital International (MSCI) in US dollars as our proxy for the stock market portfolio. As Sweden is a small and open economy, many funds specialize in investing in the global market. The local market index is closely correlated with the global one.

Portfolios We focus on holdings of cash and risky assets, excluding defined contribution pension accounts. Cash consists of bank account balances and Swedish money market funds. The risky portfolio contains risky financial assets that are directly held stocks and risky mutual funds. Within the financial portfolio, the average participant has a risky share of 40%, owns 4 different mutual funds, and directly invests in 2 or 3 firms. These estimates are similar to the average number of stocks in U.S. household portfolios. The vast majority of risky asset participants (90%) hold mutual funds, while 60% of them own stocks directly. For every individual, the complete portfolio consists of the risky portfolio and cash. The risky share is the weight of the risky portfolio in the complete portfolio. Market participants have strictly positive risky shares. Financial wealth is defined as the sum of cash, stocks, funds, bonds, derivatives, capital insurance, and other financial wealth.

2.3 Income risk

Our measures of countercyclical income risk are built in three steps. First, we sort workers by education$\times$industry group. Second, we compute the first three moments of the distribution of income shocks for each group and year. Third, for each group, we regress these moments on stock market excess returns.

Education-industry groups We sort workers using 71 industry codes and five levels of academic achievements: 1) high school dropouts, 2) high school, 3) college (Bachelor), 4) holders of Master degrees, and 5) Doctorates. Because measuring higher moments in small samples is challenging, we ignore groups in years for which we have fewer than 100 observations or less than 10 years of data. We end up with 321 groups. We allow workers to move from one group to another, which mostly happens when they switch employment. To measure labor income risk, we further restrict our data in several ways. We exclude students, retirees, and individuals for which industry of employment or education is missing. We remove observations for which annual disposable income is below 1,000kr, drop the top 0.01% of earners, and only keep workers between age 20 and 64. For each individual, we drop the first and last year of earnings.

Moments In a second step, we compute the mean, variance, and skewness of log disposable income growth rates for each year from 1983 to 2014 and each group of workers. Disposable income is the sum of all non-financial sources of income, including social transfers. We deflate this variable using the CPI index of 2009.

We assume that workers correctly anticipate the expected growth of their log disposable income conditional on their group and age. Therefore, we start by regressing yearly changes in log disposable income on a series of age dummies. We estimate an OLS regression for each of our 321 industry×education groups, which captures the heterogeneity in life-cycle profiles of earnings across groups. Specifically, we estimate:

$$y_{it} - y_{it-1} = \dot{f}(a_{it-1}, g_{it-1}) + \hat{\varepsilon}_{it} \tag{14}$$

where $\dot{f}(a, g)$ is a third-order polynomial estimated for each group and which captures expected growth rate of earnings. We use the residuals of these regressions $\hat{\varepsilon}_{it}$ as our empirical measure of ε_{it}: the unexpected change in log disposable income.

Finally, for each year and each industry×education group, we compute the mean, variance, and skewness of the distribution of observed shocks $\hat{\varepsilon}_{it}$ using Equations (4)-(6). Overall, our methodology largely follows Guvenen et al. (2014)'s study of US workers with two differences. First, we compute cross-sectional moments within industry and educational groups, whereas these authors pool all prime age male workers. Second, we use disposable income rather than labor income, which is a better measure of what households can use for consumption. When we compute these moments using pre-tax earnings, we find skewness to vary over the business cycle in ways quantitatively similar to the United States. However, magnitudes are smaller for disposable income because of redistribution, unemployment insurance, and progressive taxation.

Cyclicality of moments In a third step, we estimate the cyclicality of each moment by regressing its time-series on contemporaneous and lagged yearly stock market returns. Denoting μ_n the n-th moment of income shock of $n \in \{1, 2, 3\}$, for each group of workers g, we estimate:

$$\mu_{n,gt} = \beta_{n,1,g} \cdot r_{s,t} + \beta_{n,2,g} \cdot r_{s,t-1} + u_{n,gt}. \tag{15}$$

We define the cyclicality of the n-th moment as $\beta_{n,1,g} + \beta_{n,2,g}$. We call Covariance the cyclicality of the first moment, Countercyclical variance the negative of the cyclicality of the second moment, and Cyclical Skewness the cyclicality of the third moment.

We include lag returns on the right-hand side because the stock market may react faster to economic news than the labor market. Indeed, we find that, in contrast to the US, the Swedish labor market tends to follow trends in the world stock market with a one year lag. For example, our economy-wide measure of skewness drops from .06 in 2007 to -.08 in 2008 and -.34 in 2009. During that recession, left-tail income risk peaked one year after the stock market crash. From an economic point of view, it makes sense that asset prices react to a change in economic risk before that risk actually materializes. From the point of view of investors, the fact that negative stock market returns precede higher labor income risk is sufficient to increase the covariance between stock and human capital returns. Indeed, news that the distribution of income shocks will worsen in the coming year immediately reduces the present value of human capital.

2.4 Human capital

Equation (10) tells us that the effect of countercyclical income risk depends on the relative importance of human capital relative to wealth. To test this prediction, we build a simple measure of human capital, defined as the present value of expected future earnings. We compute human capital as:

$$H_{it} = \sum_{k=0}^{T_i} s_{it} \frac{E[Y_{ik}]}{(1+r)^{k-t}}, \tag{16}$$

where T_i denotes the number of years before worker i retires, which we assume to be at age 64, and s_{it} is his survival probability up to year k. Survival probabilities are imputed from life tables computed by the Bureau of Statistics Sweden. Future expected earnings are determined by current earnings and age/group specific growth rates from our estimate of equation (14). Our definition of human capital does not include future pensions, which are not exposed to similar risk. Following Calvet et al. (2019), we discount future earnings at $r = 4.1\%$.

2.5 Summary statistics

Table 1 reports summary statistics for all workers (first two columns) and risky asset market participants (last two columns) at the end of 2003. The market participants do not differ from non-participants in terms of age, sex and family size. Participants have slightly higher education level compared to non-participants and are relatively wealthier. We also report summary statistics regarding our income shock moments and co-moments.

Table 1: **Summary Statistics**

	All		Participants	
	Mean	Std Deviation	Mean	Std Deviation
Income				
Non financial disposable income ($)	27,683	13,045	29,934	14,241
Entrepreneur (%)	9.13%		10.75%	
Variance	0.081	0.036	0.081	0.036
Skewness	-0.007	0.008	-0.007	0.008
Covariance	-0.013	0.024	-0.013	0.024
Countercyclical Variance	0.013	0.018	0.013	0.018
Cyclical skewness	0.023	0.029	0.021	0.029
Wealth				
Financial wealth ($)	21,437	68,816	33,497	87,873
Real estate wealth ($)	77,496	153,831	100,107	182,144
Gross wealth ($)	99,163	193,402	133,984	233,699
Debt ($)	40,177	71,828	44,772	81,714
Equity share	0.25	0.31	0.43	0.29
Demographic characteristics				
Age	43.43	11.65	44.73	11.56
Sex	1.48	0.50	1.46	0.50
High school dummy	0.84	0.37	0.87	0.34
Post-high school dummy	0.36	0.48	0.42	0.49
Immigration dummy	0.12	0.33	0.07	0.26
Family size	2.60	1.41	2.58	1.35
Observations	3,775,902	3,775,902	2,143,906	2,143,906

This table reports the main income, financial, and demographic characteristics of all Swedish population members (columns 1 and 2) and market participants (columns 3 and 4) at the end of 2003. Financial wealth consists of cash, direct stock holding, fund holding, bond holding, derivatives, capital insurance, and other financial wealth. Gross wealth is the sum of financial and real estate wealths. Income is inflation adjusted, using the 2009 CPI index of 2009.

3 Income risk and equity investment

3.1 Group-level analysis

Figure 1 offers an overview of the relationship between cyclical skewness and stock holdings at the group level, measured by its average equity share (left panel), its participation rate (center), and the average equity share of stock market participants (right). Clearly, groups of workers facing higher cyclical skewness invest less in stocks.

Figure 1: **Cyclical skewness and equity holdings**

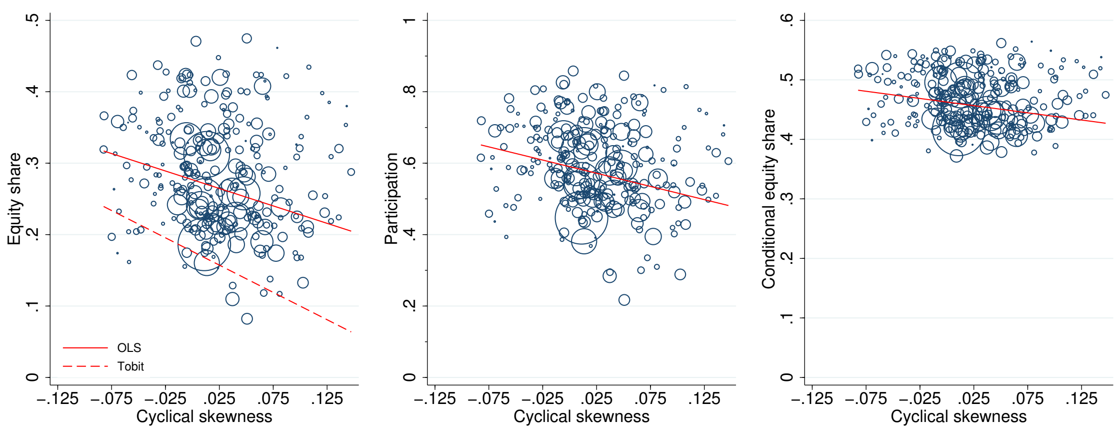

This figure reports the relationship between our measure of cyclical skewness, estimated for each education×industry group, and the average equity share of these groups (left panel), their participation rate (center), and the average equity share of participants (right). Circles reflect group size. Red lines represent OLS regressions weighted by group size. On the left panel, the red dashed line adjustment displays the predicted value of a Tobit regression with a zero lower bound, which are obtained after collapsing the data by group and participation status.

Do we observe similar patterns for other co-moments? Columns (2) and (3) of Table 2 show that our measures of covariance risk and countercyclical volatility are not correlated with equity shares. Theoretically, these moments are important for optimal portfolios, but previous studies have already argued that, in the overall population, the correlation between individual income growth and stock market returns is close to zero (Cocco et al., 2005), and the variance of income shocks is not cyclical (Guvenen et al., 2014). Heterogeneity across different groups of workers does not seem to matter. After all, workers in groups with the highest covariance risk (≈ 0.075) expect their earnings to fall by only 3% when the stock market loses 40%. The lack of correlation between equity share and covariance risk is consistent with previous work on Swedish data (Calvet and Sodini, 2014). In column (4), we include all three co-moments, and the coefficient for cyclical skewness remains significant and is quantitatively close to that of the univariate regression.

Table 2: **Equity share and countercyclical risk at the group level**

	(1)	(2)	(3)	(4)
Cyclical skewness	-0.311			-0.317
	(-3.41)			(-3.12)
Countercyclical variance		-0.348		-0.269
		(-1.97)		(-1.44)
Covariance			-0.154	0.137
			(-1.06)	(0.83)
Observations	321	321	321	321
Adjusted R^2	0.032	0.009	0.000	0.033

This table reports the result of tobit regressions of the equity share on measures of countercyclical income risk. All variables are collapsed by industry$\times$education group. Regressions use the number of observations per group as analytical weights. T-stats are reported in parenthesis.

3.2 Micro-level analysis

In the rest of this section, we run micro-level regressions where the left-hand side variable is the equity share, a stock market participation dummy, or the risky share of participants. The right-hand side variables are our measures of income risk co-moments as well as year fixed effects and demographic and economic control variables. We estimate equations of the type:

$$\begin{aligned} \text{Stock Holdings}_{it} = \;& b_1 \cdot \text{Covariance}_{g(it)} + b_2 \cdot \text{Countercyclical variance}_{g(it)} \\ & + b_3 \cdot \text{Cyclical skewness}_{g(it)} + b_c \cdot \text{Controls}_{it} + v_t + \varepsilon_{it}. \end{aligned} \tag{17}$$

Our control variables include the average variance and unscaled skewness of income growth (see Equations (5)-(6)), age, gender, household size, and dummy variables identifying entrepreneurs and immigrants. We also control for the composition of workers' overall endowment by including the value of human capital, real-estate, financial assets, and debt. We assume that workers have constant relative risk aversion and therefore scale these variables by total wealth, inclusive of human capital. Indeed, with CRRA preferences, the optimal equity share is a function of the wealth composition of the agent in relative terms, as in Equation (10). As portfolio weights sum to one, we omit one of the components of wealth: the human capital share. We examine the role of financial wealth, in absolute terms, in Section 5.2.

Table 3: **Equity share and countercyclical income risk**

	(1)	(2)	(3)	(4)	(5)
Cyclical skewness	-1.113			-1.366	-0.320
	(-2.92)			(-4.62)	(-2.09)
Contercyclical variance		-0.647		-0.781	0.434
		(-0.84)		(-1.46)	(1.48)
Covariance			-0.517	-0.430	0.655
			(-0.70)	(-0.91)	(2.10)
Skewness				-4.310	-3.572
				(-3.97)	(-6.72)
Variance				0.826	-0.407
				(3.15)	(-2.92)
Age				-0.004	-0.002
				(-12.38)	(-10.98)
Sex				-0.028	-0.033
				(-2.96)	(-7.76)
Immigrant				-0.206	-0.203
				(-29.28)	(-34.62)
Household size				0.013	0.011
				(12.73)	(17.35)
Entrepreneur				-0.034	-0.022
				(-4.42)	(-4.60)
Financial/Total Wealth				1.660	1.609
				(51.60)	(49.97)
Real Estate/Total Wealth				0.512	0.483
				(27.25)	(32.40)
Debt/Total Wealth				-0.219	-0.237
				(-13.36)	(-13.08)
Education FE					Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	32,934,044	32,934,044	32,934,044	32,933,774	32,933,774
Pseudo R^2	0.006	0.004	0.004	0.112	0.133

This table reports the result of tobit regressions of the equity share on measures of countercyclical income risk, controlling for worker and households characteristics. T-statistics reported in parenthesis are clustered by industry$\times$education group.

3.2.1 Equity share

Table 3 reports the results of Tobit regressions in which the dependent variable is the unconditional equity share. First, we regress the unconditional equity share against each co-moment separately. As reported in columns (1)-(3), only cyclical skewness is significantly correlated with equity shares. In column (4), we include all labor income risk moments, including unconditional variance and skewness, in the regression as well as controls, except for education. Neither the point estimate nor the significance of cyclical skewness falls substantially. However, including education dummies in column (5) causes our main coefficient of interest to drop from -1.614 to -.425.

There are several possible explanations. First, education can reduce equity shares for reasons unrelated to countercyclical income risk, such as financial literacy. In that case, it is important to control for education to obtain the right estimate of the relationship between cyclical skewness and equity shares. Second, education can increase equity shares because it reduces countercyclical income risk. Education significantly increases re-employment rates of the unemployed (Riddell and Song, 2011), and Alan (2012) shows that, in a calibrated model, differential exposure to countercyclical tail risk partially explains portfolio differences between education groups. In principle, this should not be a concern because we control for measures of countercyclical income risk. In practice, these variables are measured with error, as they are based on a handful of recessions. Finally, we would expect the correlation between cyclical skewness and equity shares to be lower within education groups if workers with higher risk aversion choose industries with lower income risk (Bonin et al., 2007; Fouarge et al., 2014). If risk aversion plays an important role in determining an individual's industry of employment but not his education level, then endogeneity could be more problematic within education groups, and introducing education fixed effects would bias our results towards zero. In terms of economic magnitudes, a two-standard-deviation change in cyclical skewness reduces equity shares by 2 to 8 percentage points.

Interestingly, without education fixed effects, there is a positive and statistically significant relationship between equity shares and the variance of income shocks. This relationship becomes negative and significant with education fixed effects. In Sweden, the variance of income shock increases with education (Calvet et al., 2019) and so does stock market participation. This generates a spurious correlation between variance and equity holdings that disappears when we control for education.

3.2.2 Participation and conditional equity share

As reported in Table 4, the effect of cyclical skewness on conditional equity shares is three times smaller than on the unconditional equity shares. Hence, our findings appear to be mostly driven by the extensive margin: the decision not to participate. This would be consistent with a model with fixed stock market participation costs. In theory, cyclical skewness disproportionately affects workers with low financial wealth. For these workers, the presence of a fixed cost creates a discontinuity in the optimal investment policy: when their optimal equity share falls to the point at which holding stocks is not worth paying the participation cost, they choose not to participate at all. Panel A reports the results of OLS regressions where the dependent variable is a participation dummy, which confirm the existence of a strong relationship between cyclical skewness and participation.

Table 4: **Participation and conditional equity shares**

Panel A. Stock market participation

	(1)	(2)	(3)	(4)	(5)
Cyclical skewness	-0.921			-1.185	-0.346
	(-2.69)			(-4.96)	(-2.54)
Contercyclical variance		-0.471		-0.505	0.431
		(-0.67)		(-1.15)	(1.68)
Covariance			-0.377	-0.312	0.584
			(-0.59)	(-0.79)	(2.07)
Demographics				Yes	Yes
Wealth composition				Yes	Yes
Education FE					Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	32,934,044	32,934,044	32,934,044	32,933,774	32,933,774
Adjusted R^2	0.005	0.002	0.002	0.127	0.148

Panel B. Conditional equity share

	(1)	(2)	(3)	(4)	(5)
Cyclical skewness	-0.354			-0.435	-0.036
	(-3.26)			(-3.50)	(-0.57)
Contercyclical variance		-0.295		-0.481	0.020
		(-1.32)		(-2.32)	(0.20)
Covariance			-0.224	-0.210	0.188
			(-1.06)	(-1.15)	(1.86)
			(1.03)	(1.34)	(3.62)
Demographics				Yes	Yes
Wealth composition				Yes	Yes
Education FE					Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	19,038,061	19,038,061	19,038,061	19,038,005	19,038,005
Adjusted R^2	0.027	0.026	0.026	0.084	0.096

This table reports the results of OLS regressions in which the dependent variable is a stock market participation dummy (Panel A) or the equity share of market participants (Panel B). The explanatory variables are measures of countercyclical income risk. Demographic controls include age, sex, household size, and dummies identifying immigrants and entrepreneurs. Wealth composition variables control for the share of human capital, financial wealth, real-estate, and debt in total wealth. Education fixed effects control for five educational levels. T-statistics reported in parenthesis are clustered by industry×education group.

3.3 Economic interpretation

We have documented a negative relationship between cyclical skewness and stock holdings. Is this relationship evidence of hedging?

3.3.1 Portfolio differences between identical twins

A possible interpretation of our findings is that employment and stock holdings are jointly determined by omitted variables such as risk preferences, economic literacy or general intelligence. To address this concern, we study differences in equity shares within pairs of identical twins. Identical twins share their genes and early lives, which means that they are highly similar when their first industries of employment are determined. We refer to Calvet and Sodini (2014) for a detailed discussion of the Swedish twin sample and the use of twin pairs as an empirical strategy to study portfolio choices.

To study differences within twin pairs, we estimate the following Tobit model:

$$\begin{aligned} \Delta\pi^*_{jt} = \;& b_1 \cdot \Delta\text{Covariance}_{jt} + b_2 \cdot \Delta\text{Countercyclical variance}_{jt} \\ & + b_3 \cdot \Delta\text{Cyclical skewness}_{jt} + b_c \cdot \Delta\text{Controls}_{jt} + u + \varepsilon_{jt} \end{aligned} \tag{18}$$

$$\begin{aligned} \Delta\pi_{jt} = \;& 1 - \pi_{1,jt} && \text{if } \Delta\pi^*_{jt} > 1 - \pi_{1,jt} \\ \Delta\pi_{jt} = \;& \Delta\pi^*_{jt} && \text{if } -\pi_{1,jt} < \Delta\pi^*_{jt} < 1 - \pi_{1,jt} \\ \Delta\pi_{jt} = \;& -\pi_{1,jt} && \text{if } \Delta\pi^*_{jt} < -\pi_{1,jt}, \end{aligned}$$

where $\Delta x_{jt} = x_{2,jt} - x_{1,jt}$ is the difference in variable x between two identical twins in year t.

Table 5 shows that our findings are robust. In particular, columns (5) from Tables 3 and 5 are the most directly comparable and show that the coefficients for cyclical skewness are quantitatively similar, whereas none of the coefficients associated with other measures of income risk are close to being statistically significant. This confirms that cyclical skewness is remarkable as a measure of income risk in its ability to predict equity holdings and that our findings would be robust if we controlled for a host of potential latent variables.

Table 5: **Equity shares within identical twin-year pairs**

	(1)	(2)	(3)	(4)	(5)
ΔCyclical skewness	-0.311			-0.464	-0.385
	(-1.80)			(-2.44)	(-2.00)
ΔCountercyclical variance		0.129		0.297	0.422
		(0.46)		(0.92)	(1.30)
ΔCovariance			0.108	-0.025	0.121
			(0.50)	(-0.09)	(0.45)
ΔSkewness				-0.677	-0.626
				(-1.21)	(-1.12)
ΔVariance				-0.101	-0.212
				(-0.52)	(-1.08)
ΔHousehold size				-0.001	-0.001
				(-0.37)	(-0.36)
ΔEntrepreneur				-0.020	-0.019
				(-1.30)	(-1.24)
ΔFinancial/Total Wealth				1.187	1.189
				(19.29)	(19.35)
ΔReal Estate/Total Wealth				0.297	0.298
				(8.57)	(8.62)
ΔDebt/ Total Wealth				-0.117	-0.119
				(-4.00)	(-4.04)
ΔEducation FE					Yes
Observations	34,460	34,460	34,460	34,460	34,460
Pseudo R^2	0.000	0.000	0.000	0.035	0.035

This table reports the result of tobit regressions where the dependent variable is the within twin-year difference in equity shares and the independent variables are the within twin-year differences in measures of income risk, composition of total wealth, and households characteristics. T-statistics reported in parenthesis are clustered by industry$\times$education group pairs.

3.3.2 Moderating role of the human capital share

To further support the causal interpretation of our findings, we now turn to the moderating role of the share of human capital in households' total wealth. Equation (10) shows that, in theory, the effect of countercyclical income risk depends on the importance of human capital relative to other forms of wealth. Specifically, the hedging demand for equity should be proportional to the human capital-to-wealth ratio. In particular, households for which human capital represents a negligible share of future consumption do no need to hedge against labor income risk and thus constitute a theoretical placebo group.

To test these predictions, we cut our sample into five groups based on the share of human capital in an individual's total wealth. Then, we run Tobit regressions of the equity shares within each subsample and report the coefficients for cyclical skewness in Figure 2.

Figure 2: **Effect of cyclical skewness by human capital share level**

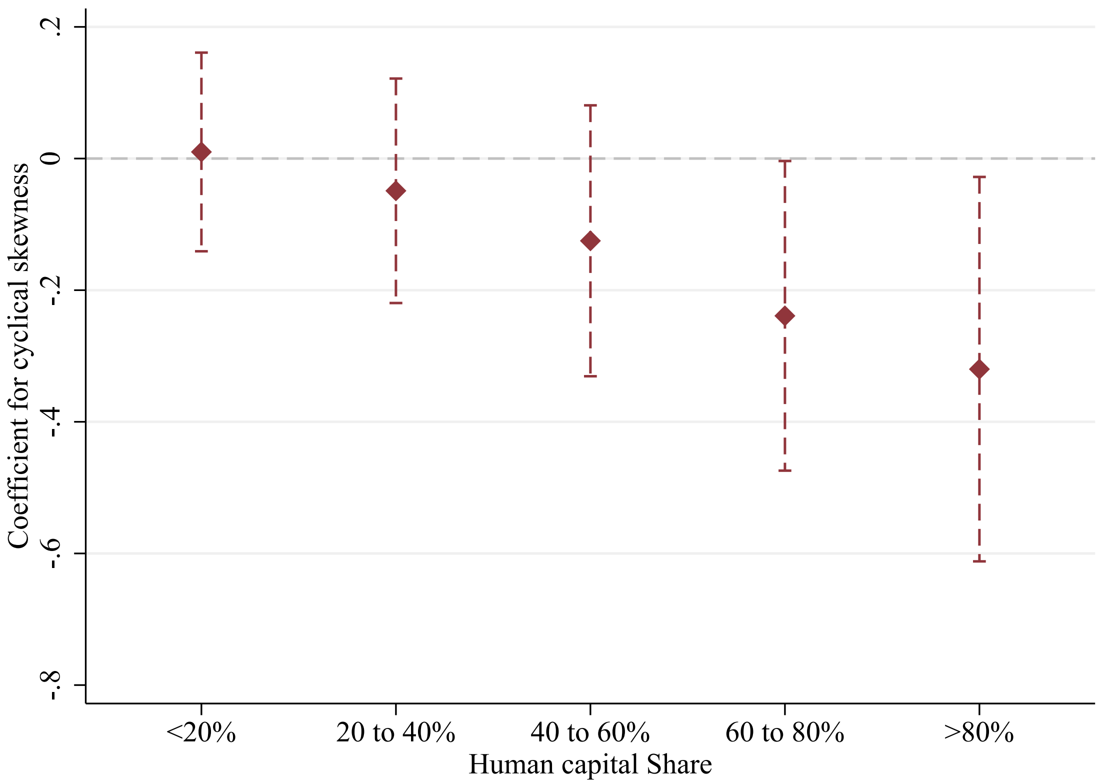

This figure reports the regression coefficients of the equity share on cyclical skewness when we run the same Tobit regressions as in column (5) of Table 3 in sub-samples of workers with different levels of human capital-to-wealth ratios. Dashed lines represent 95% confidence intervals. Standard errors are clustered by industry$\times$education group.

We find that, when human capital represents less than 20% of total wealth, cyclical skewness and equity share are not correlated. In fact, our point estimate is close to zero with a small standard error. We think of this group as our placebo group. As predicted by Equation (10), the relationship between cyclical skewness and the equity share becomes steeper and statistically significant as we move to subsamples with higher human capital shares. It is the strongest for households whose human capital-to-wealth ratio exceeds 80%, which represents 65% of our sample.

Figure 2 is key to the economic interpretation of our results. If the relationship between cyclical skewness and equity holdings is explained by hedging motives, then the strength of this relationship should be proportional to the relative weight of human capital. This is what we observe in the data.

4 Consumption risk and equity investment

For our findings to be consistent with theory, the relationship between cyclical skewness and equity holdings must be an increasing function of the human capital-to-wealth ratio and must be driven by permanent income shocks. In this section, we rely on these economic intuitions to build measures of labor-market-implied consumption risk and test whether these measures predict portfolios.

4.1 Theory

Our measure of labor-market-implied consumption risk can be understood by rearranging Equation (10) as:

$$\pi W = \left(\frac{\mu - r}{\gamma \sigma_s^2} - \beta_C \right) (W + H) \tag{19}$$

$$\beta_C = \frac{\beta_H H}{W + H} \tag{20}$$

where β_C is the market beta of consumption implied by income shocks. β_C represents the linear relationship between the percentage change in the certainty equivalent of total wealth $W + H$ and stock market returns implied by the market beta of human capital. If the agent has isoelastic utility, his consumption at a given age is a linear function of total wealth $H+W$. Therefore, if future labor earnings represent 50% of total wealth, an unexpected 1% shock to human capital translates into a .5% shock to total wealth, and therefore to lifetime consumption. In other words, β_C is a measure of a worker's exposure to the stock market through labor income risk adjusted for his reliance on labor to finance his lifetime consumption.

4.2 Labor-market-implied consumption risk

We construct labor-market-implied consumption risk measures in two steps. First, we build measure of permanent income risk. Second, we adjust these measures to account for the diversification effect of other forms of wealth.

4.2.1 Permanent income shocks

Because we cannot observe the permanent and transitory components of earnings in the data, we build an approximate measure of permanent disposable income. Previous studies have used a rolling average of log income to build such proxy (Bonaparte et al., 2014; Kopczuk et al., 2010) by computing the average log disposable income over a k-year window. More recently, Busch et al. (2021) also use this method to decompose cyclical skewness. We follow this methodology using a 3-year window and show this proxy skims away most of transitory shocks from our measures of countercyclical income risk. Our proxy for log permanent income is:

$$\hat{z}_{it} = \frac{y_{it-1} + y_{it} + y_{it+1}}{3}. \tag{21}$$

Assuming the income process of Section 1.1, our empirical measure of permanent income shocks is therefore:

$$\hat{\eta}_{it} = \hat{z}_{it} - \hat{z}_{it-1} = \frac{\eta_{it+1} + \eta_{it} + \eta_{it-1} + \epsilon_{it+1} - \epsilon_{it-2}}{3}. \tag{22}$$

If η_{it-1} to η_{it+1}, ϵ_{it-2}, and ϵ_{it+1} are independently distributed, and μ_n denotes the n-th moment of a distribution, then for $n \leq 3$:

$$\mu_{n,t}(\hat{\eta}) = \frac{\mu_{n,t+1}(\eta) + \mu_{n,t}(\eta) + \mu_{n,t-1}(\eta) + \mu_{n,t+1}(\epsilon) - \mu_{n,t-2}(\epsilon)}{3^n} \tag{23}$$

As before, we assume a linear relationship between cross-sectional moments and contemporaneous and lag market returns, that is, for permanent shocks:

$$\mu_{n,t}(\eta) = \beta_{\eta,n,1} \cdot r_{s,t} + \beta_{\eta,n,2} \cdot r_{s,t} + u_{\eta,n,t}, \tag{24}$$

and similarly for transitory shocks. Assuming that the distribution of income shocks does not depend on returns from two or more years prior, we can replace all moments in Equation (23) as linear functions of returns to obtain:

$$\mu_{n,t}(\hat{\eta}) = \frac{(\beta_{\eta,n,1} + \beta_{\eta,n,2} + \beta_{\xi,n,2}) \cdot r_{s,t} + (\beta_{\eta,n,1} + \beta_{\eta,n,2}) \cdot r_{s,t-1}}{3^n} + v_{\eta,n,t}, \tag{25}$$

where $v_{\eta,n,t}$ is a random variable independent of contemporaneous and lag returns. Equation (25) shows that if we measure permanent cyclical skewness by regressing our moments of $\hat{\eta}$ on market returns and lag market returns, we need to multiply the resulting coefficients by $3^3/2$. More generally, our measure of the cyclicality of the n-th moment will be:

$$\frac{3^n}{2}\left(\beta_{\hat{\eta},n,1} + \beta_{\hat{\eta},n,2}\right) = \beta_{\eta,n,1} + \beta_{\eta,n,2} + \frac{\beta_{\xi,n,2}}{2}. \tag{26}$$

This measure should capture 100% of the cyclicality of permanent shocks ($\beta_{\eta,n,1}+\beta_{\eta,n,2}$) and, relative to our previous countercylical risk measures, skims away roughly three quarters of transitory shocks. Indeed, the remaining term $\frac{\beta_{\xi,n,2}}{2}$ represents one quarter of transitory cyclical skewness ($\beta_{\xi,n,1}+\beta_{\xi,n,2}$) provided that $\beta_{\xi,n,1} \approx \beta_{\xi,n,2}$.

Transitory income shocks are the difference between total and permanent shocks.

4.2.2 Countercyclical consumption risk

Our measures of permanent income risk can be used to examine the role of consumption risk. Indeed, a permanent drop in income reduces lifetime consumption in proportion to the share of human capital in total wealth. This shock to consumption might be attenuated or amplified by returns on other assets, depending on a worker's portfolio. For this reason, actual consumption risk is an endogenous regressor: among other things, it crucially depends on equity holdings. Therefore, we build our measures of consumption risk under the assumption that wealth is entirely invested in the risk-free asset. In that case, the n-th moment of the distribution of lifetime consumption shocks is:

$$\mu_{n,it}(\dot{c}) = \left(\frac{H_{it}}{W_{it}+H_{it}}\right)^{n} \mu_{n,g(i)t}(\eta), \tag{27}$$

where $\dot{c}_{it}$ is the unexpected change in log lifetime consumption. The same logic applies to co-moments. For example, the cyclical skewness of consumption shocks would be:

$$\text{Cyclical Skewness}(\dot{c})_{it} = \left(\frac{H_{it}}{W_{it}+H_{it}}\right)^{3} \text{Cyclical Skewness}(\eta)_{g(i)} \tag{28}$$

4.3 Equity holdings

With our proxy for permanent income shocks, we can run Tobit regressions of the equity share on measures of countercyclical consumption risk. Table 6 reports our findings. The coefficients of these regressions should theoretically be higher, as consumption shocks reflect returns on an agent's entire wealth, inclusive of human capital. We find this prediction is verified in each specification. One explanation could be that we multiplied the right-hand side variables by numbers between 0 and 1, which would mechanically increase the associated coefficient. However, we note that our t-statics for cyclical skewness also increase substantially.

Table 6: **Equity share and labor-market-implied consumption risk**

	Equity share			Participation			Conditional equity share		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Cyclical skewness	-2.519	-0.440	-0.579	-2.104	-0.421	-0.546	-0.747	-0.083	-0.114
	(-8.38)	(-2.74)	(-4.66)	(-8.53)	(-2.91)	(-4.99)	(-7.53)	(-1.78)	(-3.13)
Countercyclical var.		0.456	0.027		0.350	0.130		0.153	-0.097
		(1.07)	(0.06)		(0.94)	(0.39)		(1.39)	(-0.60)
Covariance		1.243	1.735		1.100	1.592		0.373	0.512
		(4.16)	(5.71)		(4.10)	(5.59)		(4.27)	(4.96)
Demographics		Yes	Yes		Yes	Yes		Yes	Yes
Wealth composition		Yes	Yes		Yes	Yes		Yes	Yes
Education FE		Yes	Yes		Yes	Yes		Yes	Yes
Industry FE			Yes			Yes			Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	32,936,703	32,933,774	32,933,774	32,936,703	32,933,774	32,933,774	19,039,493	19,038,005	19,038,005
Pseudo R^2	0.017	0.133	0.142	0.018	0.148	0.159	0.031	0.097	0.100

This table reports the results of regressions in which the dependent variable is a measure of equity holdings and the explanatory variables are measures of countercyclical consumption risk and other worker characteristics. The dependent variable is the equity share in columns (1)-(3), a participation dummy in columns (4)-(6), and the conditional equity share in columns (7)-(9). We run Tobit regressions in columns (1)-(3) and OLS regressions in columns (4)-(9). Demographic controls include age, sex, household size, and dummies identifying immigrants and entrepreneurs. Wealth composition variables control for the share of human capital, financial wealth, real-estate, and debt in total wealth. Education fixed effects control for the five levels of educational attainment used to sort workers into industry$\times$education groups. T-statistics reported in parenthesis are clustered by industry$\times$education group.

5 Implications for quantitative models

Countercyclical income risk has received attention in the finance literature because it helps quantitative models solve important empirical puzzles. First, life-cycle models with countercyclical income risk can generate an equity share that does not decrease with age (Storesletten et al., 2007; Lynch and Tan, 2011; Catherine, 2021). Second, asset pricing models with countercyclical risk generate a higher equity premium and greater volatility (Constantinides and Ghosh, 2017; Schmidt, 2016). In this section, we expand our empirical analysis to shed light on the predictions of these models.

5.1 Portfolio choices

Portfolio choice models generally struggle to match three aspects of the data: (i) the low average equity share; (ii) the lack of a strong positive relationship between equity shares and the model main state variable, the human capital-to-wealth ratio; and (iii) an equity share that does not fall with age. Our findings show that cyclical skewness can help models match the data along these three dimensions.

We have already documented that cyclical skewness helps solve the first two issues. In particular, workers with high human capital-to-wealth ratio may be discouraged to own stocks because because cyclical skewness is a greater concern to them.

To explore how countercyclical income risk might shape the life-cycle profile of the equity share, we run our main Tobit regression by age group. Figure 3 shows that the coefficients for cyclical skewness display a U-shaped pattern, being close to zero for households in their twenties, reaching a trough in their early forties, and converging back to zero as they get closer to retirement.

Figure 3: **Cyclical skewness and equity share over the life cycle**

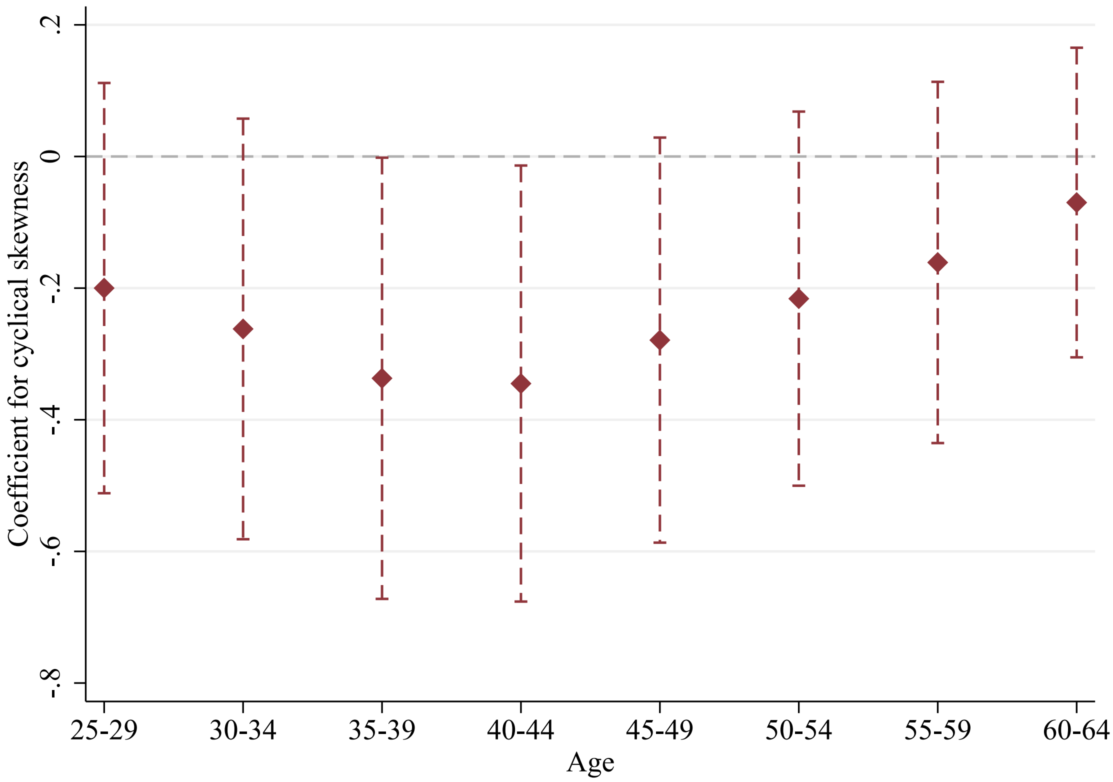

This figure reports the regression coefficients of the equity share on cyclical skewness when we run the same Tobit regressions as in column (5) of Table 3 for different age groups. Dashed lines represent 95% confidence intervals. Standard errors are clustered by industry×education group.

Catherine (2021) shows that, in a realistic life-cycle model, the effect of cyclical skewness follows this pattern[2] if households must pay a fixed cost to participate in the stock market. In the presence of a fixed cost, young households do not invest in stocks regardless of countercyclical income risk. As they start accumulating financial wealth, paying the fixed cost becomes worthwhile for workers without cyclical skewness who want to invest most of their savings

[2]See Panel B.2. of Figure 7 of Catherine (2021).

in stocks but not for those with cyclical skewness who have lower optimal equity shares. Thus, a fixed participation cost can explain why the role of cyclical skewness initially grows with age. However, at some point, the decline in the human capital-to-wealth becomes the dominant force and explains why cyclical skewness reduces stock holdings less as workers get closer to retirement.

5.2 Asset prices

Recent theoretical studies by Constantinides and Ghosh (2017) and Schmidt (2016) argue that countercyclical income risk can explain the equity premium and excess volatility puzzles. A common assumption in these papers is that consumption risk is equally distributed across households. This is true if they face the same labor income risk and have the same human capital-to-wealth ratio. Catherine (2021) shows that, in a realistic life-cycle model in which the human capital-to-wealth ratio varies greatly, countercyclical consumption risk is not equally distributed. Rather, it is concentrated among young households who have too little wealth for their portfolio choice to affect prices. For older households with substantial financial wealth, countercyclical income risk plays a lesser role in determining optimal portfolios because tail income shocks do not translate into tail consumption shocks.

To test this hypothesis, we run our main Tobit regression by decile of financial wealth. Figure 4 shows that cyclical skewness has no statistically significant effect on the equity share of households in the three highest deciles of financial wealth. In addition to being statistically insignificant, our point estimate is also close to zero and 25 times lower than in the first decile. As the top three deciles posses 88% of total financial wealth, it is unlikely that cyclical skewness could have large implications for asset prices. Fagereng et al. (2018) find that the portfolio response to their measure of uninsurable wage risk also vanishes as financial wealth is accumulated and argue that income risk is therefore unlikely to impact stock prices. Our findings complement theirs in the sense that their measure of income risk is orthogonal to stock market returns, which makes our measure a priori more likely to generate a hedging motive.

Importantly, high-wealth individuals are exposed to other forms of tail shocks during recessions: for example entrepreneurial risk. Salgado et al. (2020) document a high level of cyclical skewness in various business performance metrics, which may reduce wealthy private businesses owners demand for publicly traded stocks. Heaton and Lucas (2000) show that entrepreneurial risk affects their portfolios.

Figure 4: **Cyclical skewness and equity share by decile of financial wealth**

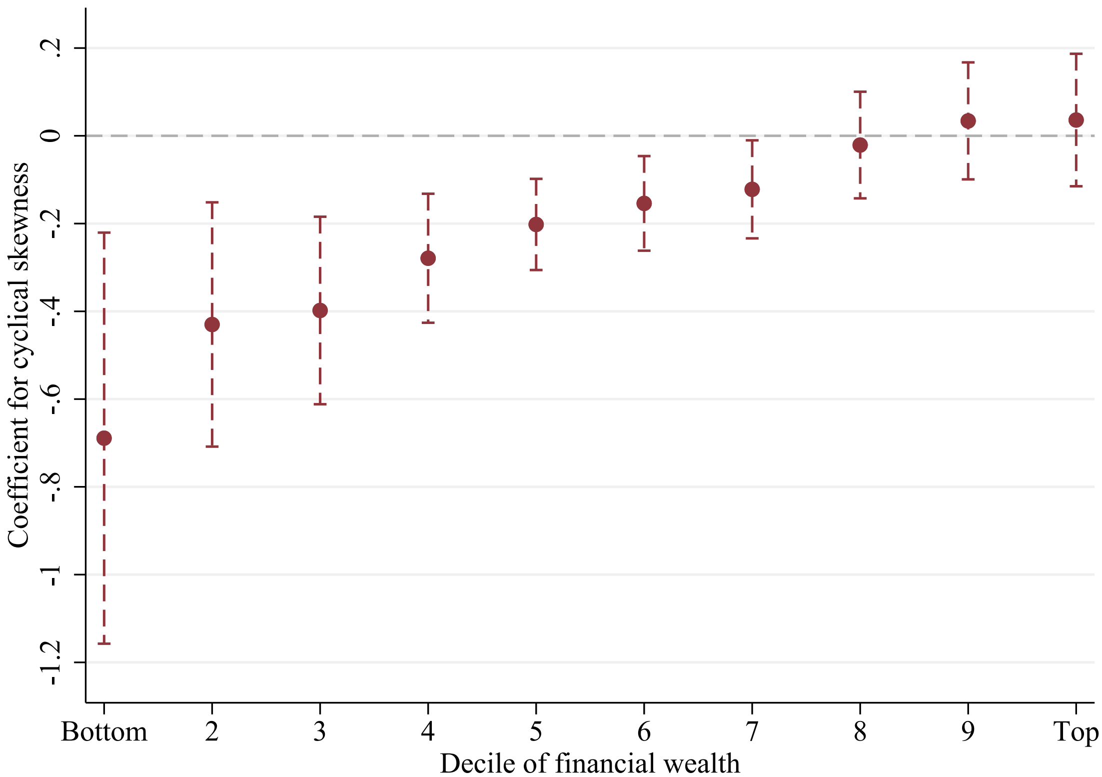

This figure reports the regression coefficients of the equity share on cyclical skewness when we run the same Tobit regressions as in column (5) of Table 3 for different deciles of financial wealth. Dashed lines represent 95% confidence intervals. Standard errors are clustered by industry$\times$education group.

6 Conclusion

In this paper, we document that workers who face more left-tail income risk following low stock market returns are less likely to participate in the stock market and, when they participate, invest less in stocks. We show that the relationship between cyclical skewness and equity shares is stronger when human capital represents a larger share of total wealth, as predicted by theory.

Cyclical skewness has been proposed as a solution to important empirical puzzles in household finance. The first one is the low stock market participation rate. We find that cyclical skewness reduces the willingness of workers to invest in stock. A second puzzle is that the share of wealth households invest in stocks does not fall with age as standard life-cycle models predict. We show that cyclical skewness reduces the stock holdings of young workers more, and thus contribute to explaining the life-cycle profile of the equity share.

On the other hand, our findings suggest that countercyclical labor income risk has no effect on the equity share of wealthy investors and is therefore unlikely to explain asset pricing puzzles on its own. Importantly, many wealthy investors are exposed to other sources of tail risk during recessions, in particular capital losses on undiversified investments in private businesses. It is therefore possible that, despite our findings, tail consumption risk is important for asset pricing.

References

Alan, Sule, "Do disaster expectations explain household portfolios?," *Quantitative Economics*, March 2012, *3* (1), 1–28.

Barro, Robert J., "Rare Disasters and Asset Markets in the Twentieth Century," *The Quarterly Journal of Economics*, 08 2006, *121* (3), 823–866.

_ **and José F. Ursúa**, "Rare Macroeconomic Disasters," *Annual Review of Economics*, 2012, *4* (1), 83–109.

Benzoni, Luca, Pierre Collin-Dufresne, and Robert S. Godlstein, "Portfolio Choice over the Life-Cycle when the Stock and Labor Markets Are Cointegrated," *Journal of Finance*, October 2007, *62* (5), 2123–2167.

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

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

Bonhomme, Stéphane, Micole De Vera, Laura Hospido, and Siq Wei, "Income Risk Inequality: Evidence from Spanish Administrative Records," Working Paper 2136, Banco de Espana October 2021.

Bonin, Holger, Thomas Dohmen, Armin Falk, David Huffman, and Uwe Sunde, "Cross-sectional earnings risk and occupational sorting: The role of risk attitudes," *Labour Economics*, 2007, *14* (6), 926–937.

Bowlus, Audra, Émilien Gouin-Bonenfant, Huju Liu, Lance Lochner, and Youngmin Park, "Four Decades of Canadian Earnings Inequality and Dynamics across Workers and Firms," Working Paper 28757, National Bureau of Economic Research May 2021.

Busch, Christopher, David Domeij, Fatih Guvenen, and Rocio Madera, "Skewed Idiosyncratic Income Risk over the Business Cycle: Sources and Insurance," *American Economic Journal: Macroeconomics*, 2021, *forthcoming.*

Calvet, Laurent E. and Paolo Sodini, "Twin Picks: Disentangling the Determinants of Risk-Taking in Household Portfolios," *Journal of Finance*, 04 2014, *69* (2), 867–906.

_ **, John Y. Campbell, Francisco J. Gomes, and Paolo Sodini**, "The Cross-Section of Household Preferences," *Working paper*, 2019.

Campbell, John Y. and Luis M. Viceira, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors* number 9780198296942. In 'OUP Catalogue.', Oxford University Press, 2002.

Catherine, Sylvain, "Countercyclical Labor Income Risk and Portfolio Choices over the Life Cycle," *Review of Financial Studies (Forthcoming)*, 2021.

Chang, Yongsung, Jay H. Hong, Marios Karabarbounis, Yicheng Wang, and Tao Zhang, "Income volatility and portfolio choices," *Review of Economic Dynamics*, 2021.

Cocco, Joao F., Francisco J. Gomes, and Pascal J. Maenhout, "Consumption and Portfolio Choice over the Life Cycle," *Review of Financial Studies*, 2005, *18* (2), 491–533.

Constantinides, George M. and Anisha Ghosh, "Asset Pricing with Countercyclical Household Consumption Risk," *The Journal of Finance*, 2017, *72* (1), 415–460.

_ **and Darrell Duffie**, "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy*, 1996, *104* (2), 219–240.

d'Astous, Philippe and Stephen Shore, "Human Capital Risk and Portfolio Choices: Evidence from University Admission Discontinuities," *Working paper*, 2021.

Ebrahimian, Mehran and Jessica Wachter, "Risks to Human Capital," March 2020, (26823).

Fagereng, Andreas, Luigi Guiso, and Luigi Pistaferri, "Portfolio Choices, Firm Shocks, and Uninsurable Wage Risk," *Review of Economic Studies*, 2018, *85* (1), 437–474.

Fouarge, Didier, Ben Kriechel, and Thomas Dohmen, "Occupational sorting of school graduates: The role of economic preferences," *Journal of Economic Behavior and Organization*, 2014, *106*, 335–351.

Gabaix, Xavier, " Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance *," *The Quarterly Journal of Economics*, 03 2012, *127* (2), 645–700.

Gourio, François, "Disaster Risk and Business Cycles," *American Economic Review*, May 2012, *102* (6), 2734–66.

Guvenen, Fatih, Serdar Ozkan, and Jae Song, "The Nature of Countercyclical Income Risk," *Journal of Political Economy*, 2014, *122* (3), 621 – 660.

Heaton, John and Deborah Lucas, "Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk," *The Journal of Finance*, 2000, *55* (3), 1163–1198.

Huggett, Mark and Greg Kaplan, "How Large is the Stock Component of Human Capital?," *Review of Economic Dynamics*, 2016, *22*, 21–51.

Kopczuk, Wojciech, Emmanuel Saez, and Jae Song, "Earnings Inequality and Mobility in the United States: Evidence from Social Security Data Since 1937*," *The Quarterly Journal of Economics*, 02 2010, *125* (1), 91–128.

Krueger, Dirk and Hanno Lustig, "When is market incompleteness irrelevant for the price of aggregate risk (and when is it not)?," *Journal of Economic Theory*, 2010, *145* (1), 1 – 41.

Lynch, Anthony W. and Sinan Tan, "Labor income dynamics at business-cycle frequencies: Implications for portfolio choice," *Journal of Financial Economics*, 2011, *101* (2), 333–359.

Mankiw, N.Gregory, "The equity premium and the concentration of aggregate shocks," *Journal of Financial Economics*, 1986, *17* (1), 211–219.

Martin, Ian W. R., "Consumption-Based Asset Pricing with Higher Cumulants," *The Review of Economic Studies*, 09 2012, *80* (2), 745–773.

Massa, Massimo and Andrei Simonov, "Hedging, Familiarity and Portfolio Choice," *Review of Financial Studies*, 2006, *19* (2), 633–685.

Merton, Robert C., "Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case," *The Review of Economics and Statistics*, August 1969, *51* (3), 247–57.

_ , "Optimum consumption and portfolio rules in a continuous-time model," *Journal of Economic Theory*, December 1971, *3* (4), 373–413.

Nakamura, Emi, Jón Steinsson, Robert Barro, and José Ursúa, "Crises and Recoveries in an Empirical Model of Consumption Disasters," *American Economic Journal: Macroeconomics*, July 2013, *5* (3), 35–74.

Paron, James D., “Heterogeneous-Agent Asset Pricing,” *Working Paper*, 2022.

Riddell, W. Craig and Xueda Song, “The impact of education on unemployment incidence and re-employment success: Evidence from the U.S. labour market,” *Labour Economics*, August 2011, *18* (4), 453–463.

Rietz, Thomas A., “The equity risk premium a solution,” *Journal of Monetary Economics*, 1988, *22* (1), 117–131.

Salgado, Sergio, Fatih Guvenen, and Nicholas Bloom, “Skewed Business Cycles,” *Working Paper*, 2020.

Schmidt, Lawrence D. W., “Climbing and Falling Off the Ladder: Asset Pricing Implicationsof Labor Market Event Risk,” *Working Paper*, 2016.

Storesletten, Kjetil, Chris Telmer, and Amir Yaron, “Asset Pricing with Idiosyncratic Risk and Overlapping Generations,” *Review of Economic Dynamics*, October 2007, *10* (4), 519–548.

Viceira, Luis M., “Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income,” *Journal of Finance*, 04 2001, *56* (2), 433–470.

Vissing-Jorgensen, Annette, “Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures,” April 2002, (8884).

Wachter, Jessica A., “Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?,” *The Journal of Finance*, 2013, *68* (3), 987–1035.