The Review of Financial Studies

Countercyclical Labor Income Risk and Portfolio Choices over the Life Cycle

Sylvain Catherine
The Wharton School, University of Pennsylvania, USA

I estimate a life cycle model of portfolio choices that incorporates the relationship between market returns and the skewness of idiosyncratic income shocks. The cyclicality of skewness can explain (a) low stock market participation among young households, (b) why the equity share of participants slightly increases until retirement, and (c) why renters invest less in stocks than do homeowners. With a relative risk aversion of 6 and yearly participation cost of \$250, the model matches the evolution of wealth, participation, and conditional equity shares over the life cycle. Nevertheless, cyclical skewness increases the equity premium by at most 0.5 percentage points. (*JEL* G11, G12, D14, D15, J24)

Received March 18, 2019; editorial decision October 12, 2021 by Editor Stijn Van Nieuwerburgh. Authors have furnished an Internet Appendix, which is available on the Oxford University Press Web site next to the link to the final published paper online.

How should human capital affect the demand for stocks? Because the correlation between stock market returns and labor income shocks is close to zero, standard life cycle models conclude that human capital is a substitute for bonds and increases the demand for equity (Viceira 2001, Cocco, Gomes, and Maenhout 2005). This conclusion creates two problems. First, as the importance of human capital declines over the life cycle, so should the share of wealth invested in equity (Jagannathan and Kocherlakota 1996).[1] Yet, in the United States, both stock market participation and the equity share of participants increase with age until retirement. Second, when human capital increases the

This paper greatly benefited from the comments of the editor Stijn Van Nieuwerburgh and two dedicated referees. I am grateful to David Thesmar, Fatih Guvenen, Amir Yaron, Francisco Gomes, Annette Vissing-Jørgensen, Jonathan Parker, Hugues Langlois, Denis Gromb, and Kim Peijnenburg for useful feedback. I am also thankful for comments from participants at the Western Finance Association conference, the Texas Finance Festival, the Northern Finance Association conference, the CEPR Household Finance conference and workshop and seminar participants at HEC Paris, the University of Minnesota, and Johns Hopkins. All errors are my own. Supplementary data can be found on *The Review of Financial Studies* web site. Send correspondence to Sylvain Catherine, scath@wharton.upenn.edu.

[1] Beside these authors, Viceira (2001), Cocco, Gomes, and Maenhout (2005), Campbell et al. (2001), Cocco (2005), Gomes and Michaelides (2005), Gomes, Kotlikoff, and Viceira (2008), and Chai et al. (2011) reach the same conclusion.

The Review of Financial Studies 00 (2021) 1–39

doi:10.1093/rfs/hhab136 Advance Access publication December 24, 2021

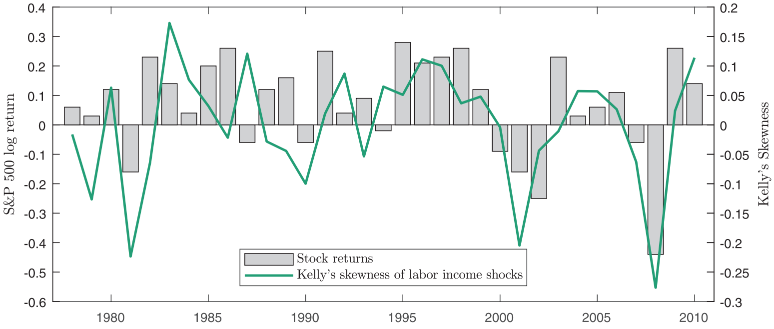

Figure 1
Skewness of income shocks and stock returns in the United States
This graph plots the evolution of the cross-sectional skewness of log wage yearly changes (line) in the United States from 1978 to 2010 and yearly S&P 500 log returns (bars). Here, skewness is defined as $\frac{(d9-d5)-(d5-d1)}{d9-d1}$, where di denotes the ith decile of the log change in wage. Those deciles were computed by Guvenen, Ozkan, and Song (2014) using Social Security panel data. Their sample is restricted to male workers between 25 and 60 years old and log wage changes are adjusted for life cycle effects. S&P 500 data are taken from Robert Shiller's website.

demand for stocks, models need unrealistically high levels of risk aversion to match the average equity share of working households.

These discrepancies are evidence that households make serious investment mistakes or that portfolio choice models do not fully capture the extent of their background risk. In particular, life cycle models ignore that the cross-sectional skewness of labor income shocks is cyclical (Guvenen, Ozkan, and Song 2014). Cyclical skewness implies that the worst human capital shocks happen during recessions, which tend to coincide with stock market crashes. As Figure 1 illustrates, negative returns on the stock market coincide with high left-tail risk on the labor market. As a consequence, workers investing in stocks take the risk of losing large portions of their earnings and savings at the same time.

In this paper, I show that introducing cyclical skewness in a life cycle model reconciles its predictions with the data and leads to more plausible estimates of preference parameters. To do so, I set up a model in which a worker with constant relative risk aversion (CRRA) invests in a risk-free asset, the stock market portfolio and housing. The main innovation is the relationship between stock market returns and the skewness of idiosyncratic labor income shocks. In a first step, I estimate the joint dynamics of labor income and stock market returns using a large set of historical moments computed by Guvenen, Ozkan, and Song (2014) using Social Security panel data. In a second step, I estimate households' risk aversion and discount factor, as well as a yearly stock market participation cost by targeting the evolution of wealth, participation and equity holdings over the life cycle in the Survey of Consumer Finances (SCF).

My main findings are as follows. First, the model matches the data well with a risk aversion of γ=6 and a participation cost of \$250. Without cyclical skewness, the same model requires a risk aversion above 12 or a participation cost above

$400 to match the average equity share and cannot explain its decline over the life cycle. Second, comparative statics show that cyclical skewness reduces the optimal equity share of young workers with little wealth from roughly 100% to 0%. Third, cyclical skewness helps explain three important aspects of the cross-section of portfolio choices: (1) equity shares are not negatively correlated with the wealth-to-earnings ratio; (2) renters have lower equity shares than homeowners; and (3) housing and equity shares are not positively correlated. Finally, cyclical skewness reduces the aggregate demand for equity by only 12% and, therefore, is unlikely to explain the equity premium as Constantinides and Ghosh (2017) suggest.

The first step of my quantitative exercise is to estimate the dynamics of labor income and yearly stock market returns. The labor income process is estimated by targeting the cross-sectional mean, variance, and skewness of idiosyncratic earnings growth for each year from 1978 to 2010. Targeting these moments for earnings growth over 1-, 3-, and 5-year periods allows me to disentangle the dynamics of transitory and persistent shocks and to estimate the persistence of the latter. I find tail shocks to be quite persistent and therefore likely to have portfolio implications. In a second step, I incorporate this income process into a life cycle model. Though the key innovation is the income dynamics, the model is also more realistic than in past studies. Indeed, the model allows for mean reversion in the persistent component of earnings and takes into account the social safety net, stock market crashes, homeownership decisions, and Social Security benefits based on an individual's lifetime earnings.

The first estimation shows that the model can closely match the evolution of wealth and equity shares over the life cycle with a discount factor of β=0.91 and a relative risk aversion of γ=7.1. By contrast, without cyclical skewness, the same model cannot match the data and generates a very high estimate of γ (12.0) and a very low estimate of β (0.67), which is consistent with previous estimates (Fagereng, Gottlieb, and Guiso 2017).[2]

In a second set of estimations, I introduce a fixed yearly participation cost and target the evolution of participation rates. With cyclical skewness, I find an estimated risk aversion of $\gamma = 6.1$, a discount rate of $\beta = .96$, and a participation cost of $250. This cost is enough to prevent the participation of young workers, because they would choose low equity shares anyway. By contrast, without cyclical skewness, young households have very high optimal equity shares, so matching their low participation rate requires a participation cost of $400. Even with this cost, the model without cyclical skewness still generates a negative relationship between age and the equity share of participants.

Policy functions show that, for households with high human-capital-to-wealth ratios, cyclical skewness reduces the optimal equity share from 100% to 0%. Even absent any participation cost, most households with less than a year of

[2] These authors estimate a similar model using Norwegian data and find γ=11 and β=0.77 with a small participation cost of $69.

salaries in wealth should barely hold any stocks. Moreover, cyclical skewness reverses the relationship between age and the equity share. While the fear of stock market crashes only reduces the equity share of young households by a dozen of percentage points when labor income shocks are normally distributed, its effect is much stronger in the presence of cyclical skewness. Tail events on the stock market and in careers matter more when they tend to coincide.

Cyclical skewness can also explain key properties of households' balance sheets. In the data, renters have lower equity share than homeowners, and, among homeowners, equity holdings are concentrated at the top of the wealth distribution. Even without fixed participation costs, the model matches these stylized facts well. However, after removing cyclical skewness, renters have higher equity shares than homeowners. Moreover, among homeowners, middle-class households have higher equity share than wealthier ones. With cyclical skewness, renters and middle-class homeowners have low equity shares because their high human-capital-to-wealth ratio makes them very sensitive to labor market tail risk.

Finally, I show that removing cyclical skewness from the model increases aggregate demand for equity by 11% to 14%, which is equivalent to a 0.5-percentage-point (pp) decrease in the equity premium. The modest magnitude of this result contrasts with findings by Constantinides and Ghosh (2017) and Schmidt (2016) that cyclical skewness in the cross-section of consumption or income shocks can solve the equity premium puzzle. In these studies, consumption risk is equally distributed across the wealth distribution. In my model, consumption is endogenous and tail income shocks only translate into tail consumption shocks for households with high human-capital-to-wealth ratios, which tend to be young households with little wealth. Wealthier households do not cut their consumption as drastically in response to the same income shocks. Second, in the data, wealth is concentrated among retirees and workers approaching retirement: households below age 50 only owned 16% of U.S. wealth in 2016. As a result, countercyclical income risk is concentrated among the households whose investment choices are the least likely to change asset prices.

Overall, the countercyclical consumption risk of the average dollar-weighted investor is small. In my model, 10% of households own close to 50% of the total wealth but bear less than 2% of countercyclical consumption risk. Taking into account that exposure to the business cycle varies with age and across the earnings distribution (Guvenen et al. 2017) does not reverse that conclusion. On the one hand, exposure to the business cycle is much higher at the top of the earnings distribution; on the other hand, it is also lower among older workers who own most of the wealth.

A handful of portfolio choice studies also argue that human capital has stock-like properties. In particular, optimal equity shares can increase with age when the labor and stock markets are cointegrated (Benzoni, Collin-Dufresne, and Godlstein 2007) or when the variance of income shocks is countercyclical (Storesletten, Telmer, and Yaron 2007; Lynch and Tan 2011). However, evidence that countercyclical variance explains the data are weak, as existing models predict high equity shares and generate life cycle patterns that do not match the data.[3] More importantly, Guvenen, Ozkan, and Song (2014) find no evidence that variance is countercyclical in U.S. administrative data. Huggett and Kaplan (2016) use the PSID to calibrate a model with cointegration and state-dependent variance and skewness but predict that fresh college graduates with a risk aversion of 10 should have an equity share of 85%, which then declines until retirement.[4] My contribution to this strand of literature is to show that, when estimated using the moments documented in Guvenen, Ozkan, and Song (2014), life cycle models with countercyclical risk can match the age profile of equity holdings and also explain the cross-section of portfolios within age groups.

My paper also contributes to the literature on housing and stock market participation. Cocco (2005) studies how investment in housing crowds out equity investment among young households. Vestman (2019) documents that renters invest less in equity than homeowners, and shows that a life cycle model with heterogenous Epstein-Zin preferences and stock market participation costs can match this stylized fact. My paper proposes an alternative mechanism. Renters and highly leveraged homeowners, whose lifetime consumption depends mostly on their future earnings, do not invest in stocks because they are the most exposed to cyclical skewness.

My study builds on the growing literature on tail income risk. Cyclical skewness has been documented in administrative data from the United States (Guvenen, Ozkan, and Song 2014), Germany, Sweden, France (Busch et al. forthcoming), Canada (Bowlus et al. 2021), and Spain (Bonhomme et al. 2021). Guvenen, Ozkan, and Song (2014) show that left-tail income risk is cyclical across the entire earnings distribution, but more so for young workers. Previous work has also explored the source of cyclical skewness. Guvenen, Ozkan, and Song (2014) show that long unemployment spells play some role in the rise in tail risk during economic downturns but cannot explain most of it. Busch et al. (forthcoming) find that cyclical skewness affects job-switchers more, but, to a lesser extent, workers who keep the same employer as well. For job-switchers, the likely explanation is the difficulty in finding opportunities paying as much as their previous employment. For workers with the same employer, lower earnings are explained by a combination of lower hours and lower wage rates.

---

[3] In Storesletten, Telmer, and Yaron (2007), the equity share rapidly rises from 0% to more than 100% between 30 and 40 years old and declines after 50. In Lynch and Tan (2011), the optimal equity share rises rapidly in the early 30s, then flattens before rising quickly again in the late 50s. With countercyclical risk only, their model predicts a declining risky share after 35.

[4] Differences with my findings probably come from discrepancies between survey and administrative data. Huggett and Kaplan (2016) assume that the third central moment of persistent income shocks falls from −0.04 in expansions to -0.06 during recessions. But Guvenen, Ozkan, and Song (2014) find that the third central moment fell from -0.07 in 2000 and 2007 to -0.29 in 2001 to -0.34 in 2008.

The fact that having to find a new job in a depressed labor market has lifelong consequences is consistent with evidence that graduating during recessions has long-term detrimental consequences on earnings (Kahn 2010).

This literature has renewed interest in Mankiw's (1986) idea that the concentration of aggregate shocks among a few individuals matters for asset pricing. Krueger and Lustig (2010) demonstrate that idiosyncratic income shocks cannot affect the equity premium if their distribution is independent from aggregate shocks. Storesletten, Telmer, and Yaron (2007) and Favilukis (2013) incorporate countercyclical variance in labor income risk in general equilibrium models and Constantinides and Ghosh (2017) does the same for cyclical skewness in consumption risk. My contribution to this literature is to show that, in a realistic model, exposure to countercyclical consumption risk is concentrated among young workers with high human capital and little wealth, which greatly mitigates its asset pricing implications.

Finally, my paper is also closely related to that of Catherine, Sodini, and Zhang (2020), which shows that, in Swedish administrative data, workers who face higher cyclical skewness invest less in stocks. This correlation is stronger in the middle of the life cycle and among households with higher human-capital-to-wealth ratio. These findings fit the predictions of my model.

1. Model

In this section, I set up a life cycle model of consumption and portfolio choices. The main novelty of this model is that the distribution of idiosyncratic labor income shocks tends to have negative skewness in years of low stock market returns.

1.1 Macroeconomic environment

Stock market returns. The log return of the stock market index on year t is

$$s_t = s_{1,t} + s_{2,t}, \tag{1}$$

where s_1 denotes the component of returns that covaries with labor market conditions. To take into account stock market crashes, I assume that s_1 follows a normal mixture distribution

$$s_{1,t} = \begin{cases} s^-_{1,t} \sim \mathcal{N}(\mu^-_s, \sigma^2_{s_1}) & \text{with probability} \quad p_s \\ s^+_{1,t} \sim \mathcal{N}(\mu^+_s, \sigma^2_{s_1}) & \text{with probability} \quad 1-p_s, \end{cases} \tag{2}$$

whereas $s_{2,t}$ is normally distributed with variance $\sigma^2_{s_2}$. Without loss of generality, I impose $\mu^-_s < \mu^+_s$ and interpret μ^-_s as the expected log return in years of stock market crashes, and p_s represents the frequency of crashes. On the other hand, μ^+_s is the expected log return during normal years.

House prices. Households are exposed to the evolution of house prices as either owners or renter. I model the dynamics of log local house prices as

$$p_t - p_{t-1} = \mu_p + \varepsilon_{p,t}, \tag{3}$$

where $\varepsilon_{p,t}$ follows $\mathcal{N}(0, \sigma^2_p)$.

Aggregate labor income shocks. Stock market returns and house prices are correlated with the growth of national wage index L_1. Specifically, the dynamics of the log of wage index is

$$l_{1,t} - l_{1,t-1} = \mu_l + \lambda_{ls} s_{1,t} + \lambda_{lp} \varepsilon_{p,t} + \varepsilon_{l,t}, \tag{4}$$

where $\varepsilon_{l,t}$ follows $\mathcal{N}(0, \sigma_l^2)$, μ_l is the average growth rate of wages, and λ_{ls} and λ_{lp} capture the correlations between the growth rate of the wage index and, respectively, stock market returns and real-estate prices.[5]

Economic intuition suggests that cyclical skewness would have greater consequences if the distribution of income shocks remained negatively skewed for a few years after stock market crashes. However, Figure 1 shows that spikes in negative skewness accompanying stock market crashes do not persist over several years.

1.2 Idiosyncratic labor income shocks

Labor earnings can be decomposed as the product of the wage index and an idiosyncratic component $L_{2,it}$:

$$L_{it} = L_{1,t} \cdot L_{2,it}. \tag{5}$$

The idiosyncratic component is further decomposed into a deterministic function of age f_{it}, a persistent component z_{it} and a transitory shock η_{it}

$$L_{2,it} = e^{f_{it} + z_{it} + \eta_{it}}. \tag{6}$$

Persistent income shocks. As is standard in the labor literature, the persistent component follows an AR(1) process, but innovations are drawn from a normal mixture. Specifically, the dynamics of z_i are given by

$$z_{it} = \rho_z z_{it-1} + \zeta_{it}, \tag{7}$$

where

$$\zeta_{it} = \begin{cases} \zeta_{it}^- \sim \mathcal{N}\left(\mu_{z,t}^-, \sigma_z^{-2}\right) & \text{with probability} \quad p_z \\ \zeta_{it}^+ \sim \mathcal{N}\left(\mu_{z,t}^+, \sigma_z^{+2}\right) & \text{with probability} \quad 1 - p_z. \end{cases} \tag{8}$$

The values of p_z, $\mu_{z,t}^-$ and $\mu_{z,t}^+$ control the degree of asymmetry in the distribution of income shocks. To match the cyclicality of skewness, I define $\mu_{z,t}^-$ as an affine function of the log growth rate of the wage index:

$$\mu_{z,t}^- = \overline{\mu_z^-} + \lambda_{zl}(l_{1,t} - l_{1,t-1}). \tag{9}$$

Since idiosyncratic shocks have an expectation of zero, I impose

$$p_z \mu_{z,t}^- + (1 - p_z)\mu_{z,t}^+ = 0 \tag{10}$$

and, without loss of generality, $p_z \leq 0.5$. In this case, and if $\sigma_z^- \gg \sigma_z^+$, p_z can be interpreted as the frequency of important events in a worker's career. If $\lambda_{zl} > 0$,

[5] Flavin and Yamashita (2002) find that the correlation of housing returns with stocks is almost zero. For the sake of parsimony, I assume that stock market returns and house prices are uncorrelated.

these events tend to be good (e.g., promotions) during expansions and bad (layoffs) in recessions. This interpretation is consistent with the data. Anecdotal evidence suggests that most workers do not experience any noticeable income shock from one year to the next. Yet, the distribution of earnings growth reveals a very high volatility. This paradox can be explained by the fact that a small share of workers, typically job-switchers, experience very large earnings shocks. Finally, note that Equation (10) guarantees that idiosyncratic shocks are uncorrelated with stock market returns.

Transitory income shocks. The transitory component of income is also modeled as a mixture of normals whose first and second components always coincide with the first and second components of the normal mixture governing the innovations to z_i

$$\eta_{it} = \begin{cases} \eta_{it}^- \sim \mathcal{N}(0, \sigma_\eta^{-2}) & \text{if} \quad \zeta_{it} = \zeta_{it}^- \\ \eta_{it}^+ \sim \mathcal{N}(0, \sigma_\eta^{+2}) & \text{if} \quad \zeta_{it} = \zeta_{it}^+. \end{cases} \tag{11}$$

1.3 Social Security and social safety net

Social Security. Social Security payroll taxes represent 12.4% of earnings up to a limit known as the maximum taxable earnings, which represents 2.5 times the national wage index:[6]

$$\mathrm{T}_{it} = .124 \cdot \min\left\{L_{it}, 2.5 \cdot L_{1,t}\right\}. \tag{12}$$

Pensioners receive a percentage of the value of the national wage index when they retire. This percentage depends on the agent's historical taxable earnings, adjusted for the growth in the national wage index. Specifically, the initial Social Security benefit B_1 is

$$\frac{B_{1,i}}{L_{1,R}} = \begin{cases} .9 \cdot S_{iR} & \text{if } S_{iR} < .2 \\ .116 + .32 \cdot S_{iR} & \text{if } .2 \leq S_{iR} < 1 \\ .286 + .15 \cdot S_{iR} & \text{if } 1 \leq S_{iR}, \end{cases} \tag{13}$$

where R is the retirement age and $L_{1,R}$ is the value of the wage index at that age. The variable S_{it} keeps track of a worker's average taxable idiosyncratic earnings:

$$S_{it} = \sum_{k=t_0}^{t} \frac{\min\left\{L_{2,ik}, 2.5\right\}}{t - t_0 + 1}, \tag{14}$$

where t_0 denotes his first year of earnings. Because the benefit formula is a concave function of lifetime earnings, individuals with low earnings enjoy higher internal rates of return on their contributions. To avoid keeping track of historical earnings, many papers model Social Security benefits as a percentage

[6] In 2014, the value of the Social Security wage index was \$46,481, and the maximum amount of taxable earnings was \$117,000.

of the agent's final permanent income. This methodology overestimates Social Security risk. In my model, a large negative shock close to retirement would have dramatic consequences on Social Security benefits, which is not the case in reality. To solve this problem, I keep track of the agent's average historical earnings.

Safety net. Welfare programs limit the impact of human capital disasters on consumption, and thus may also alleviate the portfolio consequences of cyclical skewness. To take this into account, I model the Supplemental Nutrition Assistance Program (also known as the food stamp program) and, for individuals older than 65, the Supplemental Security Income (SSI) program. Eligibility for these programs is limited to individuals with very low wealth (roughly \$2,000), which is less than 5% of the average national wage in the model. After 65 years old, eligible individuals receive supplemental income such that their total income reaches at least 20% of the national average wage. Before 65 years old, eligible individuals with earnings below 20% the wage index receive benefits equal to 6% of the wage index minus 30% of their earnings.[7] Hence, welfare benefits are defined by

$$B_{2,it} = \begin{cases} \left(.2 \cdot L_{1,t} - B_{1,i}\right)^{+} & \text{if } W_{it} < .05 \cdot L_{it} \text{ and } t \geq R \\ \left(.06 \cdot L_{1,t} - .3 \cdot L_{it}\right)^{+} & \text{if } W_{it} < .05 \cdot L_{it} \text{ and } t < R \end{cases} \tag{15}$$

where W is the wealth of the agent.

1.4 Household

I incorporate the stochastic model in the life cycle problem of an agent who decides to rent or own a house, chooses his consumption of goods C and housing H, as well as the share of his net worth W invested in equity π.

Preferences. The agent has CRRA preferences and maximizes his expected utility, given by

$$V_{t_0} = \mathrm{E} \sum_{t=t_0}^{T} \beta^{t-1} \left(\prod_{k=0}^{t-1} (1 - m_k) \right) \frac{\left(C_{it}^{1-\nu} H_{it}^{\nu}\right)^{1-\gamma}}{1-\gamma}, \tag{16}$$

where γ is the coefficient of relative risk aversion, ν the preference for housing over consumption goods, m_k the mortality rate at age k, β the subjective discount factor, and T the maximum life span.

Disposable income. Disposable income, Y_{it}, includes labor earnings, transfers, and retirement benefits, net of payroll taxes:

$$Y_{it} = L_{it} + B_{1,it} + B_{2,it} - \mathrm{T}_{it}. \tag{17}$$

Housing. The agent owns or rents his house. In the latter case, he pays a yearly rent equal to

$$X_{it} = \begin{cases} P_t H_{it} / \chi & \text{if renter} \\ 0 & \text{if homeowner,} \end{cases} \tag{18}$$

[7] In 2018, a single-person household could receive at most \$204 per month, and the average benefit was \$138.

where χ denotes the price-to-rent ratio. There are two financial frictions on the housing market. First, selling a house incurs a transaction cost $\Phi_{H,it}$ proportional to its market value:

$$\Phi_{H,it} = \begin{cases} c_H P_t H_{it-1} & \text{if selling} \\ 0 & \text{otherwise.} \end{cases} \tag{19}$$

At the beginning of each period, the key difference between renters and homeowners is that the latter need to pay this transaction cost if they decide to move.

The agent can get a 1-year mortgage on his house, but a down payment requirement limits the value of the house he can buy:

$$W_{it} \geq \phi P_t H_{it}. \tag{20}$$

The mortgage is renegotiated at the beginning of each year. Because of mortgage debt, wealth can drop below zero following a large drop in house prices. When this happens, I assume that the agent defaults on his mortgage. His house and wealth are seized by creditors, and he starts the next period as a renter with zero wealth ($W = H = 0$) and no debt.

Wealth. Wealth W_{it} is the sum of financial assets, plus housing wealth, minus mortgage debt. After receiving his income and paying his rent and transaction costs, the agent determines his consumption and invests his remaining wealth in stocks, bonds, or housing. Participating in the stock market incurs a fixed cost $\Phi_{\pi,it}$ proportional to the wage index:

$$\Phi_{\pi,it} = \begin{cases} c_{\pi,1} L_{1,t} & \text{if } \pi_{it} > 0 \\ 0 & \text{otherwise} \end{cases} \tag{21}$$

and a management fee $c_{\pi,2}$ proportional to stock holdings. Hence, wealth evolves as follows:

$$W_{it+1} = \left[W_{it} + Y_{it} - C_{it} - X_{it} - \Phi_{H,it} - \Phi_{\pi,it}\right] \cdot \left[\pi_{it} e^{s_t - c_{\pi,2}} + (1 - h_{it} - \pi_{it}) e^{r} + h_{it} e^{p_t - p_{t-1}}\right], \tag{22}$$

where π_{it} is the equity share, h_{it} denotes the housing share, and r is the risk-free rate. The agent cannot short sell or borrow against future earnings.

Dynamic optimization. The problem is solved by dynamic programming. Besides age, the state variables of the problem are the components of labor income L_1, z_i and η_i, Social Security record S_i, wealth W_i, the current home size H, and housing price p. After I drop indexes to simplify the notation, I can write the Bellman equation as

$$V_t(W, L_1, z, \eta, S, H, p) = \max_{[C, H, \pi, \mathbf{1}_{own}]} \left\{ \frac{\left(C_{it}^{1-\nu} H_{it}^{\nu}\right)^{1-\gamma}}{1-\gamma} + (1 - m_t)\beta \mathrm{E} V_{t+1} \right\}, \tag{23}$$

and the terminal value is the period utility derived from entirely consuming retirement benefits and wealth.

Table 1
SCF summary statistics

	Mean	SD	Observations
Age	48.4	16.2	144,674
Wealth	287,622	1,151,775	144,674
Financial wealth	157,614	887,177	144,674
Wage	69,534	64,760	117,845
Stock market participation	.484		144,674
Equity share	.150	.398	125,811
Conditional equity share	.283	.511	73,276
Homeownership rate	.679		64,851
Housing share	1.314	3.953	124,811
Conditional housing share	1.739	4.465	93,752

This table reports summary statistics for the SCF 1989–2016 surveys. The sample is restricted to households whose head is between 23 and 81 years old. The wage statistics are computed for households with a wage earner between 23 and 65 years old. The equity share and housing share are computed as a fraction of wealth for households with wealth of at least $1,000.

2. Estimation

The model is estimated in two steps. In the first step, I use moments estimated in previous studies using Social Security panel data and S&P 500 data to estimate the parameters controlling the dynamics of earnings and stock market returns. In a second step, I use data on the wealth and portfolios of U.S. households to estimate their preferences and a fixed yearly stock market participation cost.

2.1 Data

The estimation uses two data sources: the Social Security Master Earnings File (MEF) and the summary extract public data from the Survey of Consumer Finances. In what follows, words in italics refer to readily available SCF variables.

Labor income risk. Lacking access to the MEF, I rely on Guvenen, Ozkan, and Song (2014), who report cross-sectional moments (mean, standard deviation, skewness) of the distribution of log labor income growth, net of life cycle effects, for a representative sample of the U.S. male population between 25 and 60 years old for each year from 1978 to 2010. I also rely on Guvenen et al. (2021), who report the evolution of within-cohort wage inequalities for the same sample.

Households' portfolios. Moments relative to the wealth and portfolio choices of households are computed using the 10 waves of the SCF between 1989 and 2016. I use survey weights, adjusted to give equal importance to all surveys. I restrict my sample to households between 23 and 82 years old, who are not business owners (*bus*=0). Table 1 reports key statistics for the final sample. Households constitute the unit of analysis.

I define wealth using the SCF net worth variable (*networth*), which includes financial assets and real estate, net of debts. I normalize wealth by the average labor earnings of the year. The average earnings in each survey are computed by taking the mean of wage incomes (*wageinc*) for households whose head is between 22 and 65 and is part of the labor force (*lf*=1). The equity share is defined as equity holdings divided by wealth (*equity/networth*), for households with positive wealth. The SCF equity variable includes direct holdings as well as stocks held in mutual funds and retirement accounts. The housing share is defined as the value of the residential real estate divided by wealth (*(houses+oresre)/networth*) for households with at least $1,000 of wealth.

2.2 Stock market returns and labor income risk

The first step of the structural estimation is dedicated to the stochastic processes controlling macroeconomic and idiosyncratic shocks, which are estimated independently.

2.2.1 Macroeconomic shocks. Eleven parameters control the dynamics of stock returns, aggregate labor income shocks and house prices.

First, I set the growth rate of house prices to μ_p=0.01, which is consistent with data reported in Davis and Van Nieuwerburgh (2015). Combined with a price-to-rent ratio of χ=22, this calibration implies a return on housing close to 6%, including money saved on rent, consistent with findings of Flavin and Yamashita (2002). To estimate the volatility of house prices and their correlation with income shocks, I rely on earnings data from the Bureau of Economic Analysis and house price data from the Case-Shiller indexes for the 1977–2013 period, both of which are available at the level of metropolitan statistical areas (MSAs). I deflate these time series to obtain real log growth changes and find an MSA-level volatility of house prices of σ_h=0.09, as in Landvoigt, Piazzesi, and Schneider (2015), and a correlation of 0.276.

Second, I estimate eight parameters $-p_s$, μ_s^-, μ_s^+, σ_{s_1}, σ_{s_2}, μ_l, λ_{ls}, and $\sigma_l' = \sqrt{\sigma_l^2 + \lambda_{lp}^2 \sigma_p^2}$– by SMM using moments from the time series of the average wage growth and yearly S&P 500 returns. The first eight moments are the mean, standard deviation, and the third (skew) and fourth (kurtosis) standardized moments of each time series. The last moment is the correlation coefficient of the two time series. Stock market moments are computed using yearly data between 1900 and 2015. The time series of average wage growth goes from 1978 to 2010. For a given estimate of σ_l', the values of σ_l and λ_{lp} are then set to match the correlation between earnings and house prices.

Table 2 reports the results of the estimation and shows that, despite being slightly overidentified, the model matches all moments very well. Note that the model is conservative regarding the frequency and severity of financial crashes. In the model, log returns below $-.35$ (≈-30%) occur three times per century. In the data, this actually happened five times since 1900 (1917, 1931, 1937, 1974, and 2008). For a log return of $-.35$, workers should expect their wage to drop by approximately 5%.[8] For the 2008 crisis, the model predicts an expected

[8] This wage drop can be calculated as $0.008-0.161*0.35=0.048$.

Table 2
Estimated parameters: Market returns, aggregate income shocks, and house prices

A. Estimated parameters

	Stock market returns					Aggregate income shocks				House prices	
Model	p_s	μ_s^-	μ_s^+	σ_{s_1}	σ_{s_2}	μ_l	λ_{ls}	λ_{lp}	σ_l	μ_p	σ_p
With stock market crashes	.146	−.245	.115	.077	.114	.008	.161	.053	.017	.01	.09
Without crashes			.063	.118	.141	.003	.250	.053	.019	.01	.09

B. Moments

	log returns				Aggregate income shocks					House prices	
	Mean	SD	Skew	Kurt	Mean	SD	Skew	Kurt	Corr_{ls}	SD	Corr_{lp}
Data	.063	.185	−.623	3.323	.018	.029	−.663	3.368	.638	.09	.276
Model											
– with crashes	**.063**	**.189**	**−.596**	**3.368**	**.019**	**.030**	**−.663**	**3.433**	**.638**	**.090**	**.276**
– without crashes	**.063**	**.185**	.000	3.000	**.018**	**.029**	.000	3.000	**.638**	**.090**	**.276**

The S&P 500 log return in year t is $s_t = s_{1,t} + s_{2,t}$, where $s_{2,t} \sim \mathcal{N}(0, \sigma_{s_2}^2)$. With probability p_s, t is a year of stock market crash and $s_{1,t} = s_{1,t}^- \sim \mathcal{N}(\mu_s^-, \sigma_{s_1}^2)$. Otherwise, $s_{1,t} = s_{1,t}^+ \sim \mathcal{N}(\mu_s^+, \sigma_{s_1}^2)$. The average log change in wages is $l_{t+1} - l_t = \mu_l + \lambda_{ls} s_{1,t} + \varepsilon_t$, where $\varepsilon_t \sim \mathcal{N}(0, \sigma_l^2)$. The SMM targets the first, second, third (skew), and fourth (kurt) standardized moments of each time series, as well as their correlation coefficient. The specification "without crashes" assumes normally distributed log returns and aggregate labor income shocks by imposing $p_s = \mu_s^- = \sigma_l = 0$ and only targets the means, standard deviations, and the correlation coefficient.

log wage shock of $-.062$ versus $-.064$ in the data. Note that the correlation of .638 between aggregate labor income shocks and returns does not prevent the correlation to be close to zero at the individual level as most of the variance in income shocks is idiosyncratic.

2.2.2 Idiosyncratic shocks. Eight parameters control the distribution of idiosyncratic labor income shocks: p_z, ρ_z, $\overline{\mu_z^-}$, λ_{zl}, σ_z^-, σ_z^+, σ_η^-, and σ_η^+. To estimate these parameters, I simulate an economy receiving the same aggregate shocks as the U.S. economy between 1944 and 2010.[9] For each year between 1978 and 2010, I target the historical values of Kelly's skewness at the 1-, 3-, and 5-year horizons, as well as the standard deviation at the 1- and 5-year horizons.[10] This constitutes a first set of 155 empirical moments targeted in the SMM.

The cyclicality of skewness is identified by targeting the historical values of these moments for each of the 33 years and using the true time series of aggregate labor income shocks as an input. The decomposition of the variance between transitory (σ_η^-, σ_η^+) and persistent shocks (σ_z^-, σ_z^+) and the persistence of the latter (ρ_z) are identified by targeting these moments at different horizons. One concern is that the model could overestimate the persistence of labor income shocks if workers face different life cycle income profiles ex ante (Guvenen 2009). To address this issue, I assume that the deterministic component of y_i is of the form:

$$f_i(t)=\overline{f}(t)+\varphi_i t+\alpha_i, \tag{24}$$

where $\overline{f}(t)$ is a function of experience common to all workers, while the two other terms are individual specific and normally distributed with standard deviations σ_φ and σ_α. Note that the main role of these two parameters is to get a conservative estimate of ρ_z.[11] To avoid the multiplication of state variables, I ignore σ_φ in the life cycle model and uses σ_α to initialize the distribution of z_i.

To identify σ_α and σ_φ, I also target the evolution of within-cohort inequalities between 25 and 60 years. Hence, I use 191 moments to estimate 10 parameters. For the first set of moments from Guvenen, Ozkan, and Song (2014), the SMM procedure seeks to minimize the square of the distance between true and simulated moments, normalized by the average value of each moment over the 1978–2010 period. For the second set of moments, the life cycle profile

9 I use the NIPA tables to estimate aggregate income shocks before 1978.

10 Guvenen, Ozkan, and Song (2014) do not report the standard deviation at the 3-year horizon.

11 The literature on labor income risk offers two alternative views on the rise of within-cohort inequalities over the life cycle (Guvenen 2009). In models with "restricted income profiles" (RIP), workers face similar life cycle income profiles and inequalities result from large and highly persistent shocks. By contrast, models with "heterogeneous income profiles" (HIP) assume individual-specific deterministic drifts and require lower shock persistence to explain the data. The HIP specification is less likely to overstate the role of income risk in portfolio choices by overestimating ρ_z.

Table 3
Estimated parameters: Idiosyncratic labor income shocks

		Persistent shocks					Transitory shocks		Heterogeneous income profiles	
Model	p_z	ρ_z	$\overline{\mu_z^-}$	λ_{zl}	σ_z^-	σ_z^+	σ_η^-	σ_η^+	σ_α	σ_φ
With cyclical skewness	.136	.967	-.086	4.291	.562	.037	.895	.089	.280	.004
Without cyclical skewness		.967				.216		.341	.280	.004

This table reports parameter estimates for the dynamics of idiosyncratic income shocks. The idiosyncratic component of log wages has a transitory component, η_{it}, and a persistent component, z_{it}, with persistence ρ_z. With probability $p_z \leq 0.5$, the agent receives a transitory shock $\eta_{it} = \eta_{it}^- \sim \mathcal{N}(0, \sigma_\eta^{-2})$, and a persistent shock $\zeta_{it} = \zeta_{it}^- \sim \mathcal{N}(\mu_{z,t}^-, \sigma_z^{-2})$, where $\mu_{z,t}^- = \overline{\mu_z^-} + \lambda_{zl}(l_t - l_{t-1})$ depends on the aggregate labor income shock $l_t - l_{t-1}$, which is an exogenous input. Otherwise, the transitory and persistent shocks are, respectively, $\zeta_{it} = \zeta_{it}^+ \sim \mathcal{N}(\mu_{z,t}^+, \sigma_z^{+2})$ and $\eta_{it} = \eta_{it}^+ \sim \mathcal{N}(0, \sigma_\eta^{+2})$. Since expected idiosyncratic shocks are equal to zero, I impose $p_z \mu_{z,t}^- + (1 - p_z)\mu_{z,t}^+ = 0$. σ_α and σ_φ represent the standard deviation of fixed effects in levels and trends. The model "with cyclical skewness" is estimated by targeting moments from Social Security data reported in Figure 2. In the model "without cyclical skewness," income shocks are lognormally distributed, and the other parameters have been chosen to keep the same persistence and variance.

of earnings inequalities, the procedure minimizes the percentage difference between true and simulated moments. Collectively, the first set of moments receives three-fourths of the total weights, and the rest goes to the second set. Hence, the weight matrix is a diagonal matrix whose first 155 elements take the value 0.75/155, and the rest 0.25/36. To maintain a constant age structure within the economy, I start the simulation in 1944 and replace each cohort reaching 60 with young workers. More details on the estimation procedure are provided in Appendix B.1.

Table 3 reports the results of the SMM, and Figure 2 plots targeted and simulated moments. The model replicates correctly the cyclicality of skewness at different horizons as well as the stability of the standard deviation and, unlike Guvenen, Ozkan, and Song (2014), its high value.[12]

The empirical 1-year standard deviation of income growth above 0.5 may seem counterintuitive but can be explained by the high peakness of the distribution. Since $p_z = 0.136$, in any given year, 86.4% of workers receive shocks from normal distributions with relatively low standard deviations: 0.037 and 0.089 for the persistent and transitory components, respectively. This means that most of the variance comes from the remaining 13.6% that receives shocks with very high standard deviations: 0.562 and 0.895. A plausible economic interpretation is that a fraction of the population switches jobs or becomes unemployed, while the vast majority experience more modest shocks. For comparison, Guvenen et al. (2021) reports that the standard deviations of log earnings growth of job-switchers and job-stayers are around 0.7 and 0.3, respectively.

[12] In Guvenen, Ozkan, and Song (2014), transitory shocks are normally distributed and the model cannot match the standard deviation. This problem can be solved by modeling transitory shocks using a normal mixture as in Equation (11).

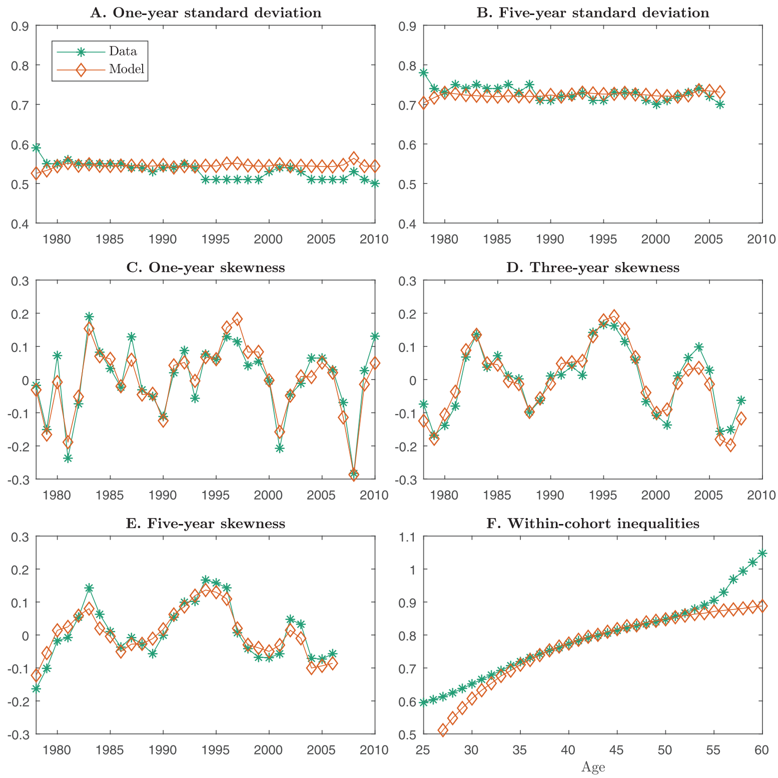

Figure 2
Fitness of idiosyncratic labor income risk model
This figure reports the empirical and simulated values of moments targeted in the estimation of the dynamics of idiosyncratic labor income shocks. Panels A to E report, for each year, cross-sectional moments from the distribution of labor income log growth, net of life cycle effects, at the 1-, 3-, and 5-year horizons among U.S. salaried male workers between 25 and 60 years old. Here, skewness is defined as $\frac{(d9-d5)-(d5-d1)}{d9-d1}$. Panel F reports the average historical within-cohort standard deviation of labor income at each age. All these moments are computed by Guvenen, Ozkan, and Song (2014) and Guvenen et al. (2021) using Social Security panel data. I simulate 67 cohorts of 7,500 individuals, with the first cohort entering the labor market in 1944. Historical aggregate income shocks between 1943 and 2010 are exogenous inputs in the simulation and come from Social Security data (1978–2010) and NIPA tables (1944–1977).

Another important result is the very high persistence of innovations to z. By increasing the volatility of the human capital stock, an overestimation of ρ_z would create spurious portfolio effects, in particular among young households. Four things can reduce this concern. First, my estimate of ρ_z is below that of the baseline specification of Guvenen, Ozkan, and Song (2014). Second, the model matches the ratio of standard deviations at the 5- and 1-year horizons, a simple measure of persistence. Third, the model does not exaggerate the evolution of within-cohort inequalities over the life cycle. Finally, the high persistence must be interpreted in conjunction with the high value of σ_η^-, which indicates that a large fraction of the most extreme shocks is transitory and recovered

Table 4
Calibrated parameters

Parameter			Value
Financial market			
	r	Risk-free rate	.02
	$c_{\pi,2}$	Proportional management fees	.01
Housing market			
	χ	Price-to-rent ratio	22
	ϕ	Down payment	.10
	c_H	Transaction cost	.10
Life cycle income profile			
	θ_1	Effect of age on log wage	.1237
	θ_2	Effect of $age^2/10$ on log wage	$-.0125$
	θ_0	Constant	-3.165
Agent			
	ν	Housing share	.30
	t_0	Age of first employment	23
	R	Age of retirement	65
	T	Maximum life span	100

within 1 year. Therefore, only shocks that do not quickly dissipate turn out to be very persistent. This is consistent with findings by Guvenen et al. (2021) that "workers whose earnings fall by 100 log points recover 25% of that decline in the first year and around 50% of the total within 10 years."

2.3 Preferences and participation cost

In the second step of the structural estimation, I use SCF data to estimate the coefficient of relative risk aversion, the discount factor and a fixed yearly participation cost. Technical details on the numerical resolution of the model and its estimation are provided in Appendix C.

2.3.1 Preset parameters. Table 4 reports preset parameters. I set the real interest rate to 2% and variable management fees to 1% of the stock portfolio. This calibration implies an equity premium slightly above 5% net of fees, in the upper range of values used in previous papers.

I approximate $\overline{f}(t)$ with a quadratic polynomial that matches the average log wage by age observed in Social Security data, from which I deduct an average 15% income tax. Households start working at 23 and retire at 65. Death probabilities are taken from Social Security actuarial life tables.

The price-to-rent ratio is set to 22. The minimum down payment and transaction costs each represent 10% of the house value, which are typical values in the literature. In the SCF, less than 5% of households have a house-to-wealth ratio above 10. Finally, the housing share in the Cobb-Douglas utility function is set to 30%, which is on the upper side of values used in previous studies but fits the data well.

In the data, households start with different levels of wealth. I compute the centiles of the distribution of wealth between 23 and 24 years old in the SCF and use them to generate 100 groups with different initial wealth in my simulation.

2.3.2 Identification. I run several sets of estimations and target the evolution over the life cycle of the following variables: wealth; conditional or unconditional equity shares; and stock market participation rate.

The empirical relationship between these variables and age is estimated as follows. First, I build 3-year age and cohort groups as in Ameriks and Zeldes (2004). Since the SCF is triennial, each cohort group moves exactly from one age group to the next between two surveys. Second, as in Fagereng, Gottlieb, and Guiso (2017), I disentangle age effects from year and cohort effects using the methodology of Deaton and Paxson (1994). Specifically, for each variable of interest, I run ordinary least squares (OLS) regressions on a set of age, year, and cohort group dummies and solve the collinearity problem by assuming that year dummies sum to zero and are orthogonal to a time trend. As a result, any time trend is captured by cohort effects. Finally, I compute the predicted values of each variable for each age group between $[23,25]$ and $[62,64]$. I adjust the SCF survey weights such that the procedure equally weights all cohorts born after 1946.

I find that equity shares increase with age until retirement. On the other hand, Ameriks and Zeldes (2004) find a slow decline in the 10 years preceding retirement using the 1989-1998 surveys and controlling for cohort fixed effects only. This decline disappears when the 2001–2016 surveys are included. Fagereng, Gottlieb, and Guiso (2017) find a slight decline in equity shares over the life cycle in Norwegian data, but this difference could be explained, for example, by lower income tail risk in Scandinavian countries.

From the SCF data, I obtain a vector $\mathbf{m}$ of $4 \cdot 14 = 56$ moments. The SMM procedure seeks the parameters γ, β, and $c_{\pi,1}$ that minimize

$$\left(\mathbf{m} - \widehat{\mathbf{m}}\left(\gamma, \varphi, \mathbf{c}_{\pi,\mathbf{1}}\right)\right)' \mathbf{W} \left(\mathbf{m} - \widehat{\mathbf{m}}\left(\gamma, \varphi, \mathbf{c}_{\pi,\mathbf{1}}\right)\right), \tag{25}$$

where $\widehat{\mathbf{m}}\left(\gamma, \varphi, \mathbf{c}_{\pi,\mathbf{1}}\right)$ is the vector of simulated moments generated by the model, and $\mathbf{W}$ is the inverse of the covariance matrix of the empirical moments. To compute the covariance matrix, I collapse the true data by year and cohort and bootstrap the collapsed data.

The identification is straightforward and well understood. A higher discount factor (β) increases the accumulation of wealth. A higher risk aversion (γ) reduces the optimal equity share. Finally, a higher fixed participation cost prevents the participation of households whose optimal equity share represents a small dollar amount.

2.3.3 Results.

Models without participation cost. I start by estimating the coefficient of relative risk aversion and the discount factor by targeting the evolution of wealth

Table 5
Estimated parameters: Preferences and participation cost

	With cyclical skewness				Without cyclical skewness		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A. Estimated parameters							
Relative risk aversion	7.136	6.066	6.259	6.090	11.952	8.007	8.819
	(.042)	(.042)	(.044)	(.028)	(.120)	(.068)	(.113)
Discount factor	0.911	0.956	0.938	0.936	0.666	0.882	0.824
	(.008)	(.005)	(.004)	(.007)	(.022)	(.010)	(.017)
Fixed participation cost		0.005	0.003	0.000		0.007	0.000
		(.000)	(.000)	(.000)		(.000)	(.000)
Fraction of never participants				0.394			0.489
				(.013)			(.012)
B. Targeted life cycle moments							
Wealth	✓	✓	✓	✓	✓	✓	✓
Equity share	✓				✓		
Participation rate		✓	✓	✓		✓	✓
Conditional equity share		✓	✓	✓		✓	✓

This table reports the estimated parameters and their standard errors in panel A, and indicates the set of targeted life cycle moments in panel B. Columns 1 and 5 assume no fixed participation cost and target the evolution of wealth and of the unconditional equity share. Columns 2 and 6 include a per-period participation cost and, besides wealth, target the life cycle profiles of the participation rate and the conditional equity share. In column 3, exposure to the business cycle depends on a worker's age and income. In columns 1–4, the model takes into account the cyclical skewness of idiosyncratic income shocks, whereas in columns 5–7 labor income shocks are lognormally distributed with identical variance, persistence and covariance with stock returns as in the model with cyclical skewness.

and that of the unconditional equity share. Columns 1 and 5 of Table 5 report estimation results for the models with and without cyclical skewness. Panel A of Figure 3 reports the simulated moments and their empirical counterparts.

Without cyclical skewness, the SMM fails to match the data and returns unlikely estimates of the discount factor (β=0.67) and relative risk aversion ($\gamma = 12$). The bond-like properties of human capital imply a positive relationship between the human-capital-to-wealth ratio and the optimal risky share, which in turn implies a decline of the equity share over the life cycle. A low β mitigates this problem by slowing the decline of the human-capital-to-wealth ratio, but, as a result, households accumulate much less wealth in the model than in the data. My estimates are close to Fagereng, Gottlieb, and Guiso (2017), who find γ=11 and β=0.77 in a similar model with stock market disasters and a small participation cost of \$69. When cyclical skewness is taken into account, human capital is no longer bond-like, and the model can match the data with more reasonable estimates of the discount factor (β=0.91) and relative risk aversion (γ=7.1). Because human capital is no longer bond-like, a positive relationship between age and equity holdings is compatible with the decline of the human-capital-to-wealth ratio over the life cycle.

Models with participation cost. In a second step, I introduce a fixed yearly participation cost and estimate it alongside β and γ. To do so, I distinguish the intensive and extensive margins by targeting the evolution of stock market participation rates and that of the equity share among participants. Columns 2

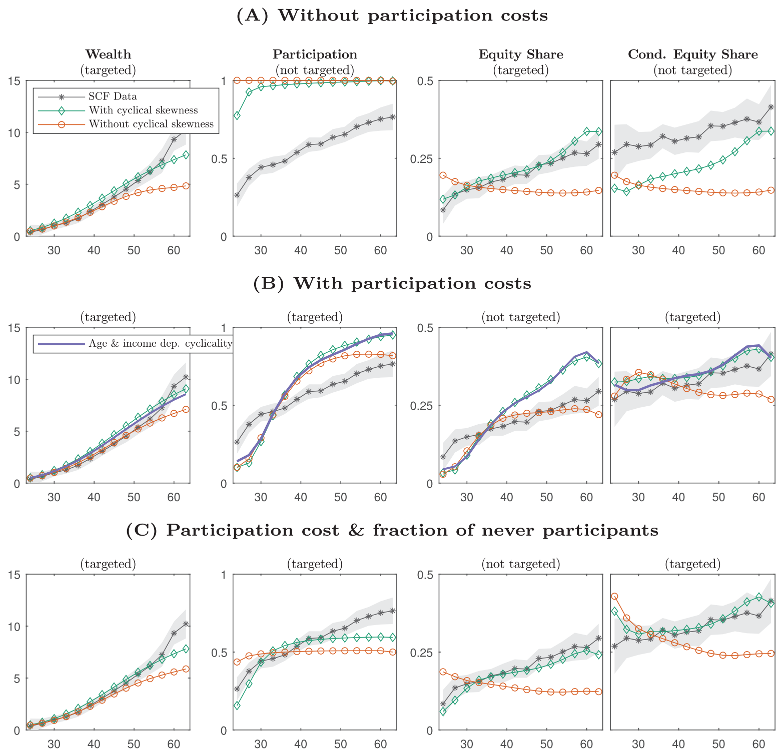

Figure 3
Fitness of estimated life cycle models
These panels plot the value of moments targeted in the SMM procedures reported in Table 5. Empirical life cycle patterns are estimated using the SCF and the methodology of Deaton and Paxson (1994). Ninety-five percent confidence intervals are reported in gray. In panel A, I assume no participation cost and target the evolution of wealth and unconditional equity shares. The models "with cyclical skewness" and "without cyclical skewness" refer to models 1 and 4, respectively. In panel B, I introduce a per-period participation cost and target the life cycle patterns of participation and that of the conditional equity share. Simulated moments refers to models 2, 3, and 6. Finally, in panel C, I assume that a fraction of the population never holds any stock, which corresponds to models 5 and 7.

and 6 of Table 5 reports the results and panel B of Figure 3 the model fitness. Without cyclical skewness, participation costs improve the model's ability to fit the data and produce more plausible estimates of β and γ. Estimated participation costs represent 0.8% of the average wage, that is, \$400 in 2016. When human capital is bond-like, matching low participation rates among young workers requires higher fixed costs because workers with low wealth want to invest all of it in the stock market. By preventing their participation, the fixed cost also reduces the average conditional equity share, which, in turn, affects the estimated value of γ, which falls from 12 to 8. The discount factor rises from 0.67 to 0.88 because wealth accumulation drives the rise in participation. Importantly, the model still fails to match the positive relationship between age and the conditional equity share.

With cyclical skewness, estimated participation costs are lower, representing only 0.5% of the average wage ($250). Because workers with low wealth would not want to invest much in the stock market anyway, a lower fixed cost is required to match low participation rates among young households. The model still generates a conditional equity share that slightly increases with age but overestimates participation among older households. With a fixed cost of $250, the model generates an annual unconditional stock market exit rate of 2.5% (see appendix Figure D1), which is consistent with the 11% exit rate over 5-year periods documented by Li (2014) using the PSID.

Both models tend to underestimate participation at the beginning of the life cycle and overestimate it at the end. This problem could be solved by introducing heterogeneity in participation costs. Arguably, some young workers face no participation costs because they are automatically enrolled in retirement plans by their employers. Other households might face much higher participation costs, for instance, because of financial illiteracy.

3. Discussion

3.1 Policy functions

Model policy function. Portfolio choice models generally predict that human capital decreases (increases) the demand for stocks when its exposure to stock market risk exceeds (is below) the agent's optimal risky share. Hence, absent participation costs, the equity share should be a monotonic function of the wealth-to-earnings ratio. This intuition does not hold in the presence of cyclical skewness. This is apparent in panel A of Figure 4, which plots the optimal equity share of a 40-year-old renter as a function of his persistent income (z) and wealth (w). Panel B displays the same policy function when income shocks are normally distributed.

In panel A, for a given level of persistent income (z), the equity share first increases with wealth, reaches a maximum, and then decreases. By contrast, in panel B, wealth always reduces the optimal equity share. As in Viceira (2001), human capital unambiguously increases the demand for stocks because the covariance between income shocks and stock market returns is negligible.

The difference between the two panels represents the effect of cyclical skewness, which is the strongest for households with high earnings-to-wealth ratios ($z-w$). What matters to the agent is the risk of severe shocks to lifetime consumption: the consequence of losing one's job during a recession is much more serious when most of his future consumption depends on future earnings. Short selling the stock market would be a simple way to hedge against this risk.

As the agent starts accumulating wealth, disastrous earnings shocks have milder implications in terms of consumption. Because left-tail consumption risk becomes less of a concern, the agent gives more attention to the low covariance between labor income shock and stock returns, and the optimal equity becomes

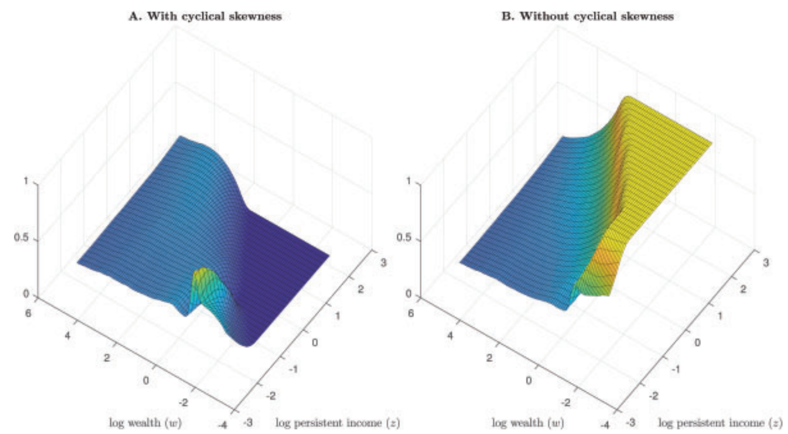

Figure 4
Optimal equity share of 40-year-old renters without participation costs
This figure represents the optimal equity share for a 40-year-old worker, whose mean historical taxable earnings has been $S=1$ so far, as a function of two state variables: the log of her wealth x and the log of her persistent income z. The transitory part of income is zero. Wealth and persistent income are scaled by the national wage index. The participation cost is set to zero. In panel A, all model parameters are set to their benchmark values, which replicate the skewness of stock returns and the cyclical skewness of idiosyncratic income shocks. In panel B, persistent income shocks are lognormally distributed with the same persistence, variance, and covariance with stock returns as in panel A. Stock market disasters occur in both specifications.

a positive function of z and a negative function of w. At the limit, when $w-z$ gets very large, the equity share converges to Merton's solution.[13]

A spike in equity share is visible in the southern part of the policy functions, in particular in the specification with cyclical skewness. This apparent abnormality coincides with the border of the region in which it becomes optimal for renters to buy a house. In this section, the effect of the SSI program are omitted.[14]

Empirical policy functions. Standard life cycle models predict that, controlling for age, the equity share should be a decreasing function of the wealth-to-earnings ratio. Figure 5 shows that the model without skewness delivers that prediction. However, there is no such relationship in the data, and it is substantially weaker in the model with cyclical skewness. This suggests that households do not behave as if human capital were a substitute for bonds but rather as if its stock-like component were close to what their optimal equity would be in a world without background risk.

Similarly, Figure 6 shows that, without skewness, the model predicts a strong positive relationship between the housing share and the optimal equity share. It

13 If returns are lognormally distributed, the solution to Merton's portfolio problem is $\frac{\mu-r}{\gamma\sigma^2}=\frac{.063-.02}{5.8\cdot.189^2}\approx .21$

14 Because SSI benefits are mean-tested, they create a discontinuity in the value function around the wealth eligibility threshold ($0.05\times L_1 \approx \$2,000$): households just below the threshold have higher utility than those above it. This break in the value function generates unrealistic behaviors for very low-wealth and low-income households. This behavior plays no role in specifications with participation costs, as these households all fall in the nonparticipation region.

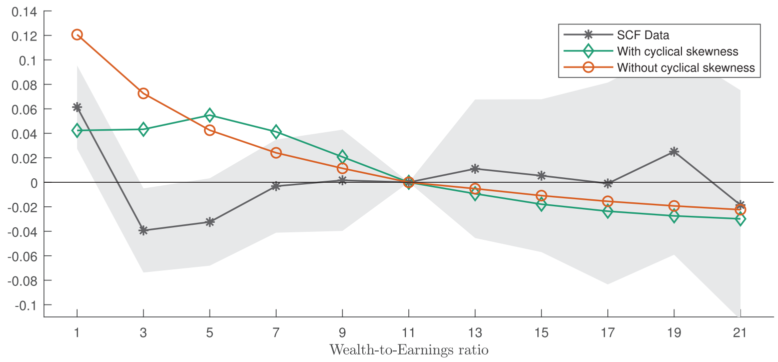

Figure 5
Conditional equity share and wealth-to-earnings ratio
This figure shows the relationship between the conditional equity share and the wealth-to-earnings ratio, in the SCF. Specifically, I regress the equity share of stock market participants on a set of wealth-to-earnings bins and age fixed effects. The sample is cut into 11 bins of wealth-to-earnings from [0-2] to [20-22]. I also include year fixed effects in the SCF specification. The figure reports the OLS coefficients associated with each wealth-to-earnings level. The specifications with and without cyclical skewness refer to the models estimated in columns 2 and 6, respectively, of Table 5.

may not seem intuitive that levered households would take more risk with their financial portfolios. However, holding age and wealth fixed, households with higher human capital buy larger houses. If human capital is bond-like, these households invest more in the stock market as well. However, there is no clear relationship between the housing and equity shares in the data or in the model with cyclical skewness.

We can compare these predictions to the Calvet and Sodini (2014) study of Swedish twins. First, this study documents that (a), controlling for human capital, the elasticity of the risky share to wealth is positive, and (b) that this elasticity is three times larger in the bottom quartile of wealth than in the top quartile. My simulated data show a relatively flat relationship, but Figure 4 suggests that their findings would be much easier to reconcile with the model when it incorporates cyclical skewness. The model without cyclical skewness predicts the opposite. As panel A shows, the optimal equity share increases in wealth until it exceeds multiple years of permanent income. Existing models explain this positive relationship by assuming a need for subsistence consumption (Wachter and Yogo 2010) or habit formation (Constantinides 1990, Campbell and Cochrane 1999), among other sources of decreasing relative risk aversion. The effect of cyclical skewness would be stronger with this type of preference.

Second, Calvet and Sodini (2014) find no relationship between the equity share and the leverage ratio, which is not consistent with the prediction of the model without cyclical skewness either. Yao and Zhang (2004) find a negative correlation between the housing and equity shares in the PSID.

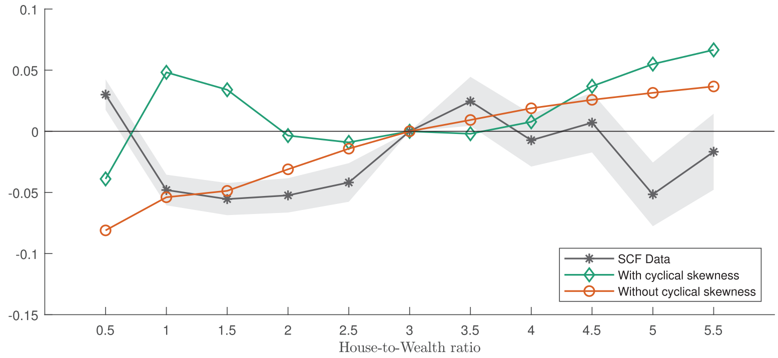

Figure 6
Conditional equity share and house-to-wealth ratio
This figure shows the relationship between the conditional equity share and the house-to-wealth ratio, in the SCF (1989–2016). Specifically, I regress the equity share of stock market participants on a set of house-to-wealth bins and age fixed effects. The sample is divided into 11 bins of house-to-wealth from [0.25–0.75] to [5.25–5.75]. The SCF data specification include year fixed effects. The figure reports the OLS coefficients associated with each house-to-wealth level. The specifications with and without cyclical skewness refer to models estimated in columns 2 and 6, respectively, of Table 5.

3.2 Decomposition of effects

In this section, I decompose the effects of cyclical skewness and stock market crashes on the life cycle profile of the optimal equity share. To do so, I start by simulating the model with normally distributed income shocks and returns, using parameter estimates reported in column 2 of Table 5. Then, I introduce or remove specific elements of the model to estimate their effect. Panel A of Figure 7 reports the evolution of the equity share over the life cycle depending on the presence of cyclical skewness and stock market crashes. Panel B reports the net effect of cyclical skewness and stock market crashes on equity shares.

Normally distributed shocks. Scenario (a) represents the evolution of the equity share when labor income shocks and stock market returns are normally distributed. This scenario is similar to those studied by Cocco, Gomes, and Maenhout (2005) and Viceira (2001), but with a lower risk aversion and higher variance of labor income shocks. In the absence of participation costs, the result is quantitatively close to Cocco, Gomes, and Maenhout (2005).

Stock market disasters. Scenario (b) shows that taking into account stock market crashes reduces the optimal equity share of young workers by only a few percentage points. In Fagereng, Gottlieb, and Guiso (2017), the introduction of left-tail risk in stock returns has a larger effect, but the difference may arise from the much higher risk aversion ($\gamma \geq 10$) considered by these authors.

Cyclical skewness. In scenario (c), households face cyclical skewness, but stock returns are lognormally distributed. Introducing cyclical skewness has a larger effect on the level of the equity share than stock market crashes for workers below age 45. The effect gets smaller as households get closer to

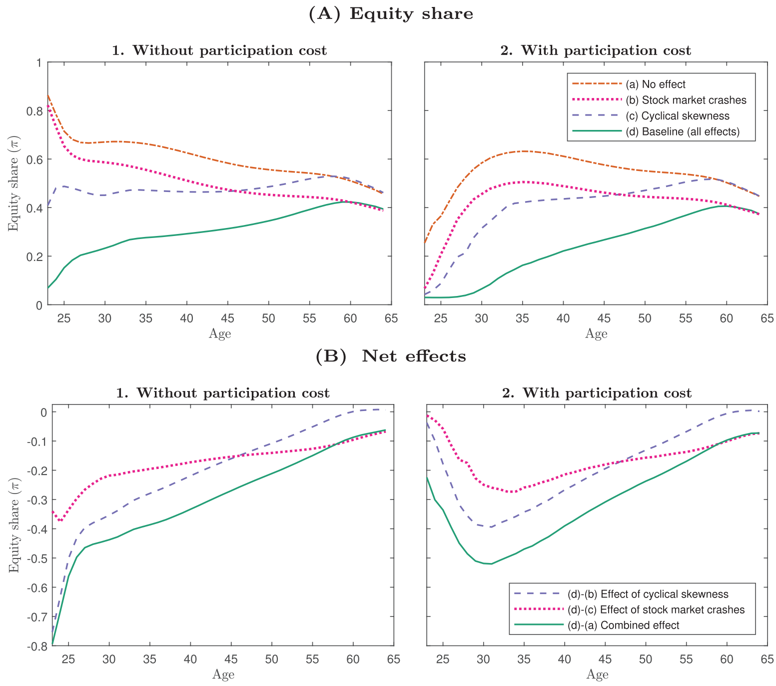

Figure 7
Decomposition of effects on the equity share
Panel A reports the simulated life cycle profile of the equity share under different scenarios. In scenario (a), log returns and idiosyncratic income shocks are normally distributed. Scenario (b) introduces stock market crashes. Scenario (c) incorporates the cyclical skewness of income shocks but assumes that log returns are normally distributed. Scenario (d) combines the effects of stock market crashes and cyclical skewness. Returns, transitory and persistent income shocks have identical means, variances and covariances in all scenarios. Panel B reports the effect of removing (a) cyclical skewness, (b) stock market crashes, or (c) both on simulated equity shares, relative to the baseline scenario.

retirement because most of their future consumption comes from wealth and Social Security entitlements.

Cyclical skewness & Stock market disasters. Scenario (d) combines the effects of scenarios (b) and (c). The difference between (c) and (d) is greater than between (a) and (b), which suggests an interaction effect between cyclical skewness and stock market crashes. For young agents with high earnings-to-wealth ratios, cyclical skewness strongly amplifies the importance of stock market tail risk. When income shocks are normally distributed, stock market tail risk reduces the optimal equity share π by roughly 9 pp over the entire life cycle. By contrast, assuming no participation cost, when cyclical skewness is taken into account, stock market tail risk reduces π by 35 (pp) at 25 years old, 20 pp at 35, 16 pp at 45, and 6 pp at 55.

Social Security. To a large extent, Social Security acts as a mandatory investment in bonds. If entitlements had perfect bond properties, and under

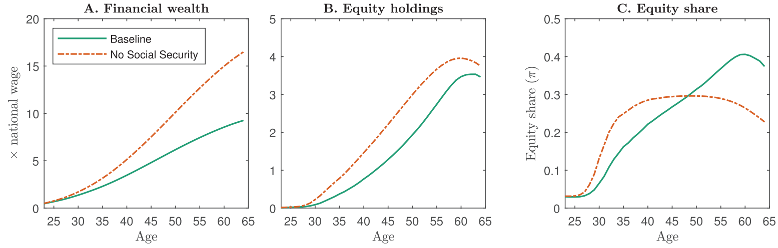

Figure 8
Effect of Social Security
These panels plot the life cycle profiles of wealth, equity holdings, and the unconditional equity share in the benchmark calibration, in the presence of participation costs, with and without Social Security. In the absence of Social Security, the agent does not receive any pension nor pay payroll taxes. Wealth and equity holdings are scaled by the national average wage.

the assumptions that (a) the mandatory yearly investment is below the optimal saving rate of the agent, and (b) the total investment in Social Security entitlements is less than what the agent wants to invest in bonds, then one dollar of investment in Social Security should reduce wealth and bond holdings by one dollar, but equity holdings should be unaffected.

Panels A and B of Figure 8 appear to be in line with these predictions at retirement age. However, equity holdings are reduced by Social Security among younger households, which is consistent with the facts that entitlements are somewhat risky, as their value evolves with the national wage index. The evolution of the equity share, reported in panel C, is quite different with Social Security. In the second half of the agent's career, the risky share is higher because entitlements reduce wealth, and therefore the denominator. In the first half of the agent's career, the risky share is lower. One possible explanation is that payroll taxes delay the accumulation of precautionary savings. Unlike wealth, Social Security entitlements cannot be used to smooth consumption in the event of a large negative income shock. This lack of liquidity could deter workers from investing in stocks and keep their precautionary savings safe.

At age 65, equity holdings are very close in the two scenarios. At that point, Social Security entitlements are risk-free and crowd out bonds. We know from Merton (1971) that optimal equity holdings represent the same fraction of *total* wealth, including the present value of entitlements. Hence, identical equity holdings in the two scenarios would indicate identical total wealth and perfect substitution between Social Security and private savings.

3.3 Stylized balance sheet statistics

Table 6 reports stylized balance sheet statistics for working-age households in the SCF and in simulated data. Overall, these statistics show that cyclical skewness helps match important stylized facts: the equity share (a) is low among renters and (b) increases with wealth among homeowners.

Table 6
Stylized balance sheet statistics

	Renters	Homeowners	
		Bottom 80%	Top 20%
A. Data			
Share of population	32%	49%	18%
Wealth	1.13	1.70	12.67
Equity share	15%	13%	25%
Housing share		214%	60%
B. Model with cyclical skewness			
Share of population	14%	66%	20%
Wealth	0.59	2.26	13.09
Equity share	0%	23%	37%
Housing share		171%	121%
C. After removing participation costs			
Share of population	14%	66%	20%
Wealth	0.61	2.31	13.75
Equity share	14%	31%	37%
Housing share		171%	121%
D. After removing cyclical skewness			
Share of population	5%	75%	20%
Wealth	0.17	2.07	12.37
Equity share	98%	51%	42%
Housing share		261%	117%

This table reports the average wealth, equity, and housing shares for renters and homeowners in and outside of the top 20% of the wealth distribution. Wealth is reported as a fraction of the average labor earnings in the economy. Panel A reports these statistics for the SCF population with positive net worth and below minimum retirement age of 62. Panel B reports the same statistics from simulated data from the baseline model with cyclical skewness and participation costs whose estimated parameters are reported in column 2 of Table 5. Panel C reports statistics for the same model and parameters, but without participation costs. Panel D reports the same statistics after removing cyclical skewness as well. The SSI program is omitted.

In the data, approximately one-third of households are renters with limited wealth, mostly invested outside of the stock market. A second group of households, representing roughly half the population, are homeowners with a leveraged position in housing. Finally, the top 20% of the population are mostly homeowners with an average equity share of 25%.

In panel B, I report the same statistics for my baseline estimation. The model performs well but slightly overestimates the share of homeowners. In the model, households with sufficient wealth have no reason to remain renters. However, in reality, households may have to move for exogenous reasons (Vestman 2019) such as job opportunities or changes in the size of their family. Renters are mostly households that cannot pay the down payment or have too little wealth to take on the risk of a mortgage. They lie in the upper-right corner of panel A of Figure 4 and therefore do not hold equity because of cyclical skewness. Among homeowners, the wealthiest have substantially higher equity shares and lower leverage, which is consistent with the data.

Simulation data in panel C use the same set of parameters, without participation costs. The equity share of renters rises to 14% but remains low. Similarly, homeowners in the top 20% keep invested a greater fraction of their wealth in stocks. Therefore, participation costs are not required to replicate the key patterns observed in the data.

Panel D uses the same set of parameters but assume lognormally distributed income shocks. Now renters are eager to invest in stocks because they have the highest human-capital-to-wealth ratio of the three groups: they are in the upper-left corner of panel B of Figure 4. Households are also willing to take up more leverage and therefore buy homes sooner, which explains the lower percentage of renters and the higher leverage among owners in the bottom 80%. Finally, among homeowners, households in the top 20% now have lower equity shares than those in the bottom 80% because their relative endowments of bond-like human capital are smaller.

3.4 Housing

Figure 9 reports how homeownership rates and the housing share of homeowners evolves with age. Panel A shows that homeownership rates grow faster in the model than in the data. In the model, virtually all households own their houses by the time they turn 40 years old. The difference with the real world could be explained by variations in housing market conditions over time and across the country. For example, the price-to-rent ratio is 10 times larger in San Francisco or Manhattan ($\chi = 50$) than in Detroit ($\chi = 5$). When the model is calibrated with a price-to-rent ratio of 50, less than 10% of households buy a house before their fifties. A richer treatment of housing that takes debt-to-income limits and exogenous moving shocks due to changes in households' professional or family life would also help the model match the data.

Panel B shows that the model matches the conditional housing share (HP/W) relatively well, though young households are, on average, slightly more levered in the model than in the data, especially without cyclical skewness. Young households remain renters longer in the model with cyclical skewness because they are more reluctant to take up a large mortgage.

Figure 10 reports the evolution of the equity share with and without the ability to buy a house in the model with cyclical skewness. Housing does not crowd out equity investment among young workers like in Cocco (2005). The reason is that young households do not want to invest in the stock market and, when they do, only invest a fraction of their wealth. The crowding out effect grows in importance as households get older and would like to invest more in stocks. After age 55, the presence of housing reduces the equity share by roughly 10 pp, or 20%.

The correlation between house price and income shocks is lower in my model (0.276) than in previous studies (Cocco 2005, Chetty, Sandor, and Szeidl 2017, Vestman 2019), but this difference has no noticeable incidence on estimated parameters or the life cycle profile of the equity share. As illustrated in Appendix Figure D2, the life cycle profile of the equity share barely changes when the correlation is increased to 0.6. This result is consistent with comparative statics reported in Chetty, Sandor, and Szeidl (2017).

Default rates are very low in the model: around 0.1% throughout the life cycle. Negligible default rates are consistent with the assumption that

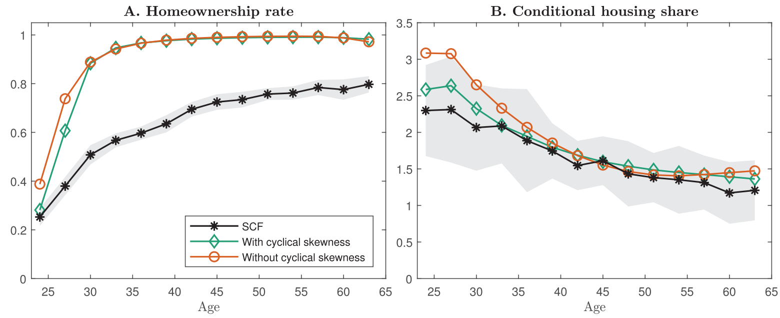

Figure 9
Housing investment over the life cycle
Panel A plots the life cycle profile of homeownership in the data and in the model with and without cyclical skewness. Panel B reports the share of housing in the household's wealth among homeowners. The specifications with and without cyclical skewness refer to the models estimated in columns 2 and 6, respectively, of Table 5. Empirical life cycle patterns are estimated using the methodology of Deaton and Paxson (1994). Ninety-five percent confidence intervals are reported in gray.

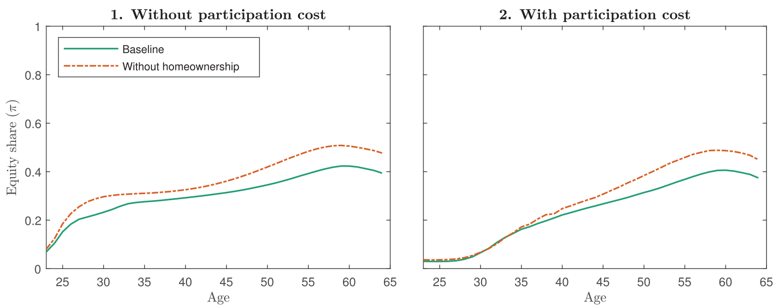

Figure 10
Effect of homeownership
These panels plot the life cycle profiles of the equity share in the baseline model, with and without the option to buy a house, using parameter estimates reported in column 2 of Table 5.

the mortgage rate equals the risk-free rate. Because of high risk aversion, households do not come close to the maximum leverage of 10 implied by the down payment constraint. Less than 3% have a housing share above 5. This means the assumption that mortgage debt is renegotiated every year is without consequence as mortgage rates do not change and the down payment constraint is never binding.

4. Aggregate Effects

4.1 Partial equilibrium results

How does cyclical skewness affect the optimal demand for equity and the equity premium? To answer this question, I run a counterfactual experiment in which labor income shocks are normally distributed, but have otherwise the same variance, persistence, and covariance with stock market returns. Holding the

Table 7
Partial equilibrium effect of cyclical skewness

Model	(1)	(2)	(3)	(4)
Effect of cyclical skewness on:				
– ln(equity share)	−.200	−.160	−.158	−.136
– ln(wealth)	.052	.039	.049	.0452
– ln(demand for equity)	−.148	−.121	−.109	−.091
Equivalent change in equity premium	−.005	−.005	−.004	−.004
Participation cost		✓	✓	✓
Age and income-dependent cyclicality			✓	
Fraction of never participants				✓

This table reports the effects of cyclical skewness on the aggregate demand for equity. Using models 1–4 of Table 5, I simulate the life of 10^6 individuals and sum their equity holdings over all ages, including retirement. Holding everything else equal, I then remove cyclical skewness from the model and repeat the computation. The first line of the table reports the log difference in aggregate demand for equity between the model without and the model with cyclical skewness. The log difference is explained by the change in aggregate wealth and the change in the aggregate equity share. I also compute the change in the equity premium (in both μ_s^- and μ_s^+) that exactly offsets the effect of removing cyclical skewness. Specifically, I use Newton's method to find the equity premium for which the aggregate demand in the model without cyclical skewness equals the demand in the model with cyclical skewness and the true equity premium.

equity premium constant, I then compute the change in the aggregate demand for equity in simulated data, including from retired households. This change represents the effect of cyclical skewness on the demand for equity. In a second step, I solve for the change in the equity premium that offsets the effect of cyclical skewness and raises the aggregate demand for stocks back to its initial value. This change corresponds to the effect of cyclical skewness on the equity premium assuming that the supply of stocks is inelastic. Therefore, it represents an upper bound on the effect of cyclical skewness on expected returns.

As reported in columns 1 and 2 of Table 7, cyclical skewness reduces the aggregate demand for stocks by only 9% to 15%, a reduction that can be offset by increasing the equity premium by half a percentage point. Columns 1 and 2 correspond to the specifications with and without fixed participation costs and when all workers are subject to identical labor income risk. The effect of cyclical skewness on the demand for equity can be decomposed into two components. First, cyclical skewness reduces the share of aggregate wealth invested in equity by 14% to 20%. However, cyclical skewness also increases aggregate wealth by roughly 5% by encouraging precautionary savings, which attenuates the first effect.

In the presence of short-sale constraints, cyclical skewness is unlikely to have a large aggregate effect in a realistic model because of two empirical facts. First, Guvenen, Ozkan, and Song (2014) document that exposure to cyclical skewness declines over the life cycle and essentially disappears around age 50 (appendix Figure D6). Second, as illustrated in Figure 11, households above age 50 own 84% of the financial wealth in the economy (and 83% of the total wealth). Total equity holdings represented \$22.4T in 2016, and other financial assets summed to \$19.7T, of which \$3.3T was held by households with a head younger than 50. Consequently, if households below age 50 invested all their financial wealth in stocks, aggregate equity holdings would increase from \$22.4T to \$25.7T.

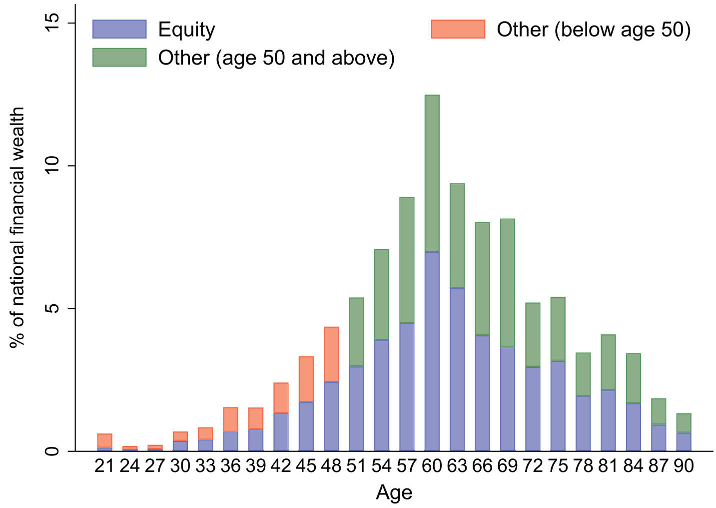

Figure 11
Distribution of financial wealth by age in 2016
This graph plots the distribution of financial wealth by age in the 2016 SCF. Total financial wealth was $42 trillion.

Taking this as an upper-bound on the plausible effect of cyclical skewness implies a reduction in the aggregate demand for equity of at most 13%. In Merton's model, this would be equivalent to a 13% drop in the equity premium: 0.8 pp. Pushing this logic even further, if households below age 50 invested all their wealth (excluding private businesses) in public equity, aggregate stock holdings would rise to $30.2T. This would still imply an upper bound of 26%, or 1.6 pp, on the direct effect of cyclical skewness on aggregate stock holdings.

4.2 Countercyclical consumption risk

Constantinides and Ghosh (2017) assume that all households face the same distribution of consumption shocks, which they calibrate to match the cross-sectional skewness of consumption growth in the Consumer Expenditure Survey. By contrast, in my model, consumption is endogenous, and, consequently, the distribution of consumptions shocks varies across the wealth distribution. Households with very low wealth-to-earnings ratios reduce their consumption drastically after a tail income shock, whereas households with higher wealth-to-earnings ratios do not. Moreover, retirees do not face labor income risk at all. Hence, the model suggests that the negative skewness in the cross-sectional distribution of consumption growth documented by Constantinides and Ghosh (2017) is probably concentrated among young and less wealthy households.

To further illustrate this point, Figure 12 plots the contribution of different deciles of wealth and age to the cyclical skewness of consumption risk in

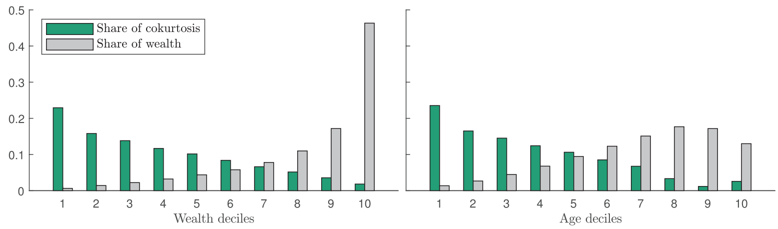

Figure 12
Decomposition of countercyclical consumption risk by deciles of wealth and age
These graphs plot the relative contribution of each decile of wealth and age to the cyclical skewness of consumption growth in the simulated data. Cyclicality is measured using cokurtosis, defined in Equation (26). Using estimated parameters reported in column 2 of Table 5, I simulate the life cycle of 10^6 individuals, including retirement. Simulated data are split by deciles of wealth (left panel) and deciles of age (right panel). The graph reports the relative contribution of each decile to the total cokurtosis and total wealth.

simulated data as well as their share of aggregate wealth. Here, I measure cyclical skewness using cokurtosis, defined as

$$\kappa = \frac{\mathrm{E}\left[(\delta_c - \mathrm{E}(\delta_c))^3 (s - \mathrm{E}(s))\right]}{\sigma_{\delta_c}^3 \sigma_s}, \tag{26}$$

where $\delta_{c,it} = \log\left(\frac{C_{it+1}}{C_{it}}\right)$ is the growth rate of consumption of household i and σ_{δ_c} its cross-sectional standard deviation. A large κ indicates that the distribution of δ_c is left (right) skewed in years of low (high) stock returns. Because κ is a sum over households, computing the contribution of each decile to κ is straightforward. As shown in Figure 12, individuals with low wealth contribute disproportionately to κ. By contrast, the top decile concentrates close to 50% of wealth and barely contributes to κ. Hence, the countercyclical consumption risk faced by the average household cannot explain the behavior of the average dollar-weighted investor.

What happens when we take into account that exposure to the business cycle varies with age and across the income distribution? To answer this question, I repeat the aggregation exercise using the model with age and income-dependent cyclicality. As reported in column 3 of Table 7, the aggregate consequences of cyclical skewness get even smaller. Two opposing forces are at play. On the one hand, greater countercyclical risk among high earners reduces the aggregate demand for equity. On the other hand, lower countercyclical risk among workers in their fifties has the opposite implication and seems to dominate.

5. Extensions and Robustness

5.1 Age and income-dependent cyclicality

The cyclicality of labor income risk declines with age and is higher at the bottom and top of the earnings distribution (Parker and Vissing-Jorgensen 2010, Guvenen et al. 2017). I take this into account in an extension of the model in

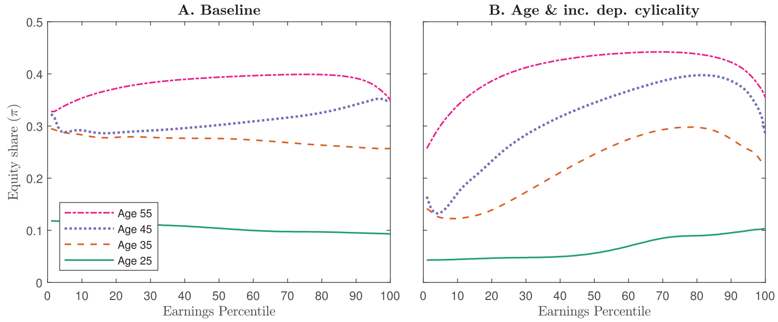

Figure 13
Introducing age and income-dependent cyclicality
These panels show the effect of introducing age and income-dependent cyclicality on the simulated equity share. Cyclicality by age and earning percentile are based on Guvenen et al. (2017) estimates of workers' "GDP beta." Simulations use parameter estimates reported in column 2 of Table 5, without participation cost.

which Equations (9) and (10) become

$$\mu^{-}_{z,t} = \overline{\mu^{-}_{z}} + b(z,t)\lambda_{zl}(l_t - l_{t-1}) \tag{27}$$

and

$$p_z\mu^{-}_{z,t} + (1-p_z)\mu^{+}_{z,t} = (b(z,t)-1)(l_t - l_{t-1}), \tag{28}$$

where $b(z,t)$ are "worker betas," measuring the sensitivity to labor market conditions as a function of persistent income z and age t. In Equation (27), cyclical variations in skewness are amplified by the coefficient $b(z,t)$. In Equation (28), the mean idiosyncratic shock is a function of the aggregate shock multiplied by excess beta.

I calibrate $b(z,t)$ using "gross domestic product (GDP) beta" estimates reported for selected age groups and percentiles of earnings in Guvenen et al. (2017).[15] The data show that exposure to the business cycles smoothly declines with age. Within each age group, exposure to the business cycle also decreases with earnings over most of the income distribution. However, this relationship reverses at the very top: business cycle variations are three times larger for the top 1% of earners than the rest of the population (see appendix Figure D4). Concretely, the model predicts that, over the 2007–2008 period, the average 40-year-old male worker in the top 1% lost 21 log points of persistent income.

Figure 13 illustrates how introducing workers' "GDP betas" changes optimal equity shares by age and across the earnings distribution. First, the life cycle profile of the equity share becomes steeper because exposure to the business cycle declines with age. Second, the model predicts a positive relationship between the optimal equity share and earnings. Calvet and Sodini (2014)

---

[15] See table A.1 of their online appendix. I approximate the missing percentiles and ages using akima interpolation and normalize the result to get an average beta of one. The correspondence between z and earnings percentiles is based on simulated data from the baseline model.

document this correlation and interpret it as evidence that human capital is a substitute for bonds. My model suggests a different interpretation: exposure to the business cycle is, overall, a decreasing function of earnings. Note that the correlation is stronger in the middle of the life cycle. At age 25, the correlation is weak because many households are in corner solutions ($\pi = 0$). At age 55, households have enough savings for cyclical skewness not to matter that much.

As reported in column 3 of Table 5, taking into account age and income-dependent exposure to the business cycle barely affects my estimate of γ. Even though Figure 13 shows that workers' "GDP betas" changes the cross-section of equity shares, it does not change the overall average equity share, which identifies relative risk aversion in the structural estimation.

5.2 Epstein-Zin utility

In this section, I consider extensions of the model in which the agent has Epstein-Zin preferences. Specifically, the value function is:

$$V_t = \left\{ (1-\beta(1-m_t))\left(C_{it}^{1-\nu} H_{it}^{\nu}\right)^{1-1/\psi} + \beta(1-m_t)\mathrm{E}\left[V_{t+1}^{1-\gamma}\right]^{\frac{1-1/\psi}{1-\gamma}} \right\}^{\frac{1}{1-1/\psi}} \tag{29}$$

where ψ is the elasticity of intertemporal substitution (EIS). I consider three specifications, for which estimated/calibrated parameters and partial equilibrium results are reported in appendix Table D1.

Baseline estimation. In the first estimation, I allow the EIS to be different from the inverse of relative risk aversion. Relative to the expected utility case, estimated risk aversion barely changes, increasing from γ=5.8 to 6.2. However, the discount factor drops from $\beta = .95$ to .93, while the EIS doubles from $1/\gamma = .17$ to $\psi = .34$. Importantly, the EIS does not seem identified by targeted moments, as combinations of lower (higher) ψ and higher (lower) β generate similar values for the SMM loss function. In the CRRA framework, the value of γ affects portfolio decisions and, through the EIS, wealth accumulation. Consequently, the model may be unable to match both aspects of the data simultaneously. In that case, breaking the link between risk aversion and the EIS helps match the data, and ψ would be well identified.

As reported in appendix Figure D5, the Epstein-Zin specification does not fit the data much better. This indicates that the new parameter is not well identified by the targeted moments, or, equivalently, that we do not need recursive preferences to match these moments. The partial equilibrium effect of cyclical skewness does not change either.

RRA of 10 and EIS of 2. I also compute the partial equilibrium effect of cyclical skewness for a standard calibration in the asset pricing literature: $\gamma = 10$ and $\psi = 2$. Under this calibration, the effect of cyclical skewness on the demand for equity is equivalent to a 0.2-pp decrease in the equity premium.

Heterogeneity in EIS. Finally, I estimate the model under the assumption that 20% of the population has an EIS of 0.1, while the rest has an EIS of 0.3, as in

Guvenen (2009). Estimated relative risk aversion is 6.2 in this specification. The partial equilibrium effect of cyclical skewness does not change significantly.

5.3 Stock market nonparticipants

The assumption that a fixed monetary cost explains low stock market participation rates is difficult to reconcile with reduced-form evidence that windfall wealth has limited effects on participation (Andersen and Nielsen 2011, Briggs et al. 2021).[16] Moreover, alternative solutions to the participation puzzle have been proposed: lack of trust (Guiso, Sapienza, and Zingales 2008), disappointment aversion (Ang, Bekaert, and Liu 2005), narrow framing (Barberis, Huang, and Thaler 2006) and ambiguity aversion (Campanale 2011, Peijnenburg 2018) among others.

To take this into account, I reestimate the model assuming that a fraction of the population never participates. Columns 4 and 7 of Table 5 report the results and panel C of Figure 3 shows the model's fitness. Estimated parameters are similar to columns 2 and 6. In both cases, the estimated share of never participants is around 46%–47%. I also report the effect of cyclical skewness on the aggregate demand for equity in column 4 of Table 7 and find previous conclusions to be robust.

5.4 Relative risk aversion

The portfolio effect of cyclical skewness increases with risk aversion. When risk aversion is below 4, the effect is relatively small because households invest most of their wealth in stocks regardless of whether they face income tail risk. Appendix Figure D7 plots the life cycle profile of the equity share for different levels of relative risk aversion, with and without cyclical skewness.

5.5 Relationship between house prices and stock returns

In the model, returns on housing and stocks are independently distributed. However, while their short-run correlation is small, some evidence indicates that these two markets are cointegrated over longer horizons. Moreover, some recessions—notably the 2008 financial crisis—are marked by large drops in house prices at the national level and greater variance across the country. Therefore, the independence assumption made in the model is a simplification.

This simplification has different consequences for renters and homeowners. As Sinai and Souleles (2005) explain, all households have to live somewhere and are in effect exposed to fluctuations in housing costs: "If rents rise, all households would experience an increase in their implicit future rent 'liability'. But homeowners would enjoy a compensating increase in the price of their

[16] Andersen and Nielsen (2011) find that receiving 134,000 euros after an unexpected inheritance raises the probability of Danish inheritors entering the stock market by only 21 pp. They also observe that the majority of households inheriting stocks actively exit the equity market. Briggs et al. (2021) find that among Swedish lottery players, a $150,000 windfall gain raises the probability of stock ownership from nonparticipants by only 12 pp.

house, since their dividend payments [implicit savings on rent] would increase commensurately. Such offsetting effects reduce the overall wealth effects from changes in house prices."

Following that logic, a positive relationship between house prices and stock market returns would increase the willingness of renters to buy stocks, because this investment would hedge their implicit exposure to house prices. On the other hand, homeowners should be relatively indifferent because their overall exposure to house prices is much lower than their balance sheet suggests: the value of their house is hedged by their need to live somewhere. If prices go up and they decide to sell their property, their capital gains would be captured by the need to rent another place.

Cointegration between the stock and housing markets could make equity more attractive to homeowners. Cointegration means that house prices would go down over the years following a permanent drop in stock market prices. Therefore, homeowners would be able to sell their property before it depreciates in value and, at the same time, enjoy a lower cost of housing over the long run.

Cyclical skewness in house prices could, on the other hand, reduce homeowners' willingness to buy stocks. Indeed, a large drop in a worker's permanent earnings reduces his future optimal consumption of housing services. In that situation, this worker should sell his house and move to a smaller, less valuable one. He could also move to another geographical area with better employment opportunities. In any case, a large depreciation in the value of his current house is not entirely offset by future housing consumption, either because his updated housing needs are smaller or because his new job will be in an area in which prices did not fall as much. Importantly, the relationship between house price risk and business cycles is nowhere as systematic as that of income risk. House prices increase during the 2001 and 2020 recessions.

5.6 Retirement years

Because this paper studies labor income risk, it focuses on households' lives before retirement. In appendix Figure D8, however, I report the evolution of wealth, real estate, and equity holdings until age 82. Around retirement age, the model predicts an increase in equity shares as Social Security payments replace volatile earnings. There is no discontinuity in the data, which could be explained by workers' ability to claim benefits early if need be.

It is well-known that, without a bequest motive, life cycle models overestimate the decline of wealth during retirement. In the model, this generates an increase in leverage, as borrowing is the only way households can consume their wealth without selling their homes. This is not observed in the data since wealth does not decrease.

Note that, to understand the implications of income risk for portfolio choices, it is important that the model matches the relative weights of human capital in workers' total wealth. This is achieved by matching the evolution of wealth. Whether households accumulate this wealth for their personal consumption or that of their descendant does not matter. Introducing a bequest motive would reduce the estimated discount factor (β) but should not affect pre-retirement portfolios as long as the evolution of the human capital-to-wealth ratio remains the same.

6. Conclusion

This paper shows that the cyclical skewness of idiosyncratic labor income shocks can reconcile life cycle models of portfolio choices with the U.S. data. Cyclical skewness can explain both the limited stock market participation among households with high human-capital-to-wealth ratios and why conditional equity shares rise with age.

Most of the existing literature has focused on life cycle patterns; however, cyclical skewness can explain other important aspect of the cross-section of equity holdings. Without cyclical skewness, the model predicts a strong positive relationship between the income-to-wealth ratio and equity shares. The model also predicts renters to have higher equity shares than homeowners. Moreover, among homeowners, simulated equity shares are positively correlated with leverage. All these predictions are rejected by the data and disappear when the model incorporates cyclical skewness.

Contrary to prior research, I find that countercyclical labor income risk has a limited effect on aggregate demand for equity because it does not significantly affect the portfolios of wealthy households. However, recent research shows that firms face substantial tail risk during recessions (Salgado, Guvenen, and Bloom 2019). Therefore, countercyclical tail risk may also deter undiversified business owners from investing in public equity, which could increase its effect on asset prices.

References

Ameriks, J., and S. P. Zeldes. 2004 How do household portfolio shares vary with age? Working Paper, Vanguard Corp.

Andersen, S., and K. M. Nielsen. 2011 Participation constraints in the stock market: Evidence from unexpected inheritance due to sudden death. *Review of Financial Studies* 24:1667–97.

Ang, A., G. Bekaert, and J. Liu. 2005 Why stocks may disappoint. *Journal of Financial Economics* 76:471–508.

Barberis, N., M. Huang, and R. H. Thaler. 2006 Individual preferences, monetary gambles, and stock market participation: A case for narrow framing. *American Economic Review* 96:1069–90.

Benzoni, L., P. Collin-Dufresne, and R. S. Godlstein. 2007 Portfolio choice over the life-cycle when the stock and labor markets are cointegrated. *Journal of Finance* 62:2123–67.

Bonhomme, S., M. De Vera, L. Hospido, and S. Wei. 2021 Income risk inequality: Evidence from spanish administrative records. Working Paper, CEMFI.

Bowlus, A., M. Gouin-Bonenfant, H. Liu, L. Lochner, and Y. Park. 2021 Four decades of canadian earnings inequality and dynamics across workers and firms. Working Paper, University of Western Ontario.

Briggs, J., D. Cesarini, E. Lindqvist, and R. Östling. 2021 Windfall gains and stock market participation. *Journal of Financial Economics* 139:57–83.

Busch, C., D. Domeij, F. Guvenen, and R. Madera. Forthcoming. Skewed idiosyncratic income risk over the business cycle: Sources and Insurance. *American Economic Journal: Macroeconomics*.

Calvet, L. E., and P. Sodini. 2014 Twin picks: Disentangling the determinants of risk-taking in household portfolios. *Journal of Finance* 69:867–906.

Campanale, C. 2011 Learning, ambiguity and life-cycle portfolio allocation. *Review of Economic Dynamics* 14:339–67.

Campbell, J., and J. Cochrane. 1999 Force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy* 107:205–51.

Campbell, J. Y., J. F. Cocco, F. J. Gomes, and P. J. Maenhout. 2001 Investing retirement wealth: A life-cycle model. In *Risk aspects of investment-based social security reform*, 439–82. Cambridge, MA: National Bureau of Economic Research.

Catherine, S., P. Sodini, and Y. Zhang. 2020 Countercyclical income risk and portfolio choices: Evidence from Sweden. Working Paper, University of Pennsylvania.

Chai, J., W. Horneff, R. Maurer, and O. S. Mitchell. 2011 Optimal portfolio choice over the life cycle with flexible work, endogenous retirement, and lifetime payouts. *Review of Finance* 15:875–907.

Chetty, R., L. Sandor, and A. Szeidl. 2017 The effect of housing on portfolio choice. *Journal of Finance* 72:1171–212.

Cocco, J. F. 2005 Consumption and portfolio choice over the life cycle. *Review of Financial Studies* 18:491–533.

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

Constantinides, G. M. 1990 Habit formation: A resolution of the equity premium puzzle. *Journal of Political Economy* 98:519–43.

Constantinides, G. M., and A. Ghosh. 2017 Asset pricing with countercyclical household consumption risk. *Journal of Finance* 72:415–60.

Davis, M. A., and S. Van Nieuwerburgh. 2015 *Housing, finance, and the macroeconomy*. In Handbook of Regional and Urban Economics, eds. G. Duranton, J. V. Henderson, and W. C. Strange, vol. 5, pp. 753âŁ"811. Amsterdam, the Netherlands: Elsevier.

Deaton, A. S., and C. Paxson. 1994 Saving and aging in Taiwan. In *Studies in the Economics of Aging*, 331–62. Cambridge, MA: National Bureau of Economic Research.

Fagereng, A., C. Gottlieb, and L. Guiso. 2017 Asset market participation and portfolio choice over the life-cycle. *Journal of Finance* 72:705–50.

Favilukis, J. 2013 Inequality, stock market participation, and the equity premium. *Journal of Financial Economics* 107:740–59.

Flavin, M., and T. Yamashita. 2002 Owner-occupied housing and the composition of the household portfolio. *American Economic Review* 92:345–62.

Gomes, F., and A. Michaelides. 2005 Optimal life-cycle asset allocation: Understanding the empirical evidence. *Journal of Finance* 60:869–904.

Gomes, F. J., L. J. Kotlikoff, and L. M. Viceira. 2008 Optimal life-cycle investing with flexible labor supply: A welfare analysis of life-cycle funds. *American Economic Review* 98:297–303.

Guiso, L., P. Sapienza, and L. Zingales. 2008 Trusting the stock market. *Journal of Finance* 63:2557–600.

Guvenen, F.. 2009 An empirical investigation of labor income processes. *Review of Economic Dynamics* 12:58–79.

Guvenen, F., F. Karahan, S. Ozkan, and J. Song. 2021 What do data on millions of U.S. workers reveal about lifecycle earnings dynamics? *Econometrica* 89:2303–39.

Guvenen, F., S. Ozkan, and J. Song. 2014 The nature of countercyclical income risk. *Journal of Political Economy* 122:621–60.

Guvenen, F., S. Schulhofer-Wohl, J. Song, and M. Yogo. 2017 Worker betas: Five facts about systematic earnings risk. *American Economic Review* 107:398–403.

Huggett, M., and G. Kaplan. 2016 How large is the stock component of human capital. *Review of Economic Dynamics* 22:21–51.

Jagannathan, R., and N. Kocherlakota. 1996 Why should older people invest less in stocks than younger people? *Quarterly Review* 20:11–23.

Kahn, L. B. 2010 The long-term labor market consequences of graduating from college in a bad economy. *Labour Economics* 17:303–16.

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

Landvoigt, T., M. Piazzesi, and M. Schneider. 2015 The housing market(s) of san diego. *American Economic Review* 105:1371–407.

Li, G. 2014 Information sharing and stock market participation: evidence from extended families. *Review of Economics and Statistics* 96: 151–60.

Lynch, A. W., and S. Tan. 2011 Labor income dynamics at business-cycle frequencies: Implications for portfolio choice. *Journal of Financial Economics* 101:333–59.

Mankiw, N. G. 1986 The equity premium and the concentration of aggregate shocks. *Journal of Financial Economics* 17:211–19.

Merton, R. C. 1971 Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory* 3:373–413.

Parker, J. A., and A. Vissing-Jorgensen. 2010 The increase in income cyclicality of high-income households and its relation to the rise in top income shares. *Brookings Papers on Economic Activity* 41:1–70.

Peijnenburg, K. 2018 Life-cycle asset allocation with ambiguity aversion and learning. *Journal of Financial and Quantitative Analysis* 53:1963–94.

Salgado, S., F. Guvenen, and N. Bloom. 2019 Skewed business cycles. Working Paper, University of Pennsylvania.

Schmidt, L. D. W. 2016 Climbing and falling off the ladder: Asset pricing implications of labor market event risk. Working Paper, MIT Sloan School Management.

Sinai, T., and N. S. Souleles. 2005 Owner-occupied housing as a hedge against rent risk. *Quarterly Journal of Economics* 120:763–89.

Storesletten, K., C. Telmer, and A. Yaron. 2007 Asset pricing with idiosyncratic risk and overlapping generations. *Review of Economic Dynamics* 10:519–48.

Vestman, R. 2019 Limited stock market participation among renters and homeowners. *Review of Financial Studies* 32:1494–535.

Viceira, L. M. 2001 Optimal portfolio choice for long-horizon investors with nontradable labor income. *Journal of Finance* 56:433–70.

Wachter, J. A., and M. Yogo. 2010 Why Do Household Portfolio Shares Rise in Wealth? *Review of Financial Studies* 23:3929–65.

Yao, R., and H. H. Zhang. 2004 Optimal consumption and portfolio choices with risky housing and borrowing constraints. *Review of Financial Studies* 18:197–239.