Richer earnings dynamics, consumption and portfolio choice over the life cycle

Julio Gálvez and Gonzalo Paz-Pardo*

January 17, 2025

Abstract

Households face earnings risk which is non-normal and varies by age and over the income distribution. We show that allowing for rich features of earnings dynamics, in the context of a structurally estimated life-cycle portfolio choice model, helps to rationalize the limited stock market participation and the low risky asset holdings of households. Because people are subject to more background risk than previously considered, the estimated model implies a substantially lower coefficient of risk aversion and lower stock market participation costs. Older workers and higher earners are exposed to negatively skewed risk and choose lower stock exposures.

Keywords: Portfolio choice, life cycle, earnings dynamics, household finance, simulated method of moments.
JEL Codes: G11, G12, D14, D91, J24, H06

---

*Julio Gálvez: CUNEF Universidad. Gonzalo Paz-Pardo: European Central Bank. We would like to thank the editor, Toni Whited, and an anonymous referee for very helpful and constructive comments. We thank seminar and conference participants at the Banco de España, the International Pension Workshop, the ECB, Netspar, the EEA-ESEM, Bristol, Rome Tor Vergata, CEMFI, CUNEF and the Bundesbank for their comments, and Manuel Arellano, Laurent Calvet, Sylvain Catherine, Julio A. Crego, Giulio Fella, Francisco Gomes, Luigi Guiso, Tullio Jappelli, Pierre-Carl Michaud, Pedro Saffi, Nawid Siassi, Paolo Sodini and Fernando Zapatero for helpful feedback. The views expressed here are those of the authors and are not necessarily the views of the European Central Bank or the Eurosystem.

1 Introduction

The risk that households face in the labor market is a key determinant of their portfolio decisions. For most workers, particularly for the young, their expected future labor market income is the largest asset they own. If this *human wealth* is risk-free, households may find it optimal to invest a large share of their financial wealth in risky, high-return investments such as stocks. If, instead, idiosyncratic income risk is large, labor market income becomes more stock-like and acts as a substitute for stocks in households' asset allocations (Viceira (2001), Huggett and Kaplan (2016)), leading them to tilt their portfolios toward safer assets.

Thus, studying household portfolios requires a good understanding of earnings dynamics, which vary by age and display non-normal and nonlinear features, as recent literature has shown (Guvenen, Karahan, Ozkan and Song (2021), De Nardi, Fella and Paz-Pardo (2020)). For instance, earnings tend to be less persistent for young workers with low incomes, who change jobs frequently. Instead, older workers with median earnings usually have very stable income flows, but face larger negative skewness driven by events which are infrequent but can be of large magnitude, such as job loss.

In this paper, we study the effect of these rich labor income dynamics on household consumption, savings, and portfolio allocations over the life cycle. We use a flexible earnings process that allows us to capture these features in a parsimonious and agnostic way (Arellano, Blundell and Bonhomme, 2017) and we compare it with the linear, canonical earnings process that is frequently used in the literature, but is restrictive. We estimate both processes using US data from the recent waves of the Panel Study of Income Dynamics (PSID) and use them as input to a life-cycle model of portfolio choice with housing, where households split their savings in financial assets into risk-free assets or risky stocks, subject to potential entry and per-period stock market participation costs. We estimate our model via indirect inference to match, separately for each earnings process, a wide set of features that characterize saving choices by US households, including stock market participation and its dynamics, wealth to income ratios, homeownership rates and the portfolio shares of stocks, exploiting the rich cross-sectional data from the Survey of

Consumer Finances (SCF) and the PSID panel. We also verify whether the estimated structural models can match features not targeted in the data, such the life-cycle profiles of wealth, the risky share, stock market participation, and the conditional housing share.

We find that the model with a nonlinear earnings process, compared to that with a canonical earnings process, can better explain the limited participation in the stock market with a much lower coefficient of risk aversion. Because human wealth is more stock-like than that implied by the canonical process, the coefficient of risk aversion that is required to rationalize household portfolio decisions drops from 11.16, which is in the ballpark of standard models that match limited participation and low risky shares (e.g., Cocco, Gomes and Maenhout (2005), Fagereng, Gottlieb and Guiso (2017)), to 6.83. This estimate is closer to microeconometric estimates that elicit the relative risk aversion coefficient via survey data, which is around 4 (Guiso and Sodini (2013)). At the same time, the nonlinear earnings process implies much lower annual participation costs into the stock market (135 dollars per year) than the canonical process (390 per year).

All layers of flexibility of our earnings process are key for our results. First, the age-dependence of earnings shocks allows us to take into account that older workers still face substantial earnings risk in the form of infrequent, but potentially large and persistent negative shocks. Second, the non-normality of income shocks further reduces the optimal portfolio share, given that, *ceteris paribus*, households want to insure against the possibility of receiving large negative shocks to their earnings (i.e., negative skewness). This feature, which is at odds with the canonical model, reduces the certainty equivalent valuation of future labor earnings at a given level of correlation between asset returns and earnings shocks, and raises the need for precautionary saving. Third, the nonlinearity in earnings shocks allows us to incorporate the fact that earnings risk is larger for relatively higher earners. As a result of this nonlinearity and of the age dependence, income risk varies endogenously over the wealth distribution. Therefore, while in the canonical process the optimal portfolio share of stocks is always decreasing in net worth-to-income ratios, that is not the case neither in the nonlinear process nor in the data.

The nonlinear process also generates a relationship between income risk and the risky share, conditional on age, income, and wealth, that looks like the one in the data, with the risky share mildly decreasing in the coefficient of variation for income as defined in Arellano, Bonhomme, De Vera, Hospido and Wei (2022). When labor income risk increases sharply, which can happen in the nonlinear process due to the presence of non-normalities and nonlinearities in persistence, households become less aggressive in their investments and reduce their risky share exposures. In the canonical process, instead, this relationship is overestimated because there is too little variation in cross-sectional income risk.

Our more realistic modelling of earnings risk also affects optimal investment advice, the welfare costs of suboptimal investment and the ability of households to insure their income fluctuations depending on their stockholding position. For instance, looking at a 50-year old homeowner with relatively low wealth (\$200,000) but median earnings, the canonical model recommends an exposure into stocks of approximately 40% of the financial portfolio. The richer nonlinear process, instead, acknowledges that the worker can still suffer sizeable income shocks and suggests a more conservative strategy of 20% into stocks. We also find renewed support for the rule-of-thumb strategy of investing $(100 - age)\%$ of one's wealth into risky assets, which turns out to be closer to optimal once we consider the relatively large standard deviation and negative skewness of earnings at later ages. We find that stockholders are better insured against income shocks as opposed to non-stockholders, both in our model and in the data, as measured by Arellano et al. (2017) and Blundell, Pistaferri and Preston (2008) partial insurance coefficients.

Recent work has emphasized the importance of non-normal features of earnings dynamics over the business cycle to explain limited household risk-taking (Catherine (2022), Catherine, Sodini and Zhang (2024), Shen (2024)). Relative to these papers, we focus on idiosyncratic earnings fluctuations over the life-cycle rather than aggregate shocks. This choice is motivated by the large costs associated with idiosyncratic shocks (equivalent to up to 25-30% of lifetime consumption according to Storesletten, Telmer and Yaron (2004) or De Nardi et al. (2020)) and allows us to reproduce the rich interaction between savings motives and earnings dynamics at different ages and points of the income distribution. At the same time, our results take into account the existence of correlation between earnings shocks and stock market returns. Our analysis highlights that non-normal, nonlinear risks over the life cycle have a quantitatively important role in explaining household portfolio decisions.

Related literature. This paper contributes to a broad literature in household finance that studies the causes of limited stock market participation (Gomes, Haliassos and Ramadorai (2021)). Several papers look at the roles of disaster risk (Fagereng et al. (2017)), housing (Cocco (2005)), trust (Guiso, Sapienza and Zingales (2008)), lack of investor sophistication (Haliassos and Bertaut (1995), Calvet, Campbell and Sodini (2007)), health risk (Rosen and Wu (2004)), wealth (Calvet and Sodini (2014), Briggs, Cesarini, Lindqvist and Östling (2015)), the presence of participation costs (e.g., Vissing-Jorgensen (2002) Alan (2006), and Bonaparte, Korniotis and Kumar (2020)) and non-homothetic preferences (Wachter and Yogo (2010), Meeuwis (2020)). We contribute to this literature by highlighting the role of age dependence, nonlinearity and non-normality in earnings risks, thus shedding new light on the link between background risk and portfolio choice decisions (see Guiso, Jappelli and Terlizzese (1996) for an early contribution).

Our analysis is focused around a life-cycle model of household portfolio choices, building on the seminal work of Cocco et al. (2005). Subsequent papers have looked at the roles of habit formation (Gomes and Michaelides (2003)), income volatility (Chang, Hong and Karabarbounis (2018)) and personal disaster risk (Nicodano, Bagliano and Fugazza (2021)). We show that the introduction of a richer earnings process yields more reasonable estimates of structural parameters in this class of models that are closer to those found in previous empirical work, while maintaining a parsimonious yet realistic model structure.

Catherine (2022), Catherine et al. (2024) and Shen (2024) point out that, because chances of large negative earnings shocks are higher in recessions, at a time in which stock returns are particularly low, households optimally reduce their equity shares. We diverge from their approach in two ways. First, our semiparametric formulation of the earnings process is very flexible and allows us to be agnostic about the specific characteristics of earnings dynamics and let the data inform our earnings process directly (De Nardi et al., 2020). In contrast, parametric processes based on a mixture of normals as in Guvenen et al. (2021), although they capture the nonlinear dynamics of the data well, require a precise specification of the structure of shocks to capture the features of income data and impose more structure (e.g., linearity conditional on each shock). Second, we focus on earnings dynamics over the life cycle and over the income distribution, rather than on business cycle variation. This choice is motivated by the large cross-sectional heterogeneity in the distribution of earnings shocks. For example, as Guvenen, Ozkan and Song (2014) show, the skewness of earnings changes varies much more over the income distribution (between -0.5 and -1.5) than it does between expansions and recessions (-1.25 to -1.75 for the median earner).

Our work also complements Athreya, Ionescu and Neelakantan (2023), who link low stock market participation at young ages with human capital accumulation decisions. Consistently with their framework, we find that expected labor market earnings are key for the portfolio decisions of the young. Although we do not explicitly model human capital accumulation, movements along the job ladder (Lise (2013)), or health shocks (e.g., Edwards (2008), Yogo (2016)), we replicate the dynamics of labor earnings across the life-cycle flexibly from the data and use them to study not only the stock market participation decision, but also conditional risky shares.

The rest of the paper is organized as follows. Section 2 discusses the models of earnings dynamics that we consider for our quantitative exercise. Section 3 presents the structural model that we estimate. We present the estimation results, the intuition underlying the underlying the structural model's estimation, and their robustness to alternative model specifications in Section 4. We analyze the implications for investment advice, the subsequent welfare costs of suboptimal investment, and consumption in Section 5. Finally, Section 6 concludes. We provide further details and robustness checks in the Appendix.

2 Earnings dynamics

Earnings dynamics are key to understand household consumption, saving, and portfolio decisions, and are a crucial ingredient in the calibration and estimation of life-cycle models. Recent empirical literature has called into question the long-established view that earnings dynamics are well-represented by a linear model. In particular, Arellano et al. (2017) and Guvenen, Karahan, Ozkan and Song (2016) present evidence that, contrary to the implications of the linear model, pre-tax household earnings exhibit deviations from log-normality, nonlinearity and age-dependence of moments.

In this section, we describe the rich features of residualized *disposable earnings*[1], as in De Nardi et al. (2020), and contrast the two models of earnings dynamics. We utilize the 1999 to 2017 waves of the PSID, as they provide information on consumption, income and assets for a representative panel of US households, which we exploit for the structural estimation. We detail the dataset construction in Appendix A.

2.1 Rich features of earnings dynamics

Higher-order moments of earnings are dependent on household age and on previous earnings. Figure 1 shows that both the conditional standard deviation (left column) and skewness (middle column) of household post-tax earnings growth become larger (in the case of skewness, more negative) as people grow older. Moreover, they show that these moments change across the income distribution. More specifically, the conditional standard deviation (left) exhibits a U-shaped pattern, decreasing until the 40th percentile and increasing from the 70th percentile upwards. Meanwhile, conditional skewness (middle) presents a decreasing pattern across the income distribution; in particular, skewness is more negative for higher earning percentiles, and for households in the age group from 54 above. These results imply that the distribution of earnings changes deviates substantially from the case of normal, age-independent shocks.

The right column of Figure 1 shows that earnings persistence is also highly nonlinear. We represent it as a function of the percentiles of the household's past earnings (τ_{init})

[1] To obtain the residualized data, we regress log disposable household earnings on a set of demographics and cohort dummies.

and the current earnings shock that the household received (τ_{shock}). Persistence for high ranked households receiving extremely negative shocks and low ranked households receiving extremely positive shocks is particularly low, in the range of 0.25. This implies that, for example, for a relatively high earning household, a large negative shock can effectively erase the memory of previous good shocks. Instead, persistence is much higher for high-ranked households consistently receiving positive shocks.

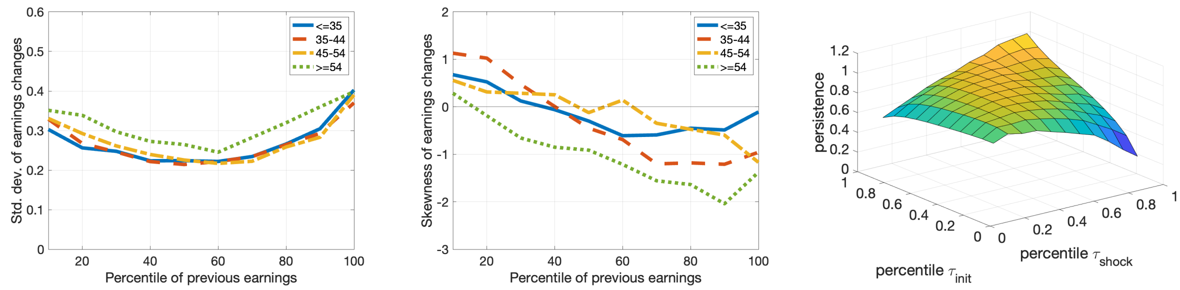

Figure 1: The figures on the left and middle columns show the standard deviation (left) and skewness (middle) of earnings changes, computed as a function of the household's position in the income distribution, divided into age groups. The right figure presents the average derivative of the conditional quantile function of household earnings y_{it} given y_{it-1}, with respect to y_{it-1}, computed from the previous percentile of the household's position in the income distribution (τ_{init}) and the shock (τ_{shock}). Data: PSID 1999-2017.

2.2 Modeling earnings dynamics

We first present the canonical model of earnings dynamics before discussing its nonlinear generalization in Arellano et al. (2017).

Consider households indexed by $i = 1, \ldots, N$ observed from age $t = 1, \ldots, T$. We decompose log earnings y_{it} as the sum of deterministic ($f(X_{it}; \theta)$) and stochastic components:

$$y_{it} = f(X_{it}; \theta) + \eta_{it} + \varepsilon_{it}, \; t = 1, \ldots, T. \tag{1}$$

The first stochastic component, η_{it}, is persistent and follows a first-order Markov process. The second component, ε_{it}, is transitory in nature, and has zero mean, independent of the persistent component, and independent over time.

The *canonical* model of earnings dynamics is described by the following process:

$$\eta_{it} = \rho\eta_{it-1} + u_{it} \tag{2}$$

$$\eta_{i0} \sim N(0, \sigma_z^2), u_{it} \sim N(0, \sigma_u^2), \varepsilon_{it} \sim N(0, \sigma_\varepsilon^2). \tag{3}$$

As emphasized by Arellano et al. (2017) and De Nardi et al. (2020), among others, the canonical process imposes the following restrictions:

1. *Linearity* of the process of the persistent earnings component. Linearity implies that the right hand side of equation (2) is additively separable to the conditional expectation and the innovation u_{it}.

2. *Normality* of the shock distributions. Normality implies that the shock distributions are symmetric, and should not exhibit skewness.

3. *Age-independence* of the autoregressive component ρ and the moments of the shock distributions, which imply the age independence of second and higher-order moments of the conditional distributions of the earnings components.

Given that these assumptions are at odds with the empirical evidence, Arellano et al. (2017) propose a general representation of the income process that allows for nonlinearity, non-normality, and age-dependence. In particular, the persistent component of income[2] is modelled as the following process:

$$\eta_{it} = Q_t(\eta_{it-1}, u_{it}), \quad (u_{it}|\eta_{it-1}, \eta_{it-2}, \ldots) \sim U[0,1], \quad t = 2, \ldots, T. \tag{4}$$

where $Q_t(\eta_{it-1}, \tau)$ is the τ-th conditional quantile function of η_{it} given η_{it-1} for a given τ. Intuitively, the quantile function maps random draws from the uniform distribution u_{it} (i.e., cumulative probabilities) into corresponding random draws (i.e., quantile) from the persistent component. We discuss the features of the *nonlinear* process in Appendix B.1.

[2]Meanwhile, Arellano et al. (2017) model the initial distribution of the persistent component η and the transitory component ε via similar quantile representations. We describe the estimation of both nonlinear and canonical processes in Appendix B.

The Arellano et al. (2017) process has direct links to structural labor market models, such as the job ladder models in Lise (2013) and Huckfeldt (2022). Consider in particular, the following example of an unusual negative shock: that of an old-age worker who receives an adverse occupation-specific shock which leads to job loss. In this case, the previous earnings history of this worker matters less long after the income shock. In this context, the nonlinear process captures the notion of "microeconomic disasters", in the tradition of the disaster risk literature. One clear difference is that, in comparison with macroeconomic disasters, microeconomic disasters happen more frequently and are easier to identify empirically.

Comparing canonical and nonlinear processes. In Appendix B.4, we compare the implications of the two processes. The results that we obtain imply that the nonlinear process is able to capture well the features of earnings data we just described, while the canonical process, by construction, cannot.

3 Model

We introduce both the canonical and nonlinear earnings processes into a standard discrete time, life-cycle portfolio choice model with housing and study their implications.

Demographics Households start working life at 25, face age-dependent positive death probabilities, and die with certainty at age 100. The model period is two years.

Preferences Households maximize:

$$\max \mathbb{E}_t \left[\sum_{t=0}^{t=T} \beta^t \mathcal{S}_t \frac{[c_t^{\nu}(1+\psi I_t^h) h_t^{1-\nu}]^{1-\gamma}}{1-\gamma} \right] \tag{5}$$

where c is nondurable consumption, h is the consumption of housing services, ψ is the relative preference for owner-occupied housing, I^h is an indicator for a households's home-ownership status, ν is the relative preference for nondurable consumption goods over housing, γ is the coefficient of relative risk aversion, β is the discount factor, and $\mathcal{S}_t$ is the probability of survival up to time t.

Earnings process As described in Section 2.2, we assume that log earnings can be decomposed to a persistent and a transitory component (Equation 1). We use alternatively, the canonical and the nonlinear specifications for both components of the earnings process. There is no earnings risk after retirement (age 65), from which households get a public pension.

Housing Households can rent or buy their house, which we denote with the indicator variable for homeownership $I^h_{it} = \{0, 1\}$. Housing is available in fixed sizes $h_{it} \in \{H_1, H_2, \ldots, H_H\}$ where H_1 denotes the smallest and H_H the largest house. The price of a house is proportional to its size:

$$p^{\mathcal{H}}(h_{it}) = p^h_{it} h_{it} \tag{6}$$

where p^h_{it} denotes the housing price per unit of housing. Every period, renters can decide whether they want to keep renting the same house, renting a house of a different size, or buy a house of the size they choose. Similarly, homeowners decide whether they want to stay in the house they own and live in, whether they want to sell it and become renters, or whether they want to sell it and buy a different-sized house. However, there are transaction costs involved with buying and selling a home, denoted by $\kappa^h(h_{t+1}, h_t)$, which we model as a fixed fraction of the house price for both buyer and seller. There are no transaction costs involved in renting.

$$\kappa^h(h_{t+1}, h_t) = \kappa_h p^{\mathcal{H}}(h_{it}) \text{ if } h_{it+1} \neq h_{it} \text{ and } I^h_t > 0 \tag{7}$$

$$\kappa^h(h_{t+1}, h_t) = \kappa_h p^{\mathcal{H}}(h_{it+1}) \text{ if } h_{it+1} \neq h_{it} \text{ and } I^h_{t+1} > 0 \tag{8}$$

House prices are risky at the idiosyncratic level and evolve according to the following process:

$$\log p^h_{it+1} = \log p^h_{it} + \epsilon^h_{it+1} \tag{9}$$

where $\epsilon^h_{it+1} \sim \mathcal{N}(0, \sigma^2_h)$. For simplicity, but also given that house prices are up to five times more volatile at the idiosyncratic and local level than at the national level (Piazzesi, Schneider and Tuzel, 2007), we abstract from aggregate house price volatility.

Renters pay yearly rent equal to a fraction ζ of national average house prices $\bar{p^h_{it}}$, which implies that rents are not risky:

$$r^h(h_{it}) = \zeta \bar{p^h_{it}} h_{it} \tag{10}$$

Mortgages Households can borrow to buy a house through short-term mortgages, up to a borrowing constraint in the form of a downpayment restriction:

$$x_{it} \geq \phi_H p^h_{it} h_{it} I^h_{it} \tag{11}$$

where x_{it} represents end-of-period net worth. Hence, the housing wealth share α_{it} can be larger than 1, but cannot exceed the inverse of the minimum required downpayment ϕ_H:

$$\alpha_{it} = \frac{p^h_{it} h_{it} I^h_{it}}{x_{it}} \leq \frac{1}{\phi_H} \tag{12}$$

Although house prices are risky, we assume that households cannot buy a house if there is a positive probability that they will have negative net worth in the following period, which avoids the need to explicitly model bankruptcy. This assumption is rarely restrictive, given that most of this risk is already prevented by the existence of a borrowing constraint.

Safe and risky financial assets. Households can save in two types of financial assets. Risk-free assets have a fixed rate of return r, while risky stocks have stochastic returns r^s_{t+1} which are i.i.d. We denote by π_{it} the share of net worth invested in the risky asset, which implies that the share of net worth invested in the risk-free asset is $1 - \pi_{it} - \alpha_{it}$.

We allow for correlation[3] between stock market returns and persistent shocks to income at the individual level:

[3] This feature intends to capture both individual bias towards owning stocks and shares of one's company or sector (see e.g., Betermier, Calvet and Sodini (2017)) and aggregate correlations between stock market returns and income shocks (see e.g., Betermier, Jansson, Parlour and Walden (2012) and Bonaparte, Korniotis and Kumar (2014)). However, as our model is a partial equilibrium one, whether this correlation is idiosyncratic or aggregate matters quite little.

$$r^s_{it+1} = (1 - \tilde{\lambda}^\eta) r^s_{t+1} + \tilde{\lambda}^\eta \eta^{shock}_{it+1} \tag{13}$$

where η^{shock}_{it+1} refers to the persistent income shock the household received between t and $t+1$, and $\tilde{\lambda}^\eta$ captures the correlation between stock returns and labor market income.

Participating in the stock market is costly, which we represent with the cost function κ^f, that depends on the households' stock market participation status $I^f_t = (\pi_t > 0)$. Following Vissing-Jorgensen (2002), these may either be per-period costs, κ^{PP} (just dependent on I^f_{t+1}), fixed but one-time κ^{FC} (only paid if $I^f_t = 0$ and $I^f_{t+1} = 1$, and zero if $I^f_t = 1$) or a combination of both:

$$\kappa^f(I^f_{t+1}, I^f_t) = \begin{cases} 0 & \text{if } I^f_{t+1} = 0 \\ \kappa^{FC} + \kappa^{PP} & \text{if } I^f_{t+1} = 1 \text{ and } I^f_t = 0 \\ \kappa^{PP} & \text{if } I^f_{t+1} = 1 \text{ and } I^f_t = 1 \end{cases} \tag{14}$$

The fixed cost can be understood as an entry cost to stock market participation, related to the time spent understanding the risks and returns associated with stocks. The per-period participation cost, meanwhile, can be understood as either the time spent in determining whether portfolio rebalancing is optimal[4] (if the household actively manages its portfolio) or the cost of delegating the investment decisions to a fund manager (if the household indirectly holds stocks via mutual funds).[5]

Budget constraint. The households' budget constraint can be expressed as follows, where x_{it} represents the sum of financial and housing wealth owned by the household at the end of period $t-1$ and before the realization of return shocks in period t:

$$\begin{aligned} c_{it+1} + x_{it+1} + \kappa^f(I^f_{it+1}, I^f_{it}) + \kappa^h(h_{it+1}, h_{it}) + r^h I(I^h_{it} = 0) = \\ \eta_{it} + u_{it} + x_{it}(r^s_{it}\pi_{it} + r(1 - \pi_{it} - \alpha_{it}) + (p_{it} - p_{it-1})\alpha_{it}) \end{aligned} \tag{15}$$

[4]An alternative rationalization of participation costs is related to psychological costs related to rebalancing stocks. One paper that considers these costs, within the context of mortgage markets, is Andersen, Campbell, Nielsen and Ramadorai (2020).

[5]There a third cost of stock market participation in Vissing-Jorgensen (2002), which is a proportional trading cost. We do not model this because neither the PSID nor the SCF provide information which allows us to identify trading costs.

We denote by z_{it} wealth after the realization of return shocks and the transitory component, i.e.:

$$z_{it} = u_{it} + x_{it}(r^s_{it}\pi_{it} + r(1 - \pi_{it} - \alpha_{it}) + (p_{it} - p_{it-1})\alpha_{it}) \tag{16}$$

Because α_{it} can be greater than 1, the second term inside the parenthesis in Equation 16 can be negative and represent mortgage costs.

Households' problem Households thus solve the following problem, where we drop the i subindex for simplicity:

$$V_t(z_t, \eta_t, I^f_t, I^h_t, h_t, p^h_t) = \max_{c_t, \pi_{t+1}, h_{t+1}, I^h_{t+1}} \left\{ \frac{[c^\nu_t(1 + \psi I^h_t)h^{1-\nu}_t]^{1-\gamma}}{1 - \gamma} + \beta \frac{\mathcal{S}_t}{\mathcal{S}_{t-1}} \mathbb{E}_t V_{t+1}(z_{t+1}, \eta_{t+1}, I^f_{t+1}, I^h_{t+1}, h_{t+1}, p^h_{t+1}) \right\} \tag{17}$$

subject to the budget constraint (15), the downpayment constraint (11) and short sale constraints $\alpha_{it} \geq 0$ and $\pi_{it} \geq 0$. Households can only borrow to buy a house, hence $\pi_{it} \leq 1$. The choices c_{it}, π_{it}, h_{it} and I^h_{t+1} imply savings x_{it}, the housing share α_{it} and the stock market participation status I^f_t. The expectation $\mathbb{E}_t$ is taken with respect to future realizations of persistent income, transitory income, stock market returns, and house prices. More specifically, the realization of stock market returns, together with the choices of x_{t+1}, π_{t+1} and h_{it+1} and the realization of the transitory component u_{it+1} implies next period's z_{t+1}, while the exogenous process for persistent labor market income determines η_{t+1} conditional on η_t, and the exogenous process for persistent local house prices determines p^h_{t+1} conditional on p^h_t.

We provide more details about the algorithm which we use to solve the households' problem and our discretization procedure in Appendix C.1.

4 Structural Estimation

We estimate our structural model via the simulated method of moments (SMM), conditional on the pre-estimated household labor income process and some externally set parameters.

4.1 Estimation strategy

4.1.1 External parameters

We set the risk-free rate to 2%, the equity premium to 4%, and the standard deviation of stock market returns to 0.157, following Cocco et al. (2005). We obtain survival probabilities from Bell, Wade and Goss (1992) and set public pensions to 70% of the average realization of earnings at retirement age (i.e., 35% of average income of workers in the economy).

We assume that transaction costs κ_h are 5% of the value of the house, annual rental costs ζ are 2.5% of the value of the house being rented, and the standard deviation of shocks to (log) housing prices σ_h is 0.1, within the range considered in Piazzesi and Schneider (2016). The share of housing in the Cobb-Douglas utility function is set to be $1 - \nu = 0.2$, and the price of the medium-sized house is 5 times average income.[6] The minimum required downpayment on a mortgage ϕ_H is 20%.

The correlation between stock market returns and labor market income shocks is set to 0.2. Although empirical measures of the correlation between idiosyncratic labor market income shocks and aggregate stock returns tend to be low and sometimes indistinguishable from 0 (e.g. Davis and Willen (2014)), earnings and stock returns might have stronger correlations over longer time periods or be cointegrated (e.g. Benzoni, Collin-Dufresne and Goldstein (2007)). Under a risk aversion parameter of 6, which is close to our baseline estimate for the nonlinear process, Huggett and Kaplan (2016) find that the correlation between stocks and earnings is such that around 20% of the total discounted valuation of human capital during the working life is equivalent to a long stock position.

We assume that 50% of our households start their working lives at 25 as homeowners, consistently with PSID data. We opt for the conservative assumption that they have minimum equity (20%) in their homes. However, because young households in our data have very little initial financial wealth, we assume that households start out their lives with zero financial wealth[7].

[6] Average house prices are endogenous and depend on the distribution of chosen house sizes.

[7] For instance, in our SCF sample, the average financial wealth of households less than 30 years old is approximately \$52,959.02, with median financial wealth at \$15,539.77. Meanwhile, the average housing wealth of households less than 30 years old is \$112,618.6, with a median at \$50,000.

4.1.2 Estimated parameters and targeted moments

We estimate γ, β, the stock market participation costs k^{FC} and k^{PP}, and the homeownership utility premium ψ within the model. We target eleven data moments for our estimation, which we obtain from the PSID and the SCF.[8] The first four moments are cross-sectional moments related to wealth, which we obtain from the SCF because it has more comprehensive information than the PSID about household wealth and household portfolios. In particular, an advantage of the SCF over the PSID is that it provides information on the richest households because they are oversampled.[9] Specifically, we target the percentage of people that own stocks (0.67), including both direct and indirect stock ownership –for example through ETFs or mutual funds–, average wealth-to-income ratios (6.1), the conditional risky share, defined as the value of risky assets divided by net worth (0.27), and the homeownership rate (0.79). Because we are also interested in that our model matches the dynamic and cross-sectional aspects of stock market participation, we turn to the panel data in the PSID and estimate an OLS regression of a stock ownership dummy on a polynomial in age, indicators of homeownership and past stock market participation, income, and wealth, and target its coefficients inside the model. The regression that we estimate can be considered as an empirical policy function for stock market participation, in the spirit of Bazdresch, Kahn and Whited (2018). We also target its parameters in our estimation.[10]

4.1.3 Estimation method

We outline the SMM estimation procedure here. Let d_{it} be the vector of data observations. Let $d^s_{it}(\theta)$ be a simulated vector from simulation s, for $s = 1, \ldots, S$, which depends on the structural parameters of the model, θ. In our context, the structural parameters are γ, β, ψ, and the κ's. Next, we define the following vector of estimating equations:

$$g(d_{it},\theta) = \begin{bmatrix} g_1(d_{it},\theta) \\ g_2(d_{it},\theta) \end{bmatrix} = \begin{bmatrix} \frac{1}{N_1}\sum_{i=1}^{N_1}\left(m_1(d_{it}) - \frac{1}{S}\sum_{s=1}^{S} m_1(d^s_{it}(\theta))\right) \\ \frac{1}{N_2T_2}\sum_{i=1}^{N_2}\sum_{t=1}^{T_2}\left(m_2(d_{it}) - \frac{1}{S}\sum_{s=1}^{S} m_2(d^s_{it}(\theta))\right) \end{bmatrix}$$

[8] We give further details about our sample selection and precise variable definitions in Appendix A.

[9] We have also alternatively conducted our estimation using exclusively moments from the PSID, and results are very similar.

[10] Appendix A.4 provides more details about the computation of these moments and the estimation of this regression in the data.

in which $g_1(d_{it}, \theta)$ corresponds to the vector of moments $m_1(\cdot)$ from the SCF and $g_2(d_{it}, \theta)$ corresponds to the parameters of the empirical policy function of stock market participation estimated from the PSID, $m_2(\cdot)$.

The SMM estimator is the solution to the minimization of the following quadratic form:

$$\widehat{\theta} = \min_{\theta} g(d_{it}, \theta)' W g(d_{it}, \theta). \tag{18}$$

wherein W is the optimal weighting matrix estimated using the influence function technique in Erickson and Whited (2002). Specifically, W is the inverse of the clustered covariance matrix Ω of the influence functions of the data moments $m(d_{it})$, $\psi_{m(d_{it})}$. We compute the clustered covariance matrix by stacking the influence functions for the elements of the data moments, and computing a clustered covariance. Because we obtain the samples from two different, independently sampled, but complementary datasets, the clustered covariance matrix is block-diagonal.[11] More details of the calculation of the weighting matrix, and the computation of standard errors are in Appendix C.2.

	Model	
Parameter	Nonlinear	Canonical
γ	6.83 (0.6630)	11.16 (0.2861)
β	0.860 (0.0335)	0.892 (0.0121)
κ^{FC}	0 (0.0194)	0 (0.0704)
κ^{PP}	0.0018 (0.0087)	0.0052 (0.0006)
ψ	0.1053 (0.0094)	0.0117 (0.0059)

Table 1: Parameter estimates (standard errors in parentheses). β is expressed in annual terms. The participation costs are expressed as fractions of average household income, which is the numeraire in the model.

4.2 Estimated parameters and model fit

The models for the nonlinear and canonical processes imply remarkably different estimated parameters (Table 1) despite fitting our data targets similarly well (Table 2). Most notably, the implied CRRA risk aversion parameter is substantially lower (6.83) under the nonlinear process than it is under the canonical process (11.16). With richer,

[11]See Arellano and Meghir (1992) and Ridder and Moffitt (2007) for a discussion of estimating standard errors when data come from two independent datasets.

		Model	
Moment	Data	Nonlinear	Canonical
Participation	0.677	0.679 [0.476]	0.677 [0.000]
Risky share	0.259	0.259 [0.000]	0.257 [-0.451]
Average W/I	5.599	5.626 [0.367]	5.572 [-0.367]
Homeownership	0.795	0.797 [0.451]	0.794 [-0.225]
OLS constant	-0.581	0.980 [0.119]	-1.559 [-0.074]
OLS, past participation	0.448	0.496 [0.043]	0.540 [0.083]
OLS, age	-0.012	-0.049 [-0.046]	0.031 [0.073]
OLS, age^2	0.0001	0.00058814 [0.068]	-0.0002778 [-0.062]
OLS, log income	0.033	-0.091 [0.217]	0.025 [-0.014]
OLS, log wealth	0.064	0.110 [0.154]	0.066 [0.007]
OLS, homeownership	-0.074	0.062 [0.102]	0.006 [0.060]

Table 2: Targeted vs. model-implied moments. $t-$statistics that report the differences between the data moment and the model-implied moment are reported in brackets.

more realistic earnings risk, the certainty equivalent valuation of future labor earnings goes down, and households optimally reduce their allocation to stocks at a given coefficient of relative risk aversion γ. As a result, the nonlinear process generates lower risky shares over the life cycle; thus, the calibrated coefficient of relative risk aversion does not need to be as large as in the case of the canonical process. This mechanism holds true even in the presence of correlation between labor market income and stock market returns, which reduces the extent to which risky labor market income is a substitute for stocks.

Following the same intuition, households on the margin of participating in the stock market are more likely to choose not to do so under richer, more realistic earnings risk than under the canonical earnings process. Hence, the model with the nonlinear earnings process can explain the observed patterns of limited stock market participation with much lower stock market per-period participation costs: in dollar terms, they amount to 135 per year, almost three times less than those estimated with the canonical process (390 per year). However, this coefficient is structurally estimated with relatively less precision compared to its magnitude.

Both models generate zero entry costs to the stock market. Given that we are targeting not only observed stock market participation but also its dynamics over time, this result suggests that nonlinear earnings dynamics do not provide additional evidence in favor of stock market entry costs, and that both of our models imply that the observed persistence in stockowner status is more related to endogenous selection (for example, higher wealth people are more likely to be higher wealth in the following period, and also more likely to be stockowners in both periods) than one-off costs associated with entering the stock market.

The discount rate β is relatively low for both versions of the model, which is common in life-cycle savings models that target observed wealth-to-income ratios and incorporate high-return assets such as stocks. This is compounded by the presence of housing as an asset that provides direct utility (see e.g. Fagereng et al. (2017) for the former case and Paz-Pardo (2024) for the latter). It is lower for the nonlinear than for the canonical process because nonlinear earnings risk generates additional precautionary saving ceteris paribus (De Nardi et al., 2020). The nonlinear process requires a higher homeownership utility premium to match observed homeownership, which is consistent with housing being a risky asset in our framework.

Table 2 shows the model fit by comparing our targets in the data (left column) with the model implications under the nonlinear (central column) and those under the canonical processes, respectively (right column). In brackets, we report the result of a t-test which that compares the data moments with the simulated moments of that particular model, following Nikolov and Whited (2014). For both the nonlinear and the canonical earnings process, the model fits its targets remarkably well given how parsimoniously parameterized it is (we estimate 5 parameters to fit 11 targets), with all of the model-implied moments not being statistically different from the empirical data moments. In particular, the model closely replicates the limited level of stock market participation that we observe in the data and the very low conditional risky share of stockholders, two crucial moments to understand the savings and portfolio decisions of US households (Alan (2012) and Bonaparte et al. (2020)). We also replicate very closely the average levels of household wealth accumulation and the homeownership rate.

With respect to the OLS regression of the determinants of stock market participation, both processes do a good job in replicating the level of persistence in stock market participation that we observe in the data (the coefficient on past participation is 0.45 in the data, 0.50 in the nonlinear model and 0.54 in the canonical model), which is key for the identification of participation cost parameters and, in particular, to determine our zero estimated entry costs.

Both models match the fact that richer households are more likely to invest in the stock market, but overestimate the extent to which this happens. The nonlinear process generates a counterfactual negative effect of income on stock market participation; however, because the coefficient is estimated with relative imprecision in the data, partially because it is very closely correlated with other covariates such as wealth, the difference between the data and the nonlinear model is not statistically significant. Neither model replicates the fact that homeowners, conditional on all other variables, are less likely to participate in the stock market.

In Appendix C.3 we give further information about how the key moments we are interested in replicating help us identify our estimated parameters.

4.3 Empirical policy functions

We now turn to describing how the estimated structural models match key facts related to stock ownership and conditional risky shares over the life cycle. These moments, which we do not explicitly target in the estimation, are also informative about the relevant features of the nonlinear process that help explain portfolio decisions and generate a lower estimated parameter of risk aversion and lower participation costs.

Life-cycle implications of both processes. The first four panels of Figure 2 show how the two structural models match the life-cycle counterparts of the moments related to household wealth accumulation and portfolio decisions whose average we explicitly target in our estimation. Overall, the figures show that both models do a good job in replicating these profiles. In particular, looking at the first row, we observe that both nonlinear and canonical processes fit very closely the profile of average wealth accumulation (left) and the homeownership rate (right).

While both processes overestimate the slope of the stock market participation profile (middle left), the nonlinear process generates a relatively flatter profile that is closer to the empirical counterpart, particularly between households aged 30 to 50. This success is driven by the relatively lower participation costs which the model with the nonlinear process needs to rationalise the observed average participation. With relatively higher participation costs, the canonical process underestimates stock market entry for the high-income young. When households are relatively older and richer, both processes imply a higher participation rate than the one we observe in the data.

With respect to the conditional risky share (middle right), both processes generate a relatively flat share during most of the working age, but slightly overestimate the amount of stocks held by the oldest households, which may be related to the fact that the model is not designed to capture relevant sources of risk during the retirement period, such as medical expense risk. While in the case of the canonical process the low risky shares are mostly driven by the high coefficient of relative risk aversion, the nonlinear process succeeds in obtaining a low and flat risky share with a much lower CRRA coefficient since it correctly replicates the earnings risk faced by households at different ages.

The bottom two figures show two moments which are not explicitly targeted in the estimation. In the bottom left panel, we represent the unconditional risky share, where we find that the canonical process displays a linear growth which is inconsistent with the data, while the nonlinear process correctly implies a relatively flat profile between ages 20 and 50, only overestimating the risky share at the latest ages, similarly to the canonical process. Both models replicate the conditional housing share relatively well. This success reassures us that the relevant underlying patterns of homeownership and home equity, which may condition the optimal allocation of financial assets, are correctly captured in our model.

Portfolio choice and income risk. To further probe into the relationship between income risk and the risky share, we estimate an empirical policy function (EPF) for the determinants of the risky share, in which we explicitly consider the role of heterogeneous income risk using the coefficient of variation (CV) measure (Arellano et al., 2022). The

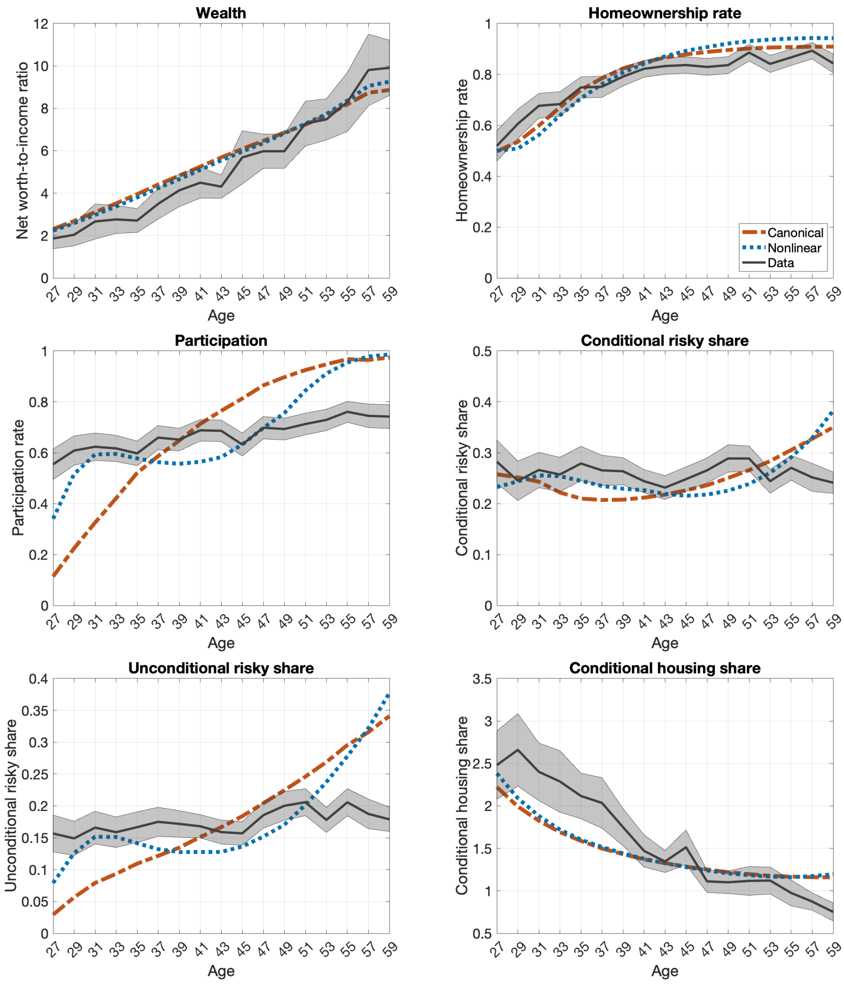

Figure 2: Life-cycle profiles implied by the structural models (canonical in red dash-dot, nonlinear in blue dotted), in comparison with data from the SCF (black solid). The empirical life cycle patterns are estimated using OLS regressions, following Deaton and Paxson (1994). The implied life cycles of the structural models are estimated using OLS regressions with age dummies. 95% point-wise confidence bands are shaded.

estimating equation is:

$$\text{Risky Share}_{it} = b_0 + b_1 \text{Coefficient of variation}_{it} + \mathbf{Z}'_{it}\gamma + \varepsilon_{it}. \tag{19}$$

The CV is a one-period ahead measure that summarizes the uncertainty in the predictive income distribution of the household. Specifically, it is the ratio of the mean absolute deviation of income (dispersion) and mean expected income (location). For example, a household with an expected income of 50,000 dollars and a CV of 0.1 expects a deviation of next year's income from its mean by ±5,000. Hence, low CV measures are associated with relatively low household income risk. We show how to compute the CV measure, which is a relatively simple prediction problem, in Appendix D.1.

The risky share is defined as the share of risky assets in net worth. We also control for demographic characteristics by introducing them into the vector $\mathbf{Z}_{it}$, which include quadratic polynomials in (log) income and (log) wealth, a homeownership dummy, and age dummies, following Guiso et al. (1996).[12] We then estimate the model via Tobit regressions. We compare the estimation results of those that are based on simulated data from the structural models with those based on the PSID data, and report the associated t-statistics in square brackets.

Note that these EPFs are close counterparts of the participation regression that we target in our estimation, with the risky share rather than stock market participation as the dependent variable, and with the addition of the coefficient of variation as a measure of income risk.

Dep. variable: risky share out of net worth	Data	Canonical	Nonlinear
Coefficient of variation	-0.087	-0.874 [10.340]	-0.058 [1.242]
Log of household income	-0.144	0.156 [-1.806]	-0.484 [7.583]
Log of household income (squared)	0.008	-0.011 [3.943]	0.015 [-4.034]
Log of household wealth	0.180	0.285 [-2.359]	-0.453 [15.198]
Log of household wealth (squared)	-0.002	-0.009 [2.478]	0.024 [-16.049]
Homeownership dummy	-0.145	-0.049 [-6.421]	-0.085 [-4.869]
Constant	-1.481	-2.358 [-10.074]	5.615 [35.316]

Table 3: Empirical policy functions: the determinants of the risky share and the role of income risk. All regressions control for age by age fixed effects. In the regressions using PSID data, we also control for time and cohort dummies, and demographics. t-statistics that report the difference between each of the models and the data are in brackets. Full estimation results of the data are in Table D1 of Appendix D.1.

The regression results, which are in Table 3, indicate that, on average, there is a negative relationship between income risk and the risky share. That is, if the perceived risk in labor income increases, at a given level of income and wealth, there is a shift towards riskless assets. Both the estimated models under the nonlinear and the canonical process exhibit this relationship; however, only the coefficient for the nonlinear process is not significantly distinguishable from the data. Instead, the canonical process overestimates this relationship, which is related to the fact that there is less variation in the distribu-

[12] In regressions using the PSID, we control for a host of household demographics, including marital status, education of both head and spouse, whether the household owns a business or not, family size, year and cohort dummies.

tion of income risk under the canonical process than under the nonlinear process. We can observe this by looking at the kernel densities of the estimated CV measure in Figure D1 in Appendix D.1.

Both processes replicate the negative correlation between homeownership and the risky share, but they both underestimate the extent to which this happens, although the nonlinear process does slightly better. Neither of the processes matches well the coefficients on the income and wealth polynomials. In the case of income, the nonlinear process replicates correctly the signs of both coefficients, but underestimates the linear term and overestimates the quadratic term. The canonical process gets the signs wrong, but is closer to the coefficients in the data in terms of magnitudes. In the case of wealth, the canonical process replicates better the marginal effects, while the nonlinear process implies a negative sign in the linear term and a positive sign in the quadratic term. These are the opposite of those in the data, but both biases will tend to cancel each other out. Overall, the conclusion from this analysis is that neither the canonical nor the nonlinear models can replicate well the marginal effects of income conditional on wealth and of wealth conditional on income, first because they are highly correlated and second because the estimation of the earnings process does not control for observed or unobserved characteristics that are potentially correlated with both of them and portfolio choices at the same time.

Conditional risky shares. We now turn to show that, despite missing some of the marginal effects of income and wealth, our model, particularly the one equipped with the nonlinear earnings process, replicates well the portfolio patterns over the wealth distribution that we observe in the data. To do so, we regress the conditional risky share on a set of wealth (and net worth-to-income) bins and age fixed effects. We then compare the implied empirical policy functions with the one we obtain from the SCF data[13].

The left panel of Figure 3 shows the conditional risky share as a function of the net worth-to-income ratio, while the right panel of Figure 3 shows the conditional risky share

[13] One concern is that the SCF variable for income is not the same as the one that we use in the PSID, which is disposable earnings. We repeat this exercise with the PSID panel data, and obtain similar results as in Figure 3.

across different bins of the wealth distribution. Standard models of portfolio choice imply that after controlling for age effects, conditional risky shares are decreasing functions of net worth-to-income ratios. The canonical model exhibits this feature, both with net worth-to-income and wealth. However, both the data and the nonlinear model display more complex, non-monotonic patterns, in which the conditional risky share is higher for people for whom their wealth to income ratios are either very low or relatively high, and lower for those whose wealth to income ratios are around or slightly below the average.

The nonlinear process can replicate these facts better than the canonical process because of the key role of labor market income in explaining optimal portfolio decisions. Namely, in the canonical process an increase in the relative weight of financial wealth in a household's portfolio always implies a lower conditional risky share because bond-like labor market income is becoming comparatively less relevant, which must be compensated with a safer financial portfolio; in the nonlinear process, where labor market income is more stock-like, this is not necessarily true and it depends on the relative risk faced by that particular household in the labor market. At relatively larger levels of the wealth-to-income ratio, households are frequently subject to relatively large labor market income risk[14]; as labor market income represents a progressively small part of those households' portfolios, households are ready to take on more risk in their financial portfolios.

While the nonlinear process does relatively well in explaining the relatively large risky share of wealthy households (right panel), it overestimates the risky share for high net worth-to-income households (left panel). This feature is probably driven by the fact that households with a net worth-to-income ratio over 10, which represent less than 10% of our sample, have specific features which the model cannot replicate (for example, they might have recently received a substantial inheritance).

In Appendix D.2, we show the results of estimating these same EPFs for three broad age groups. Although the estimates in the data are relatively noisy, we conclude that the model under the nonlinear earnings process is also better able to capture the age-dependent patterns found in the SCF data across all age groups. The canonical process,

[14]This can be observed from looking at the plots of the CV measure of Arellano et al. (2022) on the net worth-to-income ratio, which we report in the bottom right panel of Figure D2 of Appendix D.1.

meanwhile, does not match well the implied risky shares at higher levels of wealth for the different age groups.

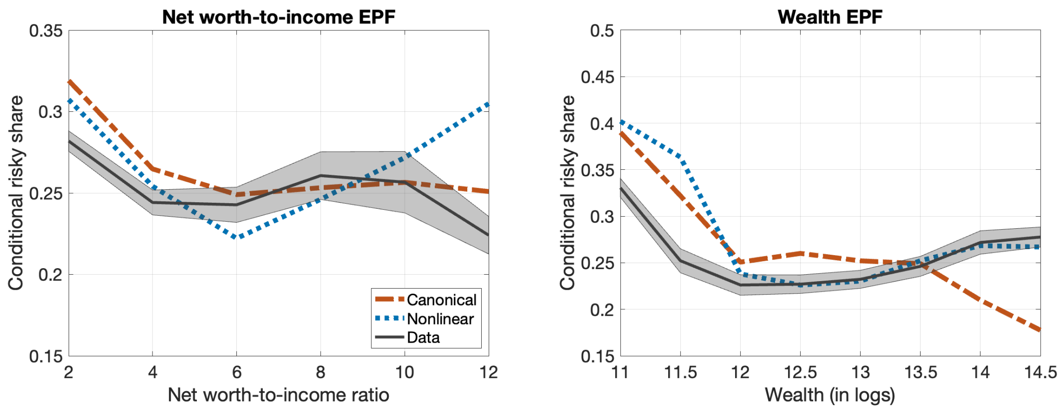

Figure 3: Empirical policy functions. The figures show the relationship between the conditional risky share and the net worth-to-earnings ratio (left) or wealth (right) that are implied by the structural models (canonical in red dash-dot, nonlinear in blue dotted), in comparison with data from the SCF (black solid). The EPFs are the predicted equity shares from a regression of the conditional risky share on bins of wealth (or the net worth-to-earnings ratio) and age fixed effects. In the data, the estimation also includes year fixed effects. 95% point-wise confidence bands are shaded, computed using robust standard errors.

4.4 Decomposing the role of earnings dynamics

As described in Section 2, our flexible nonlinear earnings process differs from the canonical process in several ways: age-dependence, non-normality of shocks, nonlinearities, etc. To gauge the relative contribution of these factors to explaining our results, in Table 4 we report the estimated parameters under a set of intermediate processes which we describe in more detail in Appendix B.3: one with age-dependence, but no non-normalities or nonlinearities, and one with age-dependence and non-normality, but no nonlinearities.

We find that allowing for age-dependent persistence and variance is not enough to generate substantial departures from the canonical model. The coefficient of risk aversion barely drops (11.16 to 11.08) and the per-period participation costs to the stock market actually increase, both suggesting that this model generates patterns of household behavior that are very similar to those under age-independent variance and persistence of earnings. It is only when we allow for the skewness and kurtosis of earnings shocks to differ from those of a normal distribution and to vary by age (third row) that we obtain a significant reduction of the coefficient of relative risk aversion, down to 7.88. When agents internalize that their earnings are subject to rare, but relatively large and negatively skewed shocks, they optimally invest less in stocks ceteris paribus, which lowers the CRRA coefficient we require to explain the observed investment patterns.

However, the non-normal age-dependent model still misses the fact that earnings risk and persistence varies substantially over the income distribution. It is only after including those in the fully-fledged nonlinear model (last row) that we obtain a significant reduction in per-period participation costs in the stock market, which are more than halved from 0.0040 to 0.0018, and a further reduction of the coefficient of relative risk aversion from 7.88 to 6.83. Thus, we conclude that, while only a realistic modelling of the full nonlinear dynamics of earnings can generate our results, allowing for age-dependent non-normal shocks can go a long way in explaining low risky shares conditional on stock market participation.

Process	γ	β	κ^{FC}	κ^{PP}	ψ
Canonical	11.16	0.892	0	0.0052	0.0117
Normal, age-dependent	11.08	0.840	0	0.0073	0.0301
Non-normal, age-dependent	7.88	0.849	0	0.0040	0.0986
Nonlinear	6.83	0.860	0	0.0018	0.1053

Table 4: Parameter estimates under alternative, intermediate earnings processes

4.5 Separating earnings risk and risk preferences

As we showed in Section 4.2, both the nonlinear and canonical earnings processes fit the data, but with different estimated parameters. To further develop the intuition on how richer earnings dynamics affects household portfolio choices, we compare the implied empirical policy functions using simulated data from the model under the nonlinear process and under the canonical process, but keeping preference parameters constant at those estimated from the nonlinear process.

We show the results of this counterfactual experiment in Figure 4. The left and middle panels of the figure show the participation and conditional risky share profiles under the two alternative models. Naturally, the canonical model equipped with parameters from the nonlinear model will result in counterfactual life cycle profiles that do not match the data. In particular, with a lower coefficient of relative risk aversion and lower participation costs into the stock market, the canonical process overestimates both stock market participation shares (practically implying that every household owns some stocks) and the conditional risky share, which is as high as 60% for many age groups. These implications highlight that the canonical process is perceived as less risky by households, who are thus more aggressive with respect to their financial investments. This finding can be further seen in the right panel of Figure 4, where we show the implied EPFs for the risky share conditional on wealth. We find that across different wealth levels, the implied conditional risky share is much higher for the canonical process with nonlinear parameters than under the nonlinear process. Appendix D.3 shows the remaining life-cycle profiles and EPFs. Overall, the results show that the canonical process misses many important features of the earnings dynamics we observe in the data, and thus a model equipped with it has a hard time matching key features of household portfolios, unless it displays a very high coefficient of relative risk aversion.

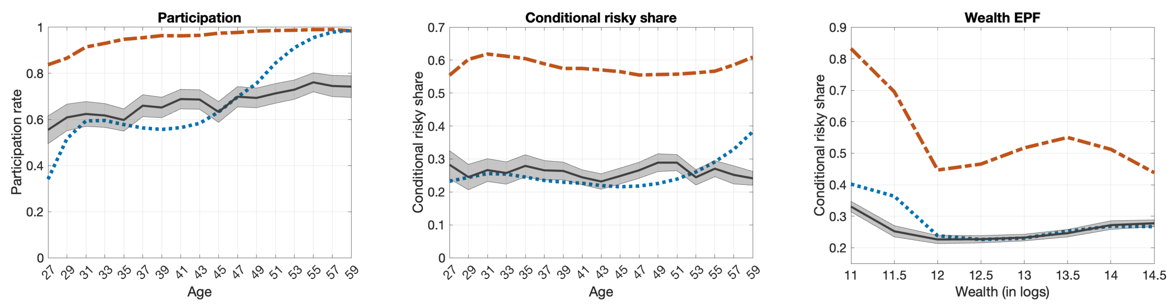

Figure 4: Empirical policy functions with estimated parameters from the nonlinear process. The left and middle panels plot the life-cycle profiles implied by the data (black solid), the estimated model under the nonlinear process (blue dotted), and the simulated model under the canonical process (red dash-dot). The right panel plots the implied relationship between the conditional risky share and wealth, which comes from a regression of the conditional risky share on different wealth bins. 95% point-wise confidence bands are shaded, computed using robust standard errors.

5 Implications

5.1 Investment advice

Life-cycle portfolio choice models are frequently used to provide investment advice or to measure the costs and benefits of different investment strategies. Given that our main results show that a realistic representation of earnings dynamics is key for their estimation and analysis, we now evaluate its effect on optimal investment strategies.

Figure 5 shows the optimal portfolio share of stocks for different income, age, and wealth groups, under the two earnings processes. As a result of our estimation, both processes match average portfolio shares exactly, but they imply remarkably different distributions. As the left hand side panel shows, under the richer earnings process young households with relatively large financial wealth holdings should invest more in the stock market than under the canonical process. This is mostly driven by their lower estimated coefficient of risk aversion, which more than compensates the additional riskiness of labor market income under the nonlinear process.[15] The difference is larger for the lowest earners; given that the nonlinear process recognises that low earners at age 30 still have a lot of upside potential later on in their lives, it recommends a relatively larger share of investment in stocks.

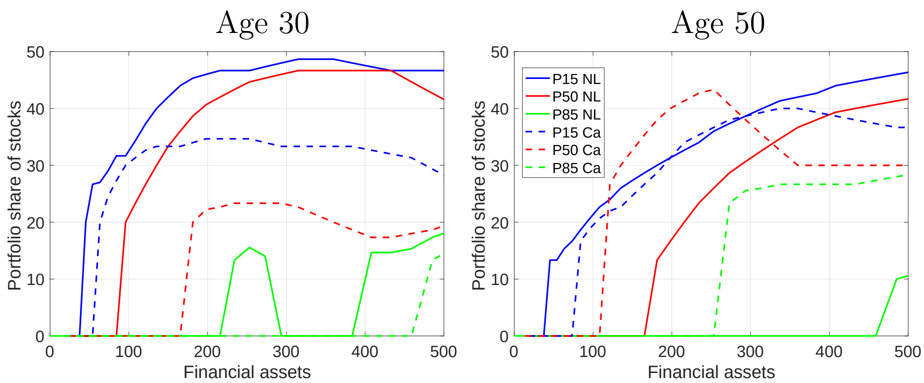

Figure 5: Optimal portfolio share of stocks by level of financial wealth (x-axis), earnings process (straight lines: nonlinear; dashed lines: canonical), and position in the income distribution (percentile 15, blue, median worker, red, percentile 85, green). Policy functions are plotted for existing homeowners who do not choose to change their house ownership status in the following period.

[15]In Appendix D.4 we show the relative contribution of the different parametrization and the different riskiness properties of the two earnings processes in delivering these results.

The picture is different when we look at the right hand side panel, which represents the optimal investment of households at age 50. Here, the additional riskiness of the earnings process, driven mostly by its negative skewness (e.g., unemployment risk), dominates the effect of the lower coefficient of risk aversion: as a result, the model suggests that older workers should invest relatively less in the stock market under the nonlinear earnings process. This effect is particularly strong for workers in the middle to upper part of the earnings distribution (red and green lines).

Because housing is discrete, some of the policy functions are non-monotonic: households that have different medium-term plans in terms of homebuying or home upsizing may take very different portfolio decisions in anticipation of paying for a downpayment and acquiring a mortgage. In Figure 5, we observe this among the young highest earners, who are at their prime homebuying ages.

5.2 Welfare costs of suboptimal investment

We also compute the utility costs under the veil of ignorance of a set of investment strategies, following Cocco et al. (2005), for the two earnings processes and both estimated levels of the coefficient of relative risk aversion γ. These computations compare the utility associated with the consumption streams that households can achieve in the baseline version of our model, in which they can optimally choose their portfolio shares, with three alternative investment strategies that we impose exogenously, namely, full participation into the stock market, no participation at all in the stock market and the common investment advice of investing (100-age)% of wealth into stocks (e.g., Malkiel (1999)). We assume that, in these alternative scenarios, households can optimally adjust their consumption and savings decisions in the light of the exogenously imposed asset strategy.[16]

We begin by comparing our baseline estimated nonlinear earnings process model with the estimated canonical process (first two rows of Table 5). In this comparison, both models differ both in their earnings process and in their preference parameters. We ob-

[16] We consider that some households start out life as homeowners and some start out as renters, but we compute the consumption compensations before they know to which group they will belong.

serve that both processes generate a similarly low welfare cost of following the strategy of not investing into stocks at all. Differences are more apparent in the other two alternative investment strategies. Investing everything into stocks is very costly under the canonical process with a risk aversion over 11 (almost 2% of consumption in every date and state), and significantly less costly (0.92% of consumption) under the nonlinear process with a risk aversion around 7. With respect to following a 100 minus age investment rule, in the nonlinear process the cost is 0.89% of consumption, versus 1.87% in the canonical process. With the nonlinear earnings process, the standard deviation and skewness of earnings shocks increases as households age (as shown in Figure 1), which leads to a lower optimal risky share as households approach retirement, thus giving additional evidence in favour of the simple heuristic rule, even if it was designed without these considerations in mind.

To understand to which extent these differences are driven by the different estimated preference and cost parameters, we also report (last row) results for a version of the model with the canonical earnings process but with the estimated parameters for the nonlinear process, which feature lower risk aversion and lower participation costs to the stock market. Two main messages emerge. First, according to intuition, the costs of not participating at all in the stock market are lower when the coefficient of risk aversion γ and the participation costs are higher. Thus, miss-specifying γ at the level implied by the canonical process also implies underestimating the costs of households not participating in the stock market by about an order of magnitude at a given earnings risk.

Second, at a given parametrization, the nonlinear process implies lower costs of not investing into stocks, which is consistent with its additional riskiness. However, it also implies (slightly) lower costs of investing everything into stocks, which suggests that its additional flexible features (age variation, nonlinearities, etc.) also play a role in determining welfare costs, apart from the average level of risk. Nevertheless, these differences across earnings processes are smaller than those implied by different specifications of the CRRA parameter. In Appendix D.5 we report how these utility costs from suboptimal investment vary over the initial income distribution.

γ	κ^{PP}	Process	No stocks	All into stocks	100 minus age
6.83	0.18%	Nonlinear	0.02	0.92	0.89
11.25	0.52%	Canonical	0.01	1.92	1.87
6.83	0.18%	Canonical	0.15	1.25	1.01

Table 5: Utility costs of alternative investment strategies, measured as consumption compensations in every date and state, under the veil of ignorance, expressed in percentage terms.

5.3 Consumption

Finally, we study the consumption implications of the two earnings processes, with a focus on stockholders vs. non-stockholders. To do so, we simulate data from both canonical and nonlinear models and compute partial insurance coefficients via the Arellano et al. (2017) approach, which we describe in Appendix D.6. The results from our estimation imply that the partial insurance coefficients for the nonlinear earnings process are much closer to the data: expresssed in terms of Blundell et al. (2008) coefficients, our results imply that 38% of persistent earnings shocks in the nonlinear earnings model are effectively insured, as opposed to 30% in the canonical earnings model, and 36% in the data. We also find that stockholders appear to self-insure their consumption better than non-stockholders, which suggests the benefits of diversification. These results are in line with the implied Blundell et al. (2008) coefficients for stockholders and non-stockholders, respectively, which we outline in Appendix D.6.

6 Conclusion

In this paper, we estimate a richer stochastic process for earnings that features a transitory component and a persistent component that allows for age-dependence in moments, nonlinearity, and non-normality. We use it as an input to an estimated life-cycle portfolio choice model with housing that features a one-time fixed entry cost and a per-period participation cost, and compare the implications of the canonical permanent/transitory linear process, with age-independent, normal shocks and the nonlinear earnings process. Our results indicate that the model with the nonlinear earnings process exhibits a lower risk aversion coefficient and lower participation costs than the canonical earnings process.

The model with the nonlinear earnings process also replicates more closely stock market participation by age, consumption insurance, and wealth accumulation patterns.

Our paper complements recent literature that shows that countercyclical skewness is important to understand limited stock market participation (Shen (2024) and Catherine (2022)). A promising direction for future work is to combine both frameworks, using the business-cycle varying earnings process proposed in Paz-Pardo (2024), and study potential complementarities between both approaches.

In our paper, we assume that the earnings process is exogenous and that households supply labor inelastically. In a model with variable labor supply (see, for example, Gomes, Kotlikoff and Viceira (2008)), households would have access to an additional margin of adjustment through the reduction of their portfolio share of human wealth in case its riskiness increases, which would mitigate the portfolio effects on financial wealth.

Finally, our model assumes that all households face similar preferences. However, as Galvez (2017) notes, households potentially have heterogeneous preferences across the wealth distribution and over the life cycle. Estimating the distribution of preferences is an exciting avenue for further research.

References

Alan, Sule (2006), 'Entry costs and stock market participation over the life cycle', *Review of Economic Dynamics* **9**(4), 588–611.

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

Andersen, Steffen, Campbell, John Y, Nielsen, Kasper Meisner and Ramadorai, Tarun (2020), 'Sources of inaction in household finance: Evidence from the danish mortgage market', *American Economic Review* **110**(10), 3184–3230.

Arellano, Manuel (2003), *Panel data econometrics*, Oxford university press.

Arellano, Manuel and Meghir, Costas (1992), 'Female labour supply and on-the-job search: an empirical model estimated using complementary data sets', *The Review of Economic Studies* **59**(3), 537–559.

Arellano, Manuel, Blundell, Richard and Bonhomme, Stéphane (2017), 'Earnings and consumption dynamics: A non-linear panel data framework', *Econometrica* **85**(3), 693–734.

Arellano, Manuel, Bonhomme, Stéphane, De Vera, Micole, Hospido, Laura and Wei, Siqi (2022), 'Income risk inequality: Evidence from spanish administrative records', *Quantitative Economics* **13**(4), 1747–1801.

Athreya, Kartik, Ionescu, Felicia and Neelakantan, Urvi (2023), 'Stock market participation: The role of human capital', *Review of Economic Dynamics* **47**, 1–18.

Bazdresch, Santiago, Kahn, R Jay and Whited, Toni M (2018), 'Estimating and testing dynamic corporate finance models', *The Review of Financial Studies* **31**(1), 322–361.

Bell, Felicitie C., Wade, Alice H. and Goss, Stephen C. (1992), 'Life tables for the United States Social Security Area: 1900-2080', (Social Security Administration, Office of the Actuary).

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

Betermier, Sebastien, Calvet, Laurent E and Sodini, Paolo (2017), 'Who are the value and growth investors?', *The Journal of Finance* **72**(1), 5–46.

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

Blundell, Richard, Pistaferri, Luigi and Preston, Ian (2008), 'Consumption inequality and partial insurance', *The American Economic Review* pp. 1887–1921.

Blundell, Richard, Pistaferri, Luigi and Saporta-Eksten, Itay (2016), 'Consumption inequality and family labor supply', *The American Economic Review* **106**(2), 387–435.

Bonaparte, Yosef, Korniotis, George and Kumar, Alok (2020), 'Income risk, ownership dynamics, and portfolio decisions'.

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

Briggs, Joseph S, Cesarini, David, Lindqvist, Erik and Östling, Robert (2015), Windfall gains and stock market participation, Technical report, National Bureau of Economic Research.

Calvet, Laurent E and Sodini, Paolo (2014), 'Twin picks: Disentangling the determinants of risk-taking in household portfolios', *The Journal of Finance* **69**(2), 867–906.

Calvet, Laurent E, Campbell, John Y and Sodini, Paolo (2007), 'Down or out: Assessing the welfare costs of household investment mistakes', *Journal of Political Economy* **115**(5), 707–747.

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

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

Chang, Yongsung, Hong, Jay H and Karabarbounis, Marios (2018), 'Labor market uncertainty and portfolio choice puzzles', *American Economic Journal: Macroeconomics* **10**(2), 222–62.

Cocco, Joao F (2005), 'Portfolio choice in the presence of housing', *Review of Financial studies* **18**(2), 535–567.

Cocco, Joao F, Gomes, Francisco J and Maenhout, Pascal J (2005), 'Consumption and portfolio choice over the life cycle', *Review of financial Studies* **18**(2), 491–533.

Davis, Steven J and Willen, Paul (2014), 'Occupation-level income shocks and asset returns: Their covariance and implications for portfolio choice', *The Quarterly Journal of Finance* **3**(03n04), 1350011.

De Nardi, Mariacristina, Fella, Giulio and Paz-Pardo, Gonzalo (2020), 'Nonlinear household earnings dynamics, self-insurance, and welfare', *Journal of the European Economic Association* **18**(2), 890–926.

Deaton, Angus and Paxson, Christina (1994), 'Intertemporal choice and inequality', *Journal of political economy* **102**(3), 437–467.

Druedahl, Jeppe and Jørgensen, Thomas Høgholm (2017), 'A general endogenous grid method for multi-dimensional models with non-convexities and constraints', *Journal of Economic Dynamics and Control* **74**, 87–107.

De Nardi, Mariacristina, French, Eric and Jones, John B. (2010), 'Why do the elderly save? The role of medical expenses', *Journal of Political Economy* **118**(1), 39–75.

Edwards, Ryan D (2008), 'Health risk and portfolio choice', *Journal of Business & Economic Statistics* **26**(4), 472–485.

Erickson, Timothy and Whited, Toni M (2002), 'Two-step gmm estimation of the errors-in-variables model using high-order moments', *Econometric Theory* **18**(3), 776–799.

Fagereng, Andreas, Gottlieb, Charles and Guiso, Luigi (2017), 'Asset market participation and portfolio choice over the life-cycle', *The Journal of Finance* **72**(2), 705–750.

Fella, Giulio (2014), 'A generalized endogenous grid method for non-smooth and non-concave problems', *Review of Economic Dynamics* **17**, 329–344.

Galvez, Julio (2017), 'Household portfolio choices under nonlinear income risk: an empirical framework', Working paper.

Gomes, Francisco and Michaelides, Alexander (2003), 'Portfolio choice with internal habit formation: A life-cycle model with uninsurable labor income risk', *Review of Economic Dynamics* **6**(4), 729–766.

Gomes, Francisco, Haliassos, Michael and Ramadorai, Tarun (2021), 'Household finance', *Journal of Economic Literature* **59**(3), 919–1000.

Gomes, Francisco J, Kotlikoff, Laurence J and Viceira, Luis M (2008), 'Optimal life-cycle investing with flexible labor supply: A welfare analysis of life-cycle funds', *American Economic Review* **98**(2), 297–303.

Guiso, Luigi and Sodini, Paolo (2013), Household finance: An emerging field, *in* 'Handbook of the Economics of Finance', Vol. 2, Elsevier, pp. 1397–1532.

Guiso, Luigi, Jappelli, Tullio and Terlizzese, Daniele (1996), 'Income risk, borrowing constraints, and portfolio choice', *The American Economic Review* pp. 158–172.

Guiso, Luigi, Sapienza, Paola and Zingales, Luigi (2008), 'Trusting the stock market', *the Journal of Finance* **63**(6), 2557–2600.

Guvenen, Fatih, Karahan, Fatih, Ozkan, Serdar and Song, Jae (2016), What do data on millions of U.S. workers reveal about life-cycle earnings risk? Working paper, University of Minnesota.

Guvenen, Fatih, Karahan, Fatih, Ozkan, Serdar and Song, Jae (2021), 'What do data on millions of us workers reveal about lifecycle earnings dynamics?', *Econometrica* **89**(5), 2303–2339.

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

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

Huckfeldt, Christopher (2022), 'Understanding the scarring effect of recessions', *American Economic Review* **112**(4), 1273–1310.

Huggett, Mark and Kaplan, Greg (2016), 'How large is the stock component of human capital?', *Review of Economic Dynamics* **22**, 21–51.

Karahan, Faith and Ozkan, Serdar (2013), 'On the persistence of income shocks over the life cycle: Evidence, theory and implications', *Review of Economic Dynamics* **16**(3), 452–476.

Lise, Jeremy (2013), 'On-the-job search and precautionary savings', *Review of economic studies* **80**(3), 1086–1113.

Malkiel, Burton Gordon (1999), *A random walk down Wall Street: including a life-cycle guide to personal investing*, WW Norton & Company.

Meeuwis, Maarten (2020), 'Wealth fluctuations and risk preferences: Evidence from us investor portfolios', *Working Paper* .

Nicodano, Giovanna, Bagliano, Fabio-Cesare and Fugazza, Carolina (2021), Life-cycle risk-taking with personal disaster risk, Technical report, University of Torino.

Nikolov, Boris and Whited, Toni M (2014), 'Agency conflicts and cash: Estimates from a dynamic model', *The Journal of Finance* **69**(5), 1883–1921.

Paz-Pardo, Gonzalo (2024), 'Homeownership and portfolio choice over the generations', *American Economic Journal: Macroeconomics* **16**(1), 207–237.

Piazzesi, Monika and Schneider, Martin (2016), 'Housing and macroeconomics', *Handbook of Macroeconomics* **2**, 1547–1640.

Piazzesi, Monika, Schneider, Martin and Tuzel, Selale (2007), 'Housing, consumption and asset pricing', *Journal of Financial Economics* **83**(3), 531–569.

Ridder, Geert and Moffitt, Robert (2007), 'The econometrics of data combination', *Handbook of econometrics* **6**, 5469–5547.

Rosen, Harvey S and Wu, Stephen (2004), 'Portfolio choice and health status', *Journal of Financial Economics* **72**(3), 457–484.

Shen, Jialu (2024), 'Countercyclical risks, consumption, and portfolio choice: Theory and evidence', *Management Science* **70**(5), 2862–2881.

Storesletten, Kjetil, Telmer, Christopher I. and Yaron, Amir (2004), ‘Consumption and risk sharing over the life cycle’, *Journal of Monetary Economics* **51**(3), 609–633.

Viceira, Luis M (2001), ‘Optimal portfolio choice for long-horizon investors with nontradable labor income’, *The Journal of Finance* **56**(2), 433–470.

Vissing-Jorgensen, Annette (2002), Towards an explanation of household portfolio choice heterogeneity: Nonfinancial income and participation cost structures, Technical report.

Wachter, Jessica A and Yogo, Motohiro (2010), ‘Why do household portfolio shares rise in wealth?’, *Review of Financial Studies* **23**(11), 3929–3965.

Yao, Rui and Zhang, Harold H (2005), ‘Optimal consumption and portfolio choices with risky housing and borrowing constraints’, *Review of Financial studies* **18**(1), 197–239.

Yogo, Motohiro (2016), ‘Portfolio choice in retirement: Health risk and the demand for annuities, housing, and risky assets’, *Journal of Monetary Economics* **80**, 17–34.

Internet Appendix for the paper "Richer earnings dynamics, consumption and portfolio choice over the life cycle"

A Data, descriptive statistics and moments for the structural estimation

We use a combination of the PSID and the SCF for the estimation of the earnings process and the calculation of the auxiliary statistics for the structural estimation.

A.1 PSID

The PSID follows a large number of US households and their potential spin-offs since 1968. While the survey was originally designed to track income and poverty, the PSID has since evolved into tracking household consumption and wealth in more recent waves. When it originally started, the PSID was composed of two main samples: the Survey Research Center (SRC) sample, which was designed to be representative of the US population, and the Survey of Economic Opportunity (SEO), which oversamples the poor.

For the purposes of this study, we focus on the biennial waves that started in 1999. This is because starting from this wave, the PSID has continuous information on household earnings, assets, and consumption.

To construct the statistics that we use for estimation, we follow the sample selection criterion in Yao and Zhang (2005) and Blundell, Pistaferri and Saporta-Eksten (2016). In particular, we consider households with heads aged 25 to 60 years old and who have continuously participated in the labor force. We also impose that households should have a minimum amount of $1,000 of net total assets, following Catherine (2022). Finally, we impose that households should have complete information on demographic characteristics. This leaves us with 30,678 household-year observations. We exclude individuals who are part of the SEO to obtain a representative sample.

A.1.1 Variable definitions

The main variables that we use for the calculation of auxiliary statistics and the earnings process are income, wealth, and the risky share.

The definition of income that we use follows De Nardi et al. (2020). In particular, we use *disposable household earnings*, which are defined as the sum of household labor income and transfers, such as welfare payments, net of taxes and Social Security contributions paid. The reason for this is due to our focus on understanding how households choose between different assets to insure their consumption against income risk. As the PSID provides us information on *pre-tax* earnings, we construct our measure of disposable income by using the tax function in Blundell et al. (2016).

Wealth is defined as net worth, which is the sum of total financial wealth, housing wealth, car values, net of debt. Total financial wealth is the sum of households' holdings in stocks, bonds, and cash, plus any amount invested in retirement accounts. To construct the share of stocks invested in retirement accounts, we follow the rule of thumb in Vissing-Jorgensen (2002). That is, if a household reports that it allocates most of its IRA/401(k)'s in stock, we record this as a 100% allocation, while if it reports that it allocates only some of it in stocks and some in bonds, we record this as a 50% allocation. Finally, the risky share is defined as the share of stocks over household net worth.

A.2 SCF

The SCF is a repeated cross-sectional survey that studies the wealth of US households. It is triennial in nature. The main advantage of the SCF as opposed to the PSID is that it is more detailed with respect to information on wealth. A disadvantage of using the SCF is that as it is a cross-sectional survey, we wouldn't be able to follow households over time; moreover, the SCF does not have information on consumption.

In order to calculate the statistics that we use for the structural estimation, we use similar criteria as in the construction of the PSID dataset. We obtain information from the 1989-2019 sample, to have a comprehensive picture of portfolio allocations for a representative set of US households. To construct the dataset, we follow the same criteria as in the PSID, which gives us around 19,952 households.

We define wealth using the SCF net worth variable (*networth*), which includes financial assets and real estate, net of debts. The income variable that we use is the sum of wage income (*wageinc*), pension income (*ssret*), and transfers (*transfothinc*). The risky share is defined as the proportion of total household equity divided by wealth, for households with positive wealth. The *equity* variable includes direct holdings in stocks, mutual funds and retirement accounts. We define the conditional risky share as the share of risky assets in total net worth, over all households that participated in the stock market.

A.3 Descriptive statistics

Table A1 reports some statistics coming from the two datasets. The units of analysis in our estimation are households. SCF households are, on average, wealthier than the PSID households. They are also more likely to participate in the stock market, as evidenced by the high participation rates; one key reason for this is that the SCF captures better indirect participation in the stock market, such as through mutual funds and retirement accounts. However, for the rest of the variables in the summary statistics, the SCF and PSID households are similar in terms of characteristics.

	SCF		PSID	
	Mean	Standard Deviation	Mean	Standard Deviation
Age	42.974	9.667	44.689	9.539
Wealth	641027	3736128	290007	741823
Household income	115855	212024	57065	66847
Stock market participation	0.677	0.470	0.401	0.490
Equity share	0.173	0.230	0.110	0.200
Housing share	1.454	4.482	1.959	3.891
Conditional housing share	1.827	4.986	2.461	3.985
Conditional risky share	0.259	0.239	0.273	0.234

Table A1: Summary statistics, SCF and PSID. This table reports summary statistics from the SCF and the PSID surveys. The statistics were calculated for households from 27-60 years old. The conditional housing and risky share was computed for households that have at least $1,000 of net worth.

A.4 Moments for the structural estimation

We use a combination of the SCF and the PSID for the construction of the moments that we calculate for the structural estimation.

Our first set of moments are related to the evolution of wealth over the life cycle. We compute these moments using the SCF, as it provides a comprehensive picture of household portfolios. In particular, we take from the SCF the conditional risky share, the average net worth-to-income ratio[17], the homeownership rate, and the stock market participation rate.

The second set of moments are related to stock market participation decisions over the life cycle. In this context, we estimate a the parameters of an OLS regression that examines the determinants of stock market participation, following Alan (2006), Bonaparte et al. (2020) and Briggs et al. (2015). As we would like to capture persistence in stock market participation, we estimate this regression using the PSID data. We consider a model wherein stock market participation is a function of the state variables of the economic model:

$$Part_{it} = f(x_{it}, y_{it}, age_{it}, I_t^h, Part_{it-1}),$$

wherein the relevant variables include income (y_{it}), wealth (x_{it}), the household's age (age_{it}), the household's homeownership status (I_{it}^h), and past participation ($Part_{it-1}$). In practice, the regression that we consider is

$$Part_{it} = b_0 + b_1 x_{it} + b_2 y_{it} + b_3 age_{it} + b_4 age_{it}^2 + b_5 I_t^h + b_6 Part_{it-1} + \mathbf{Z}'\gamma + \nu_{it}, \qquad (20)$$

wherein $\mathbf{Z}_{it}$ includes other variables that might affect participation, such as education, and year dummies that control for aggregate effects.[18] The results of this estimation, whose estimated parameters we also report in Table 2, are in Table A2. The results that we obtain from this regression are very much in line with previous results in the literature on portfolio choice.

[17]To compute this moment, we take the ratio of mean net worth to mean household income in the SCF.

[18]We also consider a regression that uses the Deaton and Paxson (1994) approach to control for age, time, and cohort effects. The results that we obtain are quite similar.

Dependent variable: stock market participation	Coefficient (Std. Err.)
Past participation	0.448***
	(0.011)
Age	-0.012***
	(0.004)
Age squared	0.0001***
	(0.000)
Log of household assets	0.064***
	(0.003)
Log of household income	0.033***
	(0.004)
Homeownership	-0.074***
	(0.010)
Constant	-0.581***
	(0.086)
Observations	18,002
R-squared	0.431

Table A2: OLS regression, participation in the stock market. This table presents results of the estimation of equation:

$$Part_{it} = b_0 + b_1 x_{it} + b_2 y_{it} + b_3 age_{it} + b_4 age_{it}^2 + b_5 I_t^h + b_6 Part_{it-1} + \mathbf{Z}'\gamma + \nu_{it},$$

the empirical policy function for stock market participation, or equation (20). Robust standard errors in parentheses. Statistical significance: ***- $p < 0.01$, **- $p < 0.05$, * - $p < 0.10$. Data: 1999-2017 PSID panel.

B Earnings processes

B.1 Estimation of the nonlinear earnings process

As discussed in the main text, the nonlinear earnings process models the persistent component as the following general first-order Markov model:

$$\eta_{it} = Q_t(\eta_{it-1}, u_{it}), \quad (u_{it}|\eta_{it-1}, \eta_{it-2}, \ldots) \sim U[0,1], \quad t = 2, \ldots, T. \tag{21}$$

where $Q_t(\eta_{it-1}, \tau)$ is the τ-th conditional quantile function of η_{it} given η_{it-1} for a given τ.

One way to understand the role of nonlinearity is in terms of a generalized notion of persistence

$$\rho(\eta_{it-1}, \tau) = \frac{\partial Q_t(\eta_{it-1}, u_{it})}{\partial \eta} \tag{22}$$

which measures the persistence of η_{it-1} when it gets hit by a current shock u_{it} with rank

τ. This quantity depends on the past persistent component η_{it-1} and the shock percentile τ. Note that while the shocks u_{it} are i.i.d. by construction, they may differ with respect to the persistence associated with them. Moreover, persistence is allowed to depend on the size and the direction of the shock u_{it}. As such, the persistence of η_{it-1} is dependent on the size and sign of current and future shocks $u_{it}, u_{it+1}, \ldots$ In particular, the nonlinear process allows current shocks to wipe out the memory of past shocks. By contrast, in the canonical process, $\rho(\eta_{it-1}, \tau) = \rho$, independent of the realization of the past persistent component η_{it-1} or the shock u_{it}. Hence, the notion of persistence in this context is that of the *persistence of earnings histories*. Because the conditional distribution of η_{it} given η_{it-1} is left unrestricted, the nonlinear process allows for conditional dispersion, skewness and kurtosis in η_{it}.[19]

Following Arellano et al. (2017), we specify the quantile functions for the persistent and transitory components as lower-order Hermite polynomials:

$$Q_t(\eta_{it-1}, \tau) = \sum_{k=1}^{K} a_k^{\eta}(\tau) f_k(\eta_{it-1}, age_{it}) \tag{23}$$

$$Q_t(\eta_{i1}, \tau) = \sum_{k=1}^{K} a_k^{\eta 1}(\tau) \tilde{f}_k(age_{i1}) \tag{24}$$

$$Q_t(\varepsilon_{it}, \tau) = \sum_{k=1}^{K} a_k^{\varepsilon}(\tau) f_k^{\varepsilon}(age_{it}) \tag{25}$$

where $a_k^{\eta}(\tau)$, $a_k^{\eta 1}(\tau)$, and $a_k^{\varepsilon}(\tau)$ are modelled as piece-wise linear splines on a grid $[\tau_1, \tau_2]$, $\ldots$, $[\tau_{L-1}, \tau_L]$, which is contained in the unit interval. f_k, $\tilde{f}_k$, and f_k^{ε}, meanwhile, are the approximating functions. We then extend the specification for the intercept coefficients $a_0^{\eta}(\tau)$, $a_0^{\eta 1}(\tau)$, and $a_0^{\varepsilon}(\tau)$ to be the quantile of the exponential distribution on $(0, \tau_1]$ (with parameter λ_-^Q) and $[\tau_L, 1)$ (with parameter λ_+^Q).

If the stochastic earnings components are observed, we could estimate the parameters of the quantile models via ordinary quantile regression. However, as these are latent variables, we proceed with a simulation-based algorithm. Starting with an initial guess of the parameter coefficients, we iterate sequentially between draws from the posterior

[19] Specifically, a measure of period t uncertainty generated by shocks to the persistent component of productivity η_{it-1} is, for some $\tau \in (1/2, 1)$, $\sigma_t(\eta_{it-1}, \tau) = Q_t(\eta_{it-1}, \tau) - Q_t(\eta_{it-1}, 1-\tau)$. Meanwhile, a measure of skewness is $sk(\eta_{it-1}, \tau) = \frac{Q_t(\eta_{it-1}, \tau) + Q_t(\eta_{it-1}, 1-\tau) - 2Q_t(\eta_{it-1}, \frac{1}{2})}{Q_t(\eta_{it-1}, \tau) - Q_t(\eta_{it-1}, 1-\tau)}$ for some $\tau \in (1/2, 1)$.

distribution of the latent earnings components and quantile regression estimation until convergence of the sequence of parameter estimates. Standard errors are computed via nonparametric bootstrap, with 500 replications.

B.2 Estimation of the canonical earnings process

The standard estimation strategy to estimate the canonical earnings process is to use minimum distance estimation, where the goal is to choose the parameters that minimize the distance between the empirical and theoretical moments[20]. An alternative, which we implement here, is to estimate the parameters via pseudo-maximum likelihood estimation, following Arellano (2003). That is, if $u_i \sim \mathcal{N}(0, \Omega(\theta))$, then the pseudo maximum likelihood estimator of θ solves:

$$\hat{\theta}_{PML} = \arg\min_c \left\{ \log\det(\Omega(c)) + \frac{1}{N}\sum_{i=1}^{N} \hat{u}_i \Omega(c)^{-1} \hat{u}_i \right\}.$$

This is equivalent to:

$$\hat{\theta}_{PML} = \arg\min_c \left\{ \log\det(\Omega(c)) + \mathrm{tr}(\Omega(c)^{-1}\hat{\Omega}) \right\},$$

where **tr** is the trace of the resulting matrix, and $\hat{\Omega} = \sum \hat{u}_i' \hat{u}_i$. We can then use the asymptotic covariance matrix to compute the standard errors.

The assumptions on the earnings process imply the following moments:

$$\eta_{it} = \rho^{t-1}\eta_{i0} + \sum_{j=2}^{t} \rho^{t-j} u_{ij} + \varepsilon_{it} \tag{26}$$

from which

$$var(\eta_{it}) = \rho^{2(t-1)}\sigma_z^2 + \sum_{j=2}^{t} \rho^{2(t-j)}\sigma_u^2 + \sigma_\varepsilon^2 \tag{27}$$

$$cov(\eta_{it}, \eta_{it-1}) = \rho^{2t-1}\sigma_z^2 + \sum_{j=2}^{t} \rho^{1+2(t-j)}\sigma_u^2 \tag{28}$$

follow, allowing us to identify the moments.

The estimation results are in Table B1. The parameters indicate that the persistence is close to 0.90. We also find that the standard deviations of the persistent component, the transitory component and the initial distribution of the persistent component are in line with the results in the literature.

[20] Identification of the canonical earnings process follows standard covariance arguments outlined in Arellano (2003).

Parameter	ρ	σ_z	σ_u	σ_ε
	0.904	0.401	0.210	0.206
	(0.136)	(0.091)	(0.064)	(0.091)

Table B1: Parameters of the linear AR(1) process. Note: We report the parameter estimates of the linear AR(1) process for earnings. Standard errors (in parentheses). Data from the PSID, 1999 to 2017. All measures are biennial.

B.3 Intermediate earnings processes

In the main text of the paper, we consider two intermediate processes to the nonlinear earnings process of Arellano et al. (2017). In particular:

1. A version of the canonical process with age-varying persistence and variance of shocks, as in Karahan and Ozkan (2013); and

2. A version of the canonical process in which shocks are allowed to be non-normal, i.e., negatively skewed and with high kurtosis, but without nonlinearities.

We describe each earnings process in this subsection.

Karahan and Ozkan (2013) earnings process. The KO earnings process decomposes the residuals of log-earnings into three components: a household-specific fixed effect, a persistent component modelled as an AR(1), and a transitory component. The specification of the model is the following:

$$y_{it} = \alpha_i + \eta_{it} + \varepsilon_{it} \tag{29}$$

$$\eta_{it} = \rho_t \eta_{it-1} + u_{it} \tag{30}$$

$$u_{it} \sim N(0, \sigma^2_{u,t}), \varepsilon_{it} \sim N(0, \sigma^2_{\varepsilon,t}) \tag{31}$$

The key innovation of this paper is that the variance of the persistent (η) and transitory (ε) shocks are age-dependent, as well as the persistence of the persistent component (ρ_t). Identification of the parameters of the income process can be obtained via covariance restriction-type arguments, and are outlined in Karahan and Ozkan (2013). We estimate the parameters of this income process via a minimum distance estimator.

Non-normal process. We also specify a non-normal process with the restriction that the dependence between η_{it} and η_{it-1} is linear. This yields the following model specification for the dynamics of the persistent component, for a given quantile τ:

$$\eta_{it} = Q_t(\eta_{it-1}, u_{it}) = b_0(\tau) + b_1(\tau)\phi_1(\eta_{it-1}) + \gamma_1(\tau)age_{it} + \gamma_2(\tau)age_{it}^2 \tag{32}$$

In this specification, we specify $b_0(\tau)$, $b_1(\tau)$, $\gamma_1(\tau)$ and $\gamma_2(\tau)$ and as piecewise-linear splines, and model $\phi_1(\cdot)$ as a first-order Hermite polynomial in η_{it-1}. We keep the same specification for the initial distribution of the persistent component and the transitory component of income, ε_{it}. Identification of the earnings process can be established following similar arguments as in Arellano et al. (2017). We also estimate the parameters of this process via the stochastic EM algorithm.

B.4 Comparing the implications of the nonlinear and canonical earnings processes

In this subsection, we compare and contrast the implications of the nonlinear and canonical earnings processes that we earlier described in the main text. We will discuss the results in terms of (i.) age dependence in the moments, (ii.) nonlinearity and (iii.) non-normality.

Starting from age dependence in the moments, the top row of Figure B1 presents the age profile of the standard deviations of the persistent and transitory components of income. By construction, there is no age variation in the standard deviations of both components under the canonical process. In contrast, we find substantial age variation in the standard deviation of the persistent component, but little or not variation in the transitory component. As in De Nardi et al. (2020), there is somewhat a U-shaped pattern in the standard deviation of the persistent shocks. The bottom left row, meanwhile, presents the age profile of autocorrelation of the two processes. As can be observed, in the canonical process we find that autocorrelation is constant over the life cycle. We also find that autocorrelation is much lower for the nonlinear process, but we find an increase between the ages of 30 to 45. Given these differences, it is not surprising that the nonlinear process is able to capture the convex pattern of the conditional variance of

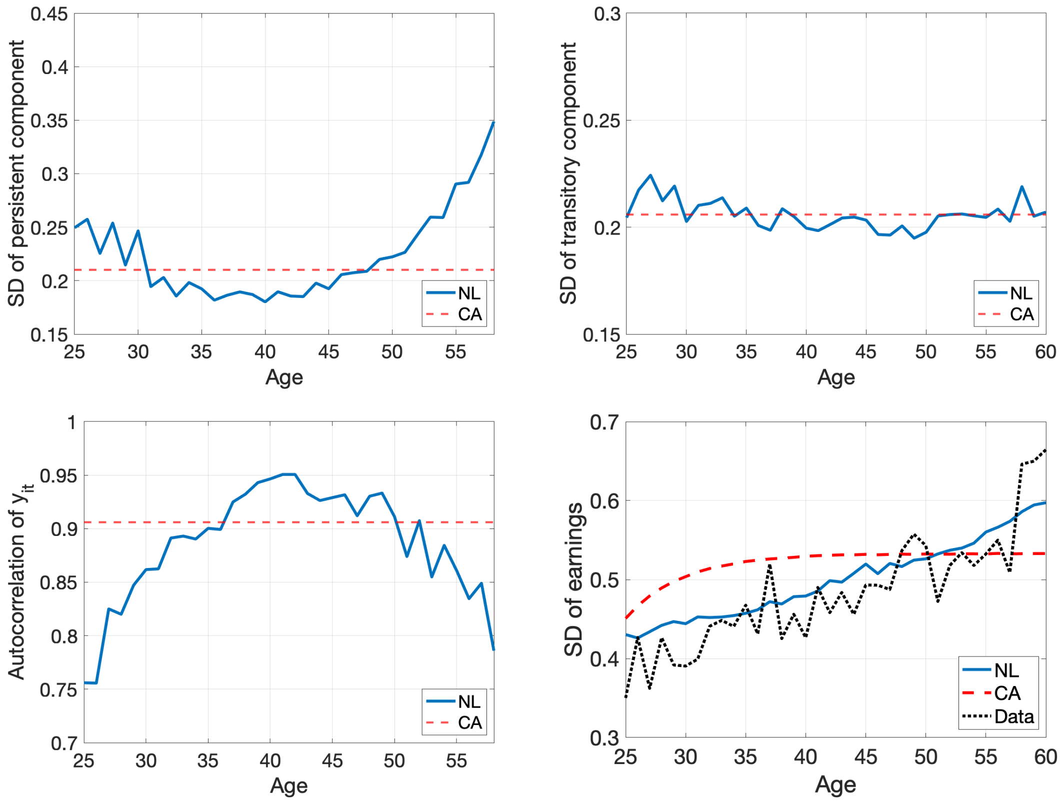

Figure B1: Age dependence of moments, canonical (red) vs. nonlinear (blue) model, PSID. The upper left figure presents the standard deviation of the persistent component of income, graphed by age. The upper right figure presents the standard deviation of the transitory component of income, graphed by age. The lower left figure presents the autocorrelation of the persistent component of income, while the lower right figure presents the cross-sectional variance of income over the life cycle.

earnings over the life cycle in our sample, which the canonical process clearly cannot.

Meanwhile, Figure B2 presents graphs of persistence as a function of the household's position in the income distribution (τ_{init}) and the shock that it receives (τ_{shock}), computed for the average age of a household in the sample (47.5 years). The upper left graph shows the estimates of the average derivative of y_{it} given y_{it-1}, with respect to y_{it-1}. The figure suggests the presence of nonlinear persistence in the data. In contrast, simulated data from the canonical earnings process implies constant persistence, which is in the bottom left panel of the figure. The nonlinear earnings process, meanwhile, is able to reproduce the empirical patterns quite well, which we show in the upper right panel. We also show in the bottom right panel the persistence of the persistent component η_{it}. As we can observe, the estimates are higher than that observed in the data, which is consistent with

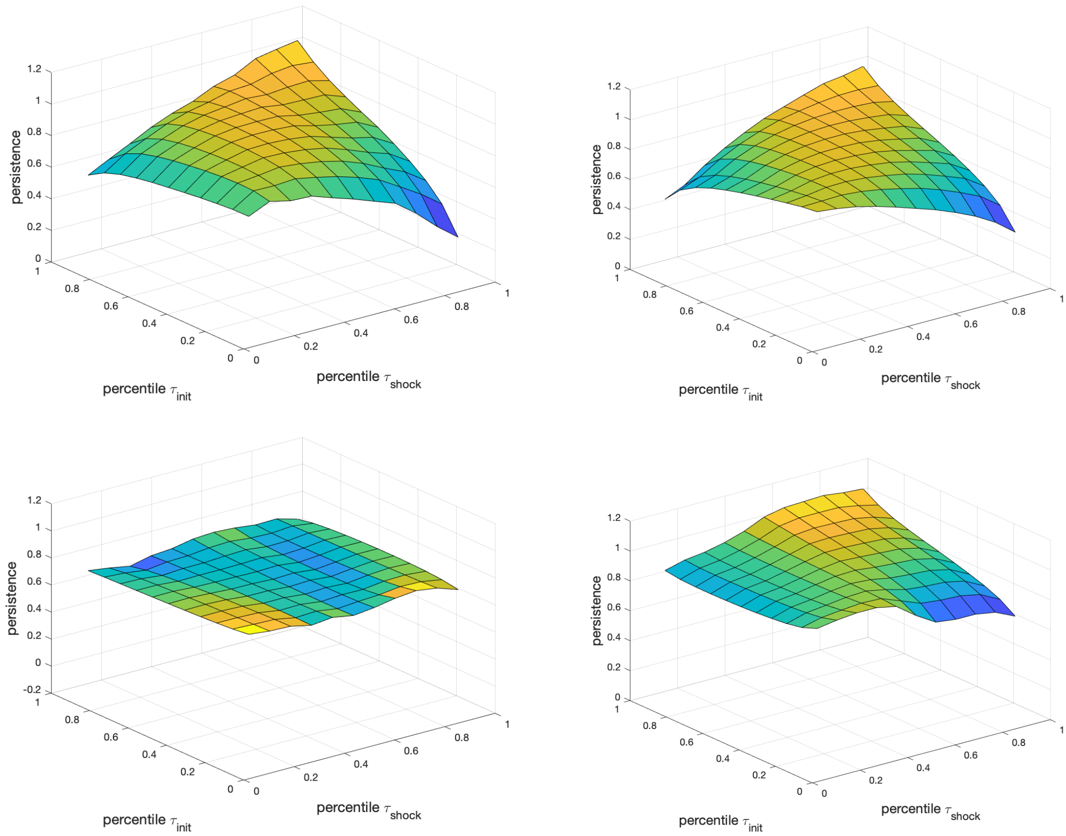

Figure B2: Persistence in the PSID. The upper left panel presents the graph of the average derivative of y_{it} given y_{it-1}, with respect to y_{it-1}, which was estimated from a quantile autoregression of y_{it} on a third-order Hermite polynomial on y_{it-1}. The upper right panel presents the same average derivative, but estimated on simulated data from the canonical earnings model. The bottom left panel presents persistence from simulated data from the canonical model. The bottom right panel presents the persistence from the persistent component of income, η_{it}.

the fact that the figure is net of transitory shocks. The associated standard errors, which are in Figure B3, are small.

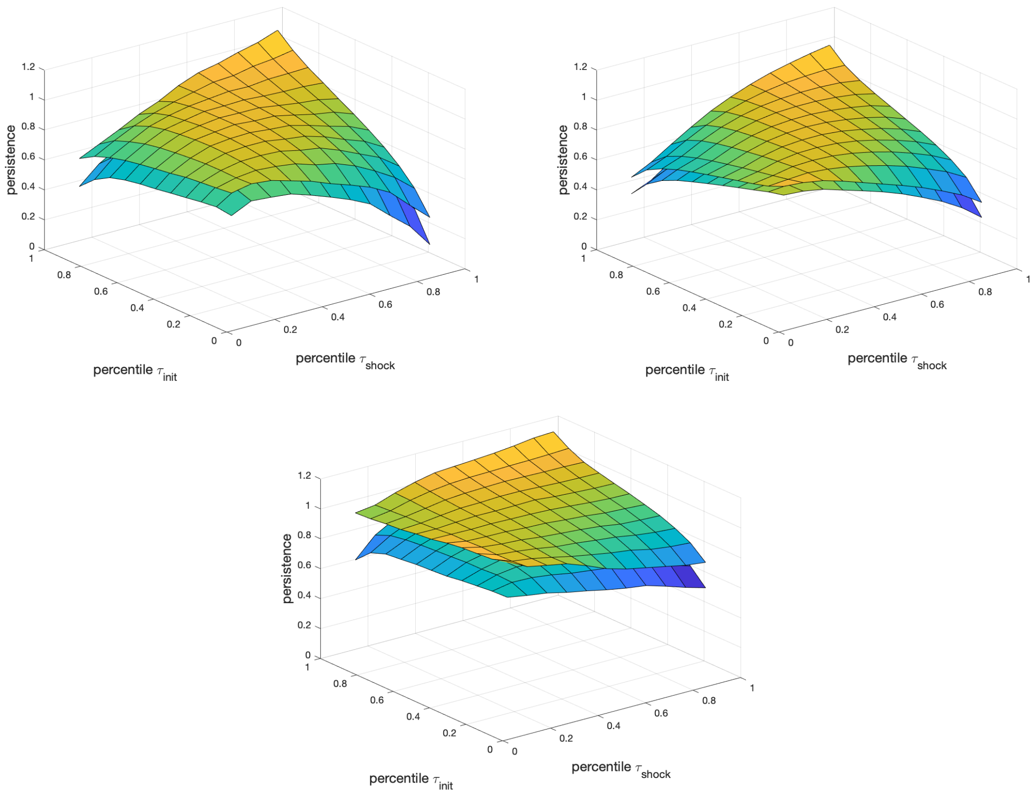

Figure B3: Persistence in the PSID, nonparametric bootstrap. The graphs presented here show the uniform 95% confidence bands calculated from nonparametric bootstraps. The top left panel presents the graph of the average derivative of y_{it} given y_{it-1}, with respect to y_{it-1}, which was estimated from a quantile autoregression of y_{it} on a third-order Hermite polynomial on y_{it-1}. The top right panel presents the average derivative based on simulated data from the nonlinear earnings model. The bottom right graph presents the average derivative of η_{it} given η_{it-1}, with respect to η_{it-1}, based on estimates from the nonlinear earnings model.

Finally, Figure B4 shows the results with respect to conditional skewness. The upper left panel shows conditional skewness as a function of the household's position in the income distribution in the data (blue) and in simulated data (green) from the nonlinear earnings model. As the results indicate, we find some evidence of conditional skewness. Moreover, skewness is positive for households with low y_{it}, and negative for households with high y_{it}. The upper right panel shows the conditional skewness based on simulated data from the canonical earnings model. As the graph indicates, the canonical earnings model predicts symmetric shock distributions. We finally show at the bottom panel the conditional skewness estimates of η_{it}; we find the same patterns, but with a larger magnitude than those for y_{it}. We compute the standard errors and show the results in Figure B5 of Appendix B.4. The results, once again, are precisely estimated.

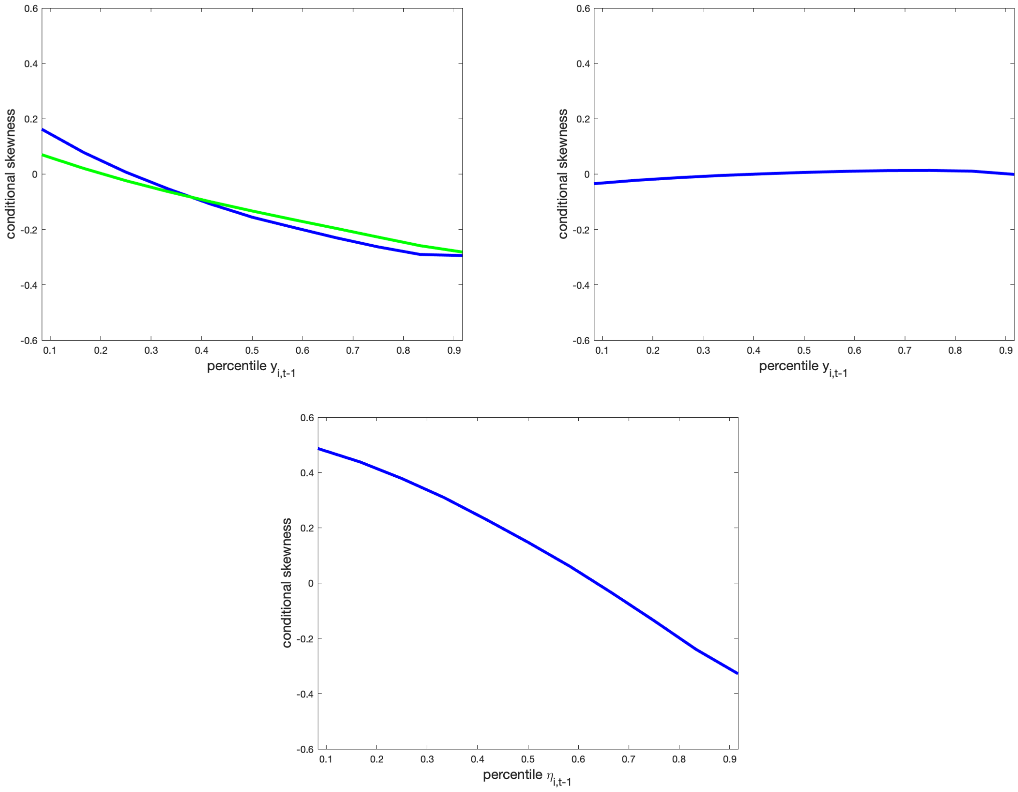

Figure B4: Conditional skewness in the PSID. The left panel presents the graph of the conditional skewness in the data (blue) and the conditional skewness of simulated data from the nonlinear earnings model (green). The right panel presents the conditional skewness based on simulated data from the canonical earnings model.

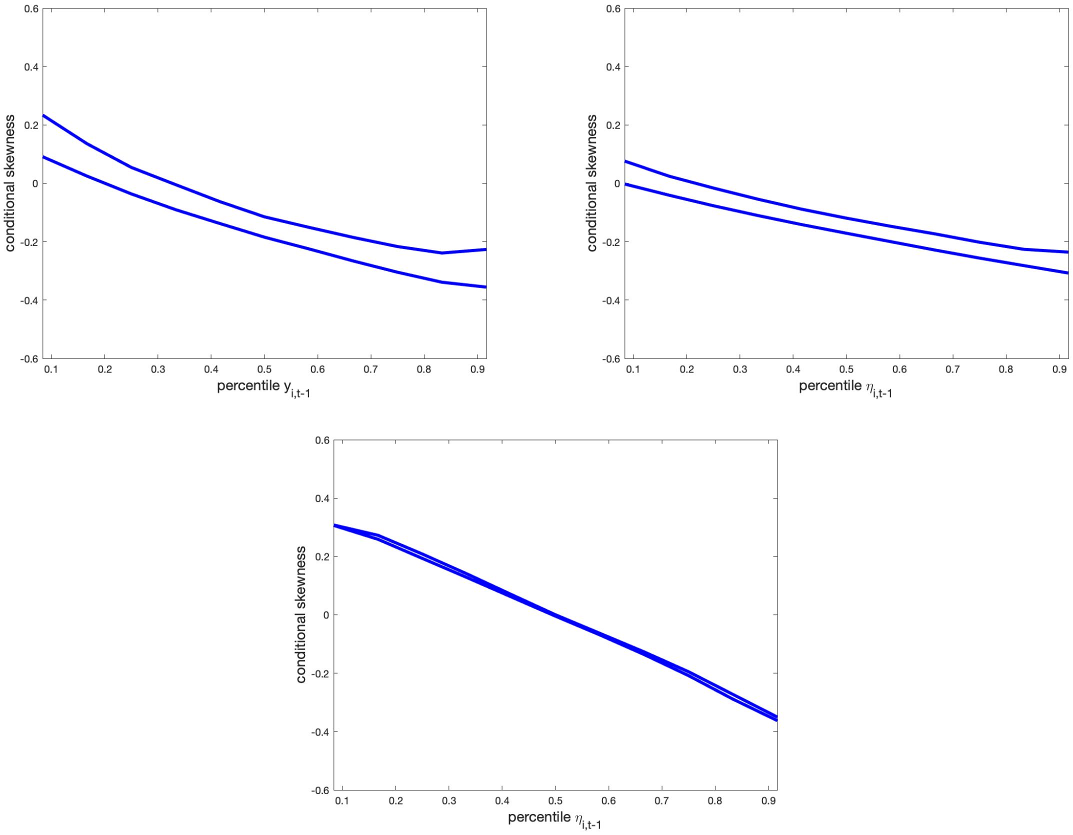

Figure B5: Conditional skewness in the PSID, bootstrap confidence intervals, nonparametric bootstrap. The graphs presented here show the uniform 95% confidence bands. The top left panel presents the graph of the conditional skewness of earnings data y_{it}. The top right panel presents the conditional skewness of earnings simulated from the nonlinear model. The bottom panel presents the conditional skewness of the persistent component η. The graphs were computed via a non-parametric bootstrap with 500 replications.

C Solving the model: computational method

C.1 Solving the households' problem

State variables and choices As described in Section 3, the state variables for the households in our model are age t, wealth z_t, persistent labor market income η_t, the stock market participation status I_t^f, the households' homeownership status I_t^h, the size of the house in which the household lives h_t, and current local house prices p_t^h. Households then choose consumption c_t, the share of risky assets π_{t+1}, the size of the house owned in the following period h_{t+1} and their homeownership status I_t^h, which jointly also determine the housing portfolio share α_{t+1}, wealth in the following period before the realization of shocks x_{t+1}, and their stockholding status.

Discretization We discretize the grid for wealth z_t with 101 points[21] and the persistent labor market income process η_t with 18 gridpoints for the nonlinear process and 8 gridpoints for the canonical process. There are two points in the stock market participation status (participant/non-participant). There are three possible house sizes h_t and two possible homeownership statuses I_t^h (owner or renter). For computational simplicity, we assume that the largest house size is only accessible through ownership, while the small and the medium-sized house can be both owned and rented. We summarise the process for idiosyncratic house price risk p_t^h with 5 gridpoints. We assume that households take portfolio decisions on their risky share π_{t+1} over a grid with 31 points. We discretize the transitory component of earnings u_t with 8 gridpoints for the nonlinear process and 4 gridpoints for the canonical process.

Solution algorithm The solution of the problem proceeds as follows.

1. Given that it is a life cycle model, we can solve it recursively, beginning from the terminal period, in which we assume that there is no utility from continuation and all agents consume their wealth and get utility from doing so. Going

[21] In practice, this grid represents wealth without taking into account the value of the house, which is dealt with separately to exploit its discreteness. An approach in which the value of the house is included as part of the wealth measured in this grid would deliver the same results.

backwards, at any given period t we will have the continuation value function $V_{t+1}(z_{t+1}, \eta_{t+1}, I^f_{t+1}, I^h_{t+1}, h_{t+1}, p^h_{t+1})$.

2. We compute the relevant expectation that enters time t's decision problem, i.e., $\mathbb{E}V_{t+1}(z_{t+1}, \eta_{t+1}, I^f_{t+1}, I^h_{t+1}, h_{t+1}, p^h_{t+1})$ conditional on time t states η_t, I^h_t, p^h_t and time t choices h_{t+1}, π_t, I^h_t. Note that, at this stage, we do not condition on the fourth choice (consumption c_t or alternatively savings z_t) because we use the endogenous gridpoint method to solve for it (see next step (3)). We also do not need to condition on h_t or I^f_t because neither matter for the expectations conditionally on time t's choices. We compute this expectation based on the known processes for labor market income (both persistent and transitory component), the process for house prices and the process for stock market returns, taking into account their correlations as well. This step is the most computationally intensive part of the solution, given the length of the grids and the high number of dimensions.

3. Based on this $\mathbb{E}V_{t+1}$, we use the endogenous gridpoint method (EGM) to compute the optimal consumption choice for each possible combination of the states in time t $z_t, \eta_t, I^f_t, I^h_t, h_t, p^h_t$, and the other three choices $\pi_t, h_{t+1}, I^h_{t+1}$. This step is relatively fast given that, in general, the EGM step does not require nonlinear maximization (although we need to adjust it for potential kinks and non-convexities induced by the several discrete choices in our framework, see e.g. Fella (2014) or Druedahl and Jørgensen (2017)). As a result, we obtain the policy function for the optimal consumption choice conditional on the states and the other choices in t $c_t(z_t, \eta_t, I^f_t, I^h_t, h_t, p^h_t | \pi_t, h_{t+1}, I^h_{t+1})$.

4. After obtaining this policy function, it is straightforward to compute the value function V_t for each state $z_t, \eta_t, I^f_t, I^h_t, h_t, p^h_t$ and choice $\pi_t, h_{t+1}, I^h_{t+1}$. To compute the two remaining household policy functions, we take the maximum over the π, h and I^h grids. As a result, we obtain the policy functions c_t, π_t, h_t and I^h_t as a function of time t states $z_t, \eta_t, I^f_t, I^h_t, h_t, p^h_t$, and also the value function $V_t(z_t, \eta_t, I^f_t, I^h_t, h_t, p^h_t)$. I^f_t follows from the π_t choice, and z_t follows from the consumption choice. All of

these choices also jointly determine the housing share α_t, which we report in the paper but which we do not explicitly need to solve the program.

5. Once we have value and policy functions for all ages, we simulate households forward throughout their lives, starting from an initial income distribution (which corresponds to the empirical income distribution at age 25) and an initial wealth distribution (as described in Section 4, all agents start out with zero financial wealth and half of the agents start out as homeowners of a medium-sized house with 20% equity on it). To do that, we draw random shocks for their transitory income, persistent income, stock returns, and house prices.

C.2 Structural estimation

The targeted moments that we choose for the structural estimation follow the literature that aims to estimate the structure of stock market participation costs (see, e.g., Alan (2006), Alan (2012), and Bonaparte et al. (2020)). We pick 5 parameters (risk aversion γ, discount rate β, and participation costs κ^{FC}, κ^{PP}, and the utility premium from homeownership ψ) to match 11 targets in the data (percentage of people that own stocks directly, mean financial wealth to income, conditional risky share, homeownership rates, and 7 parameters from the OLS regression in Table A2).

To obtain the weighting matrix W, we follow Erickson and Whited (2002) and compute the influence functions of the targeted moments. Influence functions capture how sensitive an estimator is to small perturbations in the underlying data distribution. An advantage of influence functions is that it is a computationally convenient way to estimate the variance-covariance matrix of a set of estimators, which are consistent and asymptotically efficient. Note, though, that to compute our moments, we use a combination of two datasets, wherein one is a repeated cross-section (SCF) while the other is a panel (PSID). In order to take this into account, we follow Arellano and Meghir (1992) and consider that these two datasets are independently sampled from each other. Specifically, let $M_d = (m_1(d_{it}), m_2(d_{it}))$ be the partitioned vector of moments that we match, and let

$\psi_{M(d)} = (\psi_{m_1(d_{it})}, \psi_{m_2(d_{it})})$ be the partitioned vector of influence functions[22]. Because the moments come from two different, and mutually independent samples, we write the covariance matrix of the moments as the following block-diagonal matrix:

$$\Omega = \begin{bmatrix} \Psi_1 & 0 \\ 0 & \Psi_2 \end{bmatrix},$$

wherein

$$\Psi_1 = \frac{1}{N_1} \sum_{i=1}^{N_1} \left(\psi_{m_1(d_{it})}\right) \left(\psi_{m_1(d_{it})}\right)'$$

is the clustered covariance matrix of the influence functions computed from the SCF moments, and N_1 is the number of observations in the SCF, and

$$\Psi_2 = \frac{1}{N_2 T_2} \sum_{i=1}^{N_2} \left(\sum_{t=1}^{T_2} \psi_{m_2(d_{it})} \right) \left(\sum_{t=1}^{T_2} \psi_{m_2(d_{it})} \right)'$$

is the clustered covariance matrix of the influence functions computed from the PSID moments, and N_2 and T_2 are the cross-sectional and time series dimensions of the PSID, respectively. Note that the moments and the influence functions are computed with the sample weights provided by the two datasets.

This then implies that the corresponding weighting matrix is

$$W = \Omega^{-1} = \begin{bmatrix} \Psi_1^{-1} & 0 \\ 0 & \Psi_2^{-1} \end{bmatrix}.$$

Given that the clustered covariance matrices are based on moments calculated from the data and do not depend on the parameters of the structural model, the influence functions only have to be calculated once.

Given our estimate for W, we solve the minimization problem described in Equation 18 by using a grid search algorithm, which we refine with increasingly narrow grids around previous optima. Namely, we solve the households' problem as described in Appendix C.1 for a large number of combinations of the 5 parameters of interest, find the point that minimizes the distance measure in Equation 18, and then solve the problem again for a large number of parameter combinations in the area around this new optimum. We continue until the level of numerical error introduced by discretization in the model is

[22] The dataset index corresponds to the SCF (index 1) and PSID (index 2).

large enough not to allow us to make further significant improvements in the model fit. As we describe in Appendix C.3, the link between our five parameters and the key five targets out of the 11 that we try to fit is quite tight, and we haven't found any (numerical) evidence for significant issues related to multiple local minima.

To compute the standard errors for the parameter estimates, we follow De Nardi, French and Jones (2010) and compute the following variance-covariance matrix (following their notation):

$$V = (1 + \tau)(DW'D)^{-1},$$

where D is the gradient matrix, that is, the responsiveness of our parameter estimates to change in the data moments, and τ is a ratio between the number of simulated households in the model, and the number of households in the data.

C.3 Sensitivity of moments to parameters

In Figure C1 we provide an intuitive measure of how the five estimated parameters are identified by and relate to the main moments we are interested in targeting. Namely, we represent by how much each of five key moments (average wealth to income ratios, share of participants in the stock market, conditional risky share of financial assets, persistence of stockholding status, homeownership rate) changes when we make changes to each of the parameters (CRRA coefficient, discount rate, participation costs and homeownership utility premium) while keeping all else constant. The changes in the moments are represented as absolute deviations from their levels implied by our main nonlinear calibration.

By looking at the top two panels, we observe that the average wealth to income ratio is tightly linked to both the coefficient of relative risk aversion and the discount rate, and increasing when both increase. However, both parameters can be separately identified because the risky share is decreasing in γ, while it does not move very much, or is even increasing, as we change β.

Although the discount factor is positively associated to homeownership (top left panel), the association is much stronger in the case of the homeownership utility premium (bottom panel). Both parameters can be separately identified because β is associated with increases in both the homeownership rate and wealth to income ratios, while ψ increases the homeownership rate while barely moving the wealth to income ratio.

Per-period participation costs govern the level of participation in the stock market, which responds very strongly to their changes (middle left). In contrast, entry costs to the stock market are identified from the OLS coefficient that determines the persistence of stockholding; when entry costs are very high, the stock holding status in the model is very persistent (middle right). This high persistence also explains that, counterintuitively, higher entry costs increase participation ceteris paribus: households that would find it optimal to enter and exit the stock market multiple times choose to remain stockholders if entry costs are high, thus increasing overall stock market participation rates.

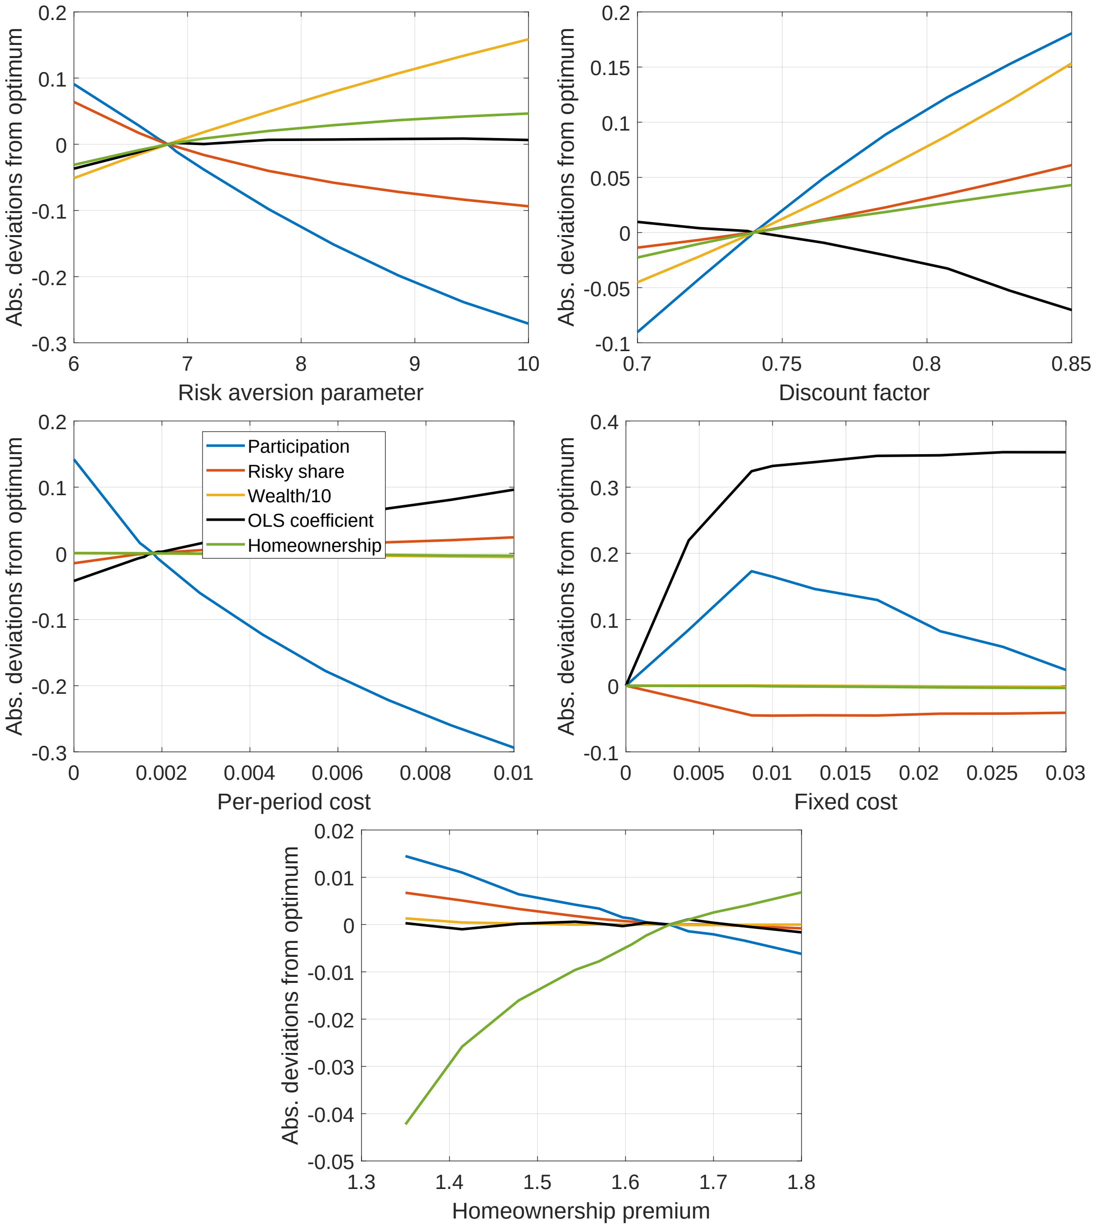

Figure C1: Deviations of the moments with respect to their targeted values, under the nonlinear earnings process, as we ceteris paribus change the coefficient of risk aversion (top left), the discount rate (top right), the fixed costs of participation in the stock market (center left), the per-period cost of participation in the stock market (center right) and the homeownership utility premium (bottom). Deviations are expressed in the same units as the moments, except for those of average wealth, which are divided by 10 for comparability with the others.

D Model implications

D.1 Coefficient of variation

The coefficient of variation (CV) measure proposed by Arellano et al. (2022) quantifies the uncertainty in the predictive distribution of income Y_{it} given covariates X_{it}. The goal is to mimic the household's prediction problem, using all available information from the data. In our case, the predictors that we use are a cubic polynomial in past log disposable income $\ln Y_{it-1}$, interacted with a linear polynomial in age. This is to ensure that we have the same predictors across the simulated data from the two structural models, and the PSID data. Note, however, that in the PSID, we also control for education and aggregate effects by including year dummies.

The CV measure is defined as:

$$CV(X_{it}) = \frac{\mathbb{E}(|Y_{i,t} - \mathbb{E}(Y_{i,t}|X_{it})||X_{it})}{\mathbb{E}(Y_{i,t}|X_{it})} \tag{33}$$

where Y_{it} is household income one period ahead, in levels. The coefficient of variation is a ratio between two measures: the mean absolute deviation, which is a measure of dispersion of the predictive distribution of income, and the mean, which is a measure of location. Estimating these two measures requires two separate prediction tasks, which we do by following Arellano et al. (2022). Specifically, we consider two parametric estimators for the two quantities:

$$\mathbb{E}(Y_{i,t}|X_{it}) = \exp(X_{it}'\beta) \tag{34}$$

and

$$\mathbb{E}(|Y_{i,t} - \mathbb{E}(Y_{i,t}|X_{it})||X_{it}) = \exp(X_{it}'\gamma) \tag{35}$$

wherein we consider exponential specifications due to the fact that income is non-negative. In practice, we estimate β and γ using two Poisson regressions. First, we regress Y_{it} on X_{it}, which delivers $\hat{\beta}$. Second, we regress $|Y_{it} - \exp(X_{it}'\widehat{\beta})|$ on X_{it}, which delivers $\widehat{\gamma}$. Finally, our estimate of the coefficient of variation is simply

$$\widehat{CV}(X_{it}) = \exp[X_{it}'(\widehat{\gamma} - \widehat{\beta})]. \tag{36}$$

We report the resulting distribution of the CV measure, and binscatter plots of the measure across wealth and the net worth-to-income ratio. Figure D1 shows the resulting kernel densities out of the estimation of the CV measure. We find that both the kernel densities of the CV estimated from the data, and from the nonlinear process exhibit similar features, while the canonical process is tightly concentrated around values from 0.23-0.29. This suggests that the nonlinear income process adequately captures the risk found in the data, which the canonical process cannot capture.

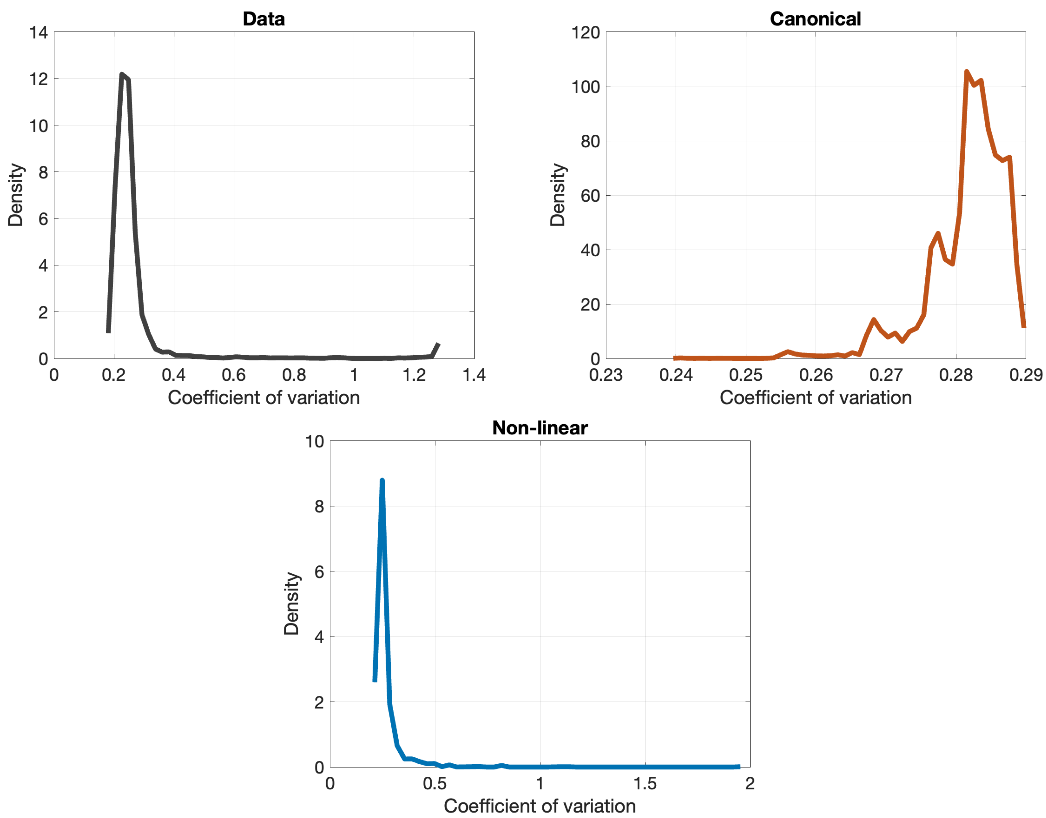

Figure D1: Kernel densities of the coefficient of variation. The kernel densities are estimated using a Gaussian kernel, with the optimal data bandwidth. Top left: PSID data. Top right: Simulated data from the canonical model. Bottom: Simulated data from the nonlinear model.

Figure D2 depicts the average coefficient of variation across wealth and the net worth-to-income ratio, calculated in the PSID data (black), and in the simulated data across the two structural models (red for canonical, blue for nonlinear). The top left panel of Figure D2 shows that the CV measure is increasing over the life cycle, which is something that the nonlinear process captures well, but the canonical process cannot. The top right figure meanwhile, shows that the CV measure has a U-shape pattern that the nonlinear process also captures. The bottom left figure shows that the CV decreases along different levels of wealth, and stays relatively flat at higher levels of wealth. The nonlinear process is able to capture this feature quite well, while overestimating risk at the very high levels of wealth. However, the canonical process implies that, on average, income risk is the same for everyone across the wealth distribution. The bottom right panel of Figure D2, meanwhile, depicts the CV measure along different levels of the net worth-to-income ratio. As can be observed, on average, the CV measure is increasing along the net worth-to-income ratio. This pattern that we observe in the data, is well-captured by the nonlinear process, though the estimates are lower at high net worth-to-income ratio levels. The implied average CV measures computed from the canonical process, however, are flat along the net worth-to-income ratio.

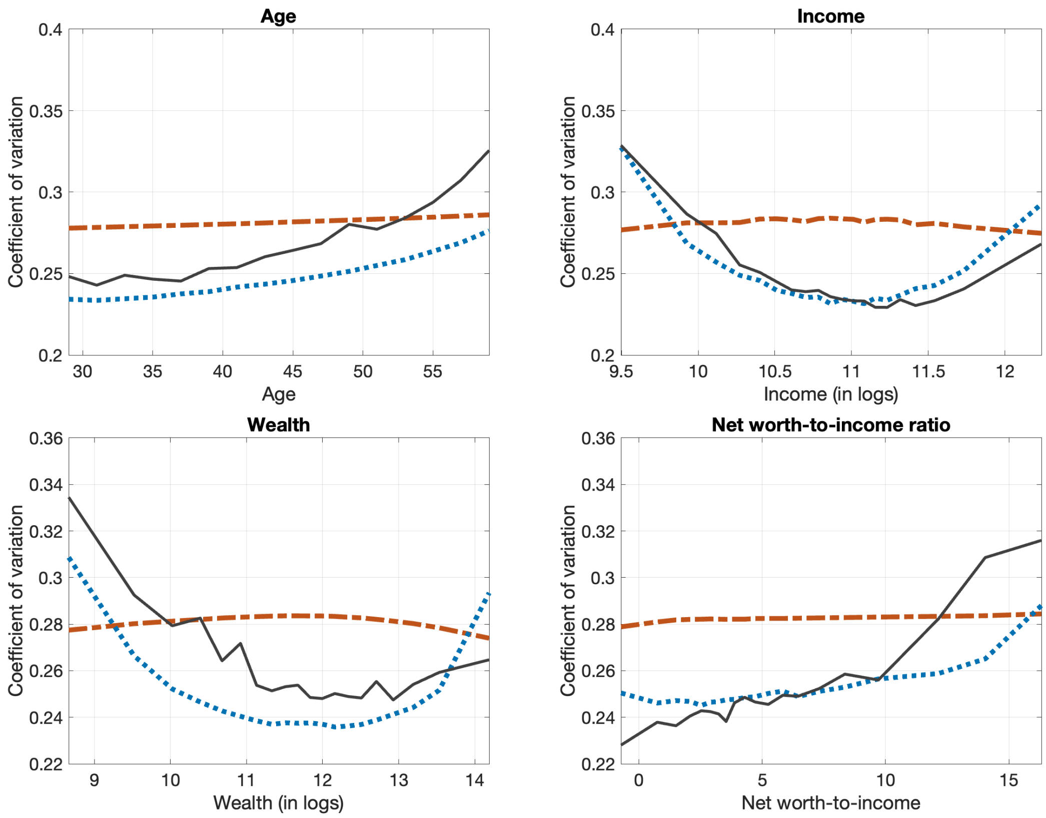

Figure D2: Average coefficient of variation measures, by different bins of age (top left), income (top right) wealth (bottom left), and the net worth-to-income ratio (bottom right). These were calculated both in the PSID data and in the simulated data from the structural models. Data: Black. Nonlinear: blue. Canonical: red.

D.1.1 Empirical policy function with the CV

Table D1 presents the estimation results of an empirical policy function of the risky share, with the coefficient of variation as the main covariate of interest. As the results show, there is a negative relationship between income risk, as measured by the coefficient of variation, and the risky share.

Dependent variable: risky share	Coefficient (Std. Err.)
Coefficient of variation	-0.0866*
	(0.0507)
Log of household assets	0.180***
	(0.0505)
Log of household assets (squared)	-0.00244
	(0.00206)
Log of household income	-0.144**
	(0.0735)
Log of household income (squared)	0.00800**
	(0.00341)
Homeownership dummy	-0.145***
	(0.0164)
Constant	-1.481***
	(0.500)
Observations	15,469

Table D1: Tobit regression, the determinants of the risky share and the role of income risk. This table presents results of the estimation of

$$\text{Risky Share}_{it} = b_0 + b_1 \text{Coefficient of variation}_{it} + \mathbf{Z}'_{it}\gamma + \varepsilon_{it},$$

the empirical policy function of the risky share, equation (19). The dependent variable is the share of risky assets in net worth. Robust standard errors are in parentheses. Controls include demographic variables, cohort and year fixed effects. Statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Data: PSID 1999-2017.

D.2 Conditional risky share over age groups

Figures D3 and D4 plot the conditional risky shares as a function of wealth and net worth-to-income, estimated along different age groups. We distinguish between three different age groups: young households (age less than 35), middle age households (35 to 45) and older households (aged 45 above). To estimate this, we regress the conditional risky share on different pre-specified bins of wealth and the net worth-to-income ratio for each age group, as in section 4.5, and then compute the predicted risky shares.

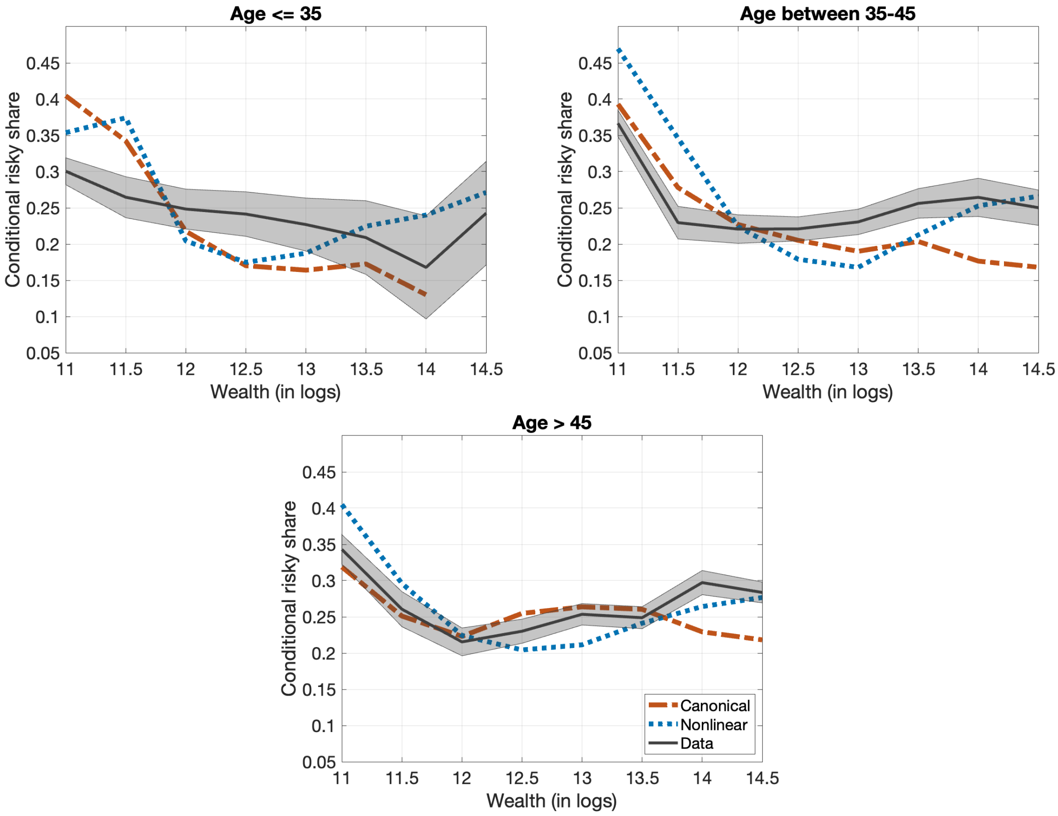

Figure D3: Empirical policy functions of wealth, divided into age groups. The figures show the relationship between the conditional risky share and wealth that are implied by the structural models, in comparison with data from the SCF. The EPFs are the predicted equity shares from a regression of the conditional risky share on bins of wealth. In the data, the estimation also includes year fixed effects. Canonical (red), nonlinear (blue), data (black). 95% point-wise confidence bands are shaded, computed using robust standard errors. Top left: Age less than 35. Top right: 35-45. Bottom: Age above 45.

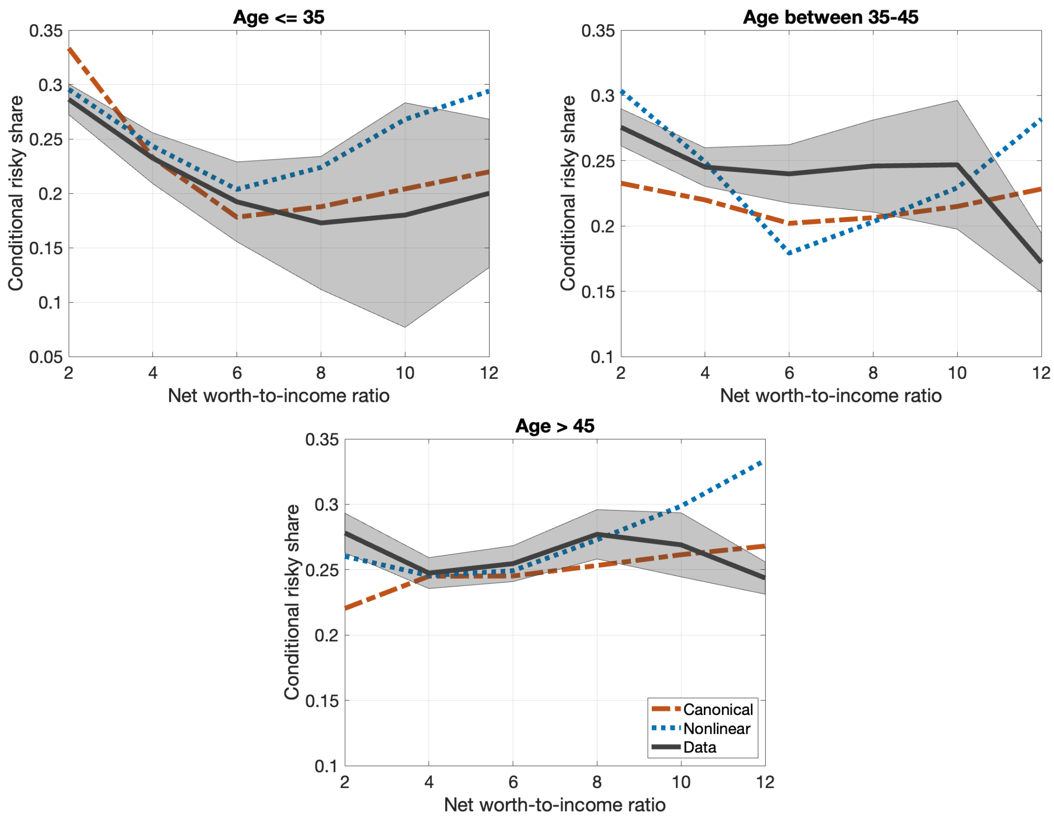

Figure D4: Empirical policy functions of net worth-to-income ratios, divided into age groups. The figures show the relationship between the conditional risky share and net worth-to-income ratio that are implied by the structural models, in comparison with data from the SCF. The EPFs are the predicted equity shares from a regression of the conditional risky share on bins of wealth. In the data, the estimation also includes year fixed effects. Canonical (red), nonlinear (blue), data (black). 95% point-wise confidence bands are shaded, computed using robust standard errors. Top left: Age less than 35. Top right: 35-45. Bottom: Age above 45.

D.3 Separating earnings risk and risk preferences

In this subsection, we report the remaining life-cycle profiles and EPFs from simulated data from the model under the nonlinear earnings process, and the model under the canonical process, but simulated with the estimated parameters from the nonlinear process. The first four figures show the life cycle profiles of wealth, the homeownership rate, the unconditional risky share and the conditional housing share. All of the profiles result in the model under the canonical process being unable to replicate the observed life cycle profile. The model under the canonical process implies a lower wealth accumulation profile (top left), and a lower homeownership rate (top right), in comparison to the data. It also implies higher unconditional risky shares (middle left), and higher conditional housing shares (middle right), though.

The bottom panel of the figure shows the implied EPF for different net worth-to-income ratio bins. As in the main text, under the canonical process but with a lower risk aversion parameter from the nonlinear model, households invest more of their wealth into risky assets.

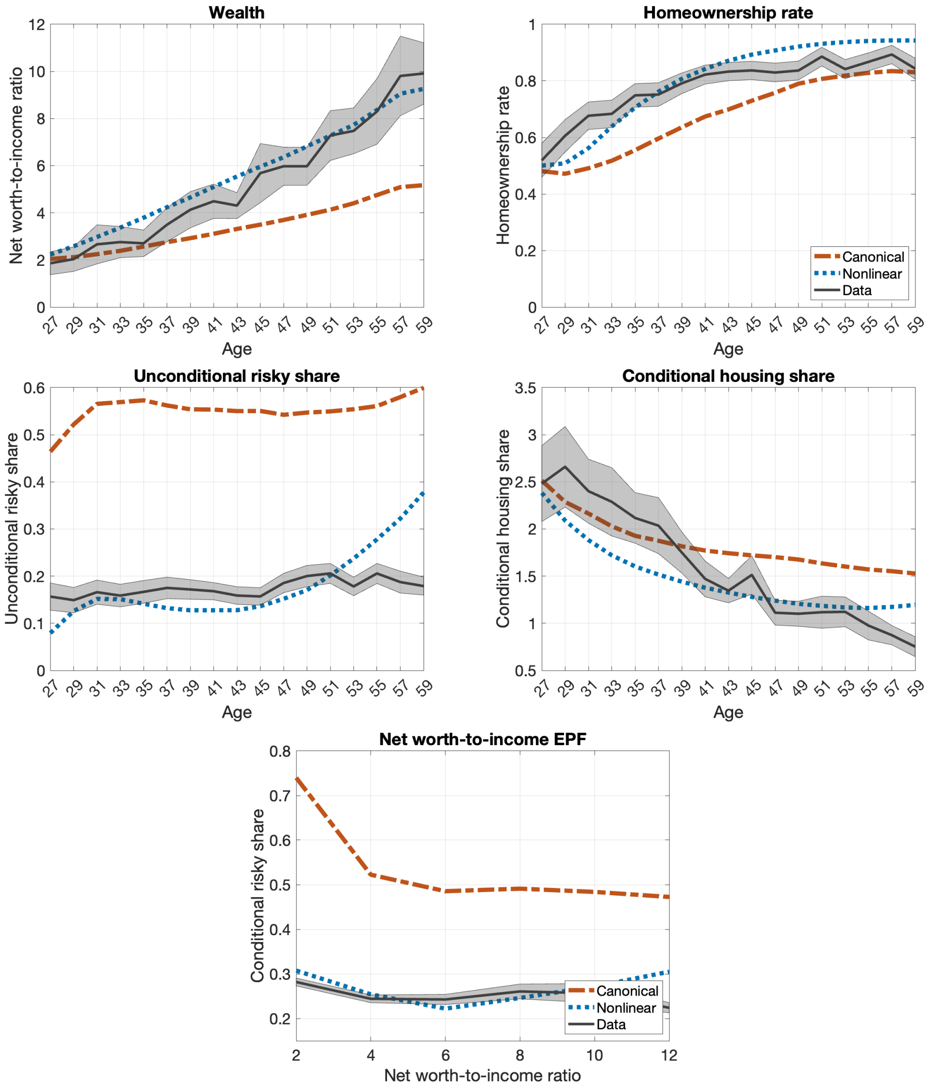

Figure D5: Empirical policy functions implied by the structural models, counterfactual experiment. The graphs plot the life-cycle profiles (first four panels) and the conditional risky share along the net worth to income distribution (bottom panel) implied by the structural models under the nonlinear and the canonical process, with the estimated parameters from the nonlinear model, in comparison with data from the SCF. The empirical life cycle patterns are estimated using OLS regressions, following the Deaton and Paxson (1994) methodology. The implied life cycles of the structural models are estimated via an OLS regression with age dummies. The EPF is the predicted equity share from a regression of the conditional risky share on bins of the net worth-to-earnings ratio and age fixed effects. In the data, the estimation also includes year fixed effects. Canonical (red), nonlinear (blue), data (black). 95% point-wise confidence bands are shaded.

D.4 Optimal investment

Figure D6 shows the relative contribution of preference parameters and earnings processes in explaining the differences in optimal investment profiles between the nonlinear and canonical processes. To do so, it represents the optimal share of stocks in financial wealth under both estimated models (top panels) and under both earnings processes, but keeping constant the estimated parameters at their level for the nonlinear process (bottom panels). Looking at the bottom panels, it is clear that the nonlinear process implies that future discounted human wealth is more stock-like than under the canonical process, which leads to lower optimal shares of stocks for all age, wealth, and income groups. The differences between processes become smaller as financial wealth increases and, as a result, human wealth has a lower weight on the household decision problem.

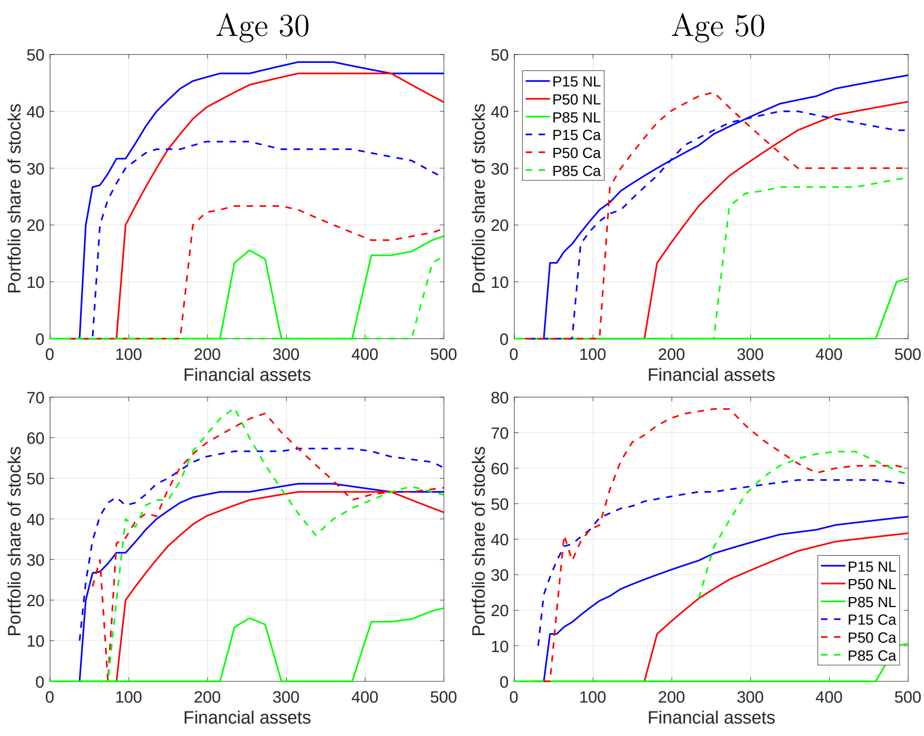

Figure D6: Optimal portfolio share of stocks by level of wealth (x-axis), earnings process (straight lines: nonlinear; dashed lines: canonical), and position in the income distribution (percentile 15, blue, median worker, red, percentile 85, green). Top: estimated parameters for each process; bottom: estimated parameters for the nonlinear process.

However, the lower estimated coefficient of risk aversion under the nonlinear process implies that households will want to invest more heavily into stocks. This effect more than compensates the additional riskiness of the process at younger ages (left hand side panels), but is not enough at older ages (right hand side panels), where optimal investment shares in stocks are still lower under the nonlinear process.

D.5 Heterogeneity in welfare costs of suboptimal investment

Figure D7 reports the welfare costs of suboptimal investment strategies over the initial income distribution. We observe that, for all three strategies, there is substantial heterogeneity, with the initially highest-income people standing to lose more from not investing their larger level of financial wealth optimally. In particular, for the case of never investing into stocks, while the average welfare costs we described in Section 5.2 are relatively low, those of the highest earners reach up to 1.2% of lifetime consumption, both in the nonlinear and the canonical process.

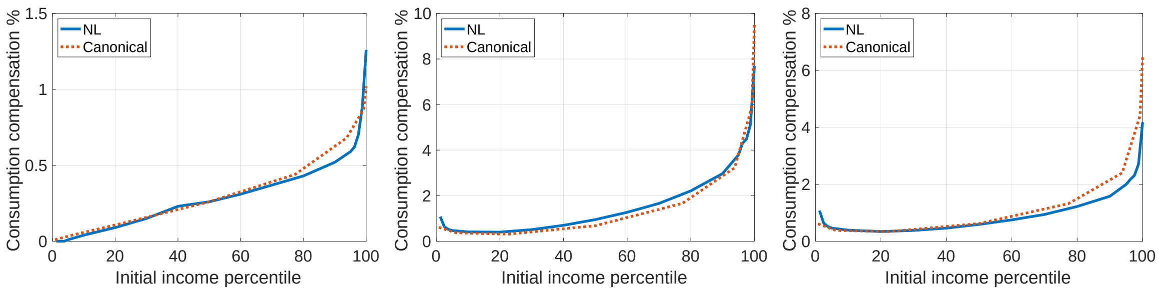

Figure D7: Consumption-equivalent compensations for suboptimal investment strategies, by initial level of earnings and for initial renters. From left to right: no stocks, all stocks, 100-age. $\gamma = 6.83, \kappa^{PP} = 0.0018$

D.6 Consumption pass-through

In this section, we discuss the implications of the nonlinear and canonical earnings processes on consumption insurance in the model with portfolio choice. To do so, we estimate semi-structural empirical consumption rules of the form:

$$c_{it} = f_t(\eta_{it}, \varepsilon_{it}, a_{it}, u_{it}), \tag{37}$$

in which c_{it} is log consumption, η_{it} and ε_{it} are the persistent and transitory components of income, a_{it} is log assets, and u_{it} is an unobserved taste shifter. The model allows us to compute consumption insurance coefficients that are a function of age and position in the asset distribution. To see this, we can write average consumption for a given observation of the earnings components and assets as:

$$\mathbb{E}(c_{it}|\eta_{it} = \eta, \varepsilon_{it} = \varepsilon, a_{it} = a) = \mathbb{E}(f_t(\eta, \varepsilon, a, u_{it})), \tag{38}$$

We can then report the average derivative effect $\phi_t(\eta, \epsilon, a) = \mathbb{E}\left(\frac{\partial f_t(\eta,\varepsilon,a,u_{it})}{\partial \eta}\right)$, and, averaging over the earnings components, $\overline{\phi}_t(a) = \mathbb{E}(\phi_t(\eta_{it}, \varepsilon_{it}, a))$. The quantity

$$\psi^{\eta} = 1 - \overline{\phi}_t(a)$$

can then be understood as the degree of partial insurance to shocks to the persistent component, as a function of age and assets. Similarly, we can define the same quantity for the transitory component.

Following Arellano et al. (2017), we approximate the consumption function with the following specification:

$$c_{it} = \sum_{k=1}^{K} a_k f_k(\eta_{it}, \varepsilon_{it}, a_{it}, age_{it}) + a_0(\tau), \tag{39}$$

where a_k are piecewise polynomial interpolating splines, and f_k's are dictionaries of functions, which are assumed to be Hermite polynomials[23]. We estimate this model on a simulated panel of households from 25 to 60 years old coming from the economy with the nonlinear earnings process, and the economy with the canonical earnings process. As this is a nonlinear regression model, we estimate the parameter estimates via OLS. Given that we can observe the otherwise latent earnings components, we do not have to resort to a simulation-based estimation algorithm.

We report estimates of the average derivative effect $\overline{\phi}_t(a)$, as a function of age and assets, for both economies. The results show that, on average, the estimated parameter $\overline{\phi}_t(a)$ lies between 0.25 to 0.75, close to the Arellano et al. (2017) result. The equivalent parameter estimates for the economy with the canonical earnings process is around 0.45 to 0.95. Both surfaces indicate that the marginal propensity to consume out of persistent income is positive, but decreasing in assets and age, consistent with theory. The implied

[23] In this application, the approximation we use is of the order (2,2,2,2), where each part of the tuple corresponds (persistent,transitory,wealth,age).

Blundell et al. (2008) coefficients, which are in the first two columns of Table D2, show that compared to the benchmark BPP estimate, consumption insurance is higher in the nonlinear economy than in the canonical economy.

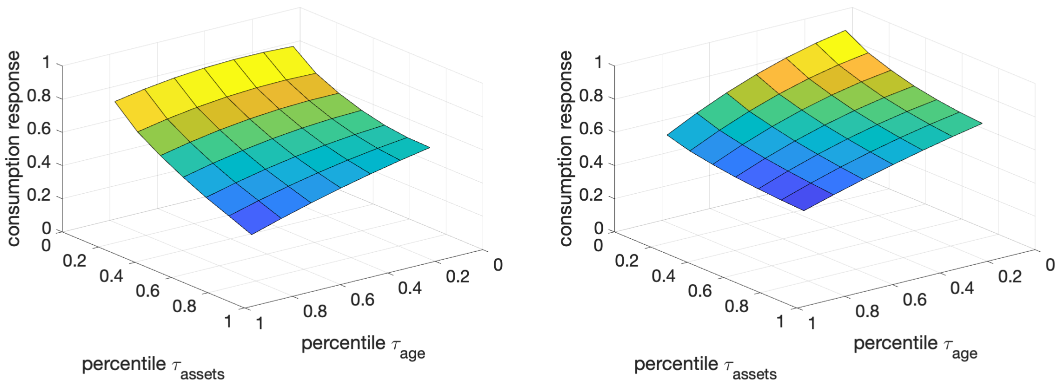

Figure D8: Consumption response to earnings shocks, nonlinear vs. linear model. Note: The graphs presented here show the average derivative effect of η_{it} on c_{it}, computed at percentiles of a_{it} and age_{it}. Data simulated from structural model of life cycle portfolio choice with the nonlinear earnings process (left) and the canonical earnings process (right).

	All		Stockholders		Non-stockholders	
	Persistent	Transitory	Persistent	Transitory	Persistent	Transitory
Canonical	0.3048	0.9100	0.2523	0.9409	0.2237	0.8207
Nonlinear	0.3785	0.8386	0.2995	0.9001	0.1866	0.8400

Table D2: Consumption insurance parameters, implied BPP coefficients.

We finally compute the implied Blundell et al. (2008) coefficients for non-stockholders and stockholders, in the case of the two economies. The results, which are in the second and third columns of Table D2, suggest that stockholders are better able to insure themsleves against income shocks than non-stockholders, both for the nonlinear and canonical economies. Moreover, as in the first column, households in the nonlinear economy are better able to insure themselves against income shocks than those under the canonical economy. These results, however, mask how insurance changes over the life-cycle. To do this comparison, we compute the implied BPP insurance parameters for stockholders and non-stockholders over the life-cycle for both economies. The results of this calculation is in Figure D9. The left panel illustrates that in the canonical economy, BPP insurance coefficients are increasing both for stockholders and non-stockholders. The right panel shows that, in the nonlinear economy, the BPP insurance coefficients are increasing for stockholders, while they are decreasing over time for non-stockholders. One reason for this is that income risk increases over the life-cycle for households in the nonlinear process, while it is constant over time for households in the canonical process, as illustrated in the top left panel of Figure D2.[24] Thus, households in the canonical economy can better insure themselves as they age because the risk they face is constant but the wealth they have accumulated increases progressively, both for stockholders and non-stockholders. However, in the nonlinear economy, income risk increases over the life cycle, making it increasingly difficult to insure against income changes. Stockholders can still better insure against it as they age because they have more sources of insurance (savings from both risky and riskless assets), which results in a BPP insurance coefficient that also increases over the life cycle. Non-stockholders, meanwhile, have comparatively less savings and also less sources of insurance, and are a progressively more selected group. Thus, their BPP insurance coefficients decrease over their lives.

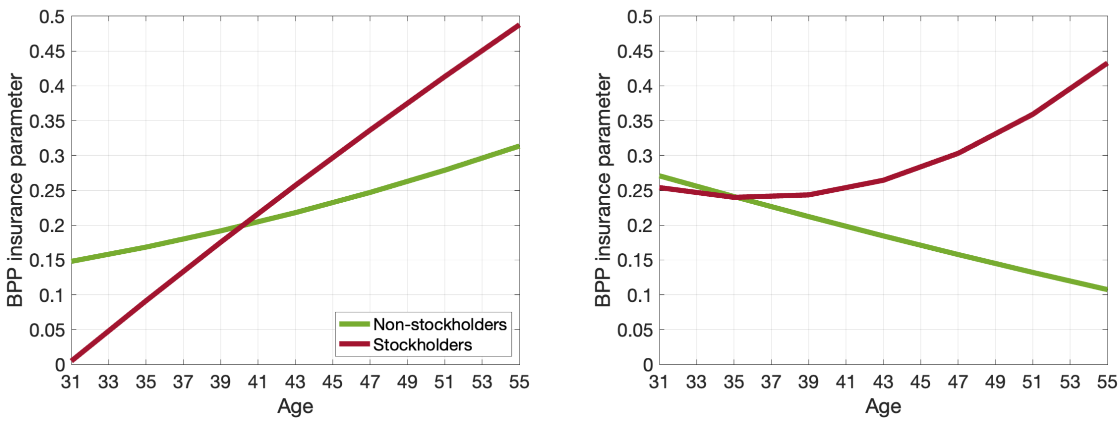

Figure D9: Consumption response to earnings shocks, stockholders vs. non-stockholders. Note: The two panels compare the implied BPP insurance parameters for stockholders and non-stockholders under the canonical (left) and under the nonlinear (right) economy. Data simulated from the structural-model of life-cycle portfolio choice under the two economies.

[24]One can also observe this from the estimated standard deviations in Figure B1.