Working Paper Series

Julio Gálvez, Gonzalo Paz-Pardo

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

Abstract

Households face earnings risk which is non-normal and varies by age and over the income distribution. We show that allowing for these rich features of earnings dynamics, in the context of a structurally estimated life-cycle portfolio choice model, helps to rationalize the limited participation of households in the stock market and their low holdings of risky assets. Because households are subject to more background risk than previously considered, the estimated model implies a substantially lower coefficient of risk aversion. We also find renewed support for rule-of-thumb investment strategies under the model with the nonlinear earnings process.

Non-technical summary

In both the US and the euro area, many households do not invest at all in the stock market and, even if they do, they tend to allocate a relatively small share of their total wealth to it. Although stocks are risky, stock market investors have faced on average much higher rates of return on their wealth over the long run. Understanding the drivers of this limited stock market exposure is key to better measure household risk aversion, but also to design policies that can encourage the accumulation of wealth, retirement savings, or investment over the long run.

This paper studies the role of labor market income risk in explaining these household portfolio decisions. More specifically, it draws from recent literature that has shown that earnings shocks are age-dependent, non-normal and non-linear. For example, younger households face less persistent shocks while older households face less frequent shocks, which are however of a larger magnitude when they occur, with large negative shocks being more likely than large positive shocks (negative skewness). Given that the labor market is a key source of income for most households, these features may have first-order implications for investment decisions. Households confronted with large amounts of asymmetric earnings risk may optimally choose to reduce their risky asset holdings because they consider that they already have substantial risk in their portfolio once their earnings are taken into account.

To quantify this effect, we use a life-cycle model of consumption, savings and portfolio decisions, which we estimate with US data from the Panel Study of Income Dynamics (PSID) and the Survey of Consumer Finances (SCF) between 1999-2017. Our model implies that, once we equip it with a realistic formulation of earnings dynamics, it is much easier to explain the relatively small risky portfolio shares of many households without assuming that they are extremely risk-averse. Furthermore, the model with flexible earnings dynamics can also better explain portfolio shares for households with different levels of income. We also study optimal investment advice and find renewed support for the portfolio rule of investing (100-age)% of one's portfolio into stocks.

We find that our results are consistent across a broad set of alternative modeling formulations, such as including disaster risk in the stock market, housing, correlations between labor market income and stock returns, richer preference specifications, etc.

Our results shed new light on the reasons behind household portfolio decisions and the optimal exposure of households to risky assets. Understanding these mechanisms is particularly

relevant in today's advanced economies, where an increasing share of national wealth is accounted for by private household savings and private or employer-provided pension funds. Furthermore, our model highlights the importance of properly measuring labor market income risk in models of the macroeconomy where household decision making is a key margin of adjustment. As a result, it can serve as a building block for macroeconomic models that study the aggregate equity premium or the aggregate equilibrium long-run real rate.

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, the expected discounted sum of 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 non-linear 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, building on Cocco, Gomes and Maenhout (2005), where households choose between saving in risk-free or risky assets, subject to potential entry and per-period participation costs to the stock market. 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, wealth to income ratios, and the portfolio shares of stocks. 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, and stock market participation.

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 assumed by the canonical process, the coefficient of risk aversion that is required to rationalize household portfolio decisions drops from 9.18, which is in the ballpark of standard models that match limited participation and low risky shares (e.g., Cocco et al. (2005), Fagereng, Gottlieb and Guiso (2017)), to 6.41. 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 process can replicate much more closely the flat pattern of risky asset holdings over the income distribution, whilst the canonical process overestimates its slope.

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 the different underlying risks households face over their working lives. As a result, our flexible model predicts lower demand for risky assets at all ages, but particularly so for the old. 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 (negative skewness). This feature, which is at odds with the canonical model, raises the need for precautionary saving, even if asset returns are uncorrelated with earnings shocks. Third, and finally, the non-linearity in earnings shocks allows us to incorporate the fact that negative skewness is larger for relatively higher earners and for older workers, who optimally choose safer portfolios.

Our more realistic modelling of earnings risk also affects optimal investment advice. For instance, looking at a 50-year old worker with relatively low wealth ($150,000) but high earnings ($140,000), the canonical model recommends a high exposure into stocks, of approximately 80% of the financial portfolio. The richer non-linear process, instead, acknowledges that the worker can still suffer sizeable income shocks and suggests a more conservative strategy of 60% into stocks. We then evaluate the welfare costs of suboptimal investment by considering alternative strategies under the veil of ignorance. We find that, despite the increased labor income risk, a portfolio that is partly invested in stocks is better than zero investment, as the welfare costs under this alternative are large, up to 1.5 percent of lifetime consumption. 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. Finally, we assess the consumption implications of our structural model. 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.

Our study complements contemporaneous work that has emphasised the importance of non-normal features of earnings dynamics over the business cycle to explain limited household risk-taking (Shen (2018), Catherine (2020), Catherine, Sodini and Zhang (2020)). Relative to these papers, our focus on the life-cycle aspects and in cross-sectional heterogeneity of households across the income distribution allow us to draw implications for optimal investment advice and welfare for different kinds of workers (Viceira (2001)), and reproduce the rich interaction between savings motives and earnings dynamics at different ages and points of the income distribution.

Although we take Cocco et al. (2005) as a starting point, we show that our implications are unchanged when we include a set of more realistic features in our structural model. In particular, we introduce housing as an additional asset and allow for the elasticity of intertemporal substitution to be different from the coefficient of relative risk aversion (Epstein and Zin, 1989). We also look at different extensions of the baseline model, such as disaster risk in the stock market, different correlations between stock market returns and earnings shocks, etc. In all of these cases, the nonlinear process better matches household portfolio decisions with lower risk aversion, lower stock market participation costs, or both. We conclude that our mechanism through increased earnings risk is complementary to many of these different approaches that have been studied in the literature as a way to rationalize 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)), and wealth (Calvet and Sodini (2014), Briggs, Cesarini, Lindqvist and Östling (2015)). 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, based 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 relatively simple model structure.

We also contribute to a literature that estimates stock market participation costs. Earlier papers obtain participation cost bounds via minimal assumptions on the structural model in the background (Vissing-Jorgensen (2002), Paiella (2007)). This was followed by a subsequent literature that infers these costs by estimating life-cycle models. Most of these papers consider either a one-time fixed entry cost (see e.g., Alan (2006)) or a per-period participation cost (see, e.g., Khorunzhina (2013), Fagereng et al. (2017)), and infer the cost structure under a canonical earnings process. In contrast, the participation costs in this paper are closer to the one in Vissing-Jorgensen (2002), who proposes modelling both fixed and per-period costs to stock market participation. We show that the estimates of these costs are closely linked to the earnings process considered. More recently, Bonaparte, Korniotis and Kumar (2020) use the new PSID waves to study household stock market entry and exit in a life-cycle portfolio choice model with a similar cost structure, but with canonical earnings dynamics, while Galvez (2017) uncovers participation cost bounds under both earnings dynamics via a semi-structural framework.

Our paper is closest to the recent studies that find a key role for countercyclical income risk in household portfolio decisions: because chances of large negative earnings shocks are larger in recessions, at a time in which stock returns are particularly low, households optimally reduce their equity shares (Shen (2018), Catherine (2020), Catherine et al. (2020)). We build on these contributions 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. 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 (negative) skewness of earnings changes. As Guvenen, Ozkan and Song (2014) show, skewness for the median earner is -1.25 in expansions (-1.75 in recessions) and it varies between -0.5 and -1.5 over the income distribution (between -0.5 and -2 in recessions). We show that these large deviations from normality and linearity are relevant to understand portfolio decisions even in the absence of business cycle fluctuations or correlations of labor market income shocks with stock market returns. Thus, we highlight a different, and complementary, channel through which richer earnings dynamics affect

portfolio decisions.

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 intuition underlying the structural model's estimation in Section 4, and present the estimation results and their robustness to alternative model specifications in Section 5. We analyze the implications for investment advice, the subsequent welfare costs of suboptimal investment, and consumption in Section 6. Finally, Section 7 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, non-linearity 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 present age dependence. Figure 1 shows that both the conditional standard deviation (left) and skewness (right) of household post-tax earnings growth become larger (in the case of skewness, more negative) as people grow older. This also implies that the distribution of earnings changes deviates substantially from the case of normal, age-independent shocks.

---

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

The bottom left row 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}) 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.
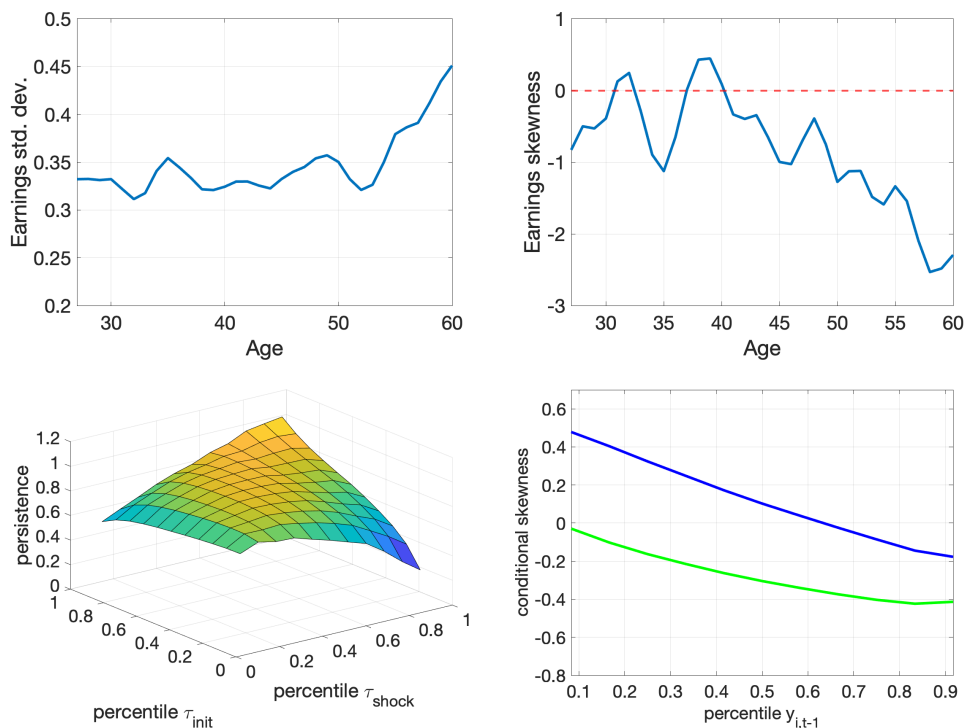

Figure 1: The top rows show the standard deviation (left) and skewness (right) of earnings changes, computed as a function of age. The bottom left figure presents the average derivative of the conditional quantile function of household earnings y_{it} given y_{it-1}, with respect to y_{it-1}, computed from the previous percentile of the household's position in the income distribution (τ_{init}) and the shock (τ_{shock}). The bottom right figure graphs conditional skewness (quantile-based) as a function of the household's position in the income distribution for heads aged 35 (blue) and 55 (green). Data: PSID 1999-2017.

Finally, the bottom right row presents conditional skewness as a function of the household's position in the income distribution. We plot this for two types of households: a household with head aged 35 (blue), and a household with head aged 55 (green). Households in the lower ranks

of the income distribution display in general more positive skewness: there is low chance they fall much further, but they can get a good unusual shock and increase their earnings substantially. Instead, for households in the upper ranks of the income distribution, particularly when they are old, the opposite is true: chances of a large negative shock are much larger than those of a large positive shock.

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 (hereafter CA) 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 CA 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 dis-

tributions, 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 non-linearity, non-normality, and age-dependence (hereafter NL). 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 NL process in Appendix B.1.

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 that receives an adverse occupation-specific shock that 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 NL 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.

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

---

[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.

3 Model

We introduce both the canonical and nonlinear earnings processes into a standard discrete time, life-cycle portfolio choice model 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:

$$\max E_t \left[\sum_{t=0}^{t=T} \beta^t \mathcal{S}_t \frac{c_t^{1-\gamma}}{1-\gamma} \right] \tag{5}$$

where c is nondurable consumption, γ is the coefficient of relative risk aversion, β is the discount factor, and $\mathcal{S}_t$ is the probability of survival at 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 CA and the NL specifications for both components of the earnings process. There is no earnings risk after retirement (age 65), from which households get a public pension.

Budget constraint Households can save in two types of financial assets:

$$c_{t+1} + s_{t+1} + a_{t+1} + \kappa^f(I_{t+1}, I_t) = (1 + r_{t+1}^s)s_t + (1 + r)a_t + y_{t+1} \tag{6}$$

where s_t is the amount of wealth invested in the risky asset and a_t the amount of wealth invested in the risk-free asset at time t, r_{t+1}^s represents the risky return of stocks, while r is the risk-free rate. κ^f represents potential costs of participation in the stock market, which depend on the households' stock market participation status $I_t = (s_t > 0)$. Following Vissing-Jorgensen (2002), these may either be per-period costs, κ^{PP} (just dependent on I_{t+1}), fixed but one-time κ^{FC} (only paid if $I_t = 0$ and $I_{t+1} = 1$, and zero if $I_t = 1$) or a combination of both:

$$\kappa^f(I_{t+1}, I_t) = \begin{cases} 0 & \text{if } I_{t+1} = 0 \\ \kappa^{FC} + \kappa^{PP} & \text{if } I_{t+1} = 1 \text{ and } I_t = 0 \\ \kappa^{PP} & \text{if } I_{t+1} = 1 \text{ and } I_t = 1 \end{cases} \tag{7}$$

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[3] (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).[4]

We define

$$x_t = (1 + r_{t+1}^s)s_t + (1 + r)a_t \tag{8}$$

as the amount of cash-on-hand that an individual has at the beginning of period t.

Finally, as in Cocco et al. (2005), we assume that the household faces borrowing and short-sale constraints:

$$a_t \geq 0, \; s_t \geq 0. \tag{9}$$

The borrowing constraint prevents the household from capitalizing against future labor income or retirement wealth. Meanwhile, the short-sales constraint ensures that the allocation to equities is non-negative.

Households' problem Households thus solve the following problem:

$$V_t(x_t, y_t, I_t) = \max_{c_t, a_{t+1}, s_{t+1}} \left\{ \frac{c_t^{1-\gamma}}{1 - \gamma} + \beta \frac{\mathcal{S}_t}{\mathcal{S}_{t-1}} E_t V_{t+1}(x_{t+1}, y_{t+1}, I_{t+1}) \right\} \tag{10}$$

subject to the budget constraint (6) and the borrowing and short-sale constraints (9), and where the expectation E_t is taken with respect to future realizations of persistent income, transitory income, and stock market returns.

4 Intuition

Before structurally estimating the model, we study the implications of both earnings processes via simulation experiments under the same parameterization. For this purpose, we follow Cocco

---

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

[4]There a third cost of stock market participation in Vissing-Jorgensen (2002), which is a proportional trading cost. As neither the PSID nor the SCF provides information on trading costs, we do not explicitly model this cost.

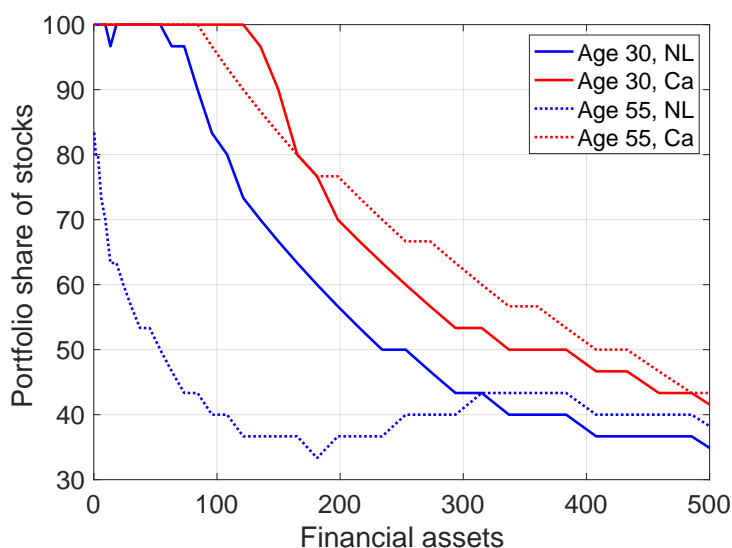et al. (2005) and set $\beta = 0.96$ and $\gamma = 10$.

Figure 2: Portfolio share of stocks, for the median earner, by cash on hand (x-axis), age and income process. Model without participation costs, Cocco et al. (2005) parameters.

Figure 2 shows the optimal risky shares for households of different ages under both income processes, assuming that there are no participation costs. The policy functions are decreasing in financial wealth. The key driver is the importance of human capital (discounted stream of future labor income) relative to accumulated wealth. During working age, since shocks to households' labor income are uncorrelated with stock returns, the *deterministic* component of labor income mimics the payoff of a riskless asset. Hence, for a given level of human capital, households with low financial wealth will tend to invest more aggressively in stocks than wealthier households. Higher financial wealth reduces the relative importance of this "bond-like" human wealth, leading households to rebalance their portfolios by investing less in stocks.

There are relevant differences between the two earnings processes. With the NL process, the *stochastic* component of labor market income is more stock-like, which reduces the optimal exposure of households to the stock market at all ages and levels of wealth. However, because the NL process captures the dependence on age and on the income distribution of earnings dynamics, the strength of this effect varies over the age, income, and wealth distribution. The age-dependence of income shocks and its non-normality in the NL process imply that household income becomes riskier throughout the working life, even at older ages. Thus, households are conservative in their investments over their working life. Meanwhile, in the CA process,

uncertainty with respect to labor income gets resolved much earlier, which implies that older households invest aggressively in stocks.

To further understand the drivers of these differences, Figure 3 compares different points over the income distribution and two additional, intermediate earnings processes that display only age-dependence (KO) and age-dependence and non-normality (KO+NN), which we describe in Appendix B.3. We observe that introducing the age-dependent features of earnings dynamics implies lower portfolio shares across all levels of financial wealth and percentiles of the income distribution, at both young and old ages, when compared with the canonical process. In general, adding non-normalities further decreases the portfolio shares of stocks. However, further adding non-linearities, which include heterogeneous effects across the income distribution, has mixed effects. The reason is that the NL process, unlike the KO+NN process, recognises that income shocks are more negatively skewed for the income-rich. As a result, the NL process recommends riskier portfolios for the poor and less risky portfolios for the rich, when compared with the KO+NN process. In general, all these features emphasize the fact that with a non-linear, non-normal earnings process, future labor income becomes more stock-like, which results in less aggressive portfolio rules.
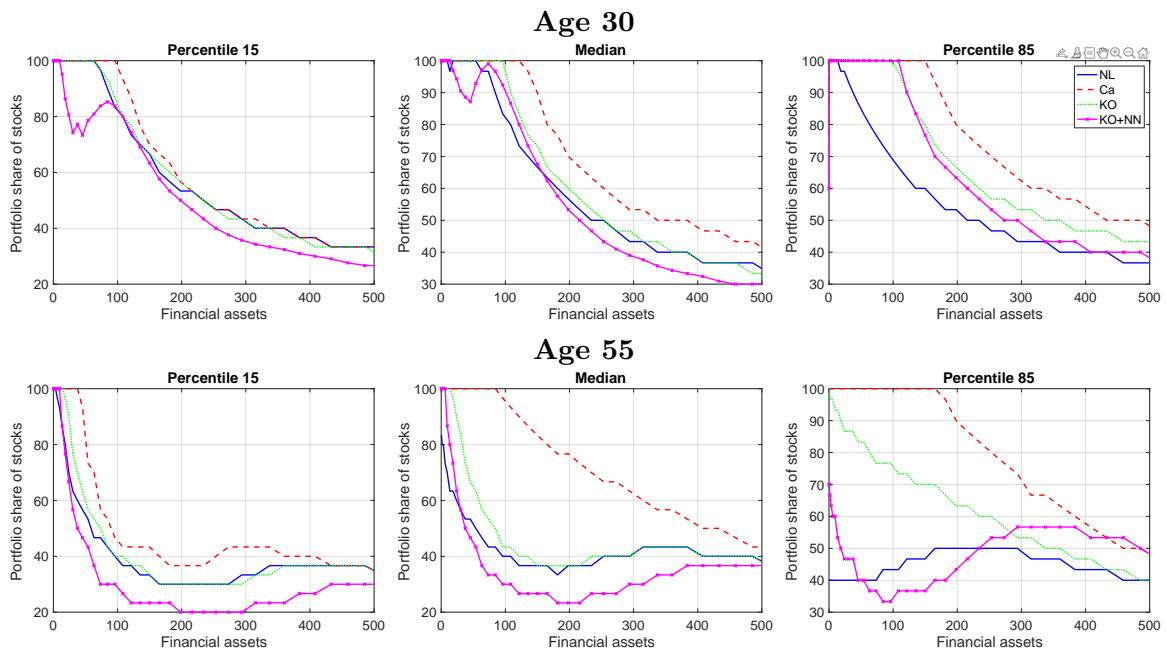

Figure 3: Portfolio share of stocks, by cash on hand (x-axis), position in the income distribution, and income process.
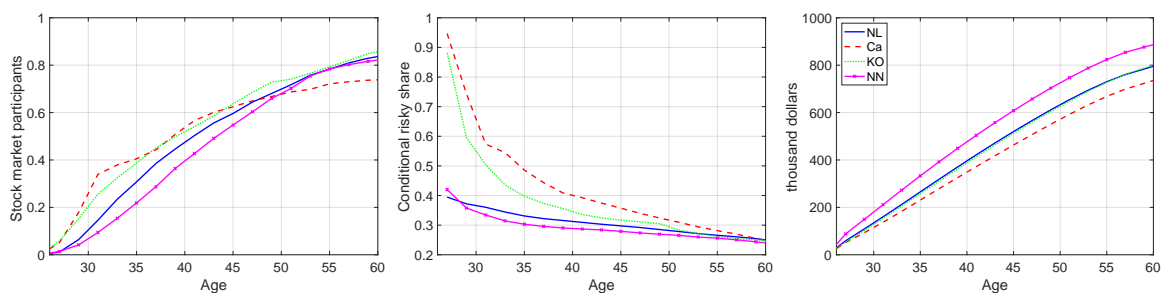
Figure 4: Stock market participation, conditional risky share and average savings, age profile under different earnings processes.

However, in this experiment, participation costs in the stock market are set to zero, which implies that (counterfactually) all households would invest into stocks. If we consider, instead, that there are costs associated with owning stocks, the optimal portfolio choices imply a wealth threshold for stock market participation, as relatively poorer households do not find it worthwhile to pay the cost to invest in stocks. Figure 4 shows stock market participation, the conditional risky share and average savings if we, instead, assume that entry costs to the stock market are 10% of average income and per-period costs are 5%. The most noticeable pattern can be seen in terms of the conditional risky share, which is much flatter in the non-linear and non-normal cases than in the cases with normal shocks, even if we allow for age dependence.

5   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.

5.1   Estimation strategy

5.1.1   External parameters

Public pensions are 70% of the average realization of earnings at retirement age (i.e., 35% of average income of workers in the economy). Meanwhile, 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).[5] We obtain survival probabilities from Bell, Wade and Goss (1992).

---

[5]In the version of the model with disaster risk, there is an ex-ante probability of 2% of stock returns being -48.5%, as in Fagereng et al. (2017).

5.1.2 Estimated parameters and targeted moments

We estimate γ, β, and the stock market participation costs within the model. We use nine data moments for our estimation. The first three moments are related to wealth. We target the percentage of people that own stocks in the PSID (0.53), median financial wealth-to-income ratios (2.0)[6], and the conditional risky share (0.55). Because we are also interested in that our model matches the dynamic and cross-sectional aspects of stock market participation, we estimate an OLS regression of a stock ownership dummy on a polynomial in age, previous stock market participation, income and wealth, and target its coefficients inside the model. Appendix C.1 provides more details about the computation of these moments and the estimation of this regression in the data.

5.1.3 Estimation method

We estimate the model via SMM, which finds the values of the parameters γ, β, and κ's that minimize the following quadratic form:

$$\Pi = \min_{(\gamma,\beta,\kappa^{FC},\kappa^{PP})} (M^s - M^d)'W(M^s - M^d). \tag{11}$$

Here, M^d is the data moments, M^s is the simulated moments from the structural model, and W is a weighting matrix. We follow De Nardi, French and Jones (2010) in the calculation of standard errors, more details of which are in Appendix C.

5.2 Estimated parameters and model fit

Table 1 shows the model fit by comparing our targets in the data (left column) with the model implications under the NL (central column) and those under the CA processes, respectively (right column). In both cases, the model fits its targets remarkably well given how parsimoniously parametrized it is (we estimate 4 parameters to fit 9 targets). In particular, the model closely replicates the limited level of stock market participation that we observe in the data (54% and 53% for the non-linear and canonical process, respectively, compared with an average of 54% in our data) and the conditional risky share of stockholders (57% and 56%, respectively, in contrast with 56% in the data), two crucial moments to understand the savings and portfolio decisions of

---

[6]In Appendix E we show that our conclusions are unchanged if we include housing in the model and target total wealth instead.

| Model | | Baseline | |
Moment	Data	Non-linear	Canonical
Participation	0.535	0.538	0.536
Risky share	0.56	0.569	0.560
Average W/I	2	2.3	2.2
OLS constant	-0.136 (0.145)	0.123	-0.329
OLS, past participation	0.329 (0.009)	0.585	0.559
OLS, age	-0.0115 (0.006)	0.0044	0.0223
OLS, age^2	6.96e-5 (6.16e-5)	-1.16e-04	-2.74e-04
OLS, log income	0.0271 (0.007)	0.114	0.169
OLS, log wealth	0.0729 (0.002)	0.165	0.105

Table 1: Targeted vs. model-implied moments.

| Model | Baseline | |
Parameter	Non-linear	Canonical
γ	6.41 (0.0343)	9.18 (0.0569)
β	0.739 (0.0112)	0.707 (0.0069)
κ^{FC}	0 (0.0005)	0 (0.0010)
κ^{PP}	0.0379 (0.0014)	0.0316 (0.0007)

Table 2: Parameter estimates (SD in parentheses). β is expressed in annual terms. The participation costs are expressed as fractions of average household income (the numeraire in the model)

US households (Alan (2012) and Bonaparte et al. (2020)). Both processes slightly overestimate the average financial wealth to income ratio. Although we only target the average stock market participation, risky share, and financial wealth, both versions of the model do a good job of fitting their profiles over the life-cycle (see Appendix C.3).

However, the model fits the data under remarkably different estimated parameters when we equip it, alternatively, with each of the earnings processes we consider (Table 2). Most notably, the implied CRRA risk aversion parameter is substantially lower (6.41) under the non-linear process than it is under the canonical earnings process (9.18). The intuition for this result follows from what we described in Section 4. Namely, at constant γ, the NL process generates lower conditional risky shares at different points of the income distribution and over the life cycle; in order to reconcile the data with the model, this implies that the calibrated coefficient of relative risk aversion does not need to be as large as in the case of the canonical process. As a result, the NL process can explain much better the relatively flat share of conditional risky shares over the income distribution (Figure 5, left); in the canonical process, instead, the

lowest income quintiles invest too little in stocks because their relative risk aversion coefficient is counterfactually large, given that it needs to match average risky asset holdings in a world in which high earners do not face negative skewness.
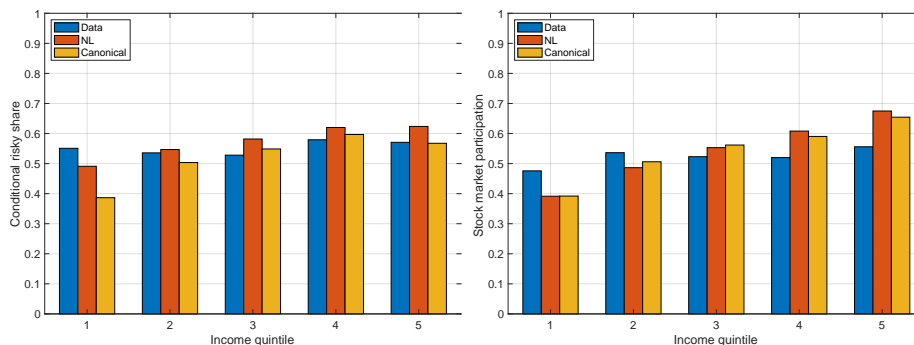

Figure 5: Share of financial assets invested in stocks, conditional on participation (left) and stock market participation (right), over the income distribution. Data: PSID.

For both cases, our estimation generates positive per-period participation costs, but zero entry costs to the stock market. The key moment in the data that governs this parameter is the persistence of the stock market participation status. As Table 1 shows, even in the absence of entry costs to the stock market both processes overestimate the persistence of stock market participation, which can be due to the fact that the estimated per-period costs are reasonably low (Briggs, Cesarini, Lindqvist and Östling (2021)), resulting into households deciding to stay in the stock market. Because both processes imply similar participation costs, the models generate a similar profile of stock market participation over the income distribution (Figure 5, right). In Appendix C.4 we give further information about the relationship between our estimated parameters and the main moments we are interested in replicating.

These implications are remarkably consistent across a wide range of alternative specifications of the model. We now turn to discuss two that are particularly relevant for the question at hand: separating the coefficient of risk aversion from the elasticity of intertemporal substitution, and explicitly considering the role of housing in household portfolios. We then discuss other robustness checks.

5.3 Epstein-Zin preferences

The use of CRRA preferences in our baseline model implies that we are restricting the elasticity of intertemporal substitution (EIS) to be exactly identical to the inverse of the coefficient of risk

aversion. This assumption might imply, for instance, that we are underestimating the actual level of risk aversion under the canonical process because the estimation is preventing γ to rise to avoid making the EIS too low. To disentangle these two effects, and to show that the key differences across processes are indeed driven by risk aversion, we estimate a version of the model with Epstein and Zin (1989), and denote the EIS, now separate from the coefficient of risk aversion, as ψ. Because the discount rate and the EIS are notoriously difficult to separately identify in the absence of e.g. intertemporal fluctuations in stock returns, we perform this experiment by alternatively fixing the discount rate or the EIS, and estimating all other parameters.

Table 3 shows that the differences in coefficients of risk aversion are, if anything, larger across processes when we use Epstein-Zin preferences. Because γ can now rise without implying a very low willingness of households to intertemporally substitute consumption, it is now higher under both processes, but much more (over 10) with the canonical process. The estimated EIS, instead, is remarkably similar across both processes, and close to 0.5.

Case	$\psi = 0.50$		$\beta = 0.84$	
Parameter	NL	CA	NL	CA
γ	7.17 (0.0511)	11.14 (0.0891)	7.22 (0.0955)	10.43 (0.1820)
β	0.806 (0.0117)	0.823 (0.0103)		
κ^{FC}	0 (0.0004)	0 (0.0007)	0 (0.0004)	0 (0.0006)
κ^{PP}	0.0267 (0.0010)	0.0222 (0.0021)	0.0333 (0.0017)	0.0222 (0.0021)
ψ			0.4784 (0.0410)	0.5044 (0.0373)

Table 3: Parameter estimates under the model with Epstein and Zin (1989) preferences. β is expressed in annual terms. The participation costs are expressed as fractions of average household income (the numeraire in the model). SDs in parentheses.

5.4 Portfolio choice under the presence of housing

Houses are often the most significant component of household balance sheets, and the decision to own or rent a home has implications for household wealth accumulation (Paz-Pardo, 2021) and financial portfolio decisions (Cocco, 2005). In this subsection, we show that our main conclusions are unchanged in a model with an endogenous housing choice, which thus explains *total wealth* rather than *financial wealth*.

In the model that we build, households receive utility both from housing and non-housing consumption, with the utility from owner-occupied housing represented by parameter u_h. Households choose to buy a house or to rent, and face transaction costs related to house purchase.

For simplicity, we assume a constant house price. Homeowners may finance their house with a mortgage, for which they need to satisfy a downpayment constraint. Given the new model elements, we now estimate five structural parameters: the previous model parameters, plus the utility gain from housing u_h. We target 11 moments to estimate the five parameters: in addition to the moments that we have discussed earlier, we also target the homeownership rate and we augment the OLS regression by including a homeownership dummy. Additional details about the model, together with model fit and other results are summarized in Appendix E.

Model	With housing	
Parameter	NL	CA
γ	5.78 (0.1180)	7.82 (0.3877)
β	0.707 (0.0439)	0.820 (0.0439)
κ^{FC}	0 (0.0009)	0.01 (0.0058)
κ^{PP}	0.0333 (0.0027)	0.0400 (0.0045)
u^h	1.375 (0.0626)	1.05 (0.0777)

Table 4: Parameter estimates, model with homeownership. β is expressed in annual terms. The participation costs are expressed as fractions of average household income (the numeraire in the model). SDs in parentheses.

Table 4 shows the associated parameter estimates. We find that the results are in line with those of the main model; in particular, the CRRA parameter is much lower under the nonlinear earnings process (5.77) than under the canonical (7.82). In both cases, it is lower than in our baseline economy because of our assumption that housing is risk-free: because households hold a significant amount of wealth in a risk-free asset, they are more willing to hold a higher share of risky assets in their financial portfolio, even at relatively low levels of risk aversion. Although the inclusion of house price risk might imply a higher CRRA coefficient for both processes, we do not expect it to affect the comparison between the nonlinear and the canonical process.

5.5 Decomposing the role of earnings dynamics

As described in Section 2, our flexible NL earnings process differs from the CA process in many ways: age-dependence, non-normality of shocks, non-linearities, etc. To gauge the relative contribution of these factors to explaining our results, in Table 5 we report the estimated parameters under the set of intermediate processes described in Appendix B.3: one with age-dependence, but no non-normalities or non-linearities (KO) and one with age-dependence and non-normality, but no non-linearities (KO+NN).

We find that the process with age-dependent persistence and variance can already generate a much lower estimated coefficient of relative risk aversion (7.13), as the KO process captures the age dependence in the persistence and the variance of earnings over the life cycle. Adding non-normalities lowers the estimated coefficient of relative risk aversion further to 6.61 - as agents internalize that shocks to their earnings are non-normal, they become less willing to invest in a risky asset ceteris paribus. However, the KO+NN process counterfactually assumes that non-normalities and, in particular, negative skewness are the same across the income distribution. Because negative skewness is larger (in absolute value) for high earners, this implies that this process overestimates tail risk at the bottom of the distribution and underestimates it at the top. As we move from the KO+NN to the NL process, this bias is corrected, which would push the coefficient of risk aversion upwards. However, this is more than compensated with other non-linearities, such as those related to persistence, implying a further fall of γ down to 6.41 in the version of the model in which closely replicate the distribution of shocks faced by each household given their age and position in the income distribution.

Process	γ	β	κ^{FC}	κ^{PP}
CA	9.18	0.707	0	0.0316
KO	7.13	0.754	0	0.0442
KO+NN	6.62	0.747	0	0.0442
NL	6.41	0.739	0	0.0379

Table 5: Parameter estimates under alternative, intermediate earnings processes

5.6 Robustness

In this section we show that the main implications of our model, most relevantly in terms of the estimated coefficient of risk aversion under both earnings processes, are unchanged under a set of additional elements that enhances the realism of our baseline model.

Namely, we first study the changes in the parameter estimates when targeting the total amount of wealth in the economy, rather than exclusively financial wealth. Naturally, this implies higher estimated discount rates, but the coefficient of risk aversion is still substantially lower under the non-linear process. Then, we introduce additional elements to the baseline Cocco et al. (2005) model. These include the possibility of a "rare disaster" in the stock market that implies large losses for stockholders, replicating the empirical distribution of wealth at age 25, an

alternative pension system in which old age benefits depend on the last realization of earnings, and correlation between labor market income shocks and stock market returns. For all of these, the canonical process implies larger coefficients of risk aversion, stock market participation costs, or both. Table 6 offers a quick summary of the parameter estimates for each of these cases, and Appendix D contains detailed descriptions of each of the experiments.

Model	NL process				Canonical process			
Baseline	6.41	0.739	0	0.0379	9.18	0.707	0	0.0316
Total wealth	6.31	0.945	0	0.1200	7.85	0.951	0	0.1137
Disaster risk	5.57	0.849	0	0.0405	7.48	0.838	0	0.0338
Initial wealth	6.30	0.806	0	0.0455	8.17	0.817	0	0.0444
Alt. pension	5.89	0.820	0	0.0714	8.67	0.755	0	0.0486
Correlation income-stocks	5.95	0.845	0	0.0457	8.37	0.811	0	0.0384

Table 6: Estimated parameters under variations of the baseline model

6 Implications

6.1 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 6 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 labour market income under the nonlinear process.[7] 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.

---

[7]In Appendix C.5 we show the relative contribution of the different parametrization and the different riskiness properties of the two earnings processes in delivering these results.
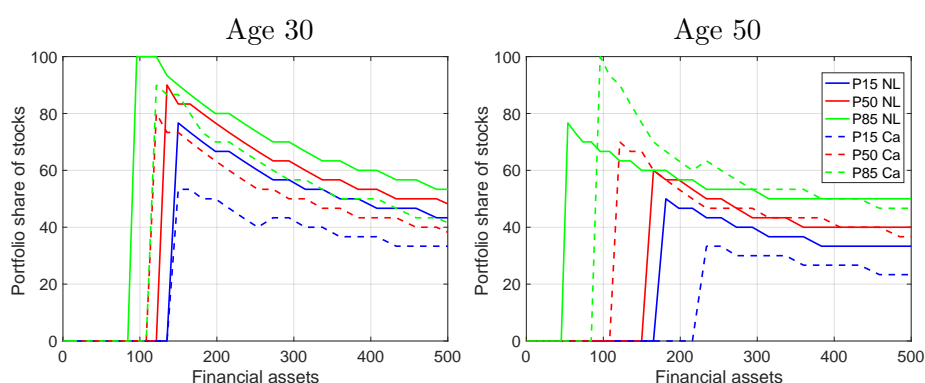
Figure 6: Optimal portfolio share of stocks by level of wealth (x-axis), earnings process (straight lines: NL; dashed lines: canonical), and position in the income distribution (percentile 15, blue, median worker, red, percentile 85, green).

The picture is different when we look at the right hand side panel, which represents the optimal investment of households 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 high earners with relatively low wealth (around $100,000), whose optimal risky shares drop from 100% to 70%.

6.2 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.

For simplicity in the comparison, we follow the Cocco et al. (2005) calibration which, most importantly, implies that we abstract from participation costs in the stock market. Table 7

shows the associated values. Several important messages stand out. First of all, and according to intuition, the costs of not participating at all in the stock market are decreasing in the coefficient of risk aversion γ. 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 between a factor of 2 (canonical earnings process) and a factor of 7 (NL earnings process).

Second, because of the additional riskiness involved in the NL process, in general it implies lower costs of not investing into stocks, and higher costs of investing everything into stocks. These costs, however, are smaller than the differences implied by different specifications of the CRRA parameter. We also show, in the last two rows, that the costs of not investing into stocks would be slightly lower if we allow for positive probability of a very large negative realisation of stock market returns.

Finally, we find that under the NL process the costs involved with following the strategy of investing 100 minus age percent of wealth in the stock market are substantially lower (around half at our preferred coefficient of risk aversion). The reason behind this is related to the age-varying features of earnings displayed in Figure 1: the standard deviation and skewness of earnings shocks increases as households age, which leads to a lower optimal risky share as households approach retirement. Moreover, independently of the portfolio rule that households follow, the heightened risk that households face under the nonlinear earnings process induces them to save for precautionary reasons. This additional reason to save reduces the welfare losses associated with not being able to optimally choose portfolio shares. These facts, which can only be established with a very flexible specification of earnings dynamics, turn out to give additional evidence in favour of that simple heuristic rule, even if it was designed without these considerations in mind.

γ	Process	No stocks	All into stocks	100 minus age
9.17	NL	0.21	0.09	0.02
9.17	Ca	1.43	0.37	0.18
6.4	NL	1.45	0.19	0.21
6.4	Ca	2.33	0.16	0.45
6.4	NL, disaster risk	1.13	0.33	0.13
6.4	Ca, disaster risk	1.83	0.27	0.30

Table 7: Utility costs of alternative investment strategies, measured as consumption compensations in every date and state, under the veil of ignorance, expressed in percentage terms.

However, these relatively low average utility costs from suboptimal investment mask sub-

stantial heterogeneity over the income distribution. In Figure 7 we show that the utility costs of not investing into stocks are as high as 6% of lifetime consumption for the initially highest earners. We also observe that the utility costs of investing everything into stocks are under-estimated under the canonical process for practically the whole income distribution, and that the investment strategy of investing 100 minus age in risky assets is substantially better for all income percentiles under the NL process.
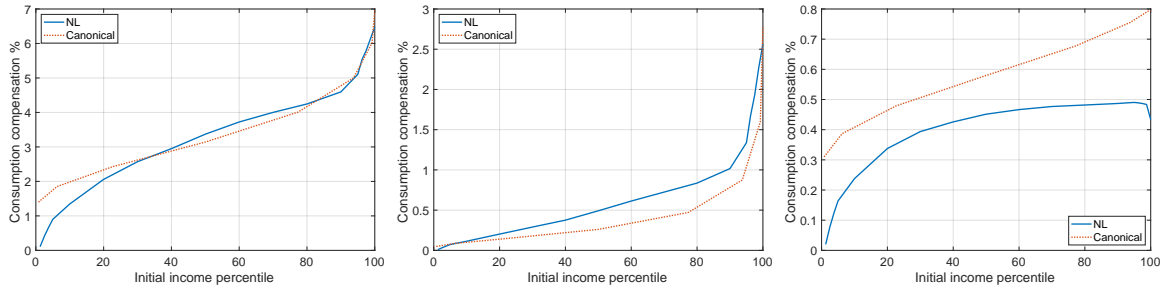
Figure 7: Consumption-equivalent compensations for suboptimal investment strategies, by initial level of earnings. From left to right: no stocks, all stocks, 100-age. $\gamma = 6.4$

6.3 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 non-linear economies and compute partial insurance coefficients via the Arellano et al. (2017) approach, which we describe in Appendix C.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 40% of persistent earnings shocks in the nonlinear earnings model are effectively insured, as opposed to 29% 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 corresponding Arellano et al. (2017) estimates of empirical consumption functions for stockholders and non-stockholders, respectively, which we outline in Appendix C.6.

7 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 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, under a variety of specifications, the model with the nonlinear earnings process exhibits a lower risk aversion coefficient 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. We finally find that the NL process implies lower costs of not investing into stocks, and higher costs of investing all wealth into stocks.

Our paper complements recent literature that showed that countercyclical skewness is important to understand limited stock market participation (Shen (2018) and Catherine (2020)). A promising avenue for future work is to combine both frameworks, using the business-cycle varying earnings process proposed in Paz-Pardo (2021), and study potential complementarities between both approaches.

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

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).

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'.

Briggs, Joseph, Cesarini, David, Lindqvist, Erik and Östling, Robert (2021), 'Windfall gains and stock market participation', *Journal of Financial Economics* **139**(1), 57–83.

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 (2020), 'Countercyclical labor income risk and portfolio choices over the life-cycle', *HEC Paris Research Paper No. FIN-2016-1147* .

Catherine, Sylvain, Sodini, Paolo and Zhang, Yapei (2020), 'Countercyclical income risk and portfolio choices: Evidence from sweden', *Swedish House of Finance Research Paper* (20-20).

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.

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.

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.

Epstein, Larry G and Zin, Stanley E (1989), 'Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework', *Econometrica: Journal of the Econometric Society* pp. 937–969.

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.

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.

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.

Kaplan, Greg and Violante, Giovanni L. (2010), 'How much consumption insurance beyond self-insurance?', *American Economic Journal: Macroeconomics* **2**(4), 53–87.

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.

Khorunzhina, Natalia (2013), 'Structural estimation of stock market participation costs', *Journal of Economic Dynamics and Control* **37**(12), 2928–2942.

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.

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

Paiella, Monica (2007), 'The forgone gains of incomplete portfolios', *The Review of Financial Studies* **20**(5), 1623–1646.

Paz-Pardo, Gonzalo (2021), 'Homeownership and portfolio choice over the generations'. ECB Working Paper no. 2522.

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

Shen, Jialu (2018), Countercyclical risks and portfolio choice over the life cycle: Evidence and theory, *in* '9th Miami Behavioral Finance Conference'.

Vestman, Roine (2019), 'Limited stock market participation among renters and homeowners', *The Review of Financial Studies* **32**(4), 1494–1535.

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.

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.

Appendix

A Data and summary statistics

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 Blundell, Pistaferri and Saporta-Eksten (2016). In particular, we consider households with heads aged 25 to 60 years old, who are continuously married, and who have continuously participated in the labor force. This leaves us with 10,655 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. Wealth is defined as total financial wealth, which is the sum of households' holdings in stocks, bonds, and cash, plus any amount

invested in retirement accounts. The risky share, then, is defined as the share of stocks in total financial wealth.

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 comparison with the PSID, we use similar criteria as in Blundell et al. (2016). We also remove households with incomplete information on education, age, and other demographic information. We also remove households that have zero labor income, and who have less than \$100 in financial assets, following Fagereng et al. (2017). This criteria gives us a sample of 54,321 households. Given that the SCF oversamples the wealthy, we use weights in the calculation of the auxiliary statistics. To have a comparable sample period as with the PSID, we work with the 1998-2016 waves.

A.3   Some summary statistics

We now compare some summary statistics that we obtain with the PSID and the SCF. In particular, we show statistics with respect to income, wealth, the conditional risky share and stock ownership, which we show in Table A1. As the table illustrates, the resulting distributions and summary statistics are similar in both datasets.

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:

$$\eta_{it} = Q_t(\eta_{it-1}, u_{it}), \quad (u_{it}|\eta_{it-1}, \eta_{it-2}, \ldots) \sim U[0,1], \quad t = 2, \ldots, T. \tag{B1}$$

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.

	PSID				
	Mean	SD	P25	Median	P75
Income	99568.92	114983.40	54396.30	80017.49	114084.40
Financial wealth	163701.90	390604.80	9573.30	46909.14	162002.00
Stock ownership	0.59	0.49	0.00	1.00	1.00
Conditional risky share	0.56	0.28	0.36	0.52	0.82
	SCF				
	Mean	SD	P25	Median	P75
Income	122378.90	195459.40	62653.93	95062.34	137097.10
Financial wealth	246234.70	1158090.00	8983.13	106.04	187084.30
Stock ownership	0.70	0.32	0.00	0.27	0.59
Conditional risky share	0.48	0.28	0.24	0.46	0.70

Table A1: Comparison, PSID vs. SCF

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

$$\rho(\eta_{it-1}, \tau) = \frac{\partial Q_t(\eta_{it-1}, u_{it})}{\partial \eta} \tag{B2}$$

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 NL process allows current shocks to wipe out the memory of past shocks. By contrast, in the CA 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 NL process allows for conditional dispersion, skewness and kurtosis in η_{it}.[8]

Following Arellano et al. (2017), we specify the quantile functions for the persistent and

---

[8]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 $\tau \in (1/2, 1)$.

transitory components as lower-order Hermite polynomials:

$$Q_t(\eta_{it-1}, \tau) = \sum_{k=1}^{K} a_k^{\eta}(\tau) f_k(\eta_{it-1}, age_{it}) \tag{B3}$$

$$Q_t(\eta_{i1}, \tau) = \sum_{k=1}^{K} a_k^{\eta 1}(\tau) \tilde{f}_k(age_{i1}) \tag{B4}$$

$$Q_t(\varepsilon_{it}, \tau) = \sum_{k=1}^{K} a_k^{\varepsilon}(\tau) f_k^{\varepsilon}(age_{it}) \tag{B5}$$

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 $[\tau_1, \tau_2]$, ..., $[\tau_{L-1}, \tau_L]$, 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 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[9]. 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:

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

---

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

This is equivalent to:

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

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:

$$\nu_{it} = \rho^{t-1}\eta_{i0} + \sum_{j=2}^{t} \rho^{t-j} u_{ij} + \varepsilon_{it} \tag{B6}$$

from which

$$var(\nu_{it}) = \rho^{2(t-1)}\sigma_z^2 + \sum_{j=2}^{t} \rho^{2(t-j)}\sigma_u^2 + \sigma_\varepsilon^2 \tag{B7}$$

$$cov(\nu_{it}, \nu_{it-1}) = \rho^{2t-1}\sigma_z^2 + \sum_{j=2}^{t} \rho^{1+2(t-j)}\sigma_u^2 \tag{B8}$$

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.

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 non-linear 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), which we call KO in the main text; 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 non-linearities, which we call KO + NN process.

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:

$$y_{it} = \alpha_i + \eta_{it} + \varepsilon_{it} \tag{B9}$$

$$\eta_{it} = \rho_t \eta_{it-1} + u_{it} \tag{B10}$$

$$u_{it} \sim N(0, \sigma_{u,t}^2), \varepsilon_{it} \sim N(0, \sigma_{\varepsilon,t}^2) \tag{B11}$$

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 we restrict the dependence between η_{it} and η_{it-1} to be linear. This yields the following restricted model specification for the dynamics persistent component, for a given quantile τ:

$$\eta_{it} = Q_t(\eta_{it-1}, u_{it}) = b_0(\tau) + b_1(\tau)\phi_1(\eta_{it-1}) + \gamma_1(\tau)age_{it} + \gamma_2(\tau)age_{it}^2 \tag{B12}$$

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.) non-linearity and (iii.) non-normality.
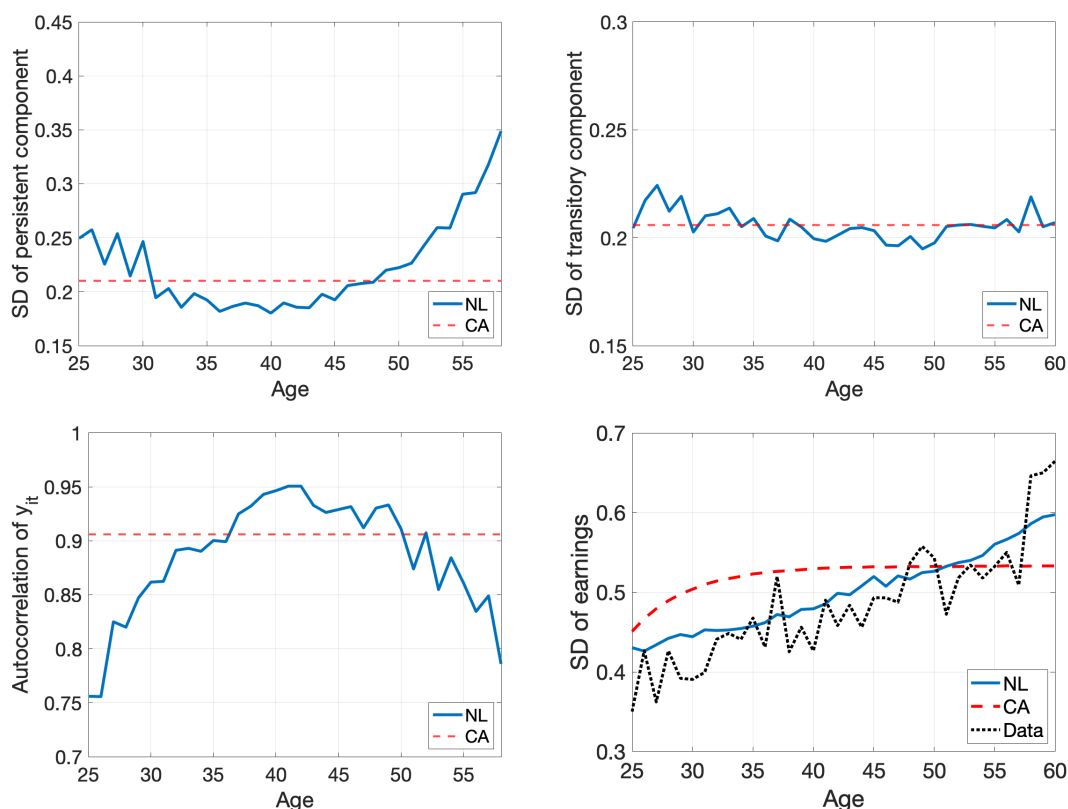

Figure B1: Age dependence of moments, canonical (red) vs. nonlinear (blue) model, PSID. The upper left figure presents the standard deviation of the persistent component of income, graphed by age. The upper right figure presents the standard deviation of the transitory component of income, graphed by 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.

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 earnings over the life cycle in our sample, which the canonical process clearly cannot.
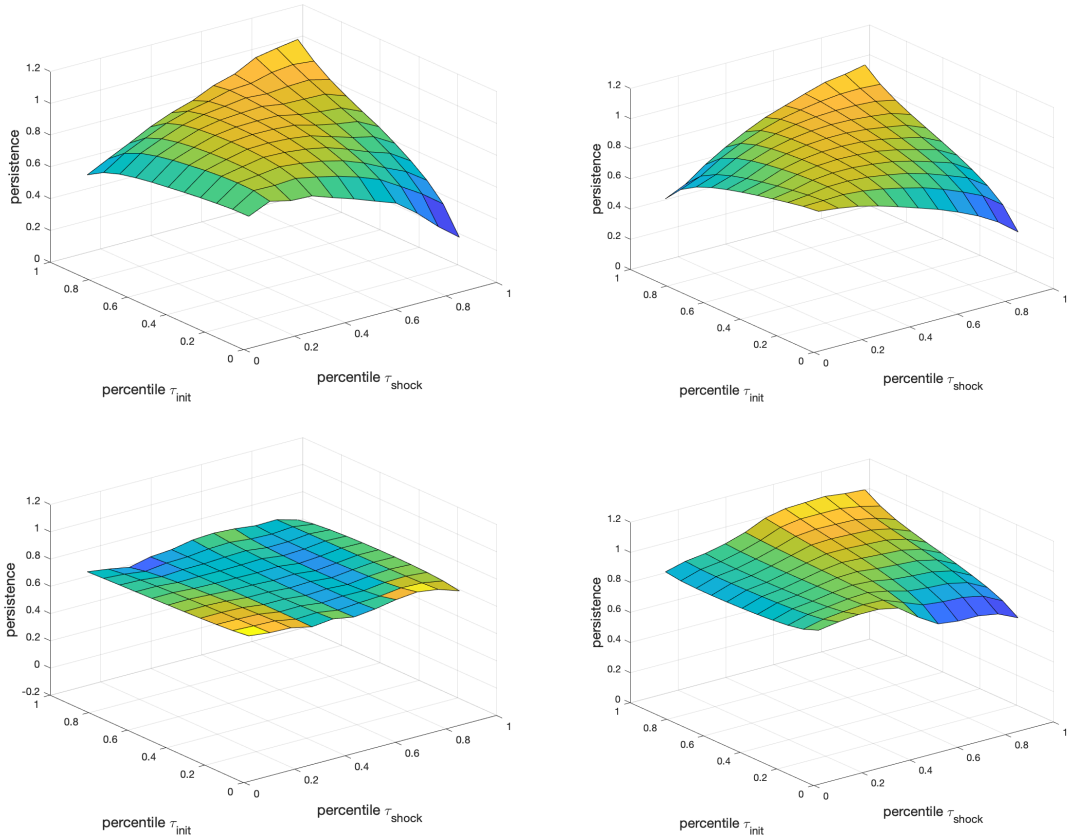
Figure B2: Persistence in the PSID. The upper left panel presents the graph of the average derivative of y_{it} given y_{it-1}, with respect to y_{it-1}, which was estimated from a quantile autoregression of y_{it} on a third-order Hermite polynomial on y_{it-1}. The upper right panel presents the same 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}.

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 the fact that the figure is net of transitory shocks. The associated standard errors, which are in Figure B3, are small.
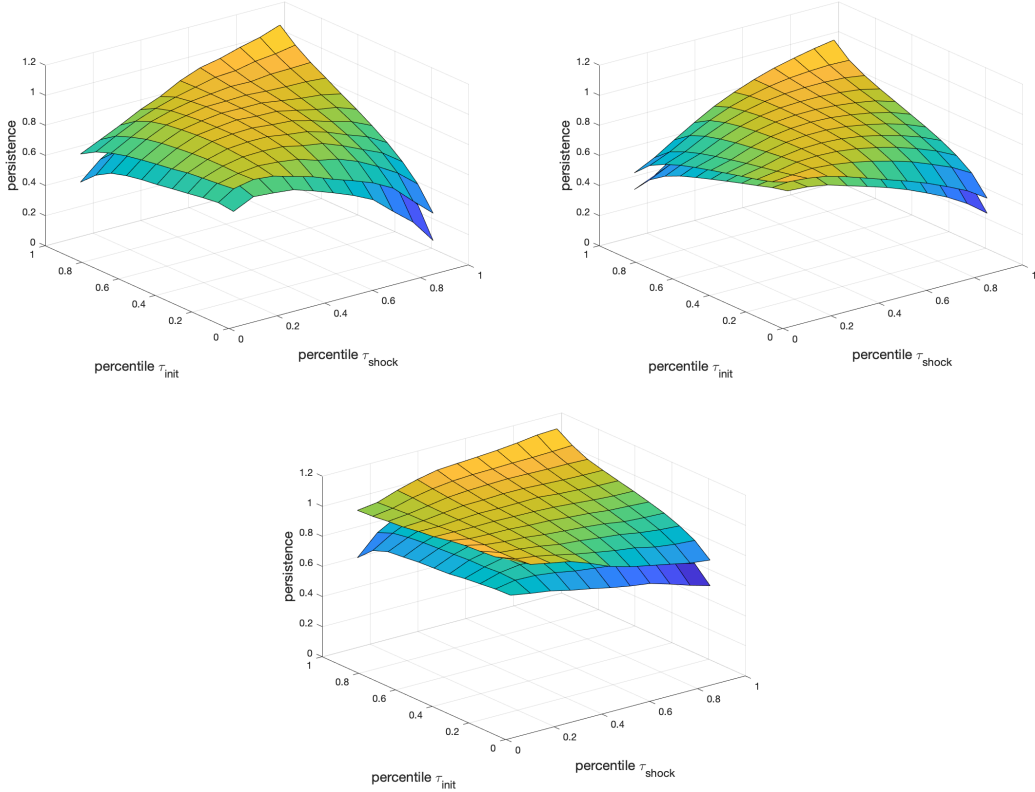
Figure B3: Persistence in the PSID, nonparametric bootstrap. The graphs presented here show the uniform 95% confidence bands calculated from nonparametric bootstraps. The top left panel presents the graph of the average derivative of y_{it} given y_{it-1}, with respect to y_{it-1}, which was estimated 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.
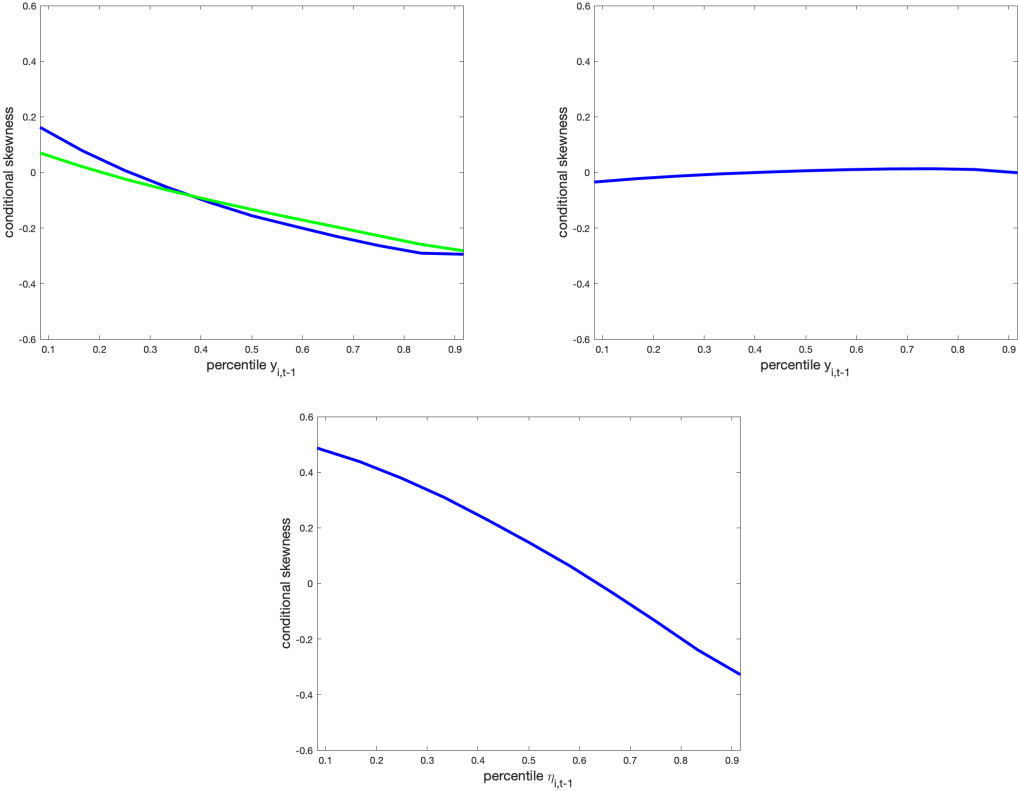
Figure B4: Conditional skewness in the PSID. The left panel presents the graph of the conditional skewness in the data (blue) and the conditional skewness of simulated data from the nonlinear earnings model (green). The right panel presents the conditional skewness based on simulated data from the canonical earnings model.
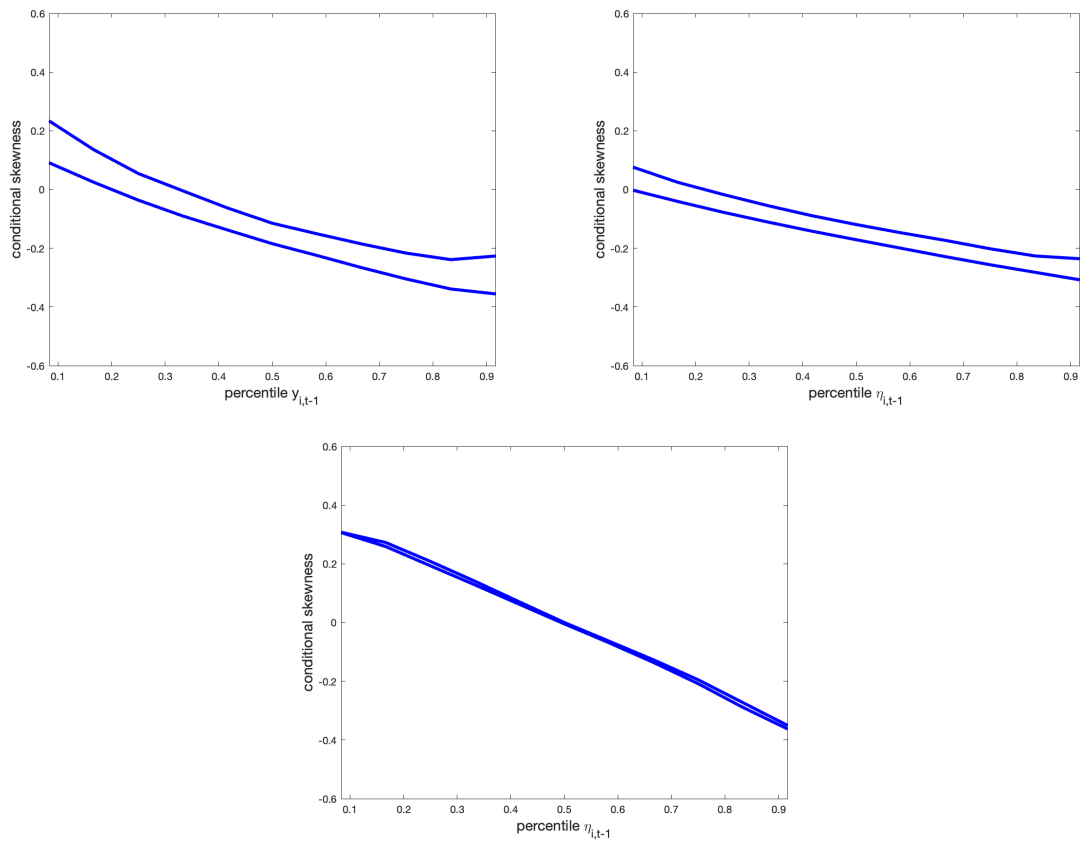
Figure B5: Conditional skewness in the PSID, bootstrap confidence intervals, nonparametric bootstrap. The graphs presented here show the uniform 95% confidence bands. The top left panel presents the graph of the conditional skewness of earnings data y_{it}. The top right panel presents the conditional 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 Model and model implications

C.1 Moments for the 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 choose nine moments to fit four parameters, which we obtain from the more recent waves of the PSID.

C.1.1 Stock market participation moments

Our first seven targeted moments are related to stock market participation decisions over the life cycle. In this context, we focus on the stock market participation rate (0.535), and the parameters of an OLS regression that examines the determinants of stock market participation. The empirical specification we pursue extends that of Alan (2006), who includes only a past participation dummy and demographics, by introducing income and wealth, which are state variables in the structural model. We report the results of the estimation in Table C1 below. The results of the regression are consistent with previous literature (see, e.g., Bonaparte et al. (2020)). In particular, the OLS coefficient on past participation suggests that there is persistence in stock market participation over the life-cycle, while increases in income and wealth lead to increase in the probability of stock market participation.

Past participation	Log financial wealth	Log household income	Age of household head	Age of head squared	Constant
0.330***	0.073***	0.021***	-0.011*	6.33e-05	-0.093
(0.009)	(0.002)	(0.006)	(0.006)	(6.35e-05)	(0.142)

Table C1: Parameters of an OLS regression on the determinants of stock market participation.

C.1.2 Wealth moments

The last two targeted moments are related to the evolution of wealth over the life cycle. In particular, we choose the conditional risky share (0.556) and the mean financial wealth-to-income ratio (2.00).

C.2 SMM estimation

In our SMM estimation, we pick 4 parameters (risk aversion γ, discount rate β, and participation costs κ^{FC}, κ^{PP}) to match 9 targets in the data (percentage of people that own stocks directly, median financial wealth to income, conditional risky share, and 6 parameters from the OLS regression in Table C1).

In order to measure the variability of these moments in the data, we use a bootstrap procedure in which we draw 1,000 samples with replacement from our PSID data. With this, we construct a data variance-covariance matrix S.

We are particularly interested that our model under both earnings processes closely replicates stock market participation, conditional risky shares, and median financial wealth to income ratios. Thus, we choose to increase the weight of these moments in our estimation procedure. Thus, our weighting matrix W is diagonal and is formed of the inverse of the standard deviations of the data moments (for the OLS parameters), 10 (for the wealth to income ratio), and 1000 (participation ratio and risky share). We have experimented with alternative values for these weights and results are very similar.

Finally, we compute a matrix D that measures the responsiveness of our parameter estimates to changes in the moments in the data. We estimate this gradient matrix numerically.

In order to compute the standard errors for our parameter estimates, we compute a variance covariance matrix V determined by (following the notation in De Nardi et al. (2010)):

$$V = (1 + \tau)(D'WD)^{-1}D'WSWD(D'WD)^{-1} \tag{C13}$$

where τ is the ratio between the number of simulated households in the model and the number of households in the data. Our results do not change if we consider that S has zeros outside of the main diagonal.

C.3 Life-cycle profiles

In Figure C1, we compare the resulting life-cycle profiles from the structural models that we have estimated with the life-cycle profiles from the data.[10] We find that both versions of the

---

[10]To compute the life-cycle profiles of stock market participation and the share of financial assets invested in risky wealth, we estimate an OLS regression with each of these variables as a function of a set of age, time, and cohort dummies, using the Deaton and Paxson (1994) restrictions to identify the age effects. To compute average financial wealth, we take the mean of financial wealth by age.

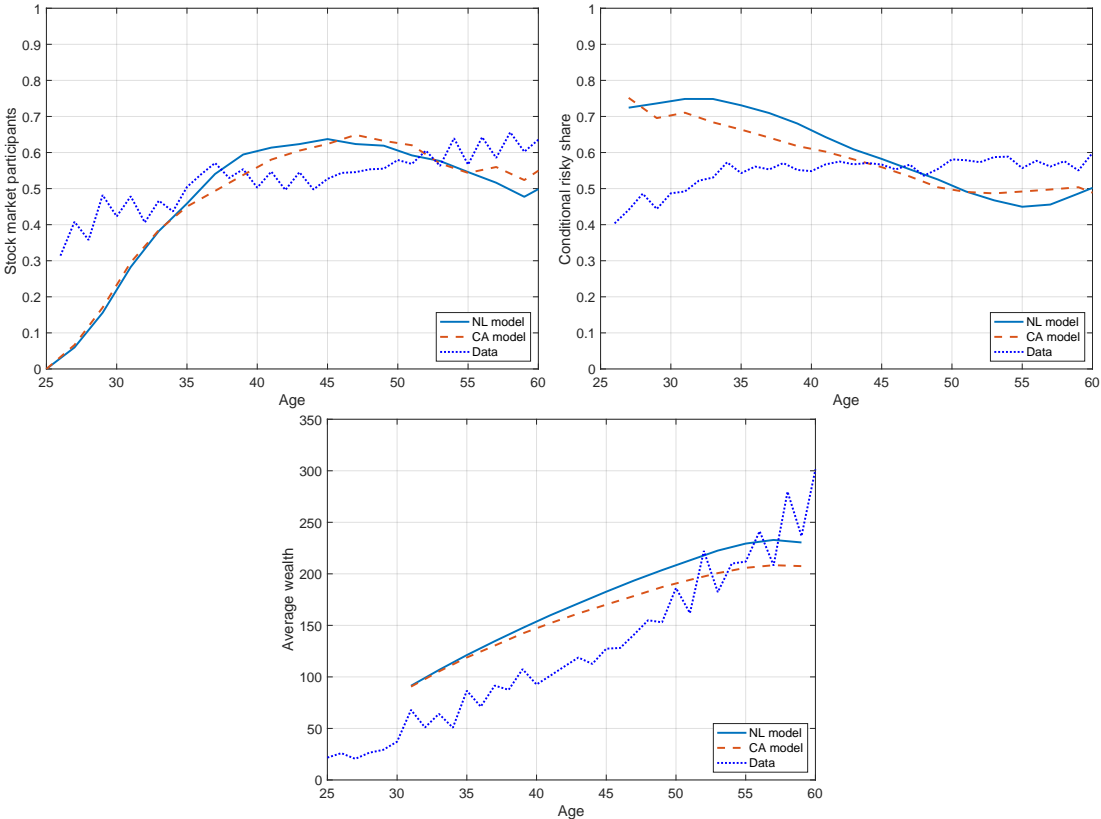model are able to fit well the life-cycle profiles.

Figure C1: Stock market participation (top left), share of financial assets invested in stocks (top right), average financial wealth in thousands of dollars (bottom), data vs. model.

C.4 Sensitivity of moments to parameters

In Figure C2 we provide an intuitive measure of how the four estimated parameters are identified by and relate to the main moments we are interested in targetting. Namely, we represent by how much each of the moments (average wealth to income ratios, share of participants in the stock market, conditional risky share of financial assets, persistence of stockholding status) changes when we make changes to each of the parameters (CRRA coefficient, discount rate and participation costs) while keeping all else constant. The changes in the moments are represented as absolute deviations from their levels implied by our main NL 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 as we change β.

The participation cost parameters govern the level of participation in the stock market, which responds very strongly to their changes. The entry cost parameter is also tightly connected to the OLS regression parameter that determines the persistence of stockholding; when entry costs are very high, the stock holding status in the model is very persistent. Our main estimated parameters, as described in Section 5.2, slightly overestimate this persistence. As the lower two panels of Figure C2 show, that happens because entry costs are already at their lower bound of zero and we can't further reduce the per-period costs without overestimating the overall level of stock market participation.
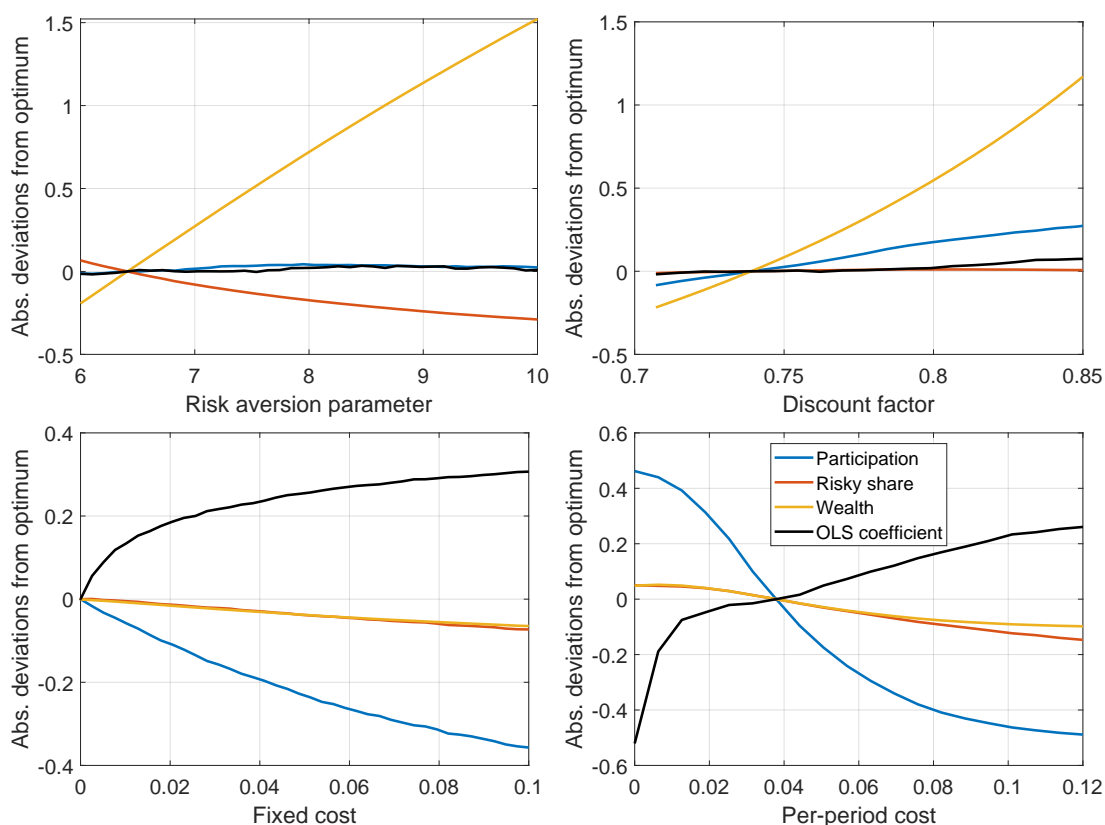
Figure C2: Deviations of the moments with respect to their targeted values, under the NL earnings process, as we ceteris paribus change the coefficient of risk aversion (top left), the discount rate (top right), the fixed costs of participation in the stock market (bottom left) and the per-period cost of participation in the stock market (bottom right). Deviations are expressed in the same units as the moments.

C.5   Optimal investment

Figure C3 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 canonical 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.
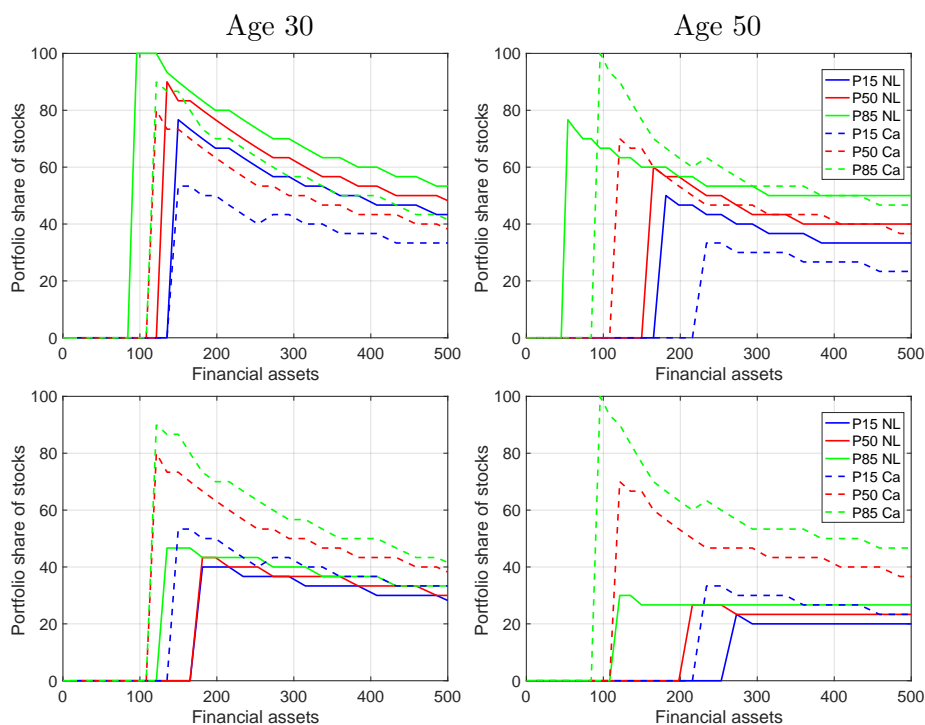

Figure C3: Optimal portfolio share of stocks by level of wealth (x-axis), earnings process (straight lines: NL; dashed lines: canonical), and position in the income distribution (percentile 15, blue, median worker, red, percentile 85, green). Top: estimated parameters for each process; bottom: estimated parameters for the canonical 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.

C.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:

$$c_{it} = f_t(\eta_{it}, \varepsilon_{it}, a_{it}, u_{it}), \tag{C14}$$

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:

$$E(c_{it}|\eta_{it} = \eta, \varepsilon_{it} = \varepsilon, a_{it} = a) = E(f_t(\eta, \varepsilon, a, u_{it})), \tag{C15}$$

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

$$c_{it} = \sum_{k=1}^{K} a_k f_k(\eta_{it}, \varepsilon_{it}, a_{it}, age_{it}) + a_0(\tau), \tag{C16}$$

where a_k are piecewise polynomial interpolating splines, and f_k's are dictionaries of functions, which are assumed to be Hermite polynomials. 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. To be consistent with Arellano et al. (2017), we use the same approximating polynomials as their paper.[11] 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 Blundell et al. (2008) coefficients, which are in Table C4, show that compared to the benchmark BPP estimate, consumption insurance is higher in the non-linear economy than in the canonical economy.
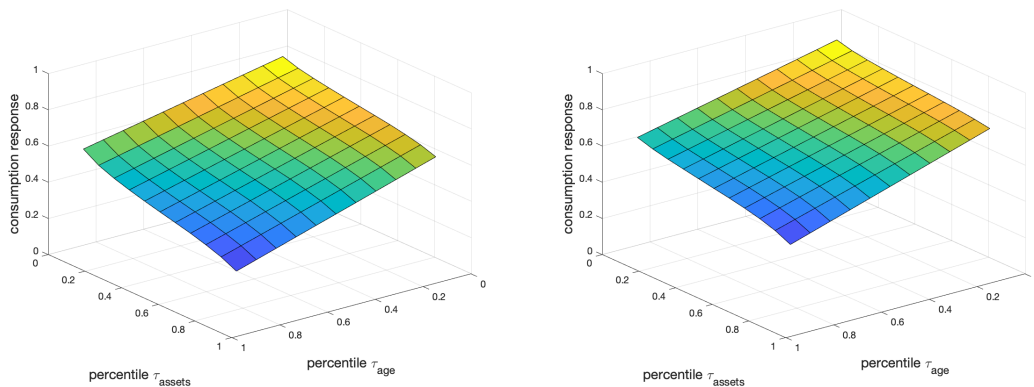

Figure C4: Consumption response to earnings shocks, nonlinear vs. linear model. Note: The graphs presented here show the average derivative effect of η_{it} on c_{it}, computed at percentiles of a_{it} and age_{it}. Data simulated from structural model of life cycle portfolio choice with the nonlinear earnings process (left) and the canonical earnings process (right).

	All		Stockholders		Non-stockholders	
	Persistent	Transitory	Persistent	Transitory	Persistent	Transitory
Canonical	0.29	0.99	0.34	0.99	0.31	0.99
Non-linear	0.41	0.99	0.44	0.99	0.38	0.99

Table C2: Consumption insurance parameters, implied Blundell et al. (2008) coefficients.

---

[11]This is (2,1,2,1), where the order is (persistent,transitory,wealth,age).

We finally compute the average derivative effects for non-stockholders and stockholders, in the case of the economy with the nonlinear earnings process. The results that we obtain are in Figure C5. The top left panel illustrates that the derivative effects with respect to consumption for non-stockholders is higher than that of stockholders. These results imply that stockholders are able to effectively insure their consumption with respect to shocks to persistent income. The MPC's with respect to wealth are also higher for households who own stocks than for non-stockholders, as shown in the top right panel of the figure. Finally, we calculate the implied Blundell et al. (2008) coefficients. As indicated in the middle and right columns of Table C2, we find that stockholders, are on average, more able to insure their consumption against persistent income shocks. Moreover, this result holds true over the life-cycle, as is shown in the bottom panel of Figure C5.
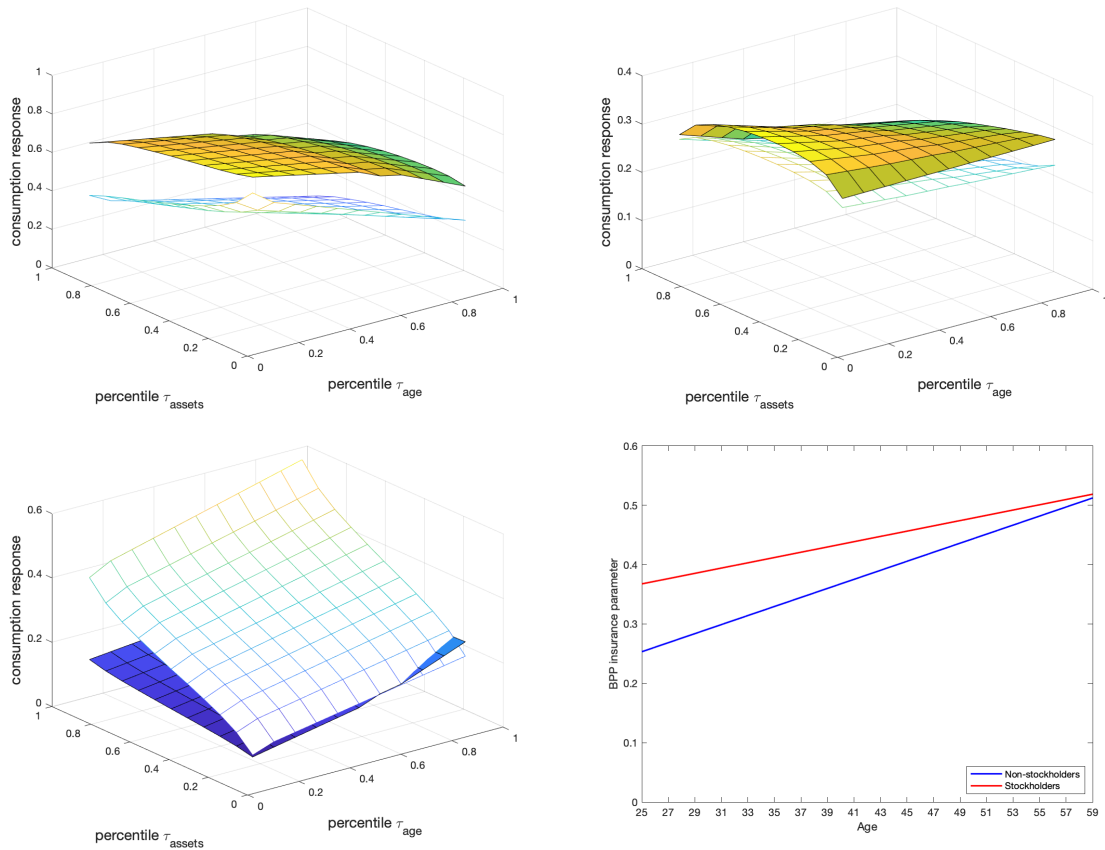
Figure C5: Consumption response to earnings shocks, stockholders vs. non-stockholders. Note: The top left panel graphs show the average derivative effect of η_{it} on c_{it} (left), computed at percentiles of a_{it} and age_{it} for non-stockholders (surface graph) and stockholders (mesh graph), estimated using simulated data from the structural model. The top right panel graphs show the average derivative effect of η_{it} on c_{it} (left), computed at percentiles of a_{it} and age_{it} for non-stockholders (surface graph) and stockholders (mesh graph), estimated using PSID data. The bottom left panel graphs show the average derivative effect of a_{it} on c_{it} (left), computed at percentiles of a_{it} and age_{it} for non-stockholders (surface graph) and stockholders (mesh graph). The bottom panel compares the implied BPP insurance parameters for stockholders and non-stockholders under the nonlinear economy. Data simulated from structural model of life cycle portfolio choice with the nonlinear earnings process.

D Robustness

D.1 Targetting total wealth

In our baseline results, we estimate the model so that it replicates the total amount of financial wealth in the economy (namely, a financial wealth-to-income ratio of 2). However, a significant amount of household wealth is held in the form of housing, raising the total wealth-to-income ratio in the United States to close to 5. As a result, in our baseline results the discount rates are remarkably low (in the region of 0.70, similarly to Fagereng et al. (2017), who also target only financial wealth).

In Section 5.4 we showed that our main implications are also present in a model in which housing is considered explicitly. Here, we perform an additional experiment that is also frequently considered in the literature: we estimate the model to the total amount of wealth in the economy, counterfactually assuming that it is all financial wealth. This also implies changing the wealth component in our OLS regression to reflect total, rather than financial, wealth.

As Table D1 shows, this alternative calibration delivers similar implications in terms of the comparison between the NL and canonical process, but a much higher discount rate. We conclude that the differences between processes are not limited to models with relatively low discount rates and low levels of aggregate wealth. However, this particular version of the model requires higher per-period participation costs to rationalize the low level of stock market participation in the data.

| **Model** | Targetting total wealth | |
	NL	CA
γ	6.31	7.85
β	0.945	0.951
κ^{FC}	0	0
κ^{PP}	0.1200	0.1137

Table D1: Parameter estimates. β is expressed in annual terms. The participation costs are expressed as fractions of average household income (the numeraire in the model)

Model		Targetting total wealth	
Moment	Data	NL	CA
Total W/I	4.8	5.1	5.0
Participation	0.53	0.504	0.528
OLS constant	-0.681 (0.177)	7.864	4.593
OLS partic	0.399 (0.011)	0.724	0.728
OLS age	-0.0188 (0.234)	-0.6523	-0.3686
OLS age2	0.00014 (7.93e-5)	2.0e-02	1.1e-02
OLS log income	0.0437	0.11288	0.095
OLS log wealth	0.0780	0.089127	0.084147
Risky share	0.56	0.552	0.550

Table D2: Targeted parameters
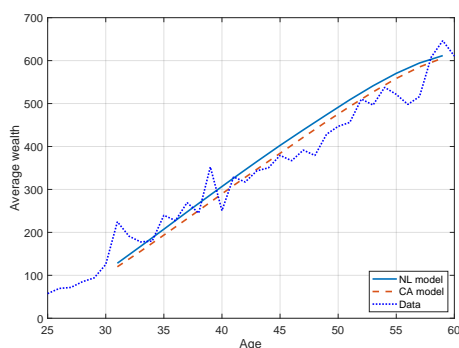
Figure D1: Average wealth over the life-cycle, targetting total wealth

D.2   Disaster risk

In this experiment, we assume that agents incoporate in their expectations the possibility that stock market returns are abnormally low (which is commonly referred to as a "disaster event" in the stock market). Namely, we adopt the quantification in Fagereng et al. (2017) and set the ex-ante probability of a large drop in the stock market to 2% in annual terms, and the amount of the average drop conditional on the abnormal return realization to -48.5%.

Model	Case with disaster risk	
	NL	CA
γ	5.57	7.48
β	0.849	0.838
κ^{FC}	0	0
κ^{PP}	0.0405	0.0338

Table D3: Parameter estimates. β is expressed in annual terms. The participation costs are expressed as fractions of average household income (the numeraire in the model)

Model		Baseline	
Moment	Data	NL	CA
Participation	0.535	0.519	0.535
Risky share	0.56	0.550	0.544
Average W/I	2	2.5	2.3
OLS constant	-0.136 (0.145)	-0.424	-0.272
OLS partic	0.329 (0.009)	0.594	0.595
OLS age	-0.0115 (0.006)	0.0323	0.0252
OLS age2	6.96e-5 (6.16e-5)	-4.43e-04	-3.44e-04
OLS log income	0.0271 (0.007)	0.152	0.170
OLS log wealth	0.0729 (0.002)	0.124	0.075

Table D4: Targeted parameters

D.3   Empirical initial wealth

In this section, we assume that households, instead of starting out life with zero wealth, they do so with a draw from the empirical wealth distribution of the 2016 SCF (20-30 year olds). For simplicity, we assume that this initial wealth draw is uncorrelated with initial income.

Model	Baseline	
Parameter	NL	CA
γ	6.30	8.17
β	0.806	0.817
κ^{FC}	0	0
κ^{PP}	0.0455	0.0444

Table D5: Parameter estimates. β is expressed in annual terms. The participation costs are expressed as fractions of average household income (the numeraire in the model)

Model		Baseline	
Moment	Data	NL	CA
Participation	0.535	0.517	0.520
Risky share	0.56	0.540	0.550
Average W/I	2	2.5	2.5
OLS constant	-0.136 (0.145)	-0.453	-0.070
OLS partic	0.329 (0.009)	0.579	0.611
OLS age	-0.0115 (0.006)	0.0335	0.0149
OLS age2	6.96e-5 (6.16e-5)	-4.59e-04	-2.34e-04
OLS log income	0.0271 (0.007)	0.140	0.170
OLS log wealth	0.0729 (0.002)	0.144	0.079

Table D6: Targeted parameters
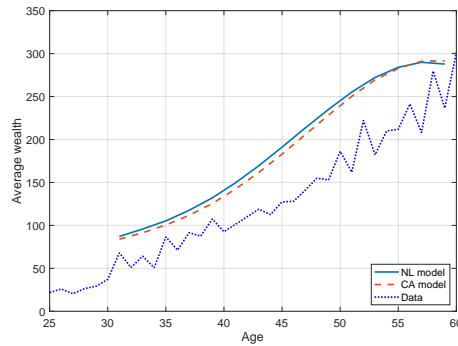
Figure D2: Average financial wealth over the life-cycle, case with empirical initial wealth.

D.4 Alternative pension system

In this section, we assume that, instead of a flat pension, households receive a replacement rate of the last realization of their earnings that is consistent with U.S. data. Namely, we follow Kaplan and Violante (2010) and assume that there is a 90% replacement rate for any earnings below 18% of the average, 32% for the earnings between 18% and 110% of the average, and 15% for the remainder. For simplicity, we only apply these replacement rates to the persistent component of earnings.

This different pension system affects the incentives to save, and thus implies a lower calibrated discount factor for both the NL and canonical process. However, the implications in terms of a lower coefficient of risk aversion and a more reasonable profile of stock market participation for the NL process still hold true.

Model	Baseline	
Parameter	NL	CA
γ	5.89	8.67
β	0.820	0.755
κ^{FC}	0	0
κ^{PP}	0.0714	0.0486

Table D7: Parameter estimates. β is expressed in annual terms. The participation costs are expressed as fractions of average household income (the numeraire in the model)

Model		Baseline	
Moment	Data	NL	CA
Participation	0.535	0.509	0.507
Risky share	0.56	0.531	0.515
Average W/I	2	3.1	2.5
OLS constant	-0.136 (0.145)	-1.235	-0.968
OLS partic	0.329 (0.009)	0.591	0.609
OLS age	-0.0115 (0.006)	0.0648	0.0487
OLS age2	6.96e-5 (6.16e-5)	-7.58e-04	-5.19e-04
OLS log income	0.0271 (0.007)	0.129	0.160
OLS log wealth	0.0729 (0.002)	0.126	0.069

Table D8: Targeted parameters
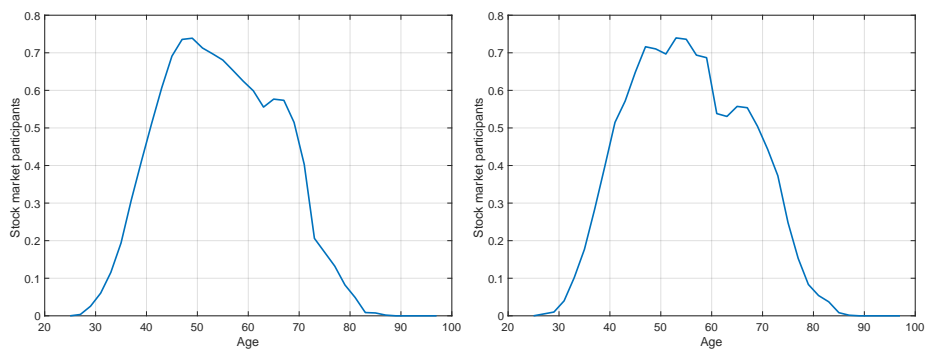
Figure D3: Stock market participation over the life cycle. Left: NL process, right: canonical process.

D.5 Correlation (0.20) between labor income and stock market returns

For this section, we assume that shocks to labor market income are correlated with returns in the stock market. This correlation is both perceived in expectation by the agents, which implies that they adjust their portfolio decisions accordingly, and occurs in the simulations.

Model	Positive correlation (0.20)		Negative correlation (-0.20)	
Parameter	NL	CA	NL	CA
γ	5.95	8.37	6.6	8.2
β	0.845	0.811	0.7746	0.80623
κ^{FC}	0	0	0	0
κ^{PP}	0.0457	0.0384	0.0444	0.0444

Table D9: Parameter estimates

E Portfolio choice under the presence of housing

This section provides a description of the life-cycle structural model of portfolio choice with an endogenous housing decision. We also discuss the moments that we use for the structural estimation.

E.1 Model

In this subsection, we describe the structural model that we take to the data.

Demographics. Agents start working life at 25, face age-dependent positive death probabilities, and die with certainty at age 100. The model period is two years and our unit of observation is the household.

Preferences. Households maximize:

$$\max E_t \left[\sum_{t=0}^{t=T} \beta^t \mathcal{S}_t \frac{[c_t \phi(h_t)]^{1-\gamma}}{1-\gamma} \right] \tag{E17}$$

where c is nondurable consumption, γ is the coefficient of relative risk aversion, β is the discount factor, $\phi(h_t)$ is a function of the current housing status, normalized such that $\phi(0) = 1$ and $\phi(1) = u_h$, and $\mathcal{S}_t$ is the probability of arriving alive at time t.

Earnings process. As in the main text, the stochastic component of earnings can be decomposed to a persistent and transitory component (Equation 1), and we use alternatively a canonical linear specification and a non-linear, non-normal specification for both components of the earnings process. There is no earnings risk after retirement (age 65), from which households get a public pension.

Budget constraint. Households can save in two types of financial assets and in housing:

$$c_{t+1} + s_{t+1} + a_{t+1} + \kappa^f(I_{t+1}, I_t) + p_h h_{t+1} + \kappa^h(h_{t+1}, h_t) + r^h I(h_t = 0) + m_{t+1} = \quad \text{(E18)}$$

$$(1 + r^s_{t+1})s_t + (1 + r)a_t + y_{t+1} + p_h h_t + (1 + r^m)m_t \quad \text{(E19)}$$

where s_t is the amount of wealth invested in the risky asset and a_t the amount of wealth invested in the risk-free asset at time t. r^s_{t+1} represents the risky return of stocks (which is i.i.d.), while r is the risk-free rate. κ^f represents potential costs of participation in the stock market, which depend on the households' stock market participation status I_t. We define

$$I_t = (s_t > 0). \quad \text{(E20)}$$

These may be per-period participation costs (just dependent on I_{t+1}), fixed participation costs (only paid if $I_t = 0$ and $I_{t+1} = 1$, and zero if $I_t = 1$) or a combination of both.

p_h represents the (fixed) price of housing, $\kappa^h(h_{t+1}, h_t)$ represents the transaction costs associated with the purchase of the house, which we model as a fixed fraction of the house price for both buyer and seller, and r^h represents the rental rate for those who do not own a house, which is a fixed fraction of house prices. m_t represents the outstanding mortgage balance ($m_t \leq 0$) and r^m the associated mortgage balance.

Households can borrow in order to buy a house. Borrowing constraints are thus:

$$a_{t+1} \geq 0, s_{t+1} \geq 0 \quad \text{(E21)}$$

for renters and

$$a_{t+1} \geq 0, s_{t+1} \geq 0, m_{t+1} \geq -\lambda_h p_h h_{t+1} \quad \text{(E22)}$$

for homeowners.

For simplicity, we assume that the borrowing constraint on mortgages is always binding $m_{t+1} = -\lambda_h p_h h_{t+1}$ (as in Vestman (2019)).

Finally, we define

$$x_t = (1 + r_{t+1}^s)s_t + (1 + r)a_t + (1 + r^m)m_t \tag{E23}$$

as the amount of cash-on-hand that an individual has at the beginning of period t.

Households' problem   Households thus solve the following problem:

$$V_t(x_t, y_t, I_t, h_t) = \max_{c_t, a_{t+1}, s_{t+1}, h_{t+1}} \left\{ \frac{(c_t\phi(h_t)^{1-\gamma}}{1-\gamma} + \beta\frac{\mathcal{S}_t}{\mathcal{S}_{t-1}} E_t V_{t+1}(x_{t+1}, y_{t+1}, I_{t+1}, h_{t+1}) \right\} \tag{E24}$$

subject to the budget constraint (E19) and the borrowing and short-sale constraints (E21) and (E22), and where the expectation E_t is taken with respect to future realizations of persistent income, transitory income, and stock market returns.

E.2   Targeted moments and model fit

All externally calibrated parameters are identical to those in the main version of the model. Additionally, we assume that the house price is fixed at $p^h = 4$ years of average disposable household income, transaction costs are $k^h(h_{t+1} = 1, h_t = 0) = k^h(h_{t+1} = 0, h_t = 1) = 0.05p^h$, annual rental costs are 2.5% of the house price, and the interest rate on mortgages is $r^m = r^a = 2\%$ annually.

We estimate the life-cycle model of portfolio choice under the presence of housing via SMM. The difference between the model we present here and the baseline model, is that we target additional moments that are informative of the homeownership status of households. Specifically, we target the following additional moments: homeownership rate (0.76), the total wealth-to-income ratio (4.8), and the parameters of an OLS regression which now includes a dummy for homeownership, following Yao and Zhang (2005).

Model fit.   The fit of the model with respect to the targeted moments is below in Table E1. As can be observed, both processes fit the data well, with the non-linear earnings process fitting the homeownership rate better than the canonical process.

Model		Baseline	
Moment	Data	NL	Ca
Participation	0.535	0.543	0.506
Risky share	0.556	0.53645	0.53202
Total W/I	4.8	4.9052	4.8899
OLS constant	-0.832 (0.137)	-0.98483	0.4019
OLS partic	0.436 (0.009)	0.4136	0.63574
OLS age	-0.0157 (0.006)	0.047838	-0.015651
OLS age2	0.00012 (6.21e-5)	-0.00054716	0.00014814
OLS log income	0.0619 (0.00662)	0.20155	0.1317
OLS log wealth	0.0642 (0.00286)	0.022024	0.00081136
OLS homeownership	-0.0720 (0.0126)	0.29955	0.21343
Homeownership rate	0.76	0.57499	0.5413

Table E1: Targeted parameters for the portfolio choice model under housing

E.3   Homeownership and life-cycle profiles of portfolio choice
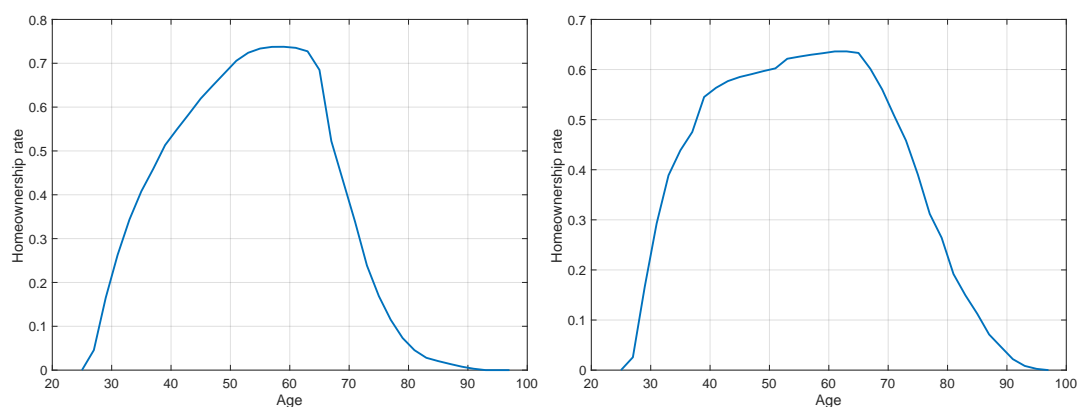
Figure E1: Homeownership rates by age (left: NL; right: CA)
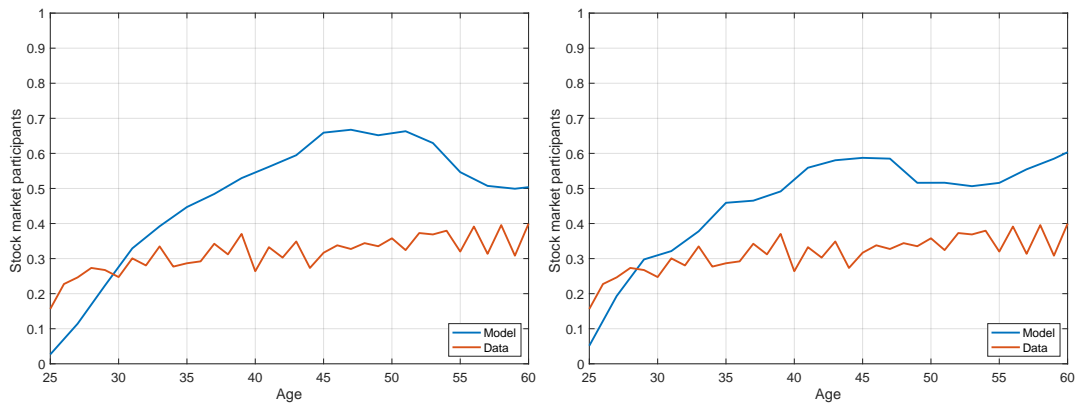
Figure E2: Stock market participation by age (left: NL; right: CA).

Julio Gálvez
Banco de España, Madrid, Spain; email: julio.galvez@bde.es

Gonzalo Paz-Pardo
European Central Bank, Frankfurt am Main, Germany; email: gonzalo.paz_pardo@ecb.europa.eu